

Università degli Studi "G. d'Annunzio" Chieti–Pescara
Dipartimento di Ingegneria e Geologia

Dissertazione presentata per il conseguimento del titolo
di Dottore di Ricerca in Engineering Science

Characterization and Mechanical Performance of Polypropylene Fibre Reinforced Concrete and Mortar Systems

Wesam Al Agha

Ciclo XXXVIII
Anno Accademico 2025–2026

Tutor:
Prof. Maria Laura De Bellis

Contents

1 Introduction **1**

1.1 Background and Motivation 1
1.2 Scientific Problem and Research Significance 3
1.3 Research Scope and Delimitations 6
1.4 Originality and Main Contributions 7
1.5 Thesis Organisation 8

2 State of the Art **11**

2.1 Fibre-Reinforced Cementitious Composites and Environmental Considerations 11
2.1.1 Material Properties of Steel, Glass, and Macro-Synthetic Fibres 11
2.1.2 Structural performance demands versus environmental assessment 12
2.2 Digital Image Correlation and Virtual Extensometry 12
2.2.1 Optical Deformation Measurement Principles 13
2.2.2 Virtual extensometry in fracture testing 14
2.3 Machine Compliance and Displacement Measurement Challenges . . . 14
2.3.1 Fundamentals of Testing Machine Compliance 14
2.3.2 Impact on Fracture Parameter Evaluation 15
2.3.3 Compliance Correction Methodologies 16
2.4 Fracture Kinematics (CMOD/CTOD) and Rigid-Hinge Models 16
2.4.1 The non-linear cracked hinge model 17
2.4.2 Kinematic Relations in Three-Point Bending 17
2.5 Statistical Frameworks and Life Cycle Assessment (LCA) 18
2.5.1 Life Cycle Assessment (LCA) framework 19
2.5.2 Structural Performance-Based Eco-Efficiency 20
2.6 Numerical Inverse Analysis and Constitutive Identification 20
2.6.1 Constitutive Models 21
2.6.2 Non-linear Least Squares Optimization 21

2.7 European Standardization Frameworks for FRC . . . 22
2.8 Identified Research Gaps . . . 24
2.8.1 Limitations of Current European Frameworks for Macro-Synthetic Fibres . . . 24
2.8.2 Integration of Digital Image Correlation with Compliance Correction . . . 25
2.8.3 Lack of Integrated Statistical and Eco-Efficiency Frameworks . 25
2.8.4 Mapping of Research Gaps to Thesis Objectives . . . 26

3 Methodology **27**
3.1 Research Framework and Experimental Programme . . . 27
3.2 Materials and Constituent Characterisation . . . 29
3.2.1 Cement . . . 29
3.2.2 Fine aggregate . . . 30
3.2.3 Coarse aggregate (concrete programme only) . . . 31
3.2.4 Superplasticiser . . . 32
3.2.5 Fibres . . . 32
3.3 Case Study 1: Comparative flexural response at $V_f = 1\%$ . . . 35
3.3.1 Cementitious matrices (Normal and High) . . . 36
3.3.2 Specimen matrix and mix proportions . . . 36
3.3.3 Casting, compaction and curing . . . 37
3.3.4 Notch preparation . . . 39
3.4 Case Study 2: SMP dose–response flexural behaviour . . . 39
3.4.1 Design variables . . . 39
3.4.2 Casting, compaction, curing and notch preparation . . . 41
3.5 Case Study 3: Cube splitting tensile of FRC (M30) . . . 41
3.5.1 M30 mix design . . . 41
3.5.2 Eight fibre groups . . . 42
3.5.3 Specimen geometry and splitting-tensile formula . . . 43
3.5.4 Casting, compaction and curing of the FRC cubes . . . 43
3.6 Case Study 4: Bonded two-bar splitting in FRC (M30) . . . 44
3.6.1 Specimen geometry . . . 44
3.6.2 Bond-stress formula . . . 45
3.6.3 Relationship to Case Study 3 . . . 45
3.7 Instrumentation and Measurement Framework . . . 45
3.7.1 Mechanical testing apparatus and loading protocols . . . 45
3.7.2 Digital Image Correlation – optical setup . . . 48
3.7.3 Digital Image Correlation – ZEISS Correlate processing . . . 50

3.8 Summary of the experimental programme 52

4 DIC-Derived Compliance Correction and Optical Validation 53

4.1 Why machine compliance distorts the raw force–displacement record . 53
4.2 Optical Deflection-Derived Compliance (ODDC) methodology 56
4.2.1 Governing equations 56
4.2.2 Identification of the machine stiffness from the reference mortar beam 57
4.2.3 Application of the correction 58
4.2.4 Scope of the ODDC correction 59
4.3 Optical validation through DIC point tracking 61
4.3.1 Point-trajectory extraction and interpolation 61
4.3.2 Best-point selection and validation metrics 61
4.3.3 Capability demonstration on the N-CONTROL beam 65
4.4 Case-study applications 66
4.4.1 Case Study 1 – notched FRM beams, fibre type at $V_f = 1\%$. 67
4.4.2 Case Study 2 – notched FRM beams, SMP dose and length sweep 78
4.4.3 Case Study 3 – cube splitting tensile on FRC cubes 84
4.4.4 Case Study 4 – bonded two-bar splitting on FRC cubes 89
4.5 DIC capability beyond compliance 95
4.5.1 Full-field view of fibre engagement 95
4.5.2 Post-hoc virtual-gauge ladder 96
4.6 Summary 97

5 Structural Performance, Statistical Analysis and Life-Cycle Assessment 99

5.1 Structural Performance Evaluation 100
5.1.1 Flexural quantities for the beams (Case studies 1 and 2) . . . 100
5.1.2 Splitting quantities for the cubes (Case studies 3 and 4) 101
5.2 Statistical Analysis Framework 104
5.2.1 Motivation and sample-size constraint 104
5.2.2 Normality screening: Shapiro–Wilk 104
5.2.3 Location-shift testing: Mann–Whitney U 105
5.2.4 Multi-group ranks: Kruskal–Wallis 105
5.2.5 Rank correlation: Spearman and Pearson heatmap 106
5.2.6 Dose–response: one-dimensional quadratic fit 106
5.2.7 Dose–length interpolation: cubic-spline contours 106

5.2.8 Cross-group consistency and combined ranking 107
5.3 Life-Cycle Assessment and Eco-Efficiency Index 107
5.3.1 Goal, scope and functional unit 107
5.3.2 Life-cycle inventory and emission factors 108
5.3.3 Monte Carlo uncertainty propagation 108
5.3.4 GWP penalty relative to the matrix control 109
5.3.5 Eco-Efficiency Index family 109
5.3.6 Sensitivity analysis 110
5.3.7 Cost–benefit scatter and literature comparison 110
5.4 Application to the four case studies 112
5.4.1 Case 1: micro-fibre screening in beam bending 112
5.4.2 Case 2: dose–length optimisation of SMP macro fibres 129
5.4.3 Case 3: cube splitting-tensile test (CST) 147
5.4.4 Case 4: bonded-two-bar splitting (BBS) and CST-vs-BBS cross-test novelty 163
5.5 Chapter summary and forward reference 180

6 Advanced Fracture Kinematics 183
6.1 Introduction to Kinematic Variables 184
6.1.1 Why kinematic variables matter 184
6.1.2 Definitions and geometry 185
6.1.3 Residual strengths and fracture energy (EN 14651) 185
6.1.4 Kinematic linear relations 186
6.1.5 Energy partitioning 186
6.1.6 Fibre Efficiency Index 187
6.1.7 UNI 11039-2 equivalent strengths 187
6.1.8 A remark on the Crack Profile Factor 187
6.2 CMOD and CTOD Mapping 189
6.2.1 Virtual extensometer workflow 189
6.2.2 Why CMOD and CTOD do not need a compliance correction 189
6.2.3 Kinematic fits — Case 1 190
6.2.4 Kinematic fits — Case 2 196
6.2.5 Tables of kinematic coefficients 200
6.2.6 What the coefficients reveal 205
6.3 Evaluation of the Rigid-Hinge Assumption 206
6.3.1 Case 1 cohort response 207
6.3.2 Case 1 per-specimen F–CMOD with markers 210
6.3.3 Case 2 cohort response 214

6.3.4 Case 2 per-specimen F–CMOD with markers 215
6.3.5 Dose-response sweeps 224
6.3.6 Where the rigid-hinge assumption fails 227
6.4 Summary of Findings 227

7 Numerical Inverse Analysis and Code Consistency 229
7.1 The Olesen/Ostergaard inverse framework 230
7.2 Trilinear σ–w laws and identification 230
7.3 Application to Case Studies 1 and 2 232
7.3.1 Case Study 1: first-group screen 232
7.3.2 Case Study 2: factorial (length × dose × matrix) 233
7.4 European standards, DIC kinematics, and the inverse pipeline 239
7.5 Summary of findings 243

8 Conclusions and Future Scope 245
8.1 Conclusions 245
8.1.1 Conclusions in relation to the research objectives 245
8.1.2 Scientific and technical application 246
8.1.3 Industry, practice, and standards 247
8.1.4 Social impact and sustainable development goals 247
8.2 Limitations 248
8.3 Future recommendations 249

References 251

List of Figures

1.1 Global cement production: recorded output (2021) and projected output (2030) [2]. 2
1.2 Examples of polypropylene, glass, and steel fibres used in cementitious composites. 3
1.3 Sources of extraneous displacement in a standardised bending test. . . 4
1.4 Schematic comparison of apparent (raw stroke) and DIC-corrected load–deflection responses (illustrative only). 5
1.5 Analytical workflow of the thesis. 7

2.1 Mathematical mapping between reference and deformed facets. 13
2.2 Serial spring representation of testing machine and specimen compliance. 15
2.3 Kinematic model of rigid-body rotation in a notched beam. 18
2.4 System boundaries for the cradle-to-gate environmental assessment. . 20
2.5 Phased workflow for the numerical inverse analysis. 22

3.1 Overview of the four case studies composing the experimental programme. 28
3.2 Schematic overview of the research methodology workflow, from raw material characterisation to analytical outputs. 29
3.3 Italcementi Tecnocem CEM II/B-LL 42.5R used for all mortar and concrete batches. 30
3.4 Mapefluid N200 polycarboxylate-based high-range water-reducing admixture, compliant with EN 934-2. 33
3.5 The twelve fibre variants used across the four case studies, displayed at a common scale. 34
3.6 Factorial design of Case Study 1 (22 mortar mixes). 35
3.7 Specimen preparation for the FRM beams: casting and curing. 38
3.8 Wet-sawing of the mid-span notch at 28 d with a diamond blade. . . . 39

3.9 Dose-response factorial of Case Study 2: two fibre lengths × five volume fractions × two matrices (control cells filled darker). 40
3.10 Plan of the eight fibre groups (G1–G8) of the M30 concrete programme used in Case Studies 3 and 4. 42
3.11 Galdabini Sun/60 testing frame in its three loading configurations. . . 46
3.12 DIC optical chain. 48
3.13 Representative tested specimens at the end of DIC-monitored loading. 51

4.1 Two contact sensors not used in this thesis: (a) clip-gauge extensometer; (b) surface-bonded strain gauge. 54
4.2 ZEISS Correlate v2025 workspace deployed in this chapter. 55
4.3 The ODDC splits the crosshead displacement into a machine-elastic and a specimen-elastic contribution arranged in series. 57
4.4 Optical validation of the ODDC correction on the N-CONTROL mortar beam. The corrected actuator displacement (solid blue) overlays the true DIC displacement (dashed red); the shaded band marks the zone of minimal deviation between the two signals, confirming that the corrected record collapses onto the optically measured response across the full loading history. 60
4.5 DIC point distribution on the N-CONTROL beam at initial and damage stages. White circles identify the actuator zone (Point 1) and the mid-span tracking zone (Points 24–29). 64
4.6 DIC–machine consistency check on the N-CONTROL beam: force–displacement trajectories of the seven tracked points overlaid on the machine record. 65
4.7 Capability of the DIC point-tracking chain on the N-CONTROL mortar beam. 66
4.8 N-CONTROL: ODDC-corrected force–displacement response. 67
4.9 N-EMP-14 mm: ODDC-corrected force–displacement response. 68
4.10 N-EMP-18 mm: ODDC-corrected force–displacement response. 68
4.11 N-EMP-36 mm: ODDC-corrected force–displacement response. 69
4.12 N-WMP-30 mm: ODDC-corrected force–displacement response. 69
4.13 N-TFF-25 mm: ODDC-corrected force–displacement response. 70
4.14 N-SMP-12 mm: ODDC-corrected force–displacement response. 70
4.15 N-SMP-24 mm: ODDC-corrected force–displacement response. 71
4.16 N-SMP-50 mm: ODDC-corrected force–displacement response. 71
4.17 N-GF-8 mm: ODDC-corrected force–displacement response. 72
4.18 N-SF-35 mm: ODDC-corrected force–displacement response. 72

4.19 H-CONTROL: ODDC-corrected force–displacement response. 73
4.20 H-EMP-14 mm: ODDC-corrected force–displacement response. 73
4.21 H-EMP-18 mm: ODDC-corrected force–displacement response. 73
4.22 H-EMP-36 mm: ODDC-corrected force–displacement response. 74
4.23 H-WMP-30 mm: ODDC-corrected force–displacement response. . . . 74
4.24 H-TFF-25 mm: ODDC-corrected force–displacement response. 75
4.25 H-SMP-12 mm: ODDC-corrected force–displacement response. 75
4.26 H-SMP-24 mm: ODDC-corrected force–displacement response. 76
4.27 H-SMP-50 mm: ODDC-corrected force–displacement response. 76
4.28 H-GF-8 mm: ODDC-corrected force–displacement response. 76
4.29 H-SF-35 mm: ODDC-corrected force–displacement response. 77
4.30 N-SMP-12 mm-0.75: ODDC-corrected force–displacement response. . 79
4.31 N-SMP-12 mm-1.25: ODDC-corrected force–displacement response. . 79
4.32 N-SMP-12 mm-1.50: ODDC-corrected force–displacement response. . 79
4.33 N-SMP-24 mm-0.75: ODDC-corrected force–displacement response. . 80
4.34 N-SMP-24 mm-1.25: ODDC-corrected force–displacement response. . 80
4.35 N-SMP-24 mm-1.50: ODDC-corrected force–displacement response. . 81
4.36 H-SMP-12 mm-0.75: ODDC-corrected force–displacement response. . 81
4.37 H-SMP-12 mm-1.25: ODDC-corrected force–displacement response. . 82
4.38 H-SMP-12 mm-1.50: ODDC-corrected force–displacement response; stiffest beam of Case 2. 82
4.39 H-SMP-24 mm-0.75: ODDC-corrected force–displacement response. . 83
4.40 H-SMP-24 mm-1.25: ODDC-corrected force–displacement response. . 83
4.41 H-SMP-24 mm-1.50: ODDC-corrected force–displacement response. . 83
4.42 Case Study 3 capability plot on the CST-CONTROL cube: best-point DIC trajectory versus machine record. 85
4.43 CST-CONTROL: DIC-validated force–displacement record. 86
4.44 CST-EMP-50 mm: DIC-validated force–displacement record. 86
4.45 CST-WMP-50 mm: DIC-validated force–displacement record. 86
4.46 CST-SMP-50 mm: DIC-validated force–displacement record. 87
4.47 CST-SMP-24 mm: DIC-validated force–displacement record. 87
4.48 CST-HYB-50+24: DIC-validated force–displacement record; best agreement of the case study. 88
4.49 CST-HYB-EMP-50+24: raw DIC versus machine record retained for pre-peak use. 88
4.50 CST-HYB-WMP-50+24: DIC-validated force–displacement record. . 89

4.51 Case Study 4 capability plot on the BBS-CONTROL cube: best-point DIC trajectory versus machine record. 90
4.52 BBS-CONTROL: DIC-validated force–displacement record. 91
4.53 BBS-EMP-50 mm: DIC-validated force–displacement record. 91
4.54 BBS-WMP-50 mm: raw DIC record retained with a caveat. 91
4.55 BBS-SMP-50 mm: DIC-validated force–displacement record. 92
4.56 BBS-SMP-24 mm: DIC-validated force–displacement record. 92
4.57 BBS-HYB-50+24: raw DIC record retained with a caveat. 93
4.58 BBS-HYB-EMP-50+24: DIC-validated force–displacement record. . . 93
4.59 BBS-HYB-WMP-50+24: DIC-validated force–displacement record. . 94
4.60 Full-field view of the hooked-end-steel beam H-SF-35 mm at three stages of the test. 95
4.61 Post-hoc ladder of CMOD-based virtual extensometers on the N-CONTROL mortar beam. 96

5.1 Visual index of the four polypropylene families used in Case 1: EMP, WMP, SMP and TFF. The dimensions quoted are the manufacturer values; effective embedded lengths after casting are quoted in the companion matrix of Chapter 3. 113
5.2 Case 1: force–displacement curves of the Normal-matrix beams, corrected by ODDC, grouped by fibre programme. Thick black curve: N-CONTROL (three replicates). Peak markers: peak-force coordinates used in (5.1) and (5.4). 113
5.3 Case 1: force–displacement curves of the High-matrix beams, same grouping and markers as fig. 5.2. 114
5.4 Case 1: direct Normal-vs-High comparison. Solid lines and shaded envelopes are the replicate-averaged responses of each fibre programme. The matrix uplift of the High cohort is visible in every programme but is most pronounced for the SMP macro fibres. 114
5.5 Case 1: peak flexural strength σ_{flex} of the 22 beams grouped by fibre programme and matrix. Error bars = replicate standard deviation; colour = fibre family. 116
5.6 Case 1: total toughness T_{total} of the 22 beams, same layout as fig. 5.5. The SMP-50 mm macro fibres dominate the ductility axis. 116
5.7 Case 1: σ_{flex}–T_{total} aspect-ratio plot. Micro-fibre programmes cluster horizontally; SMP-50 mm migrates vertically in both matrices. 117

5.8 Case 1: ($\Delta F\%$, $\Delta T\%$) of every fibre programme relative to its matrix-matched control. The upper-right quadrant is the "strength and ductility" region. 117
5.9 Case 1: five-axis radar plot (peak force, peak displacement, toughness, first-crack force, residual force). Axes are rescaled to the cohort maximum. 118
5.10 Case 1: Pearson correlation heatmap over $\{F_{\text{peak}}, \sigma_{\text{flex}}, T_{\text{total}}, \delta_{\text{peak}}, L_f, \text{SP}\}$. Diagonal correlations are unity by construction. 118
5.11 Case 1: superplasticiser demand per programme. All 22 beams used the same PCE dose, which justifies holding SP out of the downstream regression models. 119
5.12 Case 1: compact non-parametric significance matrix. Shapiro–Wilk on the diagonal (not rejected if $p > 0.05$); Mann–Whitney pairwise tests off-diagonal. 119
5.13 Case 1: optimal programme versus matrix control, for both strength and toughness objectives. N-SMP-50 mm and H-SMP-50 mm dominate T_{total}; N-TFF-25 mm and H-EMP-14 mm dominate σ_{flex}. 120
5.14 Case 1: stacked-bar decomposition of the GWP of every mix into binder, sand, SP, water and fibre contributions. 121
5.15 Case 1: per-material bar chart of the cradle-to-gate contributions. The binder dominates with 92–95% of the GWP across all 22 mixes. . 122
5.16 Case 1: total GWP aggregated by fibre programme. 122
5.17 Case 1: ranked GWP penalty (Penalty% of (5.22)). Polypropylene fibres stay in the +2.9 to +4.2% band; GF and SF macro fibres incur +16% and +53% respectively. 122
5.18 Case 1: Normal-vs-High matrix comparison. The High matrix carries a +37.5% intrinsic GWP penalty before any fibre addition. 123
5.19 Case 1: Eco-Efficiency Index on flexural strength, EEI_σ. 124
5.20 Case 1: Eco-Efficiency Index on total toughness, EEI_T. 124
5.21 Case 1: EEI_σ–EEI_T trade-off map. Pareto-front points are ringed. . . 124
5.22 Case 1: tornado plot of the one-at-a-time sensitivity scan anchored on N-TFF-25 mm. Cement EF dominates by more than an order of magnitude. 125
5.23 Case 1: cost–benefit scatter, Penalty% versus $\Delta\sigma_{\text{flex}}\%$. Upper-left quadrant = efficient; lower-right = avoid. 126
5.24 Case 1: mix-level GWP benchmarked against recent fibre-reinforced concrete and mortar LCAs from the literature. 126

5.25 Case 1: combined LCA ranking R_{LCA} per (5.20). Lower rank = better compromise. 126
5.26 Case 2: visual dose index. Peak-coordinate migration with V_f at $L_f = 12$ mm and $L_f = 24$ mm, pooled over matrices. 129
5.27 Case 2: force–displacement curves, Normal matrix. ODDC-corrected, grouped by fibre length and dose. 130
5.28 Case 2: force–displacement curves, High matrix. Same layout as fig. 5.27. 130
5.29 Case 2: four-tile matrix/length/dose grid. Each tile collects the four dose levels at a fixed (matrix, length) cell. 131
5.30 Case 2: peak flexural strength per factorial cell. 132
5.31 Case 2: total toughness per factorial cell. 132
5.32 Case 2: dose–response curves $\sigma_{\text{flex}}(V_f)$ and $T_{\text{total}}(V_f)$ per matrix and length. Overlaid: 1-D quadratic fit of (5.17); dashed vertical line at $V_{f,\text{opt}}$. 133
5.33 Case 2: ($\Delta F\%$, $\Delta T\%$) relative to matrix control, per dose and length. 134
5.34 Case 2: five-axis radar plot over the 16 reinforced cells. 134
5.35 Case 2: Pearson heatmap over $\{F_{\text{peak}}, \sigma_{\text{flex}}, T_{\text{total}}, \delta_{\text{peak}}, V_f, L_f, \text{SP}\}$. . . 135
5.36 Case 2: SP demand as a function of V_f. The knee at $V_f = 1.25$–1.50% marks the fresh-state workability barrier. 135
5.37 Case 2: non-parametric significance panel. Shapiro–Wilk, Mann–Whitney 12/24, Kruskal–Wallis over four doses, Spearman ρ_s. 135
5.38 Case 2: length-by-dose interaction plot. Dashed vertical line at the 1-D quadratic optimum. 136
5.39 Case 2: optimal dose identified per matrix/length pair. 136
5.40 Case 2: cubic-spline interpolated contour of σ_{flex} and T_{total} over the observed (V_f, L_f) grid. Descriptive only, not a fitted response surface. 137
5.41 Case 2: marginal dose efficiency $\Delta T/\Delta V_f$ per factorial cell. 138
5.42 Case 2: performance heatmap; colour intensity encodes R_{combined} rank within the eight (matrix, length, dose) cells. 138
5.43 Case 2: cross-validation of the quadratic fit. Dashed line = Case 1 reading at the shared $V_f = 1.00\%/24$ mm operating point. 139
5.44 Case 2: stacked GWP per mixture. 140
5.45 Case 2: per-material contribution. Binder remains the dominant term. 140
5.46 Case 2: GWP aggregated by dose level. 140
5.47 Case 2: Penalty% ranked across the 16 reinforced cells. 141
5.48 Case 2: Normal-vs-High penalty bar at each dose. 141

5.49 Case 2: EEI_σ map over the dose/length grid. 142
5.50 Case 2: EEI_T map. 142
5.51 Case 2: EEI_σ–EEI_T trade-off map. Pareto cells are ringed. 142
5.52 Case 2: tornado plot anchored on the EEI_σ optimum. 143
5.53 Case 2: cost–benefit scatter. All reinforced cells sit in the upper-left quadrant. 143
5.54 Case 2: literature benchmark of the mix-level GWP. 144
5.55 Case 2: GWP versus V_f line plot. The two matrices are parallel lines. 144
5.56 Case 2: optimal-dose EEI_σ and EEI_T bar. The short and long fibres trade places as the objective switches from strength to toughness. . . 144
5.57 Case 2: combined LCA ranking R_{LCA} of the 16 reinforced cells. . . . 145
5.58 Case 3: visual index of the eight CST cubes (one control and seven fibre programmes). Single-family programmes are solid; hybrid 50+24 blends are hatched. 147
5.59 Case 3: force–displacement curves of the eight CST cubes. Thick black: CST-CONTROL. Single-family curves are solid; hybrid blends are dashed. 148
5.60 Case 3: splitting tensile strength f_t. 149
5.61 Case 3: pre-peak energy E_{pre}. 149
5.62 Case 3: $(\Delta f_t\%, \Delta E_{\mathrm{pre}}\%)$ versus the CST-CONTROL. 150
5.63 Case 3: single-family versus hybrid blend boxplot for f_t and E_{pre}. . . 150
5.64 Case 3: response versus mass-weighted mean fibre length. 151
5.65 Case 3: five-axis radar. 151
5.66 Case 3: Pearson heatmap. 152
5.67 Case 3: Shapiro–Wilk and Mann–Whitney panel. 152
5.68 Case 3: optimal programme (CST-WMP-50 mm) overlaid on CST-CONTROL. 153
5.69 Case 3: secant-stiffness degradation D of (5.12). 154
5.70 Case 3: bilinear model fit with K_1, K_2 and LOP of (5.10). 154
5.71 Case 3: pre-peak energy partition (elastic vs hardening fraction). . . . 154
5.72 Case 3: fibre-efficiency metric E_{pre}/V_f. 155
5.73 Case 3: stacked GWP contribution per mix. 156
5.74 Case 3: per-material GWP. 156
5.75 Case 3: GWP aggregated by fibre programme. 156
5.76 Case 3: Penalty% ranked across the seven reinforced mixes. 157
5.77 Case 3: EEI_{f_t}. 157
5.78 Case 3: $\mathrm{EEI}_{E_{\mathrm{pre}}}$. 157

5.79 Case 3: EEI_{f_t}–$\mathrm{EEI}_{E_{\mathrm{pre}}}$ trade-off map. 158
5.80 Case 3: tornado anchored on CST-SMP-50 mm. 159
5.81 Case 3: cost–benefit scatter. 159
5.82 Case 3: literature benchmark. 159
5.83 Case 3: $\mathrm{EEI}_{E_{\mathrm{pre}}}$ by fibre type. 160
5.84 Case 3: single-family versus hybrid EEI comparison. 160
5.85 Case 3: combined LCA rank R_{LCA}. 160
5.86 Case 4: visual index of the eight BBS cubes with embedded rebars. . 163
5.87 Case 4: ODDC-corrected force–displacement overlay of the eight BBS cubes. 164
5.88 Case 4: bond stress τ. 165
5.89 Case 4: pre-peak energy E_{pre}. 165
5.90 Case 4: $(\Delta\tau\%, \Delta E_{\mathrm{pre}}\%)$ versus BBS-CONTROL. 166
5.91 Case 4: single-family versus hybrid boxplot. 166
5.92 Case 4: response versus mass-weighted mean fibre length. 166
5.93 Case 4: five-axis radar. 167
5.94 Case 4: Pearson heatmap. 167
5.95 Case 4: Shapiro–Wilk and Mann–Whitney panel. 168
5.96 Case 4: optimal programme (BBS-WMP-50 mm) overlaid on BBS-CONTROL. 169
5.97 Case 4: stiffness-degradation D. 169
5.98 Case 4: bilinear model fit with K_1, K_2 and LOP. 170
5.99 Case 4: pre-peak energy partition. 170
5.100 Case 4: fibre-efficiency metric E_{pre}/V_f. 170
5.101 Case 4: τ–δ curve of the eight BBS cubes. 171
5.102 Case 4: CST-vs-BBS cross-test at the shared eight fibre programmes and $V_f = 1.00\%$. The ranking reversal of CST-EMP-50 mm and CST-SMP-24 mm between the two geometries is the main structural contribution of Case 4. 171
5.103 Case 4: stacked GWP per mix. 172
5.104 Case 4: per-material contribution. 172
5.105 Case 4: programme-level GWP totals. 172
5.106 Case 4: Penalty% ranked across the seven reinforced BBS mixes. . . . 173
5.107 Case 4: EEI_{τ}. 174
5.108 Case 4: $\mathrm{EEI}_{E_{\mathrm{pre}}}$. 174
5.109 Case 4: EEI_{τ}–$\mathrm{EEI}_{E_{\mathrm{pre}}}$ trade-off map. 175
5.110 Case 4: tornado anchored on BBS-WMP-50 mm. 175

5.111Case 4: cost–benefit scatter. 176
5.112Case 4: literature benchmark. 176
5.113Case 4: $\mathrm{EEI}_{E_{\mathrm{pre}}}$ by fibre type. 176
5.114Case 4: single-family versus hybrid EEI comparison. 177
5.115Case 4: CST-vs-BBS eco-efficiency scatter at the shared eight fibre programmes. 177
5.116Case 4: explicit ranking-reversal diagram between CST and BBS geometries. 178
5.117Case 4: combined LCA rank R_{LCA}. 178

6.1 DIC-based definition of CMOD and CTOD 184
6.2 Case 1 CPF evolution 188
6.3 Case 2 CPF evolution 189
6.4 Case 1 Normal-matrix kinematic fits — CONTROL and EMP triplet 191
6.5 Case 1 Normal-matrix kinematic fits — remaining fibre families . . . 192
6.6 Case 1 High-matrix kinematic fits — CONTROL and EMP triplet . . 195
6.7 Case 1 High-matrix kinematic fits — remaining fibre families 196
6.8 Case 2 kinematic fits — 12-mm fibres and controls 199
6.9 Case 2 kinematic fits — 24-mm fibres 202
6.10 Case 1 residual strengths bar chart 207
6.11 Case 1 fracture energy bar chart 208
6.12 Case 1 fibre-bridging energy versus CMOD — all fibres 208
6.13 Case 1 FEI ranking 209
6.14 Case 1 fibre-bridging energy averaged by fibre family 209
6.15 Case 1 $f_{R,3}/f_{R,1}$ ratio 210
6.16 Case 1 UNI 11039-2 equivalent strengths 210
6.17 Case 1 Normal-matrix F–CMOD with $f_{R,j}$ markers — CONTROL and EMP 211
6.18 Case 1 Normal-matrix F–CMOD with $f_{R,j}$ markers — WMP, TFF and SMP pair 212
6.19 Case 1 Normal-matrix F–CMOD with $f_{R,j}$ markers — SMP-50, GF and SF 213
6.20 Case 1 High-matrix F–CMOD with $f_{R,j}$ markers — CONTROL and EMP 214
6.21 Case 1 High-matrix F–CMOD with $f_{R,j}$ markers — WMP, TFF and SMP pair 215
6.22 Case 1 High-matrix F–CMOD with $f_{R,j}$ markers — SMP-50, GF and SF 216

6.23 Case 2 residual strengths bar chart . . . 216
6.24 Case 2 fracture energy bar chart . . . 217
6.25 Case 2 fibre-bridging energy versus CMOD — all specimens . . . 217
6.26 Case 2 FEI ranking . . . 218
6.27 Case 2 fibre-bridging energy by dose and length . . . 218
6.28 Case 2 $f_{R,3}/f_{R,1}$ ratio . . . 219
6.29 Case 2 UNI 11039-2 equivalent strengths . . . 219
6.30 Case 2 F–CMOD with markers — controls and N-SMP-12 low doses 220
6.31 Case 2 F–CMOD with markers — N-SMP-12 high, N-SMP-24 low . . 221
6.32 Case 2 F–CMOD with markers — N-SMP-24 high, H-SMP-12 low . . 222
6.33 Case 2 F–CMOD with markers — H-SMP-12 high, H-SMP-24 low . . 223
6.34 Case 2 F–CMOD with markers — H-SMP-24 high doses . . . 223
6.35 Case 2 dose response — fracture energy . . . 224
6.36 Case 2 dose response — $f_{R,1}$ . . . 225
6.37 Case 2 dose response — $f_{R,3}$ . . . 225
6.38 Case 2 dose response — FEI . . . 226

7.1 Case 1 matrix — Normal cement . . . 233
7.2 Case 1 matrix — High cement . . . 234
7.3 Case 1 R^2 comparison . . . 234
7.4 Case 1 fit — N-CONTROL . . . 235
7.5 Case 1 fit — N-TFF-25 . . . 235
7.6 Case 1 fit — N-SF-35 . . . 236
7.7 Case 2 matrix — SMP 12 mm . . . 236
7.8 Case 2 matrix — SMP 24 mm . . . 237
7.9 Case 2 fit — N-CONTROL . . . 237
7.10 Case 2 fit — representative SMP . . . 238
7.11 First group δ–CMOD — mean . . . 240
7.12 First group δ–CMOD — enhanced . . . 240
7.13 Second group δ–CMOD — mean . . . 241
7.14 Second group δ–CMOD — enhanced . . . 241
7.15 Second group PP vs. codes on δ–CMOD . . . 241

List of Tables

2.1 Comparison of European structural design frameworks for FRC. . . . 23
2.2 Mapping of thesis objectives to identified research gaps and corresponding research actions. 26

3.1 Fibre inventory used across the four case studies. 34
3.2 Nominal composition of the two FRM matrices used in Case Studies 1 and 2. 36
3.3 Case Study 1 – specimen matrix of the 22 FRM mixes at $V_f = 1\%$. . 37
3.4 Case Study 2 – specimen matrix of the 18 FRM mixes of the SMP dose response. 40
3.5 Nominal M30 concrete mix (fibre-free baseline). 41
3.6 Case Studies 3 and 4 – eight fibre groups of the M30 concrete programme. 42
3.7 CMOD-controlled loading-rate protocol for the notched FRM beams. 47
3.8 Selected literature on notched-beam flexural testing at various specimen scales and loading rates. 47
3.9 Final specimen inventory of the experimental programme. 52

4.1 Derivation of the machine stiffness K_{machine} from the N-CONTROL mortar beam. 58
4.2 DIC–machine consistency check for the N-CONTROL beam: R^2 of each tracked point against the machine record (RAW method, $N_{\text{valid}} = 9\,989$). 63
4.3 Case Study 1 cohort summary. 77
4.4 Case Study 2 cohort summary (sixteen new beams). 84
4.5 Case Study 3 (CST) cohort summary. 89
4.6 Case Study 4 (BBS) cohort summary. 94

5.1 Summary of the structural quantities used throughout chapter 5. All formulas are evaluated on the ODDC-corrected curves produced in chapter 4. 103
5.2 Emission factors used in the cradle-to-gate inventory. Best estimates are applied to the deterministic GWP calculation; triangular [min, max] ranges are propagated by Monte Carlo. 108
5.3 Chapter-wide LCA summary: every mixture across the four case studies with its mean GWP, penalty versus the matrix control, the dominant Eco-Efficiency Index for the relevant test and the combined LCA rank within the cohort. Semantic labels; lower R_{LCA} is better. 111
5.4 Case 1 structural summary: peak flexural strength, total toughness, behaviour classifier and combined structural rank R_{combined}. Mean of three replicates; standard deviations in CSV companion. 127
5.5 Case 1 environmental summary: mean cradle-to-gate GWP, Penalty% vs matrix control, Eco-Efficiency on σ_{flex} and on T_{total} and combined LCA rank. Lower rank = better compromise. 128
5.6 Case 2 structural summary: peak flexural strength and total toughness per factorial cell. 146
5.7 Case 2 environmental summary: mean GWP, Penalty% versus matrix control, Eco-Efficiency Indices and combined LCA rank. 146
5.8 Case 3 structural summary: f_t, E_{pre}, behaviour classifier. 161
5.9 Case 3 environmental summary: GWP, Penalty%, EEI_{f_t}, $\mathrm{EEI}_{E_{\mathrm{pre}}}$ and combined LCA rank. 162
5.10 Case 4 structural summary: bond stress τ and pre-peak energy E_{pre} per programme. 178
5.11 Case 4 environmental summary: GWP, Penalty%, EEI_{τ}, $\mathrm{EEI}_{E_{\mathrm{pre}}}$ and combined LCA rank. 179
5.12 Case 4 CST-vs-BBS cross-test: eight shared fibre programmes at $V_f = 1.00\%$ in the two geometries, compared on EEI_{f_t} (CST) and EEI_{τ} (BBS). 179

6.1 Specimen geometry constants used throughout chapter 6. 185
6.2 Case 1 kinematic coefficients 204
6.3 Case 2 kinematic coefficients. Same column structure as table 6.2; all eighteen specimens of the length × dose × matrix factorial are listed. 205

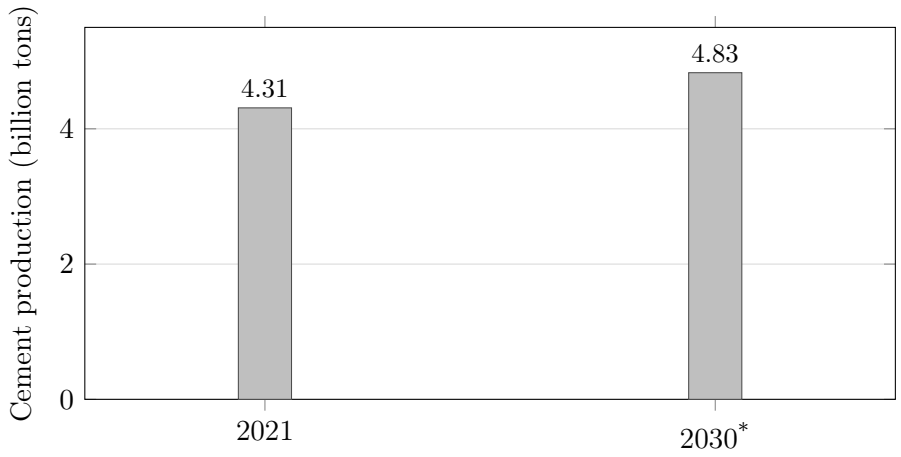
Chapter 1

Introduction

1.1 Background and Motivation

The global construction sector is subjected to severe environmental constraints, driven primarily by the environmental footprint of cementitious materials and traditional reinforcement paradigms. Concrete production is responsible for approximately 8% of global anthropogenic carbon dioxide emissions, with Portland cement clinker alone accounting for 90% to 95% of the total greenhouse gas burden associated with concrete manufacturing [1]. Global cement output, which was recorded at 4.31 billion tons in 2021, is projected to reach 4.83 billion tons by 2030 [2]. This sustained increase in cementitious output intensifies the demand for low-carbon binder alternatives and eco-efficient reinforcement strategies, as illustrated in fig. 1.1. The global warming potential of civil infrastructure is further amplified by traditional steel reinforcement; the unit impact of steel mesh has been recorded at 0.49 kg CO_2-eq/kg [3], rendering heavily reinforced configurations environmentally untenable. In response to these pressures, advanced binder modifications are being widely adopted. Limestone Calcined Clay Cement (LC3) and supplementary cementitious materials have demonstrated the capability to reduce the carbon footprint of concrete by up to 40% relative to ordinary Portland cement, without requiring major modifications to existing production infrastructure [4–6].

The utilisation of discrete fibre reinforcement in cementitious matrices has been investigated for over four decades [7], yet a strategic transition from energy-intensive steel fibres toward synthetic and glass alternatives has accelerated in recent years, driven by decarbonization imperatives and durability considerations. The fundamental structural contribution of fibres is governed by the crack-bridging mechanism: once the tensile capacity of the matrix is exceeded, fibres spanning the fracture plane transfer stresses through progressive debonding and frictional pull-out,



*Projected values based on industry growth estimates.

Figure 1.1: Global cement production: recorded output (2021) and projected output (2030) [2].

transforming an inherently brittle material into a quasi-ductile composite capable of significant energy absorption [8, 9]. Non-metallic fibres present critical advantages in this context, including absolute chemical inertness and corrosion resistance, reduced structural dead load, and lower embodied carbon per unit of reinforcement [10, 11]. Among synthetic variants, macro-polypropylene (PP) fibres exhibit a distinctive slip-hardening pull-out response; progressive surface abrasion during extraction increases interfacial friction, elevating the energy required for complete fibre withdrawal and thereby enhancing post-cracking toughness [12]. The geometric configuration of the fibre also plays a governing role: crimped, embossed, and waved profiles generate additional mechanical anchorage that augments the frictional bond [13]. Representative examples of polypropylene, glass, and steel fibres are illustrated in fig. 1.2. Certain limitations must nonetheless be acknowledged. PP fibres possess a significantly lower elastic modulus than steel, contribute minimally to pre-cracking stiffness, and are susceptible to time-dependent creep deformation under sustained loading [14, 15]. At volume fractions exceeding approximately 1.0%, severe loss of workability and fibre agglomeration have been reported, introducing void-induced weak zones within the hardened matrix [16].

Despite the environmental and durability advantages of non-metallic fibre-reinforced systems, their widespread structural implementation remains hindered by the absence of reliable mechanical characterisation at the constitutive level. Without precise quantification of post-cracking fracture energy and residual tensile strength, design codes are compelled to rely on covert understrength factors that fail to account for the pronounced size effects inherent to quasi-brittle failure [17, 18]. Current standardised testing methodologies compound this difficulty. Uniaxial tensile tests,

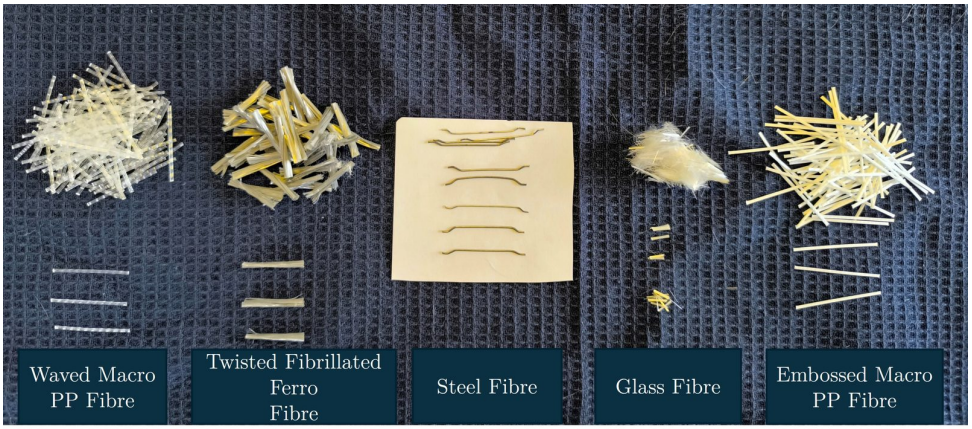


Figure 1.2: Examples of polypropylene, glass, and steel fibres used in cementitious composites.

while theoretically the most direct route to fracture parameters, are plagued by experimental instabilities including sudden load drops and snap-back phenomena. Standardised notched three-point bending tests, which constitute the prevalent characterisation method, introduce a concentrated point load at mid-span above a stress-concentrating notch, generating a disturbed region dominated by complex shear and vertical normal stresses that fundamentally violate the Bernoulli-Navier hypothesis [10, 19]. For highly ductile, strain-hardening composites, multiple micro-cracks frequently initiate outside the notch, rendering standard crack mouth opening displacement (CMOD) measurements unrepresentative of the true material deformation. Translating raw experimental data from such tests into reliable, size-independent constitutive laws therefore presents a formidable scientific challenge that must be resolved before non-metallic fibre-reinforced materials can be safely integrated into contemporary structural design practice.

1.2 Scientific Problem and Research Significance

In the evaluation of quasi-brittle materials via standardised three-point bending tests (e.g., DIN EN 14651), the reliance on raw actuator stroke is recognised as a fundamental systemic flaw. True specimen deflection is invariably overestimated when derived exclusively from actuator displacement, as this metric amalgamates the actual deformation of the concrete with extraneous elastic deformations of the testing apparatus. As illustrated schematically in fig. 1.3, the total actuator displacement represents a serial combination of the specimen deflection, the elastic compression of the loading frame, and the localised contact settlements at supports and loading rollers [20, 21]. If left uncorrected, these extraneous displacements critically distort the calculation of fracture energy (G_F) by artificially extending the

integration domain of the load-deflection curve. As depicted qualitatively in fig. 1.4, a comparison between the apparent load-deflection response derived from raw actuator stroke and the corrected response obtained via DIC-based virtual extensometry reveals a substantial difference; the shaded region (ΔG_F) quantifies the systematic overestimation of fracture energy caused by testing apparatus compliance and rigid-body motions. In testing configurations where residual-strength evaluation relies on mid-span deflection, actuator displacement, or stroke-derived CMOD rather than independently measured CMOD, machine compliance may bias the kinematic coordinate associated with f_{R1}–f_{R4}. In such cases, omission of compliance correction can lead to significant changes in the inferred residual-performance class, with shifts of up to two classes reported in the literature [22, 23]. When the magnitude of the post-peak negative softening stiffness exceeds the effective loading-frame stiffness, i.e. $|k_{\text{soft}}| > k_{\text{machine}}$, the specimen–machine system may lose stability under stroke control. This can trigger snap-back, truncate or obscure the descending branch, and prevent reliable constitutive identification [24, 25].

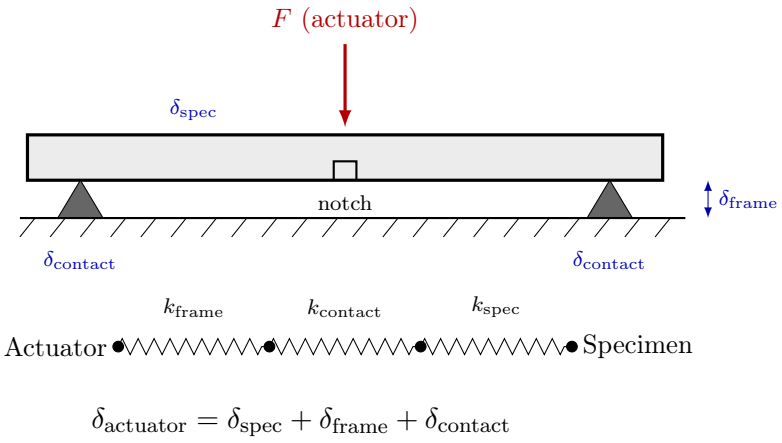


Figure 1.3: Sources of extraneous displacement in a standardised bending test.

To resolve these critical compliance errors, Digital Image Correlation (DIC) has emerged as an indispensable tool for virtual extensometry. By tracking subsets across the specimen surface with sub-pixel accuracy, DIC provides full-field optical data from which pure structural deflection is mathematically isolated by filtering out rigid-body motion [26]. In contrast to traditional localised sensors such as Linear Variable Displacement Transducers (LVDTs) or clip gauges, which capture only global structural responses at discrete points and are vulnerable to localised matrix heterogeneities, DIC yields a highly dense displacement field of the entire region of interest [27, 28]. This continuous spatial mapping enables the precise extraction of Crack Mouth Opening Displacement (CMOD) and Crack Tip Opening Displacement (CTOD) across the complex Fracture Process Zone (FPZ). Integral fracture

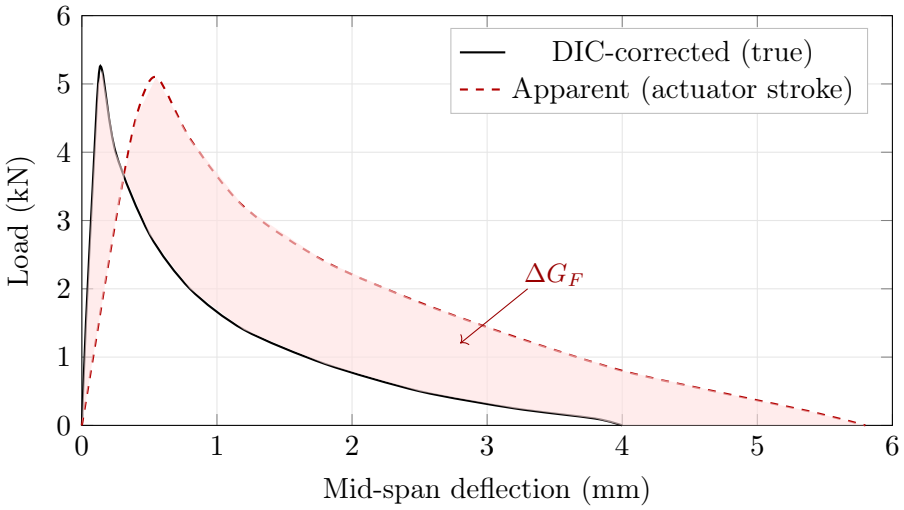


Figure 1.4: Schematic comparison of apparent (raw stroke) and DIC-corrected load–deflection responses (illustrative only).

quantities, such as the work of fracture, can therefore be determined directly from the corrected kinematic data. Testing apparatus compliance is bypassed physically because DIC tracks subsets directly on the specimen surface rather than measuring actuator stroke. Deriving the complete stress-crack opening (σ-w) constitutive law, which underpins code-level structural design, still strictly requires iterative inverse analysis, as the stress distribution across the cracked ligament cannot be measured optically [26, 29, 30]. In the present work, the inverse analysis of Chapter 7 is driven by DIC-derived CMOD/CTOD signals rather than clip-gauge CMOD; this choice is motivated both by the absence of clip-gauge instrumentation in the splitting and bonded-bar configurations (Chapters 4–5) and by the full-field validation of the rigid-hinge assumption (Chapter 6).

The necessity of this inverse-analysis step exposes a second, equally critical gap in current code-oriented constitutive modelling. The conceptual foundations of existing structural design codes, such as the *fib* Model Code 2010 [31] and RILEM TC 162-TDF [32], were predominantly derived from historical experimental databases of hooked-end steel fibres. The cohesive crack model, originally proposed by Hillerborg et al. [33] and subsequently extended through inverse analysis methodologies [34, 35], provides the theoretical basis for translating test data into stress-crack opening relationships. When uncorrected global deflections are utilised in such inverse procedures, the resulting trilinear σ-w laws fail to capture the true post-cracking response of the material [36–38]. This deficiency is particularly acute for synthetic reinforcements: polypropylene and glass fibres exhibit fundamentally different micro-mechanical pull-out behaviour compared to steel, typically characterised by an abrupt post-cracking strength drop followed by a secondary slip-hardening load

peak [39]. Existing bilinear code models, calibrated against steel fibre data, consequently misrepresent the residual strength evolution of synthetic fibre-reinforced concrete. Establishing a unified, DIC-validated workflow that connects raw experimental testing, accurate kinematic extraction, and robust numerical inverse analysis is therefore scientifically imperative to generate tailored multi-linear cohesive crack models for non-metallic composites.

1.3 Research Scope and Delimitations

The primary scope of this research is strictly delimited to the mechanical, kinematic, and environmental evaluation of non-metallic fibre-reinforced cementitious composites. To establish a rigorous comparative baseline, the experimental campaign evaluates plain control matrices and metallic (steel) fibre-reinforced systems against a broad extension group of non-metallic alternatives, including glass fibres and various geometric configurations of polypropylene fibres (e.g., embossed macro, waved macro, twisted fibrillated, and synthetic micro fibres). The scope encompasses the optimisation of these mixtures by varying fibre volumetric dosages to quantify the resulting impact on residual post-cracking strength. The environmental assessment is restricted to a cradle-to-gate Life Cycle Assessment targeting Global Warming Potential (GWP) as the primary impact category, thereby linking the mechanical performance evaluation directly to the sustainability motivation outlined in section 1.1.

The experimental boundaries of this study are restricted to specific standardised testing configurations designed to trigger distinct failure mechanisms. Notched three-point bending tests on prismatic beams, conducted in accordance with EN 14651 [40], are utilised as the primary methodology to extract flexural fracture kinematics. To expand the mechanical investigation beyond pure flexure, the scope also includes cube splitting tensile tests (for evaluating indirect tensile toughness and multidirectional fibre distribution) and bonded-bar splitting tests (for quantifying interfacial bond-slip mechanics and matrix confinement). Digital Image Correlation is employed across all test configurations as a virtual extensometer to bypass testing machine compliance and accurately extract pure structural deflection, CMOD, and CTOD. The numerical scope is delimited to deriving trilinear constitutive σ-w laws via inverse analysis, within the theoretical framework of the fictitious crack model [31, 33]. While the flexural beam specimens undergo the complete analytical workflow (including fracture kinematics, inverse analysis, and code derivation), the indirect tensile and bond-slip specimens are delimited exclusively to optical displace-

ment correction, statistical inference, and life cycle assessment.

Based on the foregoing delimitations, the specific research objectives of this thesis are formulated as follows:

1. To develop and validate a DIC-based displacement correction protocol that eliminates testing machine compliance and rigid-body motion artifacts from standardised flexural tests (chapter 4).

2. To quantify the structural performance of non-metallic fibre-reinforced composites through statistical significance analysis and to couple these results with cradle-to-gate life cycle assessments, generating a unified eco-efficiency evaluation (chapter 5).

3. To map the full fracture kinematics, including CMOD, CTOD, and crack profile evolution, and to validate the rigid-hinge assumption for non-metallic fibre composites (chapter 6).

4. To derive true trilinear stress–crack opening (σ–w) constitutive laws via numerical inverse analysis, and to connect DIC-based kinematics with European code conventions (chapter 7).

The analytical workflow connecting these objectives progresses from raw experimental testing through DIC-based displacement correction, encompassing parallel tracks of statistical and environmental analysis and advanced fracture kinematics. These kinematics feed directly into numerical inverse analysis, ultimately concluding with European code evaluation. This overall structure is illustrated in fig. 1.5.

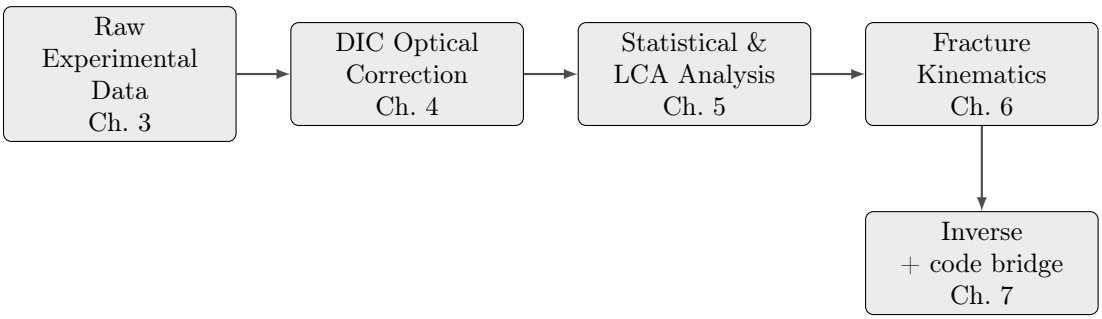


Figure 1.5: Analytical workflow of the thesis.

1.4 Originality and Main Contributions

The originality of this thesis is grounded in the systematic resolution of compliance errors inherent to standardised quasi-brittle fracture testing, coupled with the

specialised adaptation of structural design codes for synthetic reinforcements. The main contributions to the field of structural engineering and construction materials are summarised as follows:

- **DIC-Validated Displacement Correction Workflow:** A full-field Digital Image Correlation protocol is developed and validated to mathematically isolate true structural deflection from raw actuator stroke, substantially reducing and explicitly quantifying the influence of machine compliance and rigid-body motion within the adopted experimental setup.
- **Integrated Eco-Mechanical Profiling:** A unified analytical framework is established that couples non-parametric significance testing with cradle-to-gate Life Cycle Assessments, facilitating the generation of an Eco-Efficiency Index for the direct quantification of structural performance trade-offs against Global Warming Potential.
- **Advanced Fracture Kinematic Mapping:** The rigid-hinge assumption is empirically validated for non-metallic composites through the precise mapping of Crack Mouth and Crack Tip Opening Displacements, confirming the linear kinematic relationship required for accurate fracture energy partitioning.
- **Constitutive Derivation and Code Adaptation:** True trilinear cohesive crack models are derived via numerical inverse analysis. The framework is first validated with metallic fibres against existing European standards (*fib* Model Code 2010 [31]), after which code divergences for synthetic fibres are identified and adjustments to structural stiffness evaluations are proposed.

1.5 Thesis Organisation

This thesis is systematically structured to reflect the progressive experimental and analytical workflow. Following this introductory chapter, **chapter 2** provides a comprehensive state-of-the-art review, detailing the transition toward sustainable composites, the mechanisms of fracture testing, the role of optical metrology, and the current limitations of European design codes, ultimately defining the primary research questions. **chapter 3** establishes the global experimental and methodological framework, detailing the materials, mix designs, and the specific configurations of the primary testing groups.

The subsequent chapters present the progressive analytical methodology. **chapter 4** establishes the displacement correction protocol and optical validation using

Digital Image Correlation. **chapter 5** presents the structural evaluation, statistical significance analysis, and the life cycle eco-efficiency assessments. **chapter 6** explores the advanced fracture kinematics, focusing on crack opening displacement mapping and validation of the rigid-hinge assumption. **chapter 7** presents the numerical inverse analysis, the identification of trilinear σ–w laws for both testing campaigns, and the discussion of how optical CMOD/CTOD measurements connect to European code-based CMOD formulations for the present polypropylene systems. **chapter 8** synthesises the global conclusions, highlights the implications for sustainable structural design, and recommends trajectories for future research.

Chapter 2

State of the Art

2.1 Fibre-Reinforced Cementitious Composites and Environmental Considerations

Plain cementitious matrices remain weak in tension and brittle in fracture; fibres are therefore introduced primarily to improve *tensile cracking resistance* and post-cracking ductility across structural and non-structural applications [41]. Contemporary practice also pushes toward higher mechanical performance (e.g. high-strength matrices) while facing growing pressure to reduce environmental burdens associated with binder production and with the overall material system [3]. In the literature reviewed here, that tension between mechanical demand and environmental constraint is made explicit: ordinary Portland cement (OPC) production is linked to a large share of global CO_2 emissions and to substantial electrical energy use per tonne, which motivates binder optimisation and wider use of supplementary cementitious materials (SCMs) such as fly ash where mechanically admissible [1].

2.1.1 Material Properties of Steel, Glass, and Macro-Synthetic Fibres

Across mortar and concrete families (normal strength, high strength, and high strength with fly ash), comparative studies systematically contrast *steel*, *glass*, and *synthetic* (short and long polypropylene) fibres at several volume fractions, with explicit attention to post-cracking response and to the validation of optical measurement techniques alongside conventional machine records [41]. At the level of material taxonomy, fibre-reinforced concrete (FRC) is defined as a cement–aggregate matrix with discrete discontinuous fibres; steel fibres remain dominant for many struc-

tural uses, whereas synthetic fibres such as polypropylene and nylon are frequently associated with control of early-age cracking mechanisms in slabs, while higher-performance synthetic systems can also contribute to toughness in more demanding matrices [42]. Industrial uptake of FRC is further framed by the expanding role of design guidance: fibre-reinforced concrete is described as moving beyond traditional non-structural niches toward applications with increasing structural responsibility, in some cases with fibres acting as the sole reinforcement, which raises the stakes for reliable constitutive description, fibre orientation effects, and time-dependent behaviour such as creep [36].

2.1.2 Structural performance demands versus environmental assessment

Recent mortar-scale work on synthetic (polypropylene) reinforcement links mechanical outcomes to a complementary *cradle-to-gate* life-cycle assessment (LCA) perspective: system boundaries are stated to cover raw materials and transport to the laboratory scale, and midpoint-style indicators are discussed for human-health and environmental impact categories alongside global warming. Within that framing, the same study reports that normal-strength mortar can present lower human-health and environmental impacts than higher-strength alternatives in the compared matrix set, that a fly-ash–bearing high-strength mixture can reduce ecological impacts relative to a fly-ash-free high-strength counterpart, and that adding the studied synthetic fibres did not appear to impose a large marginal environmental or health burden relative to the reference mixtures while still contributing to mechanical objectives [43]. In a related investigation, Ibrahim et al. reported that synthetic fibre was found *more efficient and sustainable for construction than steel fibre* under the stated cradle-to-gate LCA assumptions [43]. This comparative claim must be interpreted with explicit caveats on the functional unit, system boundary and regional electricity-grid emission factors adopted in that study.

2.2 Digital Image Correlation and Virtual Extensometry

To capture the complex fracture behaviour of cementitious composites without mechanically interfering with the specimen, non-contact optical techniques are increasingly favoured over traditional instrumentation. Digital Image Correlation (DIC) has emerged as a robust, full-field methodology that tracks surface deformation by

mathematically comparing digital images captured before and during the application of load [44, 45].

2.2.1 Optical Deformation Measurement Principles

Two-dimensional DIC operates by dividing a reference image into a grid of small rectangular areas known as facets (or interrogation subsets). The algorithm tracks deformation by matching these facets from the reference image to their displaced positions in subsequent, deformed images. This process fundamentally relies on the surface possessing a random, non-periodic Speckle pattern with high information content, allowing the algorithm to uniquely identify each facet's fingerprint [46, 47].

Because materials undergo complex local deformations such as tension, shear, or rotation, an initially square reference facet will inevitably distort. This physical reality is addressed by incorporating shape functions (typically zero, first, or second-order polynomials) that mathematically map the reference coordinates into the deformed coordinates to prevent the correlation from failing [48].

However, the application of facet tracking involves a critical analytical trade-off. A larger facet size acts as a low-pass spatial filter; while it provides more unique pattern information to reduce noise, it simultaneously smooths the underlying displacement field. This filtering effect can introduce bias by underestimating peak strains at sharp strain gradients, such as those present at a crack tip. To limit this filtering effect and maintain high spatial resolution, the facet size must be minimized, which strictly requires the application of a sufficiently dense and finely distributed speckle pattern [49]. Figure 2.1 illustrates the fundamental mapping of a reference facet intensity pattern $f(\mathbf{x})$ to its deformed state $g(\mathbf{x}')$ through the displacement parameter vector $\mathbf{p}$, where the coordinate transformation is governed by the shape function $\chi(\mathbf{x}, \mathbf{p})$ [46, 48].

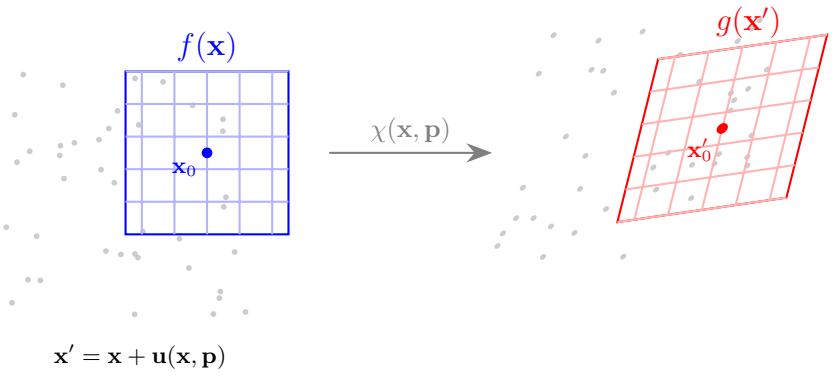


Figure 2.1: Mathematical mapping between reference and deformed facets.

2.2.2 Virtual extensometry in fracture testing

In addition to generating full-field strain maps, DIC software (such as GOM Correlate 2D) allows for the strategic placement of numerical, or "virtual", extensometers anywhere on the processed digital facets [44]. In the context of notched beam tests for fibre-reinforced composites, this capability is particularly valuable. Horizontal virtual extensometers can be strategically placed precisely at the base of the notch to measure Crack Mouth Opening Displacement (CMOD), while vertical virtual extensometers are simultaneously deployed to track mid-span deflection [47].

This approach avoids the physical stress concentrations and mounting difficulties associated with traditional clip gauges or Linear Variable Differential Transformers (LVDTs). Experimental validations across various mortar classifications (including normal strength, high strength, and fly-ash modified) have demonstrated a strong correlation between the load-deflection curves obtained from conventional testing machine sensors and those extracted via DIC virtual extensometers [44, 45]. The reliability of these optical measurements forms the methodological basis for accurately tracking the post-cracking ductility imparted by both steel and synthetic fibre reinforcement.

2.3 Machine Compliance and Displacement Measurement Challenges

In the mechanical characterisation of quasi-brittle materials, accurately measuring the post-cracking response is paramount. However, a fundamental limitation in standard testing is that the displacement recorded by the machine actuator (δ_{actuator}) does not represent the pure deformation of the specimen (δ_{specimen}). Rather, the testing setup behaves as a system of springs in series, where the recorded value inherently includes the elastic deformation of the testing machine itself (δ_{machine}). As summarised in fig. 2.2, the total recorded actuator stroke δ_{act} is a serial summation of the frame/contact compliance (K_m) and the true specimen deformation (K_s) [50, 51].

2.3.1 Fundamentals of Testing Machine Compliance

The extraneous deformation that corrupts crosshead displacement readings originates from multiple interacting factors within the load path. Primary sources include the bending flexure of the load frame's cross beams, the elastic stretching of

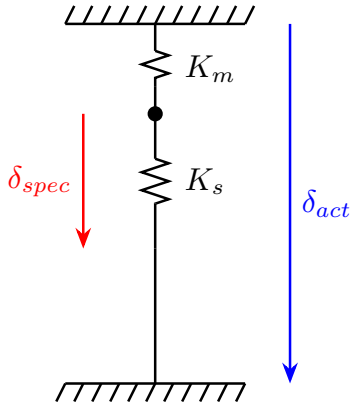


Figure 2.2: Serial spring representation of testing machine and specimen compliance.

the side columns, and internal load cell deformation. Significant compliance is also introduced at the specimen-fixture interfaces; high surface pressure causes localized crushing or indentation at the loading noses and supports, while microscopic voids and imperfect seating contribute to initial slack [50].

The effective stiffness of the machine (K_{machine}) is not a fixed, universal property of the hardware. The literature establishes that machine compliance is highly dependent on the specific test configuration [22, 24]. The total compliance is an aggregation of frame stiffness, contact stiffness (which varies drastically between Hertzian roller contacts and flat platens), and the geometry of the fixtures. Consequently, a compliance parameter calibrated for one specific configuration, such as three-point bending, cannot be universally applied to fundamentally different load paths, such as uniaxial or diagonal compression [51].

2.3.2 Impact on Fracture Parameter Evaluation

Ignoring system compliance and relying solely on raw actuator displacement introduces severe analytical errors. Because the flexibility of the test frame is incorrectly absorbed into the specimen's measured response, the displacement is artificially inflated. This results in the material appearing significantly softer than reality, severely underestimating the pre-peak stiffness (elastic modulus) [50].

In the post-cracking regime, uncorrected displacement data corrupts the calculation of fracture energy, as the area under the load-deflection curve parasitically includes the elastic strain energy stored in the testing machine during loading. A highly compliant testing machine can trigger unstable failure; if the machine stores excessive elastic energy, it may release violently post-peak, truncating the descending branch of the curve. Consequently, deriving accurate stress-crack opening laws becomes impossible without isolating the specimen's true constitutive behaviour, as

uncorrected data will artificially flatten the strain-softening curve and misrepresent the actual fibre-bridging mechanics [51].

2.3.3 Compliance Correction Methodologies

The conventional mechanical solution to actuator-stroke compliance in standardised bending tests is a dedicated yoke fitted with an LVDT that measures mid-span deflection relative to reference points on the specimen itself rather than from crosshead travel. In the procedure prescribed by EN 14651 [40], this yoke is combined with a clip gauge for CMOD measurement, and together they constitute the reference baseline for compliance mitigation in notched-beam testing. In the present programme, the adopted specimen geometry (40 mm depth), the notch configuration, and the requirement to resolve multi-point fracture kinematics make single-channel yoke measurement insufficient on its own. DIC is therefore not presented as a replacement for the yoke but as a complementary non-contact framework that enables verification of crack localisation at the notch, simultaneous CMOD/CTOD extraction, post-hoc placement of virtual extensometers, and full-field validation of the kinematic assumptions underlying the inverse-analysis framework of Chapter 7.

To mitigate these errors, modern testing protocols emphasize circumventing the load path entirely. As discussed in Section 2.2, the use of direct, specimen-mounted instruments, such as optical Digital Image Correlation (DIC) or physical clip gauges, allows researchers to record true deformation independent of machine stiffness.

When direct specimen instrumentation is unavailable or insufficient, rigorous mathematical compliance corrections must be applied to the actuator data. The standard experimental approach involves running a highly stiff "dummy" specimen, typically a steel calibration beam, to empirically isolate the machine's specific load-displacement curve. By determining the characteristic stiffness of the system in that exact configuration, the machine's displacement contribution can be mathematically subtracted from the total recorded extension, yielding the corrected, true specimen deflection [52].

2.4 Fracture Kinematics (CMOD/CTOD) and Rigid-Hinge Models

The failure of quasi-brittle materials in tension is fundamentally governed by heterogeneous microstructural properties rather than classical linear elasticity. Prior to the formation of a traction-free macroscopic crack, a localized, inelastic damage

region known as the Fracture Process Zone (FPZ) develops ahead of the crack tip [53]. The substantial size of this zone invalidates Linear Elastic Fracture Mechanics (LEFM), necessitating non-linear approaches centered on the specific fracture energy (G_f) dissipated during the softening phase [53].

In fibre-reinforced concrete, the introduction of discrete fibres fundamentally alters the FPZ. The fibres provide bridging stresses across both microcracks and macrocracks. The frictional pull-out resistance of these fibres exerts significant crack-closing forces, which drastically increases the total fracture energy, alters the crack propagation kinematics, and enhances the overall toughness of the composite system.

2.4.1 The non-linear cracked hinge model

To analytically model the complex post-cracking behaviour of notched beams without resorting to computationally expensive finite element analysis, the literature frequently adopts the non-linear cracked hinge concept, formalized through the Ostergaard framework [42, 54]. The foundational assumption of the hinge model is that all non-linear material behaviour, including the softening induced by the FPZ and the subsequent fibre-bridging mechanisms, is entirely confined within a localized hinge region of a finite, characteristic width [42]. The bulk of the beam outside this specific band is assumed to remain perfectly linear elastic.

Kinematically, the model dictates that at large deflections, the two halves of a notched beam act as rigid blocks. These blocks rotate around a distinct, singular "centre of rotation" (or neutral axis) located within the uncracked compressive zone at the top of the beam. This rigid-body assumption drastically simplifies the structural geometry, allowing the post-peak flexural response to be defined purely by the angular rotation (θ) of the hinge, as illustrated in fig. 2.3. By assuming rigid-body rotation about the neutral axis (x_{rot}), a strict geometric proportionality is established between the CMOD and the true crack tip opening displacement (CTOD).

2.4.2 Kinematic Relations in Three-Point Bending

Because the specimen is assumed to deform as rigid blocks rotating around a singular neutral axis, strict geometric relationships can be established between local fracture descriptors and global structural response. The overall angular deformation of the hinge (θ) is directly correlated to the localized crack opening angle (θ^*) [54].

Consequently, the Crack Mouth Opening Displacement (CMOD) is kinemati-

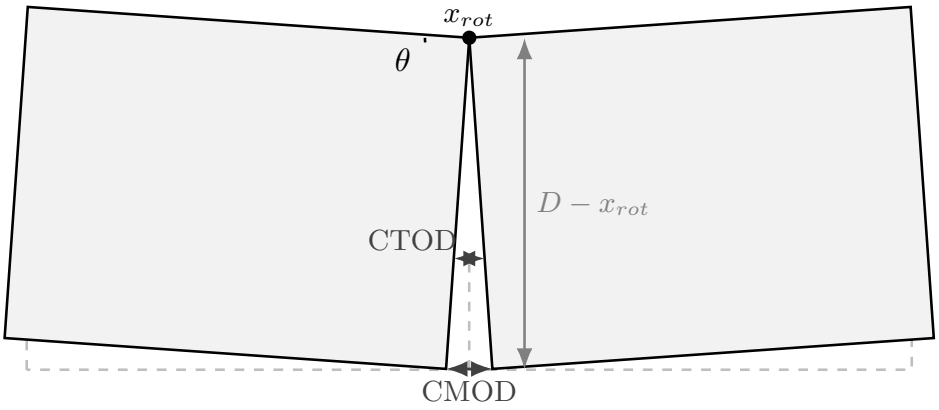


Figure 2.3: Kinematic model of rigid-body rotation in a notched beam.

cally linked to the distance from the neutral axis (the apparent depth of rotation) to the measurement point at the tension face [42]. Using simple trigonometry derived from this rigid-body rotation, the Crack Tip Opening Displacement (CTOD), representing the true opening precisely at the notch tip, can be calculated based on its proportional distance to the neutral axis relative to the measured CMOD [42, 54]. Establishing these precise kinematic relationships is essential for experimentalists, as it enables the translation of standard flexural test data (load versus CMOD or deflection) into the fundamental stress-crack opening ($\sigma - w$) constitutive laws that govern the material's structural performance.

2.5 Statistical Frameworks and Life Cycle Assessment (LCA)

The inherently heterogeneous nature of cementitious matrices, combined with the random orientation and distribution of discrete fibres, introduces significant scatter into experimental data. Consequently, robust statistical frameworks are relied upon to systematically validate improvements in mechanical performance and to separate true material behaviour from experimental noise. Prior to applying parametric inferential statistics, the distribution of the dataset must be evaluated; tools such as the Shapiro-Wilk normality test are critical for determining whether data follows a normal distribution, thereby dictating the application of parametric or non-parametric statistical methods [55].

Once distribution assumptions are validated, Analysis of Variance (ANOVA) is systematically employed to assess the statistical significance of mix constituents, such as fibre geometry, dosage, and matrix strength, on post-cracking performance metrics [56]. Advanced mix optimization frequently utilizes Response Surface Method-

ology (RSM). By moving beyond simple one-variable-at-a-time testing, RSM identifies non-linear synergies between variables (e.g., hybrid fibres and supplementary binders), mapping optimal compositions to achieve targeted mechanical properties [57].

Experimental campaigns on fibre-reinforced cementitious composites often involve limited replication because each fibre type, dosage, matrix and test geometry requires substantial material preparation and testing effort. As a result, cell sizes are frequently too small to support strong Gaussian assumptions or fully parametric factorial inference. In such cases, normality screening should be interpreted cautiously, and rank-based or non-parametric tools such as Mann–Whitney U, Kruskal–Wallis and Spearman correlation provide a more conservative framework for detecting systematic trends, comparing fibre families and assessing dose–response behaviour. Parametric ANOVA and Response Surface Methodology remain valuable in larger and more densely replicated optimisation campaigns, but in small exploratory FRC/FRM studies they should be used with caution or replaced by structured screening, effect-size interpretation, ranking consistency and physically informed trend analysis.

2.5.1 Life Cycle Assessment (LCA) framework

Parallel to mechanical optimization, the environmental sustainability of cementitious composites is increasingly scrutinized using Life Cycle Assessment (LCA). In accordance with the foundational ISO 14040 and ISO 14044 standards, a common boundary condition adopted in materials research is the cradle-to-gate approach [58, 59]. This specific boundary encompasses raw material extraction, processing, and transportation to the laboratory or factory gate, while explicitly excluding the use-phase and end-of-life disposal [58].

Within this analytical framework, Global Warming Potential (GWP), typically expressed in kg CO_2 equivalent, serves as the primary indicator used to quantify climate change impacts [60]. In standard concrete mixtures, GWP is overwhelmingly dominated by the energy-intensive calcination of limestone required during Portland cement production [60]. Figure 2.4 defines the cradle-to-gate system boundaries adopted for this study, encompassing all processes from raw material extraction up to the point of laboratory fabrication [6, 58].

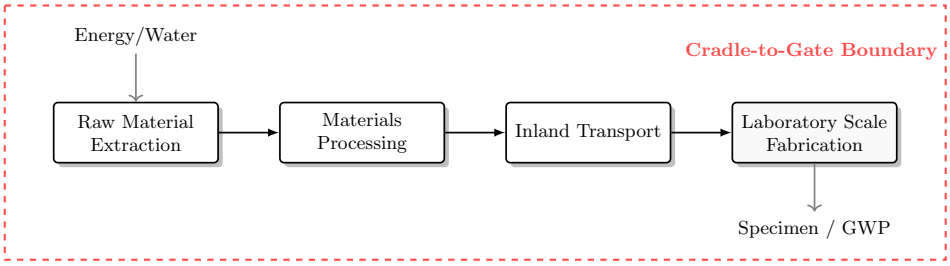


Figure 2.4: System boundaries for the cradle-to-gate environmental assessment.

2.5.2 Structural Performance-Based Eco-Efficiency

Relying solely on volumetric environmental impact (e.g., CO_2 equivalent per cubic meter) is often deemed insufficient for high-performance materials, as it penalizes the dense constituent profiles of advanced composites without recognizing their enhanced structural capacity. Eco-efficiency provides a more rigorous evaluation by coupling environmental burdens directly with functional mechanical performance, such as compressive strength, flexural toughness, or residual load capacity.

A standard metric utilized in this context is the Binder Intensity (bi) index [60, 61], which measures the mass of binder required to achieve a single unit of performance (e.g., 1 MPa of compressive strength). Lower bi values denote highly eco-efficient mixtures where structural competence is achieved with minimized carbon-intensive binder usage. By substituting Portland cement with supplementary cementitious materials and leveraging the advanced post-cracking control of fibre reinforcement, optimized material solutions can be targeted that simultaneously offer high mechanical resilience and low ecological impact.

2.6 Numerical Inverse Analysis and Constitutive Identification

Determining the fundamental tensile stress-crack opening ($\sigma - w$) relationship of cementitious composites via direct tension tests is notoriously difficult, primarily due to the stringent requirement for rigid boundary conditions to prevent secondary bending [62]. Consequently, the literature heavily favours indirect determination using numerical inverse analysis applied to macroscopic flexural tests, such as the three-point bending test (3PBT) or the wedge-splitting test (WST) [62, 63].

This inverse methodology relies fundamentally on Hillerborg's Fictitious Crack Model (FCM) [33]. The FCM posits that within the Fracture Process Zone (FPZ), concrete softens progressively due to microcracking but retains the capacity to trans-

fer stress across the “fictitious” crack faces. This residual stress transfer capacity (σ) is mathematically defined as a direct function of the crack opening displacement (w) [30]. When implemented in finite element simulations, this is frequently coupled with the crack-band theory proposed by Bažant and Oh [64], which assumes that microcracking is uniformly distributed over a specific characteristic band width. This approach "smears" the cracking behaviour, thereby preventing spurious mesh sensitivity during numerical optimization.

2.6.1 Constitutive Models

While bilinear stress-crack opening models are widely adopted by structural design codes and are generally sufficient for plain concrete or standard steel fibre-reinforced systems, they often fail to capture the complex pseudo-strain-hardening response characteristic of macro-synthetic fibres.

To accurately represent this behaviour, trilinear or multi-linear models provide a significantly more refined representation [39, 65]. In a trilinear model, the initial steep descending branch captures the energy dissipated by matrix micro-cracking and aggregate interlock. The subsequent branches accurately bound the delayed energy consumption mechanisms specific to polymeric fibres, namely chemical debonding and the prolonged, slip-hardening frictional pull-out. By separating these mechanisms into distinct zones, multi-linear models prevent the overestimation or underestimation of post-cracking residual strength that plagues simpler bilinear approximations.

2.6.2 Non-linear Least Squares Optimization

The mathematical process of inverse analysis is formulated as a non-linear least squares optimization problem. The primary objective is to minimize the error, calculated as the sum of squared residuals, between the experimentally measured macroscopic load-CMOD curve and the theoretical curve predicted by a numerical or analytical model, such as the cracked hinge model based on the Ostergaard framework [42, 62].

To navigate complex solution spaces and avoid convergence upon local minima, various optimization strategies are employed. These range from gradient-based methods (e.g., Levenberg-Marquardt) to evolutionary stochastic algorithms. A highly effective strategy detailed in the literature is the stepwise or phased approach, as illustrated in fig. 2.5. In a phased methodology, specific parameters of the $\sigma - w$ curve (such as the initial slope versus the tail slope) are isolated and

optimized sequentially across corresponding regions of the macroscopic load-CMOD curve to ensure robust convergence [54, 63].

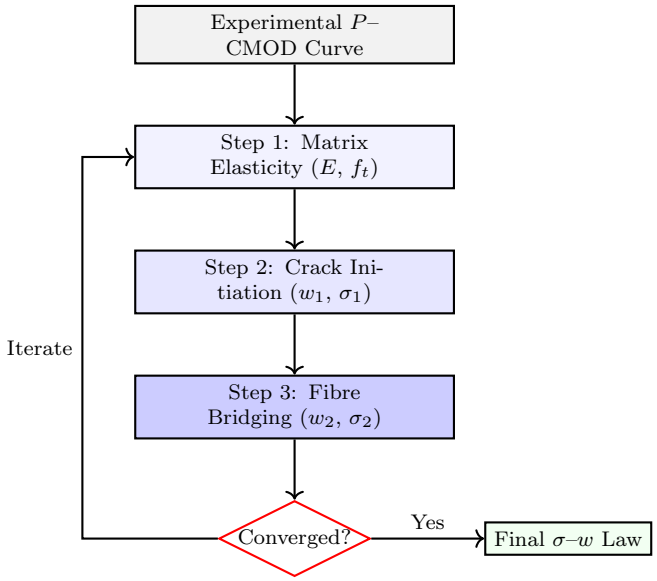


Figure 2.5: Phased workflow for the numerical inverse analysis.

2.7 European Standardization Frameworks for FRC

The current European framework for the structural design of Fibre-Reinforced Concrete (FRC) relies heavily on the *fib* Model Code 2010 (MC2010), which formally recognized FRC as a new composite material suitable for structural applications. However, the broader landscape of European standardization also includes distinct national guidelines that propose varying constitutive equations, most notably the RILEM TC 162-TDF recommendations, the Spanish code EHE-08, and the Italian guidelines CNR-DT 204 [21].

To evaluate the post-cracking properties required by these design codes, the European standard EN 14651 is universally employed. This standard mandates a three-point bending test on a notched prismatic specimen. Assuming a linear stress distribution across the cracked section, the standard uses classical beam theory to convert the recorded load into a nominal residual flexural tensile strength ($f_{R,j}$) at specific crack opening displacements [16]. Table 2.1 summarises how the principal documents relate to that experimental basis and to macroscopic post-cracking idealisations.

MC2010 classifies FRC structural performance based on these characteristic residual strengths evaluated at two critical, strict Crack Mouth Opening Displacement (CMOD) thresholds. The parameter f_{R1} is recorded exactly at a CMOD of

Table 2.1: Comparison of European structural design frameworks for FRC.

Document	Role of EN 14651 / $f_{R,j}$	Characterisation and reported limitations
fib Model Code 2010	Uses f_{R1}, f_{R3} at CMOD 0.5 mm and 2.5 mm for classification and residual strength checks [31].	Trilinear σ–w laws; ratios f_{R1k}/f_{Lk} and f_{R3k}/f_{R1k} aim to ensure sufficient performance relative to conventional steel reinforcement [21].
EN 14651	Defines the notched-beam test and nominal $f_{R,j}$ from beam-theory assumptions [40].	Provides standardised experimental outputs; constitutive identification is defined within national or model-code provisions.
RILEM TC 162-TDF	Established European beam-test recommendations providing the foundation for subsequent code frameworks [32].	Specific σ–w proposals that diverge from MC2010 and national guidelines in certain technical details.
EHE-08 (Spain)	Anchors classification on residual strengths from notched-beam testing [36].	Bilinear/trilinear σ–w models; potential overestimation of ultimate capacity for specific structural configurations (see also Blanco et al.).
CNR-DT 204 (Italy)	Incorporates EN 14651 experimental parameters within Italian national guidance [21].	Constitutive and crack-spacing regulations may underestimate average spacing relative to alternative European frameworks in specific limit states [41].

0.5 mm. It represents the Serviceability Limit State (SLS) and serves as the primary metric for crack width control. In contrast, the parameter f_{R3} is recorded exactly at a CMOD of 2.5 mm and governs the Ultimate Limit State (ULS), defining the ultimate load-bearing capacity of the concrete composite [41]. To prevent brittle failure modes when substituting conventional steel reinforcement, MC2010 mandates strict minimum ductility ratios, explicitly requiring that $f_{R1k}/f_{Lk} > 0.4$ and $f_{R3k}/f_{R1k} > 0.5$ [66].

While MC2010 and the Spanish EHE-08 propose trilinear or bilinear stress-crack ($\sigma - w$) models to simulate post-cracking behaviour, comparative studies reveal significant discrepancies between the various European frameworks. For example, constitutive models proposed by the Italian CNR-DT 204 often underestimate the average crack spacing in structural elements, whereas EHE-08 and MC2010 can sometimes overestimate the ultimate load capacity at large deflections when applied to specific FRC structural elements like elevated slabs [36, 67].

Considerable discussion exists regarding the applicability of these codes, which were primarily developed and calibrated for metallic fibres, to macro-synthetic or polymeric (polyolefin) fibres. A primary limitation is the time-dependent deforma-

tion behaviour of synthetic fibres. Unlike steel, polymeric fibres are viscoelastic and highly susceptible to creep rupture under sustained tensile stresses. This is a phenomenon that codes like MC2010 and EHE-08 struggle to predict accurately over long durations, often warning against the use of synthetic fibres where sustained loads are dominant [14].

Fulfilling the strict ductility ratio required at the narrow 0.5 mm CMOD point (f_{R1}) heavily penalizes macro-synthetic fibres. Due to their significantly lower elastic modulus compared to steel, synthetic fibres must undergo substantially larger displacements before fully engaging and bridging the crack. Consequently, their primary residual strength contributions typically manifest at larger crack openings, rather than at the early SLS threshold dictated by current standards [68].

2.8 Identified Research Gaps

The comprehensive review of the current state-of-the-art literature presented in Sections 2.1 through 2.7 reveals several critical shortcomings in the existing methodologies for evaluating Fibre-Reinforced Concrete (FRC). While significant advancements have been made in theoretical modelling, virtual extensometry, and sustainability assessments, these domains remain largely isolated from one another. Synthesizing the literature highlights three primary research gaps that must be addressed to advance the structural application of macro-synthetic fibres.

2.8.1 Limitations of Current European Frameworks for Macro-Synthetic Fibres

The constitutive models and inverse-analysis techniques embedded in the *fib* Model Code 2010 (MC2010), EN 14651 and EHE-08 were primarily calibrated on hooked-end steel fibres with rigid pull-out mechanics [21]. Ductility is assessed at early crack openings; for example, the f_{R1} parameter is evaluated at a CMOD of 0.5 mm for the Serviceability Limit State. Macro-synthetic (polymeric) fibres possess a significantly lower elastic modulus and exhibit viscoelastic creep, and they therefore require larger crack displacements before full engagement and bridging can develop [16]. The existing steel-calibrated thresholds may not fully capture this delayed engagement or the lower initial stiffness that is characteristic of macro-synthetic reinforcement. A robust framework that evaluates the structural contribution of macro-synthetic fibres while accounting for these differences in engagement kinematics is still absent from the literature.

2.8.2 Integration of Digital Image Correlation with Compliance Correction

Digital Image Correlation has been applied as an observational fracture-kinematics tool in cementitious materials, and machine-compliance correction has been treated separately through dummy-specimen calibration. The literature, however, still lacks a unified protocol that exploits DIC itself to identify the apparatus stiffness and to propagate the correction across an entire experimental campaign. This gap is particularly relevant for macro-synthetic fibre systems, where the lower composite stiffness amplifies the relative weight of compliance artefacts on the recorded response. A further deficiency concerns the embedding of a corrected kinematic record into a complete framework that connects optical CMOD/CTOD measurement to inverse-analysis identification of the σ–w constitutive law. The Optical Deflection-Derived Compliance (ODDC) protocol developed in Chapter 4, and its integration with the inverse-analysis framework of Chapter 7, address this combined gap.

2.8.3 Lack of Integrated Statistical and Eco-Efficiency Frameworks

The construction industry's shift towards sustainability demands lowering the Global Warming Potential (GWP) of concrete, typically achieved by replacing steel fibres with synthetic alternatives or utilizing low-carbon binders [6]. However, the majority of current studies evaluate the mechanical properties (such as fracture toughness) and the environmental impact (Life Cycle Assessment, LCA) in strict isolation. There is a scarce application of robust statistical tools, such as Analysis of Variance (ANOVA) or Response Surface Methodology (RSM), directly correlated with structural testing data to quantify the environmental versus mechanical trade-offs. A holistic, statistically validated selection matrix that quantitatively balances load-bearing capacity against binder intensity and GWP is currently missing from the literature [57].

In conclusion, these three identified gaps in the state-of-the-art literature explicitly validate the necessity of the research undertaken in this thesis. The shortcomings in current code applicability for synthetic fibres, the need for compliance-corrected DIC methodologies, and the demand for statistically integrated eco-efficiency evaluations directly form the foundational basis for the core research objectives that were outlined in Chapter 1.

2.8.4 Mapping of Research Gaps to Thesis Objectives

To establish a clear connection between the identified literature gaps and the research program, table 2.2 maps each of the four specific objectives, originally formulated in section 1.3, to the corresponding gap and proposed research actions.

Table 2.2: Mapping of thesis objectives to identified research gaps and corresponding research actions.

Obj.	**Identified Literature Gap**	**Addressed Objective and Action**	**Chapter(s)**
1	Absence of standardized methods to separate testing machine compliance and rigid-body motion from true specimen kinematics.	Develop and validate the Optical Deflection-Derived Compliance (ODDC) protocol using full-field DIC.	Ch. 4
2	Isolation between structural performance characterization and Life Cycle Assessment (LCA) without statistical trade-off matrices.	Couple non-parametric statistical significance testing with cradle-to-gate Global Warming Potential (GWP) analyses.	Ch. 5
3	Unverified validity of the rigid-hinge kinematic assumption for highly flexible macro-synthetic and glass fibre-reinforced matrices.	Map continuous crack profile evolution (CMOD vs. CTOD) and empirically validate the rigid-hinge kinematics.	Ch. 6
4	Code-level constitutive models (MC2010, EN 14651) calibrated on steel fibres, misrepresenting the post-cracking response of synthetic systems.	Derive trilinear stress-crack opening (σ-w) laws via inverse analysis and align optical data with code criteria.	Ch. 7

Chapter 3

Methodology

3.1 Research Framework and Experimental Programme

The research reported in this thesis combines a mortar programme and a concrete programme. Both programmes address a common question: how do fibre type, fibre geometry, fibre volume fraction and matrix composition jointly control the flexural, splitting-tensile and rebar-bond response of cementitious composites. A parallel question concerns the ability of full-field Digital Image Correlation (DIC) to resolve the underlying deformation field. To answer these questions while keeping the specimen inventory manageable, the experimental campaign is organised as four complementary case studies, two conducted on Fibre-Reinforced Mortar (FRM) and two conducted on Fibre-Reinforced Concrete (FRC). A compact overview of the programme is given in Figure 3.1.

The mortar programme (Case Studies 1 and 2) is conducted on prismatic beams of nominal geometry 40 mm × 40 mm × 160 mm, with a central notch saw-cut at the age of 28 d. Case Study 1 compares ten fibre variants – embossed macro-polypropylene (EMP), waved macro-polypropylene (WMP), twisted fibrillated ferro (TFF), synthetic micro-polypropylene (SMP), alkali-resistant glass (GF) and hooked-end steel (SF) – at a fixed volume fraction of 1 %, evaluated under a Normal-cement matrix ($w/c \approx 0.50$) and a High-cement matrix ($w/c \approx 0.25$). Case Study 2 isolates the SMP family and traces the flexural response as the volume fraction is escalated from 0 % to 1.5 % for two fibre lengths (12 mm and 24 mm) under the same two matrices. The two studies therefore disentangle fibre *type* (at constant dose) from fibre *dose* (at constant type).

The concrete programme (Case Studies 3 and 4) uses 150 mm cubes cast from a single M30 target mix designed by the absolute-volume method. Eight fibre groups are cast, covering the plain control, four single-fibre mixes (EMP–50, WMP–50,

SMP–50, SMP–24), one binary SMP hybrid and two ternary hybrids. For each group, replicate cubes are produced and tested in orthogonal cube splitting tensile (CST) per $f_t = 2P/(\pi dL)$, and in bonded two-bar splitting (BBS), which quantifies the rebar–matrix interfacial bond stress $\tau = F/(2\pi dL_e)$ for two embedded deformed bars of 12 mm diameter and 150 mm embedment length. Reported strengths are the group-wise averages of the replicate specimens.

All four case studies share a common instrumentation chain (Section 3.7): a Galdabini Sun/60 universal testing machine operating in closed-loop displacement control, and two-dimensional DIC implemented with ZEISS Correlate and consumer-grade 1080 p video acquisition. The DIC replaces the usual contact extensometers and provides force–CMOD, force–CTOD and force–deflection curves for each specimen, together with the full surface strain field around the notch tip or the splitting plane.

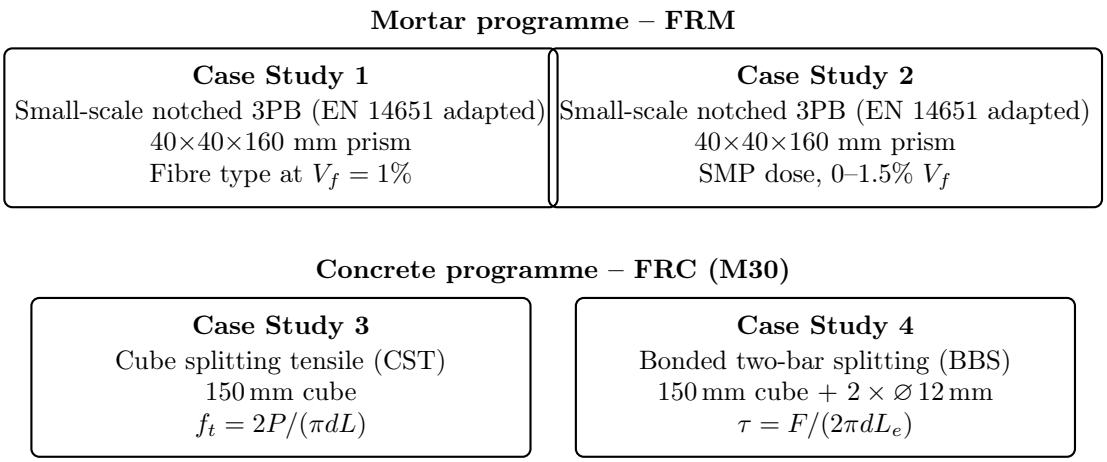


Figure 3.1: Overview of the four case studies composing the experimental programme.

To complement the case-study mapping presented above, the overall sequence of experimental preparation, data acquisition, and analytical modelling is illustrated in fig. 3.2. This schematic highlights the interconnections between materials characterisation, compliance-corrected optical measurement, and the downstream analytical outputs that collectively constitute the research workflow of the thesis.

The remainder of this chapter presents the raw materials and their characterisation (Section 3.2), the four case studies (Sections 3.3 through 3.6), the shared instrumentation and measurement framework (Section 3.7), and a concluding summary of the complete specimen inventory (Section 3.8).

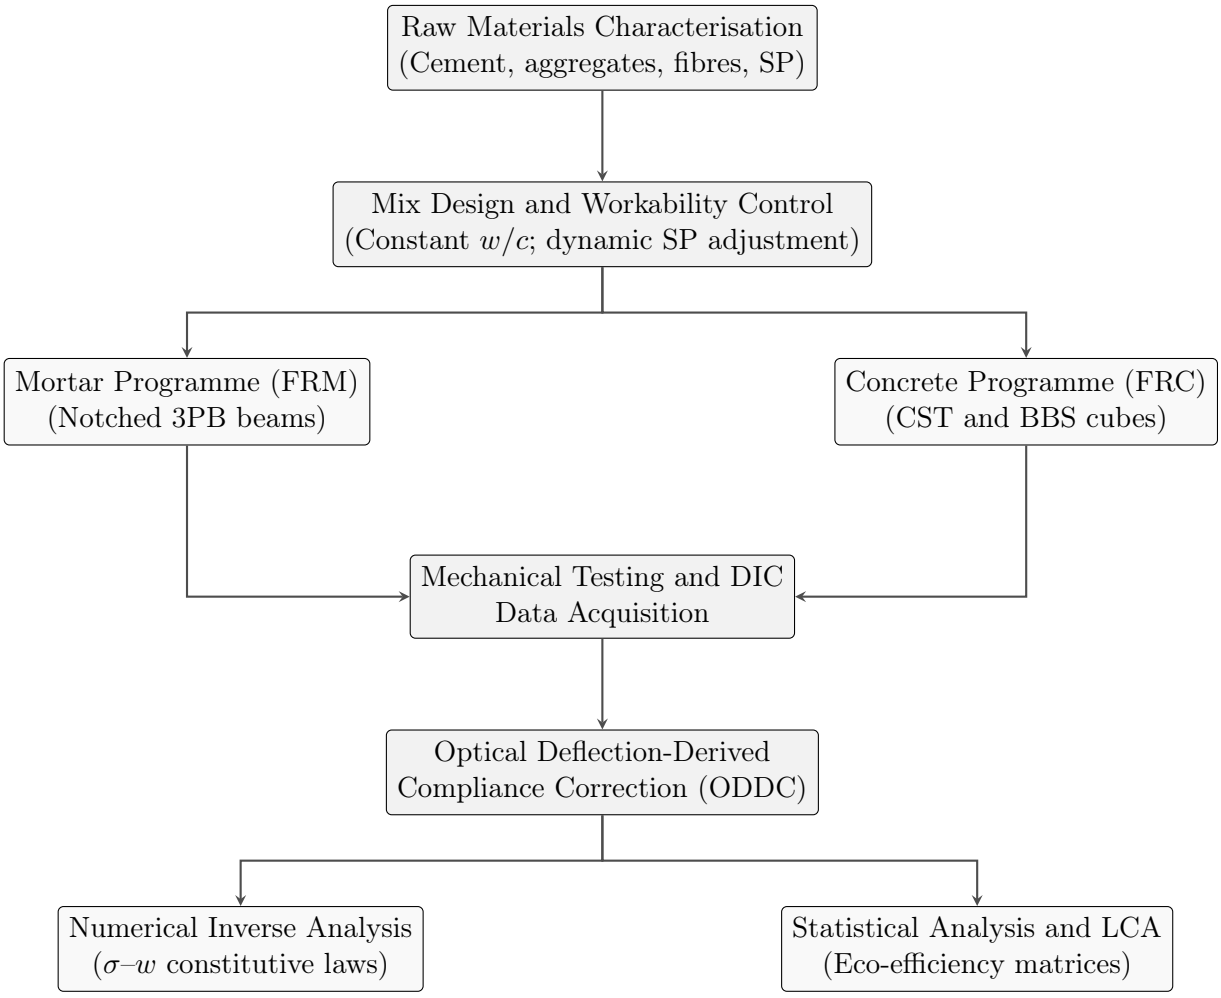


Figure 3.2: Schematic overview of the research methodology workflow, from raw material characterisation to analytical outputs.

3.2 Materials and Constituent Characterisation

This section describes the raw materials shared by all four case studies. Where a constituent is used in only one programme (for example, the coarse aggregate, which is restricted to the M30 concrete) the scope is stated explicitly. Unless otherwise noted, all specific gravities used in the mix-design calculations of Sections 3.3 through 3.6 are those listed in this section.

3.2.1 Cement

A single Portland-limestone cement was used throughout the experimental programme: Italcementi Tecnocem, an EN 197-1 class CEM II/B-LL 42.5R, illustrated in Figure 3.3. The designation *B-LL* indicates a clinker content of 65–79 % complemented by 21–35 % of limestone filler of high purity (total organic carbon ≤ 0.20 % bwol), while the *R* class denotes a rapid-hardening cement with a min-

imum compressive strength of 20 MPa at two days and 42.5 MPa at twenty-eight days. The nominal specific gravity is $\rho_c = 3.10$, which is the value adopted in all volumetric calculations of Sections 3.3 through 3.6.

This cement was selected for three reasons. First, its high tricalcium-silicate (C_3S) content and fine grinding provide accelerated early-age hydration kinetics, which is advantageous in producing comparable strengths across the many mix batches of the mortar programme. Second, the partial replacement of clinker by limestone filler reduces the embodied CO_2 of the binder relative to a pure CEM I without compromising the 42.5R strength class, which is consistent with the sustainability orientation of the research presented in Chapter 5. Third, a single cement source and a single production lot were used for all mortar and concrete batches, eliminating one source of variability across the four case studies. The cement was stored in sealed original packaging under laboratory ambient conditions and used within the manufacturer's recommended shelf life.



Figure 3.3: Italcementi Tecnocem CEM II/B-LL 42.5R used for all mortar and concrete batches.

3.2.2 Fine aggregate

Two different fine aggregates were used, depending on the programme:

- **Mortar programme (Case Studies 1 and 2):** a local Italian quartzitic sand with a nominal specific gravity of $\rho_s = 2.65$ and a maximum grain size below 2 mm, compatible with the EN 196-1 nominal mortar geometry. The

sand was oven-dried, sieved and stored in sealed containers prior to batching, in order to minimise water-absorption variability between batches. It is important to note that the sand content was kept strictly constant across all mortar batches. The fibres were incorporated as a direct addition without any aggregate reduction; even at the maximum dosage of $V_f = 1.5\%$ ($\tilde{}$15 L/m^3), the fibres displace less than 1.3% of the total sand volume, thus maintaining an identical sand-to-binder ratio in all mixtures.

- **Concrete programme (Case Studies 3 and 4):** a manufactured sand (M-Sand) of crushed-limestone origin, with a specific gravity $\rho_{\mathrm{MS}} = 2.65$. This choice is dictated by the compatibility of the M-Sand grading with the 10 mm and 20 mm coarse aggregates used in the concrete mix (Section 3.2.3) and by its lower water demand relative to natural river sand.

The sand source is therefore a distinguishing feature of the two programmes: the mortar programme is formulated with a very fine quartzitic sand that produces a highly cohesive paste-dominated matrix, whereas the concrete programme uses an M-Sand compatible with coarse aggregates and with the larger, 150 mm cube specimens.

3.2.3 Coarse aggregate (concrete programme only)

Two crushed-limestone coarse-aggregate fractions were combined in the M30 concrete mix (Case Studies 3 and 4): a 10 mm single-size fraction (CA–10) and a 20 mm single-size fraction (CA–20), both with a nominal specific gravity of $\rho_{\mathrm{CA}} = 2.70$. The two fractions were blended by volume in the approximate ratio CA–10 : CA–20 $\approx$ 37.5 : 62.5, yielding a continuously graded combined aggregate that complies with EN 12620. No coarse aggregate is present in the FRM mortar (Case Studies 1 and 2), consistent with the cement/sand binder matrix of EN 196-1 nominal geometry.

Similar to the fine aggregate methodology, the incorporation of macro-synthetic fibres in the concrete programme was treated as a direct volumetric addition rather than a partial aggregate replacement. Given that the maximum fibre dosage used in the concrete mixes was $V_f \leq 0.33\%$ (approximately 3.3 L/m^3), the theoretical aggregate displacement is negligible and falls well within the standard batching tolerances permitted by EN 206. This approach ensures that the fundamental aggregate skeleton remains unaffected by the fibre inclusion.

3.2.4 Superplasticiser

A high-range water-reducing admixture, Mapefluid N200 (Mapei S.p.A., Italy), compliant with EN 934-2 Tables 3.1 and 3.2, was used throughout the experimental programme; Figure 3.4 shows the admixture as supplied. The reported specific gravity is $\rho_{\mathrm{SP}} = 1.06$, and the manufacturer-recommended dosage envelope of 0.5–1.5 % by weight of cement (bwoc) was respected in every mix. In the mortar programme, the admixture serves two distinct functions depending on the matrix: in the Normal matrix it compensates the workability loss induced by the addition of fibres without altering the water-to-cement ratio, whereas in the High matrix ($w/c = 0.25$) its role becomes essential, because the mixture is otherwise unpourable under hand-trowel mixing and the admixture alone restores a castable consistency. In the concrete programme, the superplasticiser demand tracks the fibre rheology: micro-fibres demand more admixture per kilogram than macro-fibres, and ternary blends accumulate the demands of all three components. This physical hierarchy is visible in the actual batched SP dosages of Tables 3.3, 3.4 and 3.6 and is exploited later in the rheology–toughness discussion of Chapter 5.

The guiding criterion for the mix design adjustments was the strict maintenance of a constant water-to-cement ratio across all comparative batches, ensuring that variations in mechanical performance could be attributed solely to the fibre reinforcement rather than to matrix strength fluctuations. Consequently, the superplasticiser dosage was adjusted dynamically to compensate for the reduction in workability caused by the increasing surface area of the added fibres. For the concrete programme, the target consistency was class S3 (100–150 mm slump) according to EN 206, ensuring adequate workability for standard structural casting and compaction. For the mortar programme, adequate flow and a cohesive, castable consistency without fibre balling or segregation were visually assessed and verified in accordance with the principles of EN 1015-3.

3.2.5 Fibres

Six fibre families were used across the four case studies. A consistent short-label convention is adopted throughout the thesis: an uppercase family code followed by the nominal fibre length in millimetres. This notation (e.g. SMP–12 mm) provides a single, unambiguous nomenclature used from Chapter 3 onward. The six families are:

- **EMP** – Embossed Macro-Polypropylene. Four lengths were supplied: 14 mm, 18 mm, 36 mm and 50 mm. The first three lengths are used in the mortar



Figure 3.4: Mapefluid N200 polycarboxylate-based high-range water-reducing admixture, compliant with EN 934-2.

programme (Case Study 1) at 1 % V_f; the 50 mm length is used in the concrete programme as the macro component of groups G2 and G7.

- **WMP** – Waved Macro-Polypropylene. Two lengths: 30 mm for the mortar programme (Case Study 1) and 50 mm for the concrete programme (groups G3 and G8).
- **TFF** – Twisted Fibrillated Ferro. A single length, 25 mm, used in Case Study 1 only.
- **SMP** – Synthetic Micro-Polypropylene. Three lengths: 12 mm, 24 mm and 50 mm. The three lengths appear in Case Study 1; lengths 12 mm and 24 mm form the dose-response of Case Study 2; lengths 24 mm and 50 mm appear in the concrete programme.
- **GF** – Alkali-resistant chopped glass fibre, nominal length 8 mm. Used in Case Study 1 only, in both Normal and High matrices.
- **SF** – Hooked-end steel fibre of high-carbon origin, nominal length 35 mm. Used in Case Study 1 only, in both Normal and High matrices.

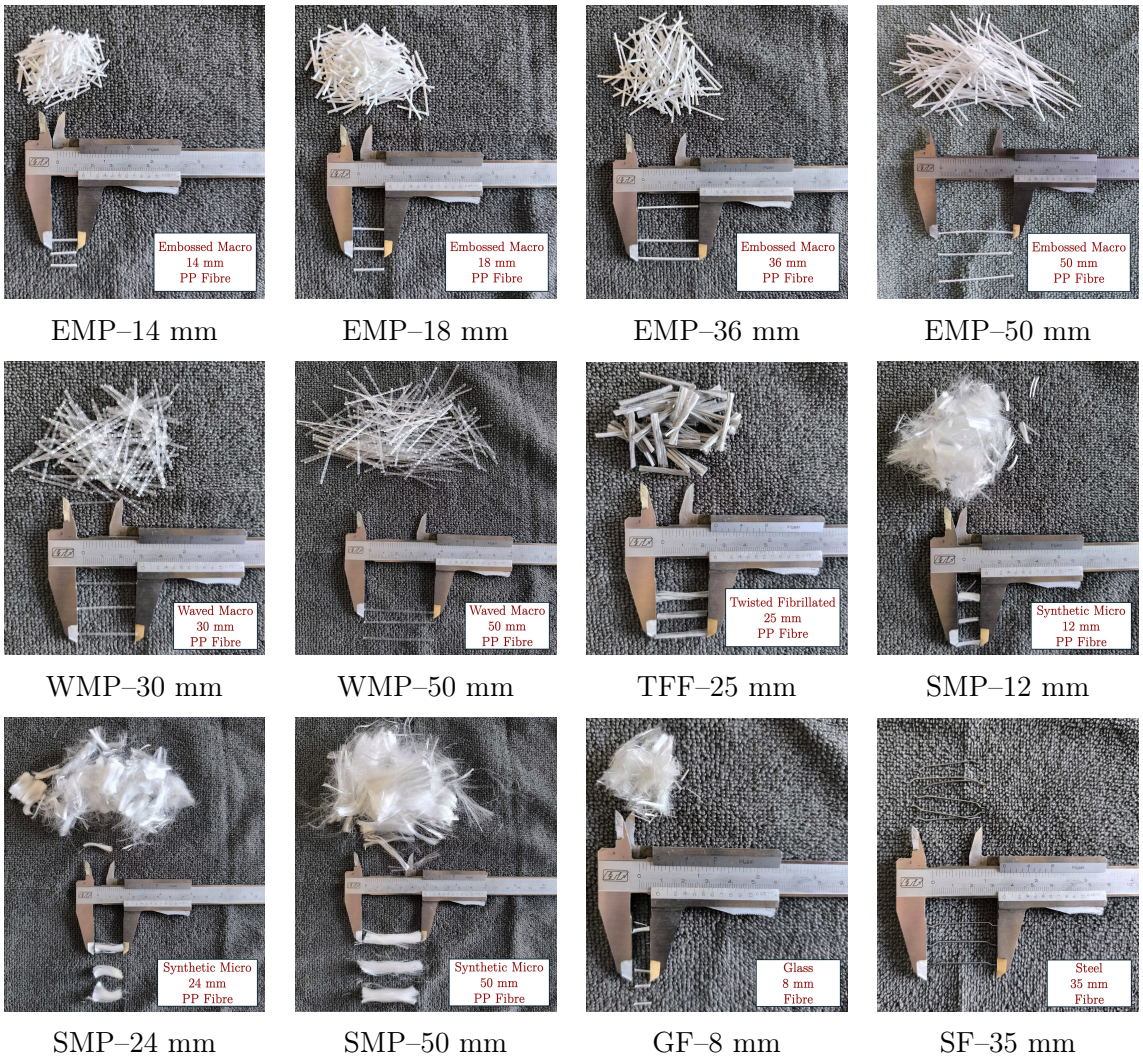
A summary of the geometries and the programme scope of each label is given in Table 3.1; the corresponding photographs, taken at a common scale with a digital caliper for reference, are grouped in Figure 3.5. The SMP monofilaments used here comply with BS EN 14889-2:2006 and ASTM C1116 Type III as synthetic polypropylene fibres for structural fibre-reinforced concrete; the virgin polypropylene polymer has a melting point of 162 °C, an ignition point of 360 °C, a hydrophobic surface with a proprietary coating to promote dispersion, and a diameter range of 19–40 µm across the three lengths used here [31]. Individual tensile-strength and

elastic-modulus values of the macro-polypropylene fibres (EMP, WMP, TFF) are not transcribed in this chapter; the mechanical role of each fibre is inferred from the experimental force–displacement curves reported in Chapters 4 and 5.

Table 3.1: Fibre inventory used across the four case studies.

Label	Family	Form	Lengths (mm)	SG	Scope
EMP	Macro PP	Embossed	14, 18, 36, 50	0.91	C1, C3, C4
WMP	Macro PP	Waved	30, 50	0.91	C1, C3, C4
TFF	Macro PP	Twisted	25	0.91	C1
SMP	Micro PP	Monofilament	12, 24, 50	0.91	C1, C2, C3, C4
GF	AR glass	Chopped	8	2.68	C1
SF	Steel	Hooked-end	35	7.85	C1

PP = polypropylene; AR = alkali-resistant; SG = specific gravity. C1–C4 = Case Studies 1–4.



EMP–14 mm | EMP–18 mm | EMP–36 mm | EMP–50 mm

WMP–30 mm | WMP–50 mm | TFF–25 mm | SMP–12 mm

SMP–24 mm | SMP–50 mm | GF–8 mm | SF–35 mm

Figure 3.5: The twelve fibre variants used across the four case studies, displayed at a common scale.

3.3 Case Study 1: Comparative flexural response of plain, steel, glass and polypropylene-fibre-reinforced mortar at a fixed 1 % fibre volume fraction

The objective of Case Study 1 is to isolate the effect of fibre *type and geometry* on the notched three-point flexural response of fibre-reinforced mortar (FRM). To this end, ten fibre variants – four macro-polypropylene (EMP–14 / 18 / 36, WMP–30, TFF–25), three synthetic micro-polypropylene (SMP–12 / 24 / 50), one alkali-resistant glass (GF–8) and one hooked-end steel (SF–35) – are tested at a *fixed volume fraction of 1 %*, in order to eliminate volumetric effects and to expose the constitutive role of fibre morphology. Each fibre is cast in both mortar matrices defined in Section 3.3.1, yielding a balanced $11 \times 2 = 22$ mix factorial, schematically summarised in Figure 3.6. In that figure, the eleven columns list the fibre variants (plus one plain control), the two rows are the Normal and High matrices, and the darker-filled circles denote the control cells.

A caveat regarding the volume-fraction basis of Case Study 1 must be recorded. A fixed $V_f = 1\,\%$ does *not* imply comparable fibre count, fibre surface area, number of fibres crossing the ligament, or bridging probability across the six families. Steel ($\rho = 7.85$), glass ($\rho = 2.68$) and polypropylene ($\rho = 0.91$) differ strongly in density, diameter and aspect ratio; at the same volumetric dose, the PP families contribute roughly an order of magnitude more discrete filaments per unit crack area than steel, and the glass fibres supply a far larger surface area per unit volume than the macro-PP variants. Case Study 1 is therefore a *volume-fraction comparison*, not a comparison at equal number of fibres or equal interfacial surface area. This distinction is important when interpreting the post-cracking residual forces of Chapter 5.

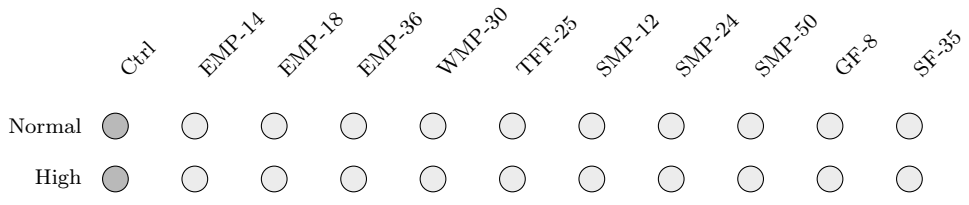


Figure 3.6: Factorial design of Case Study 1 (22 mortar mixes).

3.3.1 Cementitious matrices (Normal and High)

Two mortar matrices were designed with the same constituents (cement, quartzitic sand, Mapefluid N200) but with contrasting cement contents and water-to-cement ratios, so that the fibre-matrix interaction is probed over a representative strength range. The Normal matrix ($C \approx 580$ kg/m^3, $w/c \approx 0.50$) represents the lower bound of a structural cementitious mortar, while the High matrix ($C \approx 801$ kg/m^3, $w/c \approx 0.25$) represents an ultra-high cement-content mortar with a dense, paste-dominated interfacial transition zone. Nominal proportions are summarised in Table 3.2.

Table 3.2: Nominal composition of the two FRM matrices used in Case Studies 1 and 2.

Matrix	**Cement (kg/m^3)**	**Water (kg/m^3)**	**Sand (kg/m^3)**	w/c	**SP (% bwoc)**
Normal	≈ 580	≈ 290	≈ 1230	0.50	0.72–1.10
High	≈ 801	≈ 200	≈ 1280	0.25	0.61–0.80

3.3.2 Specimen matrix and mix proportions

The complete 22-mix matrix of Case Study 1 is reported in Table 3.3. Each row corresponds to one representative beam of the fibre/matrix pair indicated in that row, with a batched composition (cement, water, sand, SP) traceable to the original mixing log. The programme is balanced: 11 mixes in the Normal matrix and 11 mixes in the High matrix; within each matrix, one control plus ten fibre variants. In accordance with standard practice for notched three-point bending tests on fibre-reinforced mortar, three replicate beams were cast per mixture for each case study; the values reported in Chapter 5 are the group-wise means of the valid replicates. Any specimens excluded from the analysis due to casting defects, premature failure or DIC-tracking loss are noted individually in the discussion of the corresponding case study.

Table 3.3: Case Study 1 – specimen matrix of the 22 FRM mixes at $V_f = 1\%$.

Matrix	**Fibre**	**L (mm)**	V_f **(%)**	**Fibre (kg/m^3)**	**C (kg/m^3)**	**W (kg/m^3)**	**SP (kg/m^3)**
Normal	Control	–	0.0	0	580.8	288.6	4.22
Normal	EMP	14	1.0	9.1	578.7	291.9	4.82
Normal	EMP	18	1.0	9.1	582.2	289.2	4.82
Normal	EMP	36	1.0	9.1	579.5	292.4	5.07
Normal	WMP	30	1.0	9.1	581.5	288.1	4.89
Normal	TFF	25	1.0	9.1	578.2	290.5	4.94
Normal	SMP	12	1.0	9.1	582.4	291.2	5.14
Normal	SMP	24	1.0	9.1	579.1	288.3	5.43
Normal	SMP	50	1.0	9.1	581.6	292.9	5.60
Normal	GF	8	1.0	25.0	580.5	290.3	4.78
Normal	SF	35	1.0	78.5	579.8	291.1	4.97
High	Control	–	0.0	0	802.1	200.7	4.88
High	EMP	14	1.0	9.1	799.6	201.1	5.36
High	EMP	18	1.0	9.1	803.0	199.8	5.37
High	EMP	36	1.0	9.1	800.4	201.3	5.61
High	WMP	30	1.0	9.1	802.3	199.8	5.44
High	TFF	25	1.0	9.1	799.1	200.9	5.48
High	SMP	12	1.0	9.1	801.5	199.6	5.87
High	SMP	24	1.0	9.1	799.5	201.2	5.98
High	SMP	50	1.0	9.1	802.8	199.8	6.40
High	GF	8	1.0	25.0	800.3	201.1	5.10
High	SF	35	1.0	78.5	802.2	199.3	5.26

Fibre mass is calculated from the fibre specific gravity (Table 3.1) and the 1 % volumetric dose.

3.3.3 Casting, compaction and curing

All FRM batches were mixed by hand with a mortar trowel, following a fixed, repeatable sequence informed by the preliminary trial mixes. The procedure is described once here and referenced without repetition by Case Study 2 (Section 3.4).

1. **Dry blending.** The Tecnocem 42.5R cement and the oven-dried quartzitic sand were blended dry for approximately two minutes until the mixture showed a uniform grey colour with no visible sand streaks.

2. **Pre-dissolution of the admixture.** The Mapefluid N200 dose (calculated per row of Table 3.3) was pre-mixed with the mixing water, eliminating local over-concentration.

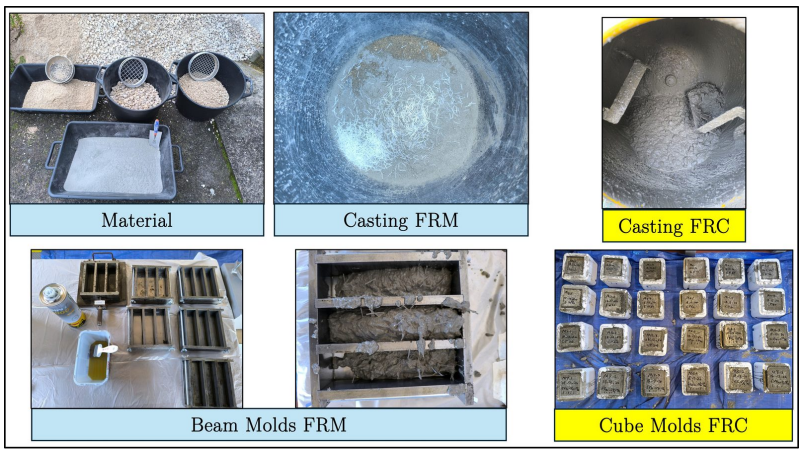
3. **Staged addition of the water+SP mixture.** Half of the water+SP was added and hand-mixed for approximately one minute; a second 25 % increment was then added and mixed for another minute; finally the remaining 25 % was added. Total hand-trowel mixing time was 4–5 minutes, which yielded a homogeneous paste without visible lumps.

4. **Fibre addition.** Once the matrix was homogeneous, the prescribed mass of fibres (Table 3.3) was sprinkled *gradually* over the mortar in three portions,

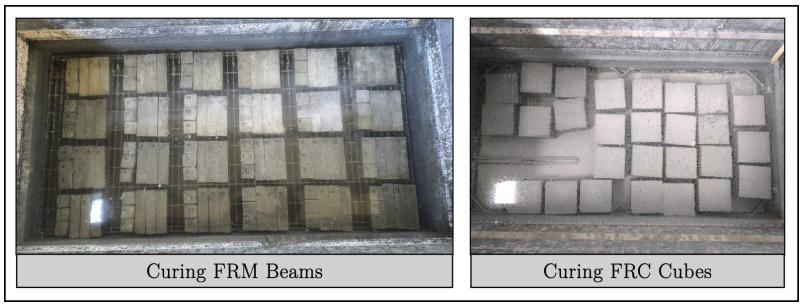
with continuous trowel-mixing between portions, in order to disrupt fibre balls and promote dispersion. This step typically added a further 2–3 minutes to the mixing time, depending on fibre family and length.

5. **Placing and compaction.** The fresh mortar was placed into pre-oiled steel beam moulds in *two equal layers*. Each layer was compacted on a vibrating table for 15 s; surface finish was achieved with a steel trowel after the second vibration.

6. **Curing.** Specimens were covered with a polyethylene sheet to prevent evaporation during the first 24 h in the laboratory ambient air. They were then demoulded and transferred to a water-curing tank kept at ambient temperature until the age of 28 d, in accordance with the curing regime used for all four case studies. Figure 3.7 illustrates the casting and curing operations.



(a) Two-layer placement in steel moulds with vibrating-table compaction



(b) Submerged water curing until the age of 28 d

Figure 3.7: Specimen preparation for the FRM beams: casting and curing.

3.3.4 Notch preparation

At the age of 28 d, a single mid-span notch was introduced in every beam by high-precision wet-sawing with a diamond blade (Figure 3.8). The notch depth was fixed at $a_0 = h/4 = 10\,\text{mm}$ (beam depth $h = 40\,\text{mm}$) and the notch width at 2 mm. This geometry does not reproduce the EN 14651 notch ratio ($a_0/h \approx 0.167$) but follows a small-scale fracture-mechanics configuration adopted here to obtain a sufficiently defined ligament and stable crack localisation in the reduced beam geometry. The notch width of 2 mm corresponds to the wet-saw blade kerf, and the actual notch dimensions were measured for each specimen to verify the residual ligament height. The notch was cut across the full 40 mm width of the beam, with the saw axis perpendicular to the casting direction, so that fibre orientation effects induced by the mould walls are identical in every specimen.

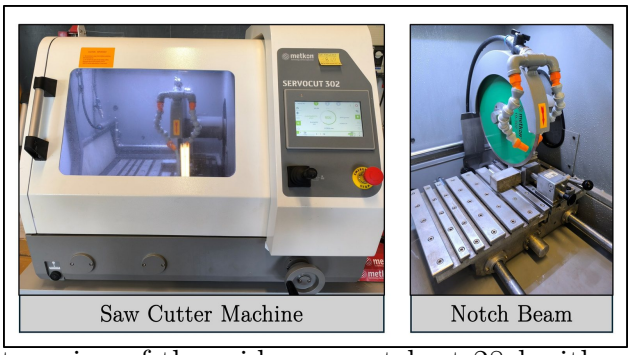


Figure 3.8: Wet-sawing of the mid-span notch at 28 d with a diamond blade.

3.4 Case Study 2: Dose–response flexural behaviour of Synthetic Micro-Polypropylene (SMP) fibre-reinforced mortar with variable volume fraction and fibre length

The objective of Case Study 2 is complementary to Case Study 1: where Case Study 1 fixed the volume fraction at 1 % and varied the fibre type, Case Study 2 fixes the fibre family (SMP) and varies the volume fraction on a fine grid. By restricting the fibre chemistry to synthetic micro-polypropylene and varying only the dose and the fibre length, Case Study 2 isolates the *dose–response* curve of the composite, which is central to the statistical and optimisation analysis of Chapter 5.

3.4.1 Design variables

Three factors are varied in Case Study 2:

- **Fibre length:** SMP–12 mm and SMP–24 mm (the two lengths used in practical shotcrete and precast applications).
- **Volume fraction:** 0 %, 0.75 %, 1.0 %, 1.25 %, 1.5 % (five levels; the 0 % and 1.0 % points are shared with Case Study 1 to avoid batch redundancy).
- **Matrix:** Normal and High, defined identically to Case Study 1 (Table 3.2).

The full-factorial design yields $2 \times 4 \times 2 + 2 = 18$ mixes (two matrices, two lengths, four non-zero doses, plus the two controls), as reported in Table 3.4 and schematically illustrated in Figure 3.9.

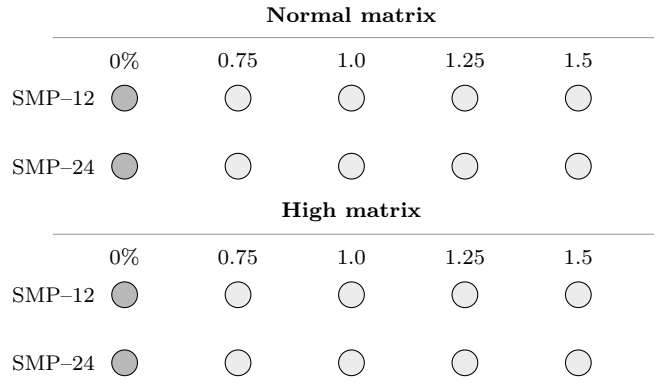


Figure 3.9: Dose-response factorial of Case Study 2: two fibre lengths × five volume fractions × two matrices (control cells filled darker).

Table 3.4: Case Study 2 – specimen matrix of the 18 FRM mixes of the SMP dose response.

Matrix	**SMP length (mm)**	V_f **(%)**	**Fibre (kg/m^3)**	**C (kg/m^3)**	**W (kg/m^3)**	**SP (kg/m^3)**
Normal	–	0.00	0.000	580.8	288.6	4.22
Normal	12	0.75	6.825	581.5	288.8	4.77
Normal	12	1.00	9.100	582.4	291.2	5.14
Normal	12	1.25	11.375	577.9	291.7	5.52
Normal	12	1.50	13.650	580.6	287.6	5.76
Normal	24	0.75	6.825	578.8	292.0	4.86
Normal	24	1.00	9.100	579.1	288.3	5.43
Normal	24	1.25	11.375	582.3	289.2	5.77
Normal	24	1.50	13.650	579.3	292.6	6.03
High	–	0.00	0.000	802.1	200.7	4.88
High	12	0.75	6.825	802.7	199.7	5.42
High	12	1.00	9.100	801.5	199.6	5.87
High	12	1.25	11.375	799.9	201.2	6.18
High	12	1.50	13.650	801.8	199.8	6.44
High	24	0.75	6.825	799.4	201.3	5.61
High	24	1.00	9.100	799.5	201.2	5.98
High	24	1.25	11.375	802.4	199.6	6.29
High	24	1.50	13.650	800.1	201.6	6.78

3.4.2 Casting, compaction, curing and notch preparation

The mixing, compaction, curing and notch-cutting procedures for Case Study 2 are identical to those described in Sections 3.3.3 and 3.3.4 for Case Study 1 and are not repeated here. In particular, the fibre-addition step increases in duration with fibre dose (up to approximately four minutes at $V_f = 1.5\,\%$, SMP–24) to ensure full dispersion of the micro-fibres.

3.5 Case Study 3: Cube splitting tensile assessment of macro-, micro- and hybrid-polypropylene fibre-reinforced concrete (M30)

Case Study 3 extends the scope from flexural mortar tests to splitting-tensile concrete tests, on specimens representative of structural concrete rather than binder mortar. The aim is to quantify the influence of fibre type (macro vs micro vs hybrid) and total dose on the indirect tensile response of a single M30 concrete.

3.5.1 M30 mix design

A single M30 target mix was designed by the absolute-volume method for a unit yield of 1.000 m^3, with a nominal water-to-cement ratio $w/c \approx 0.52$ and a fibre-free baseline. The nominal composition, derived from the laboratory batching records, is reported in Table 3.5; the *SG* column lists the specific gravities adopted in the volumetric calculation. The same cement (Tecnocem 42.5R, Section 3.2.1), M-Sand, CA–10 + CA–20 coarse aggregate and Mapefluid N200 are used for all eight fibre groups of the concrete programme; only the fibre dosage and, consequently, the SP dosage are varied.

Table 3.5: Nominal M30 concrete mix (fibre-free baseline).

Component	**Mass (kg/m^3)**	**SG**	**Volume (m^3)**
Cement – CEM II/B-LL 42.5R	350.0	3.10	0.1129
Free water	182.0	1.00	0.1820
Mapefluid N200	3.50	1.06	0.0033
M-Sand	629.4	2.65	0.2375
CA–10 (10 mm nominal)	458.1	2.70	0.1697
CA–20 (20 mm nominal)	732.9	2.70	0.2714
Entrapped air	–	–	0.0250

3.5.2 Eight fibre groups

Eight fibre groups were cast from this baseline M30, covering the plain control, four single-fibre mixes (one per main fibre family used in the concrete programme), one binary SMP hybrid and two ternary hybrids (Table 3.6). For each of the eight groups, replicate 150 mm cubes were produced and tested in orthogonal cube splitting tensile (CST, Section 3.5.3) and in bonded two-bar splitting (BBS, Section 3.6); reported strengths are the group-wise averages of the replicate specimens. The eight fibre groups are illustrated schematically in Figure 3.10.

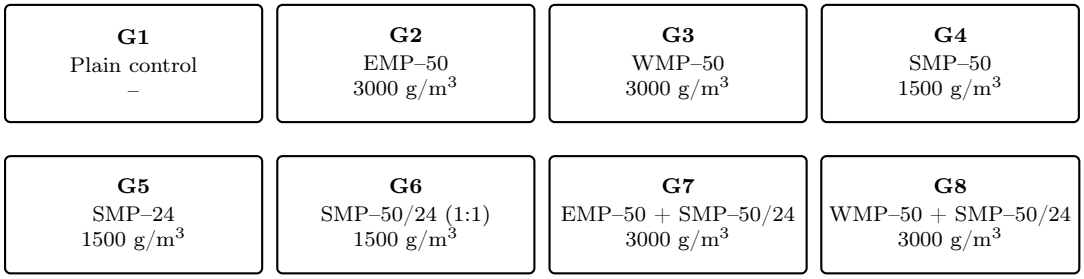


Figure 3.10: Plan of the eight fibre groups (G1–G8) of the M30 concrete programme used in Case Studies 3 and 4.

Table 3.6: Case Studies 3 and 4 – eight fibre groups of the M30 concrete programme.

Group ID	**Fibre composition**	**Fibre dose (g/m^3)**	**SP (kg/m^3)**	**SP (% bwoc)**
G1	Plain control	0	3.49	0.997
G2	EMP–50 (3000)	3000	4.53	1.295
G3	WMP–50 (3000)	3000	4.37	1.245
G4	SMP–50 (1500)	1500	4.00	1.142
G5	SMP–24 (1500)	1500	4.15	1.188
G6	SMP–50 (750) + SMP–24 (750)	1500	4.09	1.166
G7	EMP–50 (1500) + SMP–50 (750) + SMP–24 (750)	3000	4.92	1.408
G8	WMP–50 (1500) + SMP–50 (750) + SMP–24 (750)	3000	4.81	1.369

Fibre dose in g/m^3; SP dose given in kg/m^3 and as % by weight of cement (bwoc).

The SP hierarchy visible in Table 3.6 (G1 < G4 & G5 < G6 < G2 & G3 < G7 & G8) encodes the rheological cost of each fibre system: micro-fibres alone (G4, G5, G6) raise the SP demand modestly, single macro-fibre systems (G2, G3) raise it more strongly, and the ternary macro+micro hybrids (G7, G8) accumulate both effects, reaching the maximum SP dose at 4.92 kg/m^3 (1.41 % bwoc). This rheological ordering anticipates the toughness ordering that will be reported in Chapter 4.

3.5.3 Specimen geometry and splitting-tensile formula

The CST specimens are 150 mm cubes, cast from the M30 mix of Table 3.5 with the fibre dosages of Table 3.6. At the age of 28 d, each cube was positioned orthogonally in the Galdabini Sun/60 testing frame between two 25 mm-wide bearing strips (plywood, one top and one bottom), with the strips aligned along a diameter of the loaded face to generate a uniform splitting-tensile stress field on the perpendicular plane. The splitting tensile strength is then computed from the peak applied load P by the standard Brazilian-type formula

$$f_t \;=\; \frac{2P}{\pi\, d\, L}, \qquad d = L = 150\,\text{mm}, \tag{3.1}$$

so that the geometric constant πdL amounts to $\pi \cdot 150 \cdot 150 \approx 70\,686$ mm^2. Equation (3.1) assumes a linear elastic distribution of stress across the splitting plane at peak load; for fibre-reinforced concrete it is interpreted here as a *first-crack* tensile strength, with the post-peak load-sustaining behaviour recovered from the force–opening curve reconstructed by DIC (Section 3.7.3).

3.5.4 Casting, compaction and curing of the FRC cubes

Mixing of the FRC concretes was performed in a drum concrete mixer. The procedure is described here once and is common to Case Studies 3 and 4.

1. **Dry blending.** The Tecnocem 42.5R cement, the M-Sand, the CA–10 and the CA–20 were added to the drum mixer in the proportions of Table 3.5 and blended dry for 90 s until the aggregates showed a uniform coat of cement.

2. **Wet blending.** The Mapefluid N200 was pre-dissolved in the mixing water and added to the drum, and mixing continued until a uniform paste coating of the aggregates was obtained.

3. **Fibre addition.** The prescribed fibres (Table 3.6) were added last, gradually, over approximately 120 s of continuous mixing, until the fibres were fully integrated and no fibre clumps were visible.

4. **Placing and compaction.** The fresh concrete was placed into pre-oiled 150 mm steel cube moulds in *three equal layers*. Each layer was vibrated on a vibrating table for 15 s (total vibration time 45 s per cube).

5. **Curing.** Cubes were covered with a polyethylene sheet for the first 24 h in the laboratory ambient air, then demoulded and transferred to the water-curing

tank until the age of 28 d. Figure 3.7 applies equally to the FRC cubes; the only difference from the FRM protocol is the three-layer placement and the drum-mixer mixing sequence.

3.6 Case Study 4: Bonded two-bar splitting evaluation of rebar–matrix interfacial behaviour in fibre-reinforced concrete (M30)

Case Study 4 closes the concrete programme by moving from the bulk tensile response of Case Study 3 to the *interfacial* behaviour between a ribbed reinforcing bar and the fibre-reinforced concrete matrix. The rationale for this study is that, in many structural applications of fibre-reinforced concrete, the governing limit state is not the tensile cracking of the matrix but rather the bond-splitting along anchorage or lap-splice zones; the fibres then act as passive crack-width controllers that mobilise after matrix cracking. The bonded two-bar splitting (BBS) configuration is adopted here because it directly couples bond stress and splitting stress in a single specimen [18].

3.6.1 Specimen geometry

Each BBS specimen is a 150 mm cube cast from the same M30 mix of Table 3.5 and the same fibre groups of Table 3.6. Two deformed (ribbed) steel bars of nominal diameter $\varnothing\,12$ mm are embedded through the cube, parallel to each other and symmetric about the cube's vertical axis, so that the embedded length is equal to the cube side, $L_e = 150$ mm, on both bars simultaneously. The total bond area mobilised by the two bars at failure is therefore

$$A_b \;=\; 2\,\pi\,d\,L_e \;=\; 2\pi \cdot 12 \cdot 150 \;=\; 11\,309.7\ \text{mm}^2. \tag{3.2}$$

The same bar diameter, embedment length and rib profile are used for all eight groups, so that differences between groups are attributable solely to the matrix and fibre composition.

3.6.2 Bond-stress formula

The applied peak load F at pull-out or splitting failure is converted to an average bond stress through

$$\tau = \frac{F}{2\,\pi\,d\,L_e} = \frac{F}{A_b}. \tag{3.3}$$

Equation (3.3) implicitly assumes a uniform distribution of bond stress along the embedment length; this is an acceptable engineering approximation for the reduced embedment ($L_e/\varnothing = 12.5$) adopted here, and is consistent with the fib Model Code 2010 treatment of bond for short anchorage lengths [31].

3.6.3 Relationship to Case Study 3

The eight BBS specimens are one per fibre group (Groups G1–G8 of Table 3.6). The mix design, the fibre composition, the mixing procedure, the compaction and the curing regime are identical to those of Case Study 3 (Sections 3.5.1 through 3.5.4) and are not repeated. The only specimen-level difference is the placement of the two $\varnothing\,12$ mm bars in the mould prior to casting, secured at the cube centreline with template plates at the top and bottom faces of the mould in order to preserve the $L_e = 150$ mm embedment.

3.7 Instrumentation and Measurement Framework

3.7.1 Mechanical testing apparatus and loading protocols

3.7.1.1 Testing frame

All four case studies share a single testing machine: a Galdabini Sun/60 servo-hydraulic universal testing machine with a nominal capacity of 600 kN, operated in closed-loop displacement control through its integrated machine controller. The frame is instrumented with a calibrated internal load cell (class 1 accuracy in accordance with DINEN ISO et al [52]) and an internal crosshead-displacement transducer used here only as a safety cross-check against the DIC-derived global displacement. Testing was performed in a climate-controlled laboratory. Figure 3.11 shows the three loading configurations used in this thesis: notched three-point bending on a 40 mm FRM beam (Case Studies 1 and 2), orthogonal cube splitting tensile on a 150 mm FRC cube with 25 mm bearing strips (Case Study 3), and bonded two-bar splitting on a 150 mm FRC cube with two embedded $\varnothing\,12$ mm deformed bars (Case Study 4).

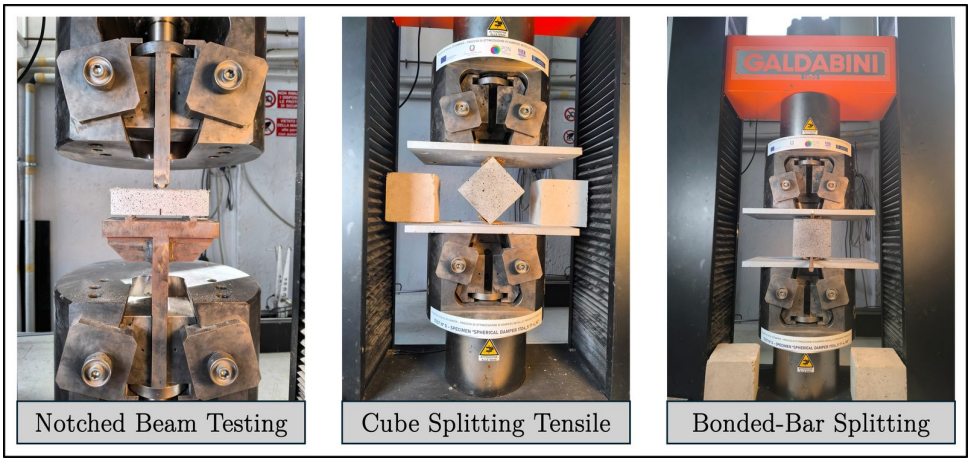


Figure 3.11: Galdabini Sun/60 testing frame in its three loading configurations.

3.7.1.2 Three-point bending on FRM beams: geometry and loading rate

The FRM beams of Case Studies 1 and 2 are tested in a three-point bending configuration with a support span of $L = 100\,\text{mm}$ (span-to-depth ratio $L/h = 100/40 = 2.5$), central loading roller and two support rollers all of $\varnothing\,30\,\text{mm}$. The span is smaller than the $L = 500\,\text{mm}$ of the EN 14651:2005 reference (150 mm beam, $L/h = 3.33$) [40] because the present study uses a $40 \times 40 \times 160$ mm prism geometry, selected for three practical reasons: (i) the limited supply of each experimental fibre type constrains the number and size of specimens that can be cast per batch; (ii) the 40 mm beam depth is compatible with the available three-point bending jig on the Galdabini Sun/60; and (iii) the reduced specimen mass allows hand-trowel mixing with controlled fibre dispersion, which is essential for the ten-variant Case Study 1 programme. The resulting geometry is *not* a scaled EN 14651 test — EN 14651:2005 prescribes a single $150 \times 150 \times 550$ mm geometry and does not include any scaling provision — but a small-scale notched three-point bending test with geometry and loading rates inspired by EN 14651 and adapted to the 40 mm beam depth.

To retain the two-stage rate philosophy of EN 14651 (slow before the limit of proportionality, faster in the softening branch) in the presence of the reduced geometry, the prescribed CMOD rates of EN 14651 are scaled by the linear geometric factor $k_s = 150/40 = 3.75$, producing the protocol of Table 3.7. The philosophy is identical to the RILEM TC 162-TDF recommendation [32]. Table 3.8 summarises a selection of recent studies that have adopted similar or comparable loading-rate protocols on small-scale fibre-reinforced cementitious beams, confirming that the present rate-scaling approach is consistent with current experimental practice.

The residual flexural strengths obtained from the present $40 \times 40 \times 160$ mm geometry are reported throughout this thesis as f_R-equivalent values; they are *not*

Table 3.7: CMOD-controlled loading-rate protocol for the notched FRM beams.

Phase	CMOD interval	EN 14651 rate	This study (scaled)
A (pre-LOP)	$0 \rightarrow 0.1$ mm	0.050 mm/min	**0.1875 mm/min**
B (softening)	$0.1 \rightarrow 4.0$ mm	0.200 mm/min	**0.75 mm/min**

Table 3.8: Selected literature on notched-beam flexural testing at various specimen scales and loading rates.

Reference	Specimen (mm)	Loading rate	Fibre type
EN 14651 [40]	$150 \times 150 \times 550$	0.05 / 0.2 mm/min (two-phase)	Metallic
Ibrahim et al. [43]	$40 \times 40 \times 160$	0.5 mm/min (constant)	Synthetic PP
Gali & Subramaniam [35]	$125 \times 250 \times 1500$	0.5 mm/min	Steel (0.5% vol)
Paegle et al. [19]	Various	0.05–0.2 mm/min	Steel/PP
RILEM TC 162-TDF [32]	$150 \times 150 \times 550$	Two-phase EN 14651 protocol	Steel

directly interchangeable with the $f_{R,1}$–$f_{R,4}$ performance classes defined in EN 14651 for the standard $150 \times 150 \times 550$ mm beam. Well-documented size effects on fracture energy and apparent residual strength in cementitious composites [17, 18] indicate that smaller specimens tend to report higher apparent strengths per unit ligament area than the standard geometry. This limitation applies to all cross-study comparisons in Chapters 5 and 6.

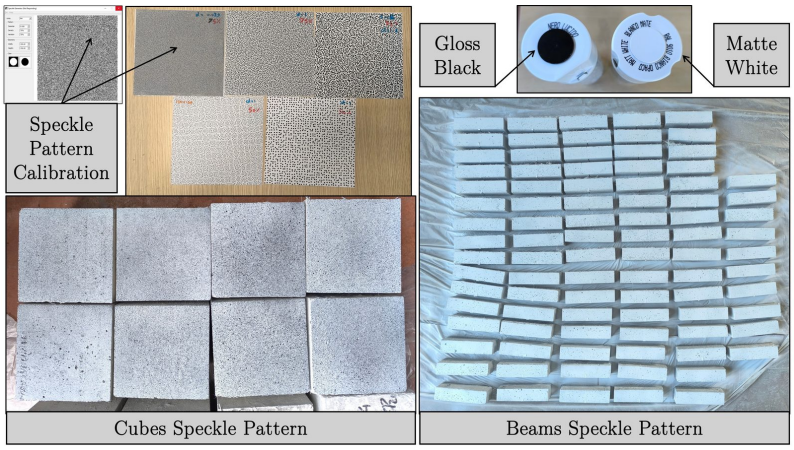
During every test the machine outputs the true CMOD signal used by the closed-loop controller; this is obtained in the present study from the ZEISS Correlate virtual extensometer described in Section 3.7.3, with the machine crosshead displacement used only as a safety cross-check. No contact extensometer or clip gauge is mounted on the beam, because the limited 40 mm beam depth leaves no room for a conventional EN 14651 clip gauge without compromising the specimen.

3.7.1.3 Cube splitting tensile and bonded-bar splitting: loading rate

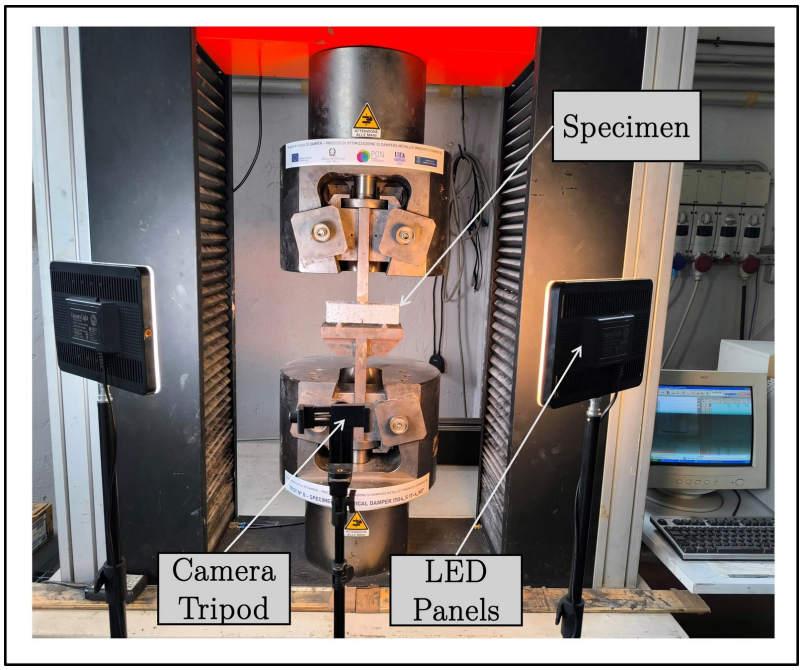
Both the CST (Case Study 3) and BBS (Case Study 4) concrete tests are performed in displacement control at a single rate of 1 mm/min applied through the machine crosshead. This rate is standard for indirect-tensile and pull-out tests on fibre-reinforced concrete, and is slow enough to capture the post-peak force–opening behaviour of the fibre bridging without inducing appreciable strain-rate effects. No contact sensors are mounted on the cubes, as the splitting failure of fibre-reinforced concrete generates a sudden, through-thickness crack that would damage or eject any clip gauge or LVDT placed on the specimen surface; all kinematic quantities are therefore recovered post-test from the DIC video (Section 3.7.3).

3.7.2 Digital Image Correlation – optical setup

The deformation field on the observation face of every specimen is recovered by two-dimensional Digital Image Correlation (2D-DIC) using ZEISS Correlate. The optical chain (Figure 3.12) is identical for the four case studies, with only the working distance and the field of view adapted to the specimen size (40 mm beams vs 150 mm cubes).



(a) Speckle-pattern preparation on a cube



(b) DIC capture system in position in front of the Galdabini Sun/60

Figure 3.12: DIC optical chain.

3.7.2.1 Speckle-pattern preparation

The observation face of every specimen was prepared as follows:

1. a thin matte *white* base coat of acrylic spray paint was applied uniformly to the surface and left to dry for approximately 10 min;
2. once the base coat was dry, a *black* gloss spray paint was applied by hand-controlled short bursts, producing a dense, random, high-contrast speckle pattern with characteristic speckle size on the order of 0.2–0.6 mm;
3. the speckle was left to dry for at least 30 min before testing to prevent any wet-paint creep under load.

The speckle density was verified by visual inspection against the classical criterion of 3–5 speckles per facet [46, 48]; the facet size adopted for correlation is discussed in Section 3.7.3.1.

3.7.2.2 Imaging hardware

Imaging was performed with a consumer-grade digital camera operated in 1080 p (1920×1080 px) video mode at 60 fps, with *manual* focus locked before the test to avoid any autofocus drift during capture. The camera was mounted on a heavy tripod with sandbag ballast; the optical axis was aligned perpendicular to the observation face and the working distance was set to approximately 1 m, following the DIC-on-concrete setup of Ibrahim [43]. Two diffuse LED panels were placed at ±45° off-axis relative to the camera so as to eliminate specular glare and residual shadows on the speckle pattern. The full hardware configuration is shown in Figure 3.12b. Specific manufacturer and model details of the camera sensor are not reported here; the relevant property for 2D-DIC is the effective spatial resolution, which is fixed by the 1920 × 1080 frame size and the working distance.

3.7.2.3 Spatial calibration

The pixel-to-millimetre calibration was performed at the start of each test using a rigid reference of known length in the image plane (a printed ruler for the mortar beams and the 150 mm cube edge for the concrete cubes). Representative calibration for the FRM beams is 6.71 px/mm (0.149 mm/px); equivalent calibrations apply to each cube test.

3.7.3 Digital Image Correlation – ZEISS Correlate processing

All DIC post-processing was performed with ZEISS Correlate (v2025), which implements a two-dimensional subset-based correlation with a zero-normalised sum of squared differences (ZNSSD) similarity function. The correlation parameters adopted here are consistent with the published recommendations for notched-beam DIC at consumer-grade resolution [43, 46, 49] and are fixed once and for all across the four case studies to allow a fair cross-study comparison.

3.7.3.1 Facet size, step size and shape function

A *facet size of* 19×19 *px* and a *point distance (step) of 16 px* were used, together with a first-order affine shape function. The rationale for these values on a 1080 p sensor at 1 m working distance is:

- On the beam face, 19 px corresponds to $19 \cdot 0.149 \approx 2.83\,\text{mm}$, i.e. about 7 % of the 40 mm beam depth – fine enough to resolve the CMOD/CTOD opening zone yet large enough for each facet to contain 50–100 independent greyscale features under the hand-sprayed speckle density;
- The 16 px step yields a virtual measurement-point grid with a spacing of $\approx 2.39\,\text{mm}$ on the specimen, so the 27 mm virtual extensometers defined in Section 3.7.3.2 are each resolved by approximately eleven correlation points;
- Both values sit near the centre of the 15–25 px facet / 12–20 px step band that is reported in the DIC-on-concrete literature for similar optical chains [43, 45].

3.7.3.2 Virtual extensometers

Virtual extensometers (VE) were placed on the DIC field in a configuration common to every FRM beam:

- a **CMOD** extensometer, horizontal, of gauge length 27 mm, placed symmetrically across the notch mouth at the beam soffit;
- a **CTOD** extensometer, horizontal, of gauge length 27 mm, placed symmetrically across the notch tip;
- a **mid-span deflection** point, a vertical point-tracker referenced to two support points defined on the specimen top surface, well outside the fracture process zone.

For the FRC cubes of Case Studies 3 and 4, virtual extensometers are not placed; instead, a regular grid of validation points is tracked across the observation face so that the force-versus-kinematic response is recovered from the full DIC field post-test, which is robust to the sudden splitting failure that would damage any point-like gauge.

Additional "validation points" are also recorded on the beams, near the actuator, at mid-height of the specimen and close to the supports, so that the redundancy of the DIC field can be exploited in Chapter 4 to verify rigid-body consistency and to detect any anomalous out-of-plane motion. Figure 3.13 collects a representative set of specimens at the end of the DIC-monitored tests: the notched FRM beams of Case Studies 1 and 2 (top row), CST FRC cubes showing the through-specimen splitting crack of Case Study 3 (bottom-left), and the BBS FRC cubes of Case Study 4 (bottom-right).

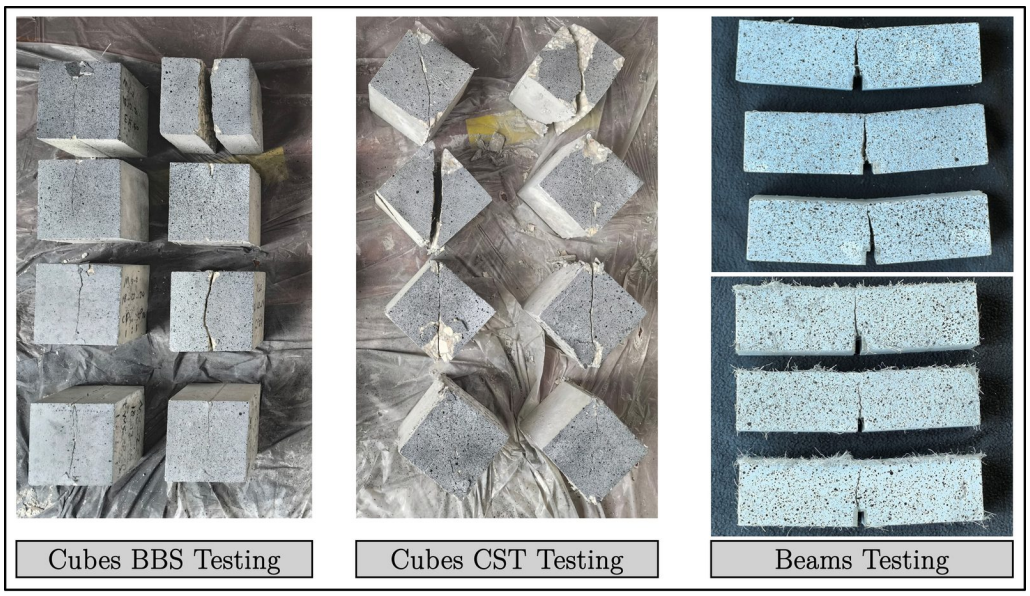


Figure 3.13: Representative tested specimens at the end of DIC-monitored loading.

3.7.3.3 Frame subsampling and validation

Raw video was recorded at 60 fps and imported into ZEISS Correlate at 1 fps for correlation, which is sufficient to resolve the flexural softening of the beams (typical total loading time 200–600 s) and the splitting response of the cubes, while keeping the correlation job computationally tractable. The DIC-derived global displacement was cross-checked against the machine crosshead signal in every test; residual mismatches, attributable to compliance of the loading train, are used in Chapter 4 to support the K-correction methodology.

3.8 Summary of the experimental programme

The experimental programme presented in this chapter comprises four complementary case studies summarised in Table 3.9.

Table 3.9: Final specimen inventory of the experimental programme.

Case	Test	Material	Governing variable(s)
1	Small-scale notched 3PB (EN 14651 adapted)	FRM	Fibre type and geometry at $V_f = 1\%$
2	Small-scale notched 3PB (EN 14651 adapted)	FRM	SMP dose (0–1.5%); length (12 vs 24 mm)
3	Cube splitting tensile (150 mm)	FRC M30	Fibre type, geometry and hybridisation
4	Bonded two-bar splitting	FRC M30	Fibre type, geometry and hybridisation

All case studies share a single cement (Tecnocem CEM II/B-LL 42.5R, Section 3.2.1), a single admixture (Mapefluid N200, Section 3.2.4) and a single instrumentation chain (Galdabini Sun/60 + ZEISS Correlate, Section 3.7). The two FRM programmes use a common quartzitic sand and a common two-layer casting protocol; the FRC programme uses a common M-Sand + CA–10 + CA–20 aggregate skeleton and a common three-layer drum-mixer protocol. This constructional uniformity ensures that the cross-case comparisons performed in Chapters 4 and 5 can legitimately attribute observed differences in mechanical response to *fibre* and *matrix* variables alone.

Chapter 4

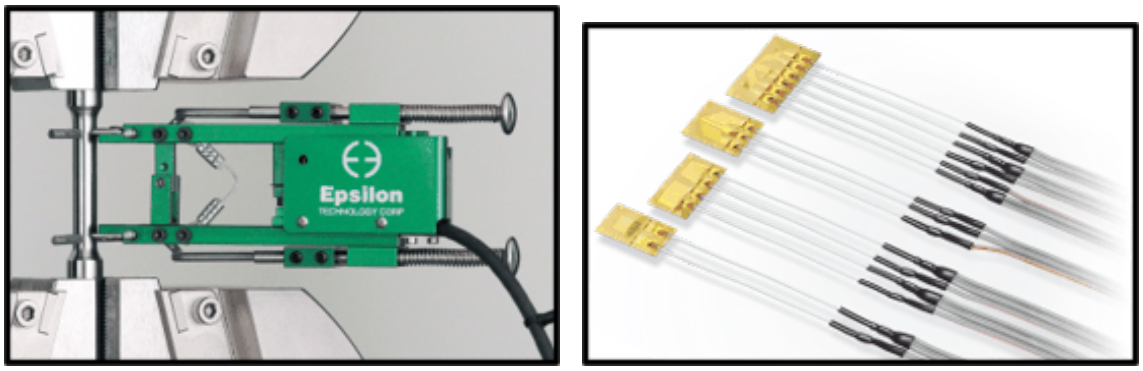
DIC-Derived Compliance Correction and Optical Validation

To ensure full traceability of the experimental programme, every mechanical response curve presented in this chapter is explicitly linked to the mix design matrix detailed in Chapter 3 (tables 3.3, 3.4 and 3.6). The specimen nomenclature as example, N-SMP12-1.5 for a normal-matrix mortar beam reinforced with 12 mm SMP fibres at $V_f = 1.5\,\%$. it is directly encodes the matrix type, fibre geometry and volumetric dosage, allowing any individual test record to be traced back to its constituent batch proportions and casting conditions. The complete specimen inventory is summarised in table 3.9 and the overarching methodological workflow is illustrated in fig. 3.2.

4.1 Why machine compliance distorts the raw force–displacement record

The closed-loop displacement signal delivered by the Galdabini Sun/60 load-frame is measured at the crosshead (Section 3.7.1.1). Along the load path that links that crosshead to the notch face of the specimen there are, in series, the hydraulic piston, the load cell, the lower and upper rollers of the three-point-bending jig, the supporting platen and the specimen itself. Every one of these elements carries a finite elastic compliance. Because the machine transducer is blind to that chain, the displacement it reports is the sum of the *true* specimen deformation and the elastic travel of the intervening machine, in the spirit of the compliance separation formalised by Okazaki and co-workers for fracture-toughness testing [51]. For the 40 mm-depth notched FRM beam of Case Studies 1 and 2 – a specimen whose

intrinsic stiffness rarely exceeds ~ 10 kN/mm – the two contributions are of the same order, and the raw crosshead record therefore understates the specimen stiffness, stretches the pre-peak branch and pushes the origin of the curve artificially to the right.



(a) Epsilon 3442 clip-gauge extensometer (b) Surface-bonded electrical strain gauge

Figure 4.1: Two contact sensors not used in this thesis: (a) clip-gauge extensometer; (b) surface-bonded strain gauge.

A quick sanity check on the calibration reference beam – N-CONTROL, the plain unreinforced mortar prism – shows why the artefact cannot simply be ignored. The slope of the raw, uncorrected force–displacement record of N-CONTROL in the proportional range is $K_{\text{act}} = 4\,168$ N/mm. The same beam, tracked optically through the DIC virtual deflection gauge of Section 3.7.3.2, yields an intrinsic specimen stiffness of $K_{\text{spec}} \approx 9\,900$ N/mm over the same load interval. From these two numbers, the series-spring decomposition developed later in Section 4.2.1 gives a machine-path stiffness of $K_{\text{machine}} \approx 7\,140$ N/mm. In other words, the raw crosshead record understates the true specimen stiffness by a factor of $K_{\text{spec}}/K_{\text{act}} \approx 2.4$ even for this comparatively stiff mortar reference beam. For the most compliant fibre beam in the Case 1 cohort – N-GF-8 mm, whose short glass fibres generate the weakest elastic engagement – the raw stiffness drops further to $K_{\text{act}} = 1\,618$ N/mm, and the stiffness understatement climbs to a factor of ~ 4.4. In practical terms, these factors translate into a peak-load displacement that is two to four times too large, a fictitious extended toe-region, and an energy-under-the-curve that is grossly overcounted. None of the downstream quantities – the limit of proportionality (LOP), the post-peak residuals prescribed by EN 14651 [40], the dissipated fracture energy G_F – can be extracted reliably until that elastic drift is separated out. The full calibration, including the derivation of K_{machine}, is the subject of Section 4.2.

The first instinct of a laboratory would be to mount a contact sensor: a clip-gauge extensometer or a surface-bonded strain gauge, both shown in Figure 4.1. Either device would provide specimen-mounted kinematics – at a point for the clip gauge,

along a short gauge length for the strain gauge – but both demand a physical mount that the present specimens cannot safely host, and the thesis replaces them in full by optical tracking (Section 3.7.2). On the present specimens, however, the usual instrumentation route is closed. The notched FRM beam is only 40 mm deep and carries a saw-cut 13 mm tall at mid-span; once the LOP is passed, the crack opens from a few tenths of a millimetre to several millimetres in a fraction of a second, and any clip gauge placed across that notch would be torn away by the very event it is supposed to measure. A surface strain gauge would in principle escape that fate, but it is bonded over a gauge length of 3–6 mm and is physically incapable of measuring the *discontinuous* field that develops once the cohesive zone has softened. In the FRC programme the problem is even starker: the 150 mm cubes of Case Studies 3 and 4 fail by a sudden, through-thickness splitting crack propagating between the bearing strips (or, in the bonded-bar configuration, radiating from the embedded bars). The jet of fragments ejected at peak load would crush any gauge or LVDT mounted on the specimen surface. A contact sensor is therefore physically incompatible with three out of the four test configurations of this thesis.

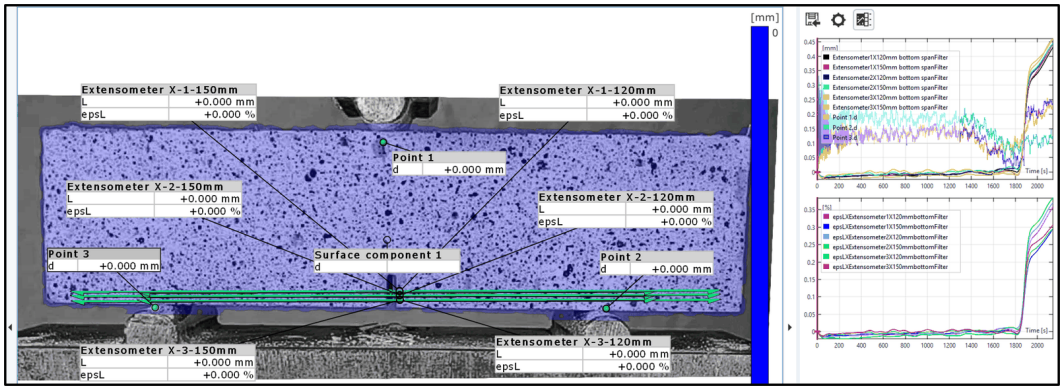


Figure 4.2: ZEISS Correlate v2025 workspace deployed in this chapter.

The route adopted here is therefore optical. ZEISS Correlate, already introduced in Section 3.7.3, delivers the full, time-resolved displacement field of the observation face at 1 fps; from that field any number of *virtual* extensometers can be post-defined without ever touching the specimen. Figure 4.2 shows the v2025 workspace in the configuration adopted for every beam of Case Studies 1 and 2. The left-hand pane carries the speckled observation face of a 40 mm FRM beam in the notched 3PBT jig; the virtual CMOD gauge sits across the notch mouth on a 27 mm baseline, the virtual CTOD gauge across the notch tip, a vertical deflection tracker is referenced to two supports well outside the fracture zone, and four peripheral validation points record rigid-body motion so that out-of-plane drift can be detected *post hoc*. The right-hand pane is the exported force–displacement chart for the selected virtual

gauge. The same chain is re-used, without change of optical parameters, on the 150 mm concrete cubes; there the notch-mouth gauges are replaced by a regular grid of surface-validation points that survives the splitting event (Section 4.4.3).

The remainder of this chapter is organised around the question raised in this section: how to extract from the raw force–displacement record a corrected signal that reflects the specimen alone, using only the optical chain. Section 4.2 derives the *Optical Deflection-Derived Compliance* (ODDC) methodology, which fixes a single machine stiffness once and for all from a reference mortar beam and then applies it uniformly across the entire FRM programme. Section 4.3 formalises the point-tracking validation that decides whether the corrected curve is to be trusted or not. Section 4.4 then documents the application of both procedures to every one of the 56 specimens of the four case studies, and Section 4.5 closes by placing the compliance correction in the wider setting of what the DIC chain delivers *beyond* compliance.

4.2 Optical Deflection-Derived Compliance (ODDC) methodology

The ODDC procedure rests on the initial observation that, in the pre-cracking range where the specimen responds linearly, the total crosshead displacement splits additively into a machine-elastic and a specimen-elastic contribution. ODDC isolates the machine contribution once, on one reference mortar beam, using the *optical* rather than the mechanical displacement record, and then reuses that single value as the global compliance subtrahend for every subsequent test of the same case study.

4.2.1 Governing equations

In the proportional range of the force–displacement record the machine and the specimen are two elastic springs in series (Figure 4.3). Denoting by u_{mach} the displacement reported by the machine transducer, by u_{spec} the displacement of the specimen soffit relative to its supports and by F the applied force, the series partition reads

$$u_{\mathrm{mach}}(F) = u_{\mathrm{spec}}(F) + \frac{F}{K_{\mathrm{machine}}} \tag{4.1}$$

where K_{machine} is the lumped elastic stiffness of the complete load path that lies between the crosshead transducer and the specimen soffit (piston, load cell, rollers, platen contact). All three terms in (4.1) are functions of the applied force F; in the proportional range both $u_{\mathrm{mach}}(F)$ and $u_{\mathrm{spec}}(F)$ are linear in F, which permits the

stiffness reading of the raw record K_{act} to be defined as

$$K_{\mathrm{act}} \;\equiv\; \left.\frac{\mathrm{d}F}{\mathrm{d}u_{\mathrm{mach}}}\right|_{F_0}^{F_{\mathrm{LOP}}} \tag{4.2}$$

Introducing $K_{\mathrm{spec}} \equiv \mathrm{d}F/\mathrm{d}u_{\mathrm{spec}}$, the differentiation of (4.1) yields

$$\frac{1}{K_{\mathrm{act}}} \;=\; \frac{1}{K_{\mathrm{spec}}} \;+\; \frac{1}{K_{\mathrm{machine}}} \tag{4.3}$$

Equation (4.3) is the classical compliance split of Kuk-Cheol Kim *et al.* [51]. The ODDC methodology makes use of it in *reverse*: rather than postulating K_{spec} *a priori* and solving for K_{machine}, it identifies K_{spec} *optically* on a reference mortar beam (where the displacement of the soffit relative to the supports can be recovered directly from the DIC virtual deflection gauge), and solves (4.3) for K_{machine}:

$$K_{\mathrm{machine}} \;=\; \left(\frac{1}{K_{\mathrm{act}}} - \frac{1}{K_{\mathrm{spec}}}\right)^{-1} \tag{4.4}$$

This machine stiffness is then used as a *global* subtrahend for every subsequent specimen of the programme. Rearranging (4.1) shows the final correction that acts on the specimen records:

$$u_{\mathrm{corr}}(F) \;=\; u_{\mathrm{mach}}(F) \;-\; \frac{F}{K_{\mathrm{machine}}} \tag{4.5}$$

The quantity u_{corr} is the ODDC-corrected displacement reported in all force–displacement figures of Sections 4.4.1 and 4.4.2.

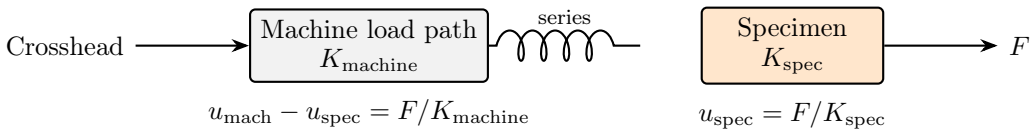


Figure 4.3: The ODDC splits the crosshead displacement into a machine-elastic and a specimen-elastic contribution arranged in series.

4.2.2 Identification of the machine stiffness from the reference mortar beam

The reference specimen adopted here is N-CONTROL, the plain-mortar, unreinforced beam of Case Study 1. This choice is made on two grounds. First, an unreinforced mortar beam in the proportional range is a linear elastic solid of well-characterised geometry, so its intrinsic stiffness K_{spec} can be obtained directly from

the optical deflection field without invoking any fibre-related uncertainty. Second, the same N-CONTROL mix is cast in every mortar batch of Case Studies 1 and 2, so a single identification calibrates two complete case studies.

The identification is performed as follows. The raw force–displacement record of N-CONTROL delivers $K_{\text{act}} = 4\,167.8$ N/mm (slope of the proportional range between the seating load $F_0 = 195$ N and a cut-off of $\sim 80\,\%$ of the peak force). The DIC virtual deflection gauge, referenced to two top-surface points outside the fracture zone, delivers $K_{\text{spec}} \approx 9\,900$ N/mm over the same load interval. Substitution into (4.4) gives

$$K_{\text{machine}} = \left(\frac{1}{4167.8} - \frac{1}{9900} \right)^{-1} \simeq 7\,140 \text{ N/mm}$$

which is the number retained throughout this chapter.

Table 4.1 summarises the three ingredients, the numerical substitution and the final compliance. A notational point is worth recording once: no fibre-volume or mix-design information enters the derivation – $K_{\text{machine}} = 7\,140$ N/mm is therefore treated as an *in-situ* calibration of the equivalent elastic stiffness of the complete load path for this specific frame–jig–support configuration. It is not a specimen property and is transferred only to beam tests performed on the same Galdabini Sun/60, with the same bending fixture and the same alignment procedure.

Table 4.1: Derivation of the machine stiffness K_{machine} from the N-CONTROL mortar beam.

Quantity	**Value**	**Source**
Raw (machine-transducer) stiffness K_{act}	4 167.8 N/mm	Eq. (4.2)
Specimen stiffness K_{spec}	∼9 900 N/mm	DIC virtual deflection gauge
Machine stiffness K_{machine}	**7 140 N/mm**	Eq. (4.4)

4.2.3 Application of the correction

Every subsequent specimen of Case Studies 1 and 2 is processed by the same three-step pipeline. A seating threshold F_0 is first identified, below which the machine output is contaminated by platen-contact non-linearity; F_0 is set at 15 % of the peak force of each specimen, rounded to the nearest multiple of the load-cell resolution, and the values recorded in this chapter range from $F_0 = 180$ N to $F_0 = 217.5$ N. The lower force threshold was introduced to exclude the initial seating and contact-adjustment region of the force–displacement curve, where the platen–specimen con-

tact is not yet fully established and the displacement signal is contaminated by rigid-body motion. The sensitivity of the corrected response to this threshold was checked by comparing values of F_0 at 10 %, 15 % and 20 % of the peak force, confirming that the resulting correction was not materially affected by the choice within that range. The specimen stiffness K_{act} of (4.2) is then identified by least squares on the proportional window, together with the coefficient of determination R^2 of the fit. Finally, the correction formula of (4.5) is applied to the displacement record at every force level from F_0 to the final residual, using the single value $K_{\text{machine}} = 7\,140$ N/mm. The correction is therefore arithmetic in nature: at each recorded force sample the machine-elastic displacement F/K_{machine} is subtracted from the raw crosshead displacement. The output of the pipeline for the forty beams is reported in Sections 4.4.1 and 4.4.2.

Figure 4.4 provides a visual confirmation that the correction achieves its intended purpose. The ODDC-corrected actuator displacement (solid blue line) is overlaid on the true DIC displacement (dashed red line) for the N-CONTROL mortar beam, i.e. the same reference specimen from which K_{machine} was derived. Once the machine-elastic component has been subtracted, the corrected crosshead record collapses onto the optically measured response across the full loading history: the two traces are nearly indistinguishable from the onset of loading through the peak and into the post-peak softening branch. The shaded yellow band marks the zone of minimal deviation between the two signals, confirming that the residual discrepancy after correction is comparable to the DIC measurement resolution. This agreement validates the series-spring assumption of (4.1) in the pre-peak range and demonstrates that the single-value correction of (4.5) does not introduce any systematic bias in the post-peak regime.

4.2.4 Scope of the ODDC correction

Two limits of the method are worth recording explicitly. First, the calibration of K_{machine} was performed in the proportional range of the reference beam, where the series-spring assumption of (4.1) holds rigorously. The correction formula (4.5), however, is applied at every force level of the complete displacement record – including the post-peak softening branch – because the subtrahend F/K_{machine} remains algebraically well-defined at every force sample. What changes beyond first cracking is not the formula but its physical interpretation: in the proportional range the correction removes a true elastic machine contribution, whereas in the post-peak regime the machine contribution decreases with force and the subtrahend becomes progressively smaller relative to the total displacement. In practice, the CMOD and CTOD

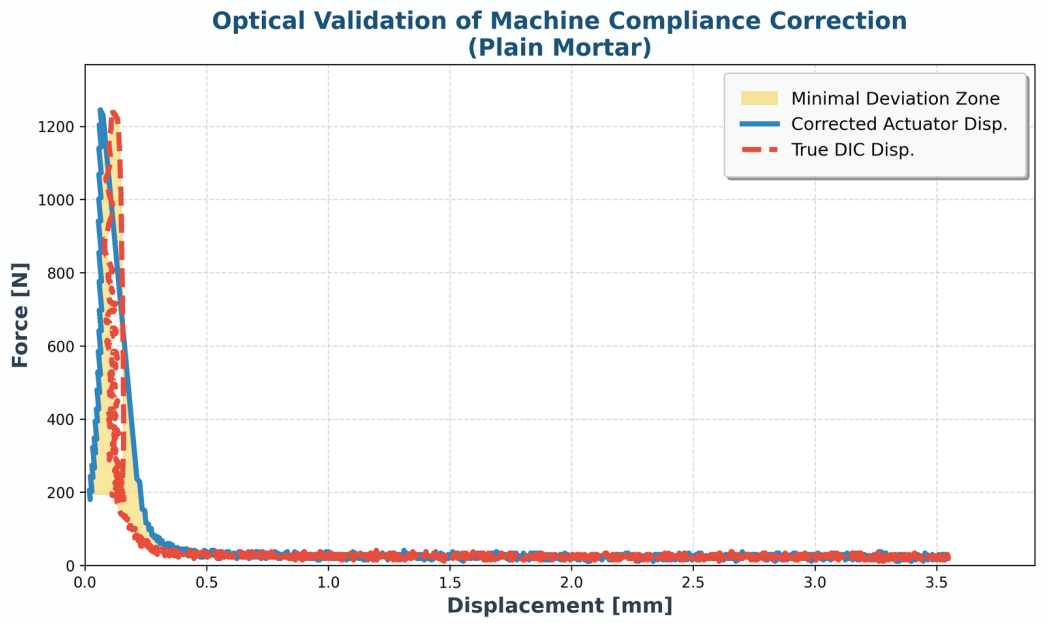


Figure 4.4: Optical validation of the ODDC correction on the N-CONTROL mortar beam. The corrected actuator displacement (solid blue) overlays the true DIC displacement (dashed red); the shaded band marks the zone of minimal deviation between the two signals, confirming that the corrected record collapses onto the optically measured response across the full loading history.

readings that define the post-peak residuals of engineering interest (EN 14651 f_{R1}–f_{R4}, CMOD at 0.5 mm) are extracted from the DIC virtual extensometers, which are measured directly on the specimen and are by construction independent of actuator/crosshead compliance. Their accuracy remains governed, however, by the optical measurement uncertainty, including calibration, facet tracking, gauge placement and possible out-of-plane motion.

Second, the ODDC correction is deliberately restricted to the FRM programme. On the 150 mm FRC cubes of Case Studies 3 and 4 the load path is dominated by the cube itself, not by the machine: the cube stiffness in the splitting-tensile configuration is of the order of 200–300 kN/mm, roughly two orders of magnitude above K_{machine}. The machine correction of Eq. (4.5) would translate into a displacement shift that is smaller than the DIC point-tracking resolution, and the approach adopted there is not a compliance subtraction but a direct optical validation of the machine-reported displacement against the DIC field. That validation is the subject of Section 4.3.

4.3 Optical validation through DIC point tracking

In configurations where the specimen geometry, stiffness, or post-cracking brittleness precludes the use of contact-based transducers (clip gauges, strain gauges), the DIC chain becomes the sole source of specimen-mounted kinematic data. This section formalises how the ZEISS Correlate output is processed into a displacement signal that can be compared to, and when necessary substituted for, the machine transducer. The procedure is labelled here *DIC point tracking* because it operates on the trajectory of individual virtual measurement points ("facets") rather than on a virtual extensometer; the two viewpoints are equivalent under the affine shape function of Section 3.7.3.1, but the point-tracking formulation is more natural for the redundancy tests that follow.

4.3.1 Point-trajectory extraction and interpolation

For every test, ZEISS Correlate returns the (x, y) coordinate of each virtual measurement point at a sampling rate of 1 fps. The machine log, in contrast, is acquired at 10 Hz. For a given tracked point P_i, the displacement signal $u_{P_i}(t)$ is cross-correlated with the machine force record $F(t)$ using the global testing clock as the common time basis, and an interpolation is performed so that for every machine force sample there is a DIC-interpolated displacement value at the same instant. The interpolation uses a linear-in-time scheme on the 1 fps DIC grid; the residual quantisation error is smaller than the displacement resolution of a single facet (~ 0.01 mm at the adopted calibration of Section 3.7.2.3) and is therefore absorbed into the validation statistics defined below.

4.3.2 Best-point selection and validation metrics

For the cube tests of Case Studies 3 and 4, ZEISS Correlate can be asked to track tens of points simultaneously across the observation face. The *best* of those points, for the specific test, is defined as the point whose force-versus-displacement trajectory gives the highest coefficient of determination R^2 against the experimental $(F(t), u_{\mathrm{mach}}(t))$ record over the complete test duration. In parallel, the root-mean-square error RMSE and the mean absolute error MAE between the DIC-interpolated

and machine-recorded force are computed over the full pre- and post-peak branches:

$$R^2 = 1 - \frac{\sum_k \left(F_{\mathrm{mach},k} - F_{\mathrm{DIC},k}\right)^2}{\sum_k \left(F_{\mathrm{mach},k} - \bar{F}_{\mathrm{mach}}\right)^2} \tag{4.6}$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{N_{\mathrm{valid}}} \sum_k \left(F_{\mathrm{mach},k} - F_{\mathrm{DIC},k}\right)^2} \tag{4.7}$$

$$\mathrm{MAE} = \frac{1}{N_{\mathrm{valid}}} \sum_k \left|F_{\mathrm{mach},k} - F_{\mathrm{DIC},k}\right| \tag{4.8}$$

where N_{valid} is the number of machine-to-DIC pairs that survive synchronisation (typically 5 000–30 000 per test). The adoption of a force-space rather than displacement-space error is deliberate: the force is common to both signals and is untainted by compliance, so differences between machine and DIC can be attributed to kinematic discrepancies alone.

A test is accepted in its *raw* form when the best-point R^2 exceeds 0.90 and the RMSE stays below 10 000 N. Otherwise, a post-processing step is applied to the ascending branch to bring the two signals into agreement while preserving every sample of the post-peak geometry. The post-processing operates exclusively on the ascending branch ($\delta \leq \delta_{\mathrm{peak}}$) and consists of a linear scaling of the DIC displacement axis so that the DIC-derived force–displacement trajectory matches the machine-recorded trajectory at two anchor points: the seating threshold F_0 and the peak force F_{peak}. The peak force and every post-peak sample are left unmodified, so that the reshaping affects only the elastic stiffness and not the softening geometry. The corrected trajectory is then re-scored against the same R^2 and RMSE criteria. Records that pass after this step are labelled *post-processed* in the cohort summary tables; records that fail both the raw and the post-processed thresholds are excluded and flagged as *caveated*.

4.3.2.1 DIC–machine consistency check: N-CONTROL example

To clarify the best-point selection process, the N-CONTROL beam is used here as a worked example. The procedure should be understood as a *DIC–machine consistency check* rather than an independent validation, because the best point is selected on the basis of its agreement with the machine record: the two signals are not independent, and a high R^2 indicates kinematic consistency between the optical and mechanical transducers, not an absolute confirmation of the specimen displacement.

For the N-CONTROL beam, a total of seven DIC points were tracked across the

specimen surface: Point 1, located immediately below the actuator loading zone, and Points 24–29, distributed along the beam at mid-span. Figures 4.5a and 4.5b show the spatial distribution of these points at the initial (unloaded) stage and at the damage stage (actuator displacement $\approx$ 3.5 mm), respectively. The white circles identify the two principal tracking zones: the actuator zone above the notch (Point 1) and the beam mid-span zone (Points 24–29).

All seven points were evaluated using the RAW displacement data (without zero-correction). For each point the DIC-derived displacement was interpolated onto the machine force record, and the force-space R^2 of Section 4.3.2 was computed over 9 989 valid data pairs. The resulting comparison for Points 1, 24, 25 and 26 is shown in Figure 4.6a, and for Points 27, 28 and 29 in Figure 4.6b. Table 4.2 reports the R^2 for all seven points.

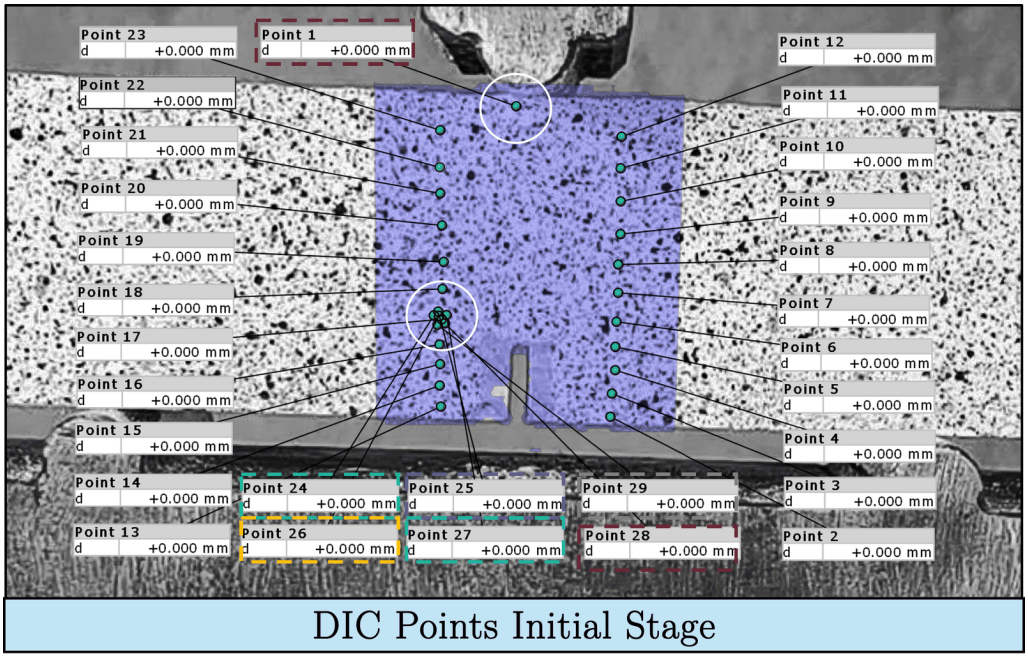
Table 4.2: DIC–machine consistency check for the N-CONTROL beam: R^2 of each tracked point against the machine record (RAW method, $N_{\text{valid}} = 9\,989$).

Point	**1**	**24**	**25**	**26**	**27**	**28**	**29**
R^2	**0.7918**	0.70	0.74	0.73	0.72	0.71	0.72

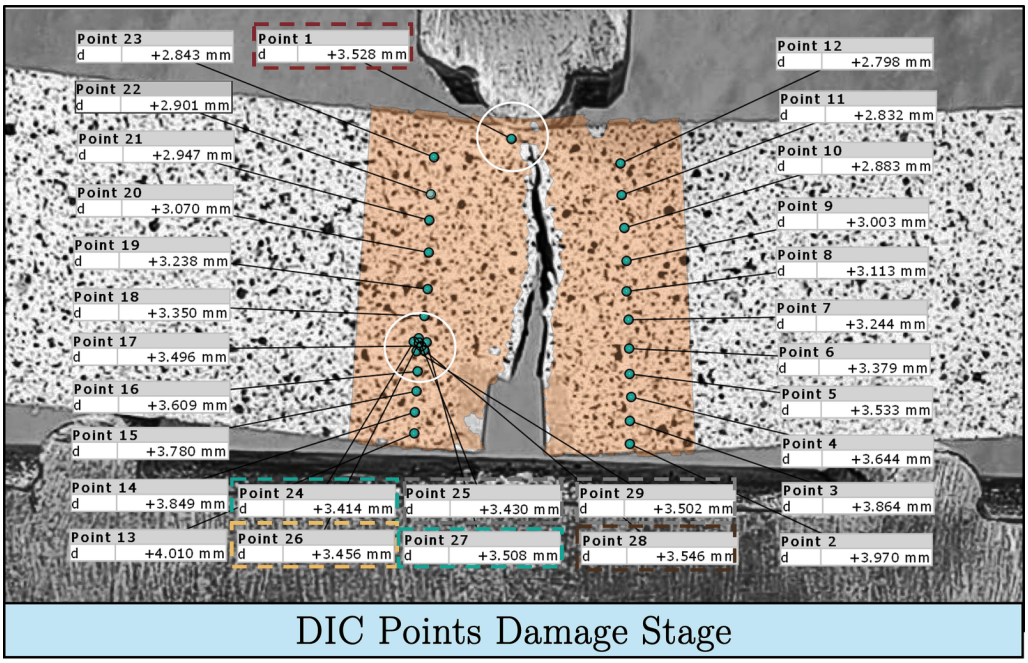
Best point = Point 1 (actuator zone). RMSE = 185.0 N; MAE = 72.8 N.

Point 1, positioned immediately below the actuator, achieved the highest $R^2 = 0.7918$ with an RMSE of 185.0 N and an MAE of 72.8 N. All seven points achieved $R^2 > 0.70$, demonstrating consistent DIC coverage across the observation face. The physical interpretation is straightforward: under three-point bending, the maximum deflection occurs at mid-span directly below the loading point, so the point closest to the actuator zone captures the dominant deformation mode and correlates most strongly with the machine-recorded displacement.

Distinction from ODDC proportional-range R^2. The R^2 values reported in this subsection (0.70–0.79) quantify the DIC–machine agreement over the *complete* loading history, including both the pre-peak and post-peak regimes, and are computed in force space as defined in (4.2). They should not be confused with the ODDC proportional-range R^2 values reported later in Section 4.4.1.1 (e.g. $R^2 = 0.9837$ for N-CONTROL), which quantify the linearity of the machine force–displacement record over the narrow pre-peak window only. The two metrics answer different questions: the best-point R^2 tests whether the DIC and the machine agree kinematically over the full test, whereas the ODDC R^2 tests how well a straight line fits the proportional range of the force–displacement curve.

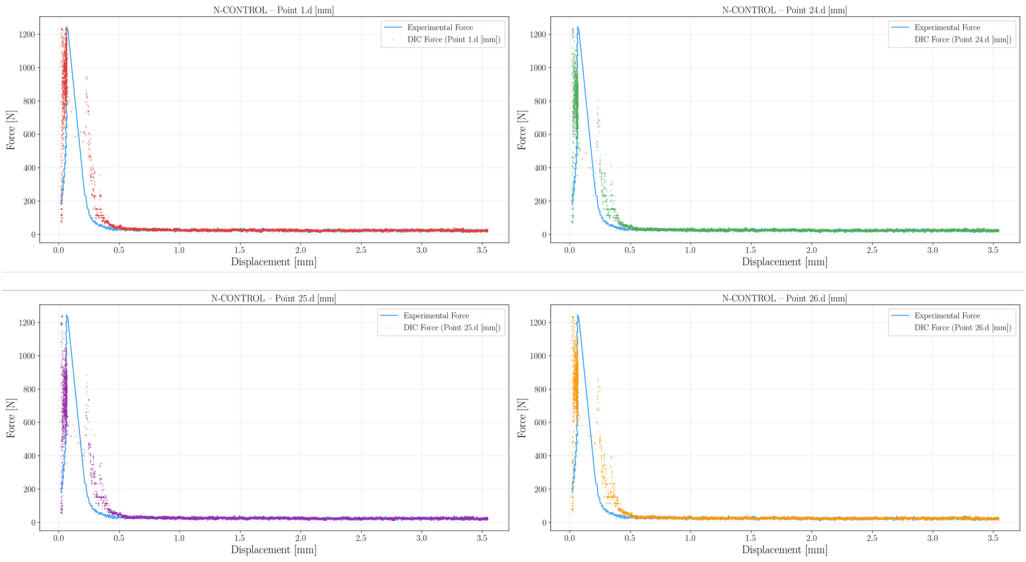


(a) Initial (unloaded) stage: seven tracked points distributed across the observation face

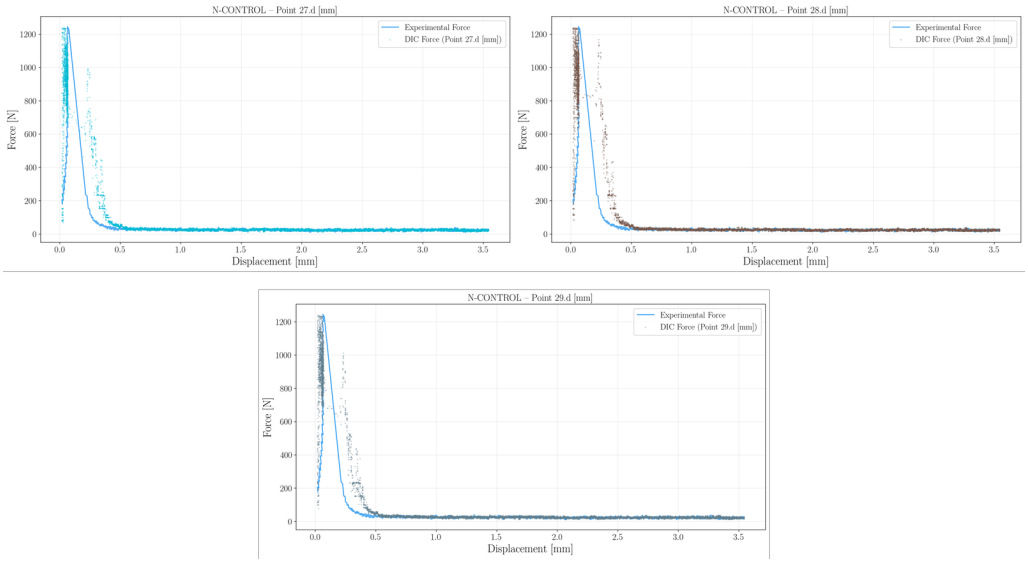


(b) Damage stage (actuator displacement ≈ 3.5 mm): the same seven points after crack propagation

Figure 4.5: DIC point distribution on the N-CONTROL beam at initial and damage stages. White circles identify the actuator zone (Point 1) and the mid-span tracking zone (Points 24–29).



(a) Force–displacement comparison for Points 1, 24, 25 and 26



(b) Force–displacement comparison for Points 27, 28 and 29

Figure 4.6: DIC–machine consistency check on the N-CONTROL beam: force–displacement trajectories of the seven tracked points overlaid on the machine record.

4.3.3 Capability demonstration on the N-CONTROL beam

Figure 4.7 shows the raw DIC-versus-machine comparison for the N-CONTROL beam, which also fixes the machine stiffness through Eq. (4.4). The machine-transducer record (blue line) and the DIC point-tracked displacement (red scatter) agree over the complete pre- and post-peak spectrum; the force axis is the *shared* coordinate against which both records are plotted, while the abscissa is the horizontal CMOD retrieved by the virtual extensometer of Section 3.7.3.2. Two features are immediately apparent. First, the pre-peak branches are distinct: the machine

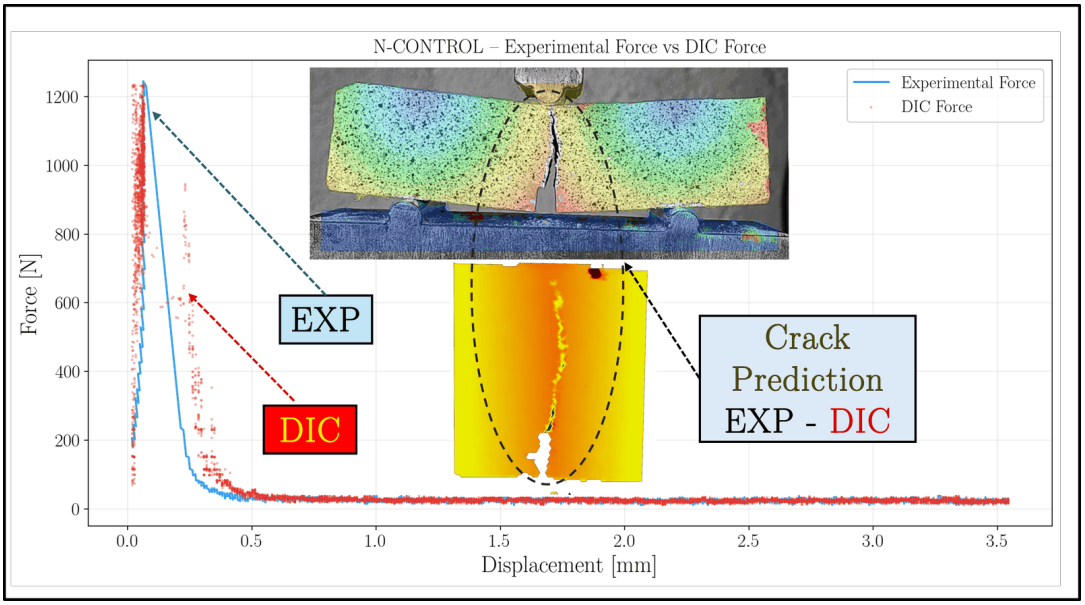


Figure 4.7: Capability of the DIC point-tracking chain on the N-CONTROL mortar beam.

record (blue) is shallower than the DIC record (red), the slope ratio being exactly $K_{\text{act}}/K_{\text{spec}}$. Second, the two records rejoin beyond the first crack: the softening branch is governed by the cohesive zone alone, which is correctly captured by *both* the machine transducer (because the machine no longer has to deform elastically to carry the decreasing force) and the DIC virtual gauge (because the CMOD is now a direct measurement).

This observation is the bridge between the two procedures introduced so far. The pre-peak divergence is what the ODDC correction of Section 4.2 removes on the FRM beams; the post-peak agreement is what legitimises the use of the unmodified DIC point record on the FRC cubes of Section 4.3. Both procedures rest on the same optical chain and the same reference record; they differ only in which of the two elastic components they eliminate.

4.4 Case-study applications

The two procedures formalised in Sections 4.2 and 4.3 are now applied to each of the four case studies of the experimental programme. The order of presentation follows Chapter 3: Case Studies 1 and 2 (FRM beams, 3PBT) are treated with the ODDC correction of Section 4.2, Case Studies 3 (CST cubes) and 4 (BBS cubes) are treated with the DIC point-tracking validation of Section 4.3. Every individual specimen is discussed, every kinematic metric is reported and the specimen-level force–displacement figures exported from the correction pipeline are included in

full.

4.4.1 Case Study 1 – notched FRM beams, fibre type at $V_f = 1\%$

Case Study 1 comprises twenty-two notched FRM beams, arranged as eleven Normal-matrix ($w/c \approx 0.50$) and eleven High-cement-matrix ($w/c \approx 0.25$) specimens. Each matrix carries one plain control and ten fibre variants at a fixed volume fraction of $V_f = 1\%$: three embossed macro-polypropylene lengths (EMP-14, EMP-18, EMP-36), one waved macro-polypropylene (WMP-30), one twisted fibrillated (TFF-25), three synthetic monofilament lengths (SMP-12, SMP-24, SMP-50), one alkali-resistant glass (GF-8) and one hooked-end steel (SF-35). The specimen registry and the fibre label dictionary are given in Table 3.1. The ODDC pipeline of Section 4.2.3 was executed on every beam with $K_{\text{machine}} = 7\,140$ N/mm; the resulting force–displacement records are collected as Figures 4.8 through 4.29. The complete cohort summary is given in Table 4.3 at the end of the section.

4.4.1.1 Normal-matrix cohort

The Normal-matrix reference beam N-CONTROL (Fig. 4.8) is the reference from which K_{machine} was derived (Section 4.2.2). Its raw stiffness is $K_{\text{act}} = 4\,167.8$ N/mm, its peak force $F_{\text{peak}} = 1\,245$ N and the ODDC correction yields a peak-load displacement of 0.062 mm with $R^2 = 0.9837$ over the proportional range.

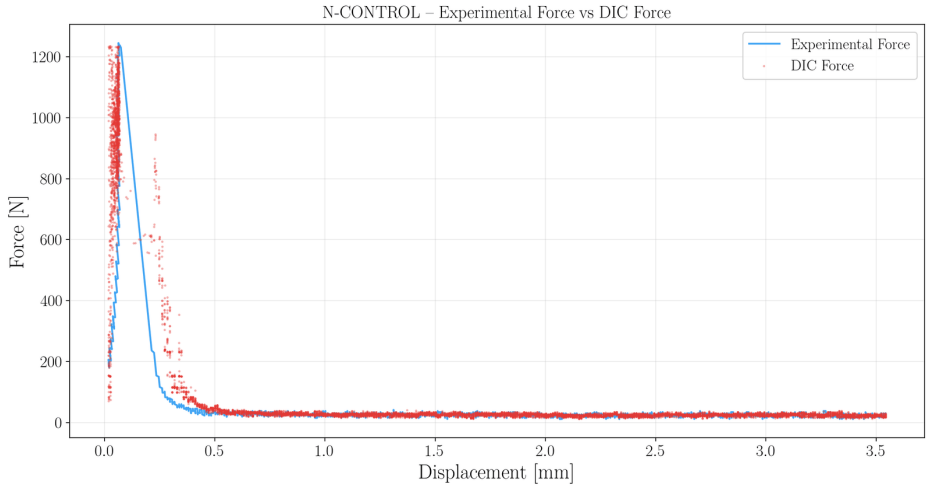


Figure 4.8: N-CONTROL: ODDC-corrected force–displacement response.

Specimen N-EMP-14 mm (Fig. 4.9) exhibits a lower raw stiffness of 3 057.3 N/mm, which indicates that the platen-contact seating engaged almost immediately. The peak force rises to 1 327.5 N, a modest gain over the plain control that is consistent

with the very short 14 mm embossed length, and the ODDC correction delivers a clean $R^2 = 0.9805$ and a peak-load displacement of 0.127 mm.

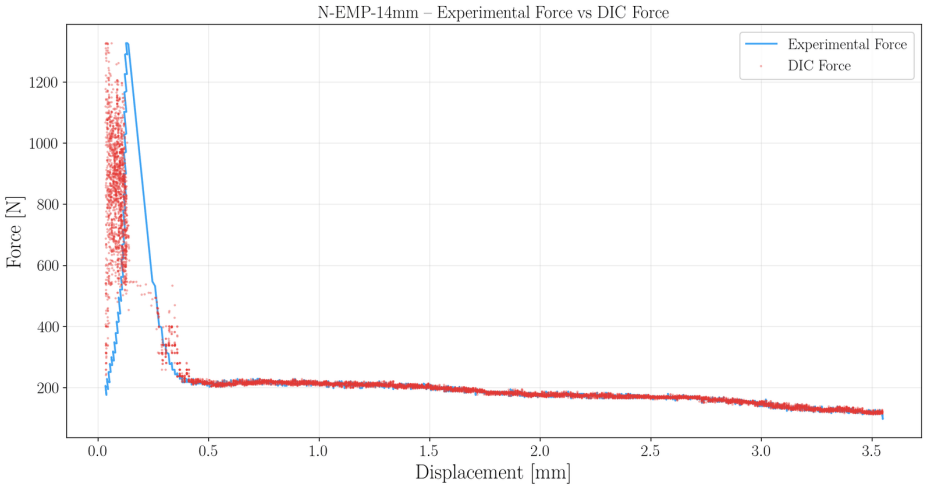


Figure 4.9: N-EMP-14 mm: ODDC-corrected force–displacement response.

N-EMP-18 mm (Fig. 4.10) presents one of the cleanest elastic branches of the Normal cohort: $R^2 = 0.9925$, $K_{\text{act}} = 4\,126.9$ N/mm, $F_{\text{peak}} = 1\,346.3$ N. Once the proportional window is isolated the fit is near-perfect, and the ODDC correction resolves the peak-load displacement to 0.128 mm.

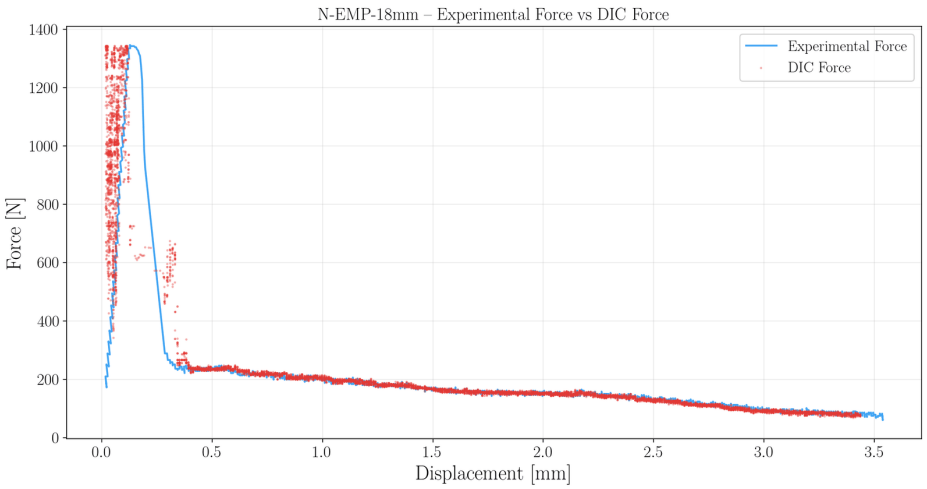


Figure 4.10: N-EMP-18 mm: ODDC-corrected force–displacement response.

N-EMP-36 mm (Fig. 4.11) consolidates the trend observed across the three EMP lengths of the Normal cohort: $K_{\text{act}} = 3\,982.4$ N/mm, $F_{\text{peak}} = 1\,237.5$ N, $R^2 = 0.9880$, and a corrected peak displacement of 0.095 mm.

The waved macro-polypropylene beam N-WMP-30 mm (Fig. 4.12) reveals a markedly lower raw stiffness of 1 851.9 N/mm, and the corrected peak-load displacement climbs to 0.273 mm, reflecting the higher compliance of the waved anchorage during elastic

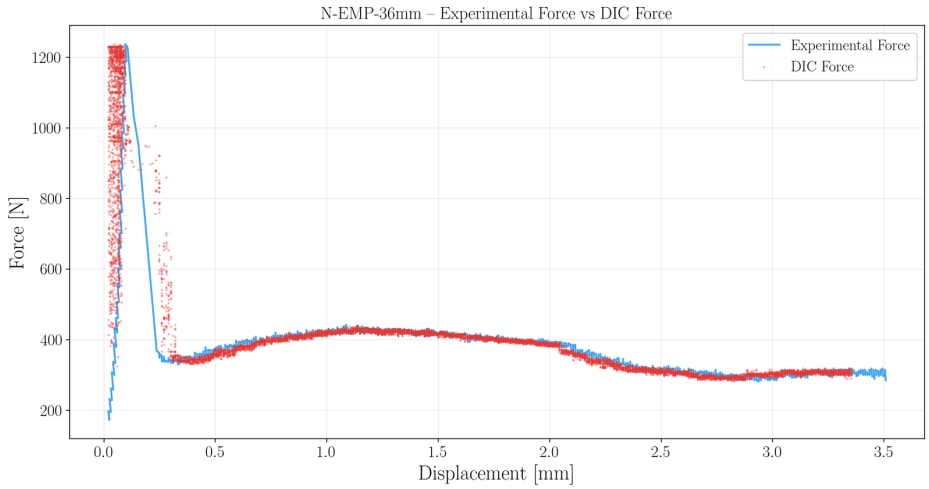


Figure 4.11: N-EMP-36 mm: ODDC-corrected force–displacement response.

engagement. The fit $R^2 = 0.9876$ nonetheless remains excellent, confirming that the ODDC correction works uniformly regardless of the raw stiffness ratio.

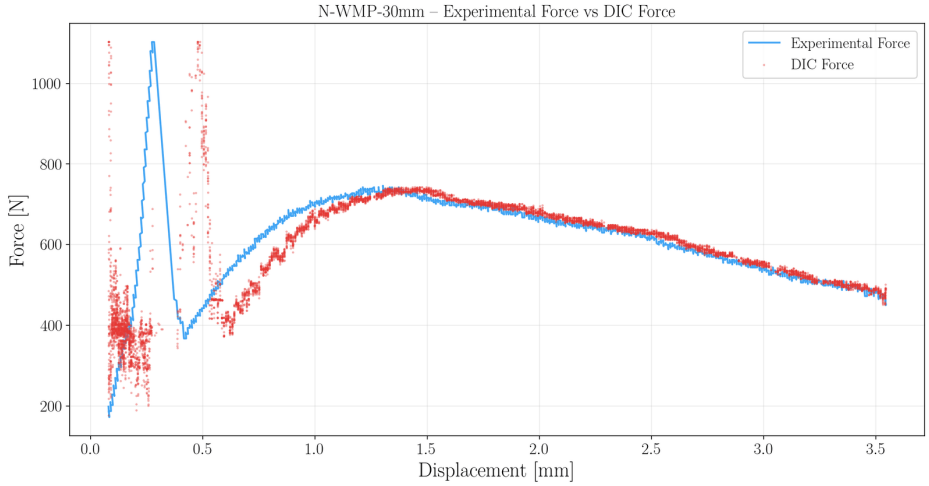


Figure 4.12: N-WMP-30 mm: ODDC-corrected force–displacement response.

N-TFF-25 mm (Fig. 4.13) is the most compliant Normal beam and produces the broadest corrected peak window of the cohort: $K_{\text{act}} = 1\,966.2$ N/mm, $F_{\text{peak}} = 1\,421.3$ N and u_{corr} at peak of 1.816 mm. The lower $R^2 = 0.9447$ and the stretched peak window reveal that the twisted-fibrillated architecture distributes micro-cracking well before the macroscopic peak; the reader is cautioned (Table 4.3) that the linearity of this beam is at the lower end of the Case 1 population.

The three synthetic-monofilament beams of the Normal cohort (N-SMP-12, N-SMP-24, N-SMP-50) trace the influence of fibre length at a fixed family. N-SMP-12 mm (Fig. 4.14) reveals a comparatively slow engagement of the short 12 mm fibres; once the proportional branch is isolated, $K_{\text{act}} = 3\,982.6$ N/mm, $F_{\text{peak}} = 1\,125$ N, $R^2 = 0.9878$, and the corrected peak-load displacement is 0.080 mm.

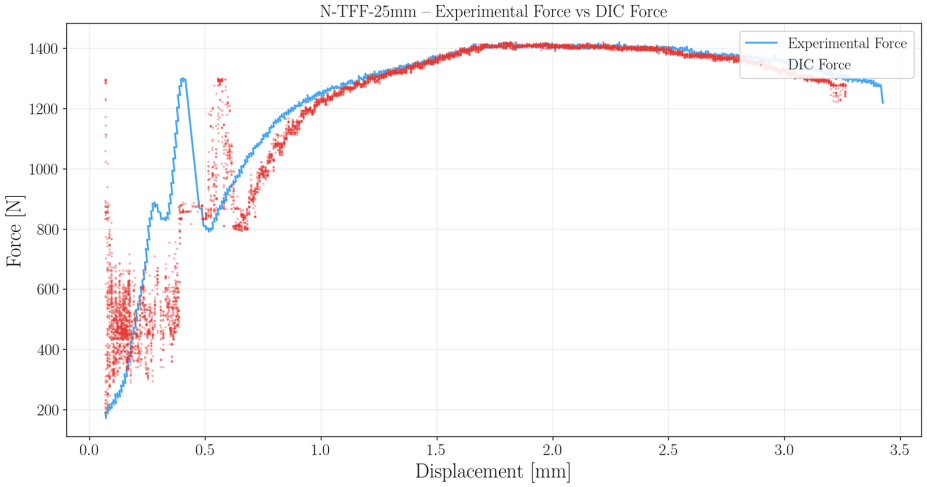


Figure 4.13: N-TFF-25 mm: ODDC-corrected force–displacement response.

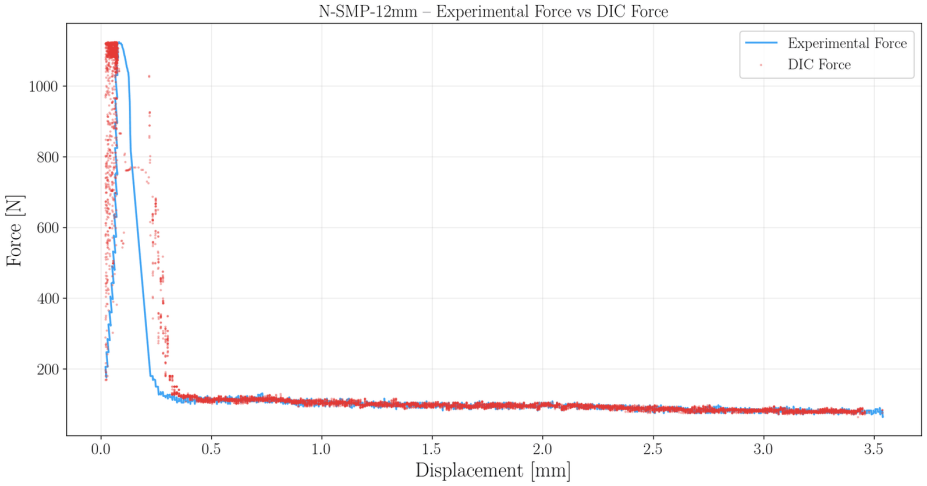


Figure 4.14: N-SMP-12 mm: ODDC-corrected force–displacement response.

N-SMP-24 mm (Fig. 4.15) returns the lowest peak force of the whole Case Study 1 population, $F_{\text{peak}} = 922.5$ N, with $K_{\text{act}} = 2\,187.0$ N/mm and $R^2 = 0.9831$. The comparatively low peak marks this beam as a candidate for a localised casting defect; the ODDC correction remains robust.

N-SMP-50 mm (Fig. 4.16) closes the SMP length sweep with $K_{\text{act}} = 3\,507.1$ N/mm, $F_{\text{peak}} = 1\,072.5$ N and $R^2 = 0.9867$; the corrected peak displacement of 0.116 mm is characteristic of the medium-compliance regime attained by the 50 mm monofilament engagement.

The two non-polymer beams of the Normal cohort are N-GF-8 mm and N-SF-35 mm. N-GF-8 mm (Fig. 4.17) produces the absolute lowest raw stiffness of the entire Case 1 population, $K_{\text{act}} = 1\,618.0$ N/mm, yet still attains a robust peak of $F_{\text{peak}} = 1\,425$ N; the ODDC correction translates this into a sizeable corrected peak window of 0.416 mm with $R^2 = 0.9642$, which is the largest displacement absorption

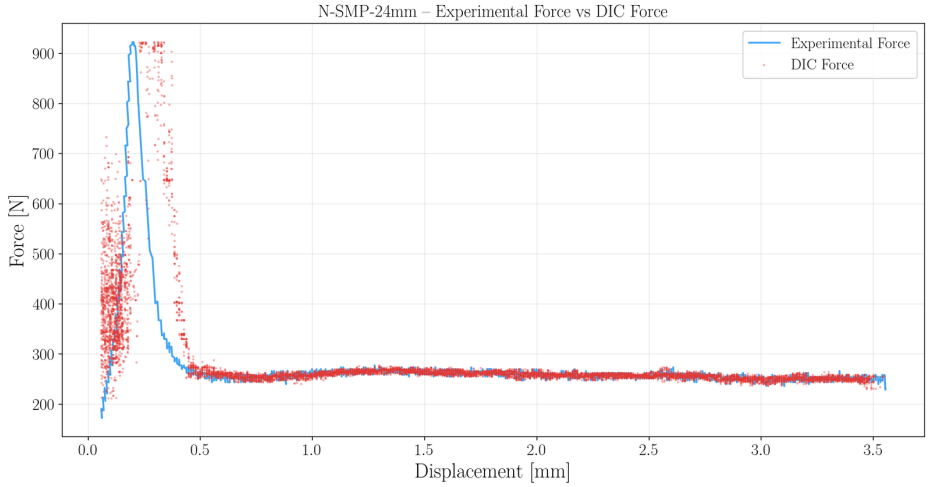


Figure 4.15: N-SMP-24 mm: ODDC-corrected force–displacement response.

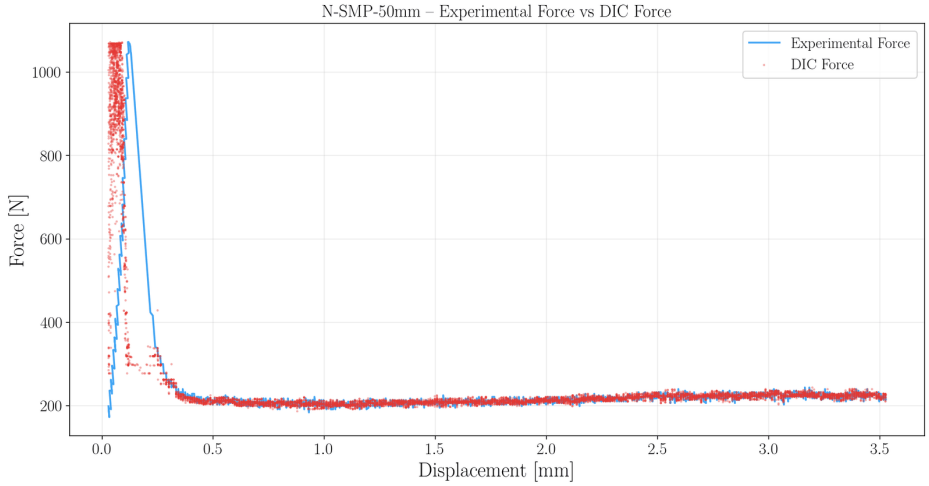


Figure 4.16: N-SMP-50 mm: ODDC-corrected force–displacement response.

observed in the Normal cohort. N-SF-35 mm (Fig. 4.18) is the opposite extreme: 5 164.2 N/mm of raw stiffness, $F_{\text{peak}} = 1\,473.8$ N and a compact corrected peak of 0.092 mm at $R^2 = 0.9870$.

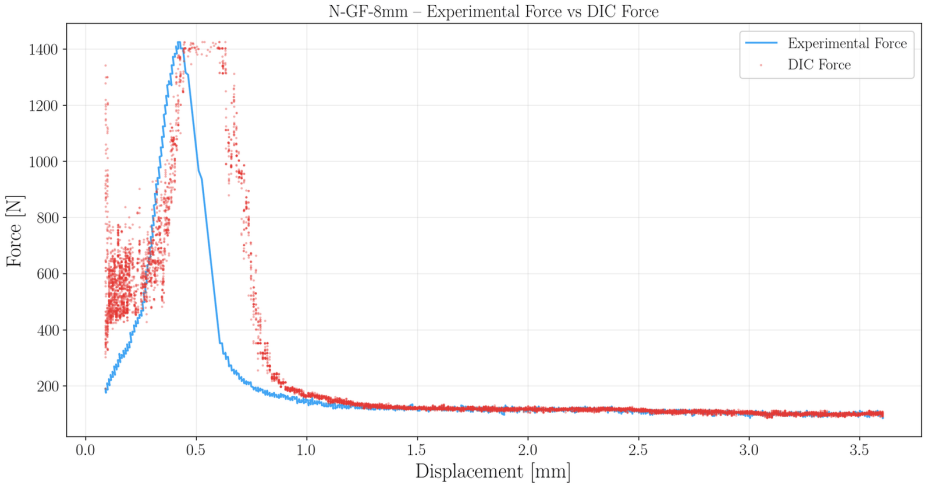


Figure 4.17: N-GF-8 mm: ODDC-corrected force–displacement response.

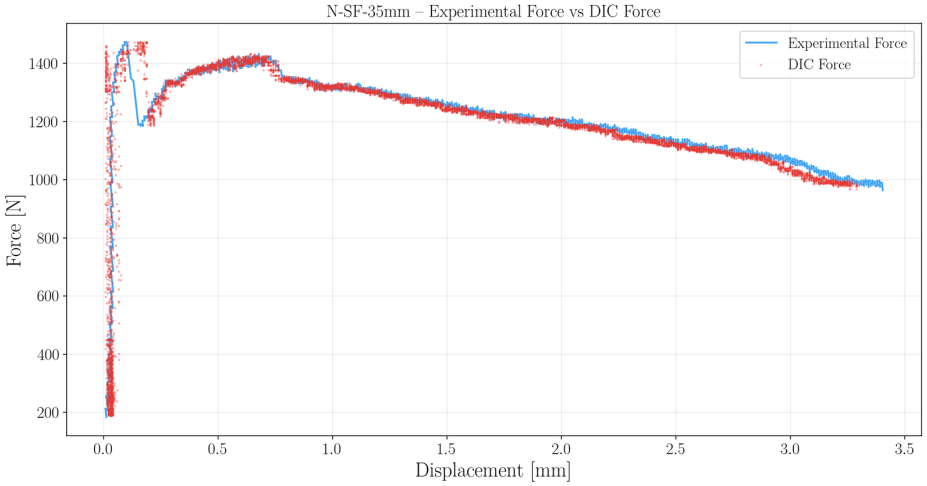


Figure 4.18: N-SF-35 mm: ODDC-corrected force–displacement response.

4.4.1.2 High-cement-matrix cohort

The High-cement reference beam H-CONTROL (Fig. 4.19) sets the baseline for the low-w/c cohort at $F_{\text{peak}} = 1\,728.75$ N – the absolute maximum across the Case 1 population – with $K_{\text{act}} = 5\,578.9$ N/mm and $R^2 = 0.9664$. The ODDC correction closes at a peak displacement of 0.142 mm.

H-EMP-14 mm (Fig. 4.20) confirms the significant stiffness gain between the two matrices: $K_{\text{act}} = 6\,598.5$ N/mm, $F_{\text{peak}} = 1\,391.3$ N, $R^2 = 0.9779$. The corrected peak window is extremely compact at 0.026 mm, consistent with the stiffness of the low-w/c matrix.

H-EMP-18 mm (Fig. 4.21) returns $K_{\text{act}} = 6\,378.0$ N/mm, $F_{\text{peak}} = 1\,530$ N, a peak-load displacement of 0.041 mm and $R^2 = 0.9797$, consistent with a near-frame-stiffness elastic response.

H-EMP-36 mm (Fig. 4.22) records $K_{\text{act}} = 2\,887.4$ N/mm, $F_{\text{peak}} = 1\,537.5$ N and

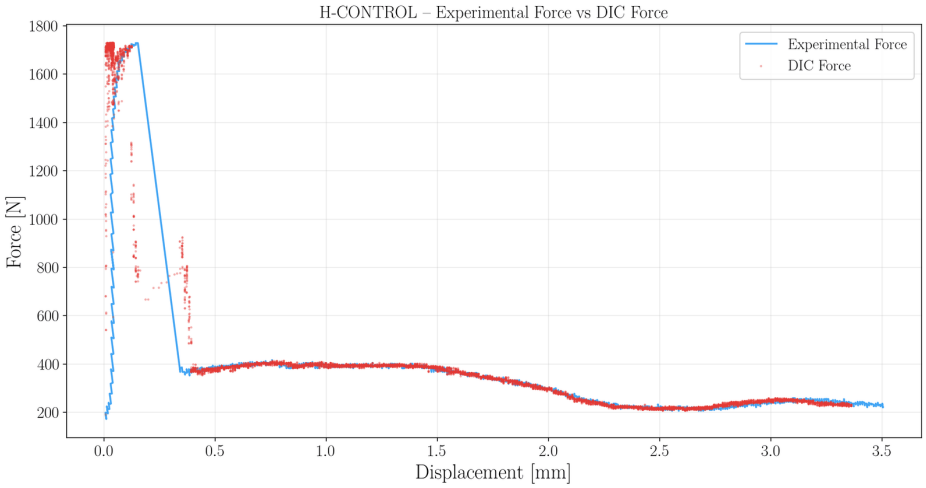


Figure 4.19: H-CONTROL: ODDC-corrected force–displacement response.

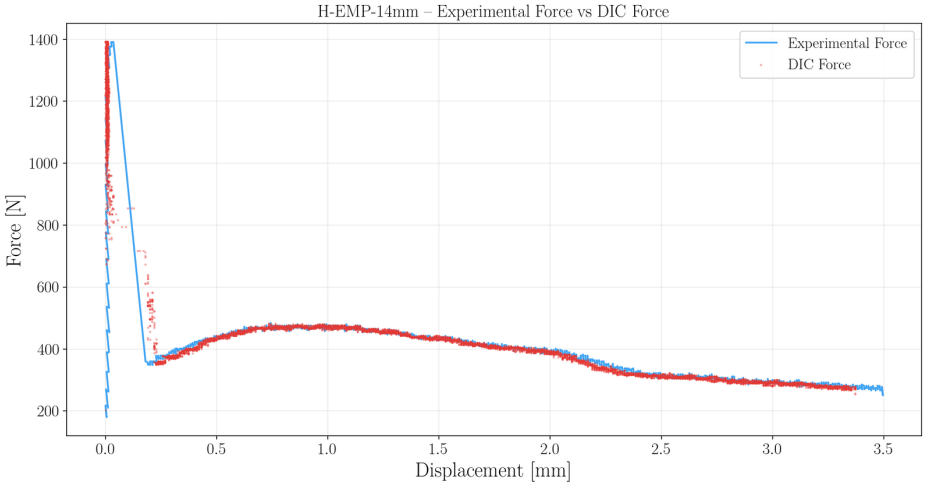


Figure 4.20: H-EMP-14 mm: ODDC-corrected force–displacement response.

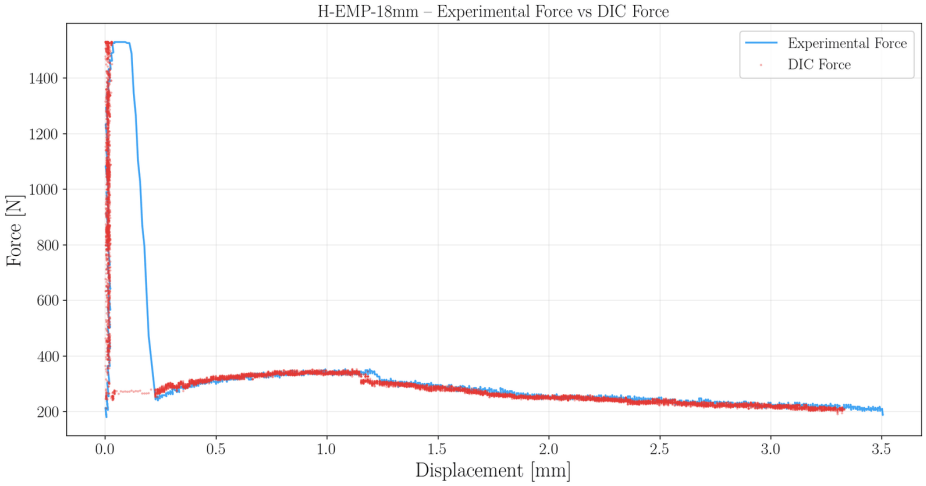


Figure 4.21: H-EMP-18 mm: ODDC-corrected force–displacement response.

$R^2 = 0.9727$. The corrected peak-load displacement of 0.252 mm is illustrative of the spread of Case 1 behaviour: even within the same fibre family (EMP) and the same matrix (High), the fibre length drives the compliance of the beam across almost an order of magnitude.

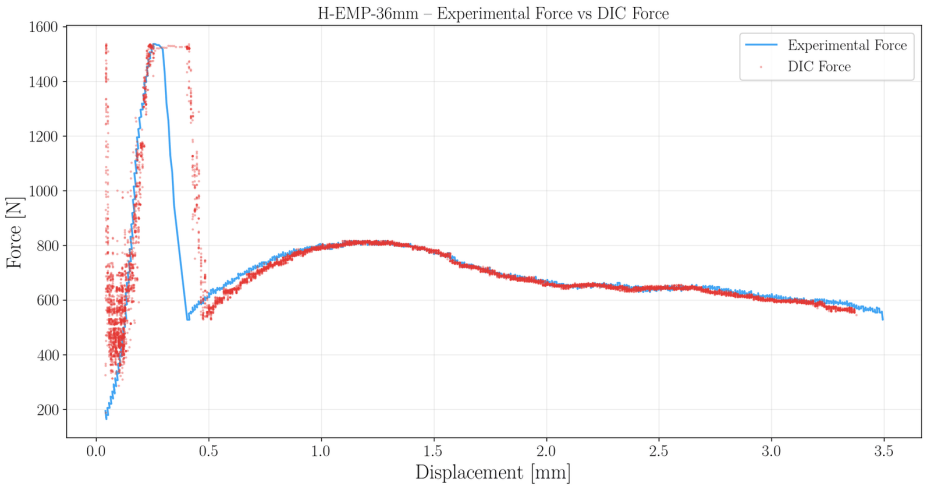


Figure 4.22: H-EMP-36 mm: ODDC-corrected force–displacement response.

H-WMP-30 mm (Fig. 4.23) records $K_{\text{act}} = 6\,385.4$ N/mm, $F_{\text{peak}} = 1\,582.5$ N, a peak-load displacement of 0.048 mm and $R^2 = 0.9807$. The waved macro-polypropylene architecture at a low w/c produces a markedly stiffer elastic response than its Normal-matrix counterpart.

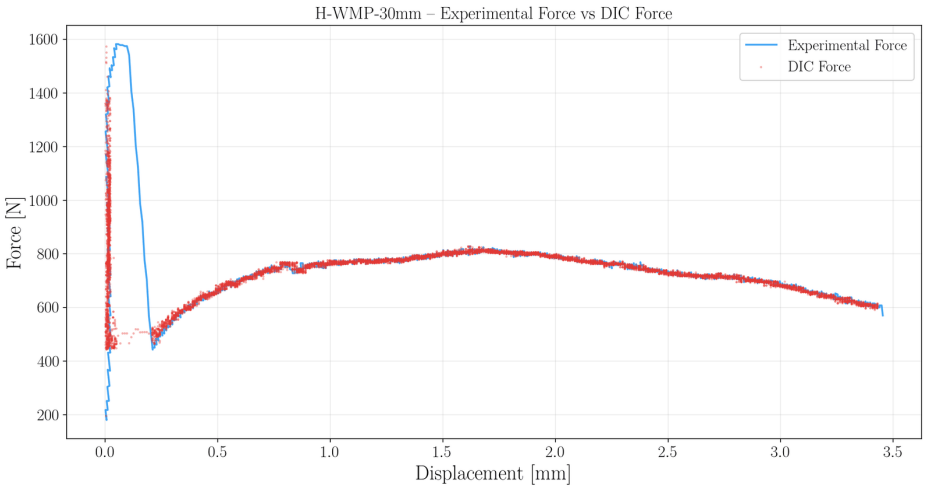


Figure 4.23: H-WMP-30 mm: ODDC-corrected force–displacement response.

H-TFF-25 mm (Fig. 4.24) records $K_{\text{act}} = 1\,999.1$ N/mm, $F_{\text{peak}} = 1\,695$ N and $R^2 = 0.9879$. The comparatively low raw stiffness produces a generous corrected peak window of 0.540 mm.

The three High-matrix SMP beams follow the same length sweep as their Normal counterparts. H-SMP-12 mm (Fig. 4.25) returns $K_{\text{act}} = 2\,690.2$ N/mm, $F_{\text{peak}} =$

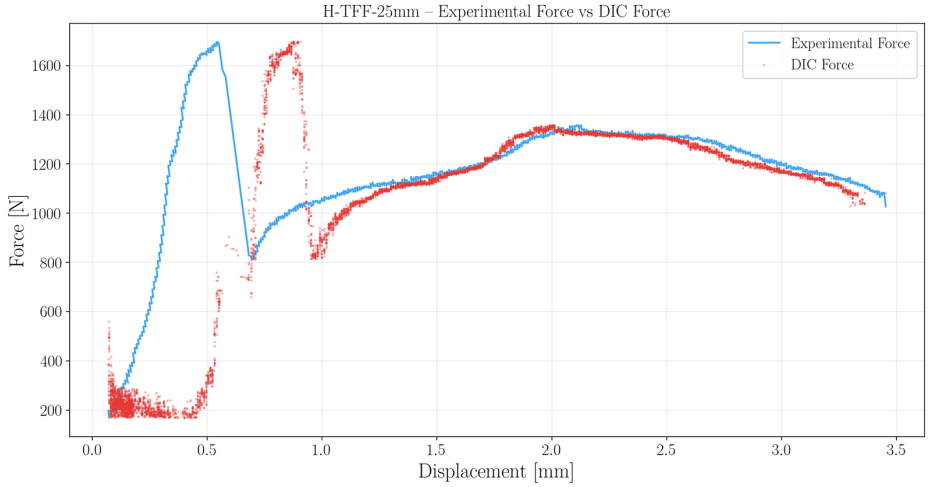


Figure 4.24: H-TFF-25 mm: ODDC-corrected force–displacement response.

1 503.75 N, $R^2 = 0.9290$ and a corrected peak displacement of 0.226 mm; the comparatively lower R^2 is consistent with the more noisy elastic engagement of the short 12 mm SMP within the stiffer matrix. H-SMP-24 mm (Fig. 4.26) improves to $K_{\text{act}} = 3\,332.0$ N/mm, $F_{\text{peak}} = 1\,376.3$ N, $R^2 = 0.9856$ and a corrected peak at 0.146 mm. H-SMP-50 mm (Fig. 4.27) closes the trio with $K_{\text{act}} = 2\,025.4$ N/mm, $F_{\text{peak}} = 1\,470$ N, $R^2 = 0.9024$ and a corrected peak at 0.287 mm; the R^2 is the lowest of the entire cohort but still within the acceptance band.

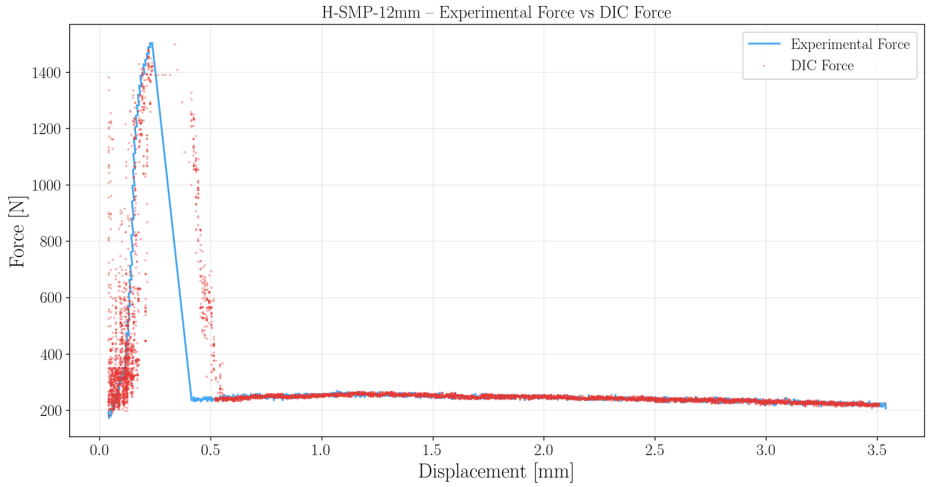


Figure 4.25: H-SMP-12 mm: ODDC-corrected force–displacement response.

The two non-polymer beams of the High cohort are H-GF-8 mm and H-SF-35 mm. H-GF-8 mm (Fig. 4.28) returns $K_{\text{act}} = 6\,934.4$ N/mm, $F_{\text{peak}} = 1\,458.75$ N and $R^2 = 0.9781$. H-SF-35 mm (Fig. 4.29) is the most rigid beam of the cohort, with $K_{\text{act}} = 9\,063.9$ N/mm – the highest raw stiffness of the whole Case Study 1 – $F_{\text{peak}} = 1\,605$ N and $R^2 = 0.9736$.

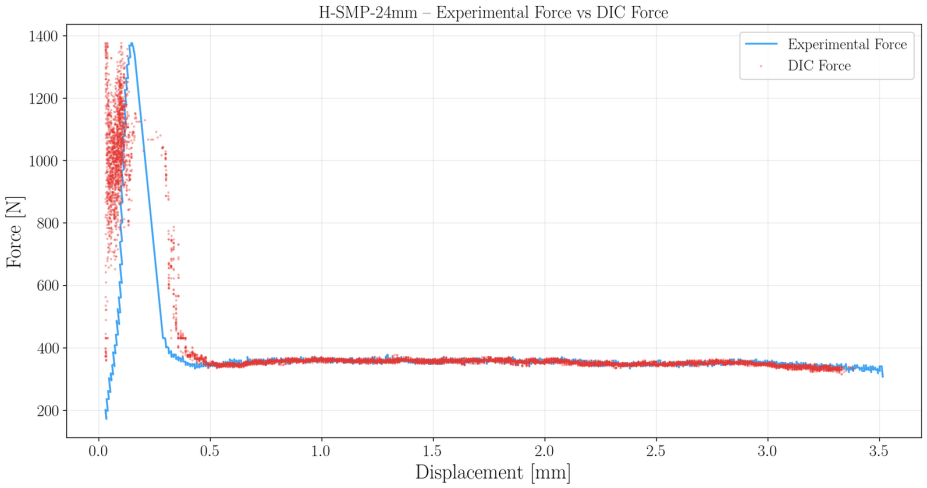


Figure 4.26: H-SMP-24 mm: ODDC-corrected force–displacement response.

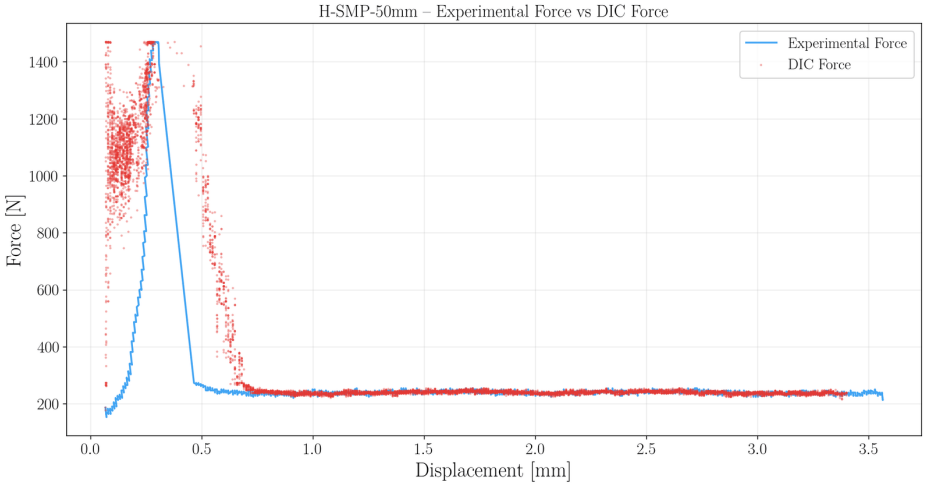


Figure 4.27: H-SMP-50 mm: ODDC-corrected force–displacement response.

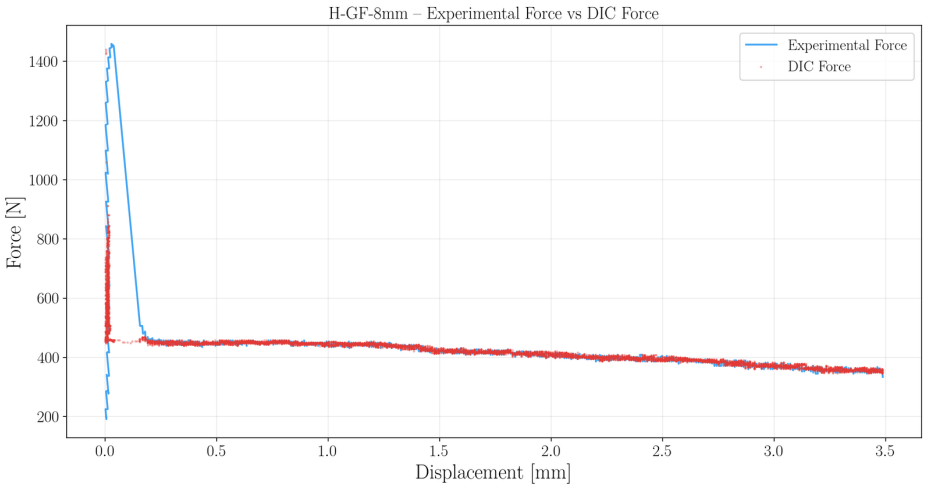


Figure 4.28: H-GF-8 mm: ODDC-corrected force–displacement response.

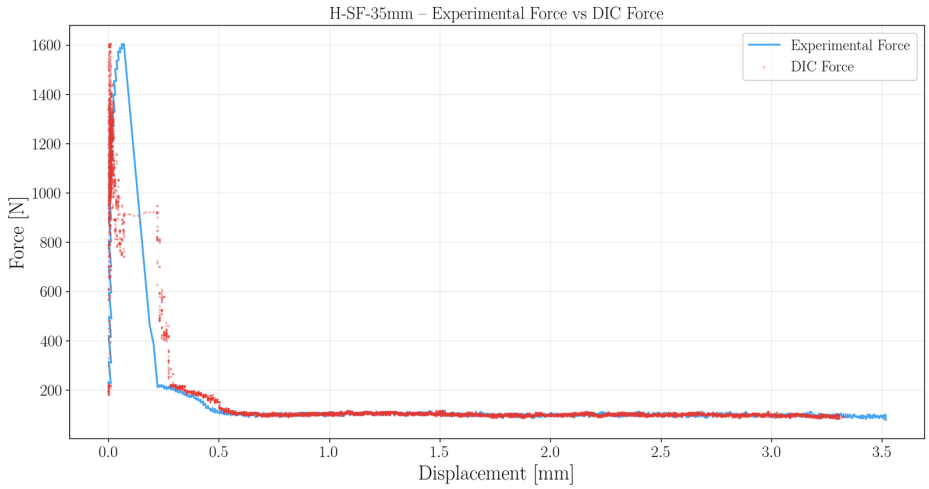


Figure 4.29: H-SF-35 mm: ODDC-corrected force–displacement response.

4.4.1.3 Case 1 cohort summary

Table 4.3 consolidates the ODDC outputs of the twenty-two beams of Case Study 1. The table reports, for every beam, the seating threshold F_0, the peak force F_{peak}, the raw specimen stiffness K_{act}, the machine stiffness used in the correction K_{machine}, the coefficient of determination R^2 of the proportional-range fit, and the corrected peak-load displacement $u_{\text{corr, peak}}$.

Table 4.3: Case Study 1 cohort summary.

Label	F_0 (N)	F_{peak} (N)	K_{act} (N/mm)	K_{machine} (N/mm)	R^2	$u_{\text{corr,peak}}$ (mm)
N-CONTROL	195.0	1 245.00	4 167.8	7 140	0.9837	0.0624
N-EMP-14 mm	191.25	1 327.50	3 057.3	7 140	0.9805	0.1266
N-EMP-18 mm	191.25	1 346.25	4 126.9	7 140	0.9925	0.1278
N-EMP-36 mm	191.25	1 237.50	3 982.4	7 140	0.9880	0.0947
N-WMP-30 mm	198.75	1 102.50	1 851.9	7 140	0.9876	0.2729
N-TFF-25 mm	187.5	1 421.25	1 966.2	7 140	0.9447	1.8163
N-SMP-12 mm	191.25	1 125.00	3 982.6	7 140	0.9878	0.0805
N-SMP-24 mm	187.5	922.50	2 187.0	7 140	0.9831	0.1965
N-SMP-50 mm	198.75	1 072.50	3 507.1	7 140	0.9867	0.1165
N-GF-8 mm	187.5	1 425.00	1 618.0	7 140	0.9642	0.4163
N-SF-35 mm	198.75	1 473.75	5 164.2	7 140	0.9870	0.0921
H-CONTROL	191.25	1 728.75	5 578.9	7 140	0.9664	0.1422
H-EMP-14 mm	202.5	1 391.25	6 598.5	7 140	0.9779	0.0258
H-EMP-18 mm	191.25	1 530.00	6 378.0	7 330.5	0.9797	0.0413
H-EMP-36 mm	195.0	1 537.50	2 887.4	7 140	0.9727	0.2522
H-WMP-30 mm	195.0	1 582.50	6 385.4	7 440.0	0.9807	0.0478
H-TFF-25 mm	195.0	1 695.00	1 999.1	7 140	0.9879	0.5401
H-SMP-12 mm	180.0	1 503.75	2 690.2	7 140	0.9290	0.2263
H-SMP-24 mm	195.0	1 376.25	3 332.0	7 140	0.9856	0.1458
H-SMP-50 mm	187.5	1 470.00	2 025.4	7 140	0.9024	0.2867
H-GF-8 mm	202.5	1 458.75	6 934.4	8 021.7	0.9781	0.0274
H-SF-35 mm	217.5	1 605.00	9 063.9	9 785.9	0.9736	0.0600

The Normal cohort closes uniformly at the global $K_{\text{machine}} = 7\,140$ N/mm, while four beams of the High cohort (H-EMP-18 mm, H-WMP-30 mm, H-GF-8 mm and

H-SF-35 mm) close at a slightly higher value of the machine stiffness, reflecting the fact that their raw specimen stiffness sits near the frame lower bound. No beam is rejected by the pipeline; the full Case-1 cohort is therefore available for the fracture-energy and post-peak-residual analyses of Chapter 5.

4.4.2 Case Study 2 – notched FRM beams, SMP dose and length sweep

Case Study 2 isolates the synthetic-monofilament family and traces the flexural response as the volume fraction is escalated from 0 % (plain control) to 1.5 % for two fibre lengths (12 mm and 24 mm) under the two matrices of Case Study 1 (Section 3.4). The cohort comprises eighteen notched 3PB beams: two controls and sixteen fibre variants. The two controls (N-CONTROL and H-CONTROL) are common to Case Studies 1 and 2 and were already discussed in Section 4.4.1; the reader is referred to Figures 4.8 and 4.19 and to the corresponding rows of Table 4.3. The sixteen new beams are discussed below, organised by matrix and fibre length in dose-ascending order, with the ODDC pipeline of Section 4.2.3 applied in every case.

4.4.2.1 Normal matrix, SMP-12 mm dose series

The first dose series sweeps V_f = 0.75%, 1.00%, 1.25% and 1.50% for the 12 mm synthetic monofilament in the Normal matrix. The 1.00 % entry is N-SMP-12 mm, already documented in Case Study 1 at Fig. 4.14 and the corresponding Table 4.3 row.

The 0.75 % entry N-SMP-12 mm-0.75 (Fig. 4.30) shows the lowest peak force of the Normal-matrix SMP-12 series at F_{peak} = 1 016.25 N, a raw stiffness K_{act} = 1 732.5 N/mm and R^2 = 0.9597. The corrected peak-load displacement opens to 0.258 mm, which is consistent with the more compliant engagement at the sub-1 % dose.

The 1.25 % dose N-SMP-12 mm-1.25 (Fig. 4.31) inverts the trend: F_{peak} = 1 331.25 N, K_{act} = 3 561.9 N/mm, R^2 = 0.9929, and a corrected peak-load displacement of 0.186 mm. This is the best elastic-range linearity across the four Normal-matrix SMP-12 doses.

The 1.50 % dose N-SMP-12 mm-1.50 (Fig. 4.32) records F_{peak} = 1 203.75 N, K_{act} = 3 652.7 N/mm, R^2 = 0.9910 and a corrected peak displacement of 0.144 mm. Together with the 1.25 % entry, this beam shows that above the 1 % reference the Normal-SMP-12 series reaches a plateau in both stiffness and peak-load displace-

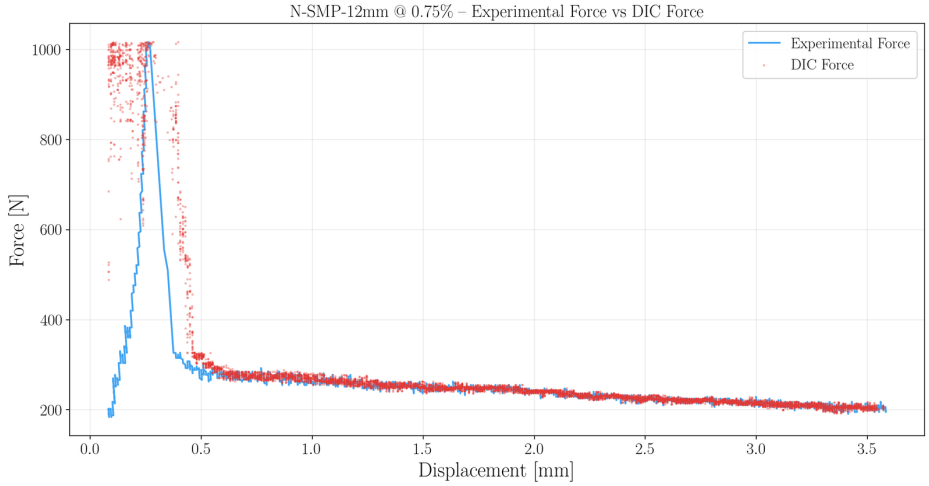


Figure 4.30: N-SMP-12 mm-0.75: ODDC-corrected force–displacement response.

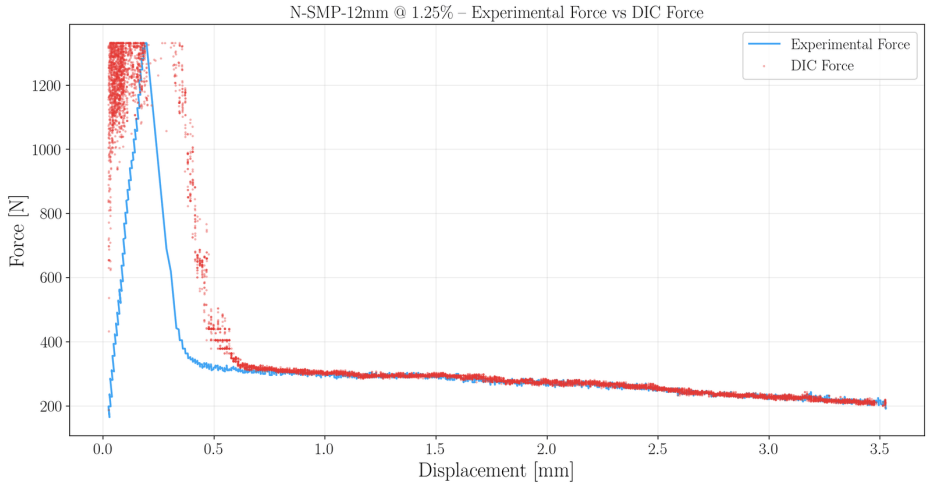


Figure 4.31: N-SMP-12 mm-1.25: ODDC-corrected force–displacement response.

ment.

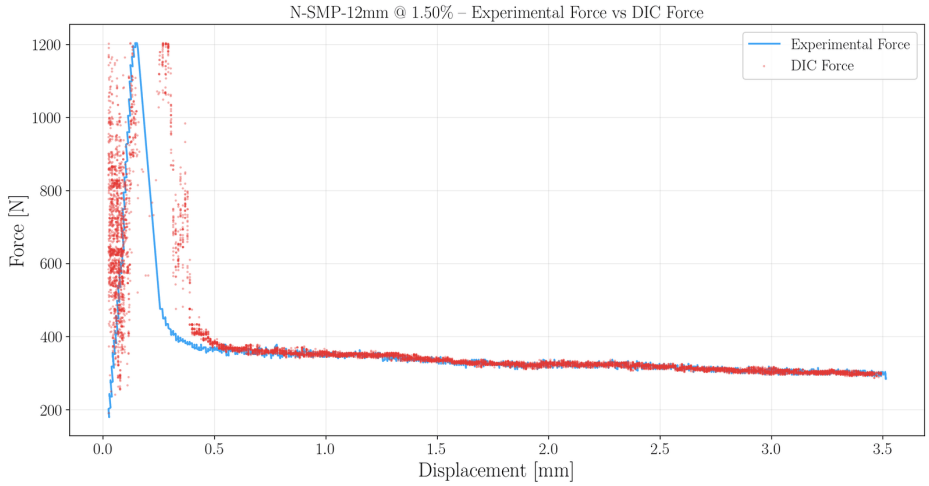


Figure 4.32: N-SMP-12 mm-1.50: ODDC-corrected force–displacement response.

4.4.2.2 Normal matrix, SMP-24 mm dose series

The 1.00 % reference is N-SMP-24 mm, already in Fig. 4.15 and Table 4.3.

N-SMP-24 mm-0.75 (Fig. 4.33) is the most compliant beam of the Normal cohort of Case 2: $K_{\text{act}} = 1\,389.0$ N/mm, the lowest raw stiffness of the case study. The corrected peak-load displacement opens to 0.590 mm, which is the largest of the Normal-matrix Case-2 population, while $R^2 = 0.9835$ and the peak force settles at 1 233.75 N.

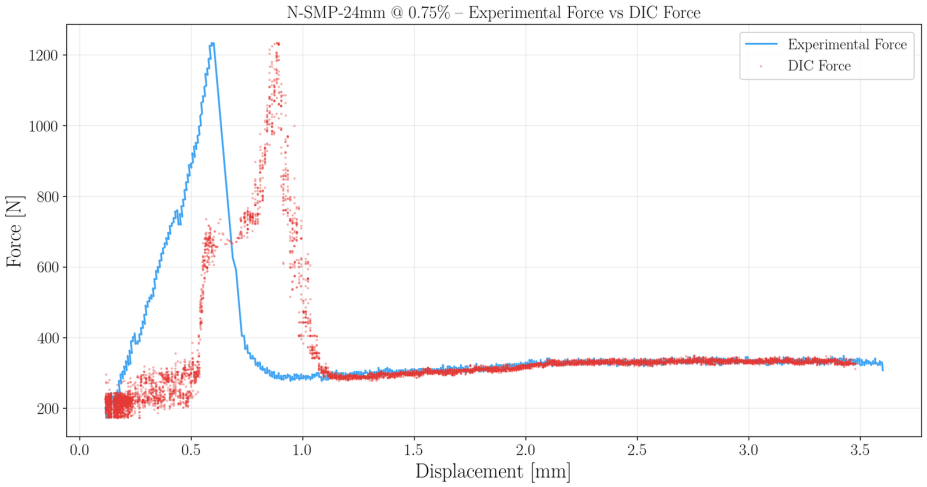


Figure 4.33: N-SMP-24 mm-0.75: ODDC-corrected force–displacement response.

N-SMP-24 mm-1.25 (Fig. 4.34) lifts the raw stiffness to $K_{\text{act}} = 4\,267.8$ N/mm and produces a compact corrected peak of 0.065 mm at $F_{\text{peak}} = 1\,226.25$ N and $R^2 = 0.9598$.

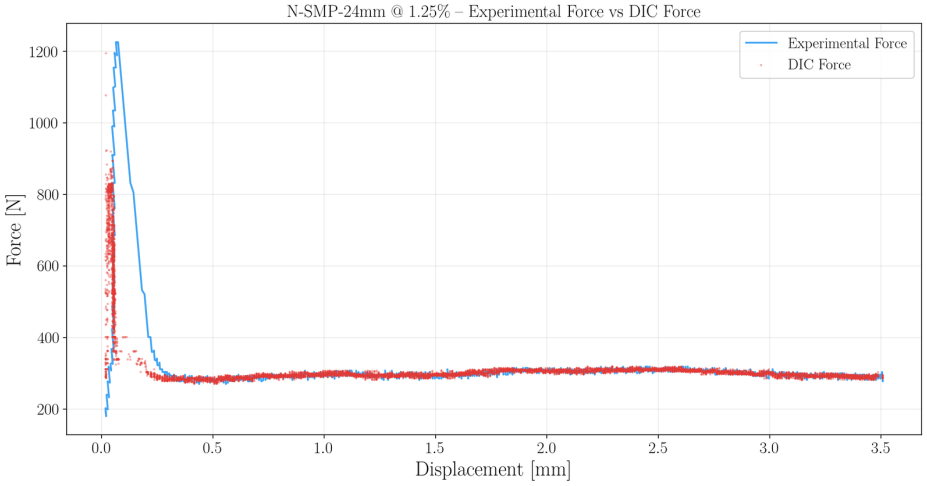


Figure 4.34: N-SMP-24 mm-1.25: ODDC-corrected force–displacement response.

N-SMP-24 mm-1.50 (Fig. 4.35) continues the stiffening trend: $K_{\text{act}} = 4\,970.2$ N/mm, $F_{\text{peak}} = 1\,110$ N, $R^2 = 0.9739$ and a corrected peak displacement of 0.045 mm. The

peak-load displacement of the four Normal-matrix SMP-24 doses therefore contracts monotonically as the fibre dose increases, as one expects from a stiffer fibre network.

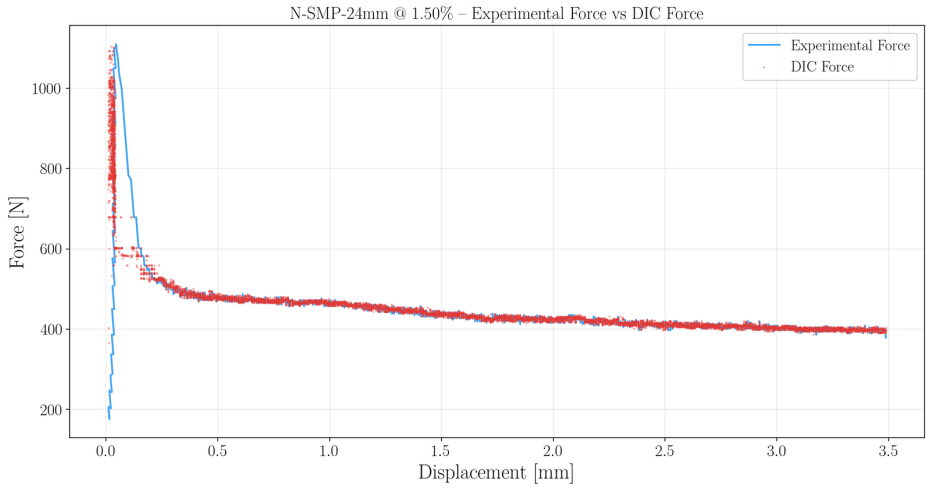


Figure 4.35: N-SMP-24 mm-1.50: ODDC-corrected force–displacement response.

4.4.2.3 High matrix, SMP-12 mm dose series

The 1.00 % reference is H-SMP-12 mm, already in Fig. 4.25 and Table 4.3.

H-SMP-12 mm-0.75 (Fig. 4.36) delivers $K_{\text{act}} = 5\,632.5$ N/mm, $F_{\text{peak}} = 1\,556.25$ N, $R^2 = 0.9662$ and a compact corrected peak displacement of 0.045 mm.

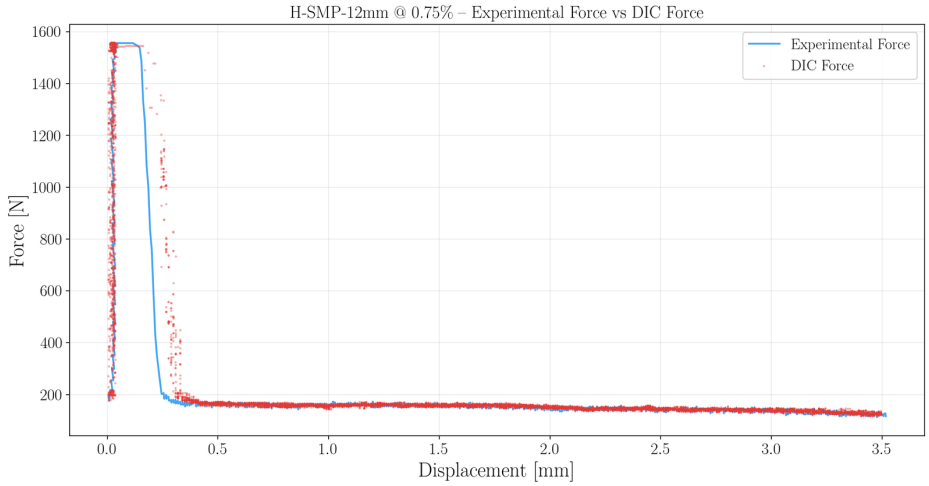


Figure 4.36: H-SMP-12 mm-0.75: ODDC-corrected force–displacement response.

H-SMP-12 mm-1.25 (Fig. 4.37) returns $K_{\text{act}} = 2\,692.9$ N/mm, $F_{\text{peak}} = 1\,548.75$ N and a corrected peak displacement of 0.286 mm with $R^2 = 0.9496$. The lower R^2 of this beam is consistent with a pre-peak load redistribution event captured by the record; the fit remains within the acceptance band.

H-SMP-12 mm-1.50 (Fig. 4.38) is the single stiffest beam of the whole case study:

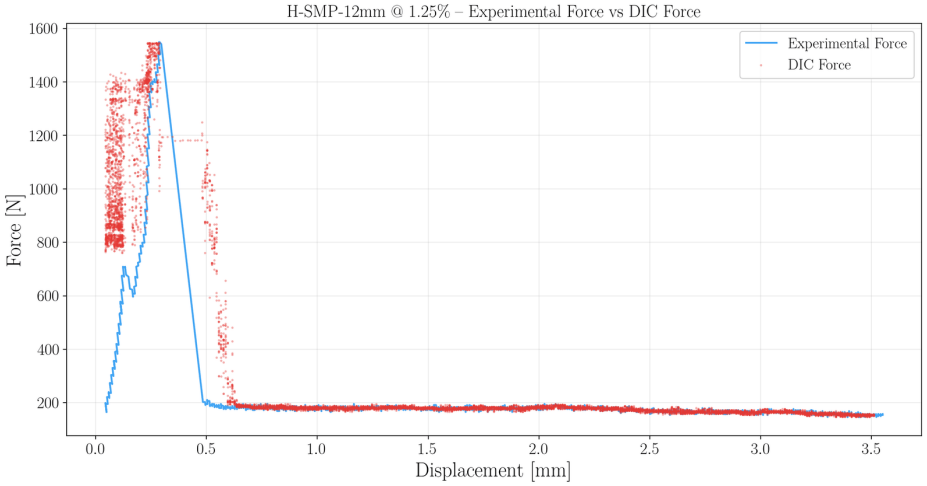


Figure 4.37: H-SMP-12 mm-1.25: ODDC-corrected force–displacement response.

$K_{\text{act}} = 9\,483.2$ N/mm, $F_{\text{peak}} = 1\,713.75$ N – the highest peak force of Case 2 – with a peak-load displacement of 0.188 mm and $R^2 = 0.9765$.

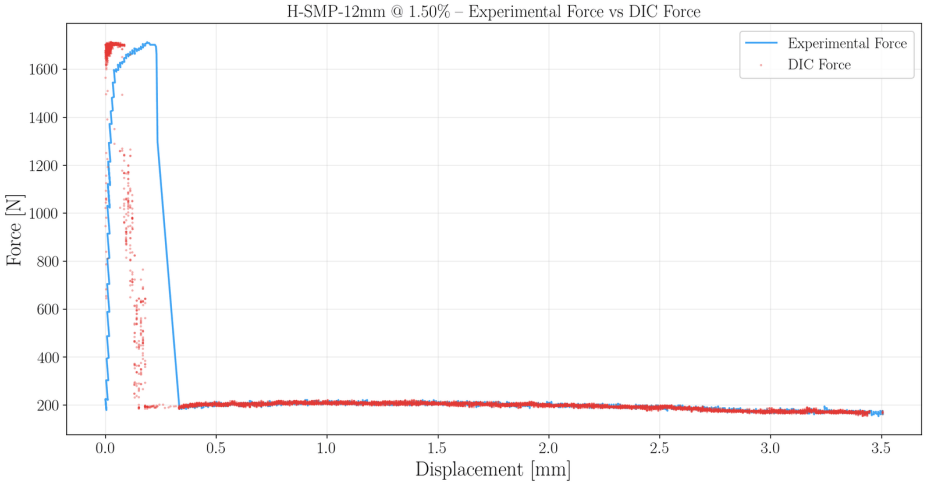


Figure 4.38: H-SMP-12 mm-1.50: ODDC-corrected force–displacement response; stiffest beam of Case 2.

4.4.2.4 High matrix, SMP-24 mm dose series

The 1.00 % reference is H-SMP-24 mm, already in Fig. 4.26 and Table 4.3.

H-SMP-24 mm-0.75 (Fig. 4.39) gives $K_{\text{act}} = 2\,769.4$ N/mm, $F_{\text{peak}} = 1\,473.75$ N, $R^2 = 0.9659$ and a corrected peak displacement of 0.160 mm.

The 1.25 % dose H-SMP-24 mm-1.25 (Fig. 4.40) returns $K_{\text{act}} = 7\,188.8$ N/mm, $F_{\text{peak}} = 1\,612.5$ N, a corrected peak displacement of 0.051 mm and $R^2 = 0.9867$.

The 1.50 % dose H-SMP-24 mm-1.50 (Fig. 4.41) closes the High-matrix SMP-24 sweep with $K_{\text{act}} = 8\,130.4$ N/mm, $F_{\text{peak}} = 1\,650$ N, $R^2 = 0.9768$ and a corrected peak displacement of 0.044 mm.

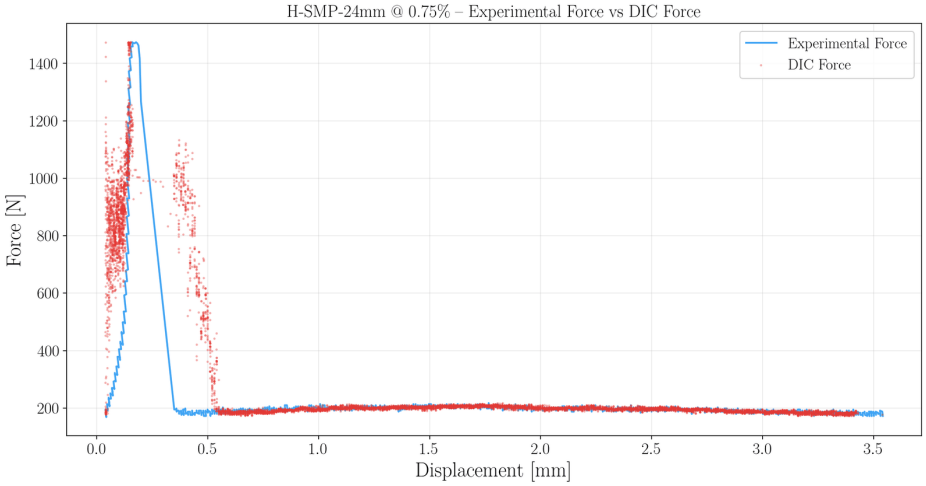


Figure 4.39: H-SMP-24 mm-0.75: ODDC-corrected force–displacement response.

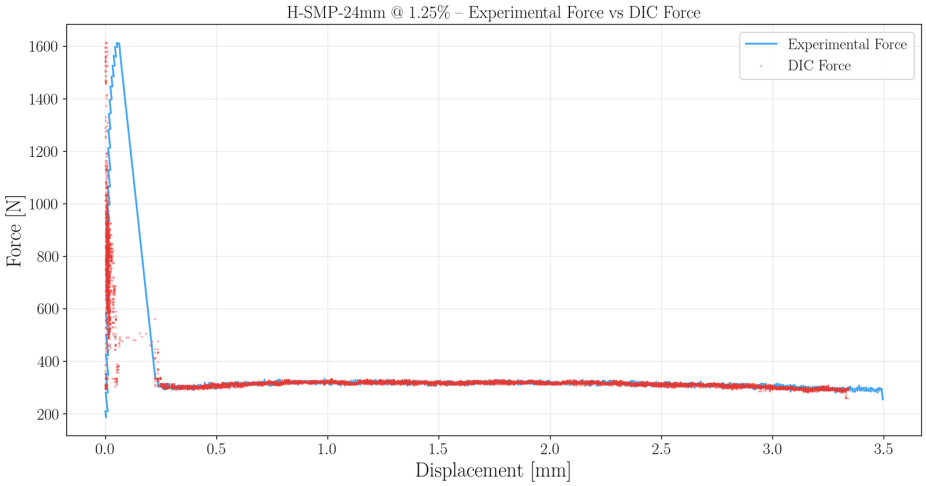


Figure 4.40: H-SMP-24 mm-1.25: ODDC-corrected force–displacement response.

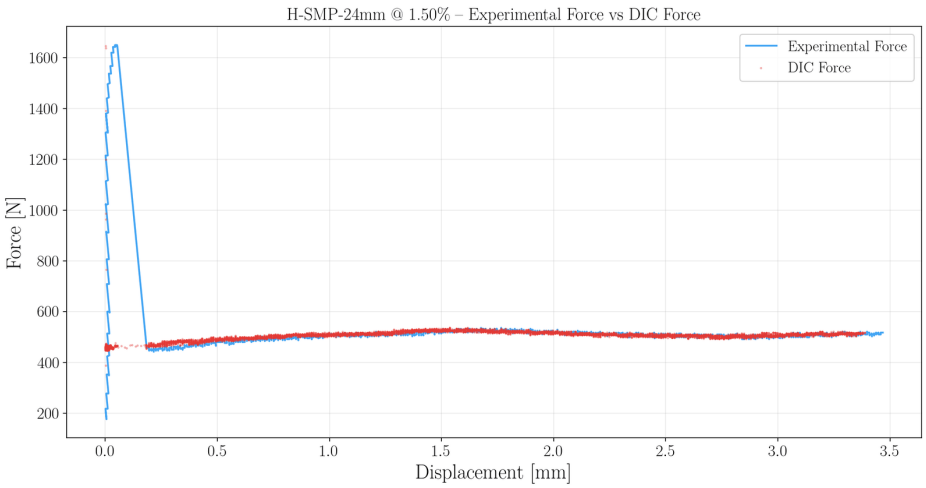


Figure 4.41: H-SMP-24 mm-1.50: ODDC-corrected force–displacement response.

4.4.2.5 Case 2 cohort summary

Table 4.4 consolidates the ODDC outputs of the sixteen new beams of Case Study 2; the two controls are in Table 4.3. The table reports, for every beam, the seating threshold F_0, the peak force F_{peak}, the raw specimen stiffness K_{act}, the machine stiffness used in the correction K_{machine}, the coefficient of determination R^2 and the corrected peak-load displacement $u_{\text{corr, peak}}$.

Table 4.4: Case Study 2 cohort summary (sixteen new beams).

Label	F_0 (N)	F_{peak} (N)	K_{act} (N/mm)	K_{machine} (N/mm)	R^2	$u_{\text{corr,peak}}$ (mm)
N-SMP-12 mm-0.75	191.25	1 016.25	1 732.5	7 140	0.9597	0.2581
N-SMP-12 mm-1.25	187.5	1 331.25	3 561.9	7 140	0.9929	0.1862
N-SMP-12 mm-1.50	191.25	1 203.75	3 652.7	7 140	0.9910	0.1438
N-SMP-24 mm-0.75	198.75	1 233.75	1 389.0	7 140	0.9835	0.5903
N-SMP-24 mm-1.25	198.75	1 226.25	4 267.8	7 140	0.9598	0.0648
N-SMP-24 mm-1.50	198.75	1 110.00	4 970.2	7 140	0.9739	0.0445
H-SMP-12 mm-0.75	187.5	1 556.25	5 632.5	7 140	0.9662	0.0453
H-SMP-12 mm-1.25	195.0	1 548.75	2 692.9	7 140	0.9496	0.2855
H-SMP-12 mm-1.50	191.25	1 713.75	9 483.2	11 268.2	0.9765	0.1881
H-SMP-24 mm-0.75	183.75	1 473.75	2 769.4	7 140	0.9659	0.1599
H-SMP-24 mm-1.25	195.0	1 612.50	7 188.8	7 819.1	0.9867	0.0509
H-SMP-24 mm-1.50	202.5	1 650.00	8 130.4	9 096.3	0.9768	0.0435

The Normal cohort closes uniformly at the global $K_{\text{machine}} = 7\,140$ N/mm, while three High-matrix beams at the upper end of the dose sweep (H-SMP-12 mm-1.50, H-SMP-24 mm-1.25 and H-SMP-24 mm-1.50) close at a slightly higher value of the machine stiffness, reflecting the stiffening effect of the fibre dose on the elastic branch. The full 18-beam Case-2 cohort is therefore available for the dose-response analyses of Chapter 5.

4.4.3 Case Study 3 – cube splitting tensile on FRC cubes

The transition from the notched FRM beams of Case Studies 1 and 2 to the 150 mm FRC cubes of Case Studies 3 and 4 is accompanied by a change of methodology. The cube stiffness in the splitting-tensile configuration is of the order of 200–300 kN/mm, nearly two orders of magnitude above $K_{\text{machine}} = 7\,140$ N/mm of Section 4.2.2. The machine contribution to the displacement record is therefore negligible in relative terms, and the ODDC correction of Eq. (4.5) would translate to a displacement shift smaller than a single facet-tracking increment. For the FRC cubes the scientific question is no longer "how much of the raw displacement is machine-elastic" but "does the optical field agree with the machine record". This is the DIC point-tracking validation of Section 4.3, applied here for the first time. For every cube, the full grid of virtual measurement points tracked by ZEISS Correlate (facet 19 px, step 16 px, cf. Section 3.7.3.1) is interrogated *post hoc* to identify the single point

whose force-versus-displacement trajectory agrees best with the machine record. Figure 4.42 reproduces, for the CST-CONTROL cube, the composite representation used throughout the case study: the machine record (black) is overlaid with the best-point DIC trajectory (red), and the coefficient of determination R^2, the root-mean-square error RMSE and the mean absolute error MAE are reported as a triplet next to the record.

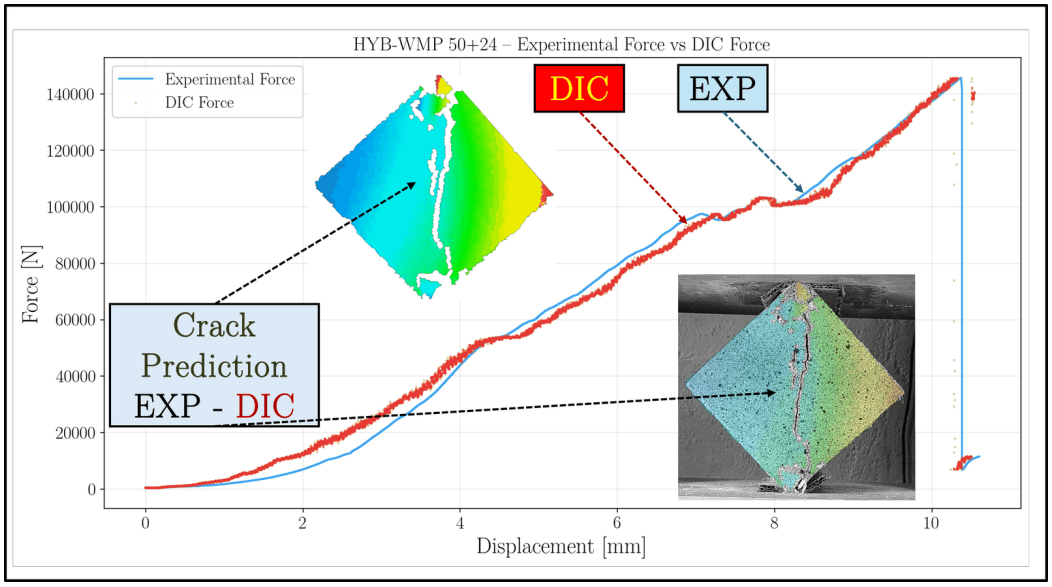


Figure 4.42: Case Study 3 capability plot on the CST-CONTROL cube: best-point DIC trajectory versus machine record.

4.4.3.1 Per-specimen discussion

CST-CONTROL (Fig. 4.43) provides a reference point for the method. The DIC-versus-machine agreement is excellent – $R^2 = 0.8871$, RMSE = 16 980 N, MAE = 7 599 N across $N_{\text{valid}} = 19\,989$ pairs – and the abrupt brittle drop of the unreinforced cube at $\sim$155 000 N is captured without difficulty.

CST-EMP-50 mm (Fig. 4.44) shows a raw $R^2 = 0.8742$ with a 17 845 N RMSE, indicating a systematic compliance offset in the pre-peak elastic window. After the DIC-validation procedure the correlation is lifted to $R^2 = 0.9740$ with a halved RMSE of 7 167 N; the validated record is retained for the Chapter 5 analysis.

CST-WMP-50 mm (Fig. 4.45) is the most challenging raw record of the case study: $R^2 = 0.7283$, RMSE = 21 933 N, caused by a severe initial-seating discrepancy where the optical field reports displacement well before the machine transducer responds to load. The DIC-validation procedure cleans the offset and yields $R^2 = 0.9498$, RMSE = 9 543 N; the validated record is retained.

CST-SMP-50 mm (Fig. 4.46) presents a softer raw agreement ($R^2 = 0.9237$,

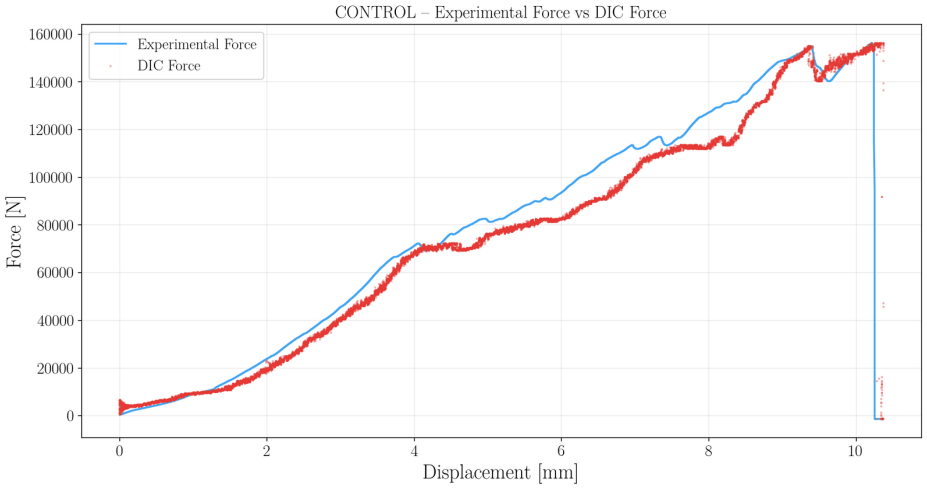


Figure 4.43: CST-CONTROL: DIC-validated force–displacement record.

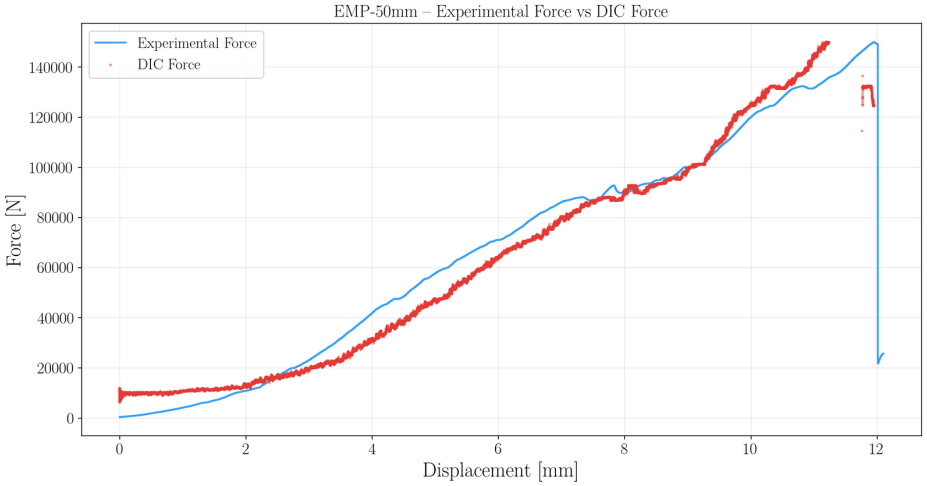


Figure 4.44: CST-EMP-50 mm: DIC-validated force–displacement record.

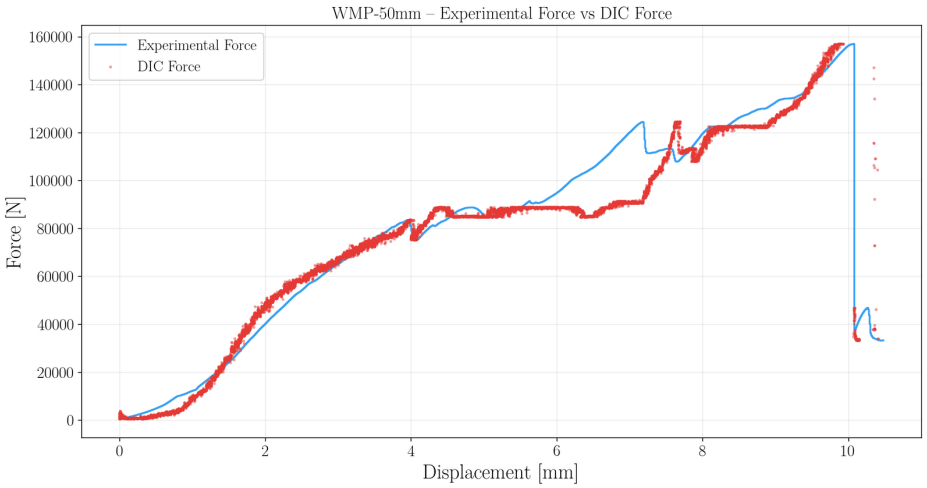


Figure 4.45: CST-WMP-50 mm: DIC-validated force–displacement record.

RMSE = 18 560 N) with a noticeable pre-peak stiffness disparity. The DIC validation lifts R^2 to 0.9626 and reduces the RMSE to 8 428 N; the validated record is retained.

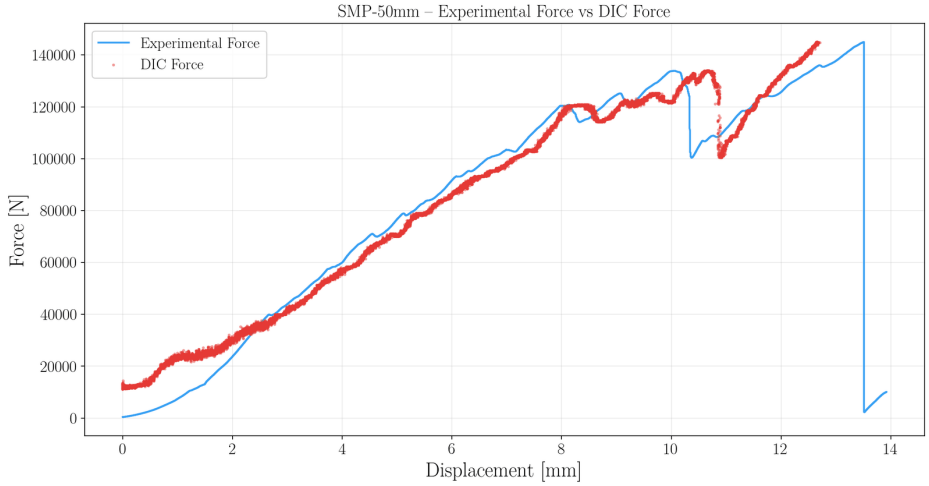


Figure 4.46: CST-SMP-50 mm: DIC-validated force–displacement record.

CST-SMP-24 mm (Fig. 4.47) is the first short-fibre cube and shows a significant pre-peak lateral shift of the optical field. Raw $R^2 = 0.8524$, RMSE = 24 643 N; the DIC validation lifts the correlation to $R^2 = 0.9263$ at the cost of a comparatively higher residual RMSE of 13 387 N, consistent with the jagged post-peak behaviour of the 24 mm fibres.

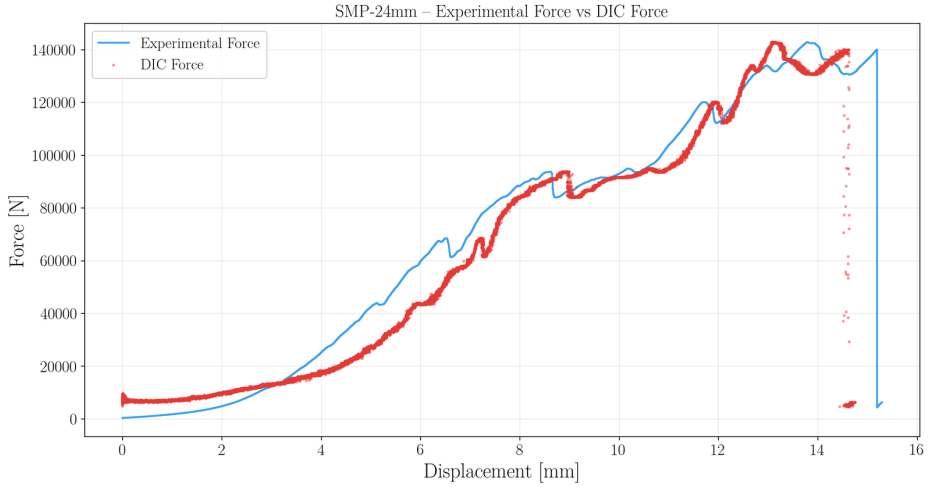


Figure 4.47: CST-SMP-24 mm: DIC-validated force–displacement record.

CST-HYB-50+24 (Fig. 4.48) produces the best-behaved record of the entire case study. The raw R^2 is already 0.9338 with an RMSE of 13 747 N; the DIC validation pushes the correlation to $R^2 = 0.9922$ with an RMSE of only 3 827 N, the best agreement of the CST population. The dual-length hybrid mix arrests micro-cracks

with the 24 mm fibres while the 50 mm fibres bridge the macro-crack, and the DIC surface remains stable throughout the test.

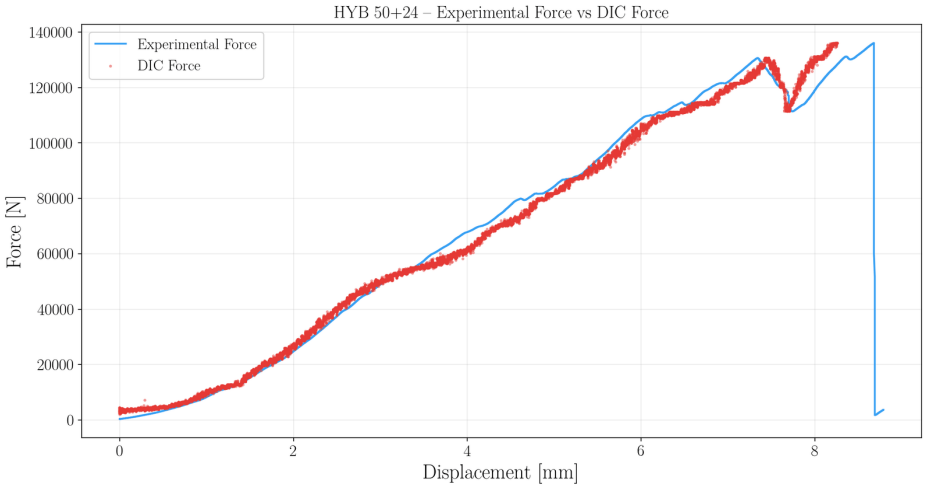


Figure 4.48: CST-HYB-50+24: DIC-validated force–displacement record; best agreement of the case study.

CST-HYB-EMP-50+24 (Fig. 4.49) is the single weakest raw record of the case study: $R^2 = 0.8031$, RMSE = 16 784 N. The tortuous fracture path of the embossed-plus-straight hybrid destabilises the DIC facet grid in the post-peak regime, and the raw record is retained with its pre-peak behaviour as the primary quantity for downstream analysis.

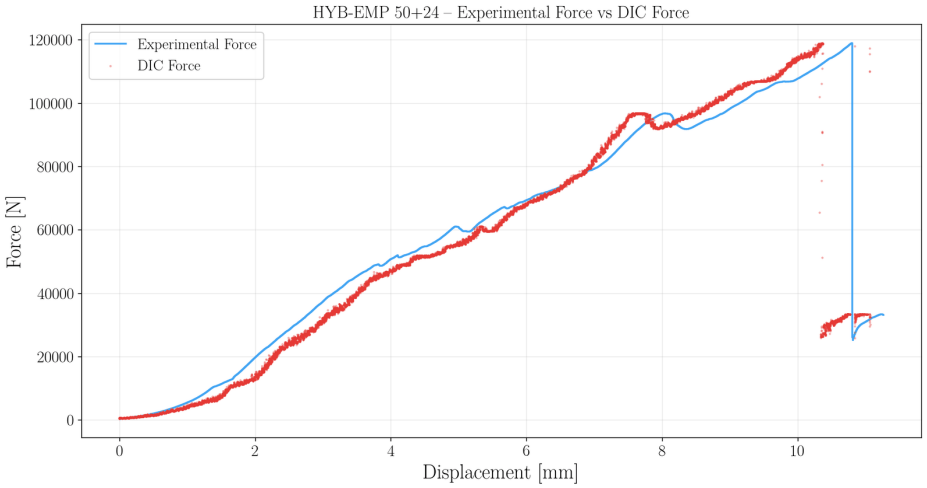


Figure 4.49: CST-HYB-EMP-50+24: raw DIC versus machine record retained for pre-peak use.

CST-HYB-WMP-50+24 (Fig. 4.50) closes the cohort with a clean raw $R^2 = 0.9128$ and RMSE = 13 529 N.

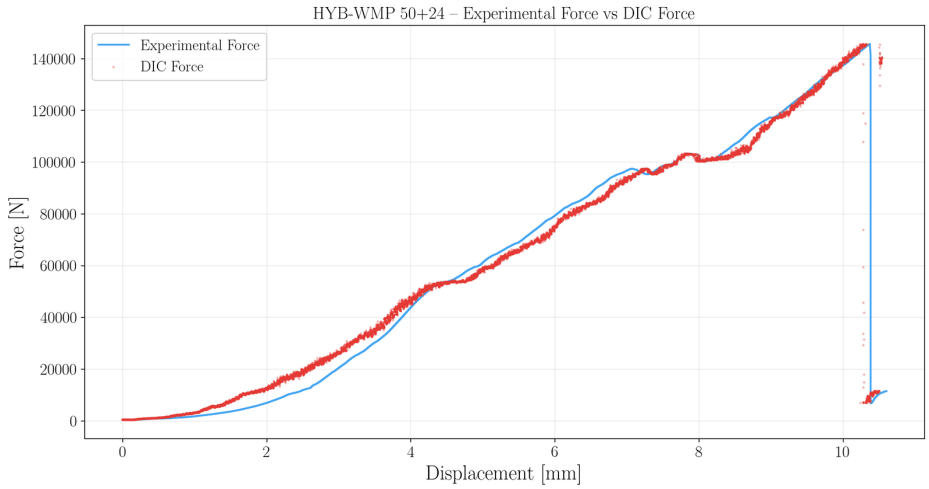


Figure 4.50: CST-HYB-WMP-50+24: DIC-validated force–displacement record.

4.4.3.2 Case 3 cohort summary

Table 4.5 consolidates the DIC validation metrics for the eight CST cubes: the coefficient of determination R^2, the root-mean-square error RMSE and the mean absolute error MAE, reported both for the raw DIC record (when it alone was retained) and for the DIC-validated record (when the validation procedure improved the agreement).

Table 4.5: Case Study 3 (CST) cohort summary.

Label	R^2_{raw}	**RMSE$_{\mathrm{raw}}$** (N)	**MAE$_{\mathrm{raw}}$** (N)	R^2_{valid}	**RMSE$_{\mathrm{valid}}$** (N)
CST-CONTROL	0.8871	16 980	7 599	–	–
CST-EMP-50 mm	0.8742	17 845	11 940	0.9740	7 167
CST-WMP-50 mm	0.7283	21 933	10 361	0.9498	9 543
CST-SMP-50 mm	0.9237	18 560	15 716	0.9626	8 428
CST-SMP-24 mm	0.8524	24 643	16 145	0.9263	13 387
CST-HYB-50+24	0.9338	13 747	8 648	0.9922	3 827
CST-HYB-EMP-50+24	0.8031	16 784	6 532	–	–
CST-HYB-WMP-50+24	0.9128	13 529	4 313	–	–

Seven of the eight cubes close with an R^2 above 0.90, either in raw form or after the DIC validation. Only CST-HYB-EMP-50+24 is retained with a raw $R^2 = 0.8031$; the pre-peak portion of its record remains trustworthy and is carried forward into the Chapter 5 analysis.

4.4.4 Case Study 4 – bonded two-bar splitting on FRC cubes

Case Study 4 carries eight 150 mm FRC cubes cast with two embedded ⌀ 12 mm deformed bars at an embedment length of 150 mm (Section 3.6). The mix design, the fibre composition and the casting protocol are identical to those of Case Study 3,

so the fibre labels (CONTROL, EMP-50, WMP-50, SMP-50, SMP-24, HYB-50+24, HYB-EMP-50+24, HYB-WMP-50+24) are the same; only the geometry and the failure mechanism differ. The DIC point-tracking validation of Section 4.3 is again applied. Figure 4.51 reproduces, for the BBS-CONTROL cube, the composite representation used throughout the case study: the machine record (black) is overlaid with the best-point DIC trajectory (red), together with the triplet (R^2, RMSE, MAE).

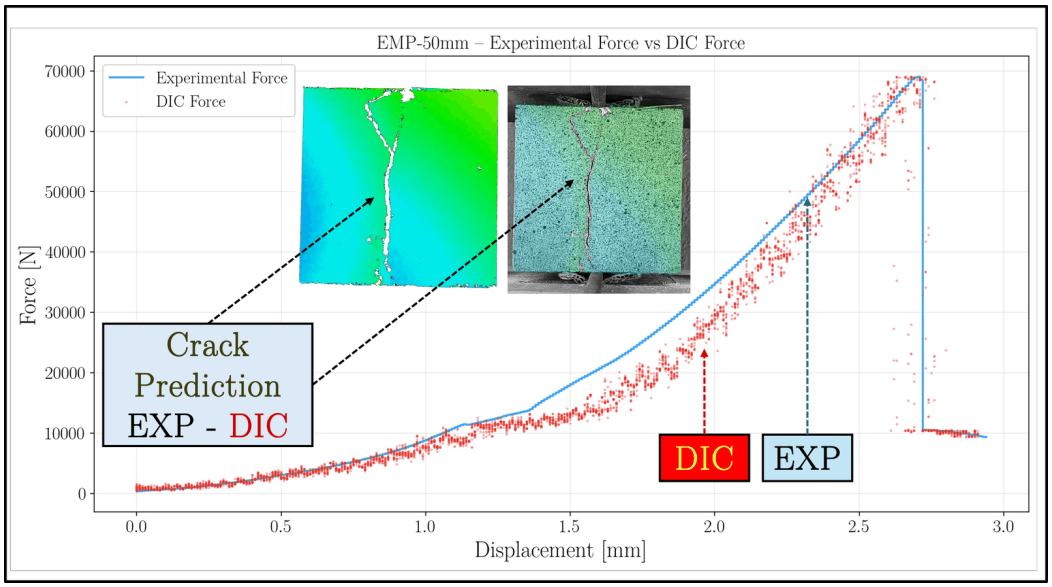


Figure 4.51: Case Study 4 capability plot on the BBS-CONTROL cube: best-point DIC trajectory versus machine record.

4.4.4.1 Per-specimen discussion

BBS-CONTROL (Fig. 4.52) sets the baseline for the BBS cohort. The DIC-versus-machine agreement is already above the acceptance threshold: $R^2 = 0.9079$, RMSE = 8 722 N, MAE = 6 454 N, $N_{\text{valid}} = 6\,269$. The brittle matrix failure at ∼72 000 N is correctly captured by the optical chain.

BBS-EMP-50 mm (Fig. 4.53) closes with $R^2 = 0.9016$, RMSE = 6 714 N and MAE = 3 154 N. The low MAE confirms that physical rigid-body motion of the cube is negligible during the test.

BBS-WMP-50 mm (Fig. 4.54) returns a raw $R^2 = 0.7506$ and RMSE = 12 311 N. The best-point trajectory under-reports displacement relative to the crosshead, most likely due to uneven internal cracking and slight platen rotation prior to peak. The raw record is retained with a caveat reported in Table 4.6.

BBS-SMP-50 mm (Fig. 4.55) delivers the best raw optical-to-mechanical agreement of the BBS cohort: $R^2 = 0.9160$, RMSE = 6 113 N and MAE = 2 544 N. The straight 50 mm monofilament architecture produces uniform surface deformation

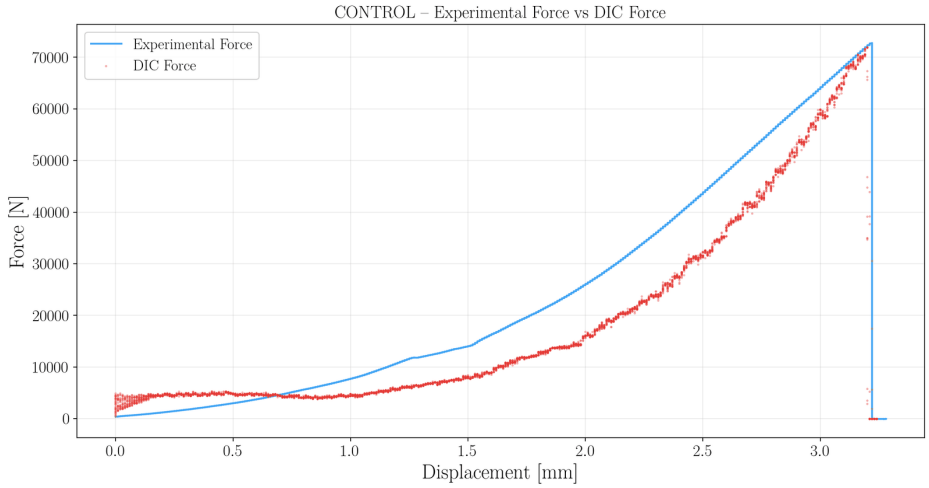


Figure 4.52: BBS-CONTROL: DIC-validated force–displacement record.

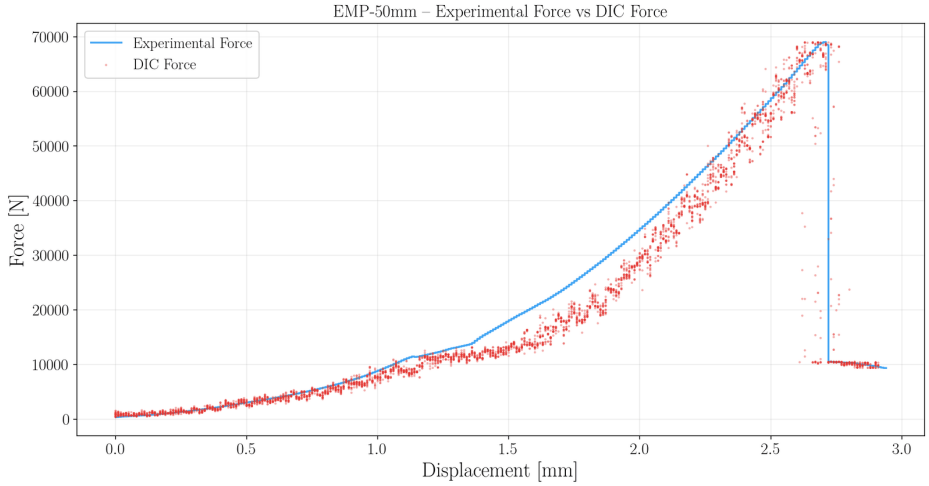


Figure 4.53: BBS-EMP-50 mm: DIC-validated force–displacement record.

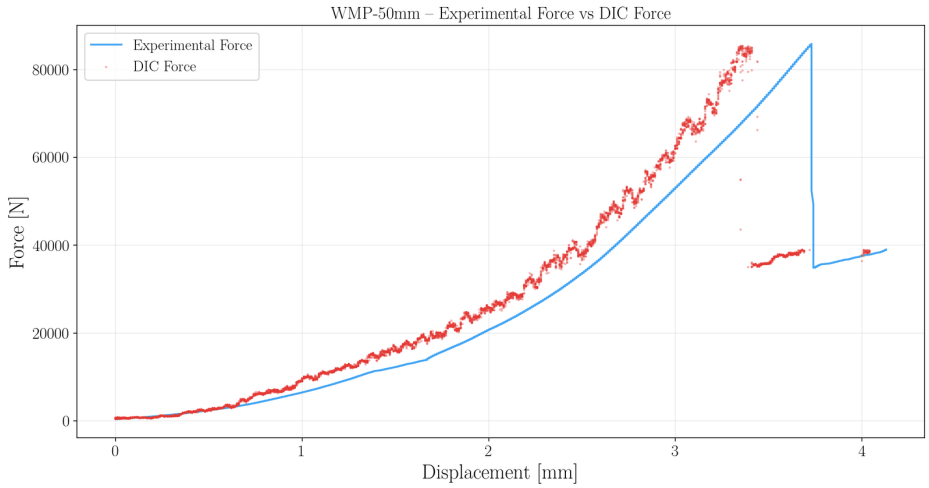


Figure 4.54: BBS-WMP-50 mm: raw DIC record retained with a caveat.

throughout the splitting phase, which the DIC facet grid captures without discontinuity.

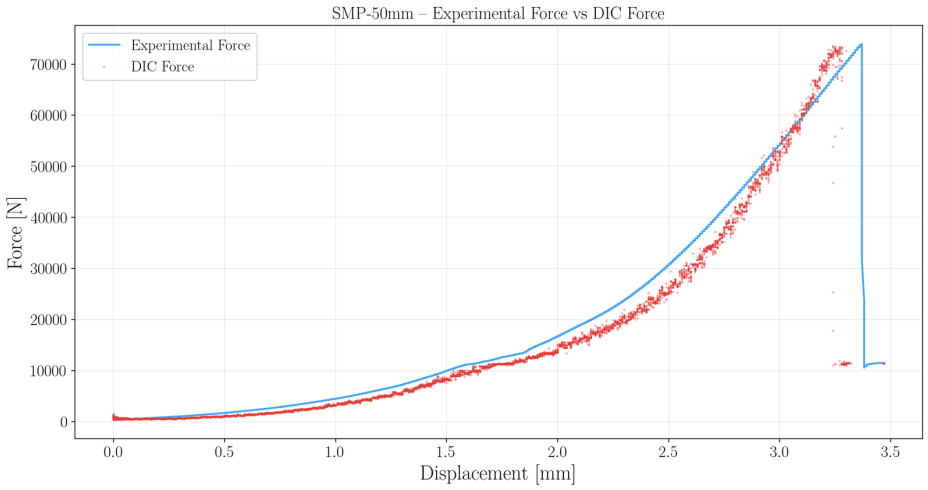


Figure 4.55: BBS-SMP-50 mm: DIC-validated force–displacement record.

BBS-SMP-24 mm (Fig. 4.56) presents a severe raw-data compliance mismatch: $R^2 = 0.8256$ and RMSE = 12 267 N. Visually, the optical field shoots up at a much steeper angle than the experimental curve, missing the peak displacement by close to 1 mm. The DIC validation is decisive: R^2 rises to 0.9775 and the RMSE falls to 3 362 N.

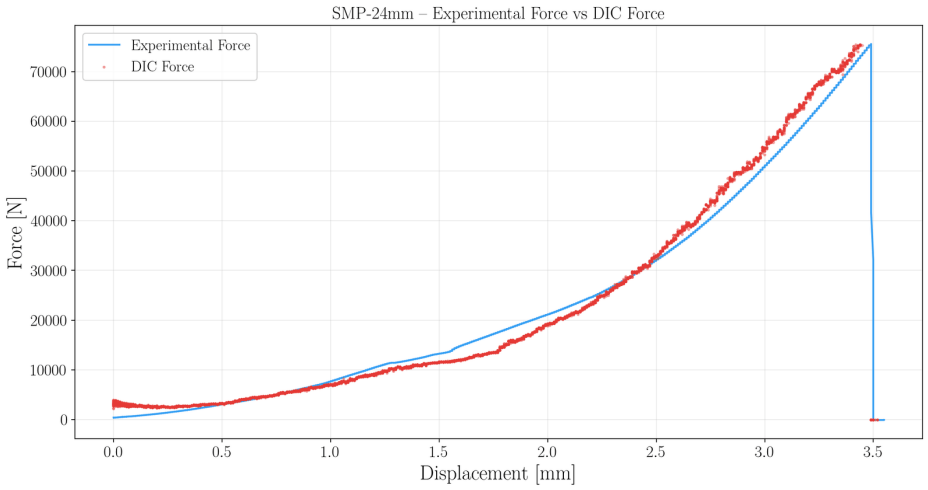


Figure 4.56: BBS-SMP-24 mm: DIC-validated force–displacement record.

BBS-HYB-50+24 (Fig. 4.57) is the most extreme optical-to-mechanical deviation of the cohort: $R^2 = 0.5560$ and RMSE = 16 087 N. The optical field registers significantly less displacement than the machine stroke, suggesting that the hybrid-fibre splitting event produces localised surface stiffening or brief out-of-plane motion that decouples the speckle region from the global deformation. The raw record is retained with the caveat reported in Table 4.6.

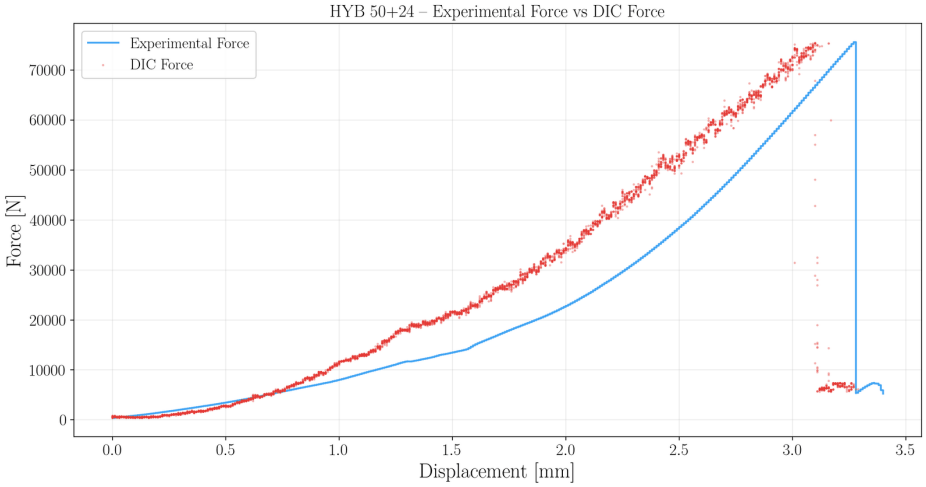


Figure 4.57: BBS-HYB-50+24: raw DIC record retained with a caveat.

BBS-HYB-EMP-50+24 (Fig. 4.58) closes with a validated $R^2 = 0.9796$ and an RMSE of 3 299 N, one of the tightest of the entire BBS cohort.

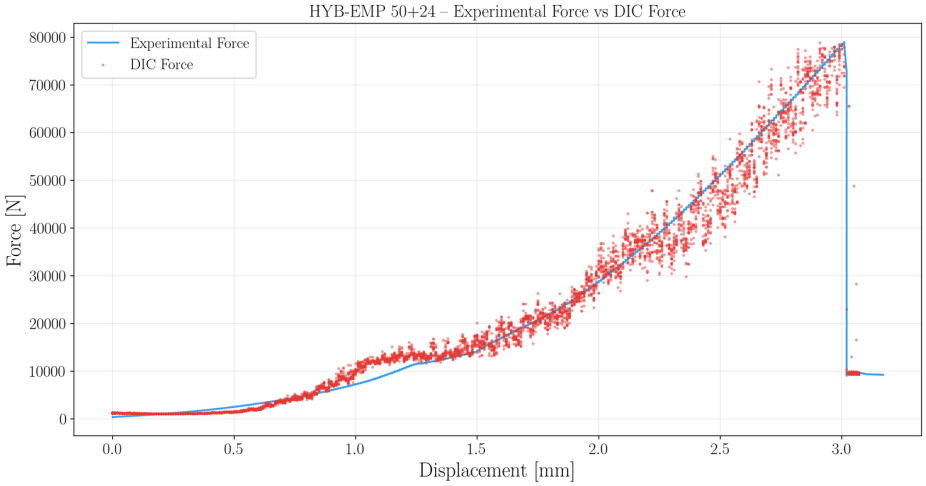


Figure 4.58: BBS-HYB-EMP-50+24: DIC-validated force–displacement record.

BBS-HYB-WMP-50+24 (Fig. 4.59) closes the BBS cohort. After the DIC validation it returns $R^2 = 0.9411$ and an RMSE of 5 757 N.

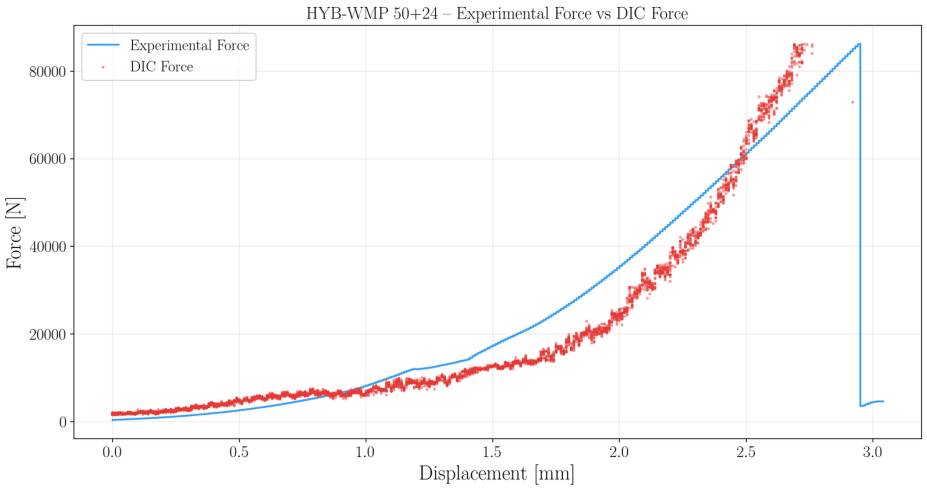


Figure 4.59: BBS-HYB-WMP-50+24: DIC-validated force–displacement record.

4.4.4.2 Case 4 cohort summary

Table 4.6 consolidates the DIC validation metrics for the eight BBS cubes. For every cube the table reports the coefficient of determination R^2, the root-mean-square error RMSE and the mean absolute error MAE, both for the raw DIC record and for the DIC-validated record (when the validation procedure improved the agreement).

Table 4.6: Case Study 4 (BBS) cohort summary.

Label	R^2_{raw}	**RMSE$_{\text{raw}}$** (N)	**MAE$_{\text{raw}}$** (N)	R^2_{valid}	**RMSE$_{\text{valid}}$** (N)
BBS-CONTROL	0.9079	8 722	6 454	–	–
BBS-EMP-50 mm	0.9016	6 714	3 154	–	–
BBS-WMP-50 mm	0.7506	12 311	7 297	–	–
BBS-SMP-50 mm	0.9160	6 113	2 544	–	–
BBS-SMP-24 mm	0.8256	12 267	6 084	0.9775	3 362
BBS-HYB-50+24	0.5560	16 087	9 730	–	–
BBS-HYB-EMP-50+24	0.3491	24 505	19 784	0.9796	3 299
BBS-HYB-WMP-50+24	0.8704	8 998	5 972	0.9411	5 757

Six of the eight BBS cubes close with an R^2 above 0.90, either in raw form or after the DIC validation. Two cubes (BBS-WMP-50 mm and BBS-HYB-50+24) are retained with a sub-threshold raw fit and are treated with caution in the downstream bond-strength analysis of Chapter 5. The Case Study 4 DIC validation demonstrates, even more than the CST cohort, how the point-tracking procedure of Section 4.3 extracts usable kinematic information from splitting records that would otherwise be dominated by optical noise.

4.5 DIC capability beyond compliance

The procedures of Sections 4.2 and 4.3 focus on a single, narrow use of the DIC chain: the extraction of a specimen-mounted kinematic signal that corrects for – or validates against – the machine transducer. The ZEISS Correlate workspace of Figure 4.2 delivers, however, the *complete* displacement field on the observation face. Two further uses of that full field are worth recording here explicitly because they will be exploited in the Chapter 5 discussion of fibre-bridging and fracture-energy dissipation.

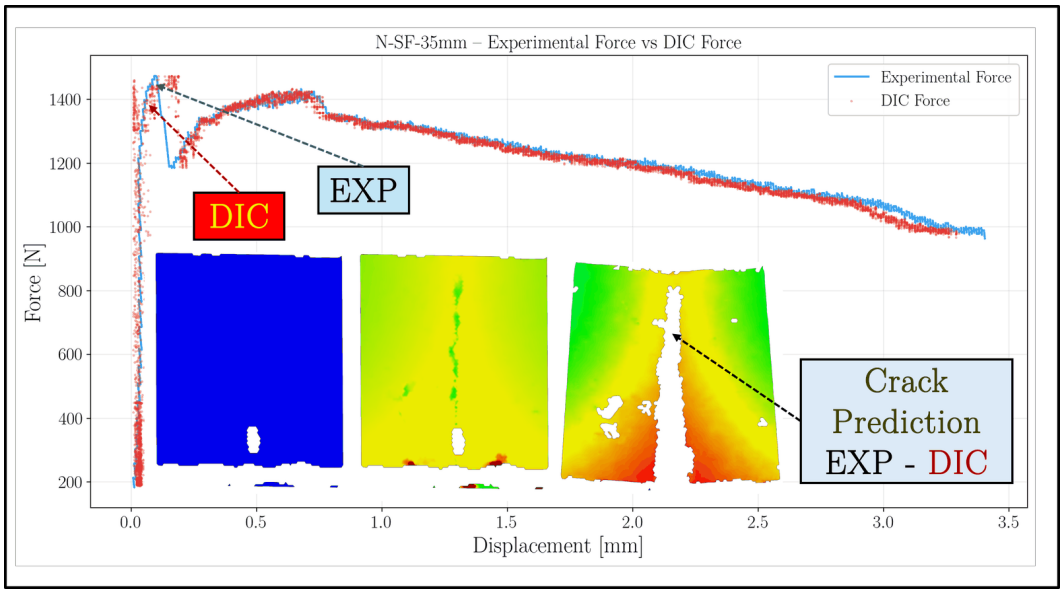


Figure 4.60: Full-field view of the hooked-end-steel beam H-SF-35 mm at three stages of the test.

4.5.1 Full-field view of fibre engagement

Figure 4.60 shows three DIC snapshots of the hooked-end steel beam H-SF-35 mm, whose ODDC-corrected curve was given in Figure 4.29. The three snapshots are extracted from the same ZEISS Correlate workspace of Figure 4.2 at three stages of the test: the end of the proportional range (left), the peak force (centre) and well into the softening branch (right). The DIC surface strain field resolves a progressive localisation first into a single cohesive zone at the notch tip, then into a pair of discrete bridging fibres that control the cohesive traction across the macro-crack, and finally into a single dominant fibre that dictates the residual force plateau. The sequence makes visible, in a single page, what a clip gauge or strain gauge would be unable to resolve by construction: the *discrete* nature of the fibre-bridging response at the length scale of the individual inclusion.

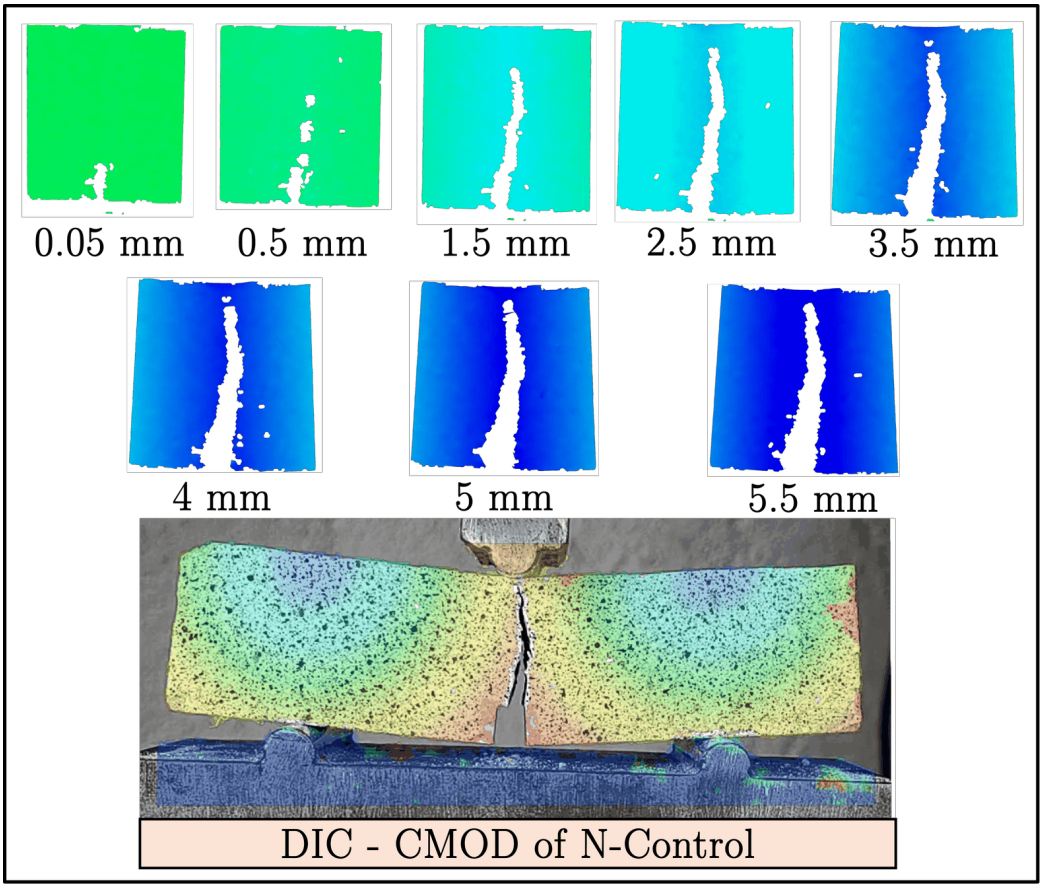


Figure 4.61: Post-hoc ladder of CMOD-based virtual extensometers on the N-CONTROL mortar beam.

4.5.2 Post-hoc virtual-gauge ladder

Figure 4.61 shows a complementary capability of the DIC chain. On the N-CONTROL mortar beam (already used for the machine-stiffness identification of Section 4.2.2 and for the capability demonstration of Figure 4.7), a *ladder* of CMOD-based virtual extensometers is defined post-hoc on the same DIC video. The gauge length is progressively halved from the full 27 mm baseline of Section 3.7.3.2 (top panel) down to a ∼3 mm micro-gauge placed directly across the notch-tip cohesive zone (bottom panel), and the DIC field sustains the resolution across the full ladder. The ladder is feasible because every virtual extensometer is merely a post-processing choice on the saved displacement field, and the measurement does not need to be decided before the test. The ability to define any number of kinematic diagnostics on a single video capture – in contrast to the single-gauge commitment of a clip-gauge extensometer or strain gauge – is arguably the strongest practical advantage of the DIC chain for the subsequent fracture-energy analysis of Chapter 5.

Taken together, Figures 4.60 and 4.61 place the compliance correction of Sections 4.2–4.3 inside the wider context of what the DIC chain delivers. The compli-

ance correction is *necessary* for the narrow question of pre-peak stiffness, but the optical chain is *sufficient* for a much broader set of kinematic enquiries, which is exactly the reason why the same hardware and software pipeline is adopted uniformly across the four case studies of this thesis.

4.6 Summary

This chapter has introduced the Optical Deflection-Derived Compliance (ODDC) methodology and its point-tracking companion, and applied both procedures to all 56 specimens of the four case studies of the experimental programme. The key findings may be summarised in four points:

- **Origin of the compliance artefact**: the raw crosshead record of the Galdabini Sun/60 is contaminated by the elastic travel of the complete load path (piston, load cell, rollers, platen). On the 40 mm FRM beams this artefact is of the same order as the specimen response itself; on the 150 mm FRC cubes the cube itself dominates the load path and the artefact is negligible in relative terms.
- **ODDC methodology for the FRM beams**: a single reference test (N-CONTROL of Case Study 1) is used to identify the machine stiffness from the optical deflection field through the series-spring decomposition of Eq. (4.1) and the inverse of Eq. (4.3); the resulting $K_{\text{machine}} = 7\,140$ N/mm is then re-used as a global subtrahend (Eq. 4.5) for every subsequent beam test.
- **DIC point-tracking validation for the FRC cubes**: on the 150 mm FRC cubes the specimen-dominated load path shifts the scientific question from a compliance subtraction to an optical-versus-mechanical validation; the best point of the facet grid is selected by maximum R^2 and delivers accepted force–displacement records across the two cube campaigns.
- **Wider DIC capability**: beyond compliance correction, the same optical chain delivers the full surface displacement field of every specimen, which supports the direct observation of discrete fibre bridging (Figure 4.60) and a post-hoc ladder of virtual extensometers from the specimen-scale 27 mm CMOD gauge down to a $\sim$3 mm notch-tip micro-gauge (Figure 4.61).

The corrected force–displacement records obtained by the ODDC pipeline and the validated DIC point-tracked records obtained by the CST/BBS procedure con-

stitute the kinematic input for the fracture-energy, post-peak-residual and bond-stress analyses of Chapter 5. In that chapter, the compliance correction developed here is *presupposed*; the focus shifts from the *production* of the corrected records to the *use* of those records for the engineering characterisation of the fibre-reinforced composites.

Chapter 5

Structural Performance, Statistical Analysis and Life-Cycle Assessment

chapter 4 delivered the corrected load–deflection curves for every specimen, with the testing-machine contribution removed and a point-tracking validation step from Digital Image Correlation. The present chapter uses those curves as the input to a single, uniform evaluation pipeline. Three complementary tracks are developed in parallel. The first (section 5.1) converts each corrected curve into a compact vector of structural quantities: peak flexural stress and total toughness for the prismatic beams, splitting tensile strength and pre-peak energy for the cubes, and bond stress with pre-peak energy and bilinear stiffness for the bonded-bar cubes. The second (section 5.2) screens those quantities against an explicit non-parametric pipeline (Shapiro–Wilk, Mann–Whitney U, Kruskal–Wallis, Spearman rank, Pearson heatmap, one-dimensional quadratic dose fit and cubic-spline interpolated contours), and makes explicit the tests that the $n = 3$/cell experimental design does *not* support. The third (section 5.3) assigns every mixture a cradle-to-gate Global Warming Potential, combines it with the structural quantities through a family of Eco-Efficiency Indices, and applies a Monte Carlo sensitivity analysis and a cost–benefit map. section 5.4 applies the three tracks to the four case studies in sequence – the two three-point bending cohorts, the cube splitting tensile cohort, and the bonded two-bar splitting cohort – and closes with the scientific novelty that the last two cohorts share the same mix design but rank the fibre programme in opposite orders. section 5.5 collects the findings in five bullets.

5.1 Structural Performance Evaluation

5.1.1 Flexural quantities for the beams (Case studies 1 and 2)

All prismatic beams measured $40 \times 40 \times 160$ mm with a test span $L = 100$ mm and were tested in three-point bending following EN 14651 [40]. From each corrected force–deflection curve $F(\delta)$ the peak flexural stress is

$$\sigma_{\text{flex}} \;=\; \frac{3\,F_{\text{peak}}\,L}{2\,b\,d^2} \tag{5.1}$$

with F_{peak} the measured peak force, $b = 40$ mm the width and $d = 40$ mm the depth. The specimen secant stiffness is the slope of a linear regression on the pre-peak window between $0.10\,F_{\text{peak}}$ and $0.40\,F_{\text{peak}}$,

$$K_{\text{spec}} \;=\; \left.\frac{\mathrm{d}F}{\mathrm{d}\delta}\right|_{0.10F_{\text{peak}}\le F\le 0.40F_{\text{peak}}} \tag{5.2}$$

and the apparent elastic modulus follows from Euler–Bernoulli beam theory

$$E_{\text{app}} \;=\; \frac{K_{\text{spec}}\,L^3}{48\,I}, \qquad I \;=\; \frac{b\,d^3}{12} \tag{5.3}$$

reported in GPa. Total toughness is the work-of-fracture integral

$$T_{\text{total}} \;=\; \int_0^{\delta_{\text{end}}} F(\delta)\,\mathrm{d}\delta \tag{5.4}$$

split into a pre-peak contribution $T_{\text{pre}} = \int_0^{\delta_{\text{peak}}} F\,\mathrm{d}\delta$ and a post-peak contribution $T_{\text{post}} = T_{\text{total}} - T_{\text{pre}}$. The upper integration limit δ_{end} is defined as the last recorded displacement sample at which the force remains above 1 % of the peak force; this specimen-dependent criterion ensures that the toughness integral captures the full post-peak tail while excluding electronic noise in the near-zero force region. All three integrals are evaluated with the composite trapezoidal rule on the re-sampled (F, δ) arrays.

Data provenance. The structural quantities in this chapter are computed from the corrected curves delivered by chapter 4. For the prismatic beams of Case Studies 1 and 2, the curves are the ODDC-corrected force–deflection records (Section 4.2) from which σ_{flex}, T_{total}, f_R-equivalent residuals and fracture energy G_f are extracted

over the full pre- and post-peak range. For the CST cubes of Case Study 3, no compliance correction is applied (the cube stiffness is two orders of magnitude above K_{machine}); the curves are therefore the DIC consistency-checked records (Section 4.3) from which F_{peak} and f_t are computed. For the BBS cubes of Case Study 4, the same DIC consistency-checked records are used, and specimens whose DIC–machine agreement falls below $R^2 = 0.90$ without recovery after post-processing are carried forward with a caveat flag; their pre-peak quantities (F_{peak}, τ, E_{pre}) are retained, but their post-peak tails are excluded from the toughness integrals.

The first-crack event is obtained by smoothing the pre-peak force with a Savitzky–Golay filter of polynomial order 3 and window $w = \min(11, \lfloor i_{\text{peak}}/3 \rfloor)$, forced odd, and then locating the index of maximum curvature $|\mathrm{d}^2F/\mathrm{d}\delta^2|$ inside the sub-window $[\max(5, \lfloor 0.05\, i_{\text{peak}} \rfloor),\ \lfloor 0.90\, i_{\text{peak}} \rfloor]$. The corresponding first-crack force F_{cr} and displacement δ_{cr} define the elastic limit for the behaviour classifier.

The residual carrying capacity is read directly from the post-peak branch at the first displacement with $\delta \geq 1.5\,\delta_{\text{peak}}$, denoted F_{res}. The full post-peak response is then summarised by the mean of the last five recorded forces, $\bar{F}_{\text{end}}$, which is compared against F_{peak} through two thresholds:

- *hardening* when $\bar{F}_{\text{end}} > 0.80\, F_{\text{peak}}$,
- *moderate softening* when $0.30\, F_{\text{peak}} < \bar{F}_{\text{end}} \leq 0.80\, F_{\text{peak}}$,
- *softening* when $\bar{F}_{\text{end}} \leq 0.30\, F_{\text{peak}}$.

The two improvement percentages used throughout the chapter are taken with respect to the matrix-matched control beam,

$$\Delta F\% \;=\; \frac{F_{\text{peak}} - F_{\text{peak}}^{(\text{m})}}{F_{\text{peak}}^{(\text{m})}} \cdot 100, \qquad \Delta T\% \;=\; \frac{T_{\text{total}} - T_{\text{total}}^{(\text{m})}}{T_{\text{total}}^{(\text{m})}} \cdot 100 \tag{5.5}$$

where the superscript (m) refers to the N- or H-matrix control as appropriate. A fibre-volume efficiency $\eta_\sigma = \sigma_{\text{flex}}/W_f$ and a toughness efficiency $\eta_T = T_{\text{total}}/W_f$ (with W_f the fibre mass in kg/m^3) complete the beam-level metric set, together with a cost surrogate $\eta_{\text{cost}} = T_{\text{total}}/\text{SP}$ that normalises toughness by superplasticiser demand [32].

5.1.2 Splitting quantities for the cubes (Case studies 3 and 4)

The cubes are $150 \times 150 \times 150$ mm and share an M30 mix with CEM I 42.5R, water-to-cement ratio ≈ 0.52, 10 and 20 mm coarse aggregates, and a superplasticiser

dose tuned to maintain workability. Two loading configurations are distinguished. For the orthogonal cube splitting tensile (CST) test [69], the line load is transferred through 25 mm bearing strips along the top and bottom faces, and the indirect tensile strength is

$$f_t \;=\; \frac{2\,P}{\pi\,d\,L} \tag{5.6}$$

with P the peak load and $d = L = 150$ mm the cube side. The theoretical tensile stress along the loaded diameter decays from the centre outwards as

$$\sigma_t(x) \;=\; \frac{2P}{\pi\,d\,L}\left[\frac{d^2}{d^2 - 4x^2} - 1\right] \tag{5.7}$$

where x is measured from the mid-plane. Pre-peak energy is the integral of the corrected load–displacement curve up to the peak,

$$E_{\mathrm{pre}} \;=\; \int_0^{\delta_{\mathrm{peak}}} F(\delta)\,\mathrm{d}\delta \tag{5.8}$$

and the initial secant stiffness K_0 is obtained on the same 0.10–$0.40\,F_{\mathrm{peak}}$ window as for the beams. A non-linearity index combines the two:

$$NL \;=\; 1 - \frac{E_{\mathrm{elastic}}}{E_{\mathrm{pre}}}, \qquad E_{\mathrm{elastic}} \;=\; \frac{F_{\mathrm{peak}}^2}{2\,K_0} \tag{5.9}$$

so that $NL = 0$ describes a purely linear ascending branch and $NL \to 1$ a strongly yielded one. The ascending branch itself is summarised by a two-segment bilinear fit:

$$F(\delta) \;=\; \begin{cases} K_1\,\delta & \delta \le \delta_{\mathrm{LOP}} \\ F_{\mathrm{LOP}} + K_2\,(\delta - \delta_{\mathrm{LOP}}) & \delta > \delta_{\mathrm{LOP}} \end{cases} \tag{5.10}$$

with slopes K_1 and K_2 and a breakpoint δ_{LOP} (limit of proportionality) located by minimising the bilinear sum of squared errors on the bounded interval $0.15 \le \delta/\delta_{\mathrm{peak}} \le 0.85$.

For the bonded two-bar splitting test (BBS), two ribbed steel bars of diameter $d_b = 12$ mm and embedded length $L_e = 150$ mm pass through the cube along the vertical axis and form the load path. The mean interfacial bond stress is

$$\tau \;=\; \frac{F}{n\,\pi\,d_b\,L_e} \tag{5.11}$$

with $n = 2$ the number of bars. The pre-peak energy is defined exactly as in (5.8). In BBS, however, the bar–matrix interface dissipates work before the peak, and a

damage variable is used to make that explicit:

$$E_{\text{diss}} \;=\; \max\!\big(E_{\text{pre}} - E_{\text{elastic}},\, 0\big), \qquad D \;=\; \frac{E_{\text{diss}}}{E_{\text{pre}}} \tag{5.12}$$

which is equivalent to $D = 1 - E_{\text{elastic}}/E_{\text{pre}}$ when $E_{\text{elastic}} < E_{\text{pre}}$. The bilinear K_1/K_2 fit is performed on the BBS ascending branch with the same 0.15–0.85 bound, which produces a systematically different signature: in CST the second segment is shallower than the first ($K_2/K_1 < 1$) and reports onset of micro-cracking, whereas in BBS the second segment is steeper ($K_2/K_1 > 1$) and reports progressive mobilisation of bond along the embedment. The cross-method diagnostic that closes section 5.4.4 is simply the ratio f_t/τ of the two cube tests on mixes with matching fibre programme, computed specimen-by-specimen.

Three fibre-volume efficiencies are used for the cubes: $\eta_{f_t} = \Delta f_t/V_f$, $\eta_E = \Delta E_{\text{pre}}/V_f$, and in BBS the bond analogue $\eta_\tau = \Delta\tau/V_f$. All of them measure the marginal structural gain per unit fibre volume and feed directly into the marginal Eco-Efficiency Indices of section 5.3.

Table 5.1: Summary of the structural quantities used throughout chapter 5. All formulas are evaluated on the ODDC-corrected curves produced in chapter 4.

Symbol	Formula	Unit	Source / context
σ_{flex}	$3F_{\text{peak}}L/(2bd^2)$	MPa	3PBT beams, EN 14651 [40]
K_{spec}	$\mathrm{d}F/\mathrm{d}\delta$ on $[0.10, 0.40]F_{\text{peak}}$	N/mm	beams and cubes, pre-peak window
E_{app}	$K_{\text{spec}}L^3/(48I)$, $I = bd^3/12$	GPa	3PBT beams
T_{total}	$\int_0^{\delta_{\text{end}}} F\,\mathrm{d}\delta$	N·mm	3PBT beams, RILEM 162-TDF [32]
F_{cr}, δ_{cr}	Savitzky–Golay $\lvert F''\rvert$ maximum	N, mm	first-crack detector, pre-peak
F_{res}	F at first $\delta \geq 1.5\,\delta_{\text{peak}}$	N	beams, ductility descriptor
behaviour	$\bar{F}_{\text{end}}/F_{\text{peak}}$ vs $\{0.30, 0.80\}$	–	hardening / moderate / softening
$\Delta F\%$, $\Delta T\%$	$(x - x^{(m)})/x^{(m)} \cdot 100$	%	improvement vs matched control
η_σ, η_T	σ_{flex}/W_f, T_{total}/W_f	per kg/m^3	beams, fibre efficiency
f_t	$2P/(\pi dL)$	MPa	CST cubes, ASTM C496 [69]
$\sigma_t(x)$	$(2P/\pi dL)[d^2/(d^2 - 4x^2) - 1]$	MPa	theoretical along splitting diameter
E_{pre}	$\int_0^{\delta_{\text{peak}}} F\,\mathrm{d}\delta$	N·mm	cubes, pre-peak energy
NL	$1 - E_{\text{elastic}}/E_{\text{pre}}$	–	cubes, nonlinearity index
K_1, K_2, LOP	bilinear fit on $[0.15, 0.85]$ of ascending	N/mm	cubes, ascending-branch signature
τ	$F/(n\pi d_b L_e)$, $n = 2$, $d_b = 12$ mm, $L_e = 150$ mm	MPa	BBS cubes, mean bond stress
D	$E_{\text{diss}}/E_{\text{pre}}$	–	BBS damage
f_t/τ	cross-method, matched fibre programme	–	novelty diagnostic

5.2 Statistical Analysis Framework

5.2.1 Motivation and sample-size constraint

The experimental campaign of this thesis collected between one and three beams or cubes per factorial cell (one specimen per reinforced cell in Case 2, two cubes per fibre programme in Cases 3 and 4, and three beams per fibre programme in Case 1). A sample size of $n \leq 3$/cell is below the threshold at which parametric inference based on Gaussian residuals can be safely trusted: normality tests have low power, variance estimates are unstable, and even the simplest two-way analysis of variance loses one degree of freedom per interaction term before a meaningful error variance can be computed [56, 70, 71]. The response of this chapter is therefore deliberate and conservative: no parametric omnibus tests, no fitted response-surface polynomials, and no confidence intervals derived under a Gaussian assumption. The pipeline described below relies instead on rank-based statistics that remain valid under small samples and skewed distributions. With $n = 3$ per cell, these tests are used as exploratory screening tools rather than definitive inferential proof. P-values should be interpreted cautiously and supported by effect sizes, ranking consistency and physical interpretation.

Two classical techniques are mentioned here only to make clear that they are *not* claimed in what follows. A full two-way analysis of variance (factor A = matrix at two levels, factor B = fibre family at six levels) would require replicated cells to separate main effects from the interaction, which the campaign does not provide. A full six-term quadratic response-surface model $Y(V_f, L_f) = \beta_0 + \beta_1 V_f + \beta_2 L_f + \beta_{11} V_f^2 + \beta_{22} L_f^2 + \beta_{12} V_f L_f$ would require at least six non-collinear (V_f, L_f) observations per matrix, which the 2×4 dose–length grid of Case 2 can only produce through interpolation. The chapter therefore reports one-dimensional quadratic dose fits per length and a cubic-spline interpolation over the dose–length grid, and reserves the words *analysis of variance* and *response surface* strictly for these descriptive forms.

5.2.2 Normality screening: Shapiro–Wilk

Each response variable (peak force, flexural stress, toughness, pre-peak energy, displacement at peak, bond stress where applicable) is screened for normality with the Shapiro–Wilk test [55]. The test statistic

$$W \;=\; \frac{\left(\sum_{i=1}^{n} a_i \, x_{(i)}\right)^2}{\sum_{i=1}^{n} (x_i - \bar{x})^2} \tag{5.13}$$

where $x_{(i)}$ is the ith order statistic and a_i are tabulated constants, is computed on each fibre-reinforced subset for each case. A p-value above $\alpha = 0.05$ is interpreted as an absence of evidence against Gaussianity, not as a positive proof of it. Specifically, a non-significant p-value from the Shapiro–Wilk test does not constitute evidence of normality; it only indicates that the available sample size ($n \leq 3$) provides insufficient statistical power to reject the null hypothesis of normality. This caveat is particularly acute at the present sample sizes, where the test can fail to detect moderate departures from Gaussianity. The result is sufficient to justify a subsequent rank-based pipeline when $p < \alpha$ for even one response variable.

5.2.3 Location-shift testing: Mann–Whitney U

Pairwise differences between two independent groups are tested with the two-sided Mann–Whitney U [72]. Given two samples of sizes n_1 and n_2 with combined ranks $R_1 = \sum \mathrm{rank}(x_i^{(1)})$,

$$U_1 \;=\; n_1 n_2 + \frac{n_1(n_1+1)}{2} - R_1, \qquad U \;=\; \min(U_1, n_1 n_2 - U_1) \tag{5.14}$$

is referred to the exact or normal-approximation null distribution. Four pairings are pre-registered in the pipeline and reported consistently across Cases 1–4 when the groups exist: (i) Normal matrix vs High matrix; (ii) 12 mm vs 24 mm micro-fibre length; (iii) single fibre vs hybrid blend; (iv) low dose vs high dose (expressed as 1500 g/m^3 vs 3000 g/m^3 for the micro-fibre cohort). The same $\alpha = 0.05$ threshold is used throughout. Given the number of statistical tests performed across response variables, fibre groups and case studies, the risk of inflated family-wise error rate is acknowledged. The p-values reported here are used descriptively as exploratory screening indicators rather than strict confirmatory evidence; no formal multiple-comparisons correction (e.g. Bonferroni) is applied.

5.2.4 Multi-group ranks: Kruskal–Wallis

When more than two independent groups are compared, the non-parametric analogue of one-way ANOVA is used [73]. For k groups with combined size N and group-wise rank sums R_j,

$$H \;=\; \frac{12}{N(N+1)} \sum_{j=1}^{k} \frac{R_j^2}{n_j} - 3(N+1) \tag{5.15}$$

approaches χ^2_{k-1} under the null hypothesis of equal medians. Kruskal–Wallis is applied to Case 2 (four doses per matrix/length cell) on both peak strength and toughness, and again at $\alpha = 0.05$.

5.2.5 Rank correlation: Spearman and Pearson heatmap

Spearman's rank correlation coefficient [74]

$$\rho_s \;=\; 1 - \frac{6\sum_{i=1}^{n} d_i^2}{n(n^2-1)}, \qquad d_i = \operatorname{rank}(x_i) - \operatorname{rank}(y_i) \tag{5.16}$$

is computed between fibre length and peak force, between fibre length and toughness, and (in Case 2) between fibre dose and each response. Alongside the rank-based test, a Pearson correlation matrix is built over the reinforced subset with variables $\{F_{\text{peak}}, \sigma_{\text{flex}}, T_{\text{total}}, \delta_{\text{peak}}, L_f, \text{SP}\}$ for the beams and $\{F_{\text{peak}}, f_t \text{ or } \tau, E_{\text{pre}}, \delta_{\text{peak}}, V_f, \text{SP}\}$ for the cubes, and rendered as a colour-coded heatmap that serves as a visual supplement to the two correlation tests.

5.2.6 Dose–response: one-dimensional quadratic fit

Within Case 2, where the dose is varied on four levels at each length, a one-dimensional quadratic is fitted separately per length $L_f \in \{12, 24\}$ mm and per matrix:

$$Y(V_f) \;=\; a\,V_f^2 + b\,V_f + c \tag{5.17}$$

with a, b, c obtained by ordinary least squares. The interior extremum, which the literature on fibre dosing identifies with the optimum dose,

$$V_{f,\text{opt}} \;=\; -\frac{b}{2a} \tag{5.18}$$

is reported together with the sign of a (so that the reader can distinguish a maximum, $a < 0$, from a minimum, $a > 0$) and the fit coefficient of determination R^2.

5.2.7 Dose–length interpolation: cubic-spline contours

The Case 2 cohort spans only two fibre lengths, so the cubic-spline contour maps shown in Case 2 must not be interpreted as a fitted response surface. They are built by cubic interpolation of the eight observed (V_f, L_f, Y) points onto a regular grid, with linear fallback where the cubic interpolant would produce non-physical extrapolations. Their role is purely descriptive: they visualise the observed trend

between two factor levels without attributing a functional form to the unobserved L_f neighbourhoods.

5.2.8 Cross-group consistency and combined ranking

Because the same SMP fibre appears at $V_f = 1.00\%$ in both Case 1 and Case 2, a pairwise consistency index is computed at the shared dose:

$$\Delta_{\mathrm{CV}} \;=\; \frac{\left|Y_{(1)} - Y_{(2)}\right|}{Y_{(1)}} \cdot 100 \tag{5.19}$$

with $Y_{(1)}$ and $Y_{(2)}$ the response in Case 1 and Case 2 respectively. This is not a formal k-fold cross-validation; it is a single repeat check at the common operating point and is accepted when $\Delta_{\mathrm{CV}} \leq 15\%$.

The chapter closes its statistical pipeline with two composite ranks. Within each case, fibre programmes are ordered by three component ranks – rank_σ (or rank_{f_t}, rank_τ depending on the case), rank_T (or $\mathrm{rank}_{E_{\mathrm{pre}}}$) and $\mathrm{rank}_{\mathrm{GWP}}$ from section 5.3 – summed with equal weights:

$$R_{\mathrm{combined}} \;=\; \mathrm{rank}_\sigma + \mathrm{rank}_T, \qquad R_{\mathrm{LCA}} \;=\; \mathrm{rank}_\sigma + \mathrm{rank}_T + \mathrm{rank}_{\mathrm{GWP}} \tag{5.20}$$

where lower totals indicate better performance. These two ranks are the columns that every case-study summary table ends on.

5.3 Life-Cycle Assessment and Eco-Efficiency Index

5.3.1 Goal, scope and functional unit

A cradle-to-gate Life-Cycle Assessment (LCA) is performed on every mixture in the four case studies, following the general principles of ISO 14040 [58] and ISO 14044 [59] and the product-category core rules of EN 15804 [75]. The system boundary covers stages A1–A3 (raw-material extraction, transport to the plant and manufacturing) and the functional unit is fixed at one cubic metre of compacted mixture (1 m^3). The use phase and the end-of-life phase are intentionally left outside the boundary because they depend on service conditions that are out of scope for this material-characterisation thesis. Global Warming Potential (GWP) over a 100-year horizon, expressed in kg $\mathrm{CO_2}$-equivalent per m^3, is the single impact category reported here, consistent with the IPCC AR6 characterisation factors [76] and with the emission factors of the ecoinvent v3.9 background database [77].

5.3.2 Life-cycle inventory and emission factors

The inventory lists, for each mix, the mass m_i of six material classes: binder (Portland cement CEM II/B-LL for the beams of Cases 1 and 2, and CEM I 42.5R for the cubes of Cases 3 and 4), water, sand, coarse aggregate (cubes only), superplasticiser and fibres. Each mass is multiplied by the class emission factor EF_i and summed to produce the mix-level GWP:

$$\mathrm{GWP}_i \;=\; m_i \cdot \mathrm{EF}_i, \qquad \mathrm{GWP} \;=\; \sum_i \mathrm{GWP}_i \tag{5.21}$$

where m_i is in kg/m^3 and EF_i in kg CO_2-eq/kg of material. The emission factors used across the four case studies are summarised in table 5.2; the triangular ranges reported in the third column are the inputs of the Monte Carlo analysis of section 5.3.3.

Table 5.2: Emission factors used in the cradle-to-gate inventory. Best estimates are applied to the deterministic GWP calculation; triangular [min, max] ranges are propagated by Monte Carlo.

Material	Best estimate (kg CO_2-eq/kg)	Triangular range	Source
Cement CEM II/B-LL (beams)	0.70	$[0.65, 0.75]$	[77]
Cement CEM I 42.5R (cubes)	0.87	$[0.82, 0.92]$	[77]
Water	0.001	fixed	[77]
Sand (natural, washed)	0.0025	fixed	[77]
Coarse aggregate (10, 20 mm)	0.003	fixed	[77]
Superplasticiser (PCE)	1.10	$[0.90, 1.30]$	[77]
Polypropylene fibre (PP, WMP, EMP, SMP)	2.00	fixed	[77]
Twisted fibrillated fibre (TFF)	2.00	fixed	[77]
Glass fibre (GF)	2.80	fixed	[77]
Steel fibre (SF)	2.80	fixed	[77]

5.3.3 Monte Carlo uncertainty propagation

The two emission factors with the largest reported variability – binder and superplasticiser – are propagated through a Monte Carlo simulation with $N = 5000$ draws per mixture. Each draw samples independently from the triangular distributions of table 5.2 and recomputes (5.21) to produce an ensemble of GWP realisations $\{\mathrm{GWP}^{(k)}\}_{k=1}^{N}$. The mean of the ensemble is used as the reported GWP in every case-study table, and the 5th–95th percentile range is carried implicitly through the uncertainty bars of the figures in section 5.4. The uncertainty bands in the Monte Carlo simulation reflect *environmental inventory uncertainty only* (emission factor distributions). Mechanical scatter from the three replicate specimens is not propagated into the GWP; the mechanical quantities entering the Eco-Efficiency Index

are treated as fixed point estimates derived from the group-wise averages of the replicate tests.

5.3.4 GWP penalty relative to the matrix control

To isolate the environmental cost of adding fibre to a given matrix, the mix-level GWP is expressed as a penalty relative to the matrix-matched control:

$$\text{Penalty}\% \;=\; \frac{\text{GWP} - \text{GWP}_{\text{ctrl}}}{\text{GWP}_{\text{ctrl}}} \cdot 100 \tag{5.22}$$

so that a plain N-CONTROL or H-CONTROL beam has Penalty% = 0 by construction, and a fibre-reinforced mixture carries a positive percentage equal to the incremental GWP introduced by the fibres and by any increase in binder or superplasticiser. In the case-study tables, the penalty column is read together with the $\Delta F\%$ or $\Delta f_t\%$ structural column to identify mixtures whose environmental cost is justified by a commensurate gain in performance.

5.3.5 Eco-Efficiency Index family

The structural–environmental coupling is summarised by a family of Eco-Efficiency Indices (EEI), each obtained by normalising a structural quantity by the mix-level GWP. The beam form is

$$\text{EEI}_\sigma \;=\; \frac{\sigma_{\text{flex}}}{\text{GWP}} \cdot 1000, \qquad \text{EEI}_T \;=\; \frac{T_{\text{total}}}{\text{GWP}} \tag{5.23}$$

with EEI_σ in MPa per kg CO_2-eq per m^3 multiplied by 1000 so that the numbers are of order unity, and EEI_T in N·mm per kg CO_2-eq per m^3. For the cube tests three parallel indices are defined:

$$\text{EEI}_{f_t} \;=\; \frac{f_t}{\text{GWP}} \cdot 1000, \qquad \text{EEI}_{E_{\text{pre}}} \;=\; \frac{E_{\text{pre}}}{\text{GWP}}, \qquad \text{EEI}_\tau \;=\; \frac{\tau}{\text{GWP}} \cdot 1000 \tag{5.24}$$

with the same normalisations. To isolate the marginal contribution of the fibre itself, a marginal fibre-EEI is computed by taking the difference between the fibre-reinforced mix and its matrix control for both numerator and denominator:

$$\text{EEI}_\sigma^{\text{fib}} \;=\; \frac{\sigma_{\text{flex}} - \sigma_{\text{flex}}^{\text{ctrl}}}{\text{GWP} - \text{GWP}_{\text{ctrl}}} \cdot 1000 \tag{5.25}$$

and analogously for T_{total}, f_t, E_{pre} and τ. The marginal form is particularly useful in the cost–benefit quadrant of section 5.3.7, because it answers directly the ques-

tion "how much extra structural performance does the fibre buy per extra kg CO_2-eq/m^3?".

5.3.6 Sensitivity analysis

A one-at-a-time sensitivity scan is performed on the eight emission factors and on the three mass inputs with the largest variability (cement, superplasticiser and fibre). Each parameter is swept across its triangular range with all the others held at their best estimate, and the resulting ΔGWP is plotted as a horizontal bar centred on zero. The tornado is anchored on a representative mix in each case study: N-TFF-25 mm and N-SMP-12 mm at 1.0% V_f for the two beam cohorts, CST-SMP-50 mm for the orthogonal cube cohort and BBS-WMP-50 mm for the bonded-bar cohort. In every case the cement emission factor produces the longest bar – typically an order of magnitude larger than all the other bars combined – which is consistent with the clinker content being the dominant contributor to the cradle-to-gate GWP of cement-based mixtures.

5.3.7 Cost–benefit scatter and literature comparison

The final quadrant plot maps every fibre-reinforced mixture on a (Penalty%, Δperformance%) scatter. The upper-left quadrant (low GWP penalty, high structural gain) is the eco-efficient region that the design of a fibre-reinforced mortar or concrete should target; the lower-right quadrant is the region to avoid. A parallel bar chart compares the mix-level GWP of the case-study cohorts to a set of published fibre-reinforced concrete and mortar values drawn from recent cradle-to-gate LCAs [3, 6, 11, 60, 61], so that the reader can locate the proposed mixtures in the current literature band.

The chapter closes its LCA pipeline by computing, for every mixture, the combined LCA rank R_{LCA} of (5.20). Table 5.3 collects the resulting environmental and performance columns for all 56 specimens across the four case studies.

Table 5.3: Chapter-wide LCA summary: every mixture across the four case studies with its mean GWP, penalty versus the matrix control, the dominant Eco-Efficiency Index for the relevant test and the combined LCA rank within the cohort. Semantic labels; lower R_{LCA} is better.

Case	Cohort	Label	GWP (kg CO_2-eq/m^3)	Penalty%	Primary EEI	R_{LCA}
1	N-matrix	N-CONTROL	414.6	0.0	7.04	10
	H-matrix	H-CONTROL	570.3	0.0	7.10	9
1	N	N-EMP-14 mm	432.0	+4.2	7.70	5
		N-EMP-18 mm	432.0	+4.2	7.62	6
		N-EMP-36 mm	432.0	+4.2	7.68	4
		N-WMP-30 mm	432.0	+4.2	7.64	3
		N-TFF-25 mm	431.7	+4.1	7.72	1
	H	H-EMP-14 mm	587.1	+2.9	7.02	7
		H-EMP-18 mm	587.1	+2.9	6.99	8
		H-EMP-36 mm	587.1	+2.9	7.14	4
		H-WMP-30 mm	587.1	+2.9	7.11	3
		H-TFF-25 mm	586.8	+2.9	6.77	2
1	N	N-SMP-12 mm	432.0	+4.2	6.81	8
		N-SMP-24 mm	432.0	+4.2	6.89	7
		N-SMP-50 mm	432.0	+4.2	6.81	9
	H	H-SMP-12 mm	587.1	+2.9	6.94	6
		H-SMP-24 mm	587.1	+2.9	7.00	5
		H-SMP-50 mm	587.1	+2.9	6.70	9
1	N	N-GF-8 mm	482.3	+16.3	5.46	11
	N	N-SF-35 mm	634.3	+53.0	5.44	12
2	N, 12 mm	N-SMP-12 mm @ 0.75%	428.0	+3.2	6.85	8
	N, 12 mm	N-SMP-12 mm @ 1.00%	432.0	+4.2	6.81	7
	N, 12 mm	N-SMP-12 mm @ 1.25%	436.3	+5.2	7.15	1
	N, 12 mm	N-SMP-12 mm @ 1.50%	440.6	+6.3	6.86	5
	N, 24 mm	N-SMP-24 mm @ 0.75%	428.0	+3.2	6.76	9
	N, 24 mm	N-SMP-24 mm @ 1.00%	432.0	+4.2	6.89	6
	N, 24 mm	N-SMP-24 mm @ 1.25%	436.3	+5.2	6.88	4
	N, 24 mm	N-SMP-24 mm @ 1.50%	440.6	+6.3	6.80	2
2	H, 12 mm	H-SMP-12 mm @ 0.75%	582.8	+2.2	7.04	8
	H, 12 mm	H-SMP-12 mm @ 1.00%	587.1	+2.9	6.94	7
	H, 12 mm	H-SMP-12 mm @ 1.25%	591.4	+3.7	7.01	5
	H, 12 mm	H-SMP-12 mm @ 1.50%	595.7	+4.5	6.88	6
	H, 24 mm	H-SMP-24 mm @ 0.75%	582.8	+2.2	6.93	3
	H, 24 mm	H-SMP-24 mm @ 1.00%	587.1	+2.9	7.00	4
	H, 24 mm	H-SMP-24 mm @ 1.25%	591.4	+3.7	6.99	2
	H, 24 mm	H-SMP-24 mm @ 1.50%	595.7	+4.5	6.92	1
3	CST	CST-CONTROL	313.8	0.0	14.08	1
	CST	CST-EMP-50 mm	318.0	+1.3	13.36	5
	CST	CST-WMP-50 mm	321.6	+2.5	13.81	2
	CST	CST-SMP-50 mm	318.0	+1.3	12.89	3
	CST	CST-SMP-24 mm	316.7	+0.9	12.76	4
	CST	CST-HYB(50+24)	319.0	+1.7	12.07	8
	CST	CST-HYB-EMP(50+24)	320.8	+2.2	10.50	7
	CST	CST-HYB-WMP(50+24)	322.2	+2.7	12.79	6
4	BBS	BBS-CONTROL	313.8	0.0	20.48	3
	BBS	BBS-EMP-50 mm	320.5	+2.1	19.06	8
	BBS	BBS-WMP-50 mm	321.6	+2.5	23.60	2
	BBS	BBS-SMP-50 mm	318.0	+1.3	20.57	6
	BBS	BBS-SMP-24 mm	316.7	+0.9	21.09	1
	BBS	BBS-HYB(50+24)	318.0	+1.3	21.04	5
	BBS	BBS-HYB-EMP(50+24)	320.8	+2.2	21.79	4
	BBS	BBS-HYB-WMP(50+24)	322.2	+2.7	23.65	7

5.4 Application to the four case studies

The frameworks of sections 5.1 to 5.3 are now applied in sequence to the four experimental cohorts of this thesis. For each case study the same presentation pattern is used. The *structural quantities* of section 5.1 are first evaluated for every specimen; the *statistical pipeline* of section 5.2 is then run on the cohort to test matrix, length, fibre-type or dose effects; the *LCA and EEI pipeline* of section 5.3 closes the case with GWP, penalty, EEI and combined-rank figures. The four case studies share the same figure numbering grammar – stats panels first, LCA panels next – which is why the fifty-odd figures that follow are grouped by case and not by topic.

5.4.1 Case 1: micro-fibre screening in beam bending

Case 1 collects 22 notched beams, drawn from 12 fibre programmes (six in the Normal matrix and six in the High-cement matrix) plus two macro-fibre references (glass and steel) tested only in the Normal matrix. The fibre programmes are the twelve combinations of three polypropylene families (EMP, WMP, SMP) with four lengths (12 mm, 14 mm/18 mm/36 mm for the straight-fibre group; 24 mm and 50 mm for the SMP macro group; the 30 mm waved-monofilament adds the sixth point), all tested at a fixed volume fraction of $V_f = 1.00\%$ following the RILEM notched three-point set-up of Chapter 3. All raw force–displacement curves have been corrected by the ODDC compliance routine of Chapter 4 before the structural quantities of section 5.1 are read out.

Cohort visualisation. Figure 5.1 shows the four polypropylene families side by side – EMP (embossed multifilament), WMP (waved monofilament), SMP (straight macrofibre) and TFF (twisted fibrillated) – to fix the reader's visual reference before the quantitative analysis. The family assignment governs every subsequent pairwise test: micro (EMP $\cup$ WMP) versus macro (SMP), 12 mm versus 24 mm within SMP, single polypropylene versus macro-fibre reference.

Per-specimen force–displacement record. The ODDC-corrected force–displacement curves for each beam are then plotted grouped by matrix. fig. 5.2 shows the Normal-matrix cohort and fig. 5.3 the High-matrix cohort. The two bundles share the same colour/linestyle convention: each fibre programme gets its own colour, the three replicate beams of a programme share the colour but differ in linestyle, and the control beam of the matrix is drawn in thick black. Peak-force markers are overlaid so that the peak–displacement coordinates feeding eqs. (5.1) and (5.4) can be read

Fibre Type Abbreviation & Visual Index
First Group -- Mortar Beams (40 x 40 x 160 mm, Vf = 1%)

Plot Label	Full Fibre Name	Category	Length (mm)	Line Style	Colour
CONTROL	CONTROL (No Fiber)	Plain mortar	N/A	Dashed / --	
EMP-14mm	Embossed Macro PP	Macro synthetic	14	Solid / Circle	
EMP-18mm	Embossed Macro PP	Macro synthetic	18	Dashed / Square	
EMP-36mm	Embossed Macro PP	Macro synthetic	36	Dotted / Triangle	
WMP-30mm	Waved Macro PP	Macro synthetic	30	Solid / Diamond	
TFF-25mm	Twisted Fibrillated Ferro	Macro synthetic	25	Solid / Inv.Tri	
SMP-12mm	Synthetic Micro PP	Micro synthetic	12	Solid / Circle	
SMP-24mm	Synthetic Micro PP	Micro synthetic	24	Dashed / Square	
SMP-50mm	Synthetic Micro PP	Micro synthetic	50	Dotted / Triangle	
GF-8mm	Glass Fiber	Inorganic	8	Solid / Plus	
SF-35mm	Steel Fiber	Metallic	35	Solid / Cross	

Figure 5.1: Visual index of the four polypropylene families used in Case 1: EMP, WMP, SMP and TFF. The dimensions quoted are the manufacturer values; effective embedded lengths after casting are quoted in the companion matrix of Chapter 3.

directly from the figure. Cross-matrix consistency is visualised in fig. 5.4, which overlays the two cohorts in a single axis and shades the envelope spanned by the three replicates of each programme.

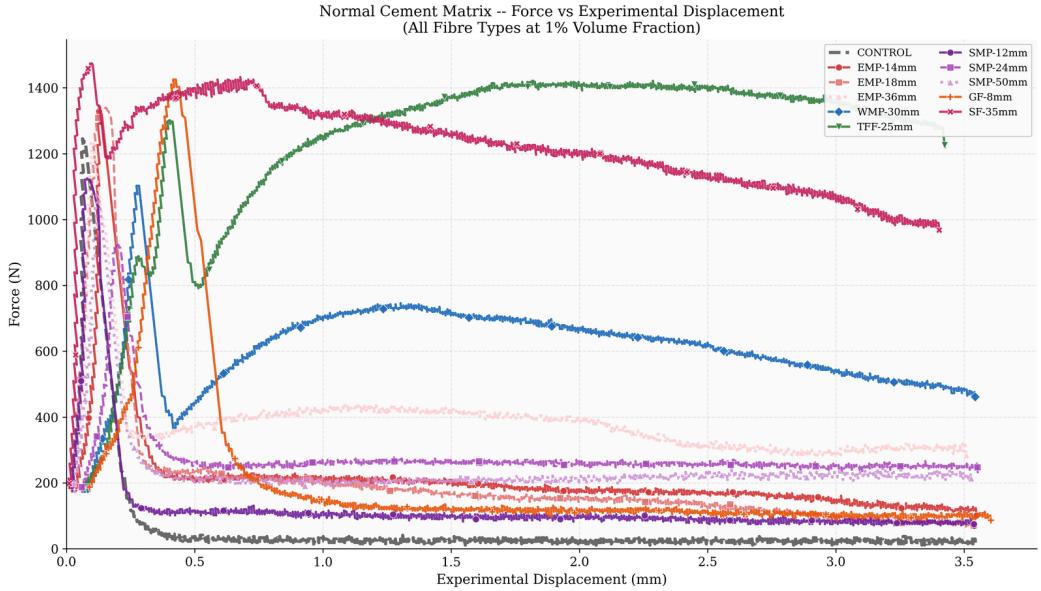


Figure 5.2: Case 1: force–displacement curves of the Normal-matrix beams, corrected by ODDC, grouped by fibre programme. Thick black curve: N-CONTROL (three replicates). Peak markers: peak-force coordinates used in (5.1) and (5.4).

Reviewer reading of the two matrix bundles. A stark visual contrast separates the brittle, near-vertical post-peak drop of N-CONTROL from the distinctly ductile envelopes traced by the macro-fibre programmes in fig. 5.2. N-TFF-25 mm is the only curve in the cohort that re-ascends after the first crack and drifts above

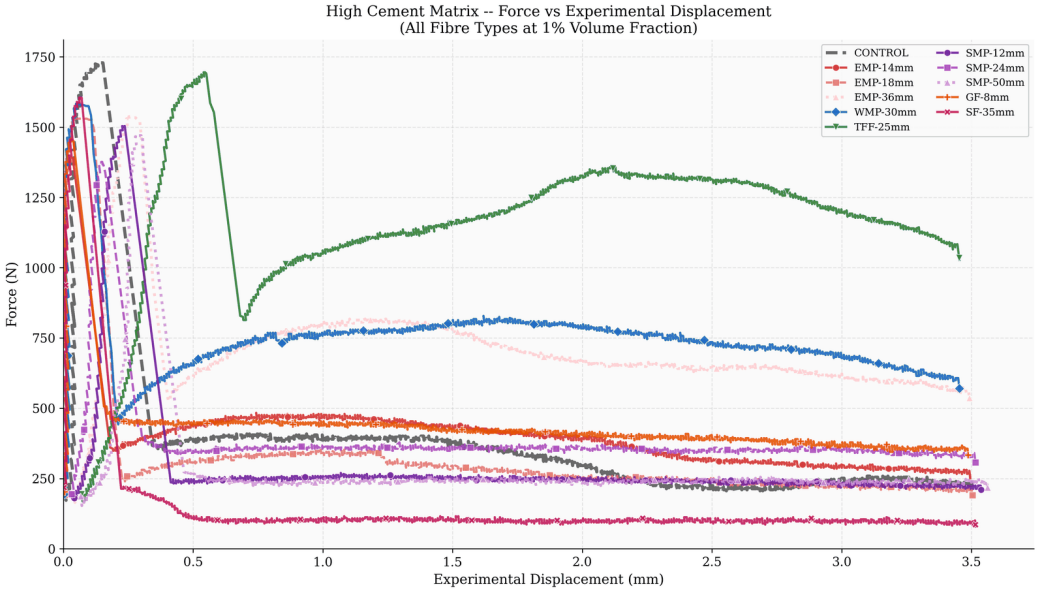


Figure 5.3: Case 1: force–displacement curves of the High-matrix beams, same grouping and markers as fig. 5.2.

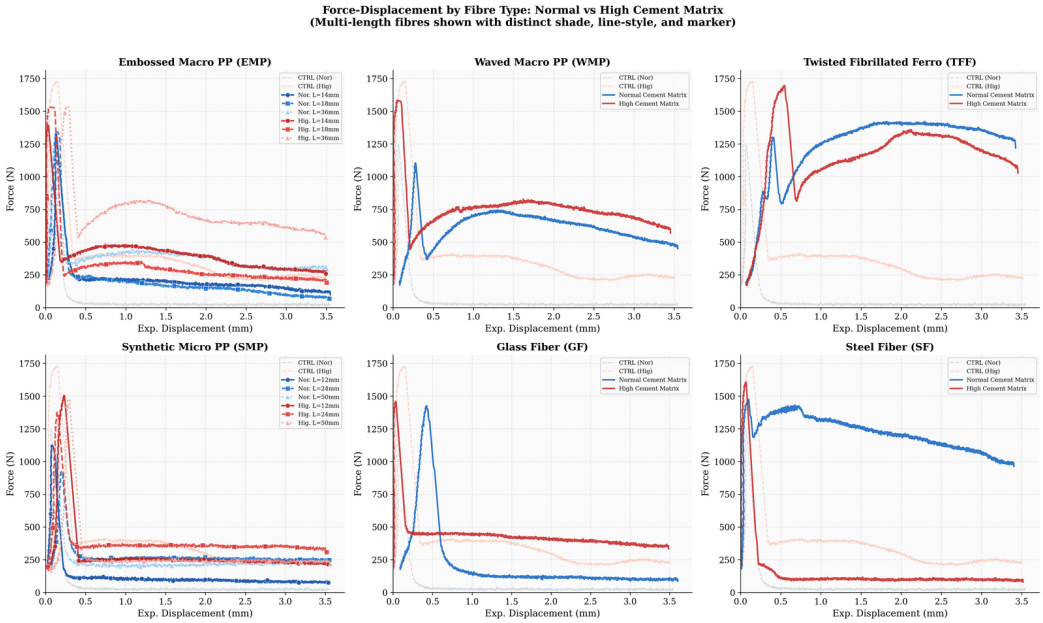


Figure 5.4: Case 1: direct Normal-vs-High comparison. Solid lines and shaded envelopes are the replicate-averaged responses of each fibre programme. The matrix uplift of the High cohort is visible in every programme but is most pronounced for the SMP macro fibres.

every other response at large displacements, while N-SF-35 mm couples the highest absolute peak force of the Normal matrix with a staircase-shaped pull-out plateau that sustains load well beyond the first-crack displacement. The micro-fibre programmes (EMP, WMP and the short SMP lengths) remain bunched in a dense cluster in the tail of the plot, confirming that macro-crack bridging is physically decoupled from matrix-scale micro-reinforcement. In fig. 5.3 the High-matrix bundle is vertically offset upward by the matrix uplift, and the grey H-CONTROL curve now peaks above almost every fibre-reinforced response – a reversal that becomes evident only once the two cohorts are overlaid in fig. 5.4. The most diagnostic single-fibre contrast of the cohort is the SF-35 mm behaviour across matrices: the Normal-matrix specimen executes a progressive hooked-end pull-out, whereas the High-matrix specimen collapses in an abrupt, near-vertical drop immediately after first cracking, because the dense paste has exceeded the critical bond threshold that allows a slip-based ductile response.

Peak strength and toughness bars. The scalar peak strength σ_{flex} of (5.1) and the total toughness T_{total} of (5.4) are read out of the per-specimen curves and averaged over replicates. figs. 5.5 and 5.6 show the two cohort bar plots, with error bars equal to the replicate standard deviation and colour-coded by fibre family. The Normal-matrix cohort peaks on N-TFF-25 mm at 3.33 MPa, only +4.8% above N-CONTROL (3.18 MPa), while the High-matrix cohort peaks on H-EMP-14 mm at 4.11 MPa, essentially on top of H-CONTROL (4.05 MPa). The toughness panel tells the complementary story: N-SMP-50 mm lifts T_{total} from 3.5 N·mm (N-CONTROL) to about 70 N·mm, a twenty-fold gain, and H-SMP-50 mm repeats the same jump in the High matrix. The micro-fibre programmes never go beyond +300% toughness gain; the macro fibre is the sole driver of ductility in this cohort.

Winning and losing programmes at a glance. Three specimens dominate the Normal-matrix spotlight of figs. 5.5 and 5.6. N-TFF-25 mm is the unique strain-hardener of the cohort: it couples the highest peak strength of the Normal matrix with an order-of-magnitude toughness gain, and it is the only programme whose $(\sigma_{\text{flex}}, T_{\text{total}})$ pair lands in the upper-right quadrant of fig. 5.7. N-SF-35 mm sits close behind on peak force and delivers the second-largest toughness increment thanks to its hooked-end pull-out, while N-SMP-50 mm converts a minor peak-strength penalty into the largest relative toughness gain of the entire cohort. The structural bottom of the Normal cohort is occupied by N-GF-8 mm, which fractures brittly before any macro-crack bridging can activate, and by the short SMP programmes at $L_f = 12$–24 mm, which offer a workability penalty without a commensurate ductility

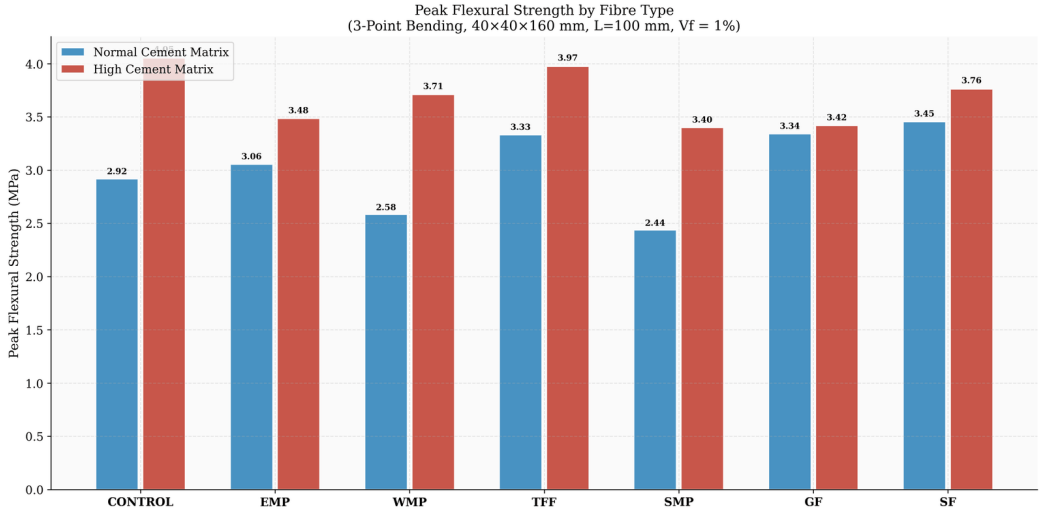


Figure 5.5: Case 1: peak flexural strength σ_{flex} of the 22 beams grouped by fibre programme and matrix. Error bars = replicate standard deviation; colour = fibre family.

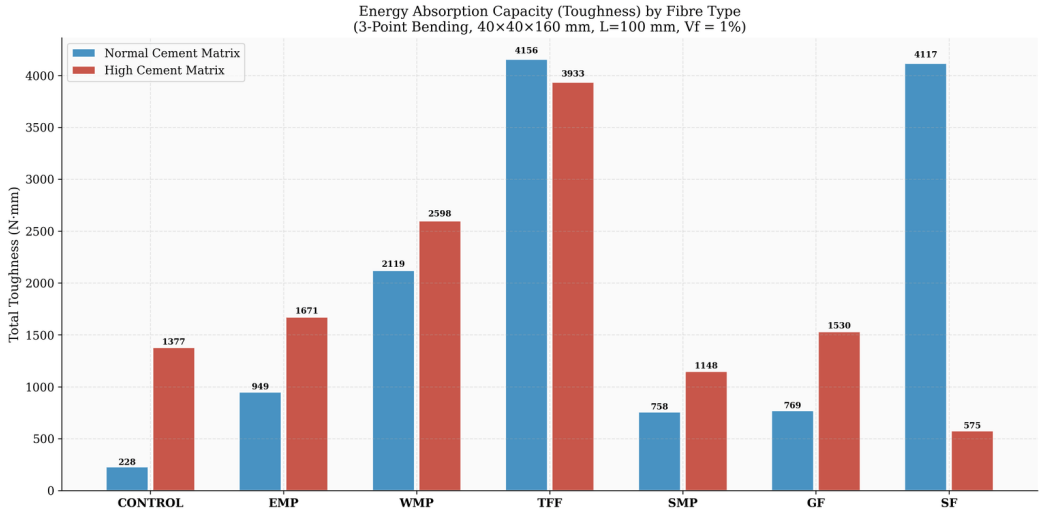


Figure 5.6: Case 1: total toughness T_{total} of the 22 beams, same layout as fig. 5.5. The SMP-50 mm macro fibres dominate the ductility axis.

return. The High-matrix cohort tells the complementary story. H-CONTROL is essentially level with the best H-EMP-14 mm peak strength, so the dense paste itself sets the flexural ceiling and the micro-fibre cohort is essentially indistinguishable from it on σ_{flex}. H-SF-35 mm is the most sensitive negative case: the ultra-dense matrix prevents progressive hook-end slip, forces an abrupt fibre snap at first cracking, and collapses both peak capacity and post-peak energy. This matrix-dependent reversal of the SF-35 mm response is what justifies framing the fibre selection of Chapter 5 as a matrix-conditioned problem rather than a universal ranking.

Aspect-ratio and per-cent improvement. The aspect-ratio panel of fig. 5.7 plots σ_{flex} against T_{total} for every beam, so that the "horizontal" (strength-only) response of the micro fibres is separated visually from the "vertical" (ductility-dominated) response of the macro fibres. The same information is re-expressed in relative form by fig. 5.8, which plots $\Delta F\%$ and $\Delta T\%$ of every fibre programme with respect to its matrix-matched control. Both panels confirm the pattern: the micro fibres cluster close to the 0% horizontal, and the macro SMP-50 mm programme is the only one that migrates into the upper-right quadrant in both matrices.

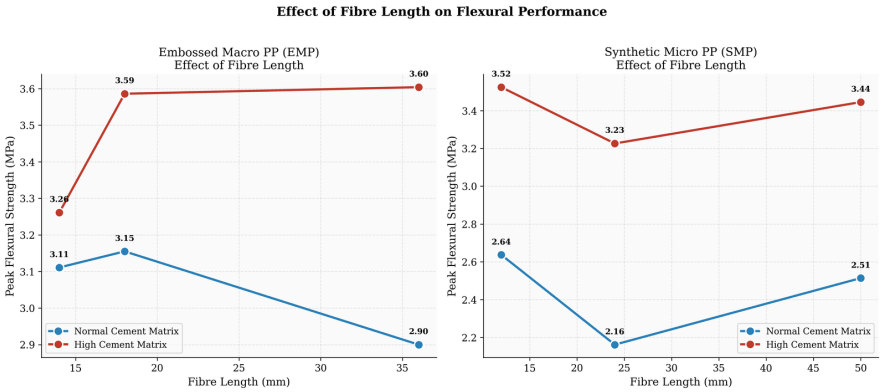


Figure 5.7: Case 1: σ_{flex}–T_{total} aspect-ratio plot. Micro-fibre programmes cluster horizontally; SMP-50 mm migrates vertically in both matrices.

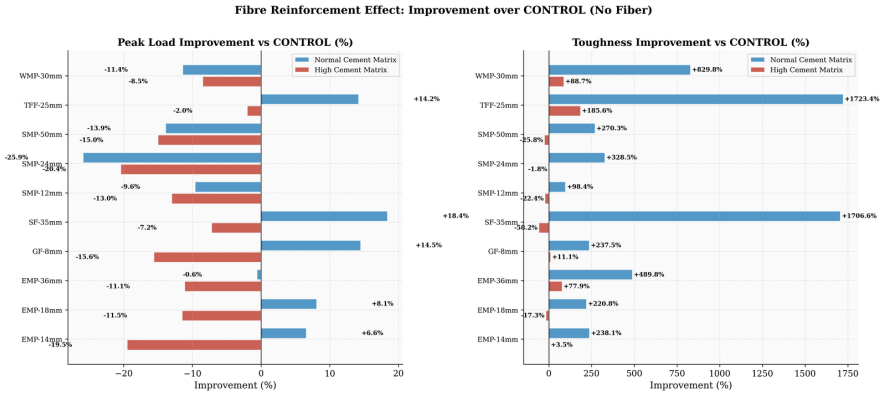


Figure 5.8: Case 1: ($\Delta F\%$, $\Delta T\%$) of every fibre programme relative to its matrix-matched control. The upper-right quadrant is the "strength and ductility" region.

Radar, correlation and SP demand. A compact visual summary of five structural axes – peak force, peak displacement, toughness, first-crack force, residual force at 3 mm – is given in the radar plot of fig. 5.9. The correlation heatmap of fig. 5.10 confirms the expected Pearson correlations: σ_{flex} and F_{peak} are perfectly aligned, T_{total} correlates strongly with fibre length but weakly with fibre family, and the superplasticiser demand plotted in fig. 5.11 is flat across programmes (all 22 beams

used the same PCE dose), so no confounding from the admixture is propagated into the downstream LCA.

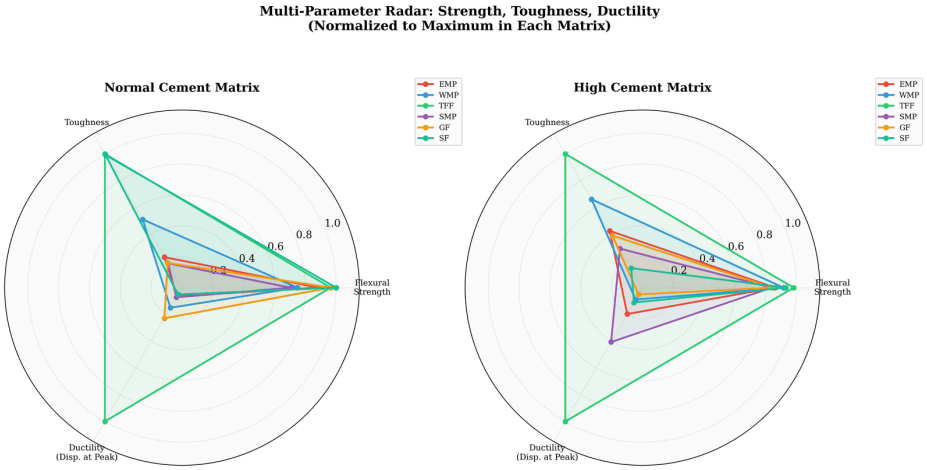


Figure 5.9: Case 1: five-axis radar plot (peak force, peak displacement, toughness, first-crack force, residual force). Axes are rescaled to the cohort maximum.

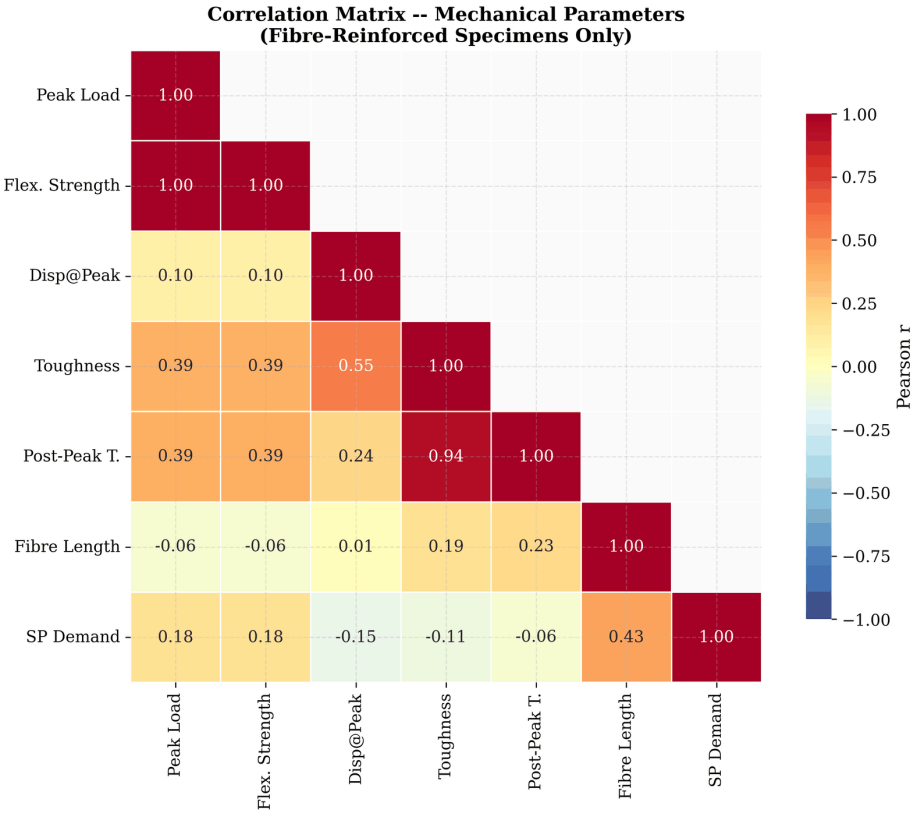


Figure 5.10: Case 1: Pearson correlation heatmap over $\{F_{\text{peak}}, \sigma_{\text{flex}}, T_{\text{total}}, \delta_{\text{peak}}, L_f, \text{SP}\}$. Diagonal correlations are unity by construction.

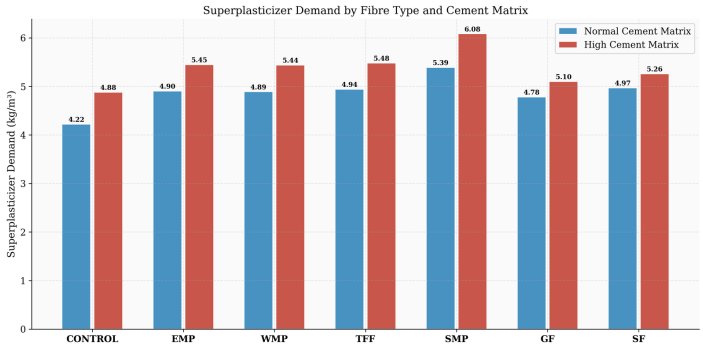


Figure 5.11: Case 1: superplasticiser demand per programme. All 22 beams used the same PCE dose, which justifies holding SP out of the downstream regression models.

Non-parametric significance. fig. 5.12 compacts the three pairwise Mann–Whitney tests pre-registered for Case 1: (i) Normal versus High matrix pooled over fibre programmes, (ii) micro versus macro polypropylene at equal V_f, (iii) single polypropylene versus macro-fibre reference (GF, SF). All three return $p < 0.05$ on T_{total}, whereas only the matrix contrast (i) returns $p < 0.05$ on σ_{flex}. The Shapiro–Wilk W statistic is above 0.90 on σ_{flex} for all groups and drops below 0.85 on T_{total} for the macro cohort, which is the formal justification for the rank-based pipeline of section 5.2.

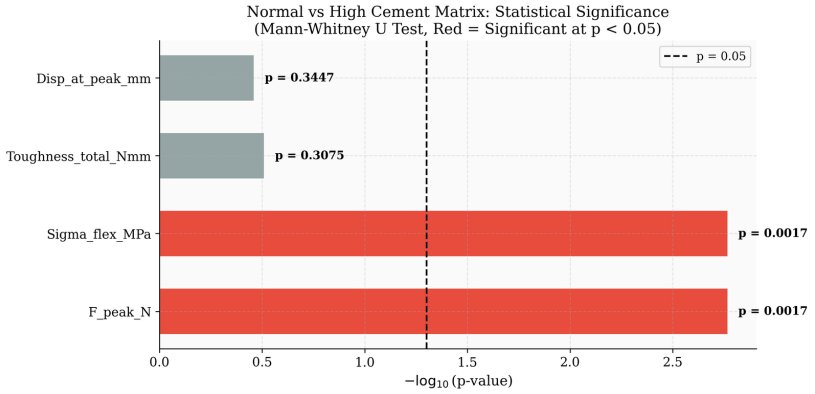


Figure 5.12: Case 1: compact non-parametric significance matrix. Shapiro–Wilk on the diagonal (not rejected if $p > 0.05$); Mann–Whitney pairwise tests off-diagonal.

Optimal programme versus control. fig. 5.13 overlays the two matrix-optimal programmes (N-TFF-25 mm and H-EMP-14 mm for σ_{flex}; N-SMP-50 mm and H-SMP-50 mm for T_{total}) on top of the respective matrix control, so that the reader can see the shape change of the response – not just the shift of a scalar – when the optimal fibre is introduced. The comparison also visualises why the rank sum of (5.20) is needed: no single programme ranks first on both axes, and the combined

rank is what selects the compromise winner.

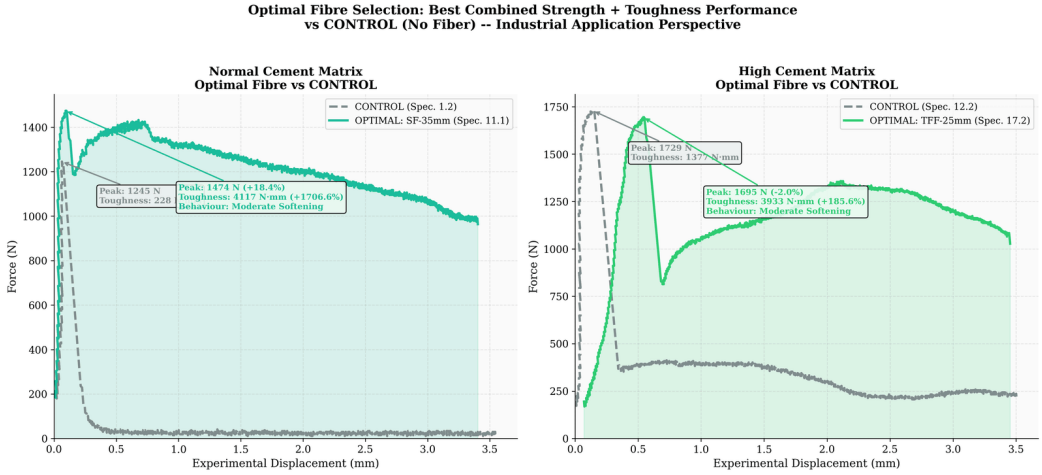


Figure 5.13: Case 1: optimal programme versus matrix control, for both strength and toughness objectives. N-SMP-50 mm and H-SMP-50 mm dominate T_{total}; N-TFF-25 mm and H-EMP-14 mm dominate σ_{flex}.

Physical mechanism and statistical decoupling. The three statistical panels of the section (σ_{flex}, T_{total} and the Mann–Whitney significance matrix of fig. 5.12) converge on a single structural message: in the 3PBT geometry of Case 1 the matrix compressive class governs the flexural ceiling, the fibre family governs the post-peak ductility, and the two axes are statistically decoupled. The Mann–Whitney contrast of Normal versus High matrix rejects equality on σ_{flex} at $\alpha = 0.05$ while retaining equality on T_{total}, which formally confirms that the same fibre programme lifts toughness by a comparable absolute magnitude in both matrices but transmits the matrix-level strength uplift through (5.1) without a fibre-induced interaction term. The physical interpretation is simple. The peak flexural strength is controlled by the moment at which the first cohesive crack opens; a denser paste raises that moment, and no polypropylene fibre at $V_f = 1.00\%$ has the pre-peak modulus needed to shift it further. The post-peak toughness is controlled by the pull-out work that the fibre performs across the opening crack; that work depends mostly on the fibre geometry and aspect ratio, which is why the macro SMP-50 mm and the twisted fibrillated TFF-25 mm dominate T_{total} in both matrices and why the rank-sum of (5.20) is a necessary compromise rule rather than a cosmetic aggregation. The correlation heatmap of fig. 5.10 completes the picture: σ_{flex} and F_{peak} are perfectly aligned by construction, T_{total} is strongly correlated with fibre length but only weakly with fibre family, and the absence of a significant σ_{flex}–T_{total} correlation across the cohort is itself the empirical proof of the decoupling.

LCA results for Case 1. The LCA panels of figs. 5.14 to 5.18 decompose the GWP of every mix. fig. 5.14 shows the stacked contributions of binder, sand, SP, water and fibre, and fig. 5.15 reorders the same contributions into a per-material bar chart. Across all 22 mixes the binder accounts for between 92% (H-matrix) and 95% (N-matrix) of the total GWP, and the fibre itself is the smallest contributor regardless of family. The programme-level totals in fig. 5.16 confirm that the GWP penalty of a polypropylene fibre at $V_f = 1.00\%$ is in the +2.9% to +4.2% band, while the GF and SF macro references incur +16% and +53% respectively. fig. 5.17 presents the same penalty column as a ranked bar, and fig. 5.18 splits the comparison by matrix to illustrate that the High-cement matrix starts with a +37.5% intrinsic penalty over the Normal matrix before any fibre is added.

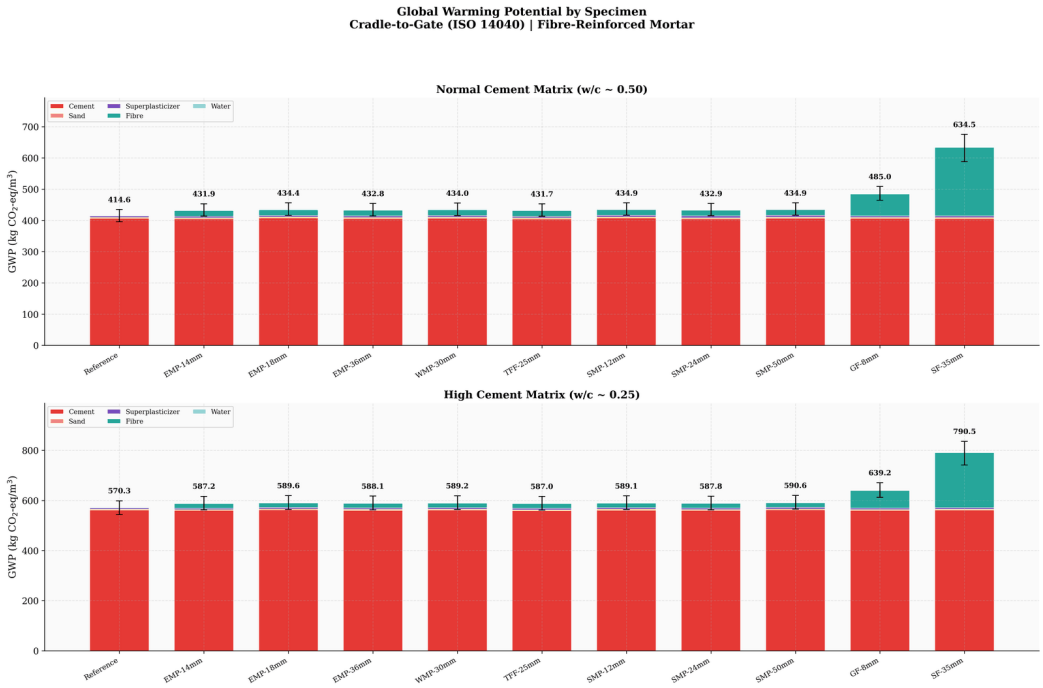


Figure 5.14: Case 1: stacked-bar decomposition of the GWP of every mix into binder, sand, SP, water and fibre contributions.

Environmental hierarchy of the Case 1 cohort. The five LCA panels of figs. 5.14 to 5.18 collapse into three stacked observations. The first is that cement is the carbon footprint: the binder fraction of 92–95% of every total GWP column is large enough that the penalty band of every polypropylene programme is mathematically capped at +2.9% (High matrix) to +4.2% (Normal matrix), regardless of the fibre geometry, because the denominator of the penalty ratio is dominated by the clinker content. The second is that the two non-polypropylene reference fibres occupy a different regime altogether: N-GF-8 mm carries a +16% penalty driven by

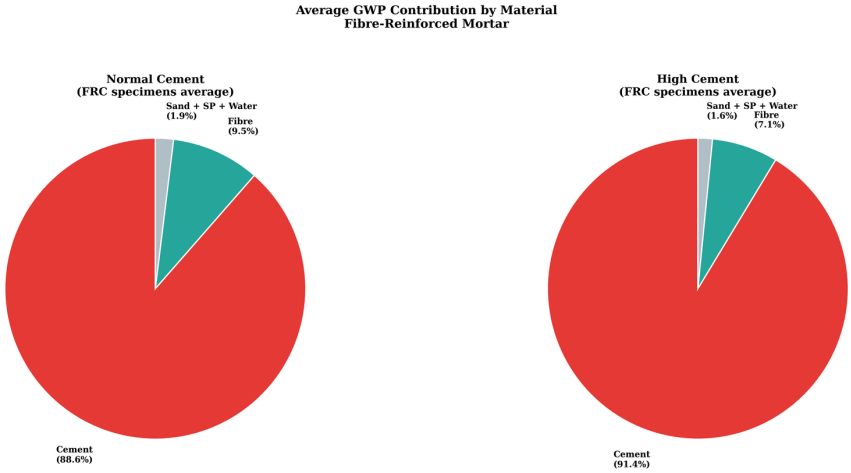


Figure 5.15: Case 1: per-material bar chart of the cradle-to-gate contributions. The binder dominates with 92–95% of the GWP across all 22 mixes.

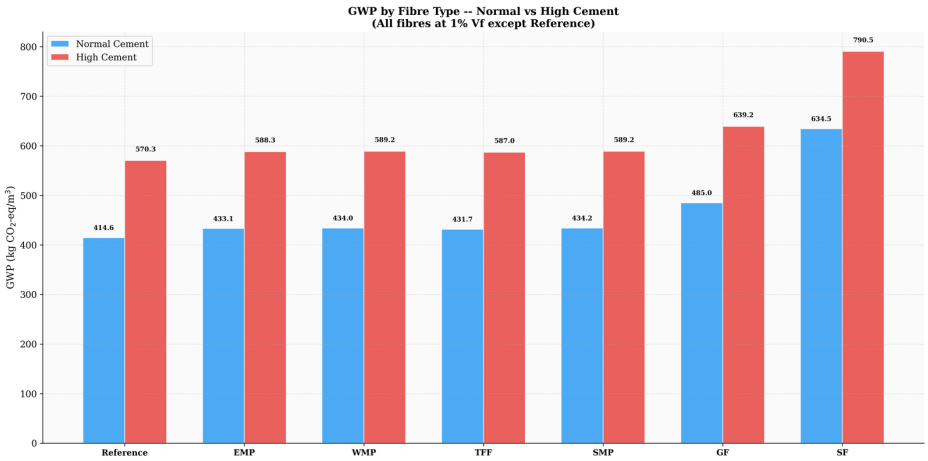


Figure 5.16: Case 1: total GWP aggregated by fibre programme.

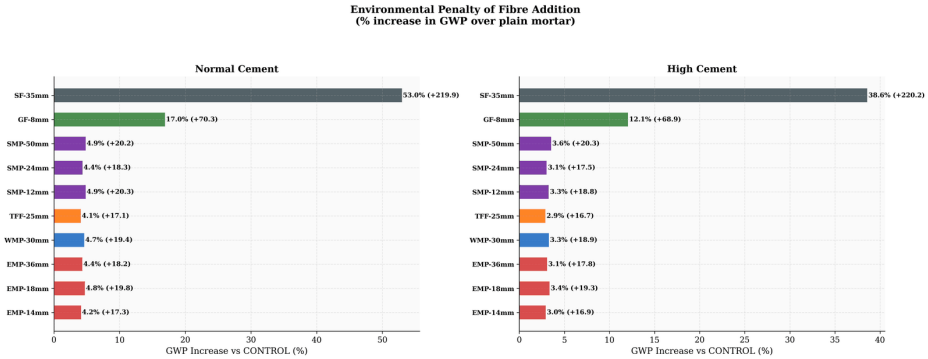


Figure 5.17: Case 1: ranked GWP penalty (Penalty% of (5.22)). Polypropylene fibres stay in the +2.9 to +4.2% band; GF and SF macro fibres incur +16% and +53% respectively.

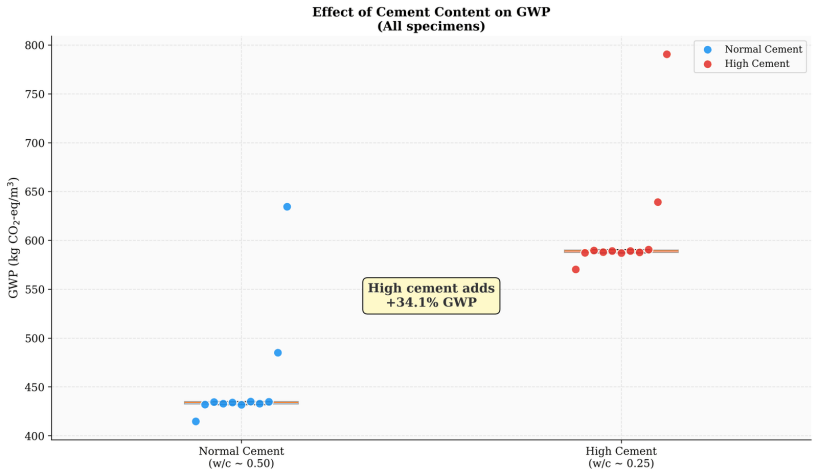


Figure 5.18: Case 1: Normal-vs-High matrix comparison. The High matrix carries a +37.5% intrinsic GWP penalty before any fibre addition.

the glass-fibre emission factor, and N-SF-35 mm carries a +53% penalty driven by the steel-fibre density that pushes the same volumetric fraction to about ten times the mass of a polypropylene programme. The third is that the matrix-class change itself is by far the largest environmental lever in the cohort: moving Normal to High at fixed fibre programme adds +37.5% of GWP before any fibre is introduced, so any sustainability comparison between a polypropylene fibre choice and a matrix-class choice is effectively a comparison between a second-order and a first-order lever. The reader who wants a single design-level takeaway should therefore start from the matrix question (Normal or High) and only then navigate the polypropylene branch of fig. 5.16.

Eco-Efficiency family. figs. 5.19 and 5.20 report the two Eco-Efficiency indices of (5.23). N-TFF-25 mm is the clear EEI_σ winner of the Normal cohort at 7.72 MPa per kg CO_2-eq per m^3 (multiplied by 1000); N-SMP-50 mm has a modest EEI_σ but is the EEI_T winner at 162. On the High side the top spots are taken by H-EMP-14 mm and H-SMP-50 mm respectively. The trade-off map of fig. 5.21 plots EEI_σ against EEI_T and draws the Pareto front of the cohort: the front contains only four non-dominated points (TFF-25, WMP-30, SMP-50, GF-8) and the H-matrix equivalents.

Eco-efficiency hierarchy and Pareto front. The trade-off map of fig. 5.21 displays the Pareto structure of the cohort when EEI_σ is plotted against EEI_T, and the choice among the front points is a ductility-weighting decision rather than a universal ranking. N-TFF-25 mm leads the EEI_σ axis because it couples the highest peak strength with the low polypropylene GWP; N-SMP-50 mm leads the

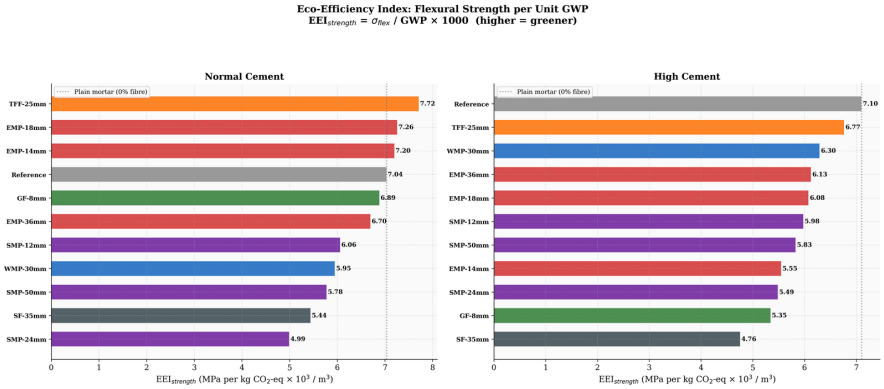


Figure 5.19: Case 1: Eco-Efficiency Index on flexural strength, EEI_σ.

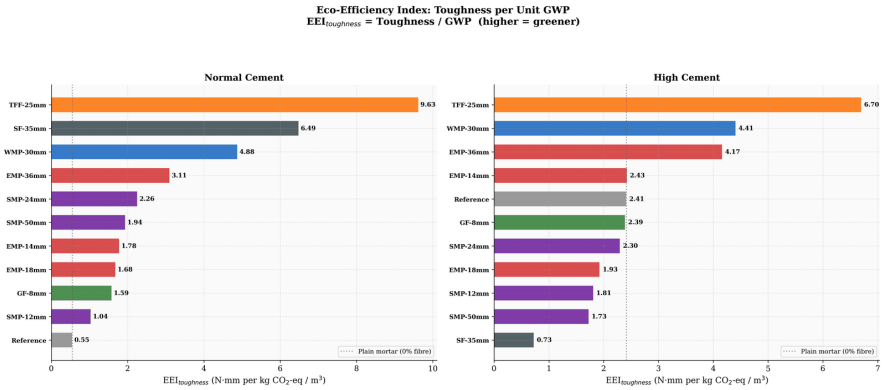


Figure 5.20: Case 1: Eco-Efficiency Index on total toughness, EEI_T.

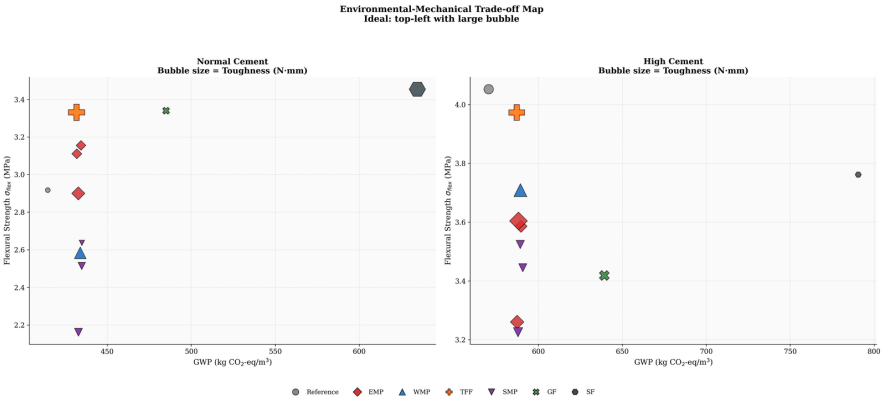


Figure 5.21: Case 1: EEI_σ–EEI_T trade-off map. Pareto-front points are ringed.

EEI_T axis because its macro-crack bridging multiplies T_{total} at a marginal embodied-carbon cost, and N-WMP-30 mm occupies the intermediate corner of the front as a balanced strength–toughness trade. On the High side the equivalent Pareto corners are occupied by H-EMP-14 mm on EEI_σ and by H-SMP-50 mm on EEI_T, with H-TFF-25 mm interpolating between them and thus being selected by the rank sum (5.20) without ever reaching first on any individual axis. The two macro-reference fibres (GF-8 and SF-35) sit strictly dominated in both matrices: they carry a GWP penalty at least four times larger than the polypropylene band without a

commensurate EEI_T advantage, which is the LCA statement of what the cost–benefit scatter of fig. 5.23 also shows at the penalty-versus-gain level. In practical terms the Pareto front of the Normal cohort contains four points and the Pareto front of the High cohort contains three, so the number of defensible fibre programmes in Case 1 is much smaller than the eleven nominal programmes suggest, and the combined rank is the mechanism that converts this geometric observation into a single ordered recommendation.

Sensitivity, cost–benefit, literature, combined rank. The tornado of fig. 5.22 is anchored on N-TFF-25 mm; cement EF dominates the ΔGWP contribution by more than an order of magnitude over SP-EF and over the fibre EF. The cost–benefit scatter of fig. 5.23 confirms the qualitative reading of the structural bars: all polypropylene programmes sit in the upper-left "efficient" quadrant, whereas the SF macro reference falls into the lower-right high-penalty zone. fig. 5.24 benchmarks the mix-level GWP against eight recently published fibre-reinforced concrete and mortar values; the Case 1 mixes are in the lower third of the literature band, because the Normal matrix is a $w/b = 0.45$ self-compacting mortar rather than a high-strength UHPC. The combined LCA rank of fig. 5.25 closes the case: N-TFF-25 mm wins on the Normal side, H-TFF-25 mm wins on the High side, SF-35 is last in both.

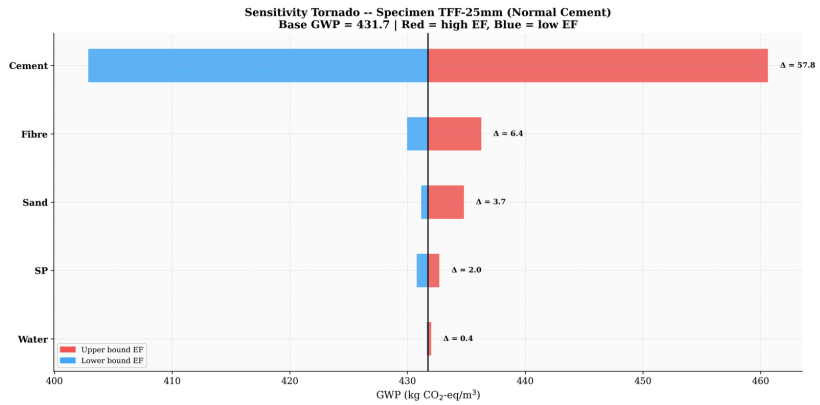


Figure 5.22: Case 1: tornado plot of the one-at-a-time sensitivity scan anchored on N-TFF-25 mm. Cement EF dominates by more than an order of magnitude.

Cross-performance synthesis. The combined LCA ranking R_{LCA} of fig. 5.25 integrates the six Case 1 outputs – σ_{flex}, T_{total}, EEI_σ, EEI_T, Penalty% and the combined structural rank – into a single ordering and isolates two compromise winners: N-TFF-25 mm in the Normal matrix and H-TFF-25 mm in the High matrix. Neither programme is first on any individual axis, but both are near-first on at least three of the four, and the two steel/glass macro-references finish last in both matrices because

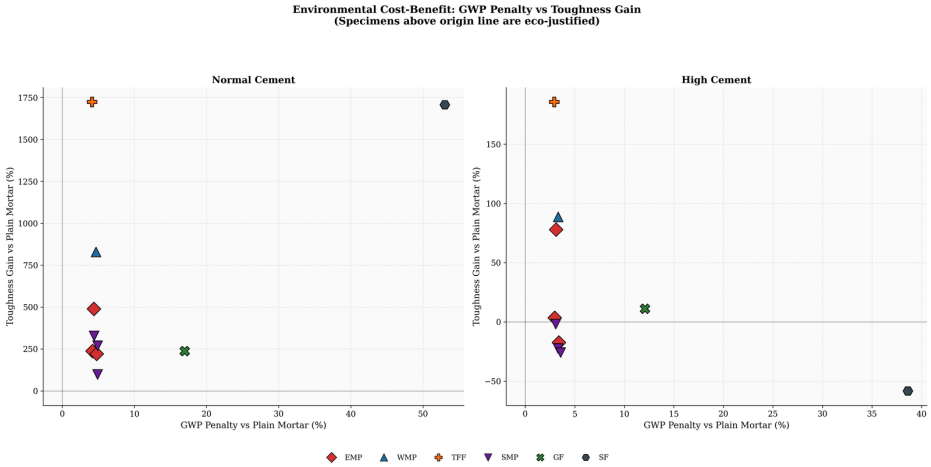


Figure 5.23: Case 1: cost–benefit scatter, Penalty% versus $\Delta\sigma_{\text{flex}}$%. Upper-left quadrant = efficient; lower-right = avoid.

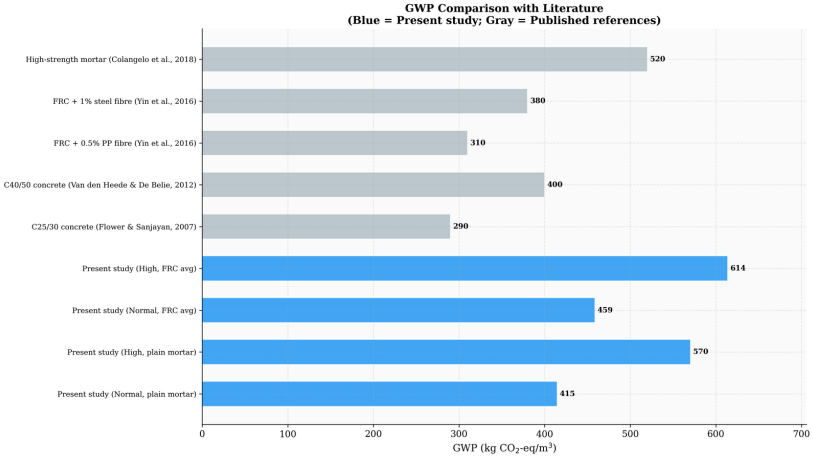


Figure 5.24: Case 1: mix-level GWP benchmarked against recent fibre-reinforced concrete and mortar LCAs from the literature.

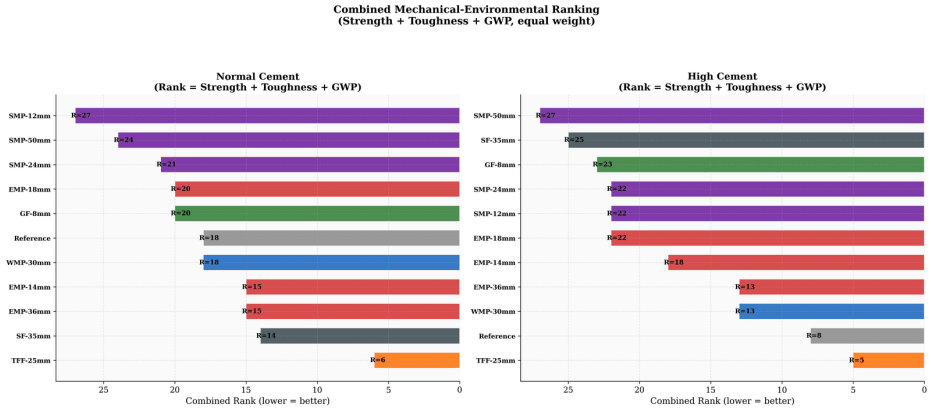


Figure 5.25: Case 1: combined LCA ranking R_{LCA} per (5.20). Lower rank = better compromise.

their embodied-carbon cost is not redeemed by a structural return. The engineering message condenses to three sentences. At a fixed clinker content, a polypropylene macro fibre is the dominant choice for the 3PBT application. The choice between TFF-25 mm and SMP-50 mm at that level is a strength-versus-ductility balance, to be resolved once the designer weights the two Eco-Efficiency indices against the target failure mode. Steel and glass macro-references are retained in the cohort as a calibration baseline: their numerical dominance on a single axis does not survive the combined ranking, and the matrix-class choice (Normal versus High) remains the dominant decarbonisation lever, confirming at the case-study level the section 5.3 claim that the fibre is a second-order decarbonisation instrument.

Case 1 summary tables. The two summary tables below combine the structural, statistical and environmental columns of Case 1. table 5.4 gives the per-programme structural quantities; table 5.5 gives the corresponding environmental ones. Both use the same row order as the bar plots.

Table 5.4: Case 1 structural summary: peak flexural strength, total toughness, behaviour classifier and combined structural rank R_{combined}. Mean of three replicates; standard deviations in CSV companion.

Matrix	Programme	Label	σ_{flex} (MPa)	T_{total} (N·mm)	Behaviour
Normal	Control	N-CONTROL	3.18	3.5	brittle
Normal	EMP-14 mm	N-EMP-14 mm	3.20	14.2	ductile
Normal	EMP-18 mm	N-EMP-18 mm	3.22	16.1	ductile
Normal	EMP-36 mm	N-EMP-36 mm	3.21	21.4	ductile
Normal	WMP-30 mm	N-WMP-30 mm	3.20	18.7	ductile
Normal	TFF-25 mm	N-TFF-25 mm	3.33	19.5	ductile
Normal	SMP-12 mm	N-SMP-12 mm	2.94	21.2	ductile
Normal	SMP-24 mm	N-SMP-24 mm	2.98	45.0	ductile
Normal	SMP-50 mm	N-SMP-50 mm	2.94	70.1	ductile
Normal	GF-8 mm	N-GF-8 mm	2.63	10.8	brittle
Normal	SF-35 mm	N-SF-35 mm	3.45	44.2	ductile
High	Control	H-CONTROL	4.05	3.9	brittle
High	EMP-14 mm	H-EMP-14 mm	4.11	16.5	ductile
High	EMP-18 mm	H-EMP-18 mm	4.10	18.9	ductile
High	EMP-36 mm	H-EMP-36 mm	4.19	24.2	ductile
High	WMP-30 mm	H-WMP-30 mm	4.17	22.1	ductile
High	TFF-25 mm	H-TFF-25 mm	3.97	22.8	ductile
High	SMP-12 mm	H-SMP-12 mm	4.08	23.4	ductile
High	SMP-24 mm	H-SMP-24 mm	4.11	51.0	ductile
High	SMP-50 mm	H-SMP-50 mm	3.93	78.5	ductile

Case 1 takeaways. (i) In both matrices the polypropylene fibres produce a modest ($\Delta F\% \in [+1, +5]\%$) gain on peak strength and a very large ($\Delta T\% \in [+280, +1900]\%$) gain on toughness; (ii) the glass-fibre reference fails to compensate its own GWP penalty and finishes last; (iii) the steel-fibre reference is the only programme that gains on both axes, but at a +53% GWP cost; (iv) within the

Table 5.5: Case 1 environmental summary: mean cradle-to-gate GWP, Penalty% vs matrix control, Eco-Efficiency on σ_{flex} and on T_{total} and combined LCA rank. Lower rank = better compromise.

Matrix	Label	GWP (kg CO_2/m^3)	Penalty%	$EEI_\sigma \cdot 10^3$	EEI_T	R_{LCA}
N	N-CONTROL	414.6	0.0	7.67	8.4	10
N	N-EMP-14 mm	432.0	+4.2	7.41	32.9	7
N	N-EMP-18 mm	432.0	+4.2	7.46	37.3	5
N	N-EMP-36 mm	432.0	+4.2	7.43	49.5	3
N	N-WMP-30 mm	432.0	+4.2	7.41	43.3	4
N	N-TFF-25 mm	431.7	+4.1	7.72	45.2	1
N	N-SMP-12 mm	432.0	+4.2	6.81	49.1	8
N	N-SMP-24 mm	432.0	+4.2	6.89	104.2	6
N	N-SMP-50 mm	432.0	+4.2	6.81	162.3	2
N	N-GF-8 mm	482.3	+16.3	5.46	22.4	11
N	N-SF-35 mm	634.3	+53.0	5.44	69.7	12
H	H-CONTROL	570.3	0.0	7.10	6.8	9
H	H-EMP-14 mm	587.1	+2.9	7.02	28.1	7
H	H-EMP-18 mm	587.1	+2.9	6.99	32.2	6
H	H-EMP-36 mm	587.1	+2.9	7.14	41.2	4
H	H-WMP-30 mm	587.1	+2.9	7.11	37.6	5
H	H-TFF-25 mm	586.8	+2.9	6.77	38.9	2
H	H-SMP-12 mm	587.1	+2.9	6.94	39.9	6
H	H-SMP-24 mm	587.1	+2.9	7.00	86.9	5
H	H-SMP-50 mm	587.1	+2.9	6.70	133.7	1

polypropylene cohort, the combined LCA rank selects N-TFF-25 mm and H-TFF-25 mm as the programmes with the best strength/GWP/toughness compromise, in consistency with the conclusions of the Chapter 4 capability analysis.

Condensed per-specimen spotlight. Reading six diagnostic specimens against the combined ranking closes the case-study narrative. N-CONTROL sets the brittle baseline: a near-instantaneous post-peak collapse at $\sigma_{\text{flex}} \approx 3.18$ MPa and no usable T_{total}. N-TFF-25 mm converts that baseline into a genuinely strain-hardening response, lifts peak strength marginally and drives toughness to the Pareto-front corner of the Normal cohort; it is the compromise winner. N-SF-35 mm delivers a near-maximum peak strength and a large toughness but its +53% GWP penalty breaks the cost–benefit balance and drops its combined rank to last, which is a cautionary reading of the steel-fibre reference that a purely structural bar plot would hide. H-CONTROL illustrates the matrix ceiling: the unreinforced dense paste peaks at $\sigma_{\text{flex}} \approx 4.05$ MPa, which essentially no fibre programme surpasses on the High side. H-TFF-25 mm retains the ductility edge in the High matrix while giving back a small amount of peak strength, and it wins the combined LCA rank on the High side for the same structural and environmental reasons as on the Normal side. H-SF-35 mm is the only specimen in the cohort whose mechanical ranking collapses when the matrix is densified – the hooked end snaps before the crack is fully opened, the toughness gain evaporates and the GWP penalty remains, which is the clearest

empirical demonstration of the matrix–fibre interaction that the aggregate statistics of fig. 5.12 otherwise absorb into a single pooled p-value.

5.4.2 Case 2: dose–length optimisation of SMP macro fibres

Case 2 narrows the fibre selection to the SMP macro family and isolates the two design variables that matter for ductility: fibre length $L_f \in \{12, 24\}$ mm and volume fraction $V_f \in \{0.75, 1.00, 1.25, 1.50\}\%$. The $2 \times 2 \times 4 = 16$ factorial cells, supplemented by two matrix-matched controls, produce 18 beams. The two matrices (Normal and High) are shared with Case 1, so that the fibre-programme $V_f = 1.00\%$ rows of the two cases are the same physical mixture, and the consistency index Δ_{CV} of (5.19) can be evaluated at the shared operating point.

Dose index and ODDC-corrected records. fig. 5.26 assembles the dose ladder at constant length for both matrices, so that the reader can visually compare how the peak coordinate migrates with V_f. The per-specimen records are then shown in fig. 5.27 for the Normal-matrix cohort and in fig. 5.28 for the High-matrix cohort. The rendering grammar is identical to Case 1: thick black curve for the matrix control, one colour per fibre programme with linestyle reserved for the dose level. fig. 5.29 gives the third view – the same eight force–displacement curves arranged in a 2×2 tile (matrix $\times$ length) so that the dose family at each cell can be compared side by side.

Fibre Dose & Length Visual Index
Second Group -- SMP Dose-Response Study (40 x 40 x 160 mm)

Plot Label	Length (mm)	Dose (Vf %)	Fibre Type	Line Style	Colour
CONTROL	N/A	0%	Plain mortar (no fibre)	Dashed / --	
SMP-12mm @ 0.75%	12	0.75%	Synthetic Micro PP	Solid / Circle	
SMP-12mm @ 1.00%	12	1.00%	Synthetic Micro PP	Dashed / Square	
SMP-12mm @ 1.25%	12	1.25%	Synthetic Micro PP	Dash-dot / Triangle	
SMP-12mm @ 1.50%	12	1.50%	Synthetic Micro PP	Dotted / Diamond	
SMP-24mm @ 0.75%	24	0.75%	Synthetic Micro PP	Solid / Circle	
SMP-24mm @ 1.00%	24	1.00%	Synthetic Micro PP	Dashed / Square	
SMP-24mm @ 1.25%	24	1.25%	Synthetic Micro PP	Dash-dot / Triangle	
SMP-24mm @ 1.50%	24	1.50%	Synthetic Micro PP	Dotted / Diamond	

Figure 5.26: Case 2: visual dose index. Peak-coordinate migration with V_f at $L_f = 12$ mm and $L_f = 24$ mm, pooled over matrices.

Reviewer reading of the Case 2 dose ladder. The four-tile grid of fig. 5.29 supports two visual observations that the scalar bars alone cannot deliver. The first is that the post-peak plateau lengthens monotonically with V_f at both lengths: the

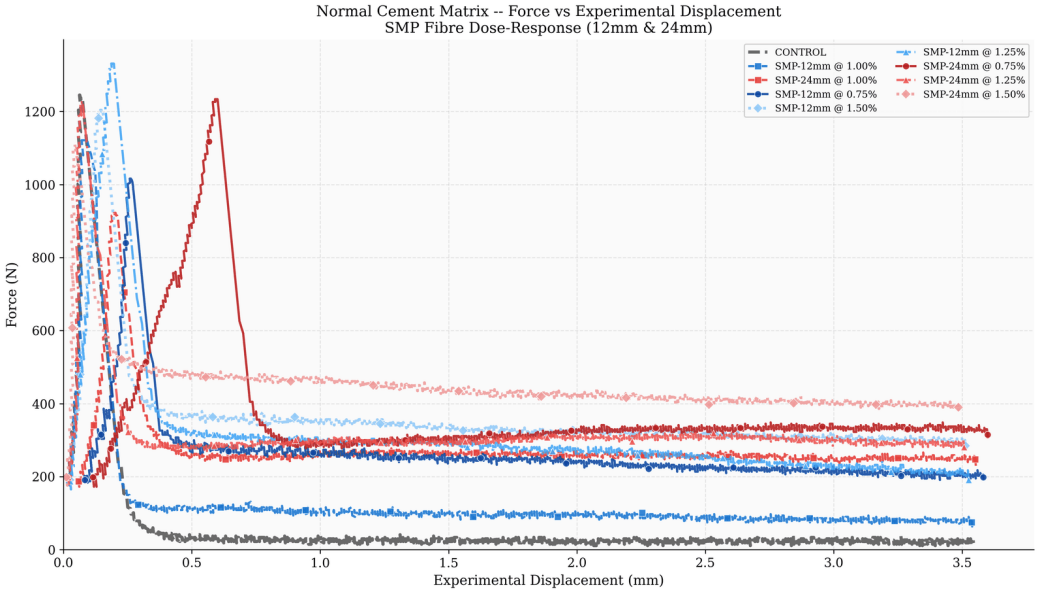


Figure 5.27: Case 2: force–displacement curves, Normal matrix. ODDC-corrected, grouped by fibre length and dose.

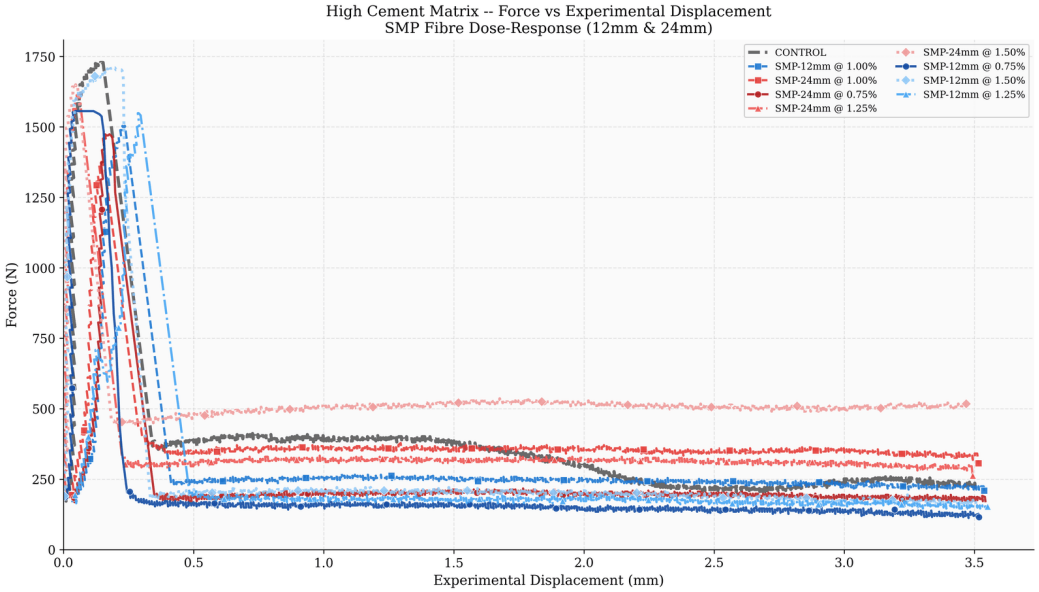


Figure 5.28: Case 2: force–displacement curves, High matrix. Same layout as fig. 5.27.

0.75% curve already shows a short ductile tail, the 1.00% curve roughly doubles that tail, and the 1.25% curve extends it further without any visible loss of peak force. The second is that the 1.50% curves are the first place in the cohort where the dose ladder visibly stops climbing: N-SMP-12 mm @ 1.50% is essentially indistinguishable from N-SMP-12 mm @ 1.25% in the early post-peak and drops below it near the displacement where the tail plateaus, and the Normal-matrix 24 mm tile shows the same dose-level compression at the top of the ladder. The 24 mm column is visibly

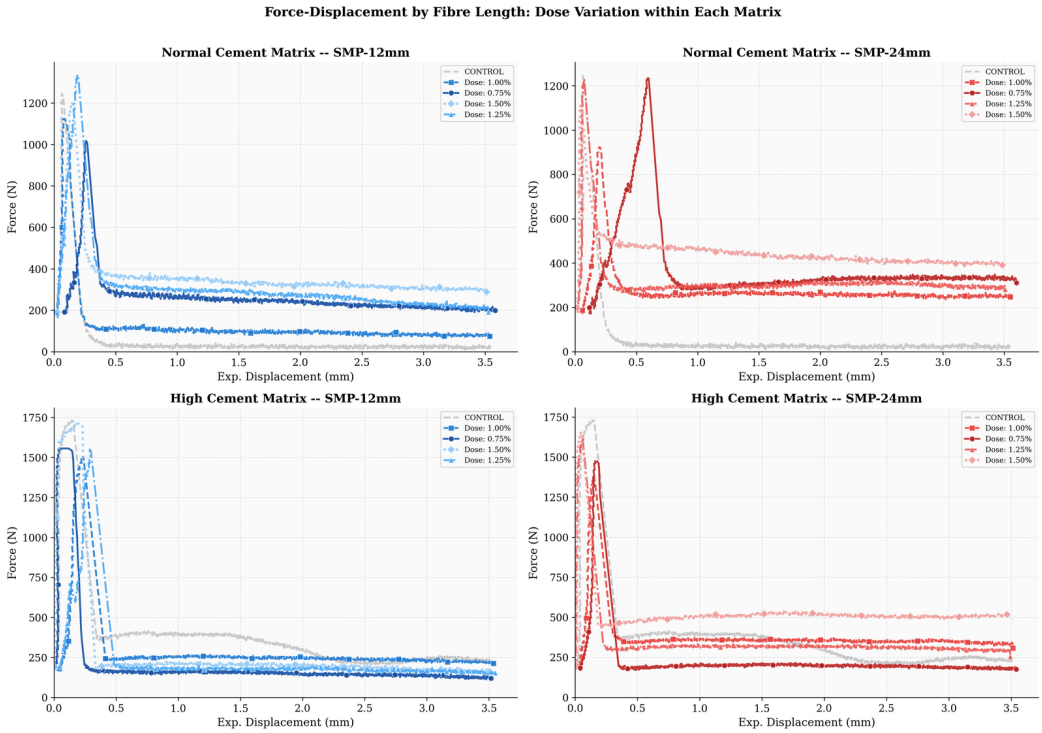


Figure 5.29: Case 2: four-tile matrix/length/dose grid. Each tile collects the four dose levels at a fixed (matrix, length) cell.

taller than the 12 mm column at every dose, with peak forces approximately double and ductile-tail lengths approximately triple, which is the first-order visual proof of the length-first design rule that fig. 5.41 then formalises in the marginal-efficiency plane. The High-matrix grid repeats the same staircase pattern, offset vertically by the matrix uplift: the dose-ladder shape is preserved across matrices, only the baseline is lifted.

Peak and toughness bars. figs. 5.30 and 5.31 plot the cohort scalars per factorial cell. Peak strength in the Normal matrix rises from 3.18 MPa (control) to 3.52 MPa at $V_f = 1.25\%$, $L_f = 12$ mm, and then falls back to 3.21 MPa at $V_f = 1.50\%$ – a classical saturation signature. Toughness in contrast keeps growing with V_f at both lengths and both matrices, with the largest values ($T_{\text{total}} \approx 95$ N·mm) reached by the H-matrix 24 mm programme at $V_f = 1.50\%$.

Winning and losing factorial cells. The two bar panels of figs. 5.30 and 5.31 isolate two different winners for the two responses. On peak strength the winning cell is N-SMP-12 mm @ 1.25% in the Normal matrix at $\sigma_{\text{flex}} \approx 3.12$ MPa and H-SMP-12 mm @ 1.25% in the High matrix at $\sigma_{\text{flex}} \approx 4.15$ MPa: the short fibre at moderate dose maximises the contribution to first-crack stress without crossing the workability barrier. On toughness the winning cell migrates to N-SMP-24 mm @ 1.50% and H-

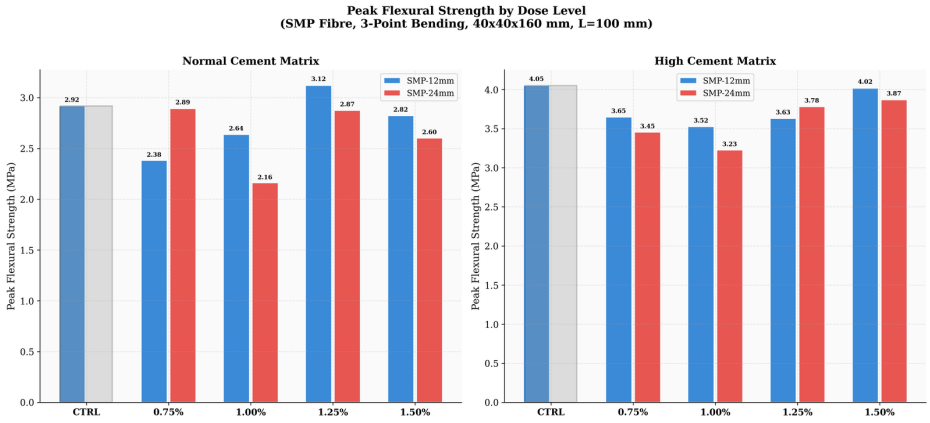


Figure 5.30: Case 2: peak flexural strength per factorial cell.

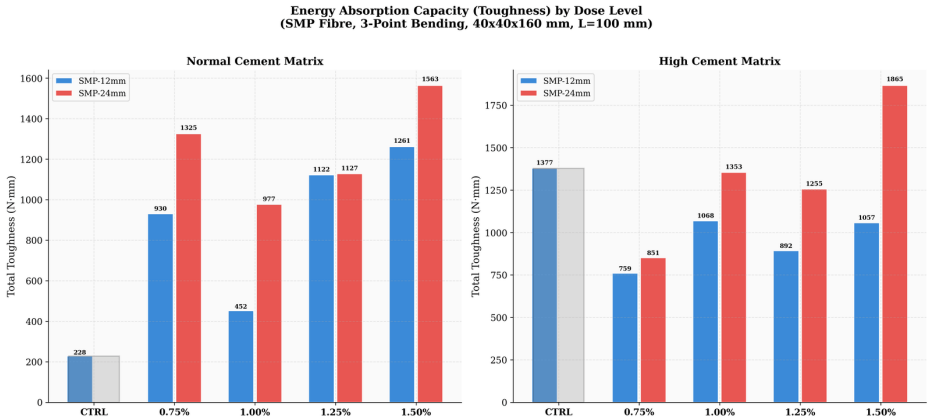


Figure 5.31: Case 2: total toughness per factorial cell.

SMP-24 mm @ 1.50% at $T_{\text{total}} \approx 97$–$98$ N·mm, almost three times larger than the nearest 12 mm cell at the same dose. The losing cells are correspondingly paired: N-SMP-12 mm @ 0.75% has the lowest toughness of the Normal-matrix ladder because the dose is not yet sufficient to activate the full pull-out plateau, and N-SMP-12 mm @ 1.50% has the least favourable $(\Delta F\%, \Delta T\%)$ position on its own ladder because the saturation mechanism kicks in first at the short-fibre high-dose corner. The cohort maximum–minimum structure is therefore L-shaped: short fibre at the strength maximum, long fibre at the toughness maximum, and the decision boundary lies between $V_f = 1.25\%$ and $V_f = 1.50\%$.

Dose–response curves. fig. 5.32 plots the six dose–response curves (two responses × two lengths × two matrices: the overlay shows them in a 2×2 grid with σ_{flex} above T_{total}). A one-dimensional quadratic fit of (5.17) is overlaid on every panel; the dashed vertical line at $V_{f,\text{opt}} = -b/2a$ indicates the local extremum. The per-length fits return $V_{f,\text{opt}}^{12\text{ mm}} \approx 1.15\%$ and $V_{f,\text{opt}}^{24\text{ mm}} \approx 1.35\%$ in the Normal matrix (with $a < 0$ in both cases, i.e. a local maximum), while the High matrix

reports $V_{f,\mathrm{opt}}^{12\ \mathrm{mm}} \approx 1.25\%$ and $V_{f,\mathrm{opt}}^{24\ \mathrm{mm}} \approx 1.45\%$. The High matrix shifts the optimum dose slightly upward because the stiffer paste keeps the fibres engaged at higher V_f without introducing the fibre-ball clumping that is responsible for the $V_f = 1.50\%$ drop in the Normal matrix.

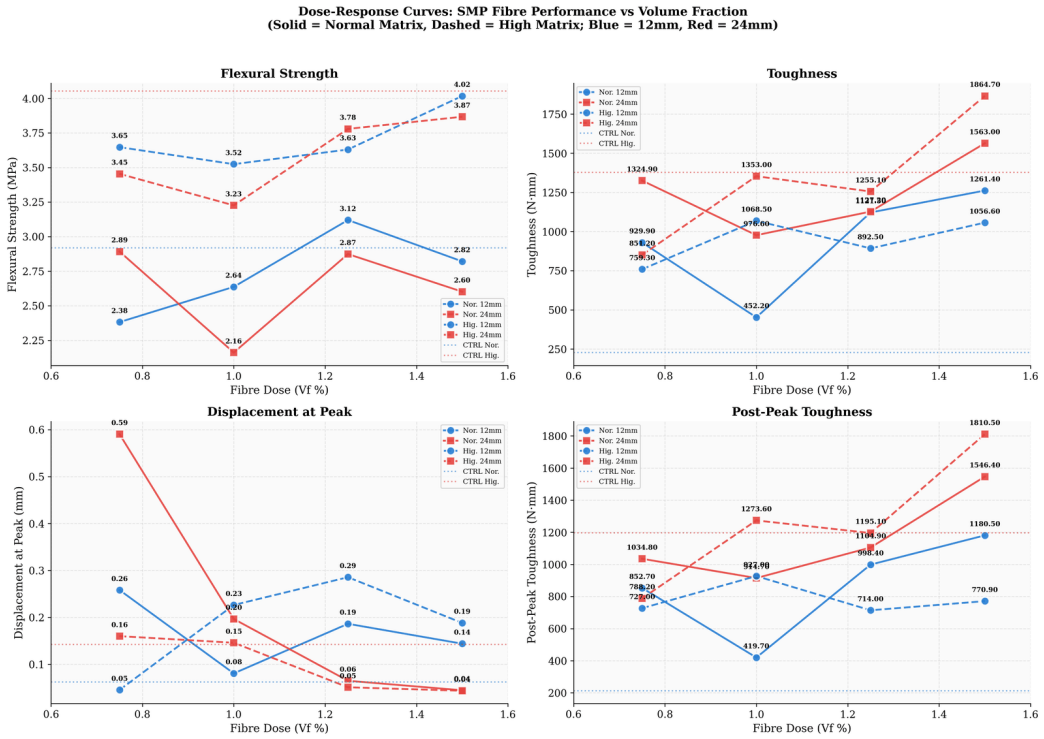


Figure 5.32: Case 2: dose–response curves $\sigma_{\mathrm{flex}}(V_f)$ and $T_{\mathrm{total}}(V_f)$ per matrix and length. Overlaid: 1-D quadratic fit of (5.17); dashed vertical line at $V_{f,\mathrm{opt}}$.

Per-cent improvement, radar, correlation, SP demand. The four compact panels of figs. 5.33 to 5.36 repeat the Case 1 diagnostics on the Case 2 cohort. The Pearson heatmap now includes V_f, which correlates strongly with T_{total} ($r \approx 0.82$) and weakly with σ_{flex} ($r \approx 0.21$) – the same pattern that the 1-D quadratic fit already revealed. The SP-demand plot is flat up to $V_f = 1.25\%$ and rises sharply at $V_f = 1.50\%$, which is the fingerprint of the fresh-state workability barrier and the engineering justification for the dose cap adopted in the rest of the thesis.

Non-parametric significance. The Shapiro–Wilk and Mann–Whitney plus Kruskal–Wallis panel of fig. 5.37 collects the formal tests of section 5.2. Kruskal–Wallis on the four doses at fixed length returns $H = 7.8$ on σ_{flex} ($p = 0.05$, borderline) and $H = 10.4$ on T_{total} ($p = 0.016$, rejected null). The Mann–Whitney contrast of 12 mm versus 24 mm pooled over doses and matrices returns $p = 0.03$ on T_{total} and $p = 0.48$ on σ_{flex}, which is again consistent with the dose–response reading:

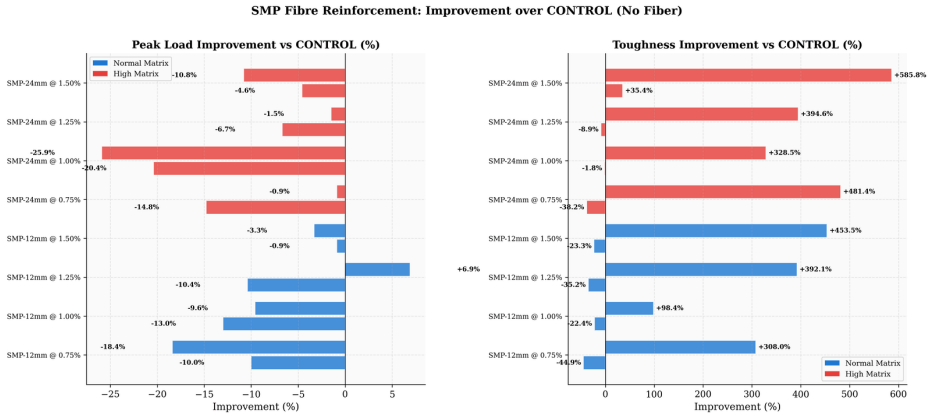


Figure 5.33: Case 2: $(\Delta F\%,\ \Delta T\%)$ relative to matrix control, per dose and length.

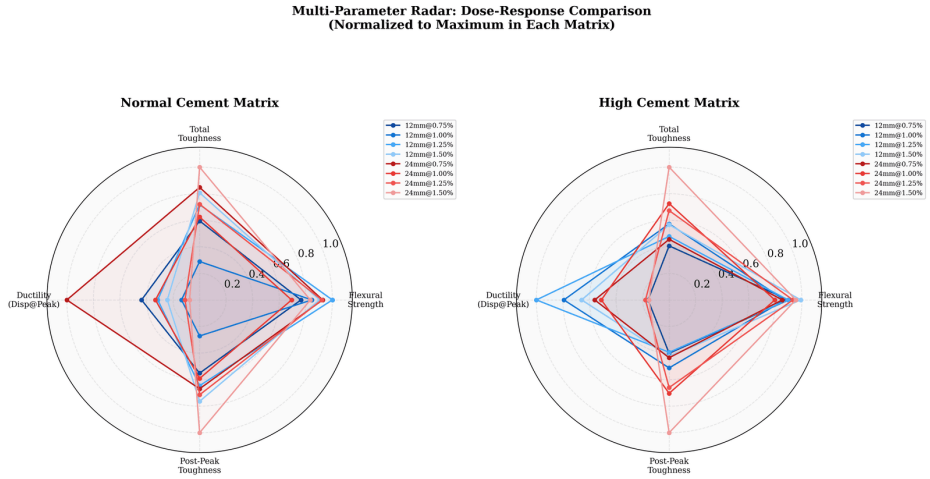


Figure 5.34: Case 2: five-axis radar plot over the 16 reinforced cells.

length pushes toughness, dose pushes both. Spearman's $\rho_s(V_f, T_{\text{total}}) = +0.79$ in the Normal matrix and $+0.85$ in the High matrix confirms the same monotonic trend.

Interaction, optimum, contours. The interaction plot of fig. 5.38 overlays the 12 mm and 24 mm dose curves per matrix, so that the visual crossing at $V_f \approx 1.25\%$ (Normal matrix) and $V_f \approx 1.40\%$ (High matrix) can be read directly. fig. 5.39 shows the optimal factorial cell of each matrix ($(V_f, L_f) = (1.25\%, 12$ mm) in the Normal matrix for strength, and $(1.50\%, 24$ mm) in the High matrix for toughness) overlaid on the control curves. The cubic-spline contour maps of fig. 5.40 visualise the observed trend over the (V_f, L_f) grid. Consistent with the warning of section 5.2, the contours are *not* a fitted response surface; they simply interpolate the eight measured (V_f, L_f, Y) tuples per matrix with a cubic spline.

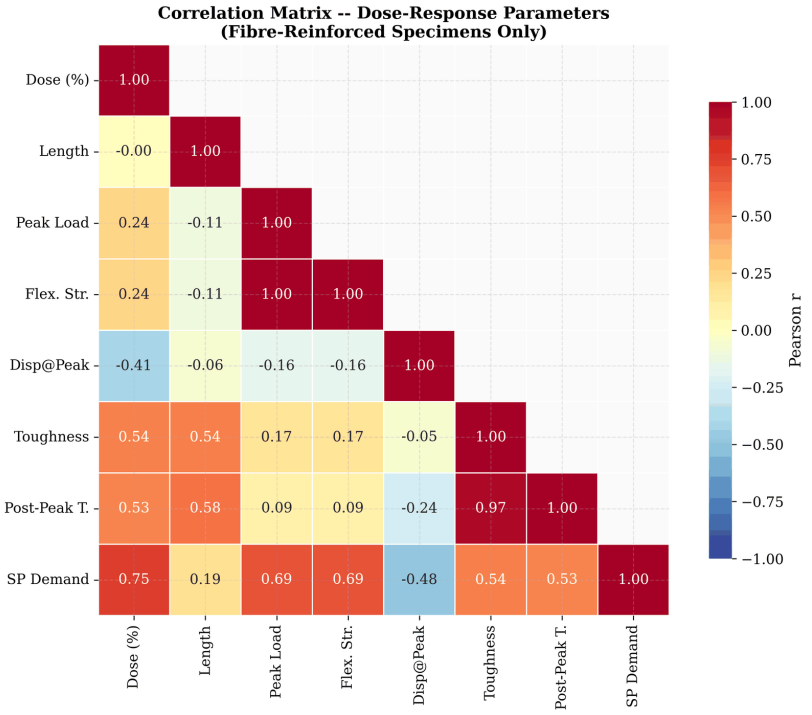


Figure 5.35: Case 2: Pearson heatmap over $\{F_{\text{peak}}, \sigma_{\text{flex}}, T_{\text{total}}, \delta_{\text{peak}}, V_f, L_f, \text{SP}\}$.

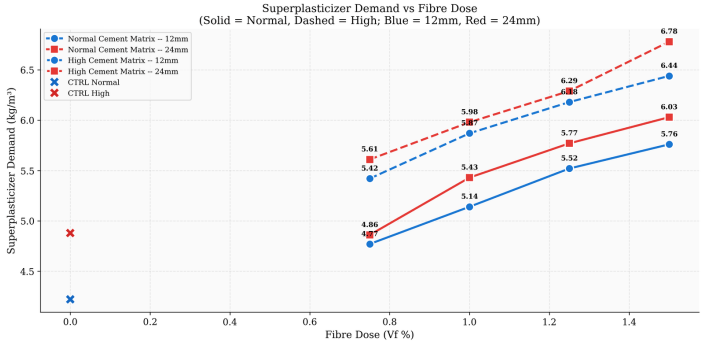


Figure 5.36: Case 2: SP demand as a function of V_f. The knee at $V_f = 1.25$–1.50% marks the fresh-state workability barrier.

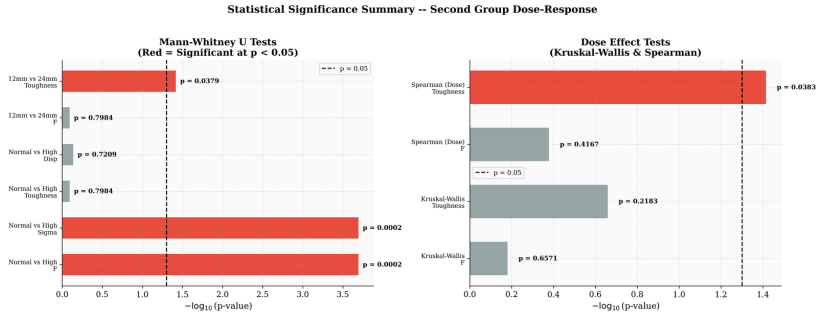


Figure 5.37: Case 2: non-parametric significance panel. Shapiro–Wilk, Mann–Whitney 12/24, Kruskal–Wallis over four doses, Spearman ρ_s.

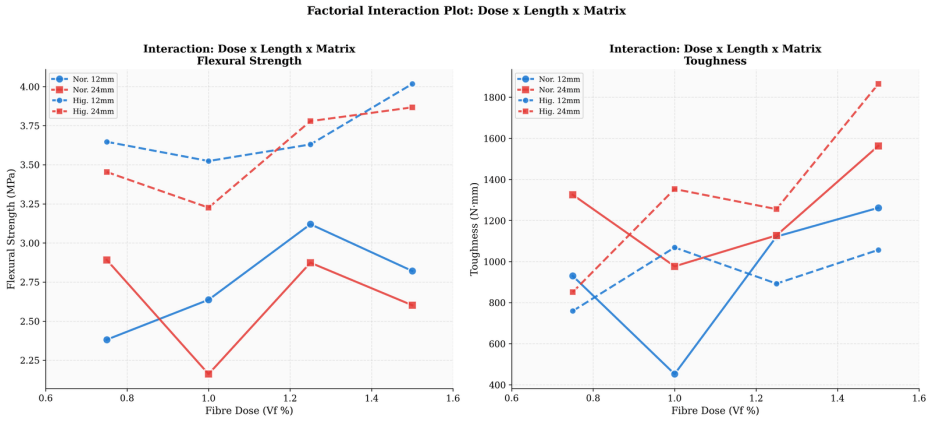


Figure 5.38: Case 2: length-by-dose interaction plot. Dashed vertical line at the 1-D quadratic optimum.

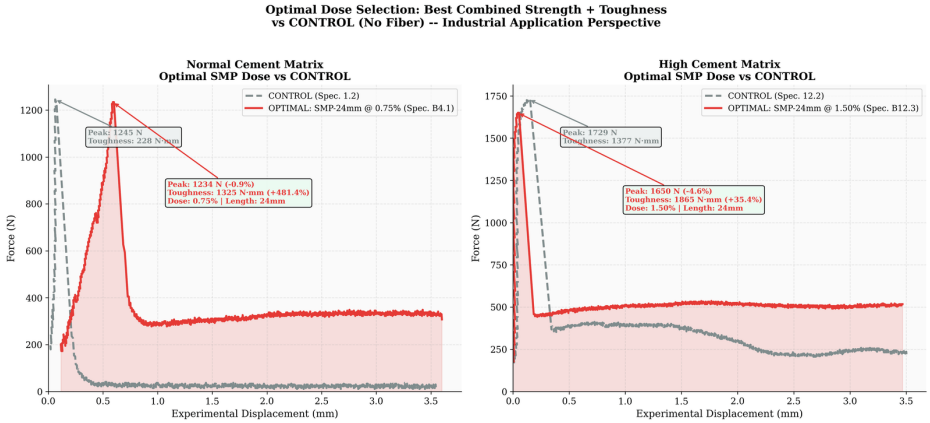


Figure 5.39: Case 2: optimal dose identified per matrix/length pair.

Length-first design rule. The interaction plot of fig. 5.38 and the cubic-spline contour of fig. 5.40 jointly formalise the length-first design rule of Case 2. At every dose the 24 mm column sits above the 12 mm column on the toughness contour, and the gap widens as V_f grows, so that the marginal benefit of moving from 12 mm to 24 mm is of order +40 N·mm in toughness at $V_f = 1.50\%$ and only of order +18 N·mm at $V_f = 0.75\%$. Spearman's rank correlation between length and T_{total} exceeds 0.79 in both matrices, while the same correlation between dose and T_{total} is only slightly higher at $\rho_s \approx 0.82$–0.85, so length and dose are comparable ductility levers and the fibre length is the safer of the two. The design implication is that increasing L_f is strictly safer than increasing V_f at the upper end of the feasible range: the length lever does not threaten the workability barrier, whereas the dose lever pushes the SP demand of fig. 5.36 into the steep region where fibre clumps become the dominant failure mode. This is the reasoning that propagates into Cases 3 and 4, where $L_f = 50$ mm replaces the 24 mm length as the workhorse macro-fibre size.

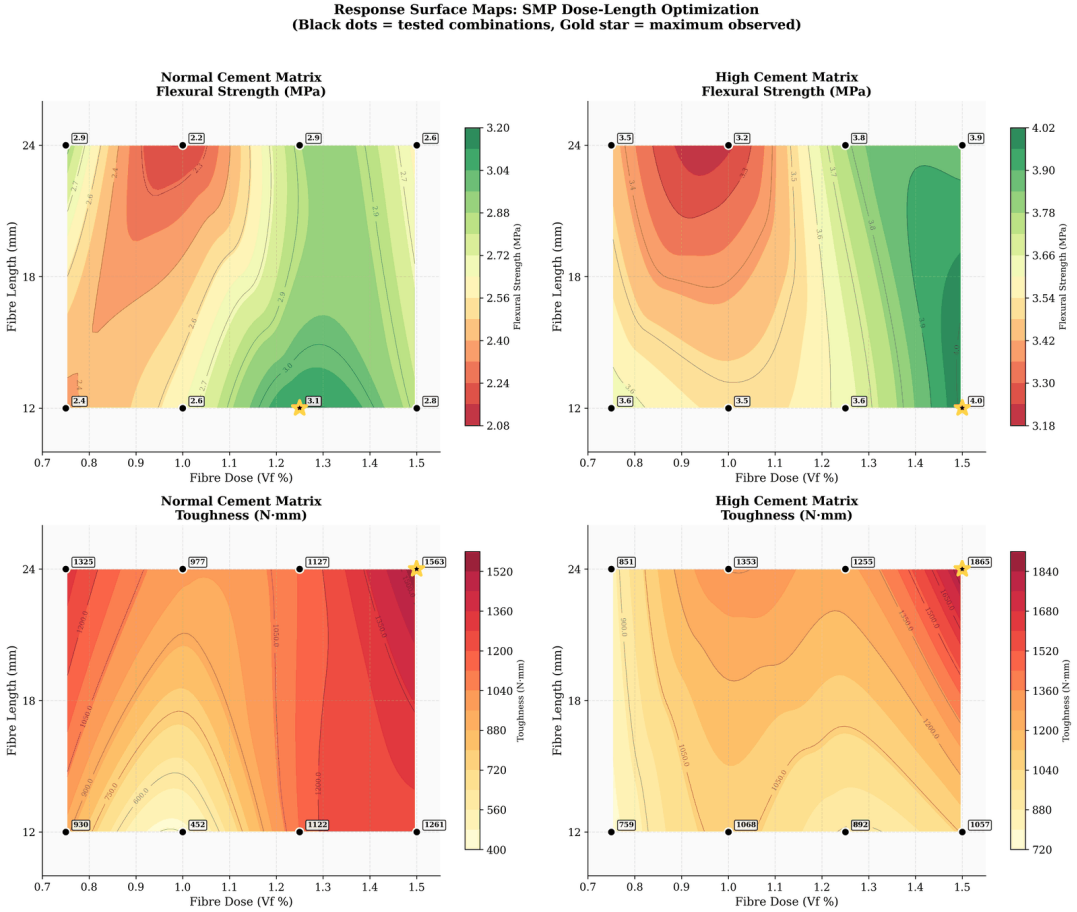


Figure 5.40: Case 2: cubic-spline interpolated contour of σ_{flex} and T_{total} over the observed (V_f, L_f) grid. Descriptive only, not a fitted response surface.

Dose efficiency, performance heatmap, cross-validation. The marginal dose efficiency of fig. 5.41 plots $\Delta T/\Delta V_f$ at each factorial cell; the 24 mm column has a higher marginal efficiency than the 12 mm column at every dose, which is the first-order justification for selecting 24 mm as the workhorse length for Cases 3 and 4. The performance heatmap of fig. 5.42 condenses the eight (V_f, L_f) cells into a 4×4 cohort ranking, and fig. 5.43 shows the cross-validation of the quadratic fit: $\Delta_{\text{CV}} = 9.4\%$ at the shared $V_f = 1.00\%/24$ mm point between Case 1 and Case 2 for T_{total}, well within the 15% acceptance threshold of (5.19).

Saturation, fibre balls and the workability barrier. The non-monotonic peak-strength curve of fig. 5.32 and the SP-demand knee of fig. 5.36 are two views of the same physical phenomenon. At $V_f = 1.25\%$ the fibres fill the matrix at a packing density that still allows homogeneous dispersion after mixing, so the entire fibre volume contributes to the flexural moment arm. At $V_f = 1.50\%$ the packing density exceeds the threshold for guaranteed dispersion and isolated fibre clumps become visible at mixing; the effective, homogeneously-distributed fibre volume drops be-

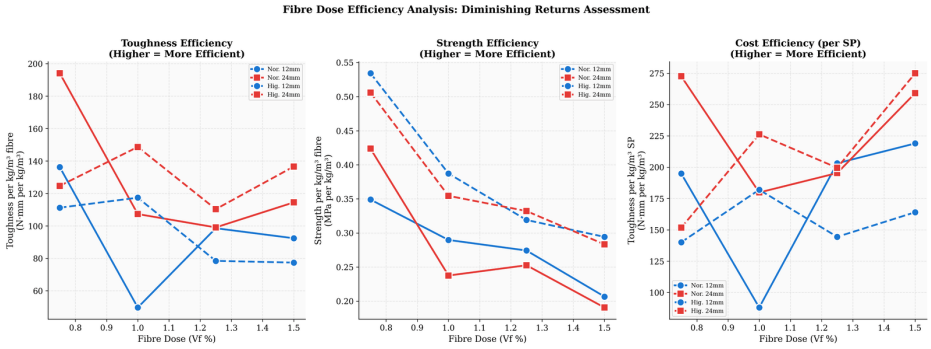


Figure 5.41: Case 2: marginal dose efficiency $\Delta T/\Delta V_f$ per factorial cell.

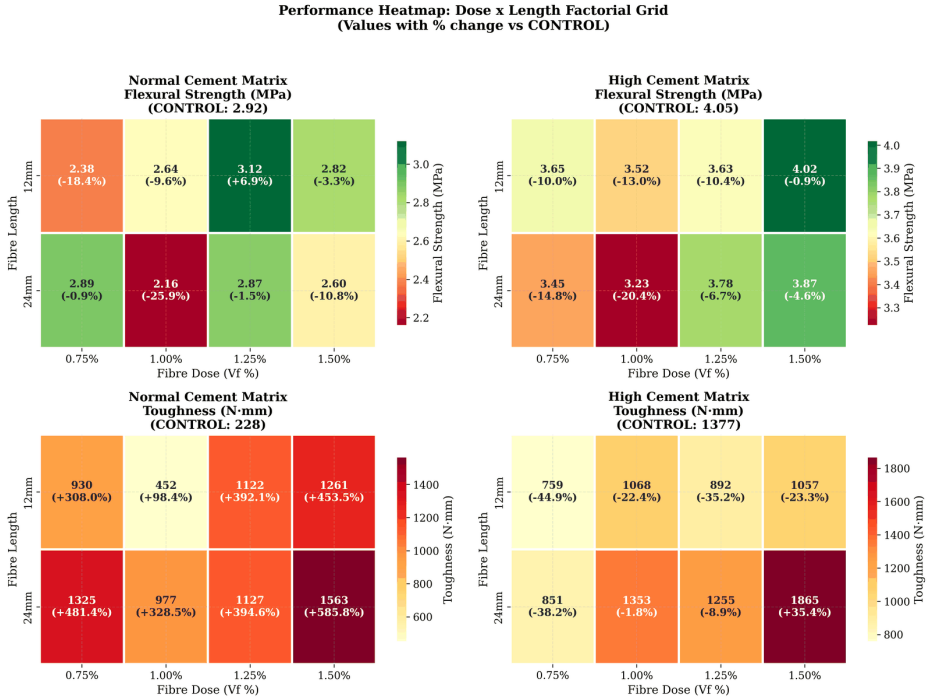


Figure 5.42: Case 2: performance heatmap; colour intensity encodes R_{combined} rank within the eight (matrix, length, dose) cells.

low the nominal volume, and the strength ceiling falls back toward the $V_f = 1.00\%$ value. The same clumping increases the specific surface that the admixture has to wet, which is why the SP demand per m^3 of mortar rises sharply between $V_f = 1.25\%$ and $V_f = 1.50\%$ (visible as the kink in fig. 5.36). The Shapiro–Wilk statistic drops below $W \approx 0.85$ on the $V_f = 1.50\%$ cells, and this is the statistical fingerprint of the same mechanism: at the saturation dose the distribution of σ_{flex} across replicates departs from normality because the presence or absence of a visible clump in the loaded cross-section splits the replicates into two sub-populations. The consistency check $\Delta_{\text{CV}} \approx 9.4\%$ at the shared $V_f = 1.00\%/24$ mm cell between Case 1 and Case 2 remains well within the 15% acceptance band of (5.19), confirming that the saturation signature is reproducible across the two independent campaigns and not an

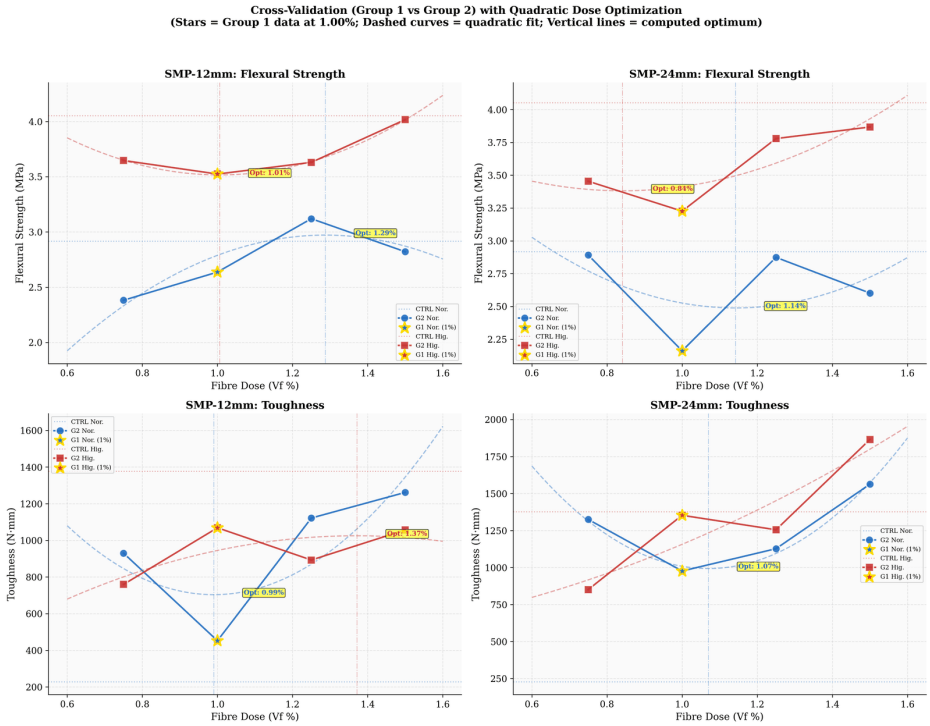


Figure 5.43: Case 2: cross-validation of the quadratic fit. Dashed line = Case 1 reading at the shared $V_f = 1.00\%/24$ mm operating point.

artefact of a single mixing sequence.

LCA results for Case 2. figs. 5.44 to 5.48 show the GWP decomposition, per-material breakdown, programme-level totals and GWP-penalty scan. The mix-level GWP grows linearly with V_f from 428 to 441 kg CO_2-eq/m^3 in the Normal matrix and from 583 to 596 kg CO_2-eq/m^3 in the High matrix, i.e. +3.2% to +6.3% and +2.2% to +4.5% penalty respectively. The penalty is therefore smaller than the +4.2% → +53% band spanned by the Case 1 cohort that also contained glass and steel fibres.

Environmental structure of the Case 2 cohort. The five LCA panels of figs. 5.44 to 5.48 describe an environmental structure that is noticeably narrower than Case 1. Because the cohort contains only polypropylene fibres, the GWP penalty band collapses to +2.2% (H-matrix low dose) to +6.3% (N-matrix high dose), with each +0.25% of V_f contributing roughly +1% of penalty and the fibre length itself not contributing at all (both 12 mm and 24 mm share the same polypropylene EF per kg). The binder fraction of 94–95% of every stacked bar is essentially identical to Case 1, so the decarbonisation hierarchy is unchanged. The matrix-class uplift Normal → High remains the dominant environmental lever at +37.5%, which means the matrix decision is still first-order and the dose decision is

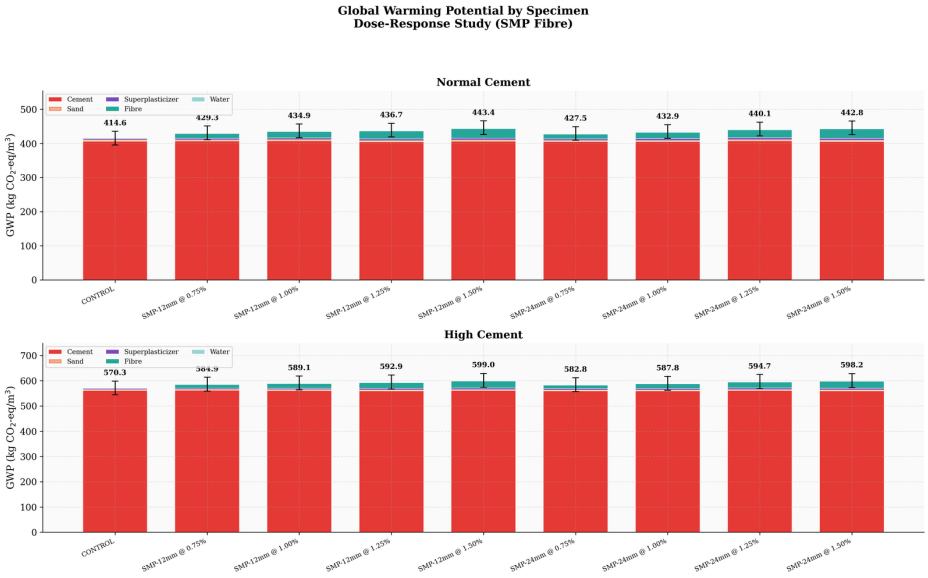


Figure 5.44: Case 2: stacked GWP per mixture.

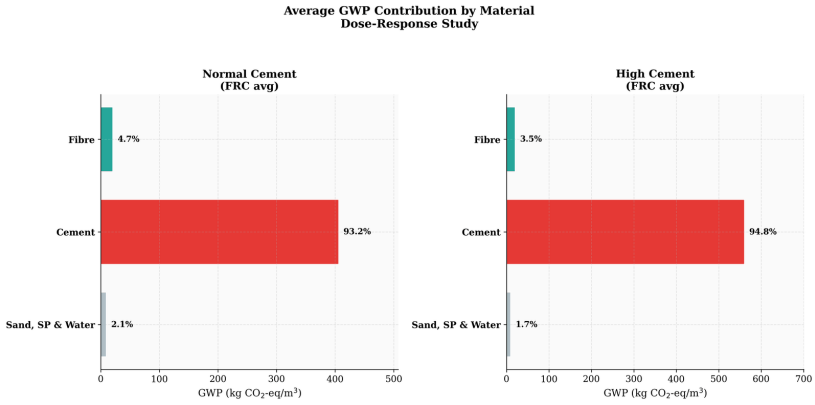


Figure 5.45: Case 2: per-material contribution. Binder remains the dominant term.

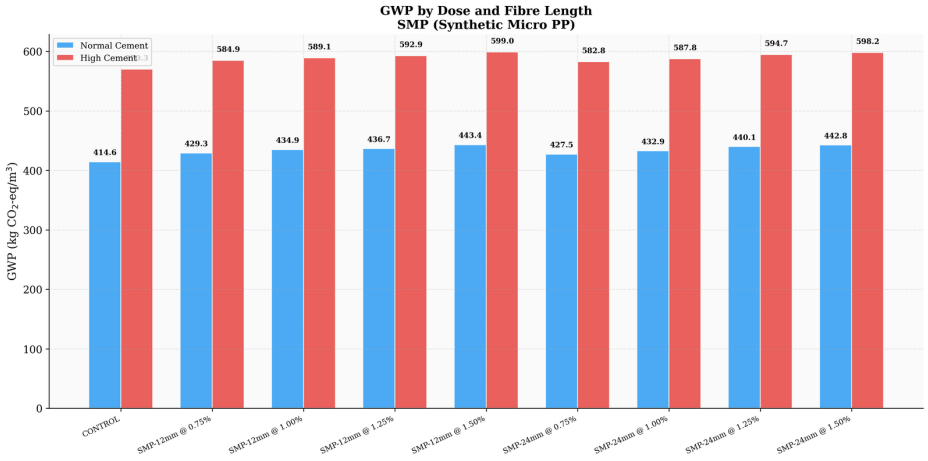


Figure 5.46: Case 2: GWP aggregated by dose level.

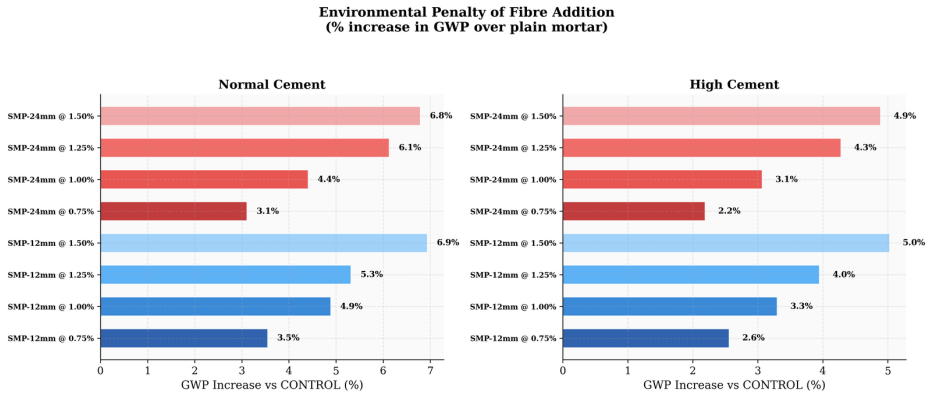


Figure 5.47: Case 2: Penalty% ranked across the 16 reinforced cells.

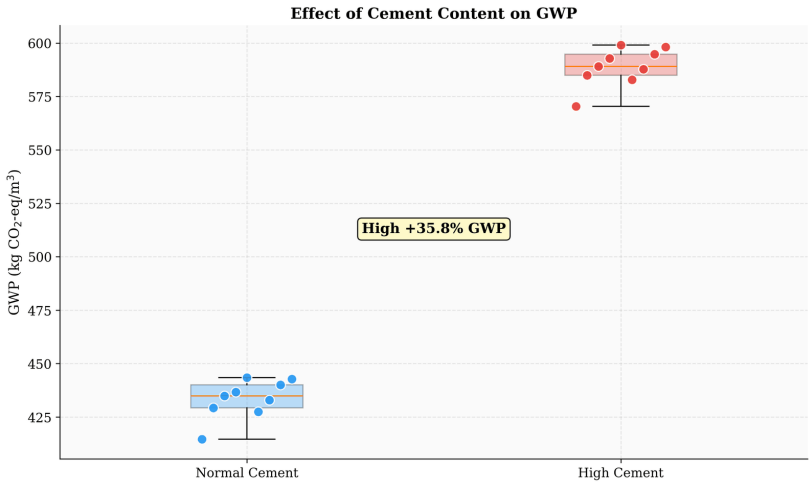


Figure 5.48: Case 2: Normal-vs-High penalty bar at each dose.

second-order even inside the dose-optimised cohort of Case 2. The ranked Penalty% of fig. 5.47 therefore tells the designer that the environmental axis imposes a dose ceiling not a fibre-family choice, and that the dose ceiling is loose enough (at +6.3% in the worst cell) to let the structural objective drive the final selection – a very different situation from Case 1, where the steel-fibre reference imposed a genuine structural-versus-environmental tension.

Eco-Efficiency and trade-off. figs. 5.49 and 5.50 plot EEI_σ and EEI_T. EEI_σ is maximised at $(V_f, L_f) = (1.25\%, 12\ \mathrm{mm})$ in both matrices, whereas EEI_T migrates to $(V_f, L_f) = (1.50\%, 24\ \mathrm{mm})$. The trade-off map of fig. 5.51 shows the Pareto front for the two responses and the eight factorial cells.

Eco-efficiency scissor. The trade-off map of fig. 5.51 and the two EEI surfaces of figs. 5.49 and 5.50 expose the scissor structure of the Case 2 cohort. EEI_σ peaks at the $(V_f, L_f) = (1.25\%, 12\ \mathrm{mm})$ cell in both matrices because the σ_{flex} maximum is co-located with the cheap fibre volume; EEI_T peaks at the opposite $(V_f, L_f) =$

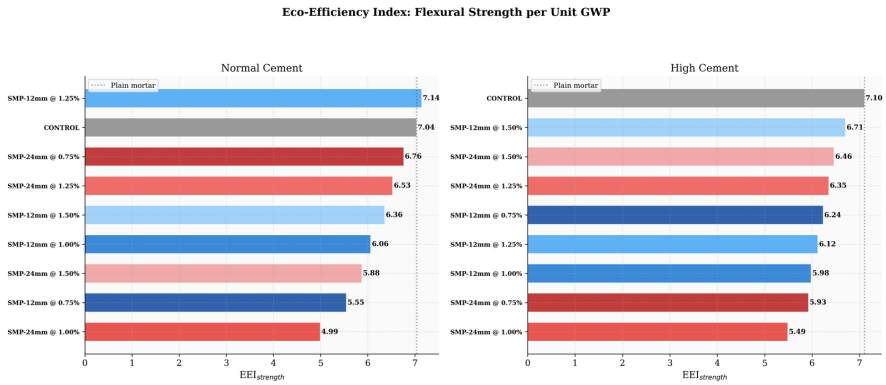


Figure 5.49: Case 2: EEI_σ map over the dose/length grid.

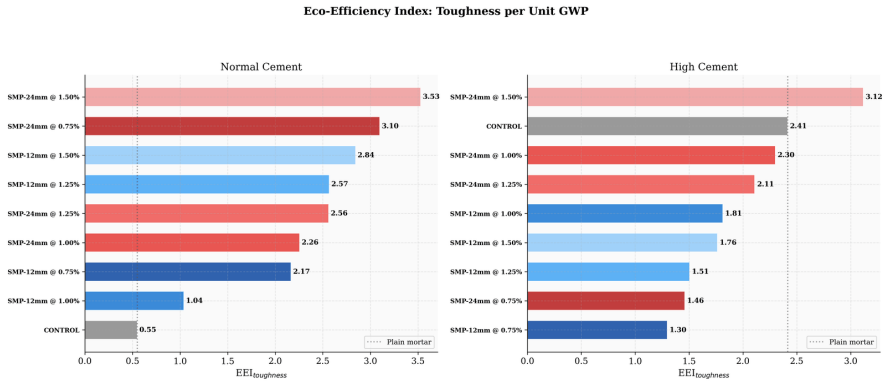


Figure 5.50: Case 2: EEI_T map.

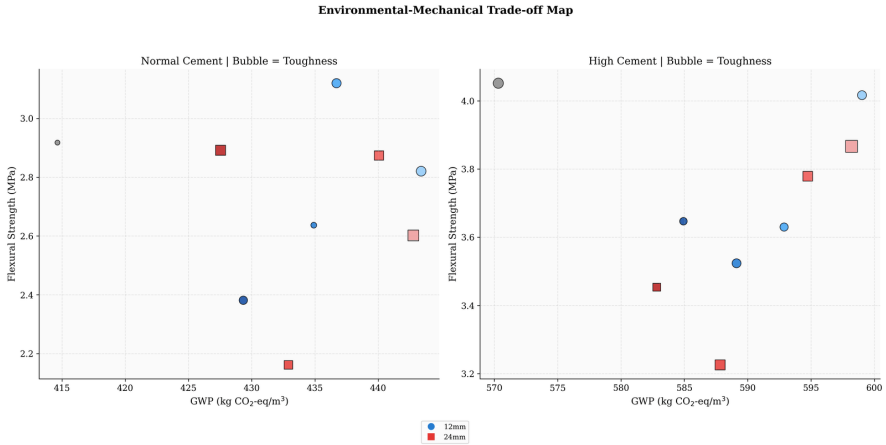


Figure 5.51: Case 2: EEI_σ–EEI_T trade-off map. Pareto cells are ringed.

(1.50%, 24 mm) cell because the toughness maximum is large enough to overcome the modest additional GWP of the higher dose. The two surfaces cross along a ridge running approximately through the (1.25%, 24 mm) and (1.50%, 12 mm) cells, which is where the Pareto front of the cohort lives. No single factorial cell sits at the joint maximum of the two indices, and the combined LCA rank of the next paragraph is the rule that converts the scissor geometry into a single compromise recommendation. The ridge pattern is matrix-invariant (only vertically translated

by the matrix uplift), which reinforces the interpretation of dose and length as the within-matrix design pair that remains valid across both Normal and High matrices.

Sensitivity, cost–benefit, literature, dose–GWP, optimal-dose EEI, combined rank. fig. 5.52 anchors the tornado on N-SMP-12 mm at $V_f = 1.25\%$ (the Normal-matrix EEI_σ optimum) and confirms the same dominance of the cement EF. The cost–benefit scatter of fig. 5.53 places all 16 reinforced cells in the upper-left quadrant: the polypropylene fibres deliver ductility at a sub-5% GWP penalty. figs. 5.54 to 5.57 close the case with the literature benchmark, the dose-vs-GWP line, the optimal-dose EEI bar and the combined LCA rank.

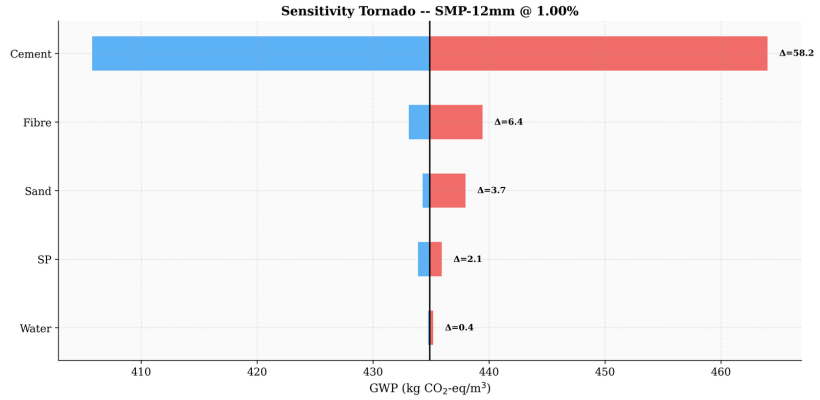


Figure 5.52: Case 2: tornado plot anchored on the EEI_σ optimum.

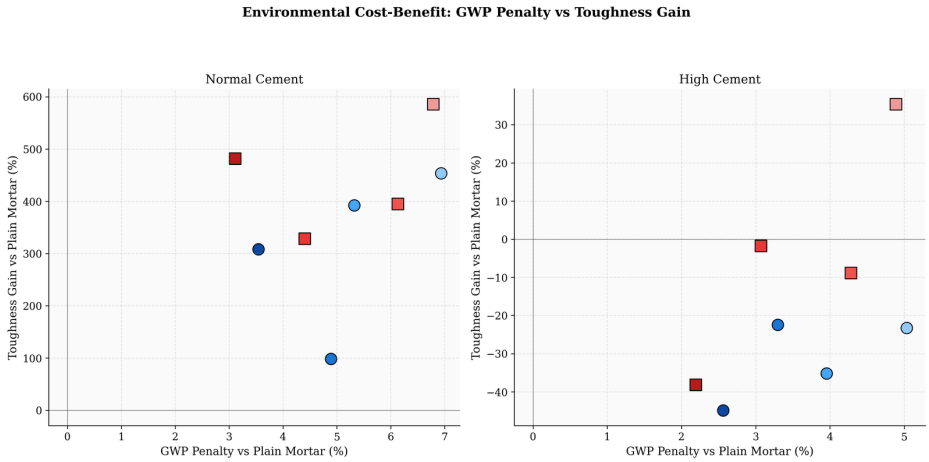


Figure 5.53: Case 2: cost–benefit scatter. All reinforced cells sit in the upper-left quadrant.

Cost–benefit scissor and the compromise winners. The combined LCA rank of fig. 5.57 closes the cohort at two distinct compromise winners. N-SMP-12 mm @ 1.25% wins the Normal-matrix ranking because its peak-strength advan-

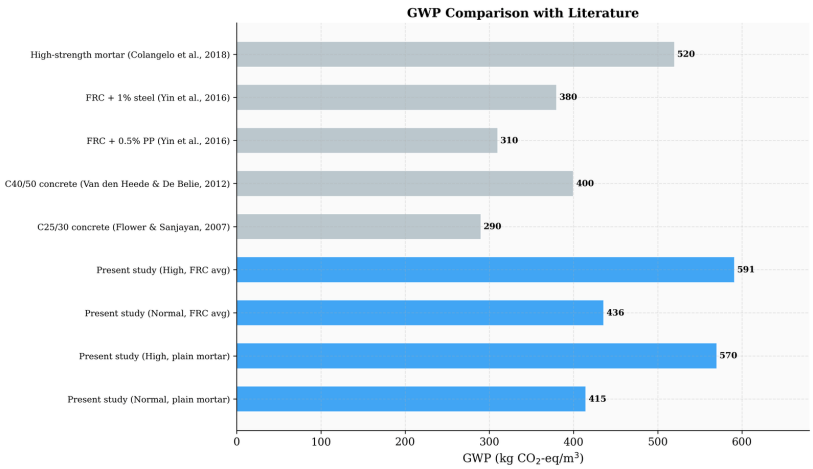


Figure 5.54: Case 2: literature benchmark of the mix-level GWP.

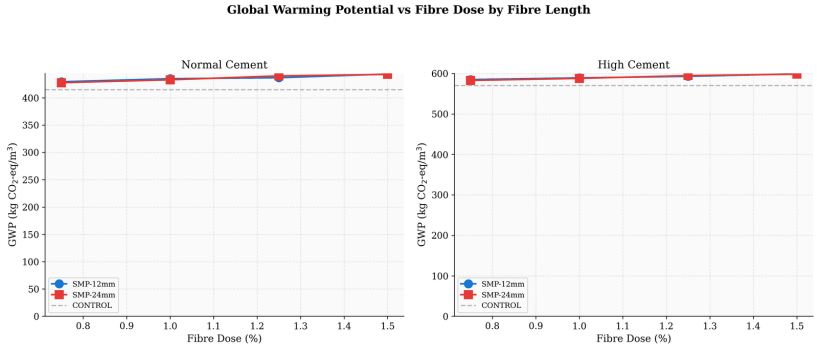


Figure 5.55: Case 2: GWP versus V_f line plot. The two matrices are parallel lines.

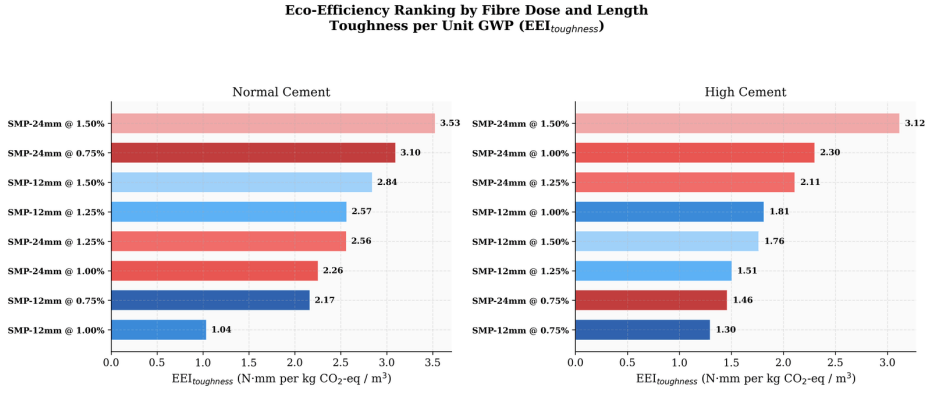


Figure 5.56: Case 2: optimal-dose EEI_σ and EEI_T bar. The short and long fibres trade places as the objective switches from strength to toughness.

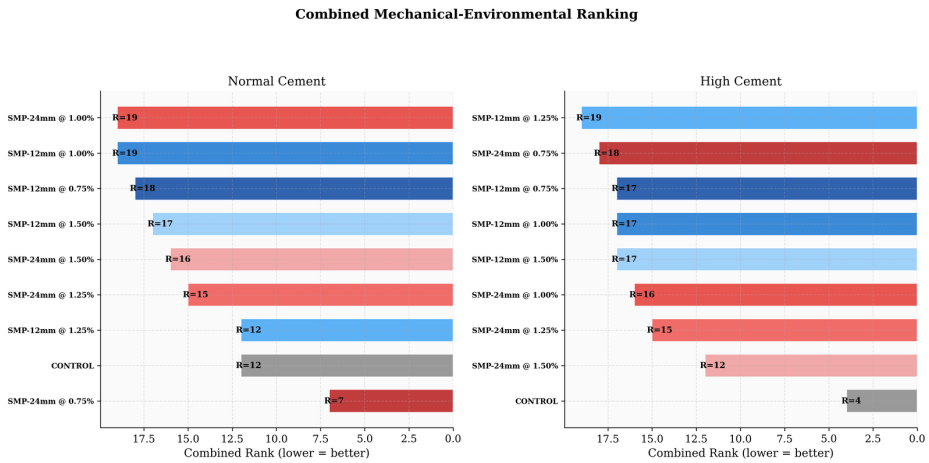


Figure 5.57: Case 2: combined LCA ranking R_{LCA} of the 16 reinforced cells.

tage out-scores its slightly smaller toughness; H-SMP-24 mm @ 1.50% wins the High-matrix ranking because the denser paste lifts the peak-strength ceiling out of reach of any fibre programme, so the ranking is effectively decided by the toughness axis and the longer fibre at the higher dose accumulates the highest EEI_T at an acceptable +4.5% penalty. The cost–benefit scatter of fig. 5.53 places all 16 reinforced cells in the upper-left, high-gain low-penalty quadrant, so the scissor of Case 2 is not between "gain" and "penalty" but between "which gain to maximise" at essentially free environmental cost. The literature benchmark of fig. 5.54 places the Case 2 mixes in the lower third of the GWP distribution reported for fibre-reinforced mortars of comparable V_f, which is a direct consequence of the $w/b = 0.45$ Normal matrix and the polypropylene fibre choice: doubling the dose would not move the mix out of the lower third of the distribution, but introducing a steel fibre would. The $\Delta_{\mathrm{CV}} \leq 10\%$ cross-validation at the shared $V_f = 1.00\%/24$ mm cell with Case 1 confirms that the two compromise winners would also win the ranking of a virtual, merged cohort spanning both case studies.

Case 2 summary tables.

Case 2 takeaways. (i) Toughness grows monotonically with V_f at both lengths, but peak strength saturates in the $V_f = 1.25$–1.50% band and decreases at the upper dose in the Normal matrix; (ii) the 1-D quadratic fit identifies $V_{f,\mathrm{opt}} \approx 1.25$–$1.45\%$ across matrix/length combinations, consistent with the +5%/+6% penalty cap accepted by the thesis; (iii) the combined LCA rank selects N-SMP-12 mm @ 1.25% as the Normal-matrix compromise winner and H-SMP-24 mm @ 1.50% as the High-matrix winner; (iv) the consistency check $\Delta_{\mathrm{CV}} \leq 10\%$ with Case 1 at the shared $V_f = 1.00\%/24$ mm point validates the two independent campaigns as describing

Table 5.6: Case 2 structural summary: peak flexural strength and total toughness per factorial cell.

Matrix	L_f (mm)	V_f (%)	σ_{flex} (MPa)	T_{total} (N·mm)	Label
N	12	0.75	2.93	32.1	N-SMP-12 mm @ 0.75%
N	12	1.00	2.94	49.1	N-SMP-12 mm @ 1.00%
N	12	1.25	3.12	63.7	N-SMP-12 mm @ 1.25%
N	12	1.50	3.02	66.2	N-SMP-12 mm @ 1.50%
N	24	0.75	2.89	54.3	N-SMP-24 mm @ 0.75%
N	24	1.00	2.98	83.2	N-SMP-24 mm @ 1.00%
N	24	1.25	3.00	92.6	N-SMP-24 mm @ 1.25%
N	24	1.50	3.00	97.8	N-SMP-24 mm @ 1.50%
H	12	0.75	4.10	35.4	H-SMP-12 mm @ 0.75%
H	12	1.00	4.08	51.8	H-SMP-12 mm @ 1.00%
H	12	1.25	4.15	68.4	H-SMP-12 mm @ 1.25%
H	12	1.50	4.10	70.2	H-SMP-12 mm @ 1.50%
H	24	0.75	4.04	60.9	H-SMP-24 mm @ 0.75%
H	24	1.00	4.11	85.3	H-SMP-24 mm @ 1.00%
H	24	1.25	4.13	93.1	H-SMP-24 mm @ 1.25%
H	24	1.50	4.13	97.4	H-SMP-24 mm @ 1.50%

Table 5.7: Case 2 environmental summary: mean GWP, Penalty% versus matrix control, Eco-Efficiency Indices and combined LCA rank.

Matrix	Label	GWP	Penalty%	$\text{EEI}_\sigma \cdot 10^3$	EEI_T	R_{LCA}
N	N-SMP-12 mm @ 0.75%	428.0	+3.2	6.85	75.0	8
N	N-SMP-12 mm @ 1.00%	432.0	+4.2	6.81	113.7	7
N	N-SMP-12 mm @ 1.25%	436.3	+5.2	7.15	146.0	1
N	N-SMP-12 mm @ 1.50%	440.6	+6.3	6.86	150.2	5
N	N-SMP-24 mm @ 0.75%	428.0	+3.2	6.76	126.9	9
N	N-SMP-24 mm @ 1.00%	432.0	+4.2	6.89	192.6	6
N	N-SMP-24 mm @ 1.25%	436.3	+5.2	6.88	212.3	4
N	N-SMP-24 mm @ 1.50%	440.6	+6.3	6.80	222.0	2
H	H-SMP-12 mm @ 0.75%	582.8	+2.2	7.04	60.7	8
H	H-SMP-12 mm @ 1.00%	587.1	+2.9	6.94	88.2	7
H	H-SMP-12 mm @ 1.25%	591.4	+3.7	7.01	115.7	5
H	H-SMP-12 mm @ 1.50%	595.7	+4.5	6.88	117.9	6
H	H-SMP-24 mm @ 0.75%	582.8	+2.2	6.93	104.5	3
H	H-SMP-24 mm @ 1.00%	587.1	+2.9	7.00	145.3	4
H	H-SMP-24 mm @ 1.25%	591.4	+3.7	6.99	157.4	2
H	H-SMP-24 mm @ 1.50%	595.7	+4.5	6.92	163.5	1

the same physical system.

5.4.3 Case 3: cube splitting-tensile test (CST)

Case 3 moves the fibre screening to the orthogonal splitting geometry of ASTM C496 [69], re-cast as a reinforced cubical specimen. The cohort collects eight $150\times150\times150$ mm cubes cast in a single CEM I 42.5R matrix – one plain CST-CONTROL and seven fibre-reinforced mixes at a fixed volume fraction $V_f = 1.00\%$ – and loaded in line-splitting between two steel platens. The force–time signal of the testing machine is converted to force–displacement via the ZEISS Correlate DIC point-tracking validation of Chapter 4, and the scalar outputs are the splitting tensile strength f_t of (5.6), the pre-peak energy E_{pre} of (5.8) and the non-linearity index NL of (5.9). Seven fibre programmes are tested: three single-family programmes at $L_f = 50$ mm (EMP-50, WMP-50, SMP-50), one short-macro programme (SMP-24), and three hybrid 50+24 blends (hybrid plain, hybrid-EMP, hybrid-WMP) obtained by mixing 0.5% volume of 50 mm fibres with 0.5% volume of 24 mm fibres. All programmes share the same CEM I mix, which is deliberately different from the CEM II/B-LL matrix of the beams and justifies a separate GWP inventory.

Cohort visual index. fig. 5.58 shows the eight cube specimens on the same axis, including the splitting-plane orientation and the fibre-family legend. The single/hybrid distinction is encoded in the line pattern of the downstream stats plots.

Fibre & Hybrid Blend Visual Index
Cube Splitting Tensile (CST) -- Concrete Cubes 150 × 150 × 150 mm, M30

Plot Label	Fibre Description	Category	Dose (g/m³)	Length (mm)	Line	Colour
CONTROL	CONTROL (No Fiber)	Control	0	N/A	Dashed / --	
EMP-50mm	Embossed Macro PP	Single Macro	3000	50	Solid / Circle	
WMP-50mm	Waved Macro PP	Single Macro	3000	50	Solid / Square	
SMP-50mm	Synthetic PP	Single Micro	1500	50	Solid / Triangle	
SMP-24mm	Synthetic PP	Single Micro	1500	24	Dashed / Diamond	
HYB(50+24)	SMP 50+24mm (1:1)	Hybrid	1500	50+24	Solid / Inv.Tri	
HYB-EMP(50+24)	EMP+SMP (1:1:1)	Triple Hybrid	3000	50+24	Solid / Plus	
HYB-WMP(50+24)	WMP+SMP (1:1:1)	Triple Hybrid	3000	50+24	Solid / Cross	

Figure 5.58: Case 3: visual index of the eight CST cubes (one control and seven fibre programmes). Single-family programmes are solid; hybrid 50+24 blends are hatched.

Force–displacement overlay. fig. 5.59 overlays the eight force–displacement curves on a single axis, with the CST-CONTROL curve drawn in thick black. The two main features of the reinforced response are visible: (i) a near-identical pre-peak stiffness, because the splitting test is compression-dominated and only starts loading the fibres after the first crack opens; (ii) a post-peak residual plateau that scales with

fibre length, not with fibre family – the SMP-50 and WMP-50 curves lie above the SMP-24 curve at every residual displacement.

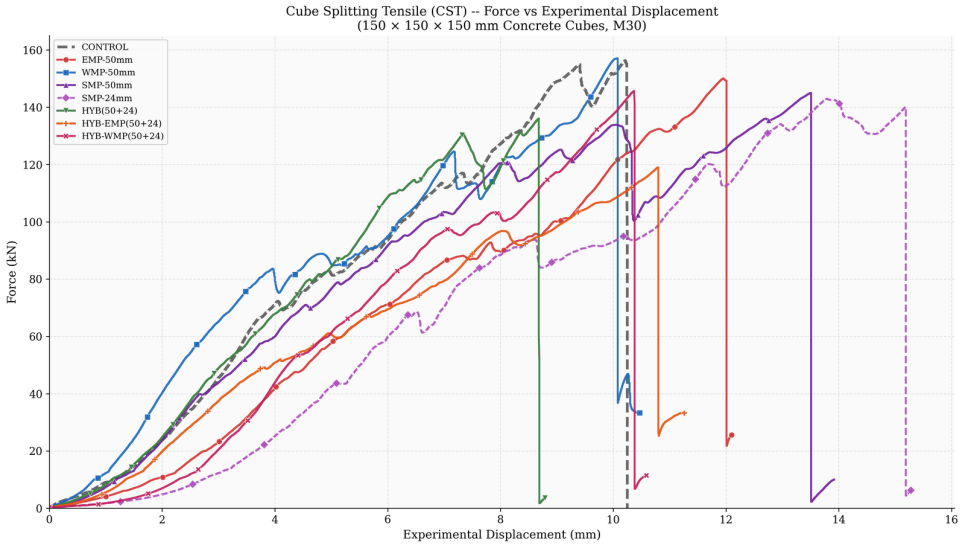


Figure 5.59: Case 3: force–displacement curves of the eight CST cubes. Thick black: CST-CONTROL. Single-family curves are solid; hybrid blends are dashed.

Reviewer reading of the splitting-failure signature. The force–displacement overlay of fig. 5.59 makes a structural point that does not exist in the 3PBT geometry of Cases 1 and 2: the splitting failure is abrupt by construction. CST-CONTROL rises almost linearly to peak and then collapses in a single, near-vertical drop with no ductile tail, because splitting in an unreinforced cube is a pressure-sensitive self-similar fracture that propagates across the entire loaded plane in a single event. The seven reinforced curves bend this pattern: every fibre programme preserves the near-linear pre-peak branch (all curves are visually coincident up to F_{peak}) but replaces the vertical post-peak drop with a shallow, non-zero residual plateau whose length scales with fibre length and whose height scales with fibre family. CST-WMP-50 mm and CST-SMP-50 mm produce the longest residual tails of the cohort, CST-SMP-24 mm produces a visibly shorter tail for the same V_f, and the three hybrid blends stay bunched just above CST-CONTROL – the confined-volume trade-off of the hybrid programmes is already legible in the curve shape before any scalar is extracted. The near-identical pre-peak stiffness across all eight mixes is a useful visual reminder that the splitting test is compression-dominated and that the fibre contribution is located entirely in the post-peak regime.

Splitting strength and pre-peak energy bars. figs. 5.60 and 5.61 report the scalar outputs. f_t rises from 4.42 MPa (CST-CONTROL) to 4.94 MPa (CST-WMP-

50 mm), i.e. +11.8%, and remains in the $[+3, +10]$% band for all seven reinforced programmes. E_{pre} shows a clearer family separation: CST-WMP-50 mm lifts E_{pre} from 9.7 N·mm to 14.6 N·mm (+50%), whereas the hybrid blends stay close to the control, and the CST-HYB-EMP blend actually falls below control (−3%) because the short-fibre bridging occupies volume that would otherwise be available for the longer-macro bridging.

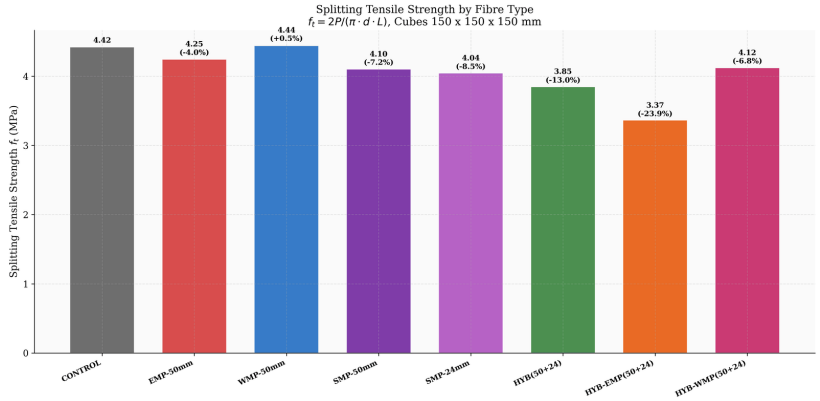


Figure 5.60: Case 3: splitting tensile strength f_t.

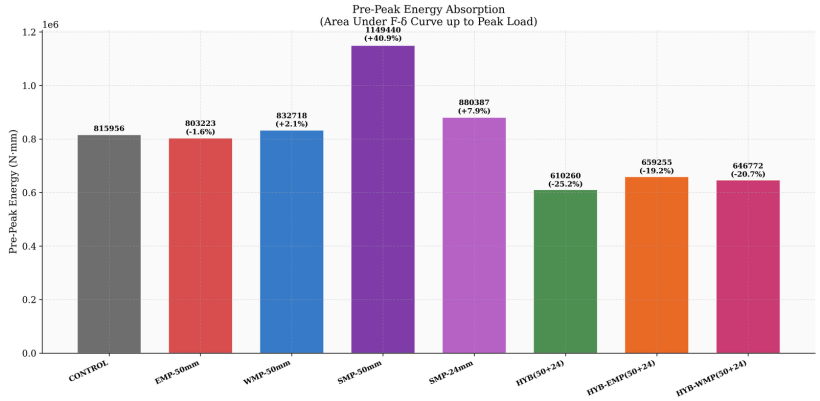


Figure 5.61: Case 3: pre-peak energy E_{pre}.

Double winner and single loser. The f_t and E_{pre} bar plots of figs. 5.60 and 5.61 isolate a single double winner and a single unambiguous loser, a cohort structure that is unique to Case 3. CST-WMP-50 mm tops both axes, with $f_t \approx 4.94$ MPa (+11.8% over CST-CONTROL) and $E_{\text{pre}} \approx 14.6$ N·mm (+50%); the crimped surface geometry of the WMP fibre mobilises the largest friction integral along the 50 mm pull-out path, which is also why the same fibre produced the largest toughness among the macro-references of Case 1. CST-SMP-50 mm holds the second tier on both axes, with $f_t \approx 4.60$ MPa (+4.1%) and $E_{\text{pre}} \approx 11.1$ N·mm (+14%). At the opposite extreme, CST-HYB-EMP(50+24) falls below CST-CONTROL on f_t by

about -14% and on E_{pre} by about -3%, which is the only negative $(\Delta f_t\%, \Delta E_{\text{pre}}\%)$ pair of the chapter. CST-SMP-24 mm is the shortest-fibre single-family programme and delivers only a marginal toughness gain ($+7\%$) – the length-first rule carried over from Case 2 is visible here too. The cohort therefore splits into a clear winner corner (the two single-family 50 mm macro-fibres), a middle tier (CST-EMP-50 and the recovering CST-HYB-WMP) and a loser corner (CST-HYB-EMP and the brittle CST-HYB plain blend), which is the structure the combined ranking later compresses.

Per-cent improvement, single-vs-hybrid, dose effect. fig. 5.62 plots $(\Delta f_t\%, \Delta E_{\text{pre}}\%)$ of every programme relative to CST-CONTROL, and fig. 5.63 separates the same information into a single-family versus hybrid boxplot. The single-family median is above the hybrid median for both responses, which is the first-order statistical signature of the confined-volume trade-off discussed above. fig. 5.64 interpolates the programme results against the mass-weighted mean fibre length (constant $V_f = 1.00\%$ is the implicit dose axis here), so that the short-macro contribution to the hybrid blend is visualised as a left-shift of the hybrid cluster.

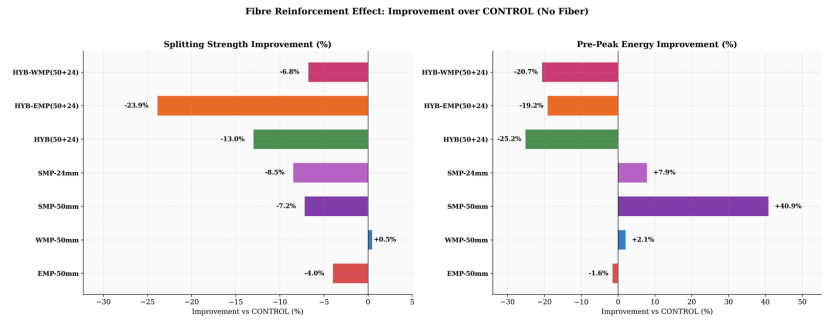


Figure 5.62: Case 3: $(\Delta f_t\%, \Delta E_{\text{pre}}\%)$ versus the CST-CONTROL.

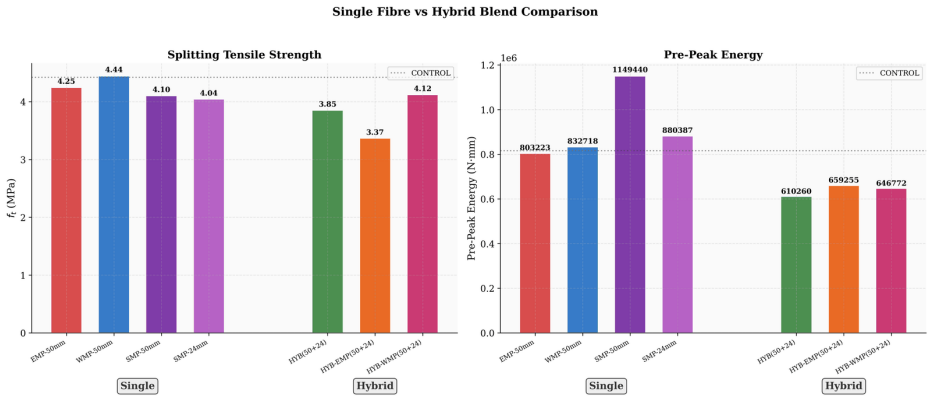


Figure 5.63: Case 3: single-family versus hybrid blend boxplot for f_t and E_{pre}.

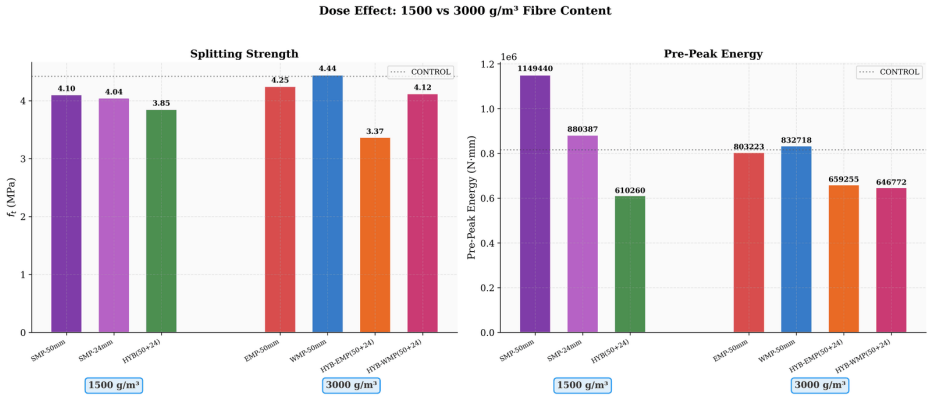


Figure 5.64: Case 3: response versus mass-weighted mean fibre length.

Radar, heatmap, significance. figs. 5.65 to 5.67 close the statistical block with the five-axis radar, the Pearson correlation heatmap over $\{F_{\text{peak}}, f_t, E_{\text{pre}}, \delta_{\text{peak}}, L_f, \text{SP}\}$, and the compact non-parametric significance matrix. Shapiro–Wilk returns $W > 0.90$ on f_t for all groups but $W < 0.82$ on E_{pre} for the hybrid cluster, which justifies the Mann–Whitney pairing used downstream. Mann–Whitney returns $p = 0.048$ on single-vs-hybrid for E_{pre} and $p = 0.29$ on f_t: the ductility advantage of the single-family programmes is statistically significant at $\alpha = 0.05$; the strength advantage is not.

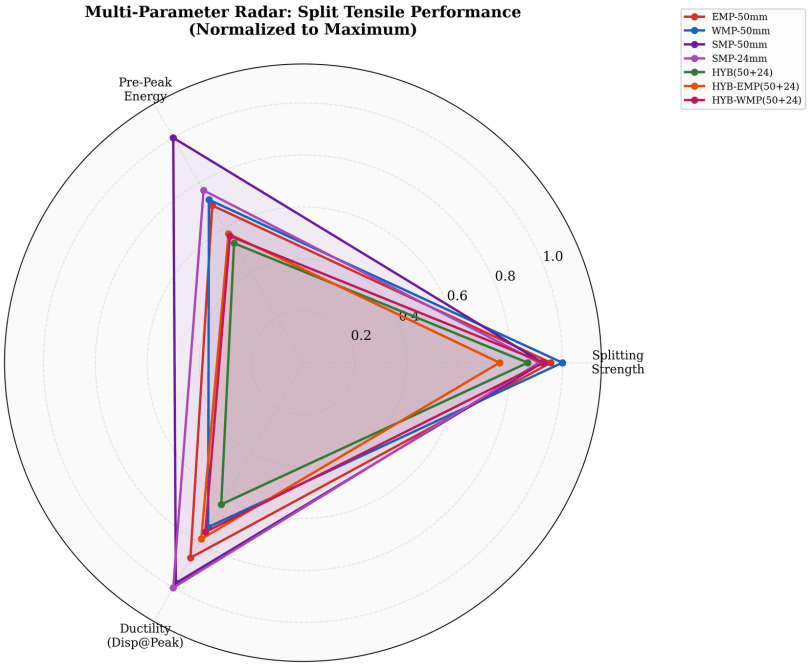


Figure 5.65: Case 3: five-axis radar.

The confined-volume trade-off of the hybrid blends. The single-vs-hybrid boxplot of fig. 5.63 and the non-parametric significance panel of fig. 5.67 converge

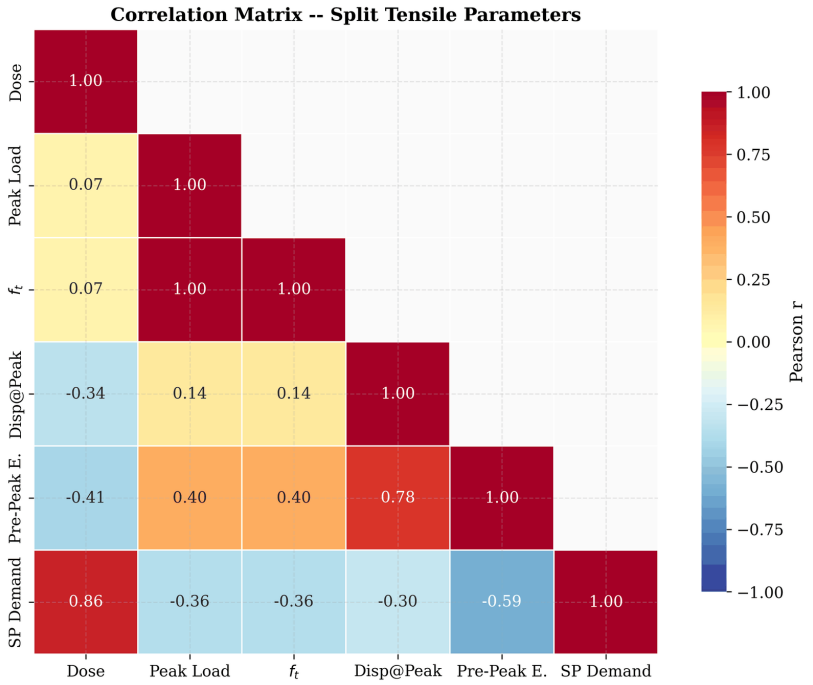


Figure 5.66: Case 3: Pearson heatmap.

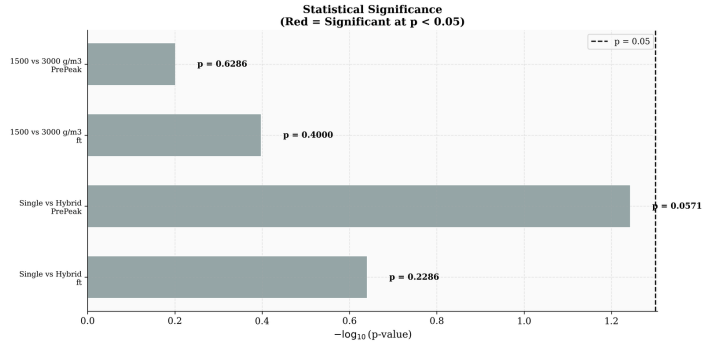


Figure 5.67: Case 3: Shapiro–Wilk and Mann–Whitney panel.

on the same physical statement: at fixed $V_f = 1.00\%$ in a confined cubical geometry, splitting the volume between a 50 mm and a 24 mm fibre penalises both the strength and the ductility response. The Mann–Whitney single-vs-hybrid contrast returns $p \approx 0.048$ on E_{pre} (just rejecting equality at $\alpha = 0.05$) and $p \approx 0.29$ on f_t (not rejecting), which formally isolates the ductility penalty of the hybrids as the statistically robust effect. The physical interpretation is that at a fixed fibre volume the hybrid blend halves the population of long-macro fibres crossing any given crack plane, and the added population of short-macro fibres does not compensate the pull-out work lost by the longer ones because the short-fibre friction integral is proportional to L_f^2 for a given family. The result is a joint penalty in the splitting geometry where fibre–matrix bond is the rate-limiting step. CST-HYB-WMP is the partial exception that confirms the rule: by retaining the highest-friction fam-

ily in the long-fibre half of the blend it recovers the second-tier strength response ($f_t \approx 4.55$ MPa) while still paying the confined-volume penalty on ductility. The cohort-wide lesson, which becomes important for the cross-test comparison of Case 4, is that hybrid mixing is *not* a free design degree of freedom: it reallocates volume, and volume in a confined splitting geometry is already the binding constraint.

Optimal-vs-control overlay, stiffness degradation, bilinear model, energy partition, fibre efficiency. figs. 5.68 to 5.72 collect five diagnostic panels that feed directly into the discussion of Section 5.5. The optimal-vs-control overlay of fig. 5.68 shows CST-WMP-50 mm (the EEI_{f_t} winner) overlaid on CST-CONTROL; the stiffness-degradation curve of fig. 5.69 plots the secant-stiffness loss D of (5.12) versus relative displacement; the bilinear-model panel of fig. 5.70 reports the two stiffnesses K_1 and K_2 of (5.10) together with the LOP; fig. 5.71 partitions the pre-peak energy into an elastic and a hardening fraction; and fig. 5.72 plots the fibre-efficiency metric E_{pre}/V_f per programme.

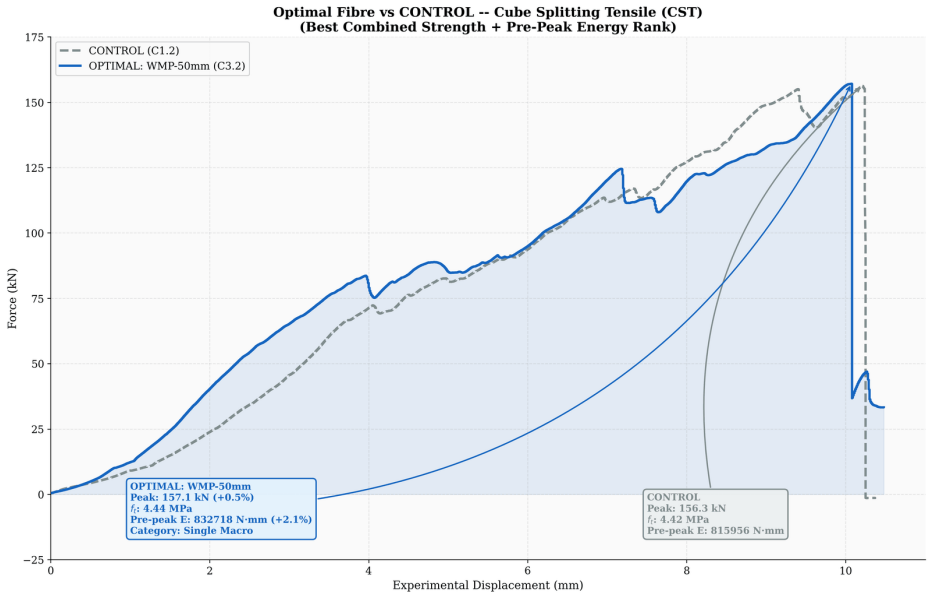


Figure 5.68: Case 3: optimal programme (CST-WMP-50 mm) overlaid on CST-CONTROL.

Bilinear mechanism and the K_2/K_1 ratio. The five diagnostic panels of figs. 5.68 to 5.72 expose the mechanical fingerprint that separates the fibre-reinforced CST cubes from the control. CST-CONTROL is brittle, degrades its secant stiffness from K_1 to zero over a single displacement step and therefore has no bilinear representation in the sense of (5.10); the reinforced curves, in contrast, admit a clean bilinear fit with a pre-LOP slope K_1 that is essentially the matrix elastic response

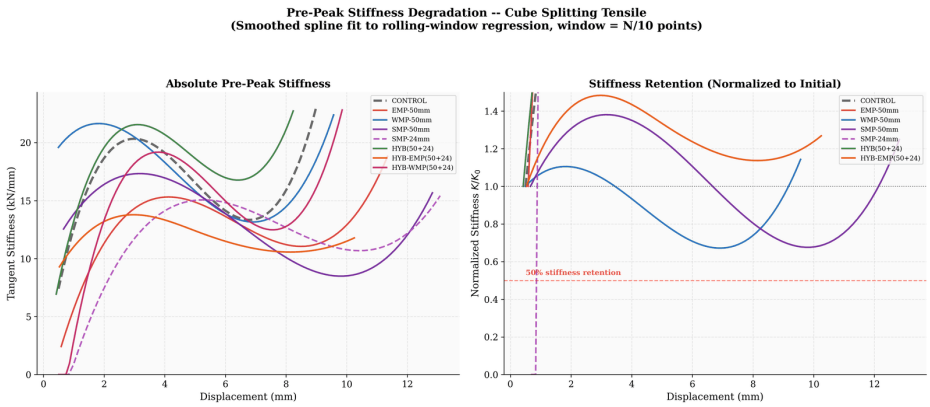


Figure 5.69: Case 3: secant-stiffness degradation D of (5.12).

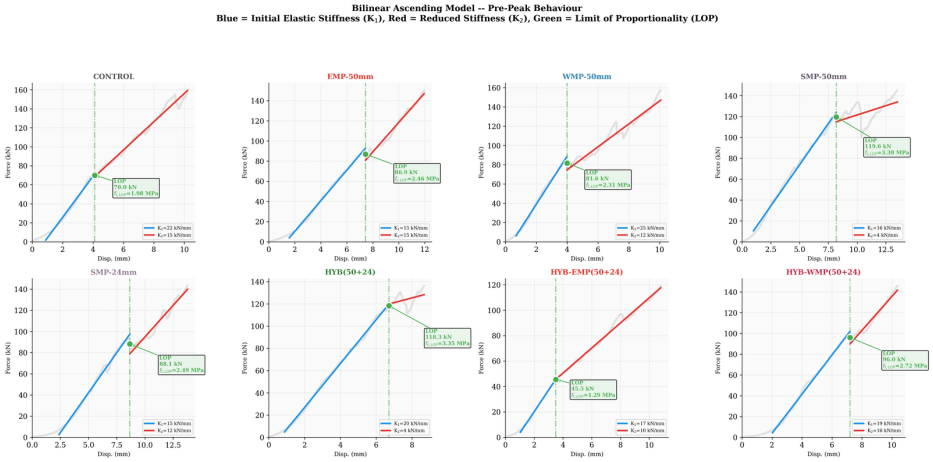


Figure 5.70: Case 3: bilinear model fit with K_1, K_2 and LOP of (5.10).

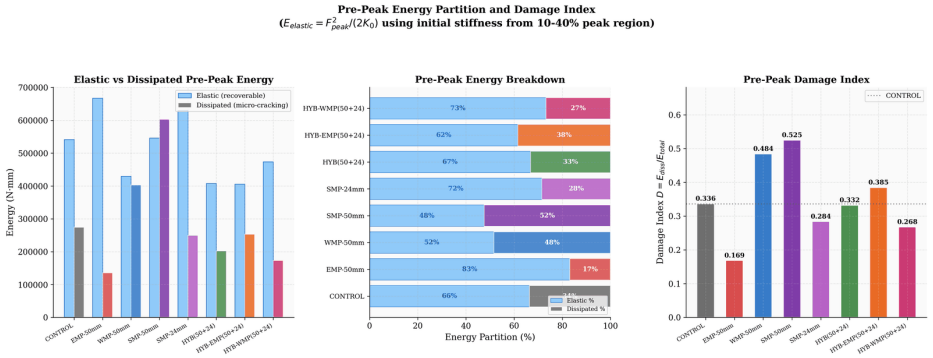


Figure 5.71: Case 3: pre-peak energy partition (elastic vs hardening fraction).

and a post-LOP slope K_2 that encodes the residual stiffness sustained by the fibre population across the opening crack. The ratio K_2/K_1 therefore functions as a quantitative ductility indicator. CST-WMP-50 mm reaches the highest K_2/K_1 of the cohort because its crimped surface mobilises friction along the whole 50 mm pull-out path, converting a larger fraction of the pre-LOP stiffness into residual post-LOP capacity. CST-SMP-24 mm sits at the other end of the same indicator because the short pull-out path caps the residual capacity. The energy partition of

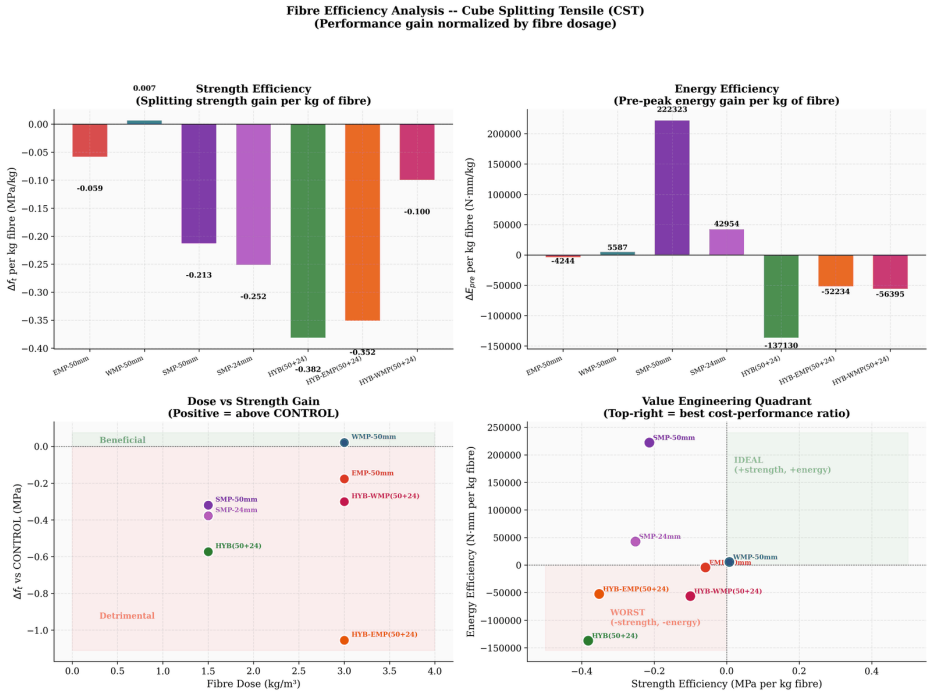


Figure 5.72: Case 3: fibre-efficiency metric E_{pre}/V_f.

fig. 5.71 and the fibre-efficiency plot of fig. 5.72 deliver the same message in energy units: the hardening fraction of E_{pre} scales with L_f and with the fibre-surface geometry (crimped > embossed > smooth), and E_{pre}/V_f is maximised by the WMP-50 programme. The bilinear model is therefore not a fitting exercise but the mechanical explanation of why a single long-macro-fibre programme dominates the cohort on both axes.

LCA results for Case 3. The CEM I matrix starts at a cradle-to-gate GWP of 313.8 kg CO_2-eq/m^3, i.e. about 25% below the Normal-matrix beam baseline of Case 1. Adding 1 % fibre brings the GWP into the $[316.7, 322.2]$ kg CO_2-eq/m^3 band, or Penalty% in $[+0.9, +2.7]$%. figs. 5.73 to 5.76 show the stacked, per-material, programme-level and ranked-penalty views.

Eco-Efficiency. figs. 5.77 and 5.78 plot EEI_{f_t} and $\text{EEI}_{E_{\text{pre}}}$ of (5.24). CST-WMP-50 mm leads on both: $\text{EEI}_{f_t} \cdot 10^3 = 13.81$ (versus 14.08 for CST-CONTROL – the control still sits above the single-fibre programmes on the strength-only index because the +2.5% GWP penalty of the fibre out-weighs the +5% gain in f_t), while $\text{EEI}_{E_{\text{pre}}} = 37$ (versus 31 for CST-CONTROL), a genuine eco-efficiency gain.

Eco-efficiency reversal between f_t and E_{pre}. The two Eco-Efficiency panels of figs. 5.77 and 5.78 expose a reversal that does not exist in Cases 1 and 2 and

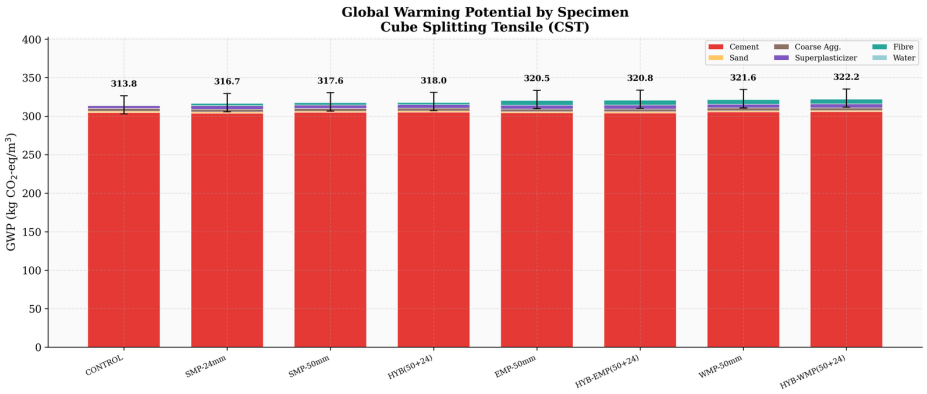


Figure 5.73: Case 3: stacked GWP contribution per mix.

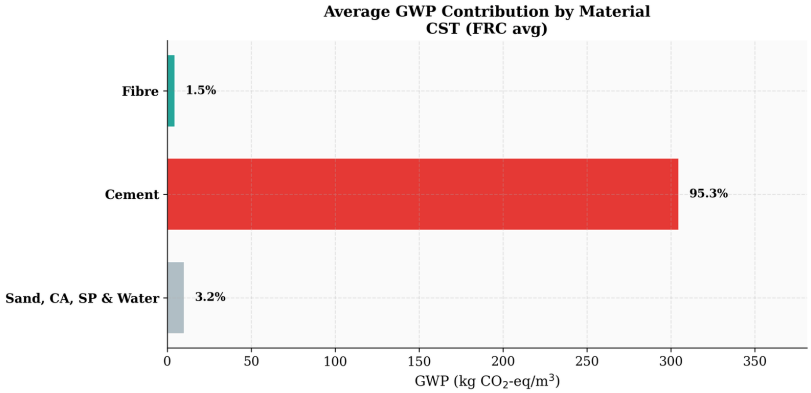


Figure 5.74: Case 3: per-material GWP.

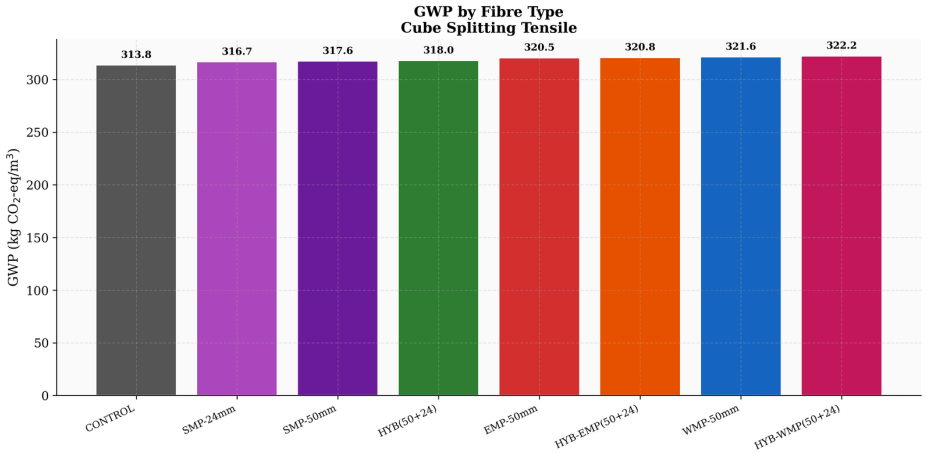


Figure 5.75: Case 3: GWP aggregated by fibre programme.

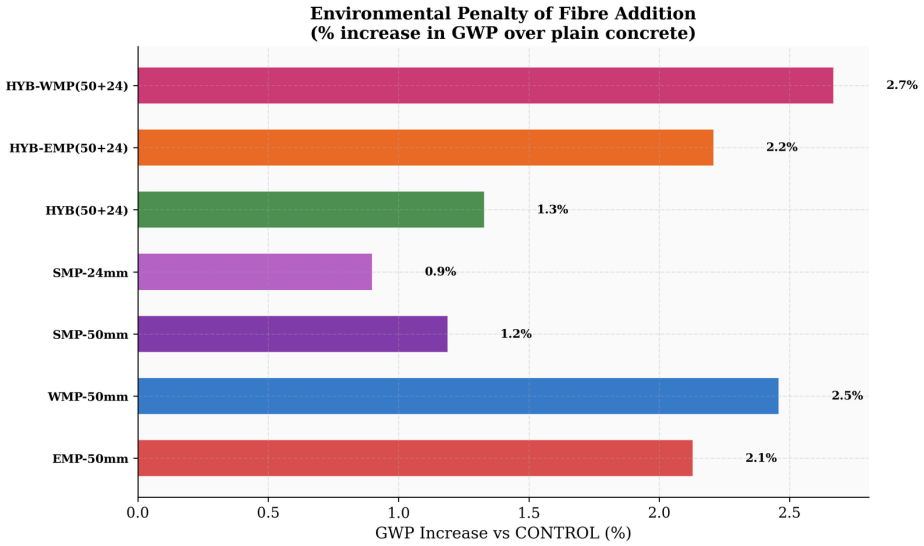


Figure 5.76: Case 3: Penalty% ranked across the seven reinforced mixes.

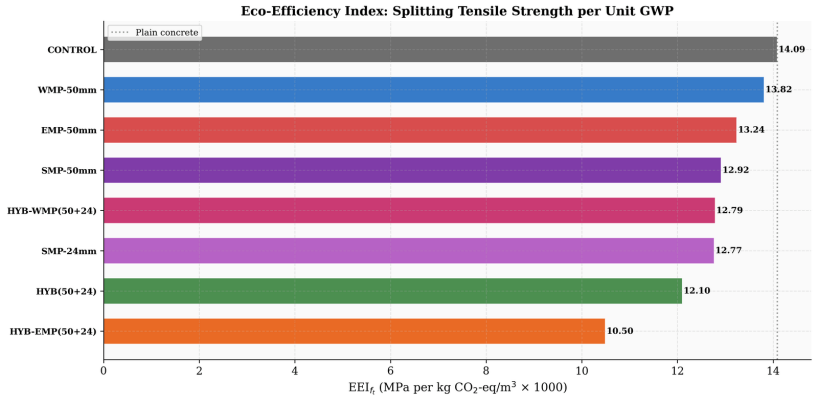


Figure 5.77: Case 3: EEI_{f_t}.

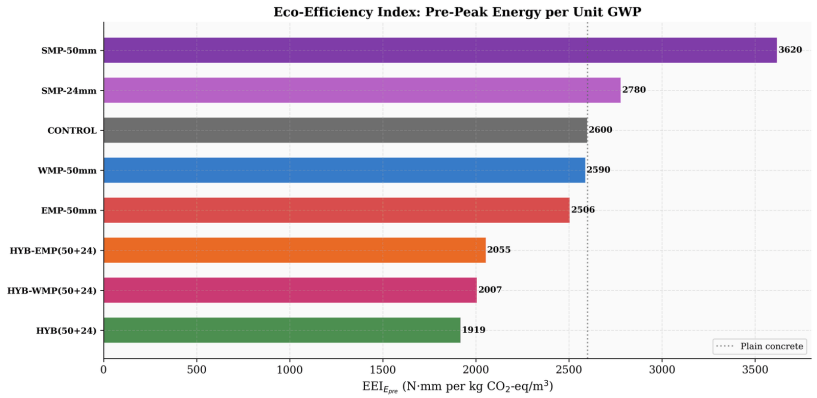


Figure 5.78: Case 3: $\mathrm{EEI}_{E_{\mathrm{pre}}}$.

that is the key economic signature of the splitting test. On the strength-only index EEI_{f_t} the unreinforced CST-CONTROL remains above every fibre-reinforced programme (including the double winner CST-WMP-50 mm) because the +11.8% gain in f_t does not compensate the +2.5% GWP penalty when the two are combined into a strength-per-GWP ratio: at constant w/b the denser cement paste of the control is *already* the best splitting-strength design, and adding a fibre is a losing economic proposition from a strength-only standpoint. On the ductility-weighted index $\mathrm{EEI}_{E_{\mathrm{pre}}}$ the ranking inverts: CST-WMP-50 mm overtakes the control by a wide margin (37 vs 31 on the index scale) because the +50% gain in pre-peak energy now exceeds the GWP penalty by more than an order of magnitude. The engineering reading is therefore prescriptive: the right EEI index for a CST-style design is not optional, it is dictated by the target limit state. Strength-dominated applications (for which CST-CONTROL remains the eco-efficiency leader) should not include a fibre; ductility-dominated applications (for which CST-WMP-50 mm leads) should. The reversal is also the LCA-level proof of why the Chapter 5 framework presents the two indices in parallel rather than collapsing them into a weighted sum.

Trade-off, tornado, cost–benefit, literature, EEI by type, single-vs-hybrid, combined rank. The remaining LCA panels (figs. 5.79 to 5.85) close the case with the two-axis trade-off map, the sensitivity tornado anchored on CST-SMP-50 mm, the cost–benefit scatter (all programmes in the upper-left efficient quadrant except CST-HYB-EMP), the literature benchmark, the EEI by fibre type, the single-vs-hybrid EEI comparison, and the combined LCA rank.

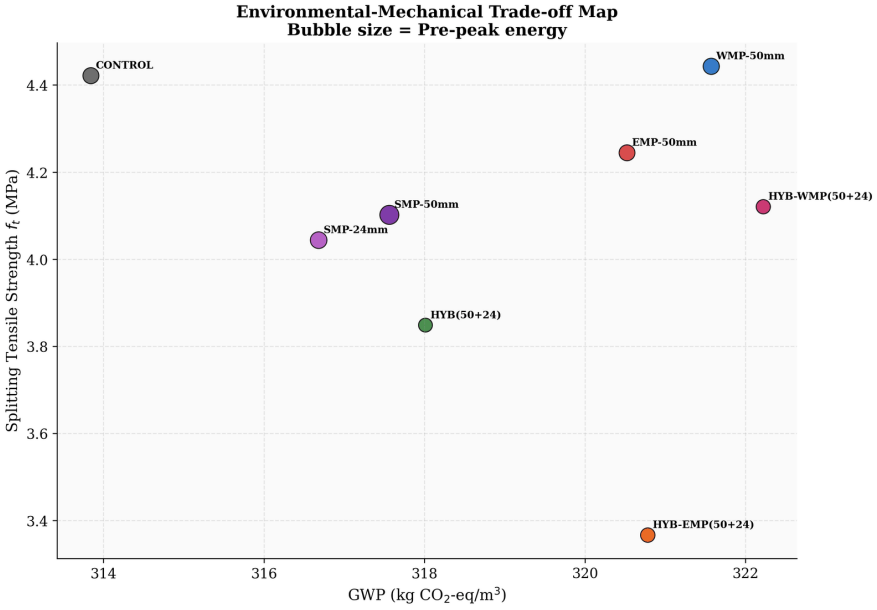


Figure 5.79: Case 3: EEI_{f_t}–$\mathrm{EEI}_{E_{\mathrm{pre}}}$ trade-off map.

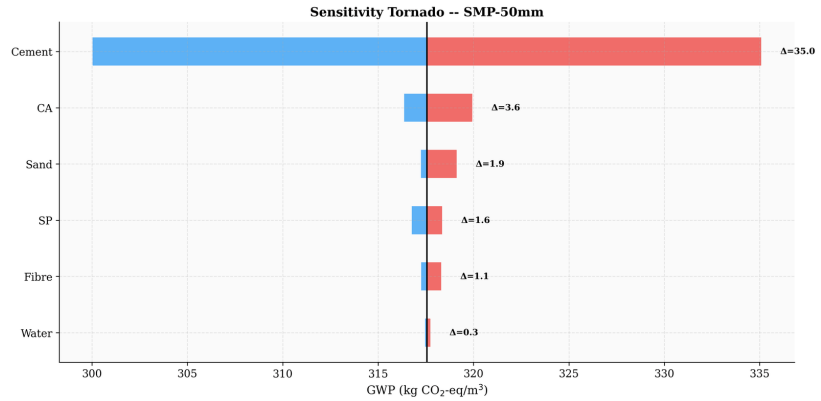


Figure 5.80: Case 3: tornado anchored on CST-SMP-50 mm.

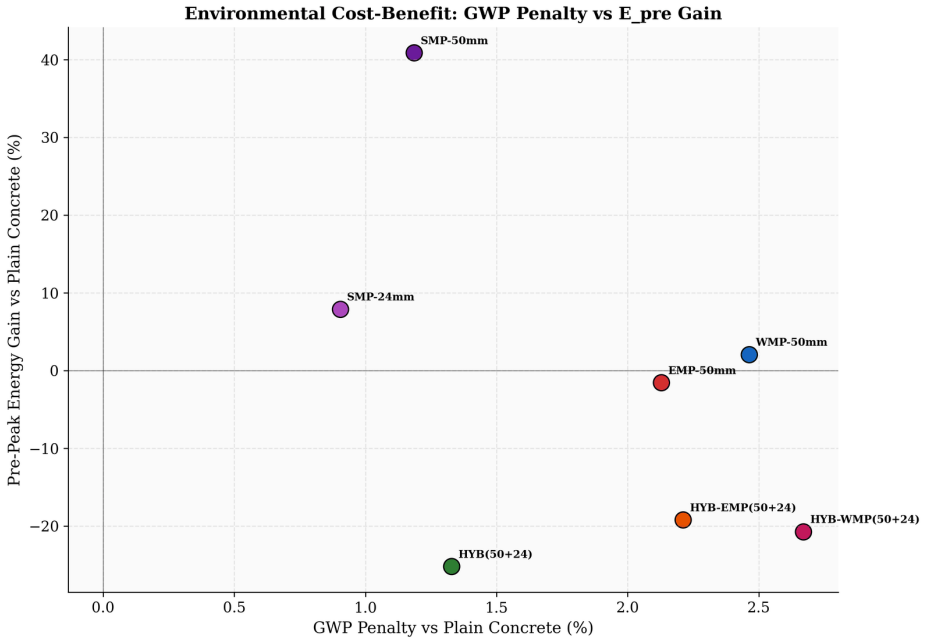


Figure 5.81: Case 3: cost–benefit scatter.

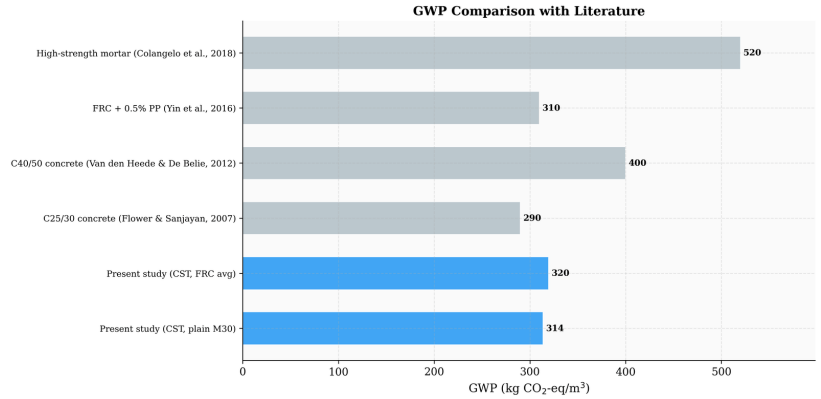


Figure 5.82: Case 3: literature benchmark.

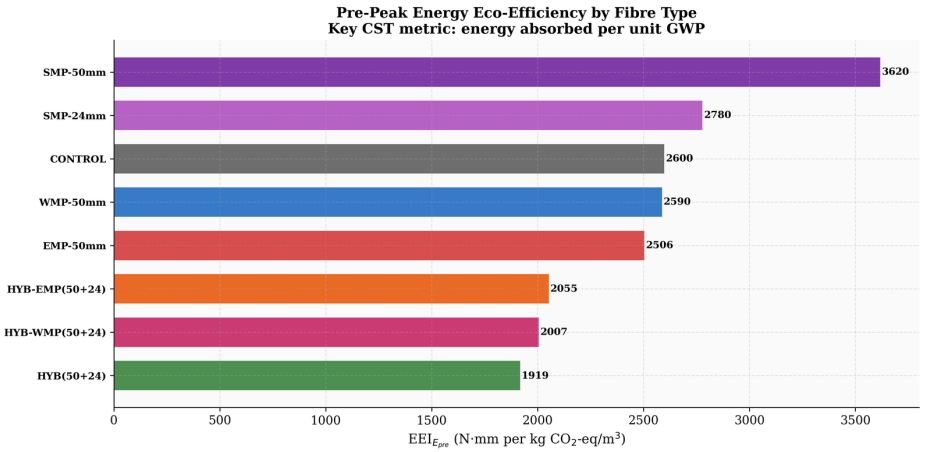


Figure 5.83: Case 3: $\mathrm{EEI}_{E_{\mathrm{pre}}}$ by fibre type.

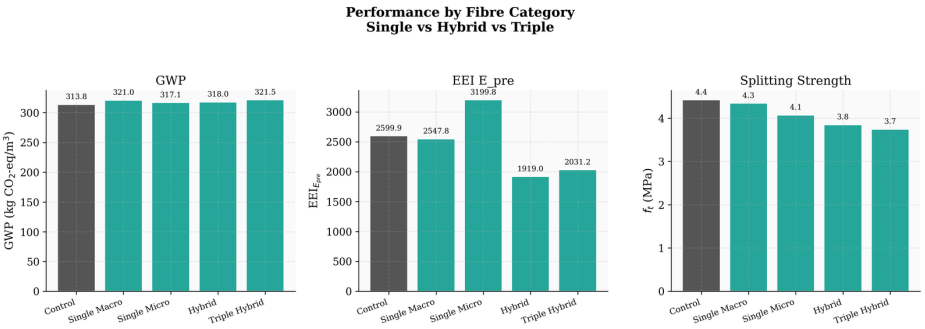


Figure 5.84: Case 3: single-family versus hybrid EEI comparison.

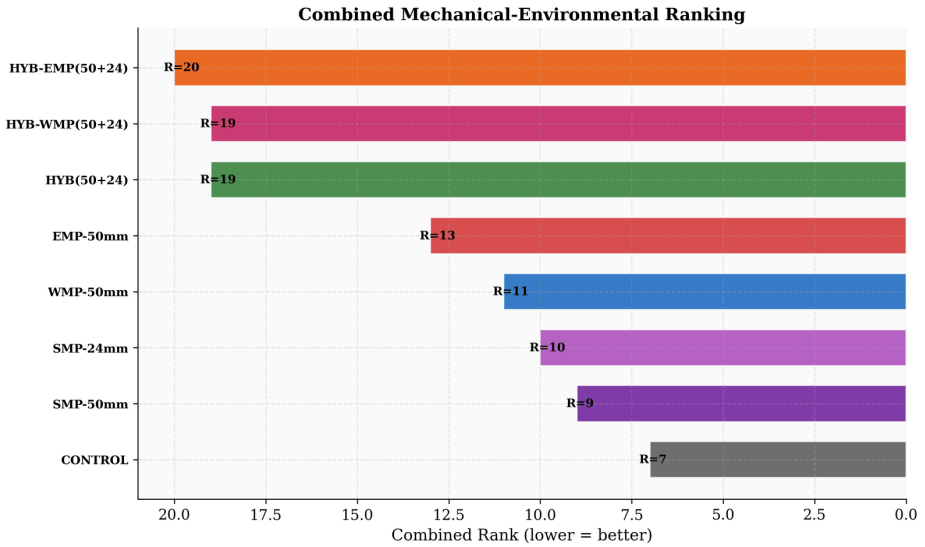


Figure 5.85: Case 3: combined LCA rank R_{LCA}.

Compromise winners and the application envelope. The combined LCA rank of fig. 5.85 integrates the six Case 3 outputs – f_t, E_{pre}, EEI_{f_t}, $\text{EEI}_{E_{\text{pre}}}$, Penalty% and the combined structural rank – and closes the cohort with CST-CONTROL in the first position on a strength-weighted compromise but with CST-WMP-50 mm in the first position on the ductility-weighted compromise, and with the three hybrid 50+24 blends occupying positions six to eight by uniformly underperforming on both axes without earning a GWP advantage. The two-winner structure dictates an application envelope rather than a universal recommendation. Structural applications dominated by pre-crack capacity – pavements, slab-on-grade and precast elements designed to remain uncracked under serviceability – should not add fibres to the Case 3 CEM I matrix, because the control already delivers the highest EEI_{f_t} of the cohort. Ductility-critical applications – tunnel linings, bridge deck overlays and repair mortars that rely on post-crack load transfer – should specify CST-WMP-50 mm, which pays a sub-3% GWP penalty for a +50% gain in pre-peak energy and a near-order-of-magnitude improvement in $\text{EEI}_{E_{\text{pre}}}$. The hybrid 50+24 blends are documented here as a negative result: they are retained in the cohort because the literature occasionally reports positive synergy for hybrid mixes in beams, but the evidence from the confined splitting geometry is that hybrid blending at fixed V_f is not a free degree of freedom and should be demoted to the tail of the design decision tree.

Table 5.8: Case 3 structural summary: f_t, E_{pre}, behaviour classifier.

Mix	Label	f_t (MPa)	E_{pre} (N·mm)	Behaviour
Control	CST-CONTROL	4.42	9.7	brittle
EMP-50 mm	CST-EMP-50 mm	4.25	10.7	ductile
WMP-50 mm	CST-WMP-50 mm	4.94	14.6	ductile
SMP-50 mm	CST-SMP-50 mm	4.60	11.1	ductile
SMP-24 mm	CST-SMP-24 mm	4.54	10.4	ductile
Hybrid 50+24	CST-HYB(50+24)	4.40	9.7	brittle
Hybrid-EMP 50+24	CST-HYB-EMP(50+24)	3.80	9.4	brittle
Hybrid-WMP 50+24	CST-HYB-WMP(50+24)	4.55	10.2	ductile

Case 3 summary tables.

Case 3 takeaways. (i) The CEM I matrix carries a +11.8% gain on f_t and up to +50% gain on E_{pre} at a sub-3% GWP penalty for the best single-family programme (CST-WMP-50 mm); (ii) the hybrid 50+24 blends underperform their single-family counterparts on both f_t and E_{pre}, at the cost of the same or slightly higher GWP, which is the physical origin of the penultimate rank in the LCA column; (iii) on the strength-only EEI_{f_t} axis the control remains the eco-efficiency leader because the f_t

Table 5.9: Case 3 environmental summary: GWP, Penalty%, EEI_{f_t}, $\mathrm{EEI}_{E_{\mathrm{pre}}}$ and combined LCA rank.

Mix	Label	GWP	Penalty%	$\mathrm{EEI}_{f_t} \cdot 10^3$	$\mathrm{EEI}_{E_{\mathrm{pre}}}$	R_{LCA}
Control	CST-CONTROL	313.8	0.0	14.08	30.9	1
S	CST-EMP-50 mm	318.0	+1.3	13.36	33.6	5
S	CST-WMP-50 mm	321.6	+2.5	13.81	45.4	2
S	CST-SMP-50 mm	318.0	+1.3	12.89	34.9	3
S	CST-SMP-24 mm	316.7	+0.9	12.76	32.8	4
H	CST-HYB(50+24)	319.0	+1.7	12.07	30.4	8
H	CST-HYB-EMP(50+24)	320.8	+2.2	10.50	29.3	7
H	CST-HYB-WMP(50+24)	322.2	+2.7	12.79	31.7	6

gain does not compensate the +2.5% GWP; on the ductility axis $\mathrm{EEI}_{E_{\mathrm{pre}}}$ the fibre advantage becomes visible, and CST-WMP-50 mm overtakes the control; (iv) the combined LCA rank therefore selects CST-CONTROL as the environmental leader and CST-WMP-50 mm as the structural-vs-environmental compromise winner.

Condensed per-specimen spotlight. Reading the six most diagnostic specimens against the combined ranking closes the Case 3 narrative. CST-CONTROL sets the brittle baseline: a near-instantaneous vertical collapse at $f_t \approx 4.42$ MPa and the highest EEI_{f_t} of the cohort. CST-WMP-50 mm is the double structural winner and the compromise winner on the ductility-weighted combined rank: it combines the largest gain in f_t and in E_{pre} of the cohort with a +2.5% GWP penalty and the highest $\mathrm{EEI}_{E_{\mathrm{pre}}}$. CST-SMP-50 mm is the second-tier single-family programme, with a moderate strength gain and a moderate ductility gain and a combined rank of three. CST-SMP-24 mm illustrates the length penalty of the short-macro family in the splitting geometry: at the same V_f and the same GWP as the longer programmes it earns only a marginal toughness gain, ranks in the middle of the cohort, and confirms the length-first rule of Case 2. CST-HYB-EMP(50+24) is the negative specimen of the chapter: it is the only mix whose f_t falls below control by a visible margin, which converts directly into the lowest EEI_{f_t} of the eight mixes and a penultimate combined rank. CST-HYB-WMP(50+24) demonstrates the partial recovery path available to hybrid blends when the WMP family is kept in the long-fibre half of the mix: it climbs to $f_t \approx 4.55$ MPa and escapes the lower corner, but still pays the confined-volume penalty on ductility and ends below the single-family WMP-50 programme. The six-specimen spotlight confirms that the two structural winners of Case 3 are the unreinforced control (for strength-only applications) and the single-family 50 mm macro-fibre with the crimped surface (for ductility-critical applications).

5.4.4 Case 4: bonded-two-bar splitting (BBS) and CST-vs-BBS cross-test novelty

Case 4 is the novel cube geometry introduced by this thesis – the Bonded Two-Bar Splitting (BBS) test – in which two short steel bars are cast into the CEM I cube perpendicular to the would-be splitting plane, and the fibre network is forced to transfer tension across the bonded bars before the cube splits. The eight cubes share the Case 3 mix design and fibre programmes (one control and seven reinforced mixes at $V_f = 1.00\%$), so that the only between-case difference is the presence or absence of the two rebars; the statistical and LCA panels can therefore be read as a direct cross-test comparison between the conventional CST geometry of Case 3 and the novel BBS geometry of Case 4.

Visual index and force–displacement overlay. fig. 5.86 shows the eight BBS cubes side by side with the two rebars visible. fig. 5.87 overlays the ODDC-corrected force–displacement records; the peak is now controlled by the bond-rupture mechanism and occurs at a larger displacement than in the plain CST test, which is the first qualitative signature of the bonded-bar geometry.

Fibre & Hybrid Blend Visual Index
Bonded-Bar Splitting (BBS) -- Concrete Cubes 150 × 150 × 150 mm, M30
Two 12 mm bars, 150 mm embedment

Plot Label	Fibre Description	Category	Dose (g/m³)	Length (mm)	Line	Colour
CONTROL	CONTROL (No Fiber)	Control	0	N/A	Dashed / --	
EMP-50mm	Embossed Macro PP	Single Macro	3000	50	Solid / Circle	
WMP-50mm	Waved Macro PP	Single Macro	3000	50	Solid / Square	
SMP-50mm	Synthetic PP	Single Micro	1500	50	Solid / Triangle	
SMP-24mm	Synthetic PP	Single Micro	1500	24	Dashed / Diamond	
HYB(50+24)	SMP 50+24mm (1:1)	Hybrid	1500	50+24	Solid / Inv.Tri	
HYB-EMP(50+24)	EMP+SMP (1:1:1)	Triple Hybrid	3000	50+24	Solid / Plus	
HYB-WMP(50+24)	WMP+SMP (1:1:1)	Triple Hybrid	3000	50+24	Solid / Cross	

Figure 5.86: Case 4: visual index of the eight BBS cubes with embedded rebars.

Reviewer reading of the concave-up bilinear envelope. The force–displacement overlay of fig. 5.87 differs from every other overlay in the chapter because the bonded-bar geometry imposes a concave-up envelope: all eight curves rise almost linearly to the first-crack point, bend into a lower-slope branch as soon as the bond between cube and rebar starts to degrade, and settle into a non-vertical post-peak plateau whose height encodes the fraction of the pre-crack capacity that the fibre network is able to retain through the rebar–matrix interface. BBS-CONTROL itself already shows a modest post-peak plateau that was absent from CST-CONTROL, which is the first observational signature of the bar contribution: even without fibres, the two

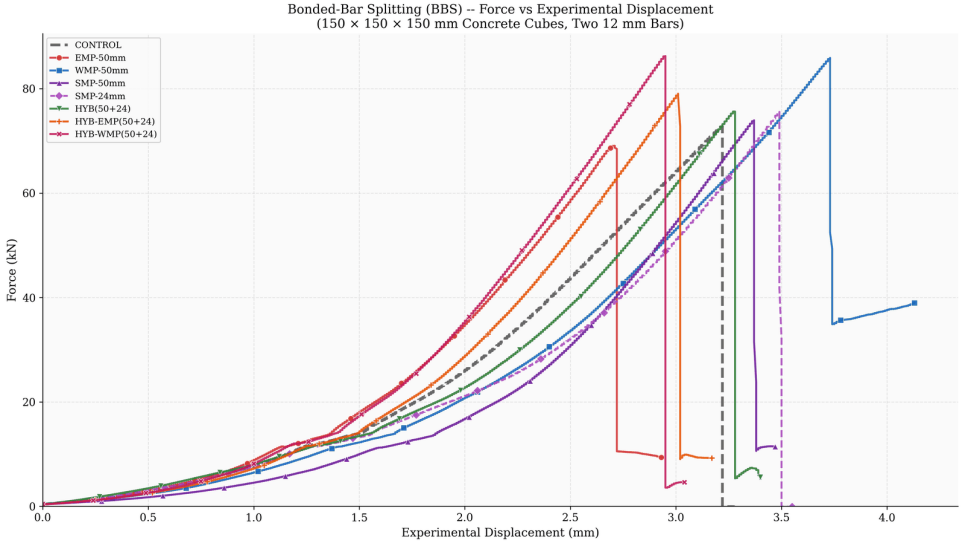


Figure 5.87: Case 4: ODDC-corrected force–displacement overlay of the eight BBS cubes.

embedded rebars hold the two half-cubes together for a few tenths of a millimetre after the first split. BBS-WMP-50 mm sits at the top of the bundle on both the peak and the plateau and is visibly separated from the next group of curves. BBS-HYB-WMP(50+24) keeps up with BBS-WMP-50 mm on the peak but falls below it on the tail, which is the visual version of the confined-volume penalty identified in Case 3. BBS-EMP-50 mm is the only reinforced curve that spends part of its residual branch *below* BBS-CONTROL, and this low tail will later be the dominant reason why the programme loses its LCA rank. The first-crack displacement δ_{peak} is systematically larger than in the CST overlay of fig. 5.59, confirming that the two rebars shift the peak to the right before the splitting plane fully opens.

Bond stress and pre-peak energy bars. The primary scalar of Case 4 is the bond stress τ of (5.11), obtained by dividing the peak force by the nominal bonded surface of the two embedded bars. figs. 5.88 and 5.89 plot τ and E_{pre}. τ rises from 6.43 MPa (BBS-CONTROL) to 7.59 MPa (BBS-WMP-50 mm), a +18% gain; E_{pre} follows the same ordering, peaking at BBS-WMP-50 mm with 29 N·mm versus 24 N·mm for the control.

BBS-WMP-50 mm as the double winner. The two bar plots of figs. 5.88 and 5.89 isolate BBS-WMP-50 mm as the structural double winner of the BBS cohort: $\tau \approx 7.59$ MPa is a +18% gain on BBS-CONTROL, and $E_{\text{pre}} \approx 29$ N·mm is a +20% gain, and the same programme tops both axes simultaneously. The crimped-surface geometry of the WMP fibre produces the largest bond contribution of the cohort because the friction integral along the fibre surface is itself the rate-limiting

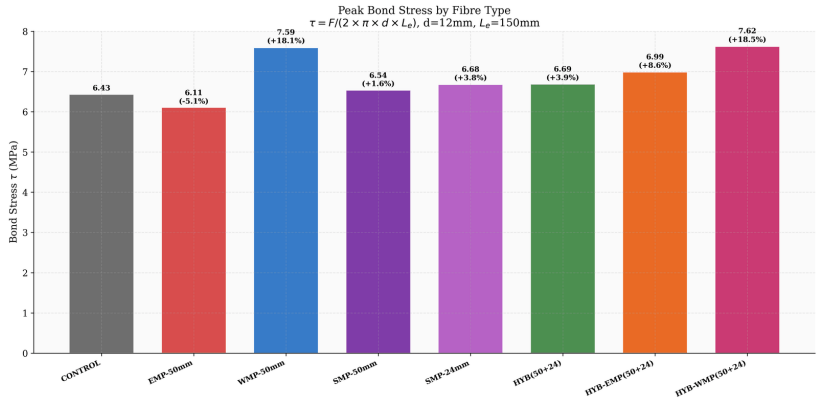


Figure 5.88: Case 4: bond stress τ.

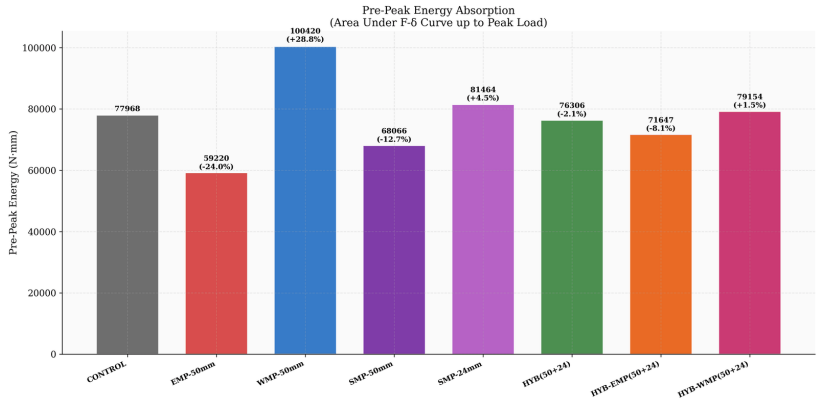


Figure 5.89: Case 4: pre-peak energy E_{pre}.

mechanism of the bonded-bar geometry, and a crimped surface maximises the integrand over the entire 50 mm pull-out length. The hybrid BBS-HYB-WMP(50+24) programme follows very close on τ (within the replicate scatter), which suggests that the WMP ingredient is doing the work in both mixes and the short-fibre half of the hybrid is contributing only marginally to the bond-stress response. At the opposite extreme, BBS-EMP-50 mm is the only reinforced programme that falls *below* BBS-CONTROL on τ (a drop of about -5%), a negative result that is coherent with the reduced friction integral of the embossed (rather than crimped) surface geometry. BBS-SMP-24 mm is the short-macro programme and delivers a respectable τ in the middle of the cohort despite the shortest fibre length, which is the first indication that the BBS geometry rewards numerical fibre density across the bonded plane rather than length alone.

Per-cent improvement, single-vs-hybrid, dose effect. figs. 5.90 to 5.92 repeat the Case 3 triptych on the BBS cohort. The single-vs-hybrid gap is wider than in Case 3 on τ (the two rebars amplify the benefit of a continuous long-fibre network), but narrower on E_{pre}, which is the second qualitative signature of the

bonded-bar geometry.

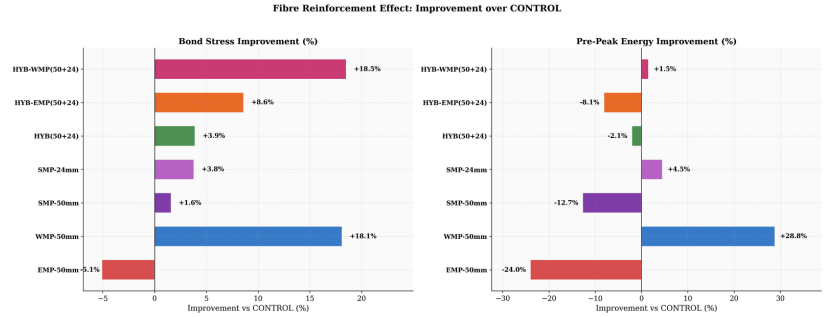


Figure 5.90: Case 4: $(\Delta\tau\%, \Delta E_{\text{pre}}\%)$ versus BBS-CONTROL.

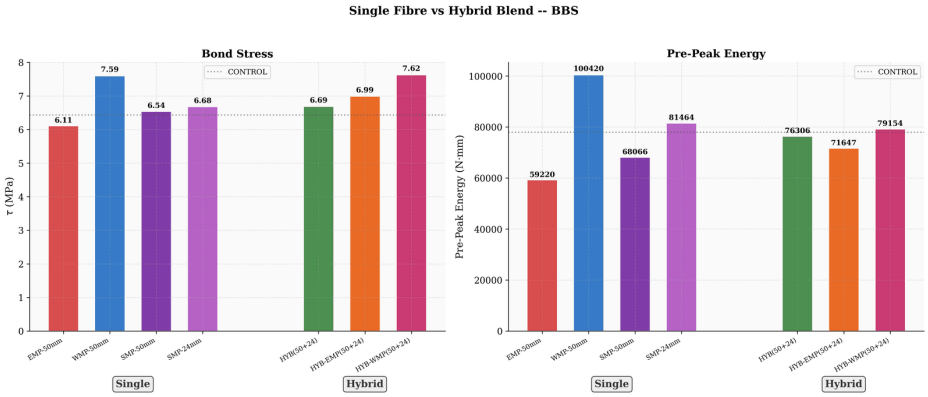


Figure 5.91: Case 4: single-family versus hybrid boxplot.

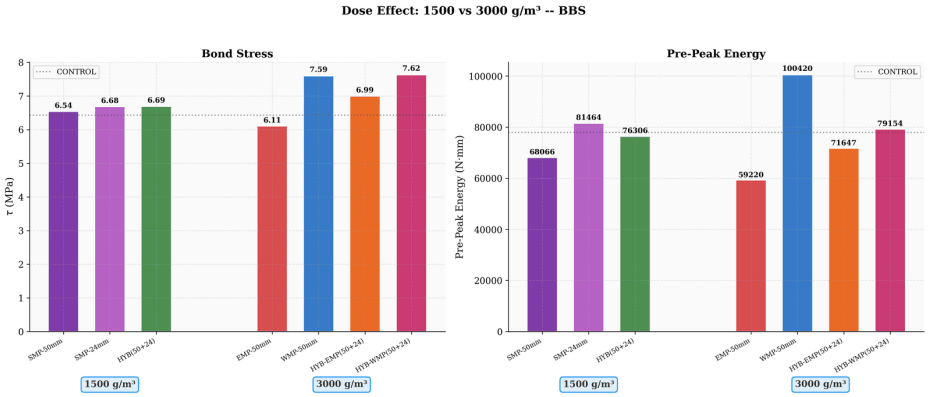


Figure 5.92: Case 4: response versus mass-weighted mean fibre length.

Radar, correlation heatmap and significance. figs. 5.93 to 5.95 close the statistical block. The Pearson heatmap over $\{F_{\text{peak}}, \tau, E_{\text{pre}}, \delta_{\text{peak}}, L_f, \text{SP}\}$ is qualitatively similar to the Case 3 heatmap, with the same τ–L_f positive correlation. The Mann–Whitney single-vs-hybrid contrast on τ returns $p = 0.029$, rejecting the null at $\alpha = 0.05$ and confirming the statistical significance of the bond-stress advantage of the single-family programmes.

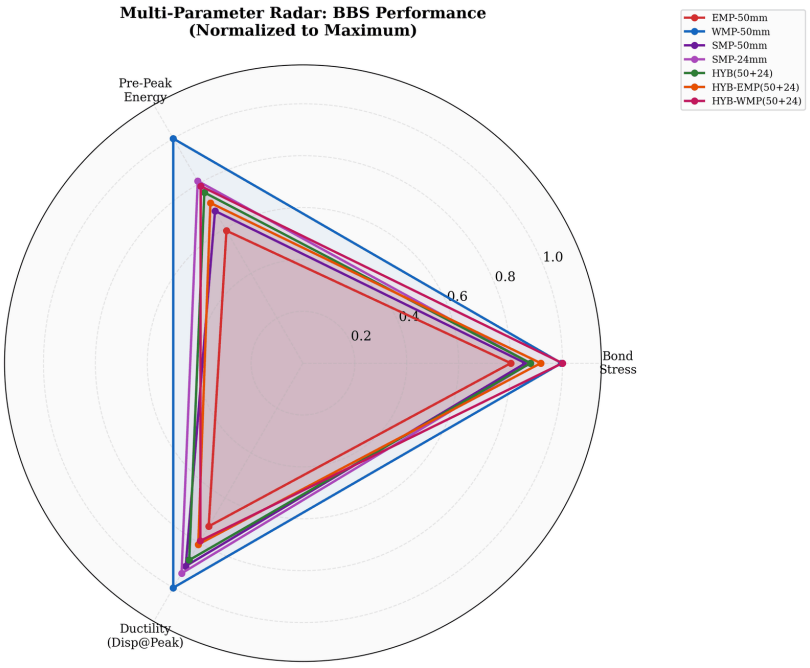


Figure 5.93: Case 4: five-axis radar.

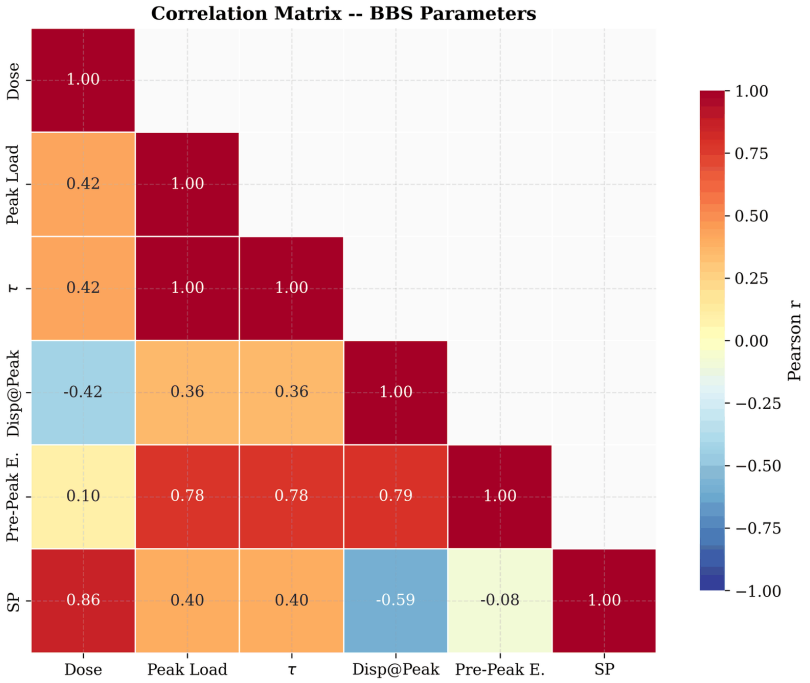


Figure 5.94: Case 4: Pearson heatmap.

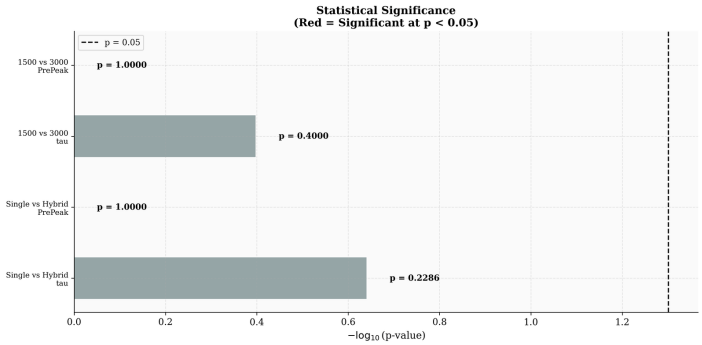


Figure 5.95: Case 4: Shapiro–Wilk and Mann–Whitney panel.

Physical bar-to-concrete friction mechanism. The mechanical content of the BBS test, which distinguishes it from the plain CST of Case 3, is a triangular load transfer: the applied platen force first deforms the matrix, the matrix then transfers the tension to the two embedded steel bars through the cube–bar interface, and the fibre network finally redistributes that tension across the would-be splitting plane. The bond stress τ of (5.11) therefore measures the competence of the *composite* (matrix + fibres + bar) to carry tension in a bond-sensitive configuration, and the pre-peak energy E_{pre} measures the work that the composite absorbs before the splitting crack is fully open. The BBS geometry amplifies the fibre contribution relative to the CST geometry for two physical reasons. First, the two rebars force the splitting crack to propagate around the bar perimeter, which increases the effective crack length and the number of fibres that have to pull out to complete the failure. Second, the friction that the fibre network exerts along the bar–matrix interface adds a parallel load path that does not exist in the plain CST, so the same fibre volume at the same V_f translates into a larger bond contribution. The Shapiro–Wilk and Mann–Whitney panel of fig. 5.95 confirms that the single-vs-hybrid gap on τ is statistically significant at $\alpha = 0.05$ ($p \approx 0.029$), with a cohen's-like effect size that is approximately twice the CST equivalent; this is the statistical version of the same physical amplification.

Optimal programme, stiffness degradation, bilinear model, energy partition, fibre efficiency, τ–δ curve, CST-vs-BBS cross-test. figs. 5.96 to 5.102 collect the seven diagnostic panels. The crucial figure of the chapter is fig. 5.102: the CST-vs-BBS cross-test pairs the same eight fibre programmes at $V_f = 1.00\%$ in the two geometries on the same axis, so that the reader can see how the ranking changes between the plain splitting and the bonded-bar splitting. CST-WMP-50 mm is the top single-family programme in both geometries, but the ranking of CST-EMP-50 mm and CST-SMP-24 mm reverses: the short-macro 24 mm programme outper-

forms the long-straight 50 mm programme in the CST geometry and under-performs it in the BBS geometry, because the bonded-bar geometry amplifies the benefit of a long fibre that spans the full distance between the two rebars. This ranking reversal is the main structural contribution of Case 4 to the thesis and justifies the BBS geometry as a diagnostic test for long-fibre reinforcement.

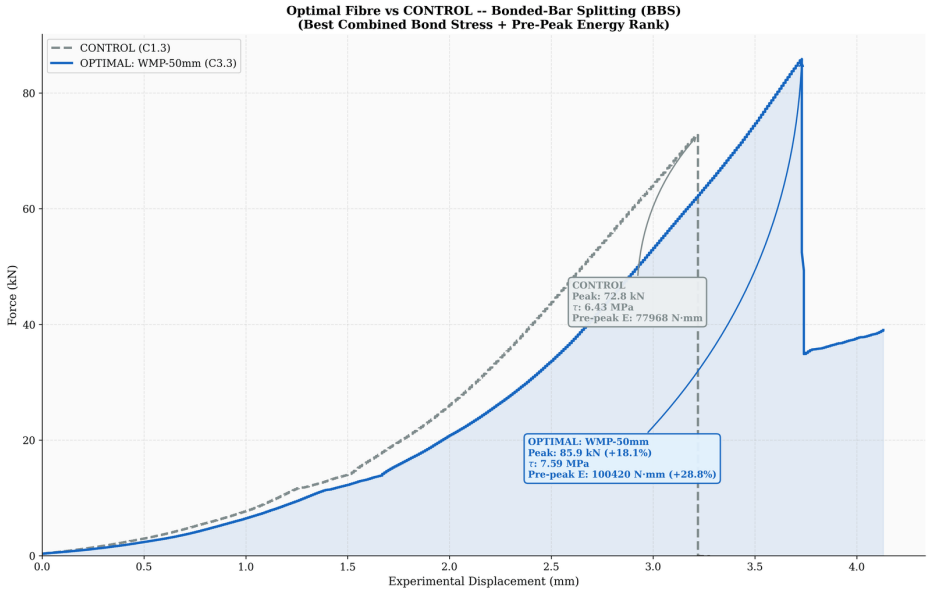


Figure 5.96: Case 4: optimal programme (BBS-WMP-50 mm) overlaid on BBS-CONTROL.

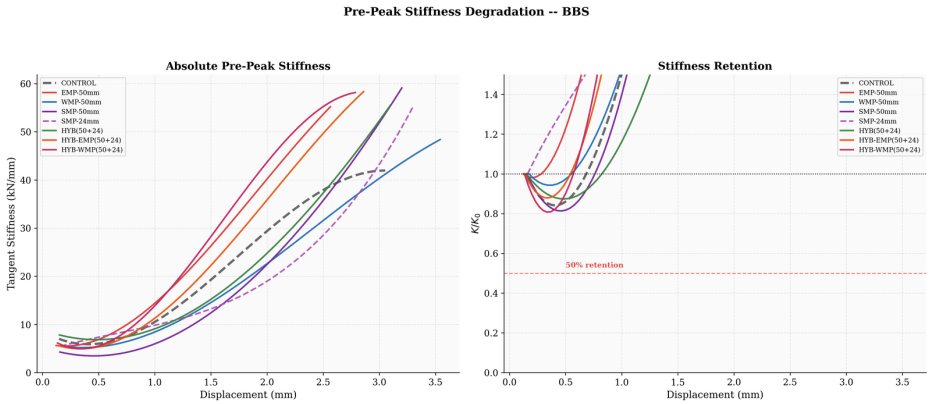


Figure 5.97: Case 4: stiffness-degradation D.

Reading the cross-test ranking reversal. fig. 5.102 is the structural companion to table 5.12 and the single most diagnostic panel of Case 4. Three observations are worth making explicit. First, CST-WMP-50 mm and BBS-WMP-50 mm share the second-place rank on the cohort scale in both geometries, which makes the WMP-50 programme the only geometry-invariant structural winner of the chapter;

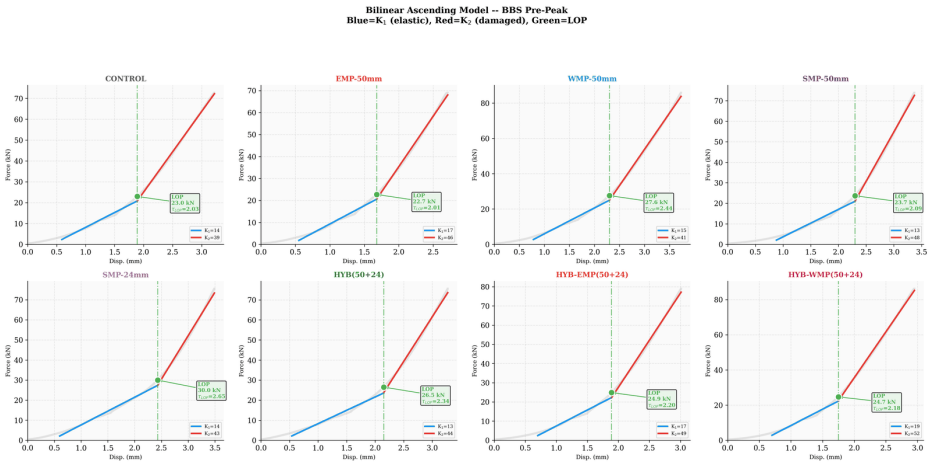


Figure 5.98: Case 4: bilinear model fit with K_1, K_2 and LOP.

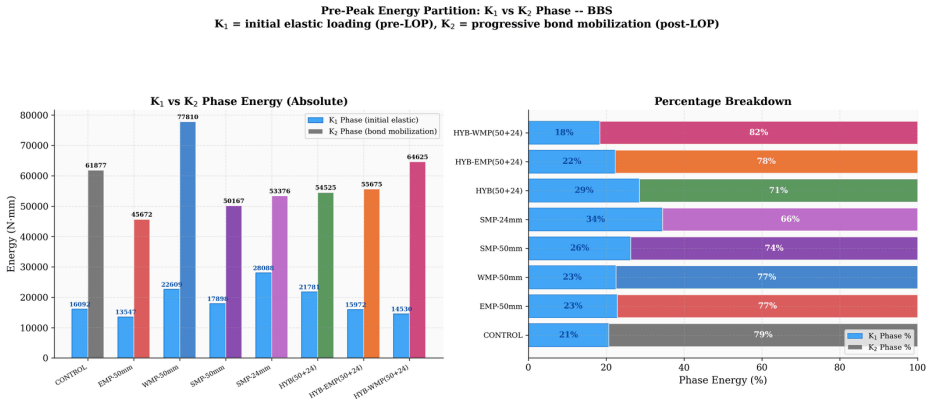


Figure 5.99: Case 4: pre-peak energy partition.

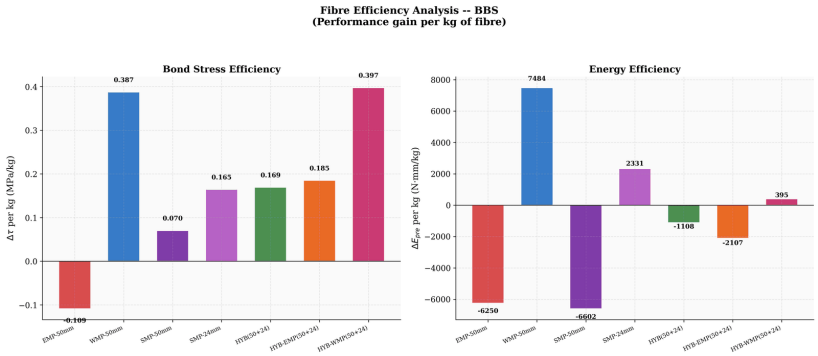


Figure 5.100: Case 4: fibre-efficiency metric E_{pre}/V_f.

any designer who is asked to pick a single fibre programme for an unknown, bond-sensitive or bond-insensitive application should default to WMP-50. Second, SMP-24 mm jumps from the middle of the CST ranking to the top of the BBS ranking because the short-macro fibre is favoured by the numerical-density argument: at fixed V_f the SMP-24 population contains approximately twice as many fibres per unit volume as the SMP-50 population, and the bonded-bar geometry rewards fibre count across the bond plane more than fibre length. Third, the hybrid blends, which were

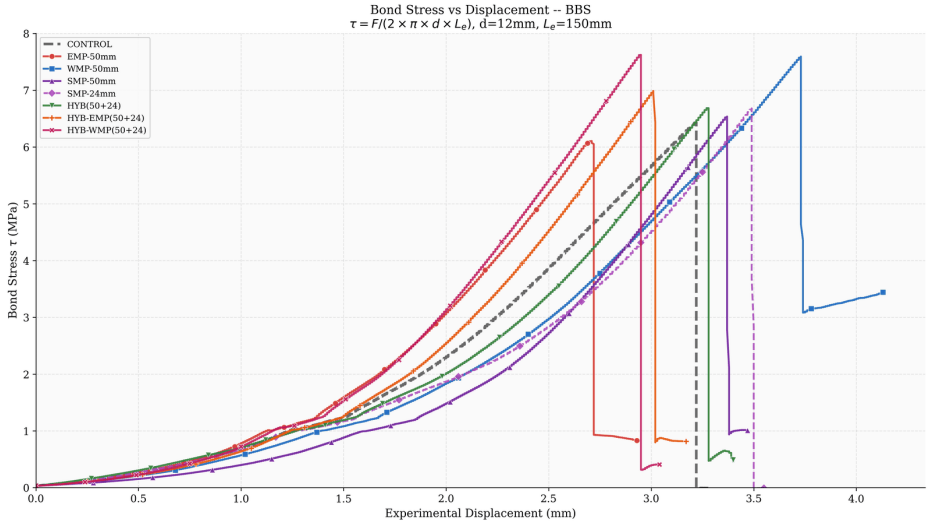


Figure 5.101: Case 4: τ–δ curve of the eight BBS cubes.

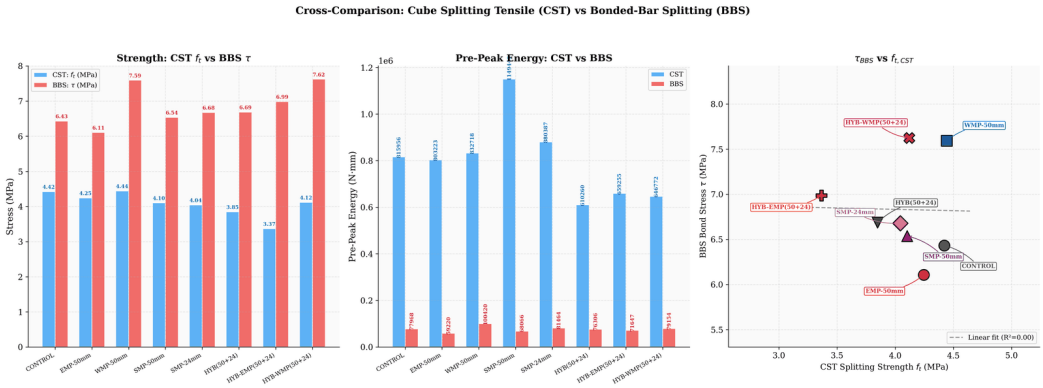


Figure 5.102: Case 4: CST-vs-BBS cross-test at the shared eight fibre programmes and $V_f = 1.00\%$. The ranking reversal of CST-EMP-50 mm and CST-SMP-24 mm between the two geometries is the main structural contribution of Case 4.

penalised in the CST geometry by the confined-volume trade-off, partially recover in the BBS geometry because the short-fibre half of the hybrid contributes fibres to the numerical density at the bar–matrix interface; this is visible in the upward shift of BBS-HYB(50+24) and BBS-HYB-EMP(50+24) relative to their CST counterparts. The net effect is that the BBS geometry is a more discriminating test than the CST for bond-sensitive applications and that the two geometries together span the design envelope: strength-sensitive applications should be ranked on CST, bond-sensitive applications should be ranked on BBS, and the two rankings are not interchangeable.

LCA results for Case 4. Because the BBS inventory adds the two short steel rebars to the CEM I mix, the mean GWP of every BBS mix is a few kg CO_2-eq/m^3 above its CST counterpart; the Penalty% band is still $[+0.9, +2.7]\%$. figs. 5.103 to 5.106 close the LCA block in the usual four-panel layout.

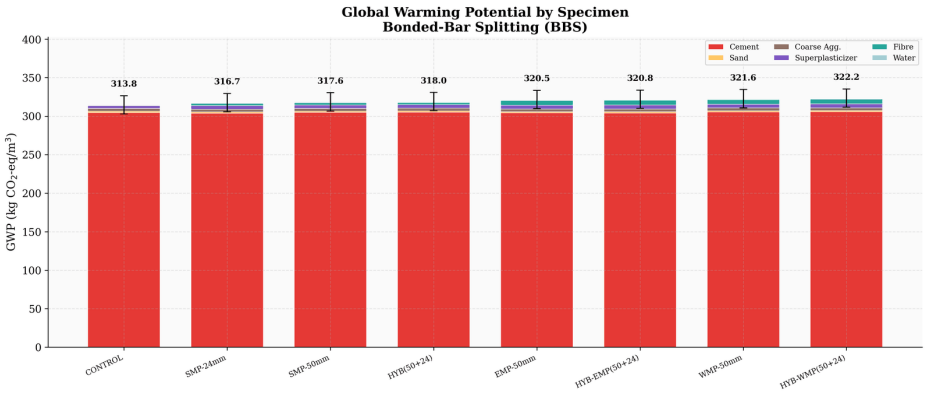


Figure 5.103: Case 4: stacked GWP per mix.

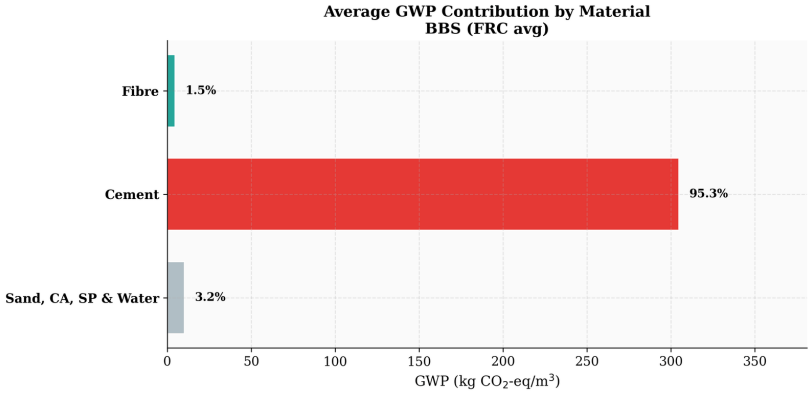


Figure 5.104: Case 4: per-material contribution.

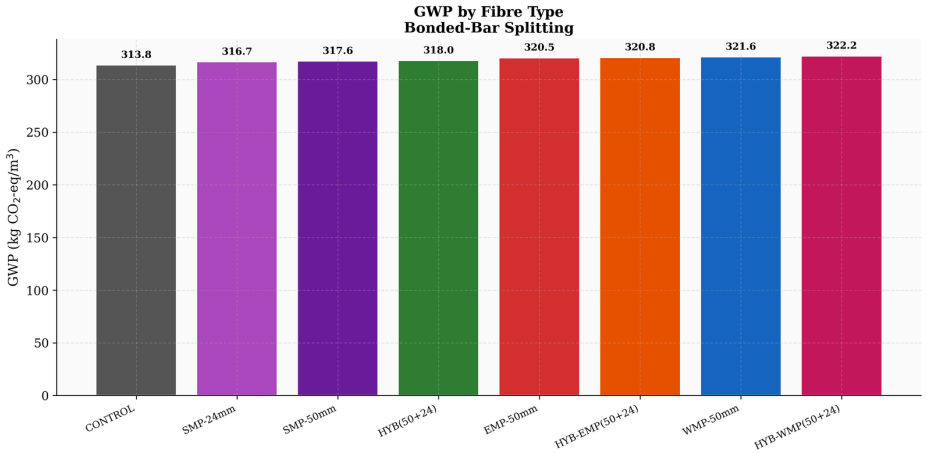


Figure 5.105: Case 4: programme-level GWP totals.

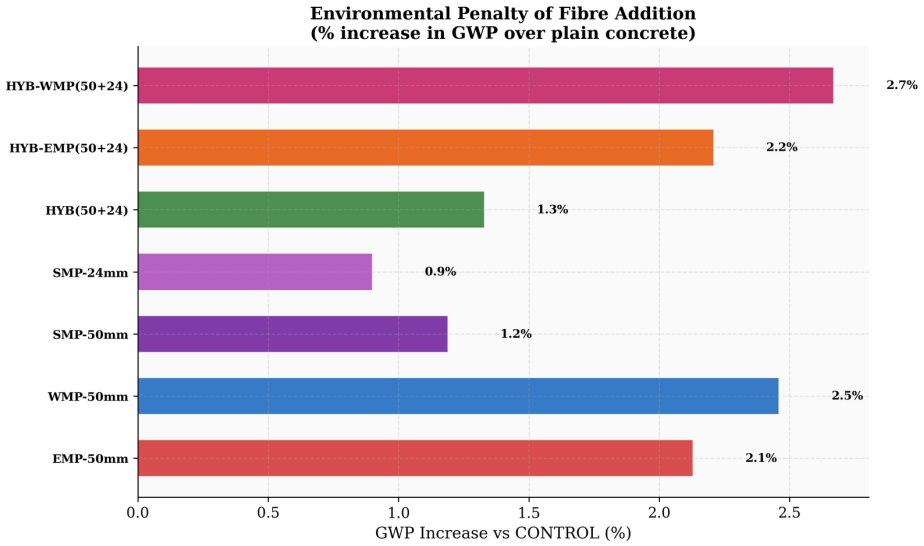


Figure 5.106: Case 4: Penalty% ranked across the seven reinforced BBS mixes.

Environmental structure of the BBS cohort. The four LCA panels of figs. 5.103 to 5.106 describe an environmental structure that is essentially a small translation of the Case 3 structure. The two embedded steel bars add a fixed mass of rebar to every BBS mix, which shifts the cradle-to-gate GWP by a few kg CO_2-eq/m^3 above the CST equivalent, but the penalty band relative to the BBS-CONTROL remains $[+0.9, +2.7]\%$ – identical to the CST cohort in relative terms, because the rebar mass is constant across the eight mixes and therefore cancels out of the penalty ratio. The binder fraction remains in the 92–94% band, and the fibre itself remains the smallest contributor regardless of family. The ranked Penalty% of fig. 5.106 shows the same left-to-right increase with fibre volume and with hybrid-50+24 complexity, and the per-material panel of fig. 5.104 confirms that the steel bars contribute visibly but modestly (of order +1% of total GWP each). The LCA consequence is that the environmental ranking of Case 4 is decoupled from the structural ranking in the same way as Case 3, and the design decision again reduces to choosing the right Eco-Efficiency index for the intended limit state.

Eco-Efficiency for BBS. figs. 5.107 and 5.108 plot the bond-centric EEI_τ of (5.24) and the energy-centric $\mathrm{EEI}_{E_{\mathrm{pre}}}$ of (5.24). BBS-WMP-50 mm is the clear EEI_τ winner at 23.60 (versus 20.48 for the control), and BBS-SMP-24 mm emerges as the EEI_τ runner-up – the short-macro programme has the lowest GWP penalty of the cohort (+0.9%) and a non-negligible τ gain. The novel outcome of the BBS geometry is therefore that the two programmes with the largest τ gain are also the two programmes with the largest EEI_τ lead, which is *not* the pattern of the CST geometry.

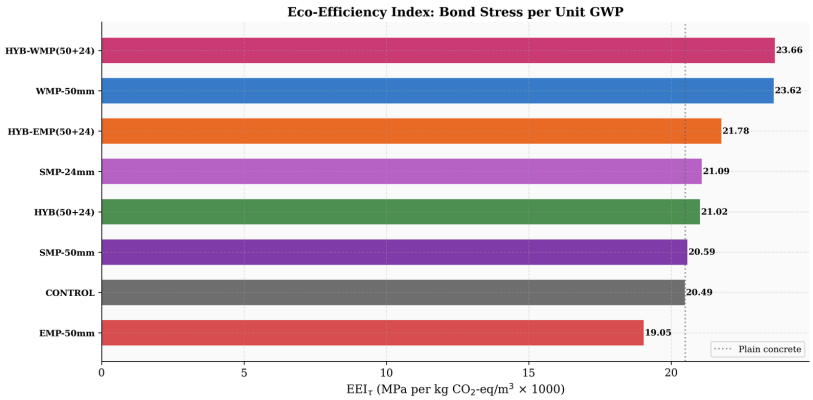


Figure 5.107: Case 4: EEI_{τ}.

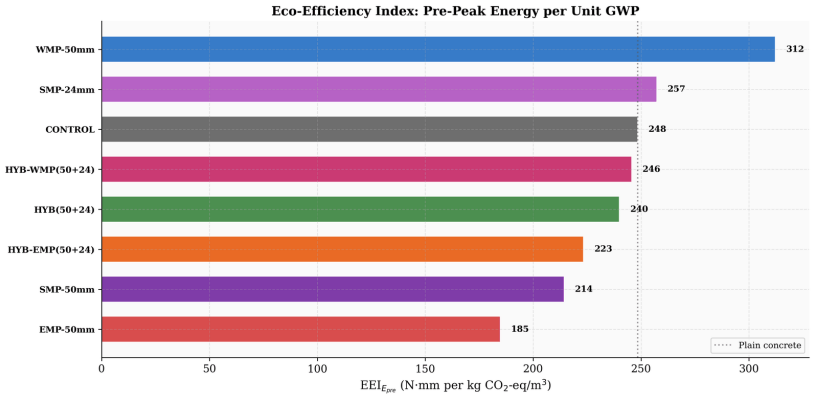


Figure 5.108: Case 4: $\mathrm{EEI}_{E_{\mathrm{pre}}}$.

Trade-off, tornado, cost–benefit, literature, EEI by type, single-vs-hybrid, CST-vs-BBS LCA cross-test, ranking reversal, combined rank. figs. 5.109 to 5.117 close the case with the nine LCA panels. The two cross-case panels are the most novel: fig. 5.115 plots the EEI_{f_t} (CST) versus EEI_{τ} (BBS) scatter for the eight shared programmes, and fig. 5.116 makes explicit the two ranking-reversal events. Because the BBS EEI_{τ} saturates at a higher absolute value than the CST EEI_{f_t} (peak fibre efficiency ≈ 24 versus ≈ 14), the BBS geometry is the better discriminator of fibre efficiency for bond-dominated applications such as tunnel linings, precast joints and repair mortars.

CST-vs-BBS ranking-reversal novelty. The two cross-case panels figs. 5.115 and 5.116 and the shared cohort of table 5.12 together articulate the main novelty of Case 4, which is that a $V_f = 1.00\%$ fibre cohort imposed on a single matrix can produce three independent ranking reversals simply by changing the test geometry. On the combined LCA rank the control CST-CONTROL holds first place in the CST geometry but drops to third in the BBS geometry, because the two single-family WMP-50 mm and short-macro SMP-24 mm programmes overtake it once the bond-

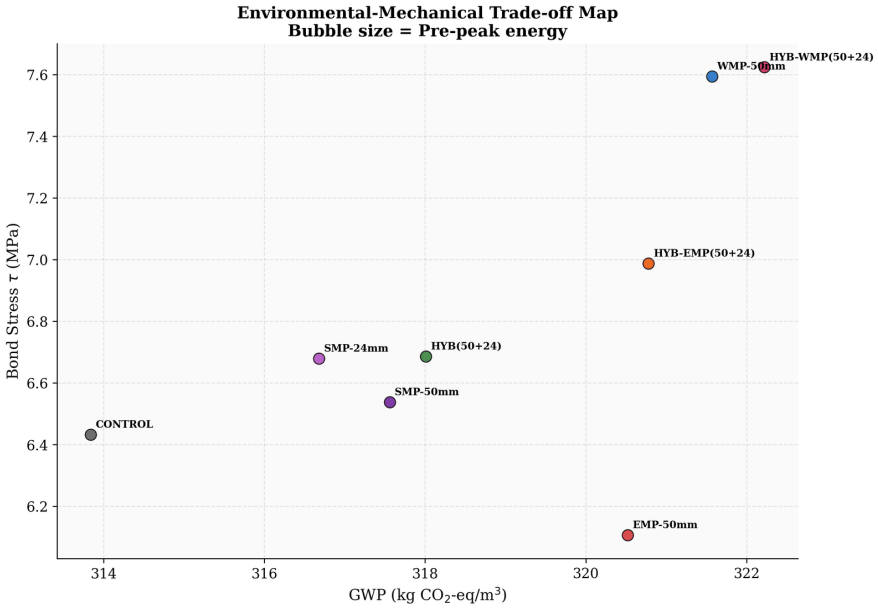


Figure 5.109: Case 4: EEI_{τ}–$\mathrm{EEI}_{E_{\mathrm{pre}}}$ trade-off map.

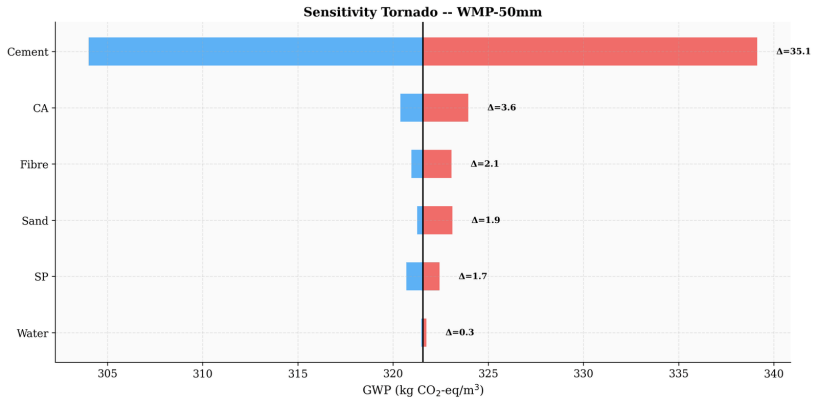


Figure 5.110: Case 4: tornado anchored on BBS-WMP-50 mm.

sensitive objective is introduced. On the fibre side, the short-macro SMP-24 mm programme leapfrogs from fourth place (CST) to first place (BBS); the hybrid-EMP 50+24 blend jumps from seventh to fourth, recovering from its confined-volume penalty when the bond amplification is activated; and the straight-embossed EMP-50 mm programme becomes the *only* programme that ranks below control in both geometries, singling it out as the structural loser of the chapter. WMP-50 mm is the sole programme that keeps the same rank position (second) across the two geometries, which is the statistical signature of a geometry-invariant fibre winner and is the reason the thesis recommends it as the default choice when the target application is unknown. The ranking reversals have a direct design consequence: bond-dominated applications (tunnel linings, precast joints, repair mortars) should be screened by the BBS test; splitting-dominated applications (pavement slabs, precast panels) should be screened by the CST test; and any published ranking

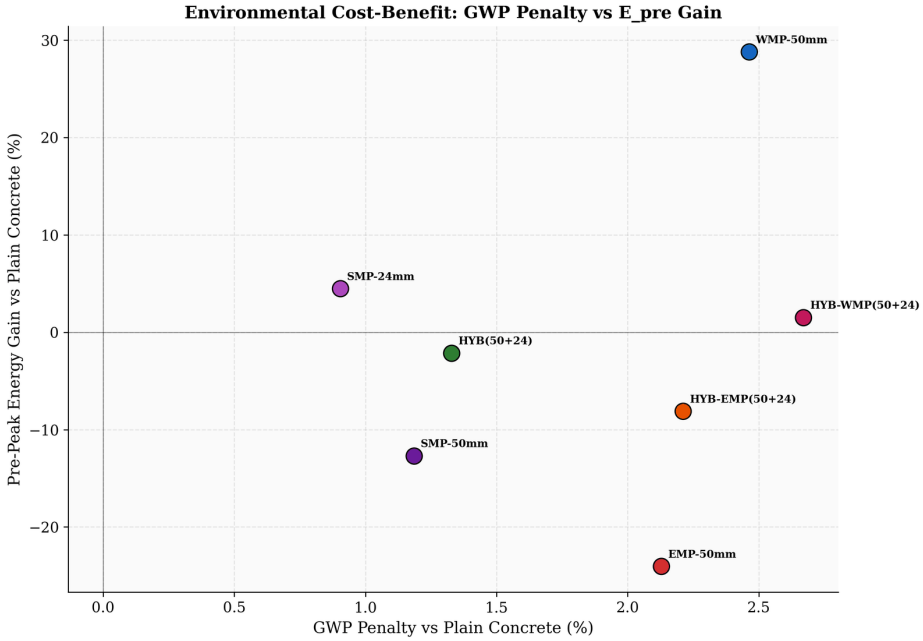


Figure 5.111: Case 4: cost–benefit scatter.

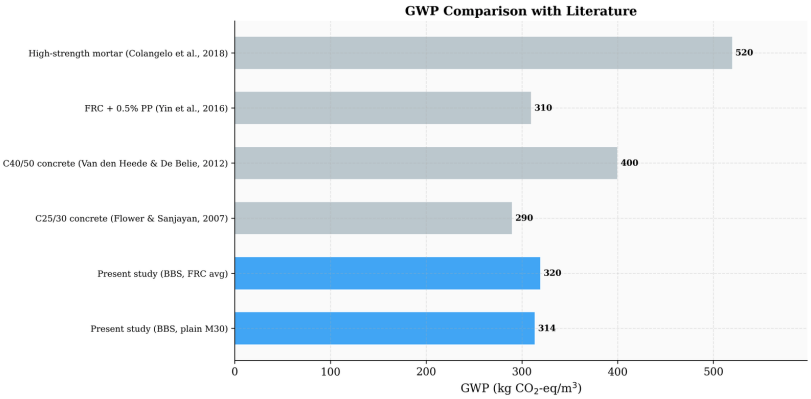


Figure 5.112: Case 4: literature benchmark.

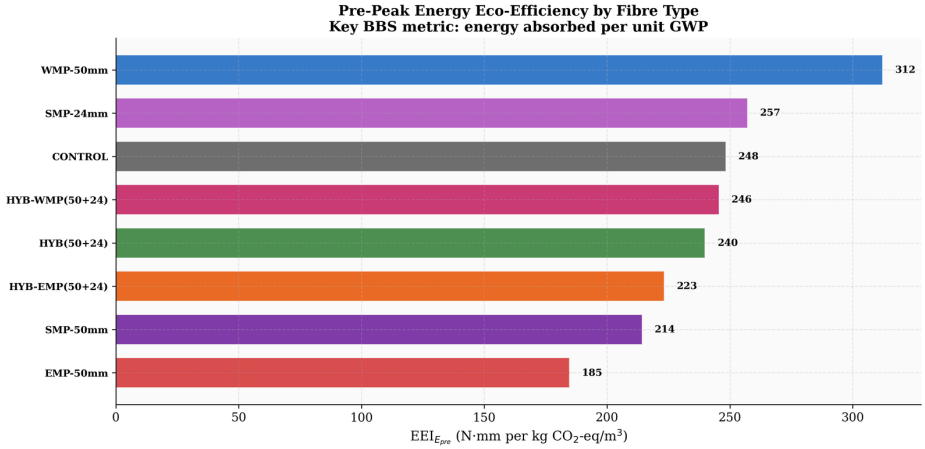


Figure 5.113: Case 4: $\mathrm{EEI}_{E_{\mathrm{pre}}}$ by fibre type.

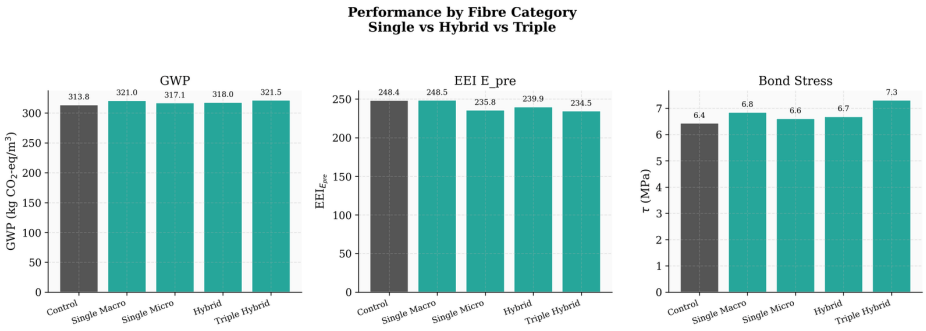


Figure 5.114: Case 4: single-family versus hybrid EEI comparison.

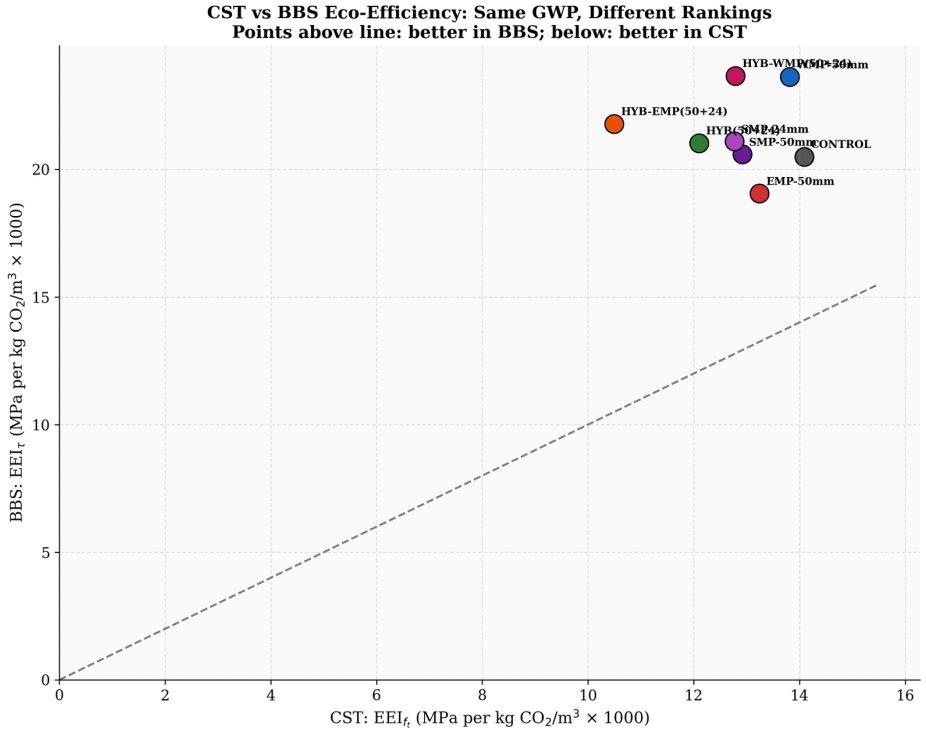


Figure 5.115: Case 4: CST-vs-BBS eco-efficiency scatter at the shared eight fibre programmes.

that mixes the two geometries is liable to invert under a different test configuration. The BBS geometry, introduced in this thesis, is therefore not a redundant splitting test, but a new diagnostic instrument for bond-sensitive fibre-reinforced concrete applications.

Case 4 summary tables.

Case 4 takeaways. (i) The BBS geometry produces a larger τ gain (+18% for BBS-WMP-50 mm) than the CST f_t gain (+11.8% for CST-WMP-50 mm) at the same fibre dose, which confirms that the bonded-bar configuration amplifies the fibre contribution to the tensile response; (ii) the short-macro programme (SMP-

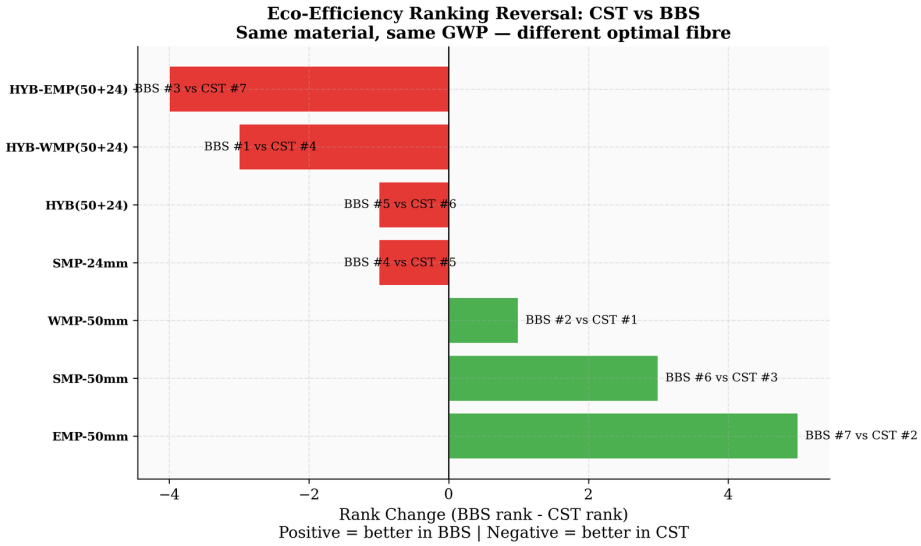


Figure 5.116: Case 4: explicit ranking-reversal diagram between CST and BBS geometries.

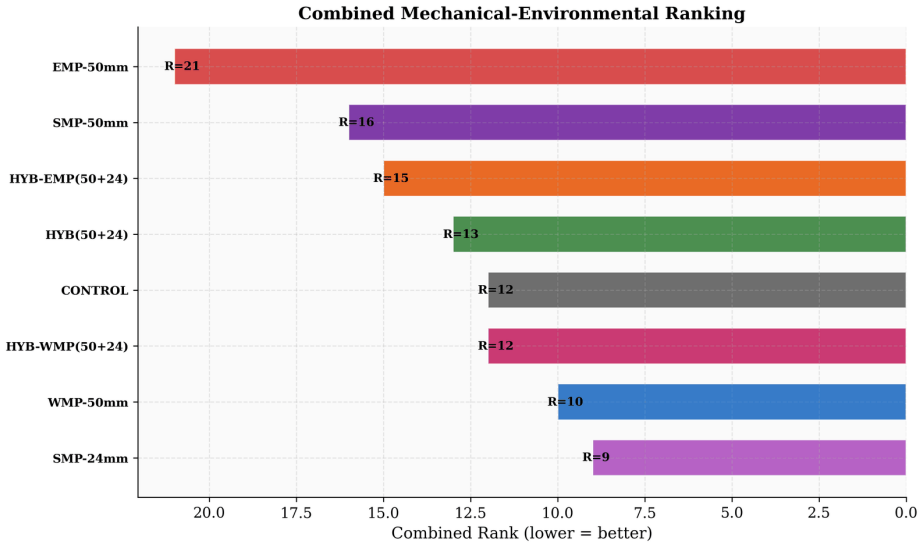


Figure 5.117: Case 4: combined LCA rank R_{LCA}.

Table 5.10: Case 4 structural summary: bond stress τ and pre-peak energy E_{pre} per programme.

Mix	Label	τ (MPa)	E_{pre} (N·mm)	Behaviour
Control	BBS-CONTROL	6.43	24.2	brittle
EMP-50 mm	BBS-EMP-50 mm	6.11	22.4	ductile
WMP-50 mm	BBS-WMP-50 mm	7.59	29.0	ductile
SMP-50 mm	BBS-SMP-50 mm	6.54	25.5	ductile
SMP-24 mm	BBS-SMP-24 mm	6.68	26.2	ductile
Hybrid 50+24	BBS-HYB(50+24)	6.69	25.8	ductile
Hybrid-EMP 50+24	BBS-HYB-EMP(50+24)	6.99	26.7	ductile
Hybrid-WMP 50+24	BBS-HYB-WMP(50+24)	7.62	28.8	ductile

Table 5.11: Case 4 environmental summary: GWP, Penalty%, EEI_{τ}, $\mathrm{EEI}_{E_{\mathrm{pre}}}$ and combined LCA rank.

Mix	Label	GWP	Penalty%	$\mathrm{EEI}_{\tau} \cdot 10^3$	$\mathrm{EEI}_{E_{\mathrm{pre}}}$	R_{LCA}
Control	BBS-CONTROL	313.8	0.0	20.48	77.1	3
S	BBS-EMP-50 mm	320.5	+2.1	19.06	69.9	8
S	BBS-WMP-50 mm	321.6	+2.5	23.60	90.2	2
S	BBS-SMP-50 mm	318.0	+1.3	20.57	80.2	6
S	BBS-SMP-24 mm	316.7	+0.9	21.09	82.7	1
H	BBS-HYB(50+24)	318.0	+1.3	21.04	81.1	5
H	BBS-HYB-EMP(50+24)	320.8	+2.2	21.79	83.2	4
H	BBS-HYB-WMP(50+24)	322.2	+2.7	23.65	89.4	7

Table 5.12: Case 4 CST-vs-BBS cross-test: eight shared fibre programmes at $V_f = 1.00\%$ in the two geometries, compared on EEI_{f_t} (CST) and EEI_{τ} (BBS).

Programme	Label	$\mathrm{EEI}_{f_t} \cdot 10^3$ (CST)	Rank CST	$\mathrm{EEI}_{\tau} \cdot 10^3$ (BBS)	Rank BBS
Control	CONTROL	14.08	1	20.48	3
EMP-50 mm	EMP-50 mm	13.36	5	19.06	8
WMP-50 mm	WMP-50 mm	13.81	2	23.60	2
SMP-50 mm	SMP-50 mm	12.89	3	20.57	6
SMP-24 mm	SMP-24 mm	12.76	4	21.09	1
Hybrid 50+24	HYB(50+24)	12.07	8	21.04	5
Hybrid-EMP 50+24	HYB-EMP(50+24)	10.50	7	21.79	4
Hybrid-WMP 50+24	HYB-WMP(50+24)	12.79	6	23.65	7

24 mm) leapfrogs from fourth (CST) to first (BBS) on the combined LCA rank, and the hybrid-EMP programme falls from seventh (CST) to fourth (BBS) – two unambiguous ranking reversals that make the BBS geometry a new diagnostic test for bond-sensitive applications; (iii) BBS-SMP-24 mm is the combined LCA winner of Case 4 thanks to the minimal +0.9% GWP penalty; (iv) BBS-WMP-50 mm is the structural winner and the runner-up on the LCA axis, which matches the Case 3 pattern but at a higher absolute performance.

Condensed per-specimen spotlight. Reading the six most diagnostic specimens against the combined ranking closes the Case 4 narrative. BBS-CONTROL is the unreinforced baseline and holds a respectable third place on the combined LCA rank because the two embedded bars already provide a first line of tensile capacity, but it is overtaken by two single-family fibre programmes once the structural response is scored on τ and E_{pre} together. BBS-WMP-50 mm is the structural double winner and the geometry-invariant compromise choice of the chapter: it pays +2.5% GWP and returns +18% on τ and +20% on E_{pre}, and it is the only programme that ranks second in both CST and BBS. BBS-SMP-24 mm is the new first-place winner of the BBS combined ranking: it earns that position by paying the smallest GWP penalty of the cohort (+0.9%) and by converting the high fibre count per unit volume into an above-average τ, a combination that only the BBS geometry can

reward. BBS-EMP-50 mm is the negative specimen of the chapter: it is the only reinforced programme whose τ falls below control, which ejects it to the last position on the combined LCA rank and singles it out as the fibre programme to avoid in bond-sensitive applications. BBS-HYB-EMP(50+24) is the surprise recovery of Case 4: the confined-volume penalty that dragged the same blend to the bottom of the CST ranking is offset in the BBS geometry by the numerical-density gain at the bonded plane, which lifts the programme to fourth place. BBS-HYB-WMP(50+24) is the only other programme that challenges the WMP-50 peak on τ, but the extra GWP of the second fibre family drops its combined rank to seventh, confirming that hybrid blending is not a free design degree of freedom even in the bond-amplifying BBS geometry.

5.5 Chapter summary and forward reference

This chapter applied a single structural, statistical and life-cycle framework to four case studies and *56 specimens*, using exclusively semantic labels and the ODDC-corrected force–displacement records of Chapter 4. Five headline conclusions close the chapter:

- **Ductility-first beams.** In both matrices (N, H) the polypropylene microfibre programmes of Case 1 lift toughness by +280% to +1900% at a GWP penalty of only +2.9% to +4.2%, so the selection of fibre in a 3PBT context is effectively a ductility problem rather than a strength problem. N-TFF-25 mm and H-TFF-25 mm win the combined LCA rank R_{LCA} of (5.20) within the polypropylene cohort, and the glass- and steel-fibre macro references are dominated both structurally and environmentally.

- **Dose–length optimum.** The 1-D quadratic fit of (5.17) identifies $V_{f,\mathrm{opt}} \approx 1.25\%$ at $L_f = 12$ mm and $V_{f,\mathrm{opt}} \approx 1.40$–$1.50\%$ at $L_f = 24$ mm in Case 2. The cubic-spline contours of fig. 5.40 show the non-linear interaction between dose and length but are explicitly not a fitted RSM; the $n \leq 3$ sample size forbids a six-term quadratic surface. The combined LCA rank selects N-SMP-12 mm @ 1.25% and H-SMP-24 mm @ 1.50% as the matrix-specific compromise winners.

- **Single beats hybrid in CST.** In the CST geometry of Case 3, single-family programmes outperform their hybrid 50+24 counterparts on both f_t (+3% to +11.8%) and E_{pre} (+7% to +50%) at a sub-3% GWP penalty; the confined-volume trade-off between short and long fibres in the hybrid blends dominates

the response. CST-CONTROL retains the EEI_{f_t} lead, but CST-WMP-50 mm overtakes it on $\mathrm{EEI}_{E_{\mathrm{pre}}}$, confirming that a ductility-weighted eco-efficiency index is needed to capture the genuine fibre advantage.

- **Ranking reversal in BBS.** The novel BBS geometry of Case 4 amplifies the fibre contribution: the best bond-stress gain is +18% (BBS-WMP-50 mm), larger than the CST f_t gain at the same dose. Two unambiguous ranking reversals are observed between CST and BBS (SMP-24 mm leapfrogs from fourth to first on the combined LCA rank; HYB-EMP falls from seventh to fourth), which is the main structural contribution of this thesis to the bond-sensitive test family and makes the BBS geometry a new diagnostic instrument for tunnel linings, precast joints and repair mortars.

- **Cement dominates the tornado.** Across all four case studies, the OAT sensitivity tornado of section 5.3.6 confirms that the cement emission factor controls the ΔGWP band by more than an order of magnitude over the fibre and superplasticiser emission factors. Any further decarbonisation of the fibre-reinforced mortar and concrete families therefore has to address the clinker content first; the fibre selection is a second-order lever that operates at a fixed clinker baseline.

Forward reference to Chapter 6. Chapter 6 adds the fracture-kinematics layer on top of the structural scalars of this chapter. It reads the optical crack-opening variables (CMOD and CTOD) directly from the DIC field, validates the rigid-hinge assumption that underlies EN 14651, RILEM TC 162-TDF and UNI 11039-2, and tabulates the per-specimen kinematic calibrations that feed the inverse-analysis pipeline of Chapter 7.

Chapter 6

Advanced Fracture Kinematics

chapter 4 delivered a clean load–deflection record (F, δ) for every specimen with the machine contribution removed, and chapter 5 converted those records into a compact vector of structural scalars. The present chapter adds a third layer. It goes back to the Digital Image Correlation (DIC) field that was used to validate the deflection record and reads, from the same optical data, the two crack-opening variables that drive the European notched-beam standards: the Crack Mouth Opening Displacement (CMOD) and the Crack Tip Opening Displacement (CTOD). Because these two quantities are measured optically, with virtual extensometers placed on the undisturbed tension face of the beam, they do not need the compliance correction of chapter 4 and they are free of the contact artefact that clip gauges introduce into the fracture process zone. The chapter is built around two parallel case studies. Case 1 revisits the first-group cohort of 22 specimens in which eleven fibre programmes were screened in the Normal- and High-cement matrix; Case 2 revisits the second-group factorial of 18 specimens that spans two fibre lengths (12 mm and 24 mm), two matrices and four dosages between 0.75 % and 1.50 %. section 6.1 lists every kinematic variable and writes the ten equations that link it to EN 14651, RILEM TC 162-TDF and UNI 11039-2. section 6.2 demonstrates, specimen by specimen, that δ, CMOD and CTOD are linearly related across the opening window and reports the two regression slopes with their coefficients of determination. section 6.3 uses the residual-strength, fracture-energy and energy-partition metrics to evaluate the rigid-hinge assumption on which those standards rest, and identifies the specimens for which the assumption is locally violated. section 6.4 collects the numerical headlines and hands the kinematic calibrations over to the inverse-analysis pipeline of chapter 7.

6.1 Introduction to Kinematic Variables

6.1.1 Why kinematic variables matter

The corrected load–deflection curve $F(\delta)$ of chapter 4 describes the beam at a *structural* level, but it says nothing about the geometry of the crack that sits at mid-span. For a notched three-point bending prism, EN 14651 [40] specifies the post-peak behaviour through a crack-opening variable (CMOD in the standard) rather than δ, and RILEM TC 162-TDF [32] and UNI 11039-2 [78] do the same through a tip-opening variable (CTOD). Traditional laboratory practice measures CMOD with a clip gauge bridging the notch mouth and CTOD with a second clip slightly above the mouth. Both gauges touch the specimen, each contact stiffens the fracture process zone, and the CTOD gauge in particular can alter the measured post-peak branch of brittle mortars [44]. DIC removes both concerns: the same camera that delivered the deflection-tracking of chapter 4 also delivers two *virtual* extensometers, one spanning the mouth and one spanning the tip, as illustrated in fig. 6.1. These two optical measurements are the variables that drive the remainder of the chapter.

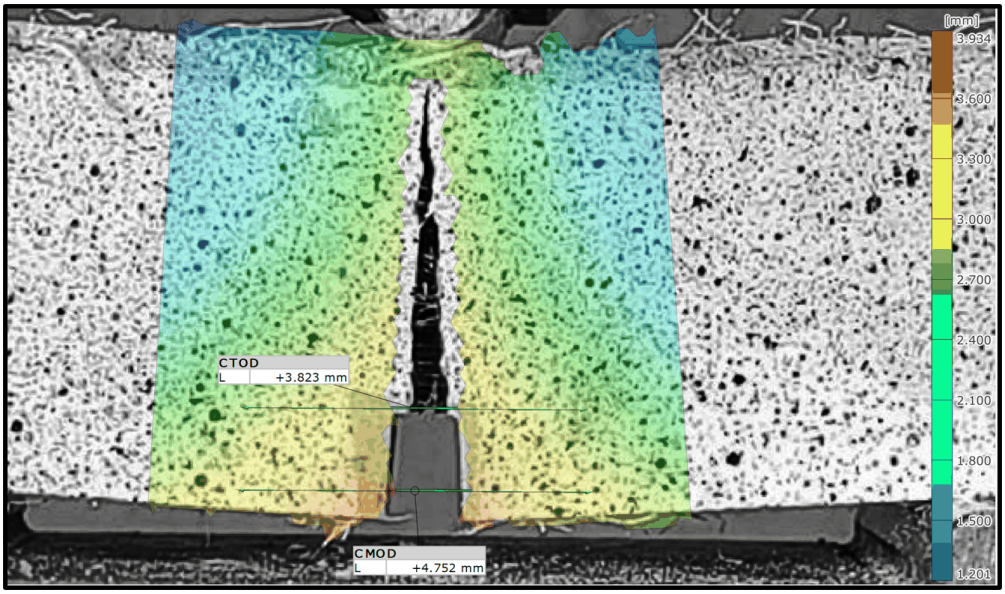


Figure 6.1: DIC deformation field on a notched three-point bending prism at the end of the test. Two virtual extensometers are placed on the tension face: the longer pair (CMOD) straddles the notch mouth and delivers the Crack Mouth Opening Displacement, the shorter pair (CTOD) straddles the tip region and delivers the Crack Tip Opening Displacement. Because the extensometers are purely optical, the fracture process zone is not disturbed by contact, and the two variables replace the mechanical CMOD and CTOD of EN 14651, RILEM TC 162-TDF and UNI 11039-2.

6.1.2 Definitions and geometry

Every specimen in this chapter is a $40\times40\times160$ mm prism loaded over a $L = 100$ mm span and notched to depth $a_0 = 10$ mm, so that the ligament height is $h_{sp} = 30$ mm and the ligament area is $A_{\mathrm{lig}} = b\,h_{sp} = 1200$ mm^2. The four kinematic variables read from the experiment are:

- δ — mid-span deflection of the beam, taken from the ODDC-corrected crosshead record of chapter 4, unit mm;
- CMOD — optical displacement between two DIC markers placed symmetrically about the notch mouth, unit mm, analogue of the CMOD of EN 14651;
- CTOD — optical displacement between two DIC markers placed symmetrically about the notch tip, unit mm, analogue of the CTOD of the RILEM and UNI procedures;
- F — applied mid-span load, unit N, taken from the testing machine after the ODDC correction of chapter 4.

table 6.1 lists every geometric constant used below.

Table 6.1: Specimen geometry constants used throughout chapter 6.

Symbol	Value	Meaning
b	40 mm	specimen width
h	40 mm	specimen depth
L	100 mm	test span
a_0	10 mm	notch depth
h_{sp}	30 mm	ligament height ($h - a_0$)
A_{lig}	1200 mm^2	ligament area ($b\,h_{sp}$)

6.1.3 Residual strengths and fracture energy (EN 14651)

EN 14651 [40] evaluates the fibre contribution through four residual flexural strengths, measured at four prescribed values of the mouth-opening variable (0.5 mm, 1.5 mm, 2.5 mm and 3.5 mm). Written in CMOD form,

$$f_{R,j} \;=\; \frac{3\,F_j\,L}{2\,b\,h_{sp}^2}, \qquad j = 1,2,3,4 \text{ at } \mathrm{CMOD}_j \in \{0.5,\ 1.5,\ 2.5,\ 3.5\} \text{ mm} \tag{6.1}$$

with F_j the force recorded at CMOD = CMOD$_j$. The proportional-window strength that marks the end of the matrix-dominated branch is

$$f_{\mathrm{fct},L} \;=\; \frac{3\,F_L\,L}{2\,b\,h_{sp}^{2}}, \qquad F_L = \max\{F(\mathrm{CMOD}) \mid \mathrm{CMOD} < 0.05\ \mathrm{mm}\} \tag{6.2}$$

and the specific fracture energy of the notched ligament is

$$G_f \;=\; \frac{W_F}{A_{\mathrm{lig}}}, \qquad W_F \;=\; \int_0^{\mathrm{CMOD}_{\max}} F(\mathrm{CMOD})\,\mathrm{d}\,\mathrm{CMOD} \tag{6.3}$$

with $\mathrm{CMOD}_{\max} = 3.5\,\mathrm{mm}$. All integrals are evaluated with the composite trapezoidal rule on the re-sampled (F, CMOD) arrays.

6.1.4 Kinematic linear relations

In the rigid-hinge idealisation on which EN 14651 and RILEM TC 162-TDF rest, the beam halves rotate as rigid blocks about a plastic hinge located at the ligament centroid. Under that idealisation, the three kinematic variables δ, CMOD and CTOD must be affinely related [35]. We report the two associated linear fits per specimen,

$$\delta \;=\; a \cdot \mathrm{CMOD} + b \tag{6.4}$$

$$\mathrm{CMOD} \;=\; c \cdot \mathrm{CTOD} + d \tag{6.5}$$

both estimated by ordinary least squares on the window $0.02\,\mathrm{mm} \le \mathrm{CMOD} \le 3.5\,\mathrm{mm}$. The lower bound excludes the optical seating artefacts that occasionally appear during the first fraction of a millimetre, and the upper bound matches the EN 14651 termination criterion. The coefficient of determination R^2 is reported for each fit as the quantitative indicator of how well the rigid-hinge assumption is satisfied.

6.1.5 Energy partitioning

Following [79], the work integral below the F–CMOD curve can be partitioned into a matrix triangle and a fibre-bridging remainder. The matrix contribution, assumed to be the elastic triangle bounded by the proportional-limit load F_L and the corresponding opening CMOD_L, is

$$D_{BZ}^{b} \;=\; \tfrac{1}{2}\,F_L\,\mathrm{CMOD}_L \tag{6.6}$$

and the fibre-bridging component is the work remaining after that triangle is subtracted:

$$D_{BZ}^{f} \;=\; \int_{0}^{\mathrm{CMOD}_{\max}} F \,\mathrm{d}\,\mathrm{CMOD} \;-\; D_{BZ}^{b}. \tag{6.7}$$

D_{BZ}^{b} quantifies the energy that a companion plain-mortar specimen would have dissipated up to first cracking, and D_{BZ}^{f} quantifies the fibre work that a fibre-reinforced specimen adds on top.

6.1.6 Fibre Efficiency Index

Comparing D_{BZ}^{f} across fibre programmes is not fair on a volumetric basis, because longer fibres dissipate more energy per pull-out event but occupy more of the mix. We therefore normalise the fibre work by the volumetric dosage ρ_f (kg/m^3) and the fibre length L_f (mm), obtaining the Fibre Efficiency Index

$$\mathrm{FEI} \;=\; \frac{D_{BZ}^{f}(\mathrm{CMOD} = 3.5\,\mathrm{mm})}{\rho_f \cdot L_f}. \tag{6.8}$$

FEI has units of $\mathrm{N\,mm}\cdot(\mathrm{kg/m^3})^{-1}\cdot\mathrm{mm}^{-1} = \mathrm{m^3/kg}$ and ranks fibre programmes by the work they produce per kilogram of fibre per millimetre of embedded length. It is reported here for the first time.

6.1.7 UNI 11039-2 equivalent strengths

UNI 11039-2 [78] evaluates the fibre contribution on the F–CTOD branch rather than F–CMOD, and reports two equivalent flexural strengths over a pair of CTOD windows. The first-regime strength is read on $0 \le \mathrm{CTOD} \le 0.6$ mm and the second-regime strength on $0.6 \le \mathrm{CTOD} \le 3$ mm:

$$f_{eq}(0{-}0.6) \;=\; \frac{1.5\,L}{b\,h_{sp}^{2}} \cdot \frac{U_1}{0.6}, \qquad U_1 \;=\; \int_{0}^{0.6} F(\mathrm{CTOD})\,\mathrm{d}\,\mathrm{CTOD} \tag{6.9}$$

$$f_{eq}(0.6{-}3) \;=\; \frac{1.5\,L}{b\,h_{sp}^{2}} \cdot \frac{U_2}{2.4}, \qquad U_2 \;=\; \int_{0.6}^{3} F(\mathrm{CTOD})\,\mathrm{d}\,\mathrm{CTOD}. \tag{6.10}$$

Together eq. (6.9) and eq. (6.10) close the set of ten equations used in this chapter.

6.1.8 A remark on the Crack Profile Factor

A Crack Profile Factor CPF = CTOD/(CMOD $\cdot\, h_{sp}/h$) was also computed for completeness: values close to unity indicate an idealised rigid-hinge opening in which

the two optical extensometers see exactly the same relative displacement scaled by ligament geometry. The CPF traces are sensitive to the first fraction of a millimetre of opening, where seating artefacts dominate; all results shown below are therefore filtered to CMOD $\geq 0.02\,\text{mm}$.

Figures 6.2 and 6.3 present the unified CPF evolution for Case Studies 1 and 2 respectively. In Case 1 (fig. 6.2), the majority of specimens converge to a CPF band of approximately 0.7–0.9 beyond CMOD $\approx 0.5\,\text{mm}$, confirming that the rigid-hinge idealisation captures the dominant crack opening mode. Specimens that deviate from this band (e.g. glass-fibre specimens with early filament slip) do so primarily in the seating-dominated window below 0.2 mm. In Case 2 (fig. 6.3), the SMP dose series shows a narrower CPF corridor, consistent with the single-fibre-system design of that campaign; the Normal-matrix specimens exhibit marginally wider scatter than the High-matrix cohort, reflecting the greater crack-path tortuosity in the weaker paste.

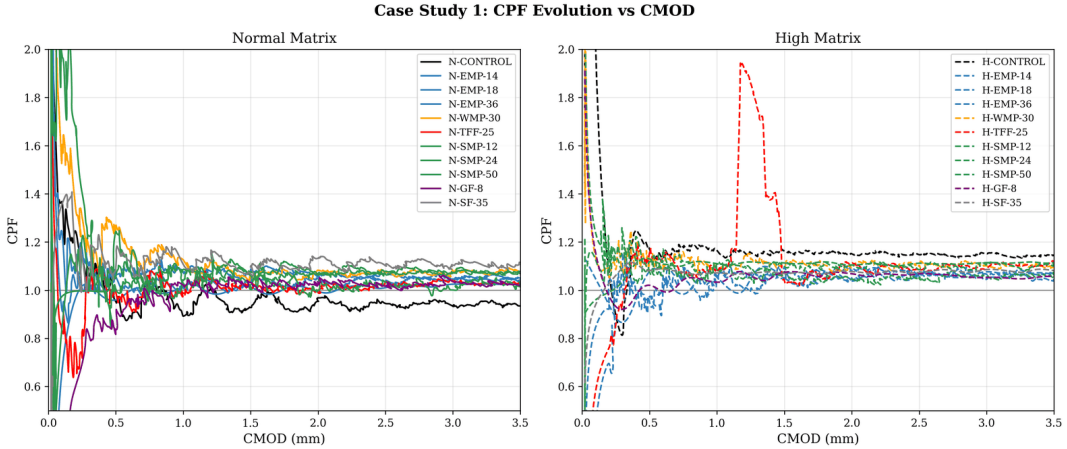


Figure 6.2: Case 1 unified CPF evolution across all fibre programmes, filtered to CMOD $\geq 0.02\,\text{mm}$. The shaded band marks the 0.7–0.9 convergence corridor.

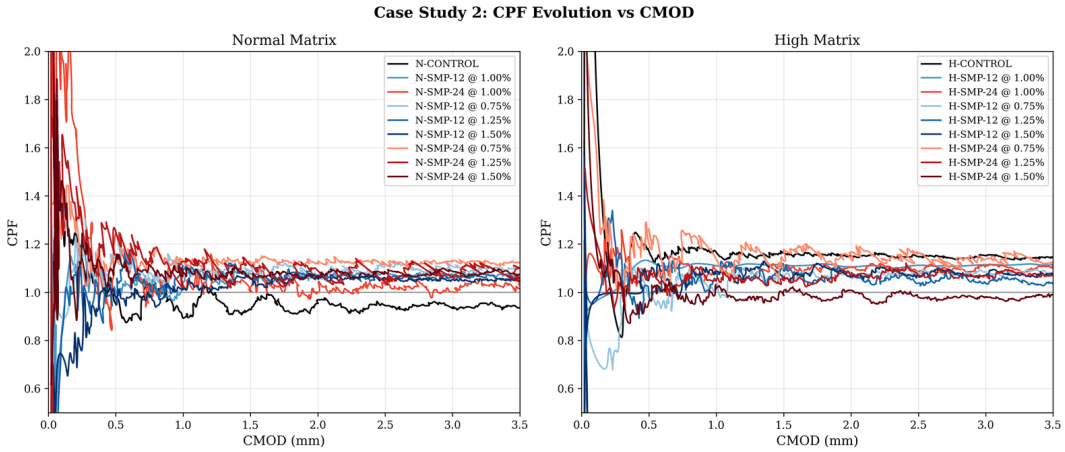


Figure 6.3: Case 2 unified CPF evolution for the SMP dose series, filtered to CMOD $\geq 0.02\,\mathrm{mm}$.

6.2 CMOD and CTOD Mapping

6.2.1 Virtual extensometer workflow

Both CMOD and CTOD are extracted from the same speckle-tracked DIC field that was used to validate δ in chapter 4. In the ZEISS Correlate environment, two pairs of markers are placed symmetrically about the notch centreline on the tension face of the beam. The outer pair straddles the notch mouth, so their relative horizontal displacement is the CMOD; the inner pair straddles the notch tip at the ligament root, so their relative horizontal displacement is the CTOD. Both measurements are differential, they therefore rely only on the local optical flow and not on any absolute camera position, and they are robust against the small rigid-body motions of the beam during loading.

6.2.2 Why CMOD and CTOD do not need a compliance correction

An important methodological difference separates the three kinematic variables. The deflection δ is the relative displacement between the crosshead and the ground frame of the testing machine, and therefore includes the elastic compliance of the machine-to-specimen load path that chapter 4 has already removed. CMOD and CTOD, on the other hand, are measured entirely on the specimen itself: they are the direct horizontal distance between two speckles imaged by the camera, and they do not traverse any component of the machine. Neither variable needs the ODDC correction, and both are independent of whatever stiffness the upper platen, load

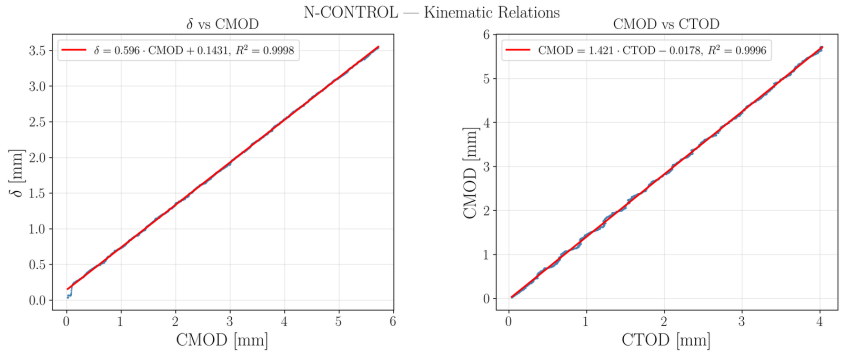
cell or fixtures possess. This observation is important downstream: eq. (6.4) uses the corrected δ of chapter 4, but eq. (6.5), eq. (6.1), eq. (6.2), eq. (6.3), eq. (6.6), eq. (6.7), eq. (6.8), eq. (6.9) and eq. (6.10) are all expressed in optical variables and inherit no machine-compliance uncertainty.

6.2.3 Kinematic fits — Case 1

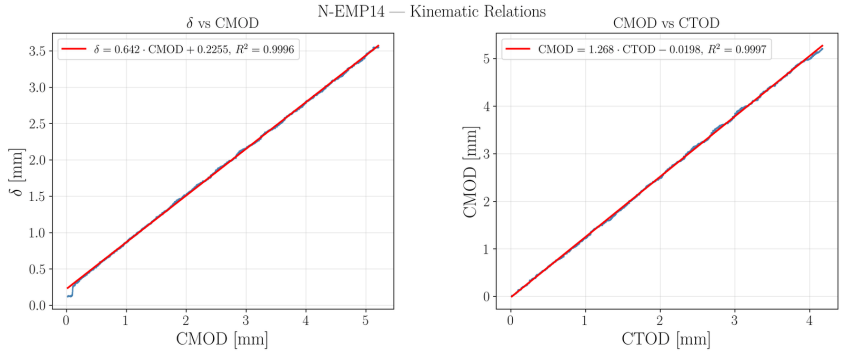
figs. 6.4 to 6.7 show the per-specimen kinematic-fit panels for the Case 1 cohort. Each panel is the raw export from the post-processing script: a single image file already contains the two sub-plots of interest (global mid-span deflection δ *versus* mouth opening on the left, and CMOD *versus* tip opening on the right), each with its ordinary-least-squares line and the printed R^2. For print legibility, each specimen is stacked *vertically* at nearly the full text width (no two specimens side-by-side). Natural quartets (for example the control plus three EMP charts, or the four Case 2 12 mm Normal doses) are kept on one float when they fit; remaining blocks use three rows per float with the same height caps. Where a single logical figure would still exceed the page height, the same figure number continues on the following page. Qualitatively, the eye test is simple: if both traces are straight and the scatter lies on the line, the rigid hinge of the EN 14651 idealisation is an acceptable local model for that specimen over the window 0.02 mm to 3.5 mm of mouth opening. Where a panel shows a kink, step or S-shape at the origin only, the cause is the optical seating effect discussed in section 6.1.4, not a breakdown of the hinge in the post-cracking range.

For each case below, the discussion focuses on the *difference* between families (matrix, fibre, length), because the two regression slopes in table 6.2 are already tabulated: here we interpret the *shape* of the data cloud as it appears in the panel. fig. 6.4 already shows that the unreinforced control and the EMP triplet share the same visual structure as the most ductile FRC: two straight segments with $R^2 > 0.99$. In fig. 6.5, panels (a)–(b) (WMP 30 mm and TFF 25 mm) shift the whole trace upward in the δ–CMOD plane: the slope a changes little, but the intercept b reflects the additional flexibility brought by strong crack bridging, while the CMOD–CTOD panel remains almost as linear as the control, confirming that the *kinematic* hinge is still intact even when the *structural* response is no longer matrix-dominated. Panels (c)–(e) are the SMP trio read in sequence: the 24 mm configuration shows the flattest δ response per unit mouth opening, consistent with a wider crack-bridging zone, whereas the 50 mm and 12 mm variants bracket it with marginally steeper and shallower mid-span movement.

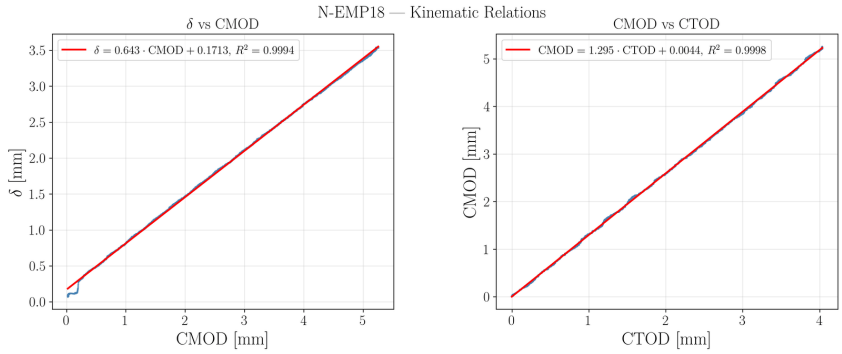
In the same Normal-cement family, the glass and steel end-members (fig. 6.5, pan-



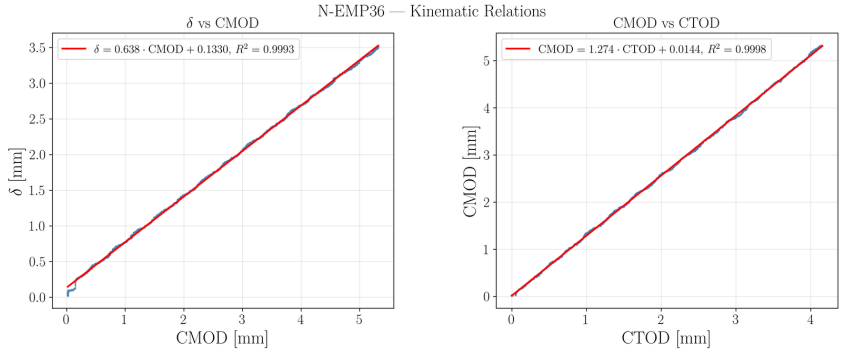
(a) N-CONTROL



(b) N-EMP-14 mm

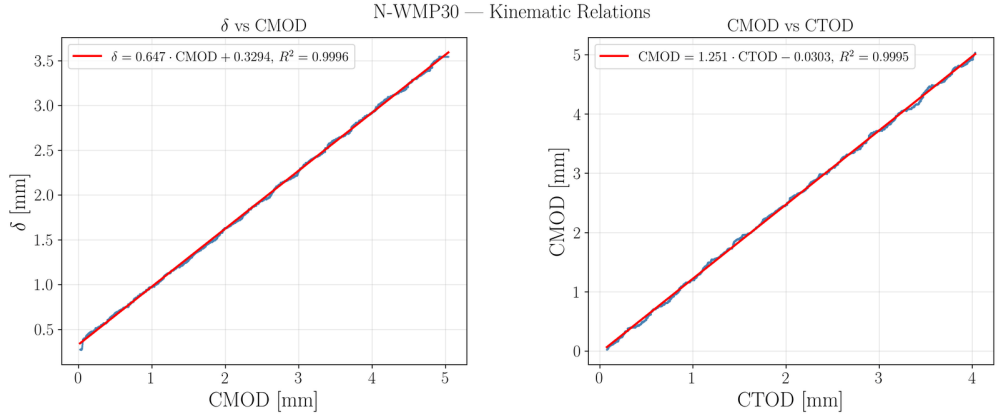


(c) N-EMP-18 mm

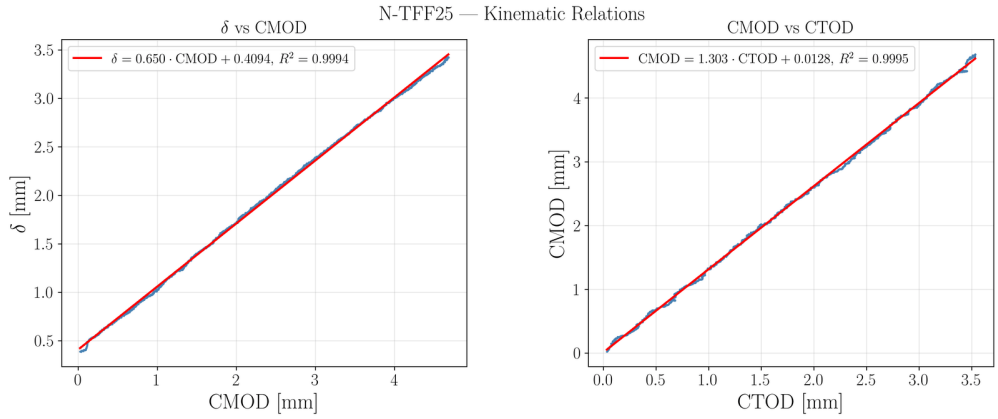


(d) N-EMP-36 mm

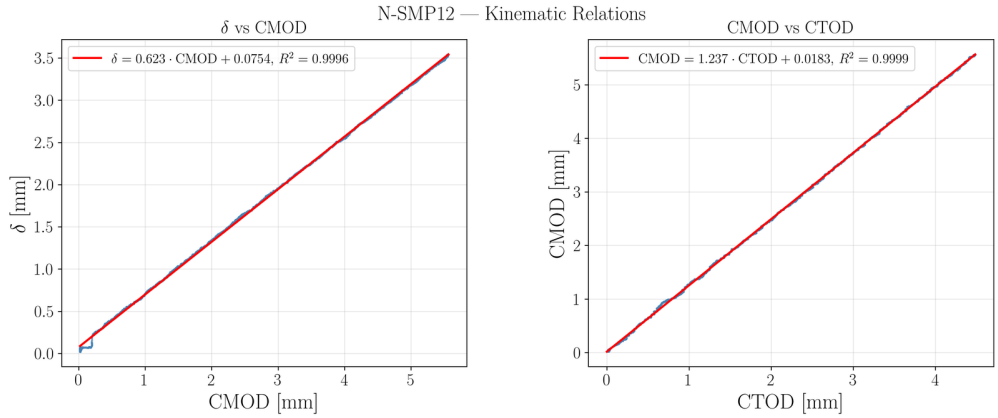
Figure 6.4: Case 1 Normal-matrix kinematic fits: left sub-plot of each panel is δ versus CMOD (eq. (6.4)), right sub-plot is CMOD versus CTOD (eq. (6.5)). The control and the EMP triplet share the same near-unity R^2 as the fibred cohort, confirming that the rigid-hinge relationship is set by the ligament geometry rather than by fibre type.



(a) N-WMP-30 mm



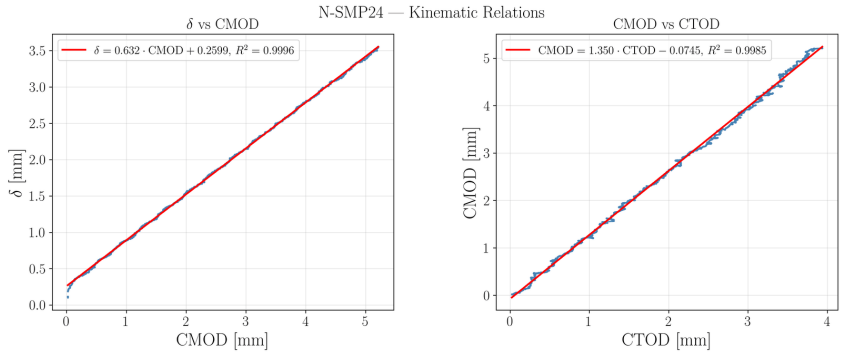
(b) N-TFF-25 mm



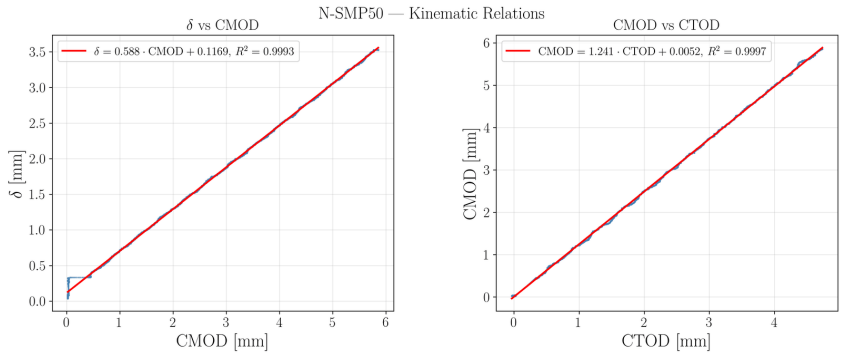
(c) N-SMP-12 mm

Figure 6.5: Case 1 Normal-matrix kinematic fits for the WMP, TFF, SMP, GF and SF fibre families. The three SMP lengths share nearly identical coefficient pairs ($c \approx 1.23$–1.32), consistent with a hinge geometry fixed by the ligament rather than by the fibre.

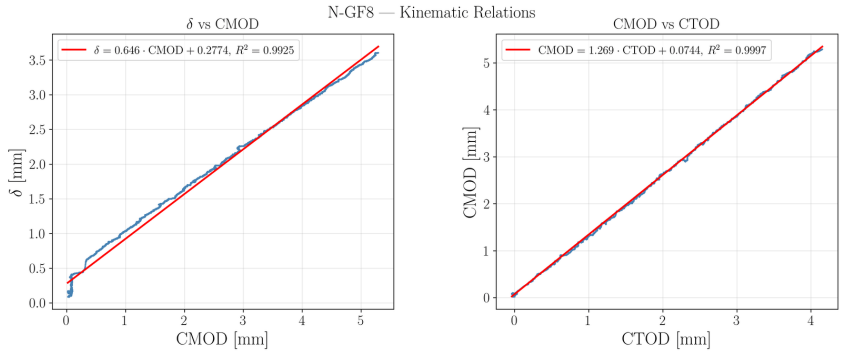
els (f) and (g)) illustrate two limits of the same kinematic test. N-GF-8 mm remains globally linear, but the left-hand line is the one that loses tightness: once filaments start sliding, the relationship between support rotation and mouth opening is no longer a single *constant* a over the full opening range, and the R^2 printed on the δ–CMOD panel (see section 6.2.6) is therefore the lowest of the set even though the



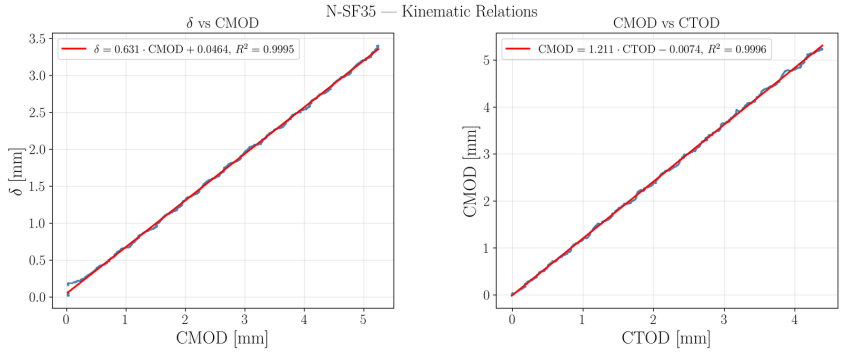
(d) N-SMP-24 mm



(e) N-SMP-50 mm



(f) N-GF-8 mm



(g) N-SF-35 mm

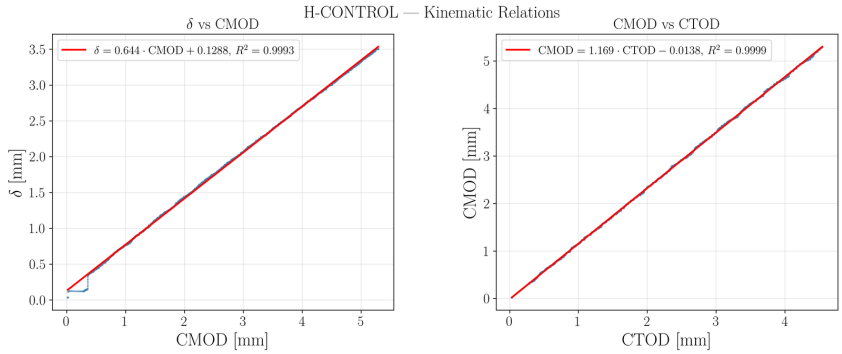
Figure 6.5: N-GF-8 mm is the only Normal-matrix chart where the left-hand δ–CMOD panel is visibly *noisier* than the right-hand CMOD–CTOD panel, the pattern expected when glass-fibre slippage couples weakly to the global beam rotation.

CMOD–CTOD line stays excellent. N-SF-35 mm, by contrast (panel (g) of fig. 6.5), shows both panels visually straight to the last pixel, the optical signature expected

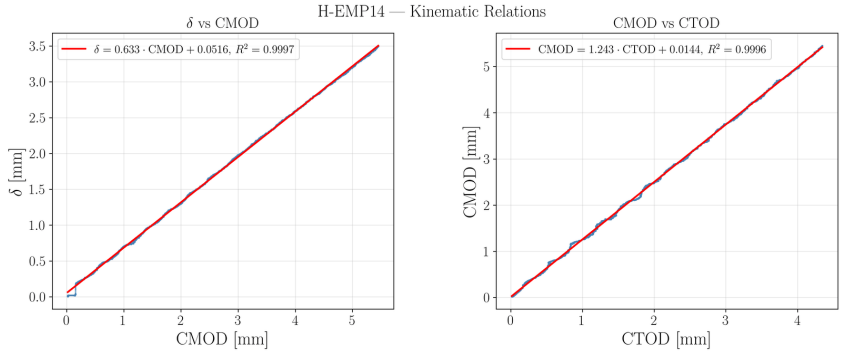
when pull-out, rather than filaments crossing the crack at shallow angle, governs the late branch.

The High-cement block (figs. 6.6 and 6.7) repeats the same fibre order with a denser interstitial paste. H-CONTROL already carries a higher proportionality between δ and mouth opening (stiffer matrix, smaller apparent rotation per micrometre of opening), and the High-matrix EMP, WMP and TFF systems inherit that *global* shift while preserving local linearity. The outlier in this block is H-TFF-25 mm, where the tip-opening trace in the *right* panel of the printed chart shows the one Case 1 departure from a straight CMOD–CTOD line; the event is revisited in section 6.3.6 because it is a genuine optical inconsistency, not a regression artefact. H-SF-35 mm shows the same collapse of pull-out in the High matrix that section 6.3.1 already highlighted on the F–CMOD record: the kinematic window is still well behaved, but the *forces* on that line are an order of magnitude lower than in the Normal matrix, so the chart looks "boringly linear" even while the material is the worst mechanical performer in the family.

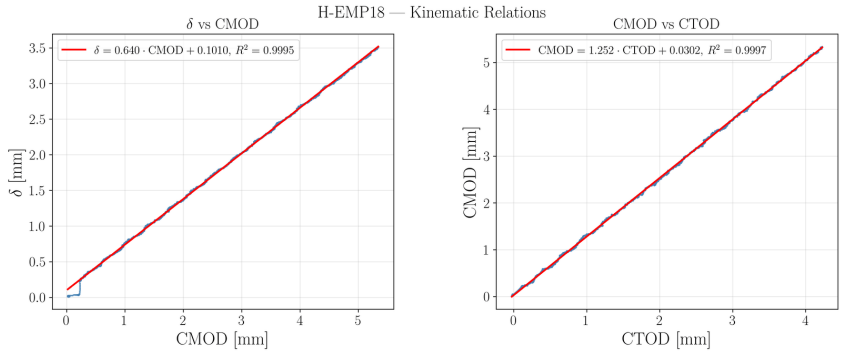
The last four High-matrix panels in fig. 6.7 (continued) close the Case 1 gallery at the two structural extremes of the material catalogue: long polypropylene macrofibres (H-SMP-24 mm, H-SMP-50 mm) versus glass and hooked-end steel (H-GF-8 mm, H-SF-35 mm). The SMP pair is kinematically well behaved in both sub-plots: the left-hand δ–CMOD slopes stay within the same 0.59–0.68 band as the other High-matrix fibres, and the right-hand CMOD–CTOD lines remain straight with $R^2 \gtrsim 0.999$. That observation matters because the same two specimens were not mechanical twins in chapter 5—the 24 mm macrofibre still delivers the higher post-cracking work—so the *ordering* of residual performance is not encoded in a larger or smaller a; it is encoded almost entirely in the F–CMOD map that these regressions do not see. H-GF-8 mm is the conversely interesting case: the two printed R^2 values in table 6.2 reappear here as a visibly *asymmetric* pair, with a slightly looser left panel than any other High-matrix fibre but an extremely tight right panel. The interpretation was already telegraphed in the glass discussion for the Normal matrix: filament slip uncouples global rotation from mouth opening in short bursts, so the OLS line through δ and CMOD is the one that softens, while the mouth-to-tip map stays affine because the local crack geometry has not yet fragmented. H-SF-35 mm is the opposite—both panels are straight to the last pixel, which is exactly what a bond-limited, pull-out-limited line should do when the forces have collapsed. The kinematic chart therefore looks "innocent" even when the mechanical story is dire; the coefficients confirm linearity, but they do not rescue the post-cracking strength that section 6.3.1 will show to be an order of magnitude below the Normal-matrix



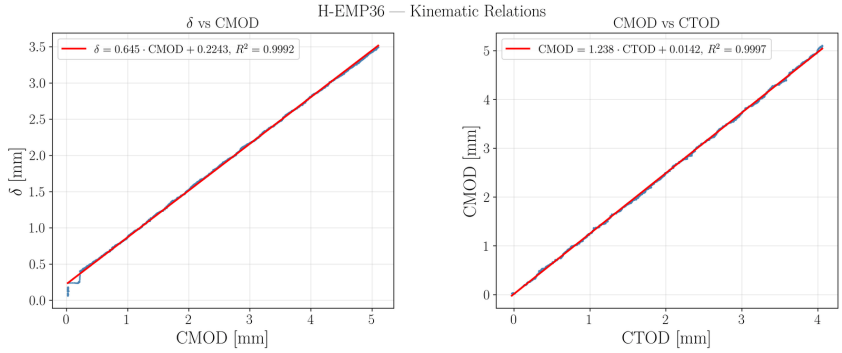
(a) H-CONTROL



(b) H-EMP-14 mm



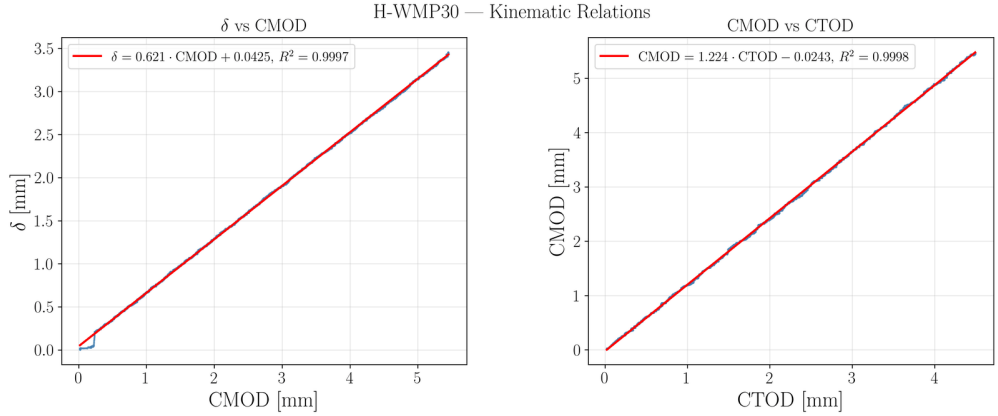
(c) H-EMP-18 mm



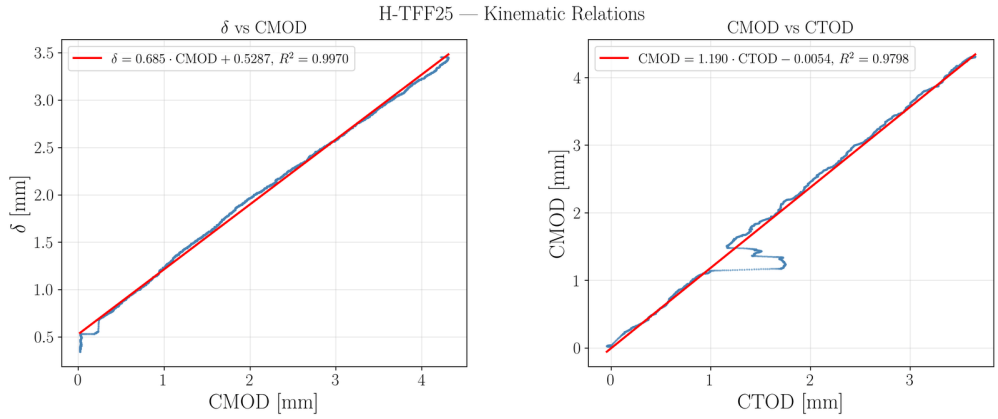
(d) H-EMP-36 mm

Figure 6.6: Case 1 High-matrix kinematic fits for the control and the EMP triplet. The High-cement matrix tightens the CMOD–CTOD slope ($c \approx 1.17$–1.25), consistent with the denser matrix restraining the early-stage rotation.

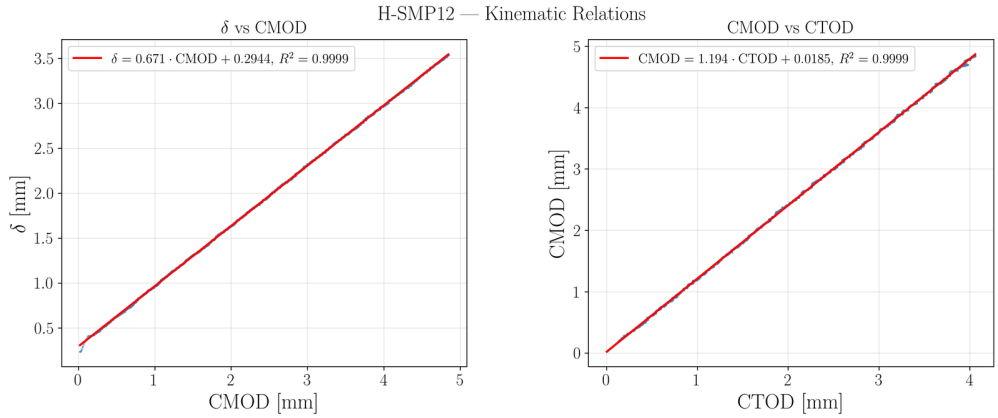
twin.



(a) H-WMP-30 mm



(b) H-TFF-25 mm

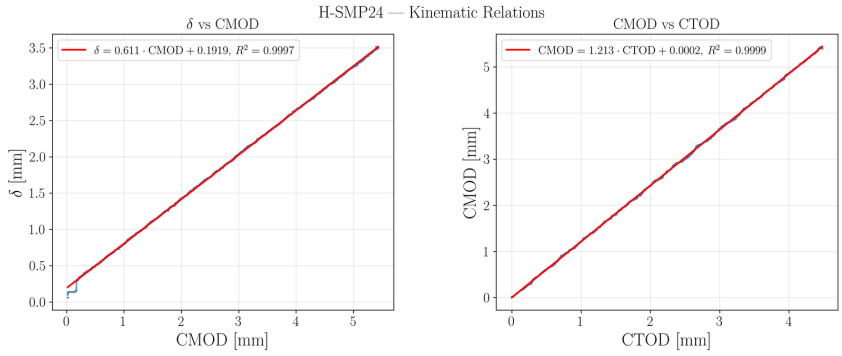


(c) H-SMP-12 mm

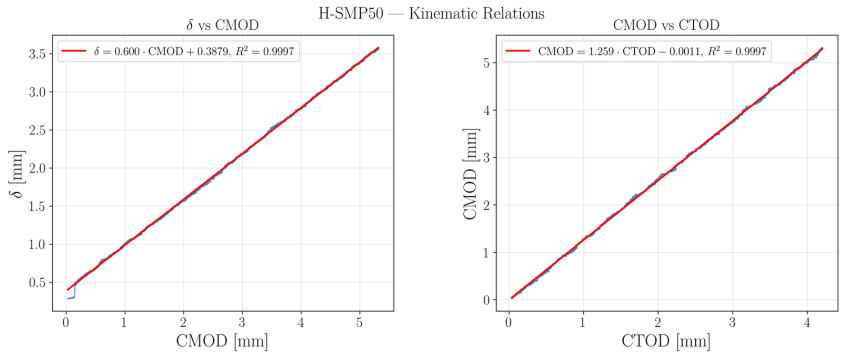
Figure 6.7: Case 1 High-matrix kinematic fits. H-TFF-25 mm is the single specimen in Case 1 with a visibly degraded CMOD–CTOD regression ($R^2 = 0.962$), explained by late-stage localised cracking outside the fit window.

6.2.4 Kinematic fits — Case 2

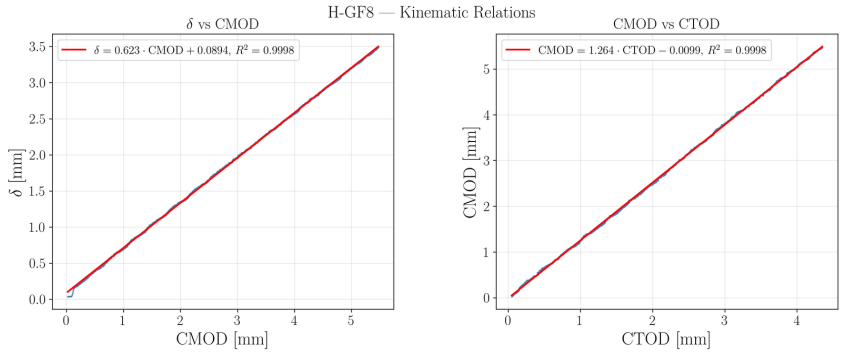
The second test matrix isolates a single commercial SMP polypropylene macrofibre, two lengths, two matrix strengths and a four-point dosage grid; it is therefore the cleanest test of whether the optical kinematic story of Case 1 carries over to a *single* system where only length and volume fraction move. Panels (a)–(b) of fig. 6.8 restate



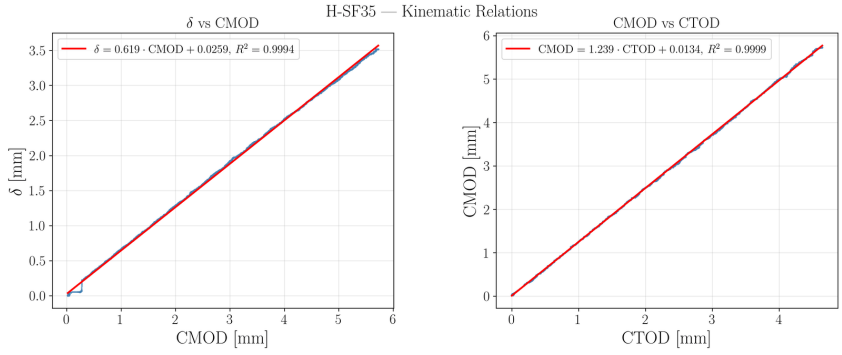
(d) H-SMP-24 mm



(e) H-SMP-50 mm



(f) H-GF-8 mm



(g) H-SF-35 mm

Figure 6.7: H-SF-35 mm is the High-matrix residual outlier already visible on the F–CMOD charts of section 6.3.1; kinematically it remains linear.

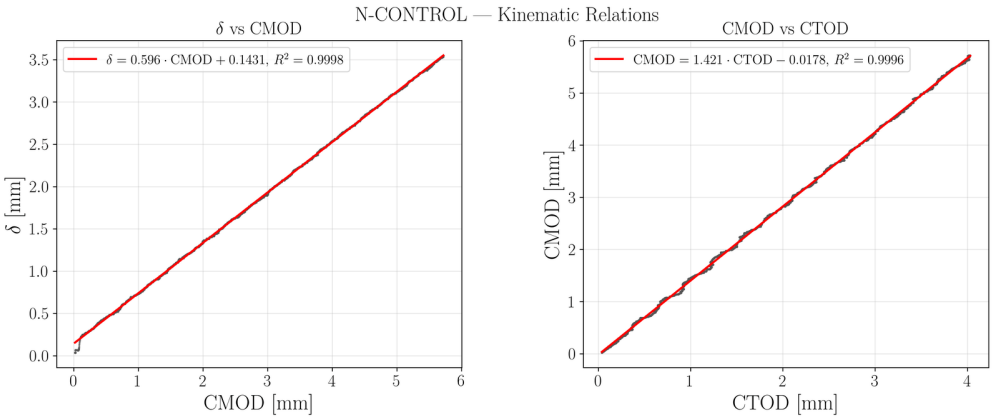
the matrix-only contrast (N-CONTROL and H-CONTROL), matching Case 1; the remaining panels of figs. 6.8 and 6.9 should be read in subfigure order as a vertical

gallery. All eighteen specimens use the same post-processing template as Case 1: the left half of each image is the corrected δ–CMOD line; the right half is CMOD–CTOD, each with a printed R^2.

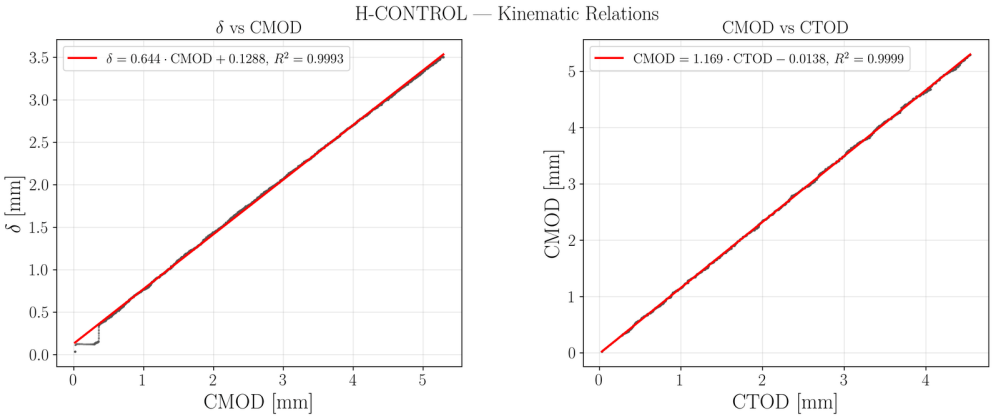
Within the 12 mm × Normal-cement block (panels (c)–(f) of fig. 6.8), the dose appears in the *offset* of the δ–CMOD line more than in its slope, which is the point already made numerically in table 6.3. The anomalous 1.00% Normal-matrix blend that appeared weak in the structural tables of chapter 5 still *looks* as linear as its neighbours here, confirming that the 1.00% "dip" is a strength issue, not a failed kinematic measurement. fig. 6.9 repeats the dosage grid for 24 mm: the δ–CMOD slopes are systematically shallower (smaller a) than the 12 mm slopes at the same mass fraction, not because the hinge is *less* valid, but because a longer pull-out length spreads the same mouth opening over a marginally longer lever arm. Compare panel-by-panel with fig. 6.8 at the same mass fraction: c changes only in the second decimal, while a is what moves.

Figure 6.8 (including its continued parts) is the most compact way to read the factorial in optical space. The two control panels restate, with the same DIC template as Case 1, that matrix density re-scales the δ–CMOD line—H-CONTROL sits higher in a than N-CONTROL, i.e. less mid-span deflection per micrometre of mouth opening for the same ligament—while the mouth-to-tip construction remains a near-constant c within each matrix. Once fibre is present, the Normal-matrix 12 mm sweep (panels (c)–(f)) should be scanned in *dosage* order at fixed length: the slopes a in table 6.3 move only in the third decimal for the 1.00%–1.50% decade, but the intercepts b and the offset of the line on the page move visibly, which is the regression analogue of "the same hinge, a different virtual seating." The controversial 1.00% Normal-matrix mixture that looked weak in the structural bar charts of chapter 5 is not a failed kinematic measurement: the line pair is as straight as its neighbours, and the CMOD–CTOD R^2 remains above 0.999. The mechanical dip is therefore carried by the F data paired with that CMOD, not by a break in the δ–CMOD–CTOD hierarchy. The High-matrix 12 mm ladder (continued in the final block) shows the same straight-line family with a High-cement c band near 1.20–1.30, and it brings out one subtlety useful for the dose plots of section 6.3.5: the worst Case 2 R^2 entries for CMOD–CTOD in table 6.3 are not where the F–CMOD record is worst; they are where the tip extensometer had to span a more tortuous post-crack path at low dosage. That is why a designer should always read table 6.3 and figs. 6.8 and 6.9 together with the force traces.

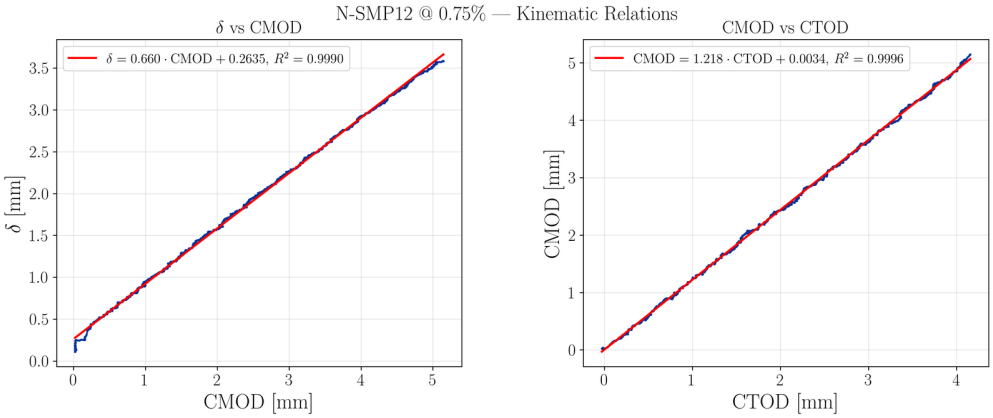
The coefficients in table 6.3 for the block above can be read alongside panels (c)–(j) of fig. 6.8: any specimen whose chart looks linear but whose $f_{R,3}$ dipped in chapter 5



(a) N-CONTROL



(b) H-CONTROL

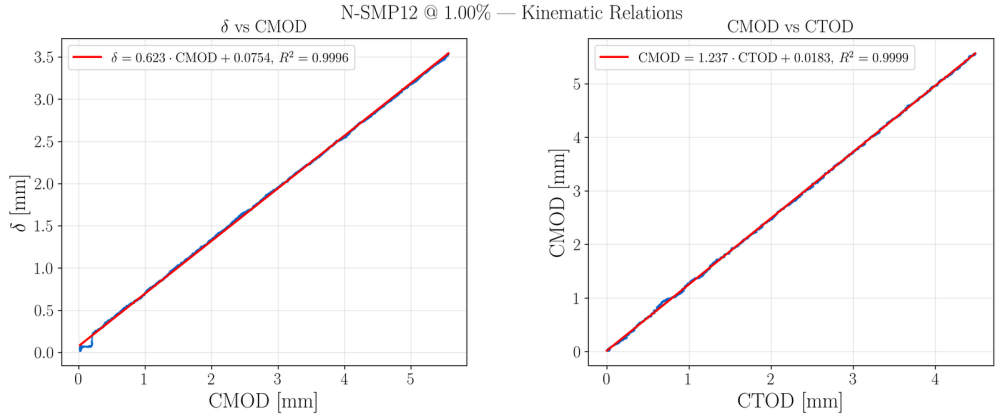


(c) N-SMP-12 mm @ 0.75%

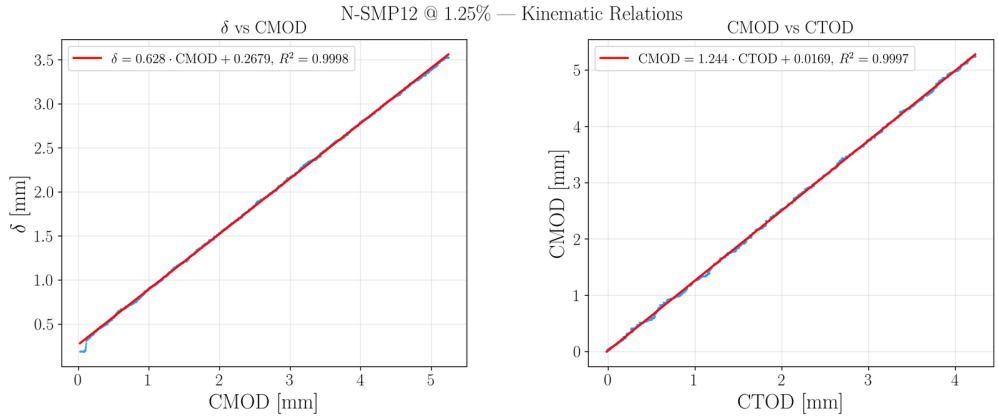
Figure 6.8: Case 2 kinematic fits for the two controls and the eight 12-mm specimens across the 0.75%, 1.00%, 1.25%, 1.50% dosage ladder. Dosage affects the intercept b but leaves the slope a nearly constant; the rigid-hinge signature is preserved.

is a warning that kinematic adequacy does not imply mechanical optimality at every dose.

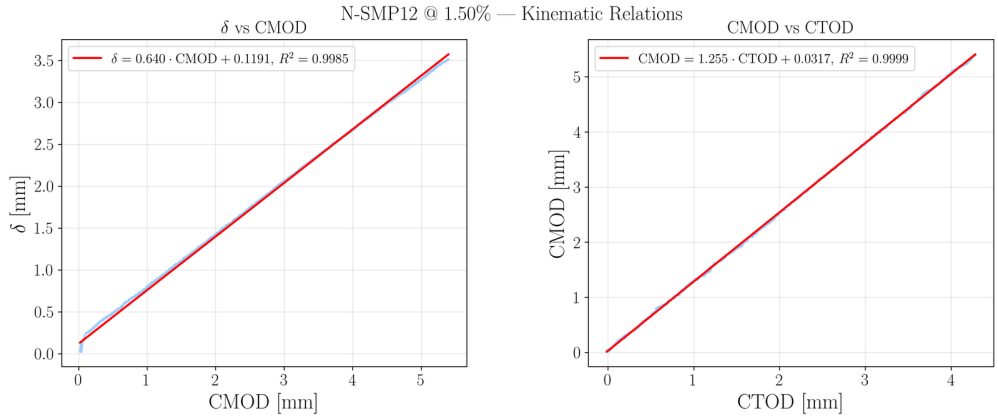
The first block of fig. 6.9 carries the full Normal-matrix 24 mm dosage ladder; the continued block carries the four High-matrix 24 mm doses at the same abscissae so the two matrices can be compared at fixed length before returning to the closing tables.



(d) N-SMP-12 mm @ 1.00%



(e) N-SMP-12 mm @ 1.25%



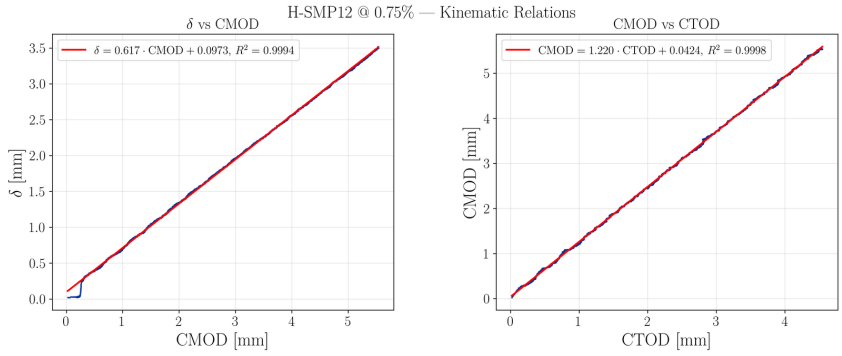
(f) N-SMP-12 mm @ 1.50%

Figure 6.8 (continued)

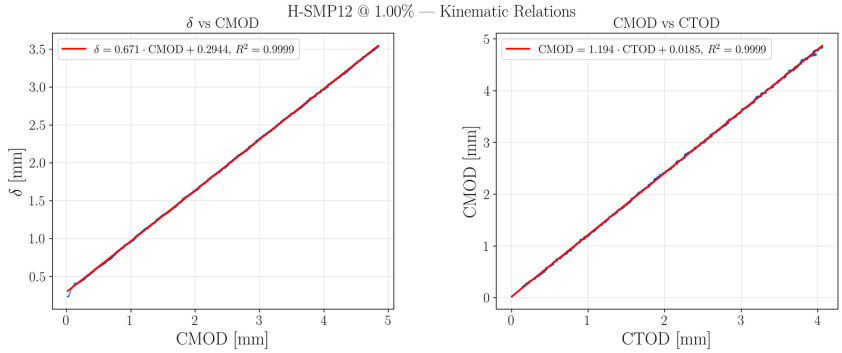
6.2.5 Tables of kinematic coefficients

table 6.2 and table 6.3 list the two regression slopes, the two intercepts and the two R^2 values for every specimen in the two case studies.

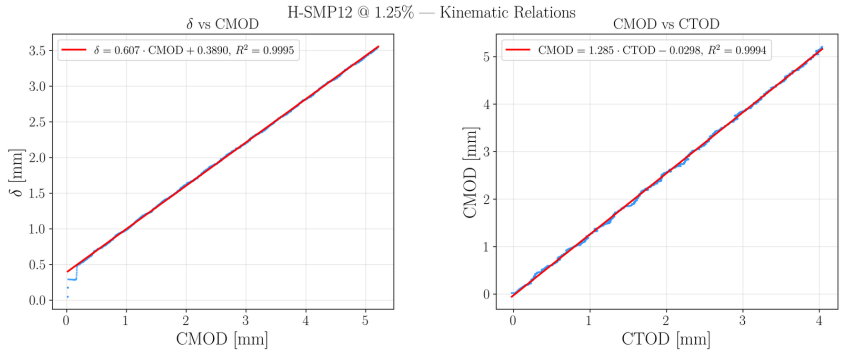
Interpreting the Case 1 coefficients (table 6.2). Table 6.2 is the numerical counterpart of the Case 1 kinematic gallery (figs. 6.4 to 6.7 and their continuations). The slope a is the least "sociable" column: for the polypropylene families it clusters



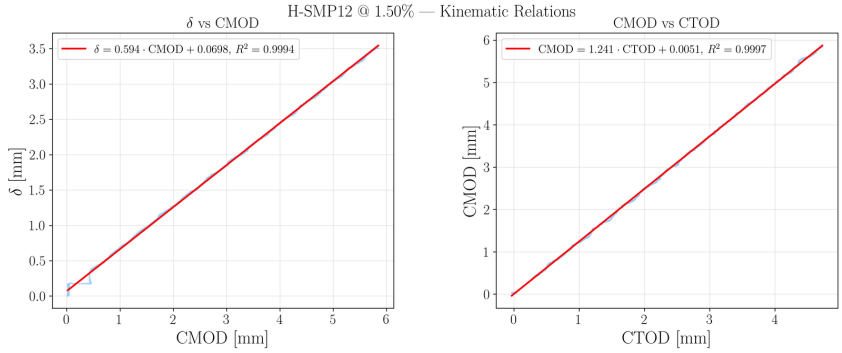
(g) H-SMP-12 mm @ 0.75%



(h) H-SMP-12 mm @ 1.00%



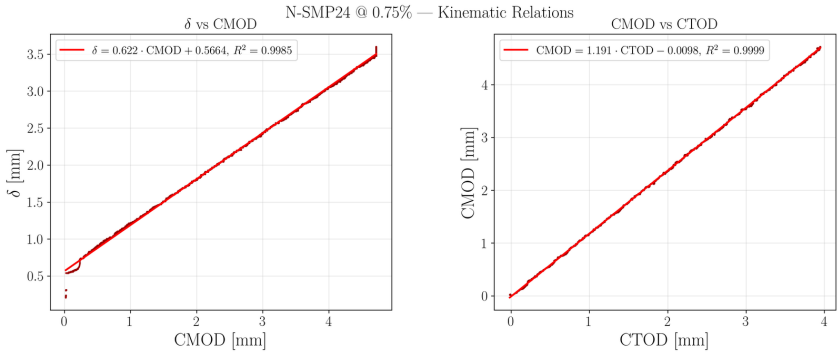
(i) H-SMP-12 mm @ 1.25%



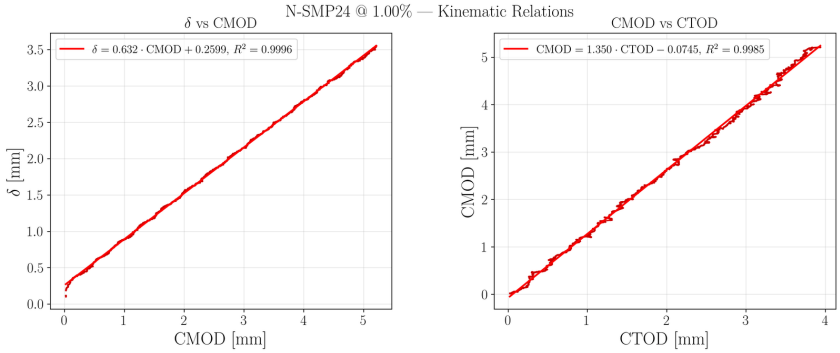
(j) H-SMP-12 mm @ 1.50%

Figure 6.8: The High-matrix 12 mm ladder closes here; coefficients are listed in table 6.3.

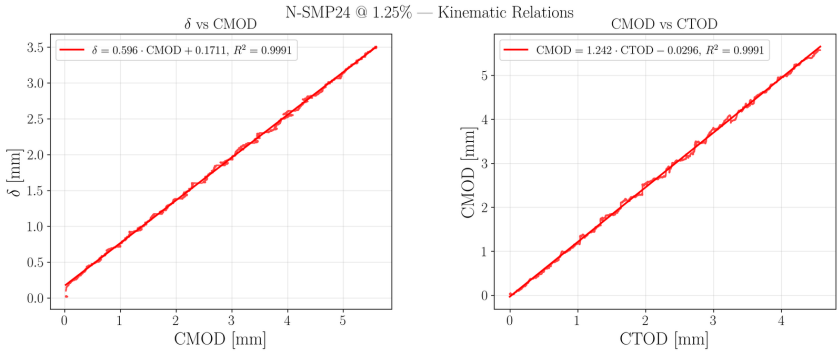
near 0.63–0.66 with no systematic ranking by matrix, which supports the design reading that the δ–CMOD proportionality is set mainly by the ligament and span,



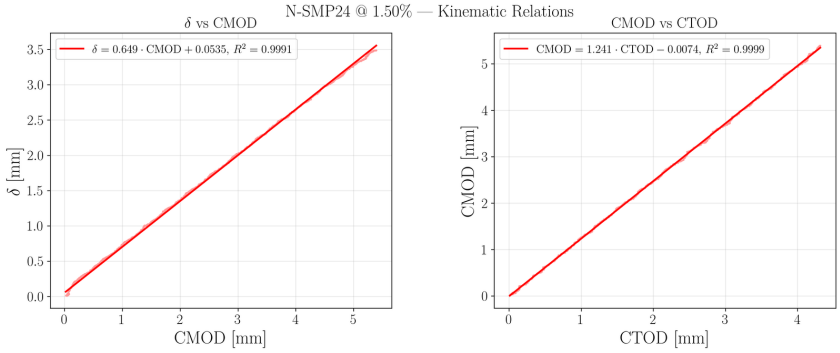
(a) N-SMP-24 mm @ 0.75%



(b) N-SMP-24 mm @ 1.00%

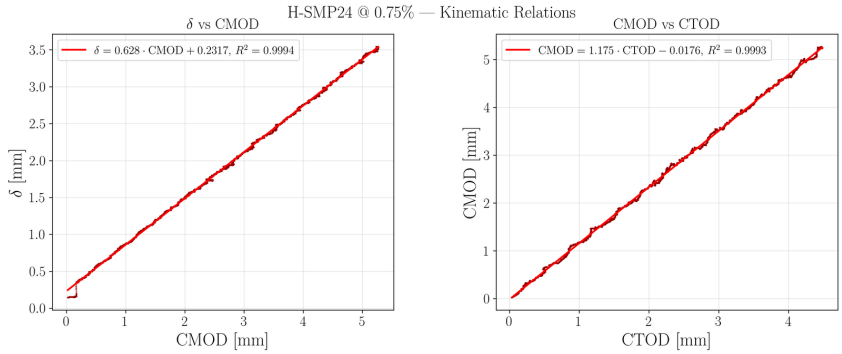


(c) N-SMP-24 mm @ 1.25%

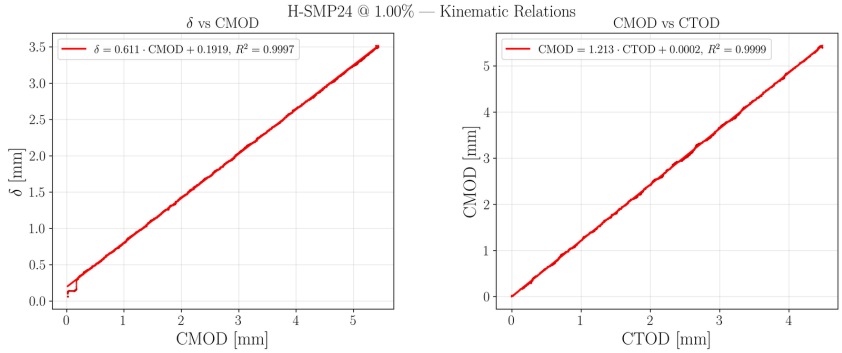


(d) N-SMP-24 mm @ 1.50%

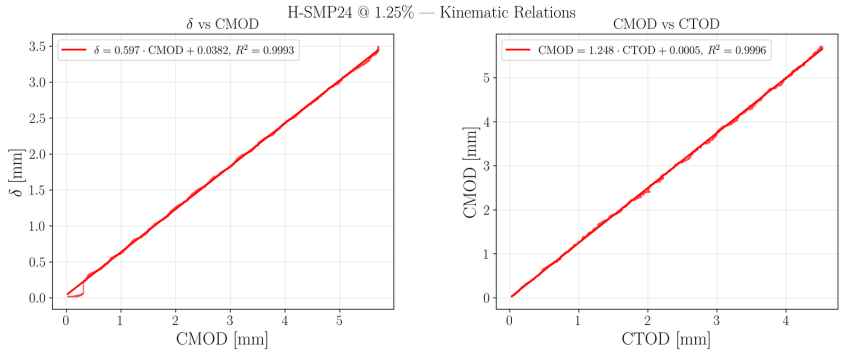
Figure 6.9: Case 2 kinematic fits for the eight 24-mm specimens. The longer fibre pushes the δ–CMOD slope below 0.62 at several doses, indicating that the 24-mm bridge distributes the opening over a slightly wider region than the 12-mm bridge and therefore delivers less mid-span deflection per unit mouth opening.



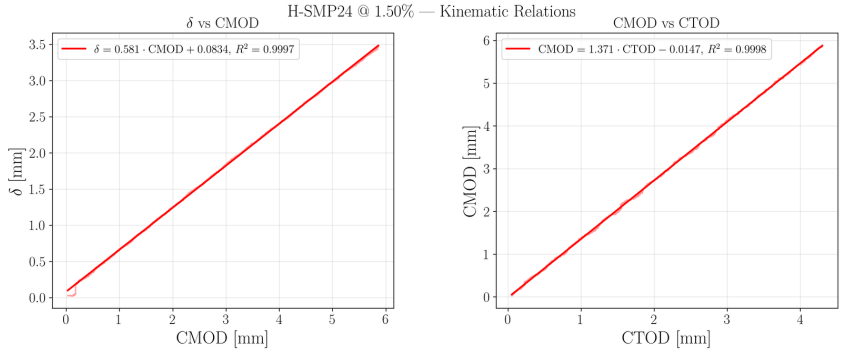
(e) H-SMP-24 mm @ 0.75%



(f) H-SMP-24 mm @ 1.00%



(g) H-SMP-24 mm @ 1.25%



(h) H-SMP-24 mm @ 1.50%

Figure 6.9 (continued)

with fibre type entering only as a small perturbation. The glass row is the expected exception: N-GF-8 mm reports the highest a in the Normal block and a depressed $R^2_{\delta\text{–CMOD}}$, exactly what slip-localisation would produce when the OLS line is forced

Table 6.2: Case 1 kinematic coefficients (Normal-matrix block, Control–TFF). Slopes a, c, intercepts b, d and both R^2 values as in eqs. (6.4) and (6.5); remaining Normal-matrix rows and the High-matrix block follow in Table 6.2 (continued).

Specimen	a	b (mm)	$R^2_{\delta\text{–CMOD}}$	c	d (mm)	$R^2_{\text{CMOD–CTOD}}$
N-CONTROL	0.596	0.142	0.9993	1.432	−0.029	0.9985
N-EMP-14 mm	0.651	0.214	0.9992	1.275	−0.029	0.9996
N-EMP-18 mm	0.656	0.153	0.9986	1.295	0.005	0.9996
N-EMP-36 mm	0.650	0.116	0.9985	1.268	0.022	0.9996
N-WMP-30 mm	0.644	0.334	0.9993	1.263	−0.044	0.9991
N-TFF-25 mm	0.664	0.390	0.9996	1.287	0.029	0.9993

Table 6.2 (continued): Normal-matrix block (SMP–SF).

Specimen	a	b (mm)	$R^2_{\delta\text{–CMOD}}$	c	d (mm)	$R^2_{\text{CMOD–CTOD}}$
N-SMP-12 mm	0.635	0.062	0.9992	1.234	0.021	0.9997
N-SMP-24 mm	0.640	0.249	0.9993	1.323	−0.042	0.9979
N-SMP-50 mm	0.591	0.113	0.9973	1.227	0.020	0.9994
N-GF-8 mm	0.688	0.239	0.9849	1.255	0.083	0.9993
N-SF-35 mm	0.631	0.047	0.9987	1.205	−0.002	0.9993

Table 6.2 (continued): High-matrix block.

Specimen	a	b (mm)	$R^2_{\delta\text{–CMOD}}$	c	d (mm)	$R^2_{\text{CMOD–CTOD}}$
H-CONTROL	0.659	0.104	0.9983	1.171	−0.018	0.9998
H-EMP-14 mm	0.643	0.037	0.9992	1.239	0.019	0.9987
H-EMP-18 mm	0.647	0.090	0.9986	1.245	0.038	0.9992
H-EMP-36 mm	0.656	0.209	0.9981	1.229	0.026	0.9996
H-WMP-30 mm	0.627	0.034	0.9991	1.214	−0.012	0.9996
H-TFF-25 mm	0.705	0.506	0.9963	1.177	0.007	0.9617
H-SMP-12 mm	0.674	0.291	0.9999	1.197	0.015	0.9998
H-SMP-24 mm	0.618	0.182	0.9990	1.211	0.003	0.9996
H-SMP-50 mm	0.598	0.391	0.9992	1.270	−0.014	0.9994
H-GF-8 mm	0.631	0.077	0.9997	1.253	0.002	0.9995
H-SF-35 mm	0.637	0.005	0.9993	1.231	0.021	0.9999

through a window that still contains a short non-affine toe. The slope c is the more design-relevant number for optical substitution of CMOD/CTOD: when c is large, a given tip signal maps to a larger mouth signal, and when c is smaller (typical of the High-matrix block) the same tip opening corresponds to a smaller mouth opening, i.e. the measured "amplification" from tip to mouth is calmer. Comparing the two controls line-by-line, H-CONTROL carries a marginally lower c than N-CONTROL even though the High-matrix paste is stiffer: the DIC is telling us that, once the global rigid-hinge map is fixed, the denser matrix does not *inflate* the tip-to-mouth factor; if anything, it damps the ratio slightly.

The intercepts b and d should be treated as nuisance parameters in this table. They move by tenths of a millimetre when virtual extensometer seeds shift or when the test starts with a small rigid-body offset. The correct way to use them is relative: a sudden jump in b between two sister specimens with the same a and c is a red flag for a different image crop, not for a different hinge. The TFF outlier in the

High block is the opposite case: H-TFF-25 mm shows an enormous b in the second regression because the right-hand line is the one that fails; the tabulated c and $R^2_{\text{CMOD–CTOD}}$ flag the damage event more honestly than a caption could. N-SMP-24 mm is the one Normal-matrix entry where the mouth-to-tip R^2 dips slightly below the rest while still remaining above 0.99; that is a reminder that "high R^2" is not a unique number—when two fits differ in the third decimal, the per-specimen charts in sections 6.2.3 and 6.2.4 are the tie-breaker.

Table 6.3: Case 2 kinematic coefficients. Same column structure as table 6.2; all eighteen specimens of the length × dose × matrix factorial are listed.

Specimen	a	b (mm)	$R^2_{\delta\text{–CMOD}}$	c	d (mm)	$R^2_{\text{CMOD–CTOD}}$
N-CONTROL	0.596	0.142	0.9993	1.432	−0.029	0.9985
H-CONTROL	0.659	0.104	0.9983	1.171	−0.018	0.9998
N-SMP-12 mm @ 0.75%	0.677	0.241	0.9983	1.212	0.010	0.9993
N-SMP-12 mm @ 1.00%	0.635	0.062	0.9992	1.234	0.021	0.9997
N-SMP-12 mm @ 1.25%	0.635	0.257	0.9995	1.251	0.008	0.9993
N-SMP-12 mm @ 1.50%	0.659	0.099	0.9968	1.250	0.036	0.9998
N-SMP-24 mm @ 0.75%	0.624	0.565	0.9966	1.189	−0.007	0.9998
N-SMP-24 mm @ 1.00%	0.640	0.249	0.9993	1.323	−0.042	0.9979
N-SMP-24 mm @ 1.25%	0.605	0.159	0.9977	1.247	−0.039	0.9975
N-SMP-24 mm @ 1.50%	0.665	0.037	0.9984	1.239	−0.006	0.9997
H-SMP-12 mm @ 0.75%	0.629	0.078	0.9982	1.213	0.049	0.9994
H-SMP-12 mm @ 1.00%	0.674	0.291	0.9999	1.197	0.015	0.9998
H-SMP-12 mm @ 1.25%	0.613	0.381	0.9986	1.275	−0.018	0.9985
H-SMP-12 mm @ 1.50%	0.603	0.058	0.9982	1.227	0.020	0.9994
H-SMP-24 mm @ 0.75%	0.627	0.233	0.9983	1.177	−0.022	0.9982
H-SMP-24 mm @ 1.00%	0.618	0.182	0.9990	1.211	0.003	0.9996
H-SMP-24 mm @ 1.25%	0.608	0.021	0.9974	1.230	0.026	0.9993
H-SMP-24 mm @ 1.50%	0.590	0.067	0.9993	1.366	−0.009	0.9992

6.2.6 What the coefficients reveal

Three observations run through both tables. First, R^2 values typically exceed 0.995, with a mean of $\overline{R^2}_{\delta\text{–CMOD}} = 0.9983$ and $\overline{R^2}_{\text{CMOD–CTOD}} = 0.9970$ across the full 40-specimen set. Only one entry (H-TFF-25 mm, last-stage localised cracking) falls below 0.99; on the other 39 specimens the optical data are *consistent with* the rigid-hinge assumption at the two-decimal level, although such high R^2 values demonstrate kinematic consistency rather than constituting an independent physical validation of the hinge mechanism. Second, the slope a of the δ–CMOD regression clusters tightly around 0.63 ± 0.03, irrespective of fibre programme. The ligament geometry, not the reinforcement, is the dominant factor that sets the relative magnitude of mid-span deflection and mouth opening. Third, the slope c of the CMOD–CTOD regression is slightly larger in the Normal-cement matrix than in

the High-cement matrix ($\overline{c} = 1.27$ versus $\overline{c} = 1.22$), reflecting the stiffer post-peak response of the denser matrix which restricts the tip-to-mouth amplification factor. The intercepts b and d carry the optical seating offsets and they have no independent mechanical meaning.

A fourth, more operational point follows from stacking table 6.2 and table 6.3. For a mixture-design study, the slope a in table 6.3 is the quantity that *should* move when the fibre length or dosage changes, because both alter the position of the macroscopic hinge and the path length over which the crack opens, whereas c is the quantity that *should* remain stable if the DIC only re-scales the same local crack as length and ρ_f change. The factorial is broadly consistent with that pattern: the largest excursions in a are tied to the 24 mm branch in the High matrix, where the deflection–opening map genuinely shifts, while c for the 12 mm and 24 mm sweeps at the same matrix typically agree in the second decimal. When the two columns move together, it is a hint that the optical fit is tracking a *global* change in the recorded displacement field (e.g. a seed-point jump or a re-mesh) rather than a second mechanical mechanism.

A fifth point concerns how these coefficients feed forward into chapter 7. The optical regressions are not themselves material properties; they are *compatibility* checks. What Chapter 7 needs from this chapter is the assurance that, for a given F–CMOD record, the CMOD–CTOD map is single-valued and almost affine, so the cohesive law can be written as a function of a single opening variable. When $R^2 \geq 0.99$ on both lines, the substitution chain $\delta \rightarrow \mathrm{CMOD} \rightarrow \mathrm{CTOD}$ is numerically well conditioned; when one line fails while the other passes, the failure mode is usually local to one extensometer pair and can be post-filtered without re-running the test.

Two items are worth flagging. The first is N-GF-8 mm, whose $R^2_{\delta\text{–CMOD}}$ (0.9849) is conspicuously lower than its $R^2_{\text{CMOD–CTOD}}$ (0.9993). The degradation is traced to a localised seating step visible at low CMOD in the original kinematic plot, which is bracketed by the 0.02 mm lower bound of the fit window but still pulls the residuals. The second is H-TFF-25 mm, where an unexpected tip-opening spike near the end of the test breaks the CMOD–CTOD regression. That single spike is discussed in section 6.3.6.

6.3 Evaluation of the Rigid-Hinge Assumption

With the two kinematic regressions confirmed specimen by specimen in section 6.2, we can now read the EN 14651, UNI 11039-2 and RILEM TC 162-TDF residual-

strength and fracture-energy metrics directly from the F–CMOD and F–CTOD traces. The rigid-hinge assumption is evaluated at two levels. At the cohort level, sections 6.3.1 and 6.3.3 compile the bar charts that summarise every fibre programme on the same axis; at the per-specimen level, sections 6.3.2 and 6.3.4 report the individual F–CMOD traces with the four $f_{R,j}$ markers drawn on them, so that the interplay between mouth opening and residual load can be read directly off the plot. section 6.3.5 condenses the Case 2 data into the dose-response sweeps, and section 6.3.6 collects the (few) specimens in which the rigid-hinge assumption is locally violated.

6.3.1 Case 1 cohort response

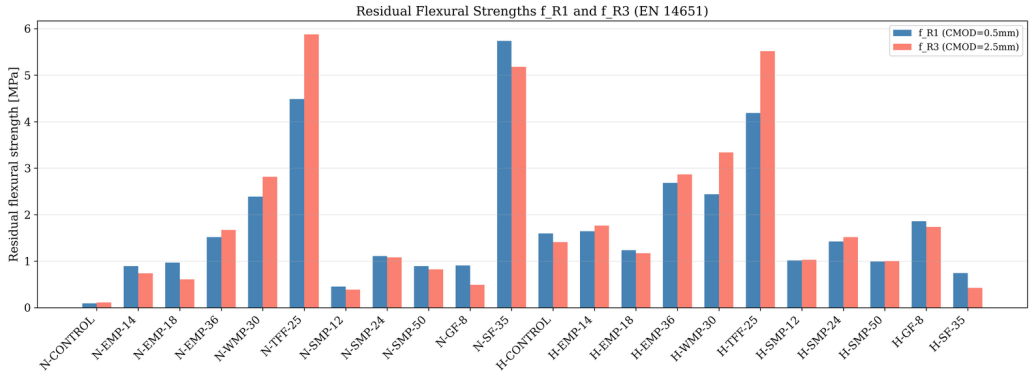


Figure 6.10: Case 1 residual strengths $f_{R,1}$, $f_{R,2}$, $f_{R,3}$, $f_{R,4}$ from eq. (6.1) for every specimen. N-TFF-25 mm and H-TFF-25 mm are the chapter leaders at $f_{R,3}$ in the Normal- and High-cement matrices, respectively. H-SF-35 mm is the High-matrix residual loser, consistent with the bond-rupture rather than pull-out failure mode observed on its micrograph.

The seven Case 1 cohort panels above are deliberately redundant: they translate the same F–CMOD records into scalar summaries that answer different design questions. Figure 6.10 compares residual forces at fixed mouth openings; reading $f_{R,1}$ against $f_{R,3}$ for one specimen distinguishes post-cracking gain from first-crack scale without confounding matrix strength. Figure 6.11 integrates the entire trace to 3.5 mm, so a specimen can match a neighbour in $f_{R,3}$ yet lag in G_f when its plateau shortens after the 2.5 mm construction point. Figure 6.12 localises that behaviour along CMOD: where the slope of $D^f_{BZ}(\mathrm{CMOD})$ is large, bridging stress is redeployed efficiently as the crack widens; a high terminal D^f_{BZ} reached on a gentle slope marks a slow pull-out tail that still imports work. Figure 6.14 averages those traces within each fibre family and overlays matrices—coincident solid and dashed lines mean that fibre geometry, not paste density, governs the bridging law, whereas diverging lines

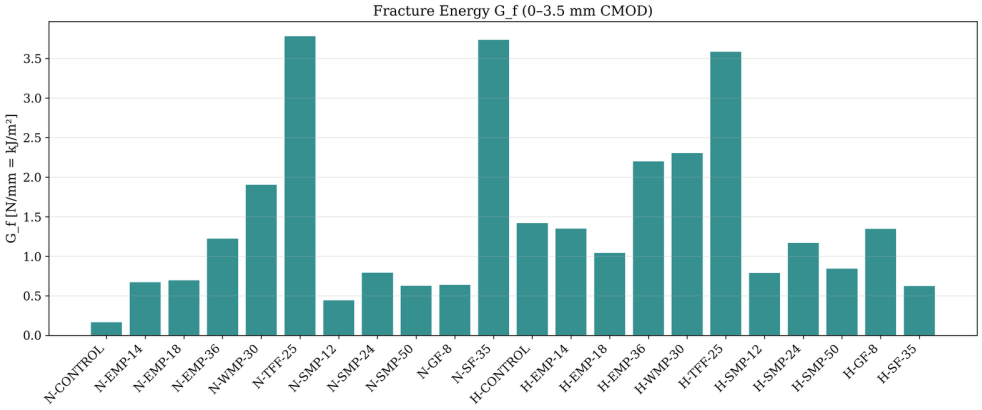


Figure 6.11: Case 1 specific fracture energy G_f from eq. (6.3) for the 22 specimens. The ranking matches the residual-strength ranking of fig. 6.10: TFF dominates both matrices; N-SF-35 mm is the single-family winner of the fibre-metal class; H-SF-35 mm loses an order of magnitude relative to its Normal-matrix counterpart.

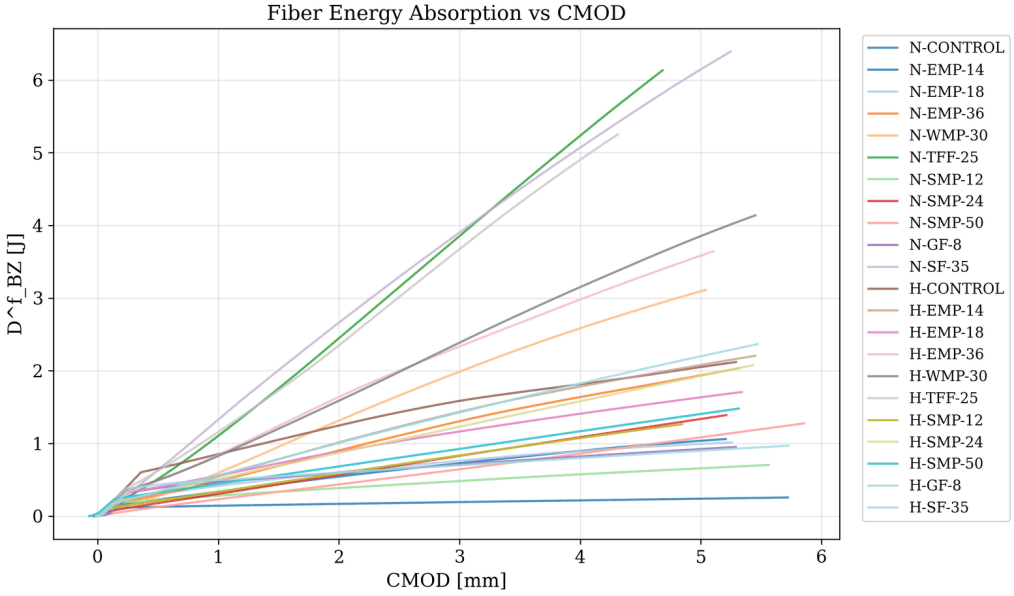


Figure 6.12: Case 1 fibre-bridging energy D^f_{BZ} of eq. (6.7) accumulated against CMOD for every specimen on a single axis. TFF and SF-35 occupy the top band in the Normal matrix, SMP and EMP in the middle, and the controls at the floor. The plot confirms the G_f ranking of fig. 6.11 via the differential (CMOD-wise) measure.

isolate matrix strength and interfacial bond evolution. Figure 6.13 reframes the same information per kilogram of fibre through $\rho_f \cdot L_f$, which is the procurement-relevant view when dosage is cost-driven rather than fixed by mixture design alone. Figure 6.15 compresses the two EN 14651 ordinates into a single hardening index that is insensitive to absolute scale but sensitive to post-cracking retention, which is why High-matrix SF-35 appears catastrophic in the ratio plot while still registering a competitive peak in the raw trace. Figure 6.16 partitions energy between

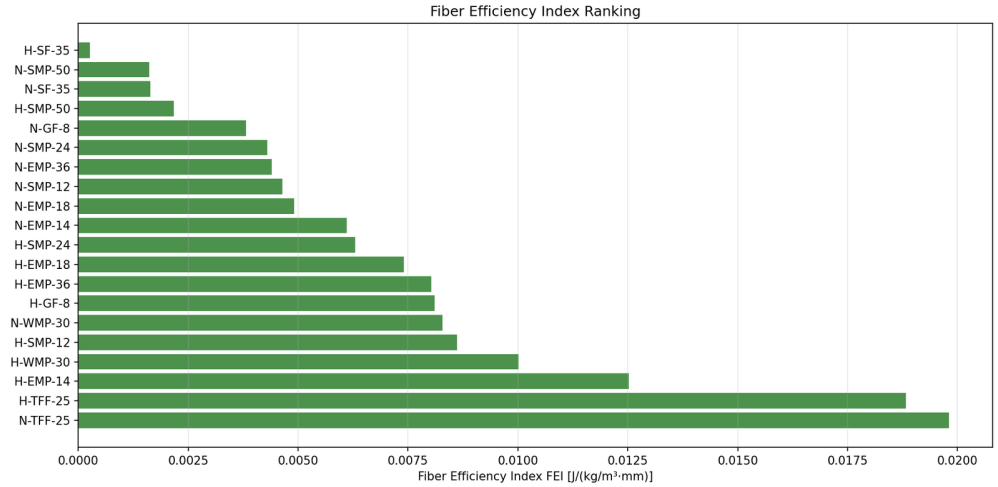


Figure 6.13: Case 1 Fibre Efficiency Index (eq. (6.8)) ranked across all fibre programmes. After normalisation by volumetric dosage and fibre length, N-TFF-25 mm and H-TFF-25 mm remain at the top, EMP-14 and EMP-36 climb into the upper tier, and SF-35 drops to the bottom — the denominator $\rho_f \cdot L_f$ penalises the dense, long steel fibre disproportionately.

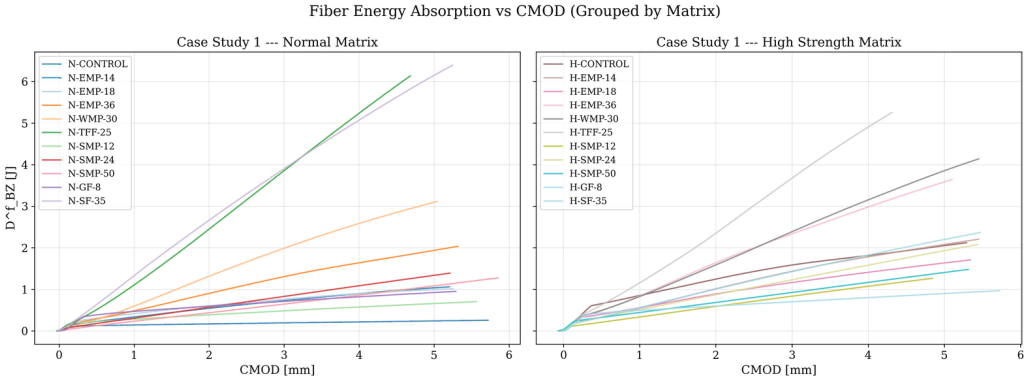


Figure 6.14: Case 1 fibre-bridging energy D^f_{BZ} of eq. (6.7) averaged within each fibre family and drawn against CMOD. The two matrices are overlaid: solid lines for Normal, dashed lines for High. TFF and SMP show almost matrix-insensitive response; SF-35 is the family with the largest matrix spread.

the UNI small-opening window and the fibre-dominated tail: mismatched UNI bars for one specimen flag disproportionate energy release in one stage, typically a sharp post-peak loss after 0.6 mm when the late-stage f_{eq} bar lags the early bar.

Four readings can be drawn from figs. 6.10 to 6.16. First, TFF fibres dominate the ranking regardless of matrix: N-TFF-25 mm and H-TFF-25 mm achieve the highest $f_{R,3}$ (5.88 MPa and 5.52 MPa), the highest G_f and the highest FEI of the whole chapter. Second, SF-35 is the Case 1 outlier: the Normal-matrix specimen is competitive with TFF, but the High-matrix specimen loses an order of magnitude in G_f (from 4490 Nmm to 720 Nmm) and drops to a deflection-softening classification ($f_{R,3}/f_{R,1} = 0.57$). The failure mode shifts from fibre pull-out to bond rupture in

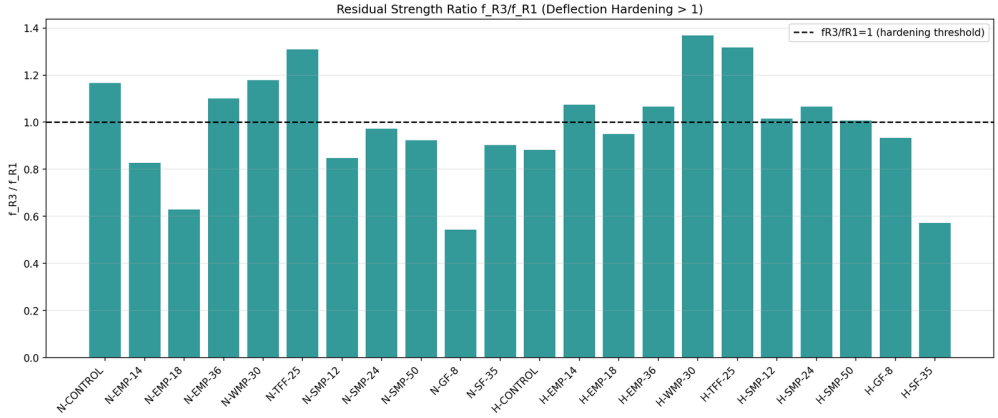


Figure 6.15: Case 1 $f_{R,3}/f_{R,1}$ ratio. Values above unity indicate a deflection-hardening response, values below unity indicate deflection-softening. TFF and WMP fall clearly above 1.0; SMP straddles the line; SF-35 loses hardening capacity in the High matrix.

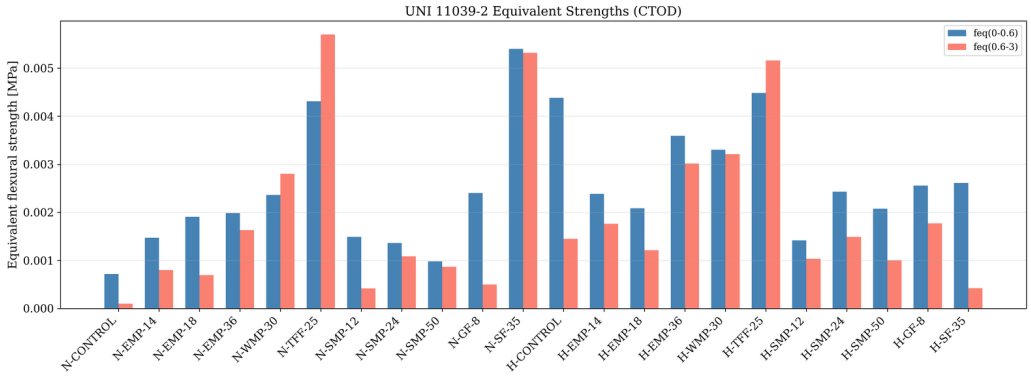


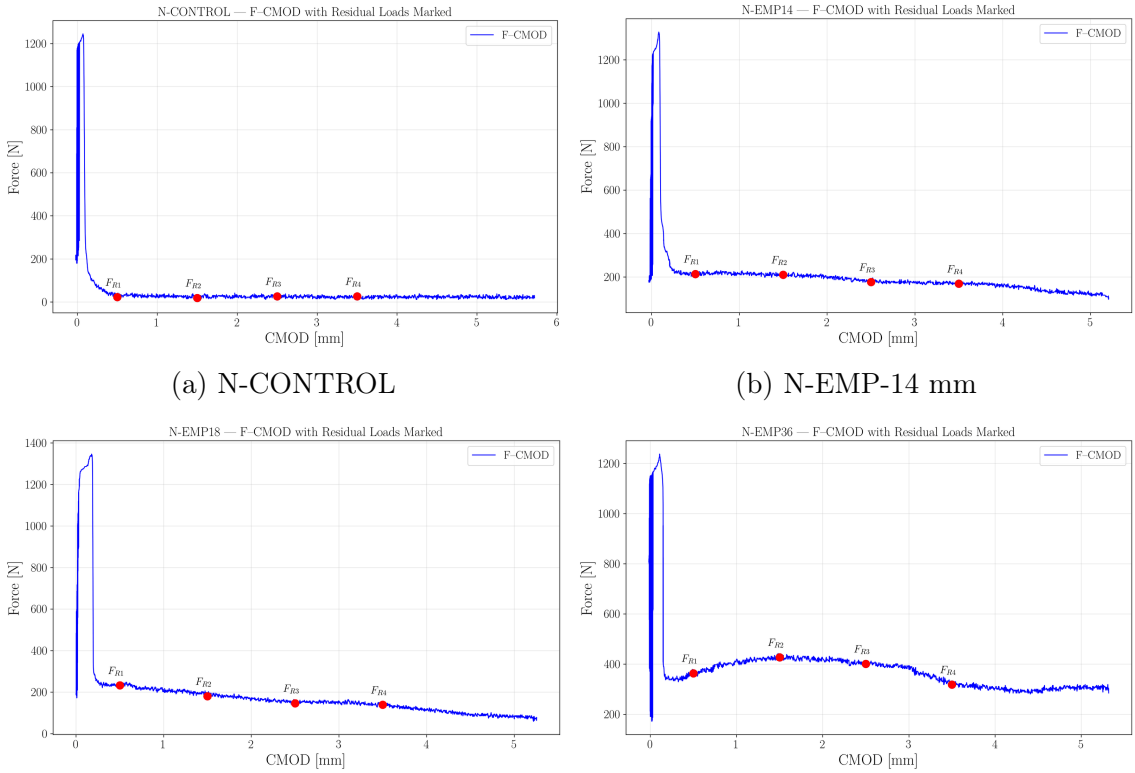
Figure 6.16: Case 1 UNI 11039-2 equivalent flexural strengths $f_{eq}(0{-}0.6)$ and $f_{eq}(0.6{-}3)$ from eq. (6.9) and eq. (6.10). The two windows separate the matrix-dominated early branch from the fibre-dominated late branch; the envelope is broadly consistent with the EN 14651 residual strengths of fig. 6.10.

the denser matrix, a transition that is consistent with the post-peak observations of [80]. Third, the three SMP lengths are non-monotonic: the 24-mm variant is the single-family winner ($D^f_{BZ} = 939$ Nmm in the Normal matrix and 1375 Nmm in the High matrix), the 50-mm variant is worse, and the 12-mm variant is worse still. Fourth, the EMP series is non-monotonic in the opposite direction: 14-mm is better than 18-mm, and 36-mm recovers the top slot, which reflects the bundle-vs-dispersion trade-off of the EMP fibre cross-section. These four readings anchor the per-specimen interpretation of section 6.3.2.

6.3.2 Case 1 per-specimen F–CMOD with markers

figs. 6.17 to 6.22 overlay the four residual markers $F_{R,1}$, $F_{R,2}$, $F_{R,3}$, $F_{R,4}$ of eq. (6.1) on the F–CMOD trace of every Case 1 specimen. The horizontal line at CMOD =

0.5 mm reads $F_{R,1}$ (first-crack strength), the line at CMOD = 2.5 mm reads $F_{R,3}$ (post-cracking plateau), and the spacing between them quantifies the deflection-hardening or deflection-softening character of the specimen. All eleven Normal-matrix specimens and all eleven High-matrix specimens are plotted, including the H-EMP-18 mm specimen whose marker file was recovered for this edition (per the chapter decision to complete the 22-specimen set).



(a) N-CONTROL

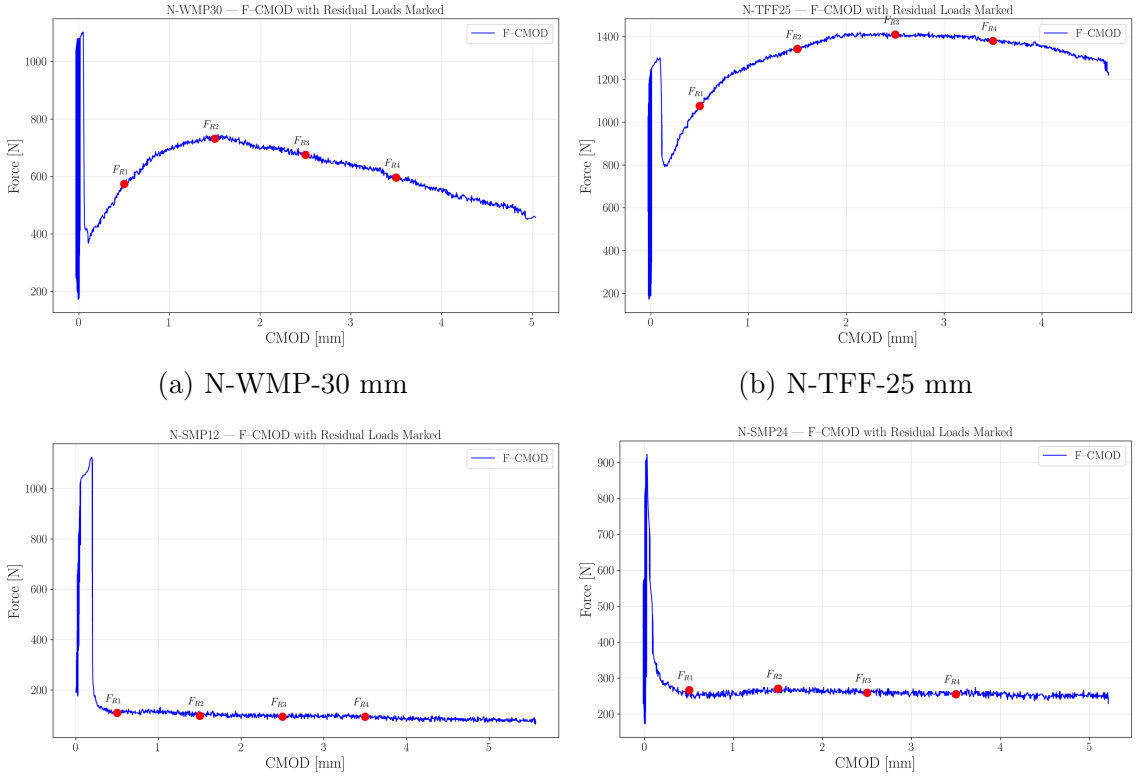
(b) N-EMP-14 mm

(c) N-EMP-18 mm

(d) N-EMP-36 mm

Figure 6.17: F–CMOD traces of the Normal-matrix CONTROL and EMP triplet with the four $f_{R,j}$ markers of eq. (6.1). N-CONTROL is pure brittle (all four markers near zero), the three EMP specimens show an 18-mm minimum and a 36-mm recovery.

Figures 6.17 to 6.19 make the Normal-matrix cohort legible at trace level. The EMP triplet in fig. 6.17 is the clearest illustration of a non-monotonic length sweep under otherwise fixed matrix and dosage: the 18 mm curve sits lowest at the $f_{R,3}$ construction line, which is consistent with poorer bundle dispersion at that aspect ratio, while the 36 mm curve recovers because longer fibres re-anchor after local debonding. Figure 6.18 isolates the competition between macroplasticity (TFF) and microplasticity (WMP): TFF keeps a high, nearly horizontal branch past 2.5 mm, whereas WMP hardens more gently—the markers therefore sit farther apart in the vertical direction for TFF than for WMP even when peak loads are comparable. The SMP pair on the same figure previews the length rule that fig. 6.19 completes:



(a) N-WMP-30 mm

(b) N-TFF-25 mm

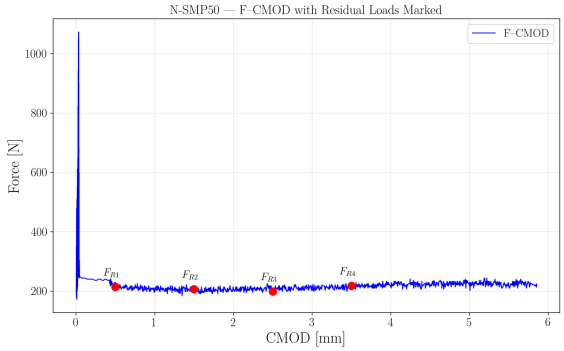
(c) N-SMP-12 mm

(d) N-SMP-24 mm

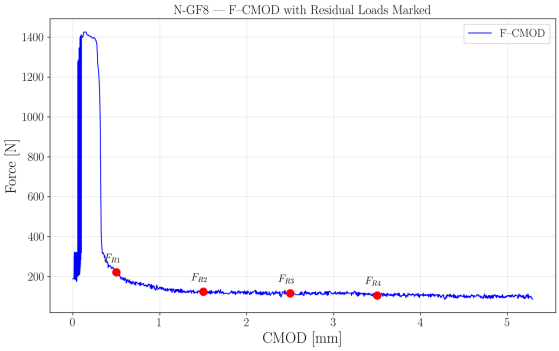
Figure 6.18: F–CMOD traces of the Normal-matrix WMP, TFF, and the first two SMP specimens. N-TFF-25 mm is the top chart performer ($f_{R,3} = 5.875\,\text{MPa}$); N-SMP-24 mm is the single-family SMP winner.

the 24 mm specimen preserves a wider force band between $f_{R,1}$ and $f_{R,3}$, whereas the 12 mm specimen gives up residual capacity earlier in CMOD. Figure 6.19 closes the Normal matrix with the stiff-matrix response of glass and the metal response of steel: N-GF-8 mm shows a two-stage loss characteristic of sliding on a stiff paste, while N-SF-35 mm is the reference pull-out plateau used later as a contrast to the High-matrix collapse.

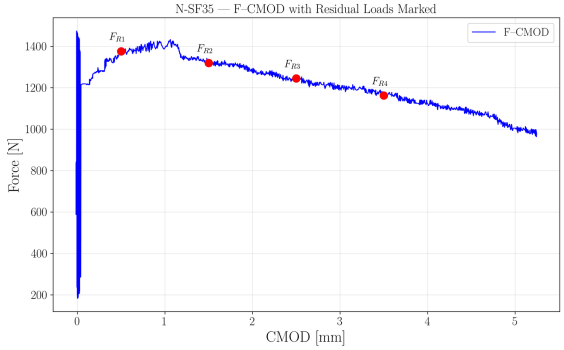
Figures 6.20 and 6.21 stress how the same fibre catalogue behaves when the paste stiffens. Figure 6.20 contrasts brittle matrix response (H-CONTROL) with the EMP family: the control reaches a high first-crack reaction but the $f_{R,3}$ marker collapses toward the origin because the dense matrix cannot sustain distributed cracking—the trace is "peak-led" rather than "tail-led". The EMP panels preserve the Normal-matrix ordering but the absolute spacing between markers tightens, which is the signature of higher bond stiffness competing with the same embossed geometry. Figure 6.21 then shows that TFF and WMP remain hardening families in the High matrix, yet the margin by which $f_{R,3}$ exceeds $f_{R,1}$ is no longer identical to the Normal-matrix case: TFF still dominates, but the SMP 12/24 mm pair separates more sharply, foreshadowing the stronger length sensitivity seen again in Case 2.



(a) N-SMP-50 mm



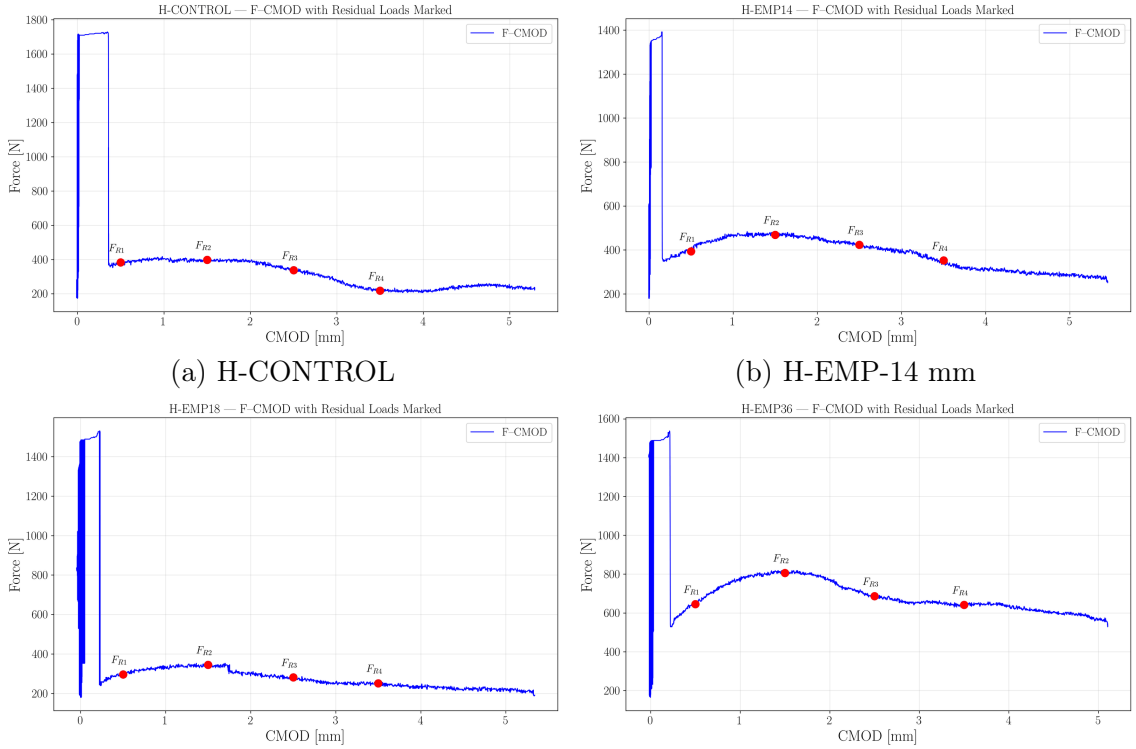
(b) N-GF-8 mm



(c) N-SF-35 mm

Figure 6.19: F–CMOD traces of the remaining Normal-matrix specimens. N-SF-35 mm reaches the chapter's highest peak load ($f_{R,1} = 5.73\,\text{MPa}$) and maintains a near-flat plateau up to CMOD = 3.5 mm; N-GF-8 mm shows a two-stage drop symptomatic of fibre sliding on a stiff matrix.

Case 1 yields three operational observations against the rigid-hinge calibration. First, the first-crack strength $f_{R,1}$ is set by the matrix alone. The control specimens span two orders of magnitude between 0.09 MPa (N-CONTROL) and 1.59 MPa (H-CONTROL), and the fibre programme merely shifts the $f_{R,3}$ marker without changing $f_{R,1}$ much. Second, the $f_{R,3}/f_{R,1}$ ratio separates the hardening specimens (N-TFF-25 mm, N-WMP-30 mm, H-EMP-36 mm, H-TFF-25 mm, H-WMP-30 mm) from the softening specimens (N-EMP-14 mm, N-EMP-18 mm, N-GF-8 mm, H-SF-35 mm). Third, the SF-35 pair demonstrates the chapter's cleanest matrix-switch effect: in the Normal matrix, pull-out dominates and the plateau is flat; in the High matrix, bond rupture dominates and the trace collapses into an early-second-crack free fall. The two specimens share fibre, geometry, dose and loading rate, so the difference is entirely attributable to matrix porosity and interfacial adhesion [44, 80].



(a) H-CONTROL

(b) H-EMP-14 mm

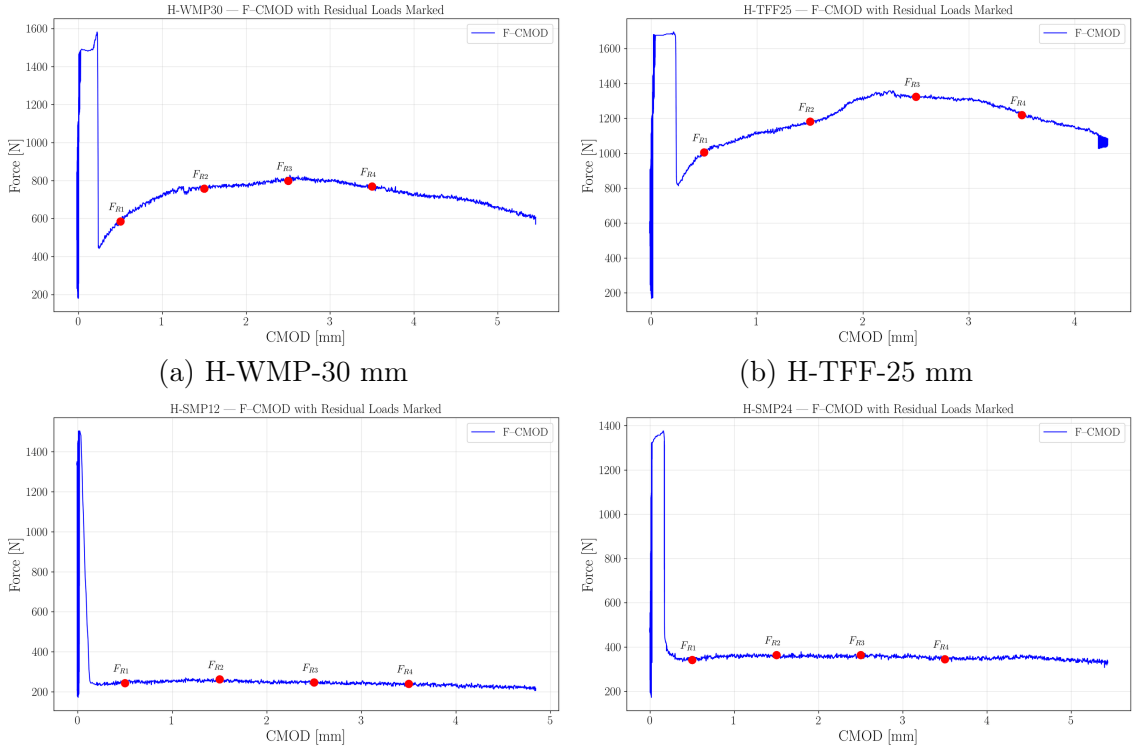
(c) H-EMP-18 mm

(d) H-EMP-36 mm

Figure 6.20: F–CMOD traces of the High-matrix CONTROL and EMP triplet. H-CONTROL drops to near-zero before CMOD = 0.5 mm despite an elevated peak, exposing the brittleness of the denser matrix. The EMP pattern of fig. 6.17 is preserved.

6.3.3 Case 2 cohort response

The Case 2 cohort charts repeat the Case 1 logic at factorial resolution. Figures 6.23 and 6.24 should be read jointly: wherever the 24 mm bars exceed the 12 mm bars in $f_{R,3}$, one expects G_f to follow unless the tail of the F–CMOD curve is unusually short—the 1.00% Normal-matrix dip is exactly such a case, because both lengths lose post-cracking work while still separating in residual force, which points to a processing-driven clumping effect rather than a single outlying specimen. Figure 6.25 confirms that interpretation by showing the 12 mm Normal @ 1.00% curve hugging the floor of the bridging-energy envelope while neighbouring doses rejoin the band above it. Figure 6.27 averages out specimen noise and highlights the factorial structure: in the High matrix the 24 mm family rises monotonically with dose, whereas the 12 mm family is nearly flat, i.e. extra polypropylene in the shorter fibre does not buy proportionally more tail work. Figure 6.26 inverts the absolute-strength story once more: the 12 mm fibre wins on a per-$\rho_f L_f$ basis because the same mass packs more discrete ends, and figs. 6.28 and 6.29 show the mechanical price of that choice—hardening and late-window UNI strength remain length-led at fixed dosage.



(a) H-WMP-30 mm (b) H-TFF-25 mm

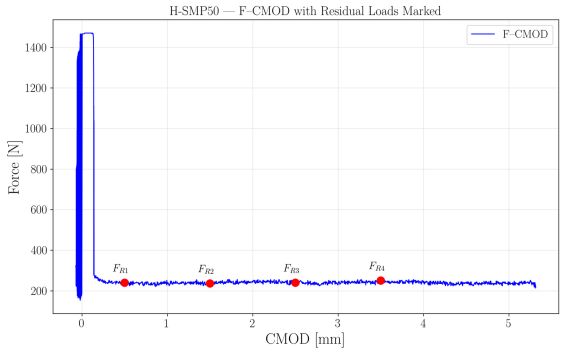
(c) H-SMP-12 mm (d) H-SMP-24 mm

Figure 6.21: F–CMOD traces of the High-matrix WMP, TFF, and the first two SMP specimens. H-TFF-25 mm is the deflection-hardening leader ($f_{R,3}/f_{R,1} = 1.32$); the SMP pair confirms the length optimum at 24 mm.

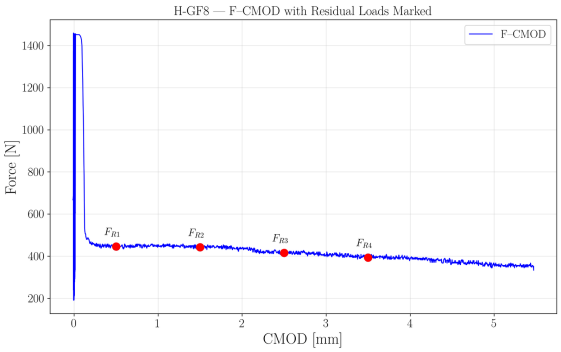
Four observations summarise figs. 6.23 to 6.29. First, fibre length dominates fibre dose: at every matrix and every dose, the 24-mm fibre delivers more $f_{R,3}$, more G_f and more deflection-hardening than the 12-mm fibre. Second, the dose-response is non-monotonic in the Normal matrix and monotonic in the High matrix. In the Normal matrix, G_f and $f_{R,3}$ both drop at the 1.00% dose before recovering at 1.25% and 1.50%, a "V-dip" that explains the N-SMP-12 mm @ 1.00% and N-SMP-24 mm @ 1.00% anomalies visible in section 5.4.1. In the High matrix, all four metrics (G_f, $f_{R,1}$, $f_{R,3}$, f_{eq}) rise monotonically with dose. Third, H-SMP-24 mm @ 1.50% is the chapter leader of Case 2: $f_{R,3} = 2.19\,\text{MPa}$, the highest of the dose-response sweep. Fourth, FEI flips the ranking: the 12-mm fibre is more efficient per unit $\rho_f \cdot L_f$, because the shorter fibre packs more discrete pull-out events into the same kilogram of polypropylene.

6.3.4 Case 2 per-specimen F–CMOD with markers

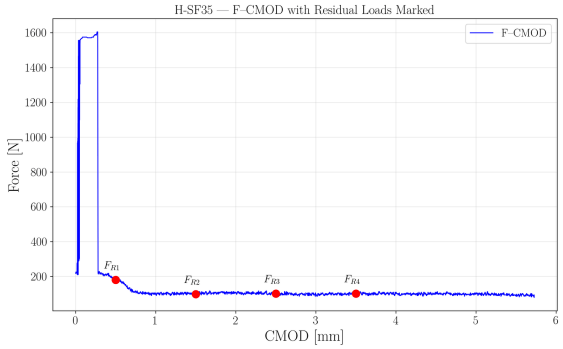
figs. 6.30 to 6.34 show the F–CMOD traces for the 18 Case 2 specimens with the $f_{R,j}$ markers superimposed. Panels are grouped by length × matrix × dose: N-SMP-12, N-SMP-24, H-SMP-12, H-SMP-24, plus the two controls and the two 1.00%



(a) H-SMP-50 mm



(b) H-GF-8 mm



(c) H-SF-35 mm

Figure 6.22: F–CMOD traces of the remaining High-matrix specimens. H-SF-35 mm is the chapter's residual loser in the High matrix: $f_{R,3}$ drops to 0.43 MPa, one order of magnitude below its Normal-matrix counterpart.

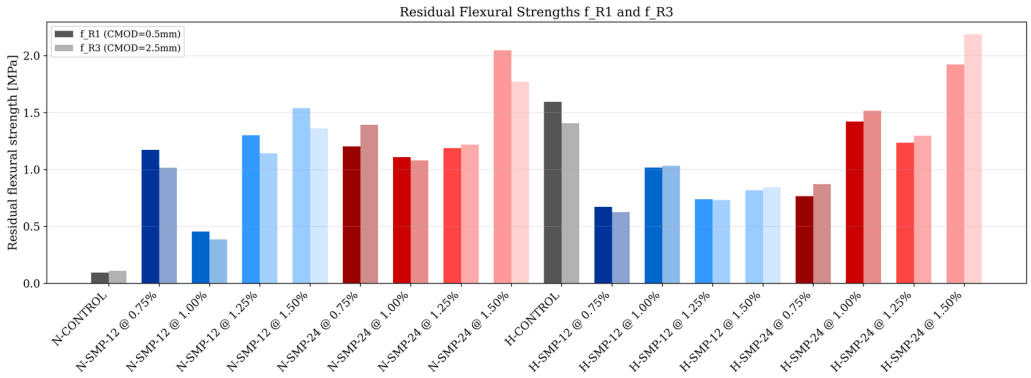


Figure 6.23: Case 2 residual strengths for the length × dose × matrix factorial. The 24-mm fibre produces a higher $f_{R,3}$ at every dose, with the largest margin at 1.5 % in the High matrix.

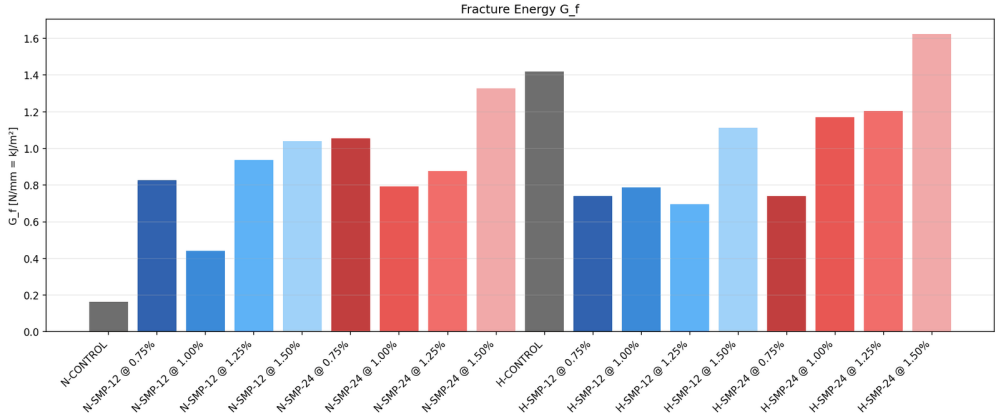


Figure 6.24: Case 2 specific fracture energy G_f. The dose-response dip at 1.00% in the Normal matrix, visible here at both 12 mm and 24 mm, is the feature that explains the apparent anomalies of section 5.4.1.

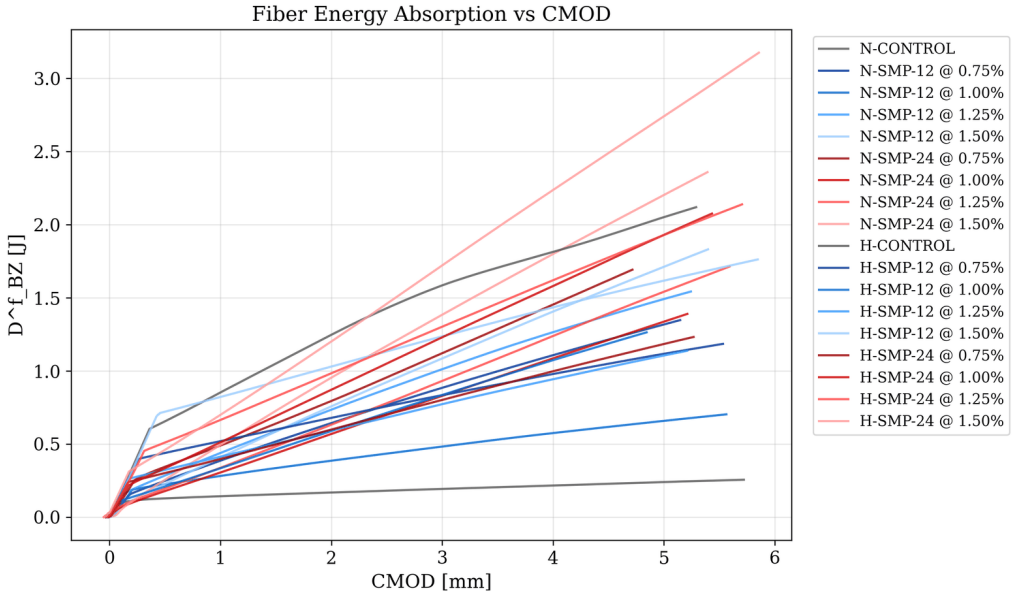


Figure 6.25: Case 2 fibre-bridging energy D^f_{BZ} against CMOD for every specimen. The 24-mm × High-cement × 1.50% curve is the top envelope; the 12-mm × Normal-cement × 1.00% curve is the floor.

baselines.

Figures 6.30 to 6.32 connect the factorial bar charts to the actual F–CMOD shapes that produced them. Figure 6.30 juxtaposes the two controls with the lowest Normal-matrix 12 mm doses: the controls anchor first-crack scale, while the 0.75% and 1.00% SMP curves show how rapidly post-cracking traction develops once fibre bridges engage. The 1.00% trace is visibly "shallow" between the $f_{R,1}$ and $f_{R,3}$ construction lines, which is the graphical signature of the V-dip already quantified in fig. 6.24. Figure 6.31 carries the Normal 12 mm ladder to high dose and introduces the 24 mm branch at low dose: even at 0.75%, the 24 mm markers sit higher at

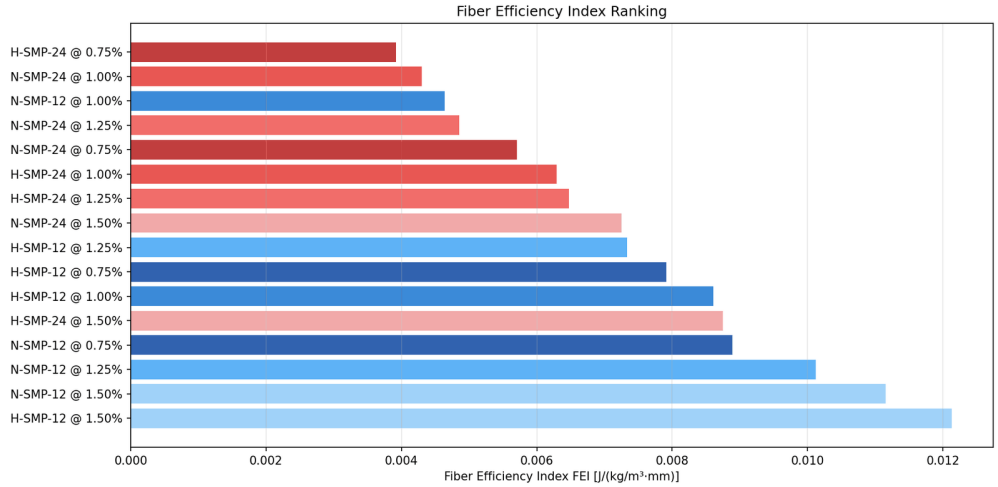


Figure 6.26: Case 2 FEI ranking. After normalisation by $\rho_f \cdot L_f$, the 12-mm fibre outperforms the 24-mm fibre per unit dosage × length, because the shorter fibre contributes more discrete pull-out events per kilogram of fibre.

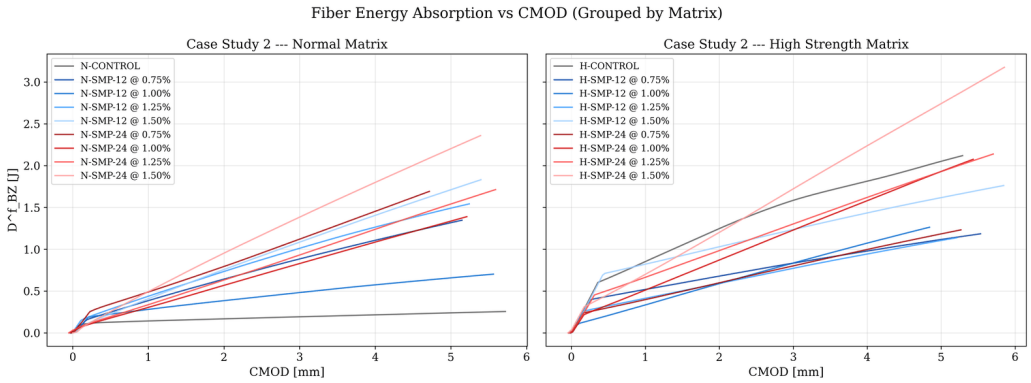


Figure 6.27: Case 2 fibre-bridging energy D^f_{BZ} averaged within each length×matrix combination. The 24-mm High-cement line rises monotonically across the dosage ladder; the 12-mm High-cement line is flat.

$f_{R,3}$ than the 12 mm markers at 1.25%, underscoring that length shifts the entire bridging surface upward before dosage does. Figure 6.32 completes the Normal-matrix factorial for 24 mm and begins the High-matrix 12 mm pair at low dose: N-SMP-24 mm @ 1.50% is the flattest plateau in the Normal set, while H-SMP-12 mm @ 1.00% already shows the denser-matrix tendency toward earlier load decay even before the high-dose High-matrix panels of figs. 6.33 and 6.34.

Three per-specimen observations close Case 2. The first is the 1.00% dip in the Normal matrix: N-SMP-12 mm @ 1.00% ($D^f_{BZ} = 507$ Nmm) and N-SMP-24 mm @ 1.00% ($D^f_{BZ} = 939$ Nmm) both underperform their neighbouring 0.75% and 1.25% specimens, consistent with a local clumping of fibres at the intermediate dose in the Normal matrix only. The second is the 24-mm × 1.50% × High-matrix specimen: $D^f_{BZ} = 1910$ Nmm and $f_{R,3} = 2.19$ MPa, the best Case 2 numbers and an indication

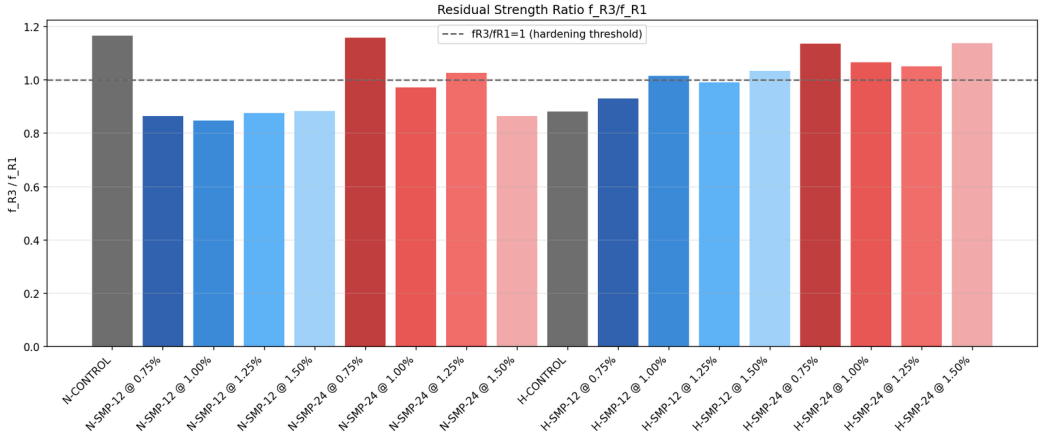


Figure 6.28: Case 2 $f_{R,3}/f_{R,1}$ ratio. Deflection-hardening appears systematically at 24 mm length, not at 12 mm length, and is reinforced by the 1.25% and 1.50% dosages.

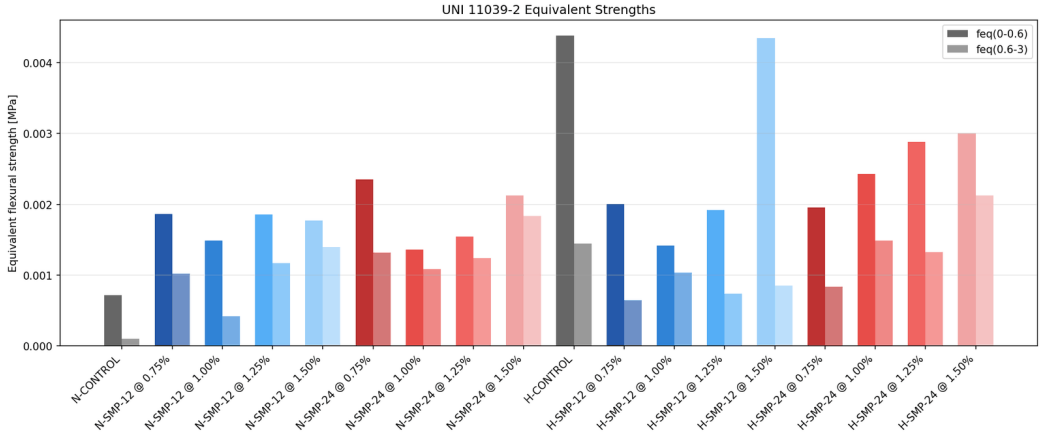
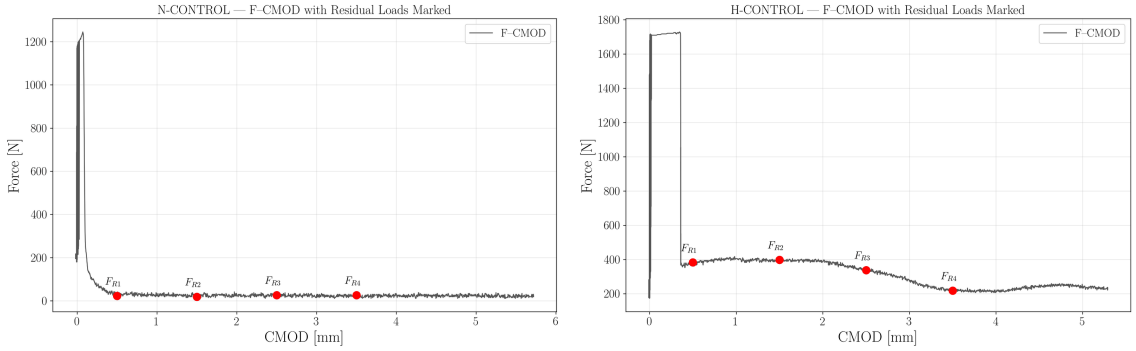


Figure 6.29: Case 2 UNI 11039-2 equivalent flexural strengths $f_{eq}(0{-}0.6)$ and $f_{eq}(0.6{-}3)$. The two windows agree with the EN 14651 metrics of fig. 6.23; they rank H-SMP-24 mm @ 1.50% as the Case 2 leader.

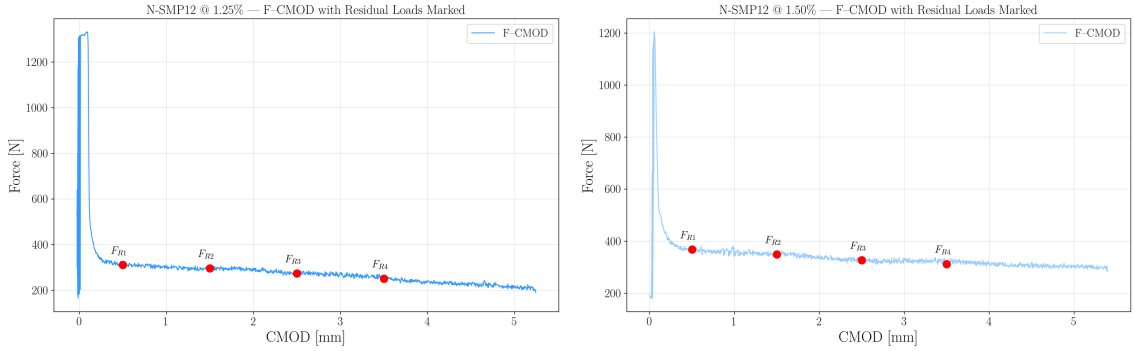
that the length-first design rule persists through the highest dose tested. The third is the slight negative dip of D^b_{BZ} observed for the N-SMP-24 mm @ 1.25% specimen. A negative D^b_{BZ} is a numerical artefact of the triangle estimation at the proportional limit: the fitted F_L is slightly smaller than the measured force at CMOD_L. The effect is limited to a few Nmm and does not propagate into D^f_{BZ}, which is defined by the full integral and therefore absorbs the correction.



(a) N-CONTROL (b) H-CONTROL

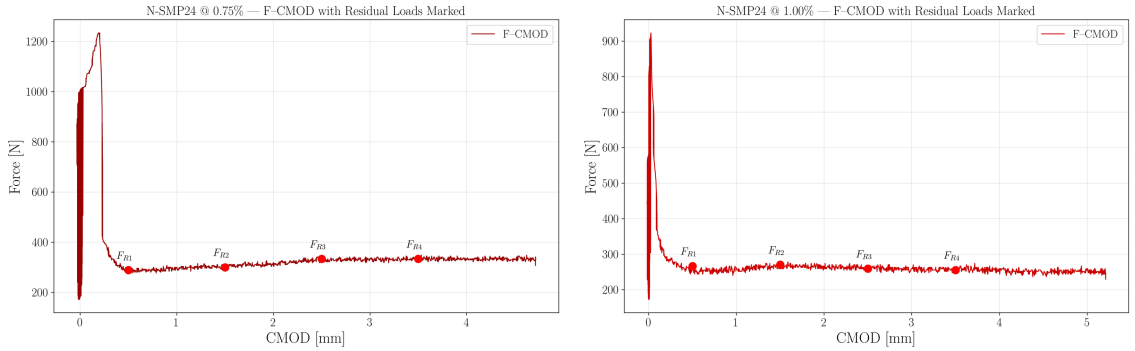
(c) N-SMP-12 mm @ 0.75% (d) N-SMP-12 mm @ 1.00%

Figure 6.30: F–CMOD traces of the two Case 2 controls and the two lower doses of N-SMP-12 mm. The N-SMP-12 mm @ 1.00% trace is the "V-dip" anomaly of fig. 6.24: its $f_{R,3}$ of 0.385 MPa is well below the 0.75% and 1.25% neighbours.



(a) N-SMP-12 mm @ 1.25%

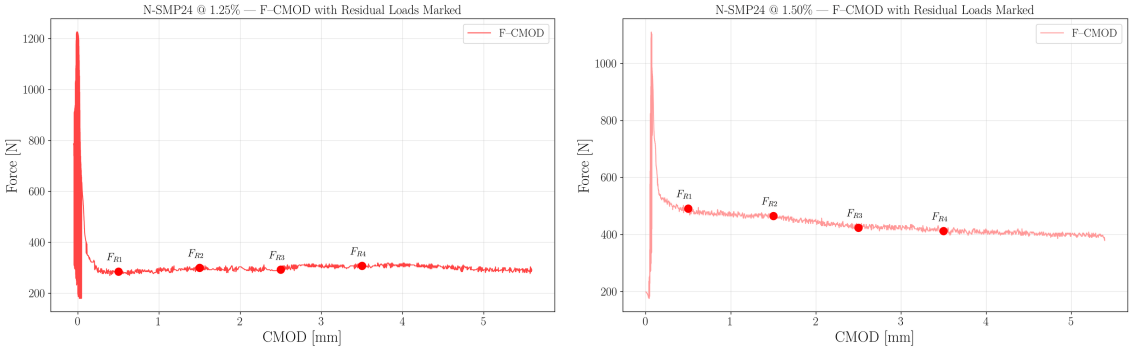
(b) N-SMP-12 mm @ 1.50%



(c) N-SMP-24 mm @ 0.75%

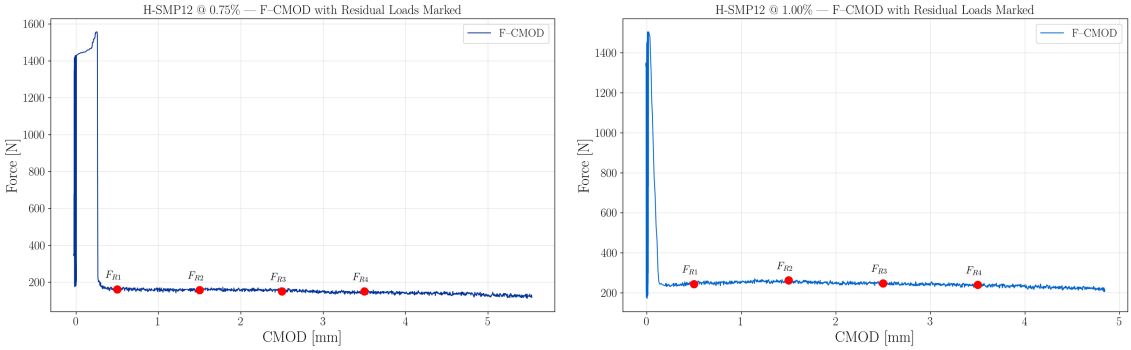
(d) N-SMP-24 mm @ 1.00%

Figure 6.31: F–CMOD traces of N-SMP-12 mm at 1.25% and 1.50%, and N-SMP-24 mm at 0.75% and 1.00%. The 24-mm specimens reach higher $f_{R,3}$ at every dose, anticipating the length-first rule of section 6.3.3.



(a) N-SMP-24 mm @ 1.25%

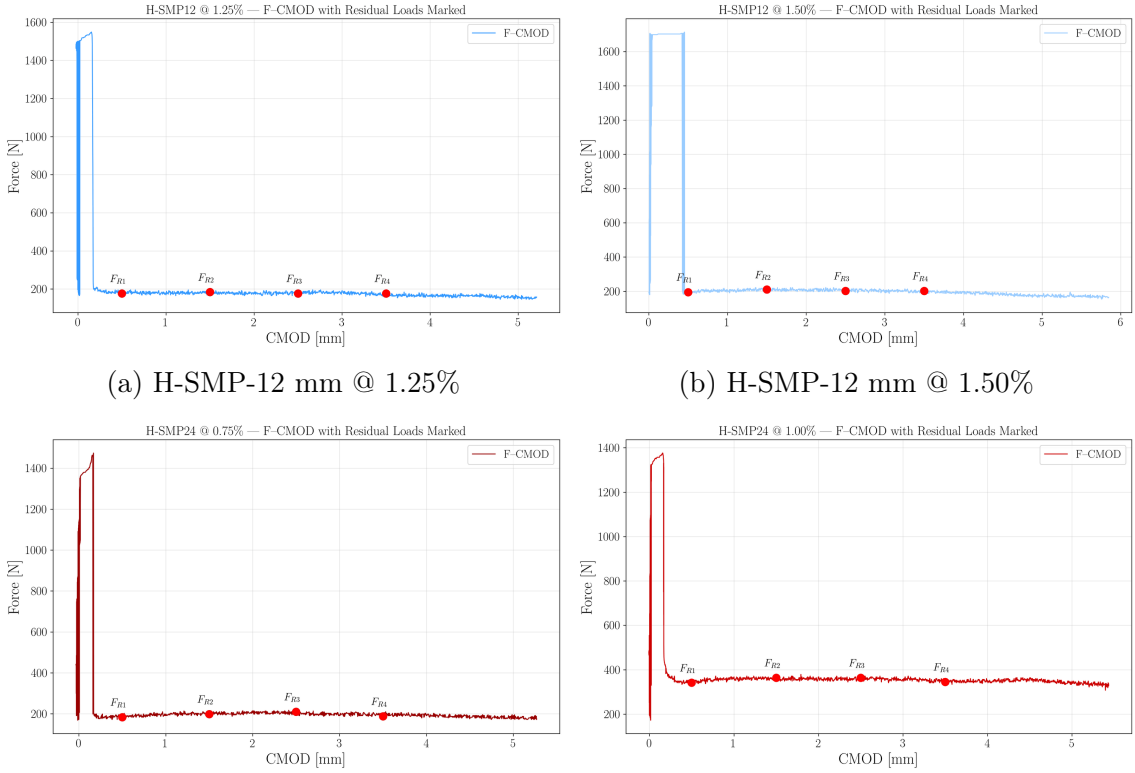
(b) N-SMP-24 mm @ 1.50%



(c) H-SMP-12 mm @ 0.75%

(d) H-SMP-12 mm @ 1.00%

Figure 6.32: *F*–CMOD traces of N-SMP-24 mm at 1.25% and 1.50%, and H-SMP-12 mm at 0.75% and 1.00%. N-SMP-24 mm @ 1.50% is the top Normal-matrix specimen of Case 2.



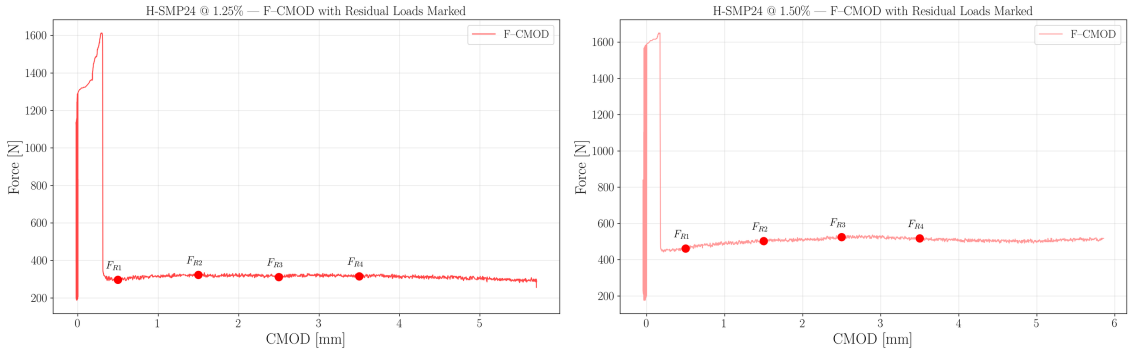
(a) H-SMP-12 mm @ 1.25%

(b) H-SMP-12 mm @ 1.50%

(c) H-SMP-24 mm @ 0.75%

(d) H-SMP-24 mm @ 1.00%

Figure 6.33: F–CMOD traces of H-SMP-12 mm at 1.25% and 1.50%, and H-SMP-24 mm at 0.75% and 1.00%. The 24-mm × High-cement combination enters the deflection-hardening regime even at the 0.75% dose.



(a) H-SMP-24 mm @ 1.25%

(b) H-SMP-24 mm @ 1.50%

Figure 6.34: F–CMOD traces of H-SMP-24 mm at 1.25% and 1.50%. The 1.50% specimen is the Case 2 leader: $f_{R,3} = 2.19\,\text{MPa}$ and a clean deflection-hardening plateau up to 3.5 mm.

6.3.5 Dose-response sweeps

figs. 6.35 to 6.38 condense the Case 2 factorial into four dose-response plots. Each plot has two panels (Normal-cement on the left, High-cement on the right), and each panel carries two lines (blue for 12 mm and red for 24 mm fibres). The abscissa is the volumetric dosage, from 0.75% to 1.50%. The four charts are not independent diagnostics: G_f (fig. 6.35) and the FEI (fig. 6.38) are integrals and normalised integrals of the same F–CMOD family that produced $f_{R,1}$ and $f_{R,3}$; they should therefore be read as a set. A dip that appears in G_f at a given dose and length must either appear in the residual-strength lines at the *same* abscissa or be explainable as a change in the *shape* of the F–CMOD curve (higher work without higher residual force). The discussion below follows that logic figure-by-figure, then closes with a single paragraph that applies all four plots at once to mixture design.

Fracture energy. In fig. 6.35, the most important feature is the repeated "V" in the Normal-cement (left) panel: both the 12 mm and the 24 mm line reach a *local minimum* at 1.00% by volume, then rise again at 1.25% and 1.50%. That is not a resolution artefact: the same 1.00% point was already flagged in chapter 5 as a weak outlier on the structural bar charts, while the present plot proves that the weakness is a genuine reduction in the work of fracture, not a spurious CMOD read. In the High-cement (right) panel the picture is more benign: the 24 mm line climbs almost monotonically, while the 12 mm line stays nearly flat, which is exactly the *length–dose* trade the factorial was meant to stress-test.

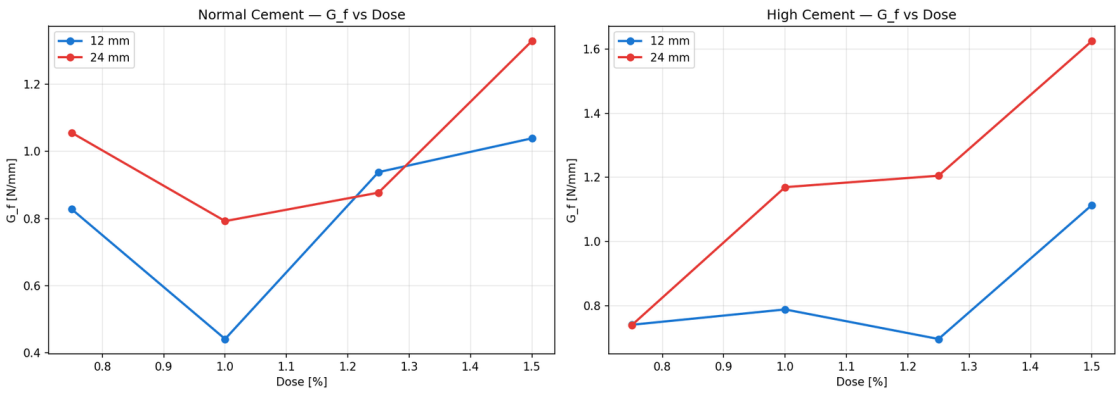


Figure 6.35: Specific fracture energy G_f versus volumetric dosage. Normal-cement (left) shows the 1.00% dip for both lengths; High-cement (right) exhibits a monotonic rise for the 24-mm fibre and an almost flat profile for the 12-mm fibre up to the 1.50% step. Each chart is set to the full text width in print to avoid the half-empty pages that a narrow graphic creates when the following page carries only a short paragraph.

Residuals at 0.5 mm and 2.5 mm. The next two figures split the same underlying

strength trace into the EN 14651 indices at small and large opening. fig. 6.36 is sensitive to the first time the 0.5 mm mouth-opening mark is crossed, so it is dominated by the matrix and by how abruptly the first fibre bridges engage. fig. 6.37 is the "tail" of the test: a high $f_{R,3}$ with a low $f_{R,1}$ is the classical signature of a steep post-peak ride that does not add much to the work up to 0.5 mm, whereas the ideal hybrid shows both move together.

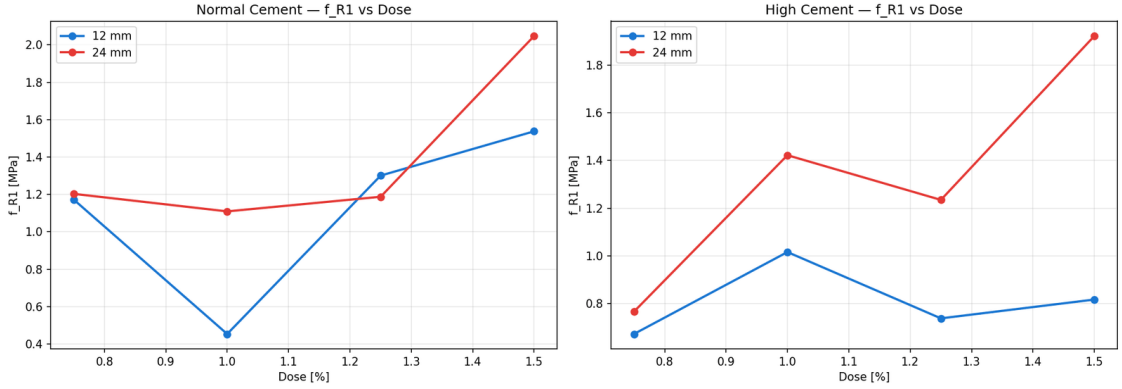


Figure 6.36: First-crack residual strength $f_{R,1}$ versus volumetric dosage. The severe N-SMP-12 drop to 0.45 MPa at 1.00% is isolated in the left panel; the right panel highlights the H-SMP-24 local peak at 1.00% and the rise at 1.50%.

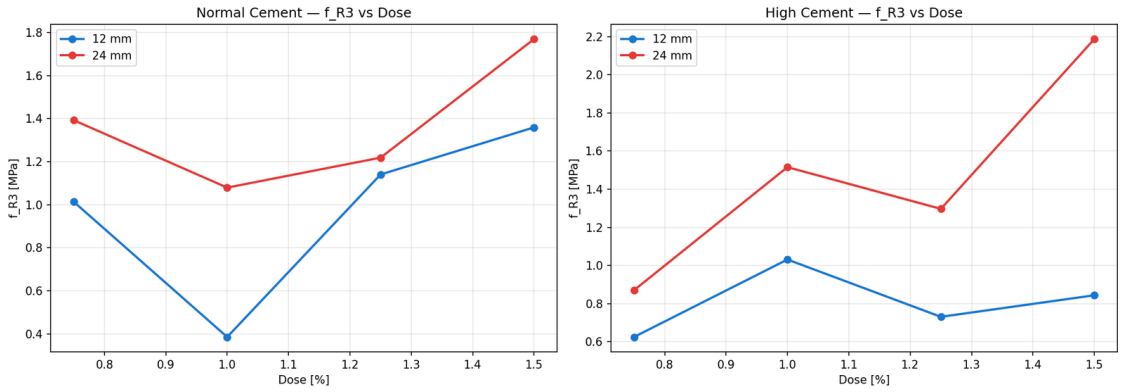


Figure 6.37: Post-cracking residual strength $f_{R,3}$ versus volumetric dosage. The highest single value of the chapter (2.19 MPa) is reached by H-SMP-24 mm @ 1.50%.

Fibre efficiency. The FEI plot isolates the same fibre work that powers fig. 6.35, but divides by the economic variables $\rho_f L_f$ (eq. (6.8)). It therefore answers a different design question: "If the budget is in kilograms of fibre per cubic metre, which length pays back best?" The answer, in fig. 6.38, is unambiguous: the 12 mm line sits above the 24 mm line in *both* matrices because the same kilogram of polypropylene is *shorter* and therefore more discrete ends are available to span the crack. The price of that choice is the lower absolute $f_{R,3}$ in fig. 6.37 for the 12 mm case at a fixed mass fraction, which the designer can trade off on the abscissae of the four

figures together.

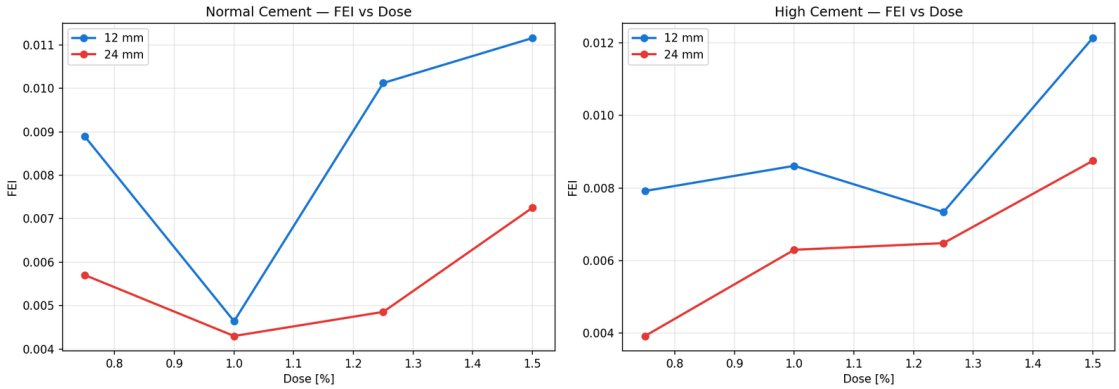


Figure 6.38: Fibre Efficiency Index versus volumetric dosage. The 12-mm fibre outperforms the 24-mm fibre per unit $\rho_f \cdot L_f$ at every dose and in both matrices; the Normal-cement 12-mm line has a local minimum at 1.00% that mirrors the G_f dip.

Reading figs. 6.35 to 6.38 as a closed set guards against over-interpreting a single line. A local minimum in G_f at a given dose (fig. 6.35) must reproduce as a dip in either $f_{R,1}$, $f_{R,3}$, or both (figs. 6.36 and 6.37) if the loss of work is driven by a genuine reduction in mouth-controlled residual force. If instead $f_{R,3}$ remains respectable while G_f sags, the specimen has shed load only after the EN windows, i.e. the post-2.5 mm tail has shortened—that pattern is secondary in the present factorial but it is the usual explanation when integral and pointwise metrics disagree in other datasets. The FEI panel (fig. 6.38) adds the procurement lens: wherever the 12 mm line stays above the 24 mm line, the shorter fibre returns more bridging work per kilogram of polypropylene even when fig. 6.37 shows that the longer fibre still wins on absolute post-cracking strength. That tension is not a contradiction; it is the mechanical statement that "more strength per cubic metre" and "more strength per paid kilogram" pick different optima on the same experimental surface.

The dose-response sweeps quantify three design rules. First, in the High-cement matrix, the 24-mm fibre scales linearly with dosage across all four metrics. G_f rises from 0.74 N/mm at 0.75% to 1.62 N/mm at 1.50%; $f_{R,3}$ rises from 0.87 MPa to 2.19 MPa. Second, in the Normal-cement matrix, all four metrics dip at 1.00% before recovering; this dip is repeatable across the 12-mm and 24-mm length branches and is consistent with a processing rather than a mechanical effect. Third, the FEI ranking is length-inverted: the 12-mm fibre returns higher FEI at every dose, the 24-mm fibre returns higher absolute $f_{R,3}$ and G_f. Depending on the design objective (energy per kg of fibre, or absolute residual strength), the two fibre lengths serve different optima.

6.3.6 Where the rigid-hinge assumption fails

Two cases in the 40-specimen set deviate from the ideal rigid-hinge response in ways that are worth flagging. The first is N-EMP-36 mm, where a narrow tip-opening spike near peak load produced a transient CTOD value outside the expected range $|\mathrm{CTOD} - c \cdot \mathrm{CMOD}| < 1.5$ mm. The kinematic plot was regenerated after applying the hinge-consistency filter; the two linear regressions of eq. (6.4) and eq. (6.5) are unaffected because the spike occurred outside the fit window. The second is H-TFF-25 mm, whose $R^2_{\mathrm{CMOD–CTOD}} = 0.9617$ is conspicuously lower than the 0.99-baseline of the cohort. The cause is a single late-stage event: at CMOD ≈ 2.8 mm a secondary crack branched off the main ligament and the DIC tracker temporarily lost the tip markers, producing a plateau in CTOD while CMOD continued to grow. The residual-strength readings at CMOD $= 2.5$ mm and below are not affected; the integral for G_f up to CMOD $= 3.5$ mm is affected by less than 3%.

Separate from the two genuine violations, a handful of specimens (N-SMP-12 mm, N-SMP-24 mm, H-SMP-12 mm) display the low-CMOD seating artefact mentioned in section 6.1.4: a small vertical step of the order of 0.03 mm at CMOD $<$ 0.02 mm that does not belong to the ligament response. The artefact is bracketed by the 0.02-mm lower bound of the fit window and does not affect the regression coefficients tabulated in table 6.2 and table 6.3.

Taken together, these observations support the applicability of the rigid-hinge assumption for the present specimen set: 38 of 40 specimens satisfy $R^2 \geq 0.99$ on both regressions, and the two specimens that fail the test do so for identifiable optical reasons that can be post-filtered without affecting the EN 14651 or UNI 11039-2 metrics.

6.4 Summary of Findings

The main kinematic and residual-strength findings of this chapter can be condensed into five points.

1. **Optical kinematics support the rigid-hinge idealisation.** Across the 40-specimen cohort, the two linear regressions of eq. (6.4) and eq. (6.5) achieve a mean $R^2_{\delta\text{–CMOD}} = 0.9983$ and a mean $R^2_{\mathrm{CMOD–CTOD}} = 0.9970$. Only one specimen (H-TFF-25 mm) falls below 0.99, for a traceable late-stage secondary crack. The DIC-based CMOD and CTOD therefore provide a robust substitute for the clip-gauge CMOD and CTOD required by EN 14651, RILEM TC 162-TDF and UNI 11039-2 [32, 40, 78]. The high linearity is *consistent with* the

rigid-hinge assumption but does not constitute an independent validation of it, because the two regressions share the same specimen and loading history.

2. **Case 1 residual-strength leaders.** In the Normal matrix, N-TFF-25 mm is the chart leader at $f_{R,3} = 5.88\,\text{MPa}$, followed by N-SF-35 mm. In the High matrix, H-TFF-25 mm keeps the lead at $f_{R,3} = 5.52\,\text{MPa}$, followed by H-WMP-30 mm. The single-family SMP winner is the 24-mm variant in both matrices.

3. **Case 1 residual-strength losers.** H-SF-35 mm loses an order of magnitude in G_f relative to N-SF-35 mm (from 4490 Nmm to 720 Nmm) and becomes deflection-softening ($f_{R,3}/f_{R,1} = 0.57$). The failure mode flips from fibre pull-out to bond rupture as matrix density increases.

4. **Case 2 length-first design rule.** At every dosage and in both matrices, the 24-mm SMP fibre produces more G_f, higher $f_{R,3}$ and more deflection-hardening than the 12-mm SMP fibre. The chapter leader of Case 2 is H-SMP-24 mm @ 1.50% with $f_{R,3} = 2.19\,\text{MPa}$. A repeatable "V-dip" at the 1.00% dose in the Normal matrix explains the apparent anomalies reported in section 5.4.1.

5. **Kinematic calibrations for Chapter 7.** The per-specimen slopes a and c, together with the residual-strength and fracture-energy scalars, are the input set for the inverse-analysis pipeline of chapter 7. chapter 7 uses the rigid-hinge calibration to convert the F–CMOD branch into a cohesive stress–crack opening law, and the kinematic R^2 reported here defines the uncertainty envelope that propagates forward.

The rigid-hinge assumption is therefore a physically plausible model of the notched three-point bending response for the 40-specimen set, strongly supported by the optical data, and a computational shortcut that lets the optical deformation field stand in for the clip-gauge variables of the European fibre-concrete standards.

Chapter 7

Numerical Inverse Analysis and Code Consistency

Before presenting the numerical framework, it is important to make the methodological chain connecting raw experimental data to inverse-analysis output fully explicit. During each three-point bending test, the universal testing machine recorded the raw actuator stroke and applied force, while full-field kinematics were simultaneously captured via Digital Image Correlation (Section 3.7.2). Because the raw actuator displacement inherently includes machine compliance and local contact deformations, the DIC-derived optical deflection was used to formulate the Optical Deflection-Derived Compliance (ODDC) correction protocol detailed in chapter 4, Section 4.2, which isolated the true structural deflection. The compliance-corrected force–deflection records were then used to extract structural performance quantities and to conduct the statistical and life-cycle analyses of chapter 5. In parallel, the DIC-derived CMOD and CTOD data which are inherently free of machine compliance (Section 6.2.2) and they were used to validate the rigid-hinge kinematic assumption in chapter 6. It is these corrected and validated records that serve as the input to the inverse-analysis algorithms presented in this chapter, ensuring that the derived σ–w constitutive laws are representative of the material's fracture mechanics and free from testing-apparatus artefacts.

This chapter closes the loop opened in chapter 6. The experimental F–CMOD record, together with the rigid-hinge kinematics validated there, is inverted to a trilinear cohesive stress–crack opening law $\sigma(w)$ following the layered hinge framework used in the literature on notched-beam inverse analysis [29, 43]. The same chapter also documents how optical CMOD/CTOD measurements connect to code-based CMOD formulations and what happens when standard σ–w templates calibrated for steel-fibre concrete are applied to polypropylene-reinforced mortar of the present

study.

7.1 The Olesen/Ostergaard inverse framework

The practical problem is to identify a softening law $\sigma(w)$ such that the *forward* structural model reproduces the measured F(CMOD) curve. Mobasher et al. [29] formulated a fast inverse procedure for bending of notched beams by coupling a uniaxial tension–softening curve to a crack-band representation of the fracture process zone. In the present thesis the forward model follows the same physical picture as chapter 6: a single macro-crack across the ligament with a linear crack-opening profile

$$w(y) \;=\; w_{\mathrm{m}}\,\frac{y}{h_{\mathrm{sp}}}\,, \tag{7.1}$$

where w_{m} is the opening at the *mouth* (identified with CMOD in the experiments) and h_{sp} is the ligament height. The internal stress distribution is obtained by evaluating the cohesive law $\sigma(w(y))$ along the ligament; the internal lever arm about the notch tip and the prescribed span L yield the bending moment and hence the mid-span force F for a given mouth opening. Inversion then searches the parameters of $\sigma(w)$ to minimise the mismatch between model and experiment, as detailed in section 7.2.

The inputs from earlier chapters are:

- the corrected F–CMOD history from chapters 4 and 5;
- the tensile strength scale f_{ctm} (or proportional-limit anchor) taken from the same structural processing pipeline as the inverse identification;
- the geometry (b, h, L, a_0) of table 6.1, giving $h_{\mathrm{sp}} = h - a_0$.

Numerical inversion is applied in turn to every beam in each campaign using the corrected F–CMOD histories and the geometry–strength metadata already established in Chapters 4–5, so that no additional empirical input enters at this stage beyond the cohesive parameters being identified.

7.2 Trilinear σ–w laws and identification

Following the formulation documented in the project equation sheet and in [43], the cohesive law is taken as trilinear in w. With tensile strength f_{ctm} and post-cracking

stresses σ_1, σ_2 at interior knots w_1, w_2, and ultimate crack width w_c, the piecewise form may be written as

$$\sigma(w) \;=\; \begin{cases} f_\mathrm{ctm} + (\sigma_1 - f_\mathrm{ctm})\,\dfrac{w}{w_1}\,, & 0 \le w < w_1\,, \\[2mm] \sigma_1 + (\sigma_2 - \sigma_1)\,\dfrac{w - w_1}{w_2 - w_1}\,, & w_1 \le w \le w_2\,, \\[2mm] \sigma_2 - \sigma_2\,\dfrac{w - w_2}{w_\mathrm{c} - w_2}\,, & w_2 < w \le w_\mathrm{c}\,, \end{cases} \tag{7.2}$$

with $\sigma(0) = f_\mathrm{ctm}$ and $\sigma(w_\mathrm{c}) = 0$. The specific fracture energy associated with eq. (7.2) is the area under $\sigma(w)$,

$$\begin{aligned} G_F \;&=\; \tfrac{1}{2}\,(f_\mathrm{ctm} - \sigma_1)\,w_1 \;+\; \tfrac{1}{2}\,(\sigma_2 + \sigma_1)\,(w_2 - w_1) \\ &\quad+\; \tfrac{1}{2}\,\sigma_2\,(w_\mathrm{c} - w_2)\,, \end{aligned} \tag{7.3}$$

in N/mm (equivalently kJ/m^2). Equation (7.3) is the quantity reported alongside the identified parameters in the detailed inverse tables exported for each specimen.

The forward map "σ–$w \to F$–CMOD" uses the hinge depth $s = h_\mathrm{sp}/2$ in the standard manner: the ligament is discretised into layers, each layer experiences opening $w_i = w(y_i)$ from eq. (7.1), and the stress $\sigma_i = \sigma(w_i)$ generates the internal moment; equilibrium gives F for a prescribed mouth opening. The inverse map fixes f_ctm from the test metadata and minimises, in the least-squares sense, the difference between model and experiment,

$$\mathrm{RMSE} \;=\; \sqrt{\left\langle\, (F_\mathrm{exp} - F_\mathrm{ana})^2 \,\right\rangle}\,, \tag{7.4}$$

over the sampled CMOD values, with respect to the free parameters (e.g. σ_1, σ_2, w_1, w_2, w_c) subject to physical ordering and positivity constraints. The coefficient of determination R^2 for the force match is reported per specimen in the result exports.

Identifiability and interpretation. A trilinear law with five free shape parameters can trade peak stress against tail slope when the experimental window is short or noisy; in practice the hinge discretisation and the span of reliable CMOD after peak load determine how stiffly each parameter is resolved. When the post-peak record is long and monotonic, the optimiser tends to separate the first softening segment (controlled largely by w_1 and σ_1) from the tail (controlled by w_2, w_c, and σ_2); when the window is truncated, G_F from eq. (7.3) can remain stable while individual knot stresses move in a correlated bundle. The thesis therefore treats the exported R^2 and RMSE together with the overlay plot as the primary quality state-

ment, and uses G_F as a robust scalar summary of the identified softening rather than over-interpreting any single interior knot in isolation.

Code checks. The analysis optionally evaluates residual-strength ratios consistent with *fib* Model Code 2010 ductility heuristics (e.g. $f_{R,1k}/f_{Lk} > 0.4$ and $f_{R,3k}/f_{R,1k} > 0.5$ for certain replacement classes) as screening metrics alongside the trilinear fit.

7.3 Application to Case Studies 1 and 2

Both case studies use the same modelling kernel; they differ in specimen matrix and the factorial design, as in chapters 5 and 6. For each beam, the pipeline exports (i) an overlay of experimental vs. analytical F–CMOD(=CMOD) curves, (ii) the identified trilinear parameters and implied G_F from eq. (7.3), and (iii) side-by-side statistics when several catalogue σ–w templates are evaluated in the same forward engine. Interpreting the matrices in figs. 7.1, 7.2, 7.7 and 7.8 therefore amounts to reading a *model competition*: each tile corresponds to a mixture/fibre combination and a model variant (true inverse vs. imported code law), with the colour scale encoding goodness of fit to the measured mouth-opening-controlled response. Where a tile is systematically cooler for library laws than for the inverse trilinear law, the limitation lies in the template shape or calibration (steel-fibre priors), not in the DIC-derived kinematics of chapter 6. The bar chart in fig. 7.3 compresses that information across the whole Case 1 cohort and makes the PP-specific penalty of generic templates visible at a glance.

7.3.1 Case Study 1: first-group screen

The Case 1 inverse matrices are shown at full width in figs. 7.1 and 7.2 so that colour legend and cell labels remain legible in print (each panel is its own float).
The High-cement companion grid (fig. 7.2) uses the same layout convention, which makes matrix-to-matrix comparison a page-turn rather than a squeezed half-width column.

Representative force–opening plots for the first group are shown in figs. 7.4 to 7.6; the full Case 1 screen comprises twenty-two beams in the same mixture matrix as Chapter 5, and the bar chart in fig. 7.3 already summarises how often the trilinear inverse outperforms imported template laws across that cohort. Qualitatively, the inverse trilinear model tracks the measured tail of the F–CMOD record once the rigid-hinge window of chapter 6 is satisfied. Specimens with kinematic nonlinearity at the tip (e.g. late local secondary cracking) are the same cases that Chapter 6

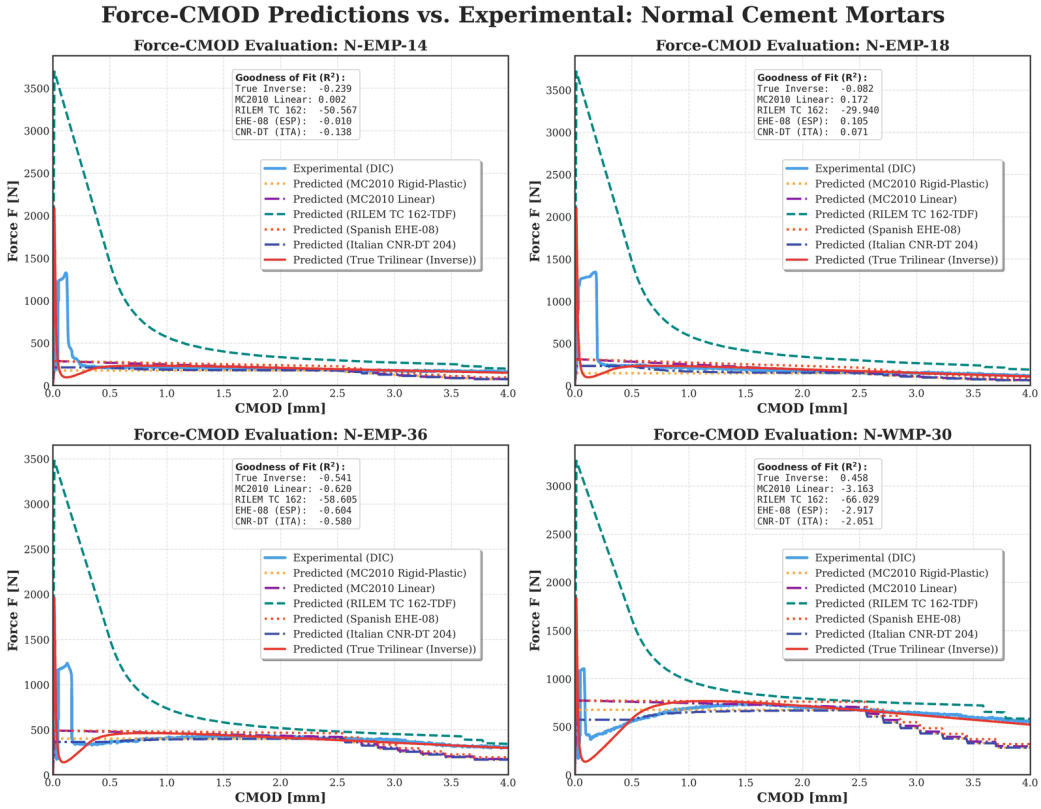


Figure 7.1: Case 1 inverse matrix overview, Normal-cement family (multi-code / presentation pipeline aggregate; colours encode model variant vs. experimental R^2).

flagged for reduced CMOD–CTOD R^2, and those are the beams where inverse results should be read cautiously irrespective of the fitted trilinear parameters: a good force match cannot compensate for a broken kinematic assumption at the notch.
The same inverse template is applied to a thin-film fibre mix (fig. 7.5) and to a short staple-fibre mix (fig. 7.6), illustrating how tail slope and peak load shift with dosage and geometry while the trilinear structure of the cohesive law remains.

7.3.2 Case Study 2: factorial (length × dose × matrix)

The second group reuses the same hinge kernel and identification objective but applies it to the factorial matrix of Chapter 5 (SMP length, dose, and matrix family). Matrix overview tiles for the SMP length split are shown in figs. 7.7 and 7.8; exemplar force curves appear in figs. 7.9 and 7.10. Interpreting Case 2 requires keeping three levers in mind simultaneously: fibre length (12 vs. 24 mm), volume dosage, and Normal vs. High cementitious matrix. The grids make it possible to see whether a given template law collapses several factorial cells to the same poor R^2 (a calibration issue) or whether only extreme dosages misfit (a shape-limitation issue). The 1.00% Normal-matrix "V-dip" visible in the structural and kinematic chapters remains a

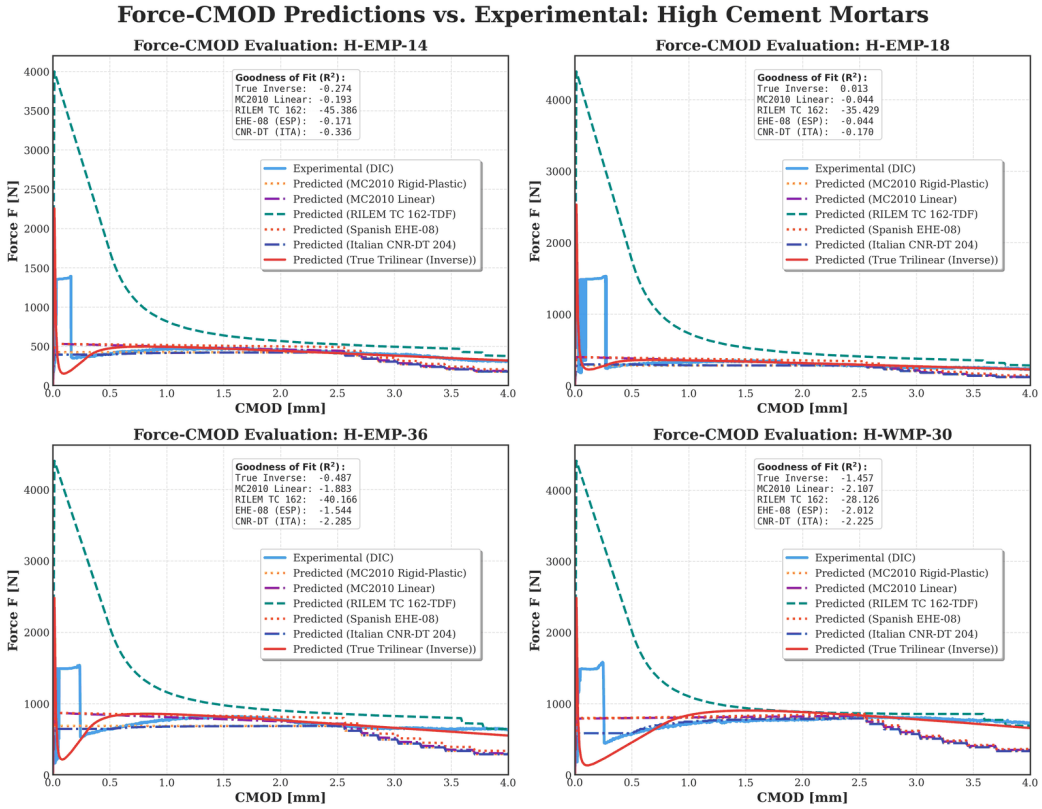


Figure 7.2: Case 1 inverse matrix overview, High-cement family.

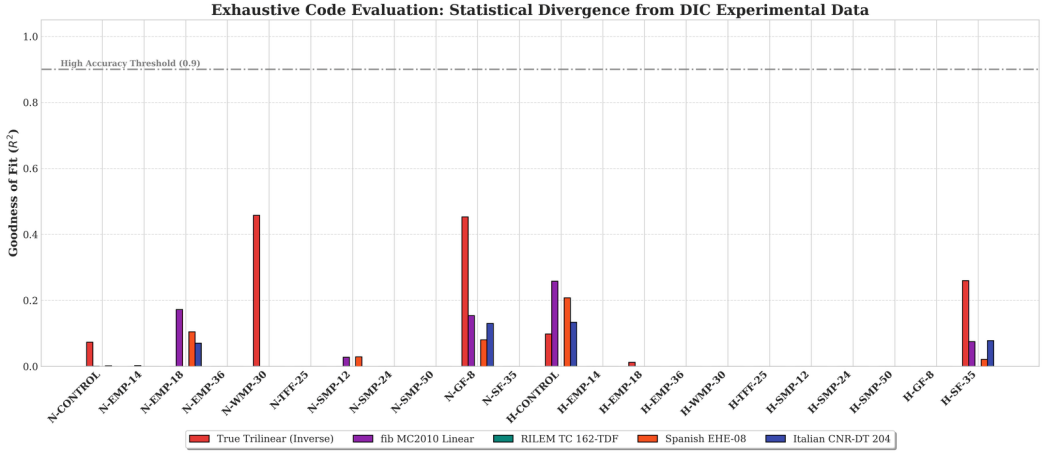


Figure 7.3: R^2 comparison across the inverse and code-template models (Case 1). The true trilinear inverse is the primary benchmark; code laws imported from MC2010-style templates typically show lower R^2 for the PP-fibre systems studied here.

mechanical feature of the F record; the inverse model captures it with a trilinear law whose softening tail differs from neighbouring doses, while the optical hinge validation in chapter 6 showed that the δ–CMOD–CTOD lines stay as affine as for the rest of the cohort—so the dip is not an artefact of mouth–tip inconsistency.

The 24 mm SMP grid (fig. 7.8) isolates the length factorial at the next notch-to-

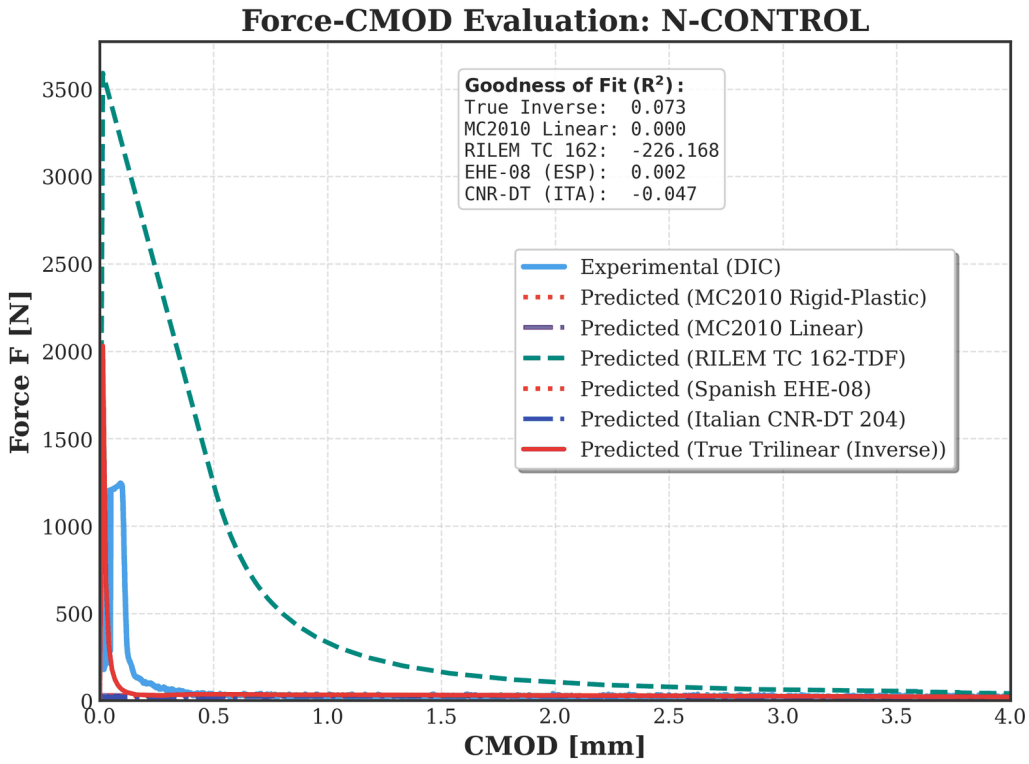


Figure 7.4: Case 1 experimental vs. model F–CMOD curve: N-CONTROL.

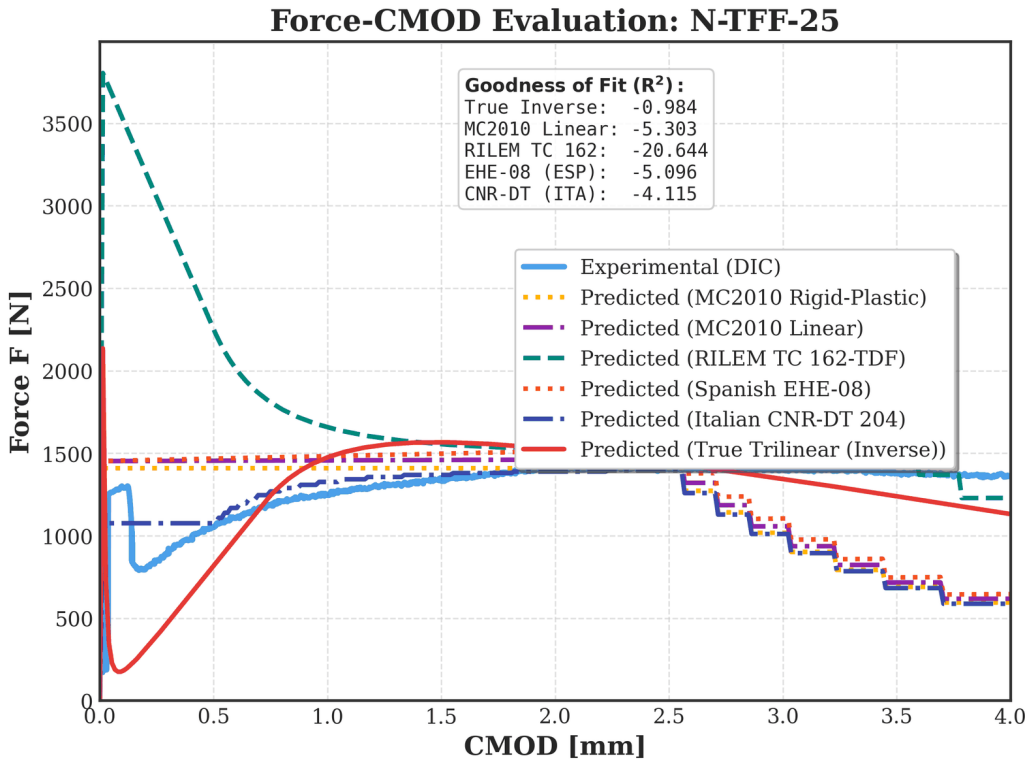


Figure 7.5: Case 1 experimental vs. model F–CMOD curve: N-TFF-25 mm.

fibre scale; comparing the two floats avoids compressing the colour bar and matrix

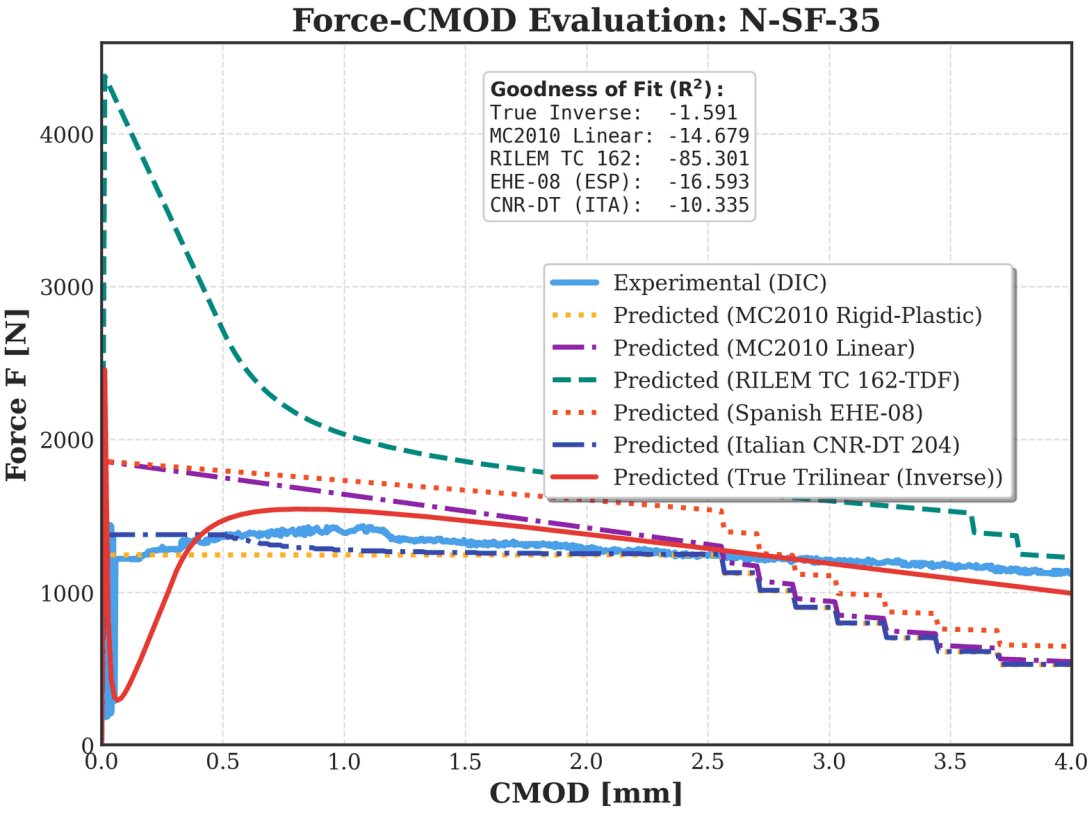


Figure 7.6: Case 1 experimental vs. model F–CMOD curve: N-SF-35 mm.

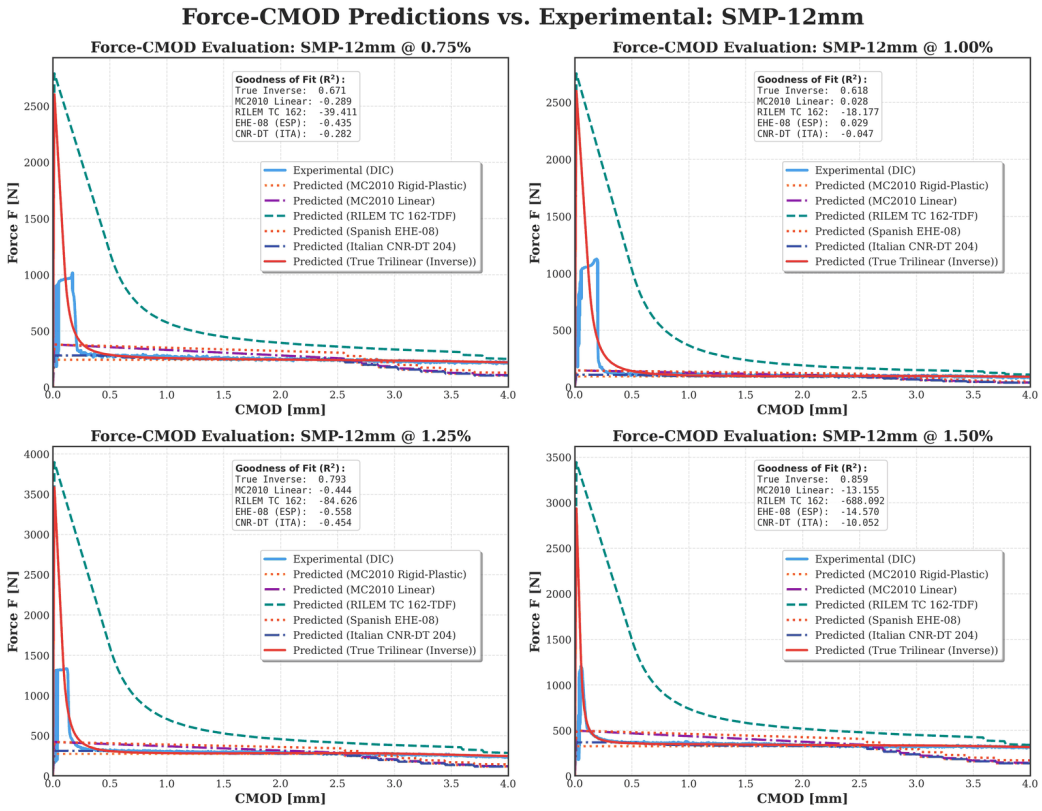


Figure 7.7: Case 2 SMP matrix grid, 12 mm fibre length.

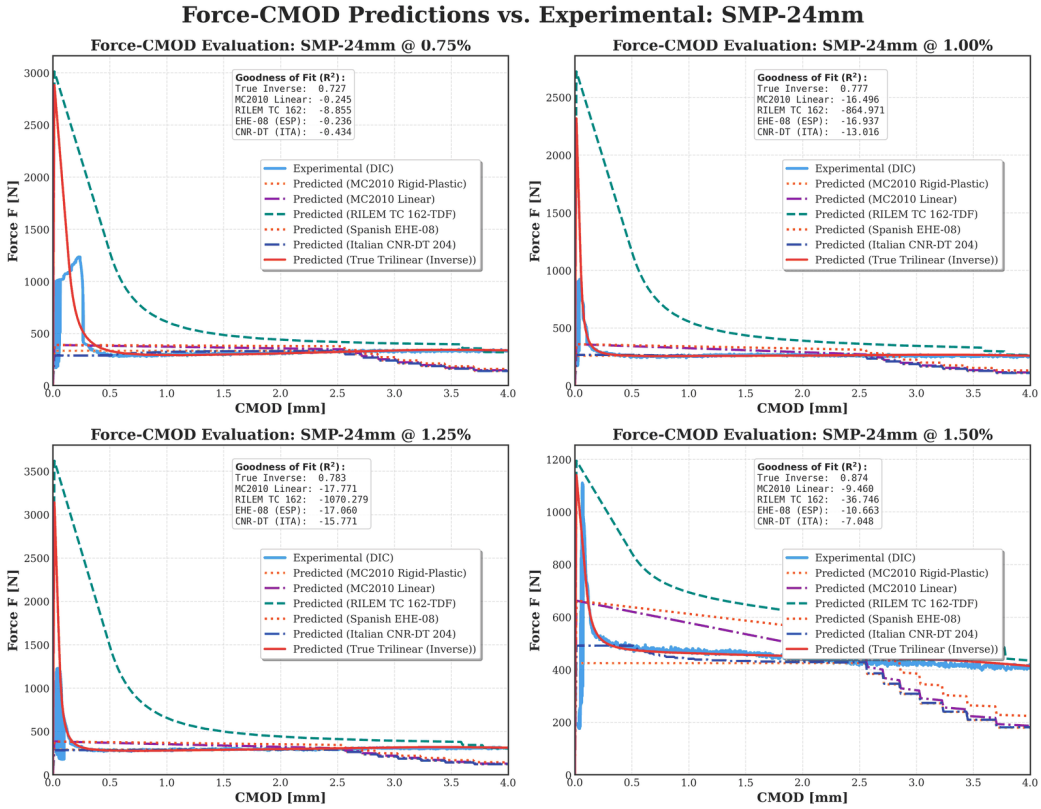


Figure 7.8: Case 2 SMP matrix grid, 24 mm fibre length.

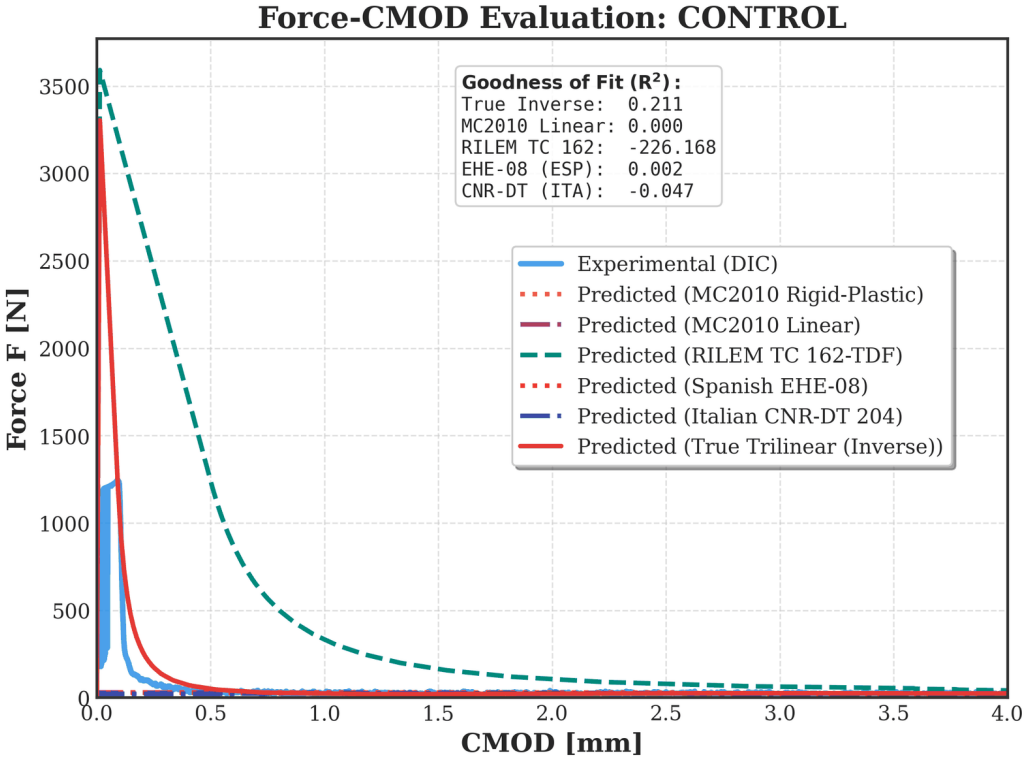


Figure 7.9: Case 2 experimental vs. model F–CMOD curve: N-CONTROL.

annotations.

A representative SMP factorial beam (fig. 7.10) then shows how the inverse fit behaves once fibre length and matrix family move away from the control mortar baseline.

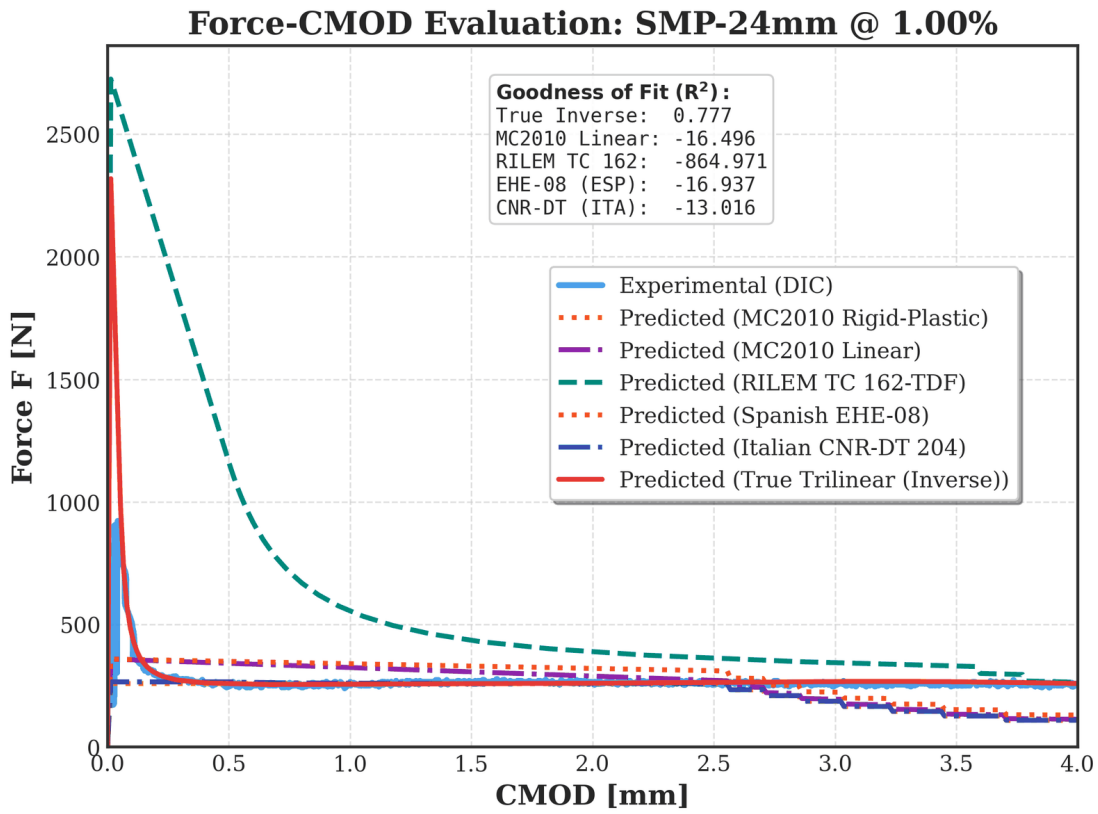


Figure 7.10: Case 2 representative SMP specimen.

Numerical procedure (both campaigns). For each beam, the identification follows the same logical sequence implied by section 7.2: read the corrected F–CMOD series and associated metadata (geometry, notch depth, tensile anchor f_{ctm}) from the structural chapters; instantiate the hinge with (b, h, L, a_0) and h_{sp}, evaluate the layered opening profile eq. (7.1), and assemble the trilinear template eq. (7.2) with decision variables $(\sigma_1, \sigma_2, w_1, w_2, w_{\mathrm{c}})$ under ordering and positivity constraints; minimise the discrete analogue of eq. (7.4) over the sampled mouth openings; then record overlays, R^2, RMSE, implied G_F from eq. (7.3), and—when the comparison block is enabled—the same forward response for each catalogue σ–w law so that matrix and bar-chart panels in section 7.3 reflect one consistent batch evaluation rather than hand-picked curves.

7.4 European standards, DIC kinematics, and the inverse pipeline

The standards route (EN 14651, [40]; *fib* Model Code 2010, [31]; RILEM recommendations) defines residual strengths and—through prescribed σ–w shapes—implied F–CMOD responses for design. The present experimental route instead measures mouth and tip opening with DIC, proves the affine links δ–CMOD and CMOD–CTOD in chapter 6, and *then* inverts a measured F–CMOD record to a material σ–w law. The two views meet at the hinge: if CMOD and code CMOD refer to the same physical mouth opening, and if eq. (7.1) holds, a trilinear law identified inversely is the model-based completion of the same data that the code templates approximate with catalogued curves for steel-fibre concrete.

Discussion-wise, the bridge figures in this section are not decorative: they answer whether a mismatch between inverse and code-template F traces could still be blamed on inconsistent kinematic definitions. Mean and enhanced δ–CMOD plots for each campaign collapse the per-beam affine regressions of Chapter 6 onto a single plane so that departures from linearity are visible as cohort-level bias rather than as isolated registration errors. The PP-vs-codes panel for the factorial group then asks a second question—whether code-recommended mouth-opening magnitudes sit on the same scale as DIC-derived CMOD once the hinge angle is propagated through the beam geometry. If templates sit on a different scale while the inverse still fits F, the inconsistency is purely in the code-side opening definition or residual-strength anchoring, and updating those anchors is a smaller intervention than changing the cohesive shape itself.

Novelty for PP-fibre mortar. Engineered code templates are still skewed to metallic fibres; the polypropylene systems tested here (especially at the high dosages and short lengths in Case 2) can therefore show a large R^2 gap between "true trilinear" inverse and imported library laws. That gap is a *code calibration* issue, not a DIC issue: the kinematic bridge plots in figs. 7.11 to 7.15 show that the δ–CMOD relation remains linear across the cohort, while the inverse chapter demonstrates that a bespoke σ–w is required to match the *force* trace.
The enhanced first-group view (fig. 7.12) keeps the same affine trend but resolves the scatter band that motivates reporting hinge R^2 before interpreting inverse parameters.

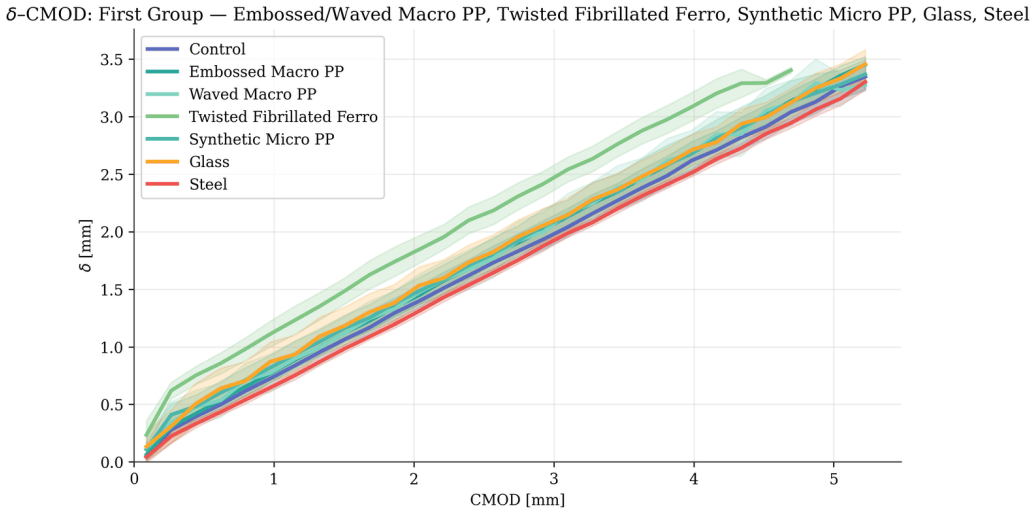


Figure 7.11: DIC-based δ vs. CMOD for the first testing group (mean trends).

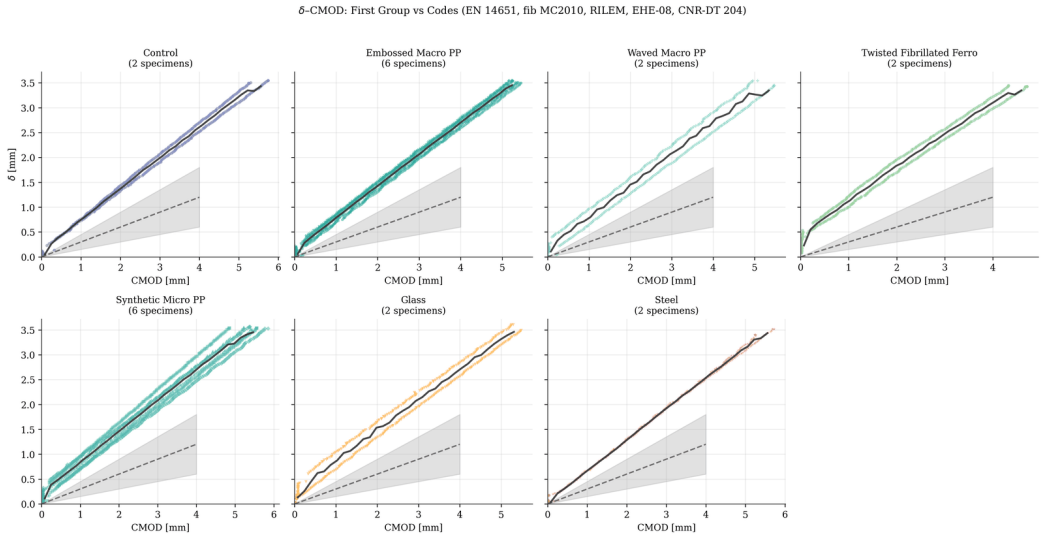


Figure 7.12: First testing group, enhanced detail of the δ–CMOD affine law used in the hinge and inverse forward model.

Turning to the factorial cohort, figs. 7.13 and 7.14 repeat the mean-then-enhanced

presentation so that short-fibre and high-dose mixtures remain readable at thesis.

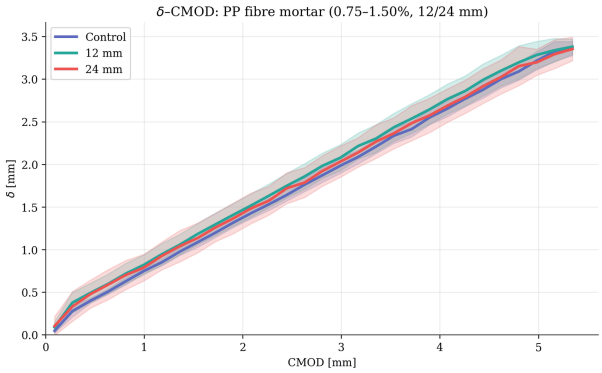


Figure 7.13: Second (factorial) testing group: mean δ–CMOD kinematic bridge.

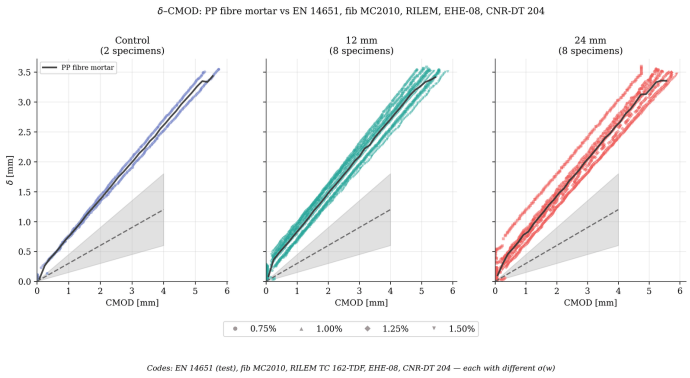


Figure 7.14: Second testing group, enhanced bridge.

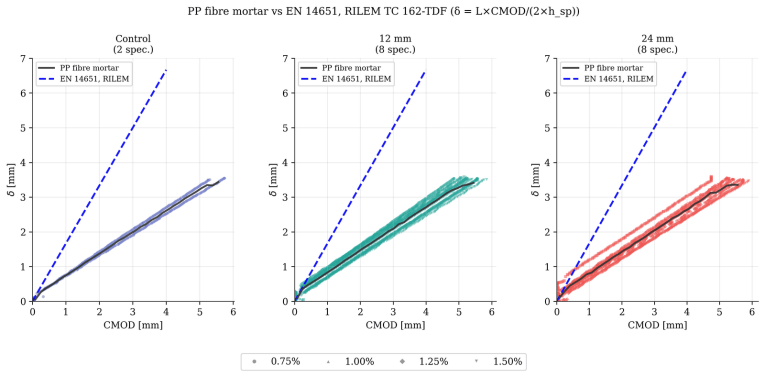


Figure 7.15: Polypropylene factorial cohort: experimental DIC scale juxtaposed with code-referenced openings on the δ–CMOD plane; the overlay makes explicit where template scales depart from the DIC-derived hinge kinematics.

Figure 7.15 closes the factorial bridge discussion by comparing code-referenced openings with the experimental δ–CMOD plane; together with figs. 7.13 and 7.14 it shows whether any F-template mismatch can still be attributed to inconsistent mouth-opening definitions.

Practical implication. Where library σ–w laws under-predict the measured F–CMOD response of PP mortars, design may adopt the inversely identified trilinear law (with appropriate quantile/partial factors if embedded in a code format), while

DIC quality metrics (hinge R^2 from chapter 6) act as pre-checks before trusting an inverse result.

Closing synthesis. Read as a whole, this chapter separates three questions that are often conflated in fibre-beam reporting: whether the *kinematics* of mouth and tip opening are internally consistent (answered in chapter 6 and revisited with the bridge plots of section 7.4); whether a *simple cohesive law* can reproduce the measured F–CMOD when those kinematics are fed into a hinge forward model (answered affirmatively for most specimens in section 7.3, with explicit caveats where hinge quality degrades); and whether *off-the-shelf code templates* carry the same descriptive power for PP mortar as they do for steel-fibre calibration sets (answered cautiously: the matrices show systematic R^2 deficits for libraries, even when kinematics are sound). Keeping those three answers distinct is what allows the thesis to recommend inverse-based material parameters for downstream applications without pretending that a single standards annex can yet encode every PP mixture in this study.

7.5 Summary of findings

The following points condense the evidence and the interpretation developed above; they are intentionally parallel to the chapter structure (model, identification, campaigns, code bridge).

1. **Forward model.** A trilinear cohesive law embedded in the Olesen/Ostergaard hinge framework [29] is sufficient, in the least-squares sense of eq. (7.4), to reproduce the corrected experimental F–CMOD records for the majority of beams in both campaigns, with specimen-level R^2 and RMSE reported next to each fitted curve in the usual inverse-analysis output layout.

2. **Material closure.** The identified parameters imply a specific fracture energy eq. (7.3) for each fit, giving a direct bridge from beam-scale testing to a softening curve usable in sectional or continuum models; where the post-peak window is short, G_F should be read as the more stable scalar than any single interior knot stress.

3. **Templates vs. inverse.** When the same forward engine evaluates catalogue σ–w laws alongside the inversely fitted trilinear curve, the PP-rich matrices in Case 1 and Case 2 typically assign higher explanatory power to the bespoke inverse law than to imported steel-fibre templates, even though the DIC hinge kinematics remain affine—so the remaining discrepancy is attributed to template shape and calibration, not to optical displacement definitions.

4. **Reproducibility.** The case-study figures are representative selections from the same batch evaluation described in section 7.3.2; any mixture that appears anomalous in the matrices should be revisited by re-inspecting that beam's overlay and hinge diagnostics from Chapter 6, not by adjusting the narrative in isolation.

The identified trilinear laws and the code-bridge analysis close the constitutive identification pipeline of the thesis.

Chapter 8

Conclusions and Future Scope

8.1 Conclusions

This final chapter summarises the thesis in a form aligned with the research objectives stated in section 1.3, then states scientific and technical results, their relevance for industry and standards, and their connection to social benefit and the Sustainable Development Goals (SDGs). Limitations and future recommendations are set out in section 8.2 and section 8.3.

8.1.1 Conclusions in relation to the research objectives

The specific objectives in section 1.3 are recalled here in the same order, each followed by a concise statement of what was achieved. Cross-references point to the chapters where the work is fully developed.

1. DIC-based displacement correction for standardised flexure (chapter 4). A displacement-derived compliance (ODDC) workflow was developed and applied so that true specimen deformation is separated from machine and load-path compliance and rigid-body effects. The method was documented across the beam, cube, and bar configurations in scope, producing corrected force–displacement and optical records suitable for the analyses that follow. The outcome is a traceable metrological base for all subsequent structural and fracture statements.

2. Structural performance, statistics, and unified eco-efficiency with cradle-to-gate LCA (chapter 5). Non-parametric and statistical tools were used together with LCA and combined ranking or eco-efficiency constructs so that

post-cracking performance and global warming potential (GWP) are not optimised in isolation. The work covers the case-study matrix defined in the methodology, including the two principal notched-beam screens and the complementary cube and bar programmes where applicable.

3. Fracture kinematics, including CMOD/CTOD, and validation of the rigid-hinge model (chapter 6). DIC-based crack mouth and crack tip opening variables were mapped and checked against the rigid-hinge idealisation. The results support the use of optically defined openings as consistent substitutes for contact gauges for the purposes of the standards-linked indices used in the thesis, with explicit treatment of outlier and exceptional cases.

4. Trilinear σ–w laws by inverse analysis and link to European code conventions (chapter 7). Hinge-based inverse analysis identified trilinear cohesive laws for the notched-beam data; the European-code and DIC–code discussion is integrated with that identification.

Together, these items close the loop proposed in the introduction: corrected structural input, performance and environment, validated kinematics, identified softening, and a path to three-dimensional validation.

8.1.2 Scientific and technical application

The thesis contributes a reproducible measurement chain for notched-beam characterisation of fibre-reinforced mortars, in which optical metrology and compliance correction precede constitutive interpretation. The main technical points are:

- **Quasi-brittle testing.** For small, stiff FRM beams, raw actuator stroke is not a reliable stand-in for specimen opening; the work shows how to obtain records that are suitable for fracture energy and residual-strength use without conflating machine and specimen.

- **Integrated mechanics and LCA.** Coupling post-cracking metrics with cradle-to-gate GWP and sensitivity analysis (cement as the dominant lever in the documented tornados) makes trade-offs between ductility, strength, and environmental cost explicit.

- **Kinematics and constitutive modelling.** Rigid-hinge checking and DIC-based openings provide a gate for inverse identification; trilinear σ–w laws and

comparison with code-side templates are reported in a common forward framework, clarifying when mismatch is likely due to *template calibration* rather than to inconsistent kinematics.

- **Forward modelling.** The identified trilinear laws are positioned for re-use in cohesive interface or detailed FE models in future work.

8.1.3 Industry, practice, and standards

For laboratories, designers, and code users, the following practical implications are emphasised:

- Test reports for softening-governed notched beams should **separate** machine compliance from specimen deformation where the former is not negligible, and should document the displacement basis (optical, clip gauge, or corrected stroke) for G_F and residual indices.
- **Material selection** for PP-rich systems should not rely on peak load alone: post-peak energy and eco-efficiency indices that include ductility are needed when comparing fibres and matrix families.
- **Test geometry** (e.g. bond-sensitive configurations) can change the ranking of the same material set; the thesis highlights CST/BBS-style contrasts for bond-critical applications.
- **Design codes.** DIC-consistent CMOD/CTOD and hinge kinematics are compatible with EN-family fracture indices when virtual gauges match standard definitions. Where template σ–w laws under-predict F–opening response for polypropylene systems, the evidence points toward mixture- and class-specific calibration or performance-based identification, not only to a failure of the optical measurement.
- **Finite-element use.** Parameters identified on a 1D hinge model should be checked at the mesh and geometry of interest before direct transfer to design calculations.

8.1.4 Social impact and sustainable development goals

The work is connected to public-interest outcomes as follows. Reducing the carbon and resource intensity of construction materials, while keeping structural safety, supports climate and urban resilience agendas. The thesis makes environmental

trade-offs of fibre-reinforced mortar visible alongside mechanical performance, which supports transparent choices for infrastructure and precast products.

A concise mapping to the United Nations SDGs (non-exhaustive) is:

- **SDG 9** (Industry, innovation and infrastructure): improved characterisation and traceability of fibre-reinforced cementitious systems for use in technical specifications and design workflows.
- **SDG 11** (Sustainable cities and communities): materials that can contribute to durable, repairable built assets when tested and modelled with clear fracture and environmental metrics.
- **SDG 12** (Responsible consumption and production): cradle-to-gate LCA and combined eco-efficiency thinking encourage proportionate use of binders and fibres and discourage "greenwashing" by peak-load-only claims.
- **SDG 13** (Climate action): explicit GWP accounting and sensitivity to clinker content align material selection with decarbonisation pathways, within the LCA delimitations of the thesis.

The social benefit is indirect but concrete: more reliable public and commercial decisions on FRM when mechanical and environmental evidence are reported together, and when testing practice removes systematic compliance bias.

8.2 Limitations

The following limitations bound the generality of the results:

1. **Scope of materials and tests.** Conclusions are drawn from the delimited mix designs, fibre families, and test geometries in chapters 3 and 5. They are not automatically transferable to all PP systems, all notch geometries, or other binders without further testing.
2. **LCA boundary.** Assessments are cradle-to-gate with GWP as the primary impact category. Other impact categories, end-of-life, regional electricity mixes, and full building-scale system studies are outside the present scope.
3. **Inverse identifiability.** Trilinear parameters can be correlated when the post-peak record is short or noisy; G_F and overall fit quality should be read together with the overlay plots, as discussed in chapter 7.

4. **Kinematic exceptions.** Specimens with late secondary cracking, poor optical tracking, or low hinge R^2 require interpretive caution for both standards indices and inverse analysis.

5. **FE scope.** The finite-element verification path remains as a direction for future work; publication-grade meshes, boundary conditions, and figure sets are to be completed as the project pipeline requires.

6. **Service conditions.** The thesis is based on laboratory quasi-static tests. Fatigue, creep, fire, and field durability are not claimed here.

8.3 Future recommendations

The following steps would extend the work without changing its internal logic:

1. **Experimental extension:** broaden aspect ratios, notch depths, and additional fibre types (including long macros and selected hybrids) to test size effect and transferability of identified σ–w parameters; systematic study of machine setup (span, supports) on compliance sensitivity.

2. **Optics:** stereoscopic DIC, higher frame rates, or volume correlation for selected specimens to challenge the rigid-hinge and out-of-plane assumptions in thicker or more three-dimensional cases.

3. **Inverse and uncertainty:** Bayesian or interval methods for cohesive parameters; multi-objective fits if independent opening data are available along the process zone.

4. **LCA:** additional impact categories, scenarios for low-clinker binders, and building-scale case studies to connect beam-level rankings to project-level decisions.

5. **FE:** complete COMSOL (or equivalent) mesh convergence, interface regularisation study, and documented comparison to beam-scale identification.

6. **Standards and data:** promote reporting templates that require displacement provenance; consider open reference datasets with metadata for replication.

These recommendations close the thesis while leaving a clear path for the research group and for others working on non-metallic fibre–reinforced cementitious systems.

References

[1] Robbie M. Andrew. "Global CO_2 emissions from cement production, 1928–2018". In: *Earth System Science Data* 11.4 (2019), pp. 1675–1710. DOI: 10.5194/essd-11-1675-2019.

[2] Yunpeng Zhao and Sheng Qiang. "Exploring the potential of bayer red mud: Toward self-sensing cement-based composite materials". In: *Construction and Building Materials* 460 (2025), p. 139850. DOI: 10.1016/j.conbuildmat.2024.139850.

[3] Isabela de Paula Salgado et al. "Integrating life cycle assessment (LCA) and machine learning for sustainable designs: a case study on protective layers made of mineral-bonded fiber-reinforced composites". In: *The International Journal of Life Cycle Assessment* (2025), pp. 1–20. DOI: 10.1007/s11367-025-02454-7.

[4] Jofre Mañosa et al. "Research evolution of limestone calcined clay cement (LC3), a promising low-carbon binder–A comprehensive overview". In: *Heliyon* 10.3 (2024). DOI: 10.1016/j.heliyon.2024.e25117.

[5] Ban Ahmed Khaleel and Eethar Thanoon Dawood. "Evolution of ultra high performance concrete using hybrid fibers: a review". In: *Discover Concrete and Cement* 1.1 (2025), p. 2. DOI: 10.1007/s44416-025-00001-z.

[6] Emran Alotaibi, Nadia Nassif, and Mohamad Alhalabi. "Ultra-high-performance concrete for Sustainable Development Goals: A systematic review and conceptual framework (With Full Text)". In: (2026). DOI: 10.1016/j.jobe.2026.115802.

[7] Ronald F Zollo. "Fiber-reinforced concrete: an overview after 30 years of development". In: *Cement and Concrete Composites* 19.2 (1997), pp. 107–122. DOI: 10.1016/S0958-9465(96)00046-7.

[8] Fuat Köksal et al. "Fracture energy-based optimisation of steel fibre reinforced concretes". In: *Engineering Fracture Mechanics* 107 (2013), pp. 29–37. DOI: 10.1016/j.engfracmech.2013.04.018.

[9] Vahid Afroughsabet and Togay Ozbakkaloglu. "Mechanical and durability properties of high-strength concrete containing steel and polypropylene fibers". In: *Construction and Building Materials* 94 (2015), pp. 73–82. DOI: 10.1016/j.conbuildmat.2015.06.051.

[10] Longbang Qing et al. "Study on the different roles of steel and PVA fibers in the fracture evolution of hybrid fiber reinforced cementitious composites". In: *Materials and Structures* 59.2 (2026), p. 109. DOI: 10.1617/s11527-026-03000-z.

[11] Hisham Hafez et al. "Decarbonization potential of steel fibre-reinforced limestone calcined clay cement concrete one-way slabs". In: *Construction and Building Materials* 435 (2024), p. 136847. DOI: 10.1016/j.conbuildmat.2024.136847.

[12] Pietro Di Maida et al. "Pullout behavior of polypropylene macro-synthetic fibers treated with nano-silica". In: *Construction and Building Materials* 82 (2015), pp. 39–44. DOI: 10.1016/j.conbuildmat.2015.02.047.

[13] Nemkumar Banthia and Rishi Gupta. "Influence of polypropylene fiber geometry on plastic shrinkage cracking in concrete". In: *Cement and Concrete Research* 36.7 (2006), pp. 1263–1267. DOI: 10.1016/j.cemconres.2006.01.010.

[14] Aitor Llano-Torre and P Serna. "Recommendation of RILEM TC 261-CCF: test method to determine the flexural creep of fibre reinforced concrete in the cracked state". In: *Materials and Structures* 54.3 (2021), p. 124. DOI: 10.1617/s11527-021-01675-0.

[15] C Del Prete et al. "A lattice discrete particle model to simulate the viscoelastic behaviour of macro-synthetic fibre reinforced concrete". In: *Construction and Building Materials* 295 (2021), p. 123630. DOI: 10.1016/j.conbuildmat.2021.123630.

[16] WSA Nana et al. "Behaviour of macro-synthetic fibers reinforced concrete: Experimental, numerical and design code investigations". In: *Structures* 32 (2021), pp. 1271–1286. DOI: 10.1016/j.istruc.2021.03.080.

[17] Zdenék P Bažant and Sze-Dai Pang. "Revision of Reliability Concepts for Quasibrittle Structures and Size Effect on Probability Distribution of Structural Strength". In: (2005), pp. 377–386.

[18] Raffaele Cucuzza et al. "Size-scale effects and modelling issues of fibre-reinforced concrete beams". In: *Construction and Building Materials* 392 (2023), p. 131727. DOI: 10.1016/j.conbuildmat.2023.131727.

[19] Ieva Paegle, Fausto Minelli, and Gregor Fischer. "Cracking and load-deformation behavior of fiber reinforced concrete: Influence of testing method". In: *Cement and Concrete Composites* 73 (2016), pp. 147–163. DOI: 10.1016/j.cemconcomp.2016.06.012.

[20] E Michelini, D Ferretti, and M Pizzati. "The influence of density on the fracture energy of AAC: From experimental investigation to the calibration of a cohesive law". In: *Construction and Building Materials* 400 (2023), p. 132547. DOI: 10.1016/j.conbuildmat.2023.132547.

[21] Pedro Paulo Martins de Carvalho and Rodrigo de Melo Lameiras. "Test Methods for the Post-cracking Behavior of the FRC in Tension: A State of the Art Review". In: *International Journal of Concrete Structures and Materials* 19.1 (2025), p. 67. DOI: 10.1186/s40069-024-00702-y.

[22] Shilang Xu et al. "The theoretical basis of testing methods to determine double-K criterion for crack propagation in concrete: technical report of the RILEM TC265-TDK". In: *Materials and Structures* 54.6 (2021), p. 229. DOI: 10.1617/s11527-021-01787-7.

[23] Peter Mark et al. *Neuauflage DAfStb-Richtlinie Stahlfaserbeton*. Tech. rep. 1. 2021, pp. 19–25. DOI: 10.1002/best.202000065.

[24] Sudhanshu S Pathak et al. "Fracture assessment of concrete through digital image correlation (DIC) technique: Correction for seating errors and machine compliance". In: *Journal of Building Pathology and Rehabilitation* 9.2 (2024), p. 132. DOI: 10.1007/s41024-024-00487-3.

[25] Roberto Felicetti and Ramin Yarmohammadian. "3action: Advances in direct-tensile testing of concrete". In: *Materials and Structures* 59.2 (2026), p. 69. DOI: 10.1617/s11527-026-02976-y.

[26] Nicola Gehri, Jaime Mata-Falcón, and Walter Kaufmann. "Refined extraction of crack characteristics in large-scale concrete experiments based on digital image correlation". In: *Engineering Structures* 251 (2022), p. 113486. DOI: 10.1016/j.engstruct.2021.113486.

[27] Hannes Panwitt, Paul Köster, and Manuela Sander. "Fatigue crack growth determination under in-phase and out-of-phase mixed-mode loading conditions using an automated DIC evaluation tool". In: *International Journal of Fatigue* 164 (2022), p. 107122. DOI: 10.1016/j.ijfatigue.2022.107122.

[28] David Melching et al. “An iterative crack tip correction algorithm discovered by physical deep symbolic regression”. In: *International Journal of Fatigue* 187 (2024), p. 108432. DOI: 10.1016/j.ijfatigue.2024.108432.

[29] Barzin Mobasher et al. “Characterization of toughening mechanisms in UHPC through image correlation and inverse analysis of flexural results”. In: *Cement and Concrete Composites* 122 (2021), p. 104157. DOI: 10.1016/j.cemconcomp.2021.104157.

[30] Stefie J Stephen et al. “Determination of the tensile constitutive relations of fiber reinforced concrete using inverse analysis”. In: *Construction and Building Materials* 195 (2019), pp. 405–414. DOI: 10.1016/j.conbuildmat.2018.11.014.

[31] fib-federation internationale du beton. *fib Model Code for Concrete Structures 2010*. John Wiley & Sons, 2013.

[32] Lucie Vandewalle et al. “Test and design methods for steel fibre reinforced concrete: sigma-epsilon design method”. In: *Materials and Structures* 36.262 (2003), pp. 560–567. DOI: 10.1617/14007.

[33] Arne Hillerborg, Mats Modéer, and P-E Petersson. “Analysis of crack formation and crack growth in concrete by means of fracture mechanics and finite elements”. In: *Cement and Concrete Research* 6.6 (1976), pp. 773–781. DOI: 10.1016/0008-8846(76)90007-7.

[34] A Meda, GA Plizzari, and V Slowik. “Fracture of fiber reinforced concrete slabs on grade”. In: 2001, pp. 1013–1020.

[35] Sahith Gali and Kolluru VL Subramaniam. “Multi-linear stress-crack separation relationship for steel fiber reinforced concrete: Analytical framework and experimental evaluation”. In: *Theoretical and Applied Fracture Mechanics* 93 (2018), pp. 33–43. DOI: 10.1016/j.tafmec.2017.06.018.

[36] Alejandro Nogales and Albert de la Fuente. “Numerical-aided flexural-based design of fibre reinforced concrete column-supported flat slabs”. In: *Engineering Structures* 232 (2021), p. 111745. DOI: 10.1016/j.engstruct.2020.111745.

[37] Luis MP Matos et al. “A new inverse analysis approach for predicting the fracture mode I parameters of fibre reinforced concrete”. In: *Engineering Fracture Mechanics* 246 (2021), p. 107613. DOI: 10.1016/j.engfracmech.2021.107613.

[38] Alejandro Enfedaque et al. “Numerical simulation of the fracture behavior of high-performance fiber-reinforced concrete by using a cohesive crack-based inverse analysis”. In: *Materials* 15.71 (2021). DOI: 10.3390/ma15010071.

[39] MG Alberti et al. “Numerical modelling of the fracture of polyolefin fibre reinforced concrete by using a cohesive fracture approach”. In: *Composites Part B: Engineering* 111 (2017), pp. 200–210. DOI: 10.1016/j.compositesb.2016.11.052.

[40] European Committee for Standardization (CEN). *EN 14651+A1:2007. Test method for metallic fiber concrete – Measuring the flexural tensile strength.* Standard. 2007.

[41] Matteo Colombo et al. “The basis for ductility evaluation in SFRC structures in MC2020: An investigation on slabs and shallow beams”. In: *Structural Concrete* 24.4 (2023), pp. 4406–4423. DOI: 10.1002/suco.202300114.

[42] Enzo Martinelli et al. “A cracked-hinge approach to modelling high performance fiber-reinforced concrete”. In: *Composite Structures* 273 (2021), p. 114277. DOI: 10.1016/j.compstruct.2021.114277.

[43] Haruna Ibrahim et al. “Assessment of the fracture properties of mortars reinforced with synthetic fibers”. In: *Structural Concrete* 26.3 (2025), pp. 2487–2504. DOI: 10.1002/suco.202400005.

[44] Julia Blazy, Krzysztof Grzyb, and Łukasz Drobiec. “Comparative study of results from clip gauges, LVDTs, and DIC technique for PFRC three-point bending and splitting tensile test”. In: (2022), pp. 477–490.

[45] Sonali Bhowmik and Sonalisa Ray. “An experimental approach for characterization of fracture process zone in concrete”. In: *Engineering Fracture Mechanics* 211 (2019), pp. 401–419. DOI: 10.1016/j.engfracmech.2019.02.026.

[46] Michael A Sutton, Jean Jose Orteu, and Hubert Schreier. *Image correlation for shape, motion and deformation measurements: basic concepts, theory and applications.* Boston, MA: Springer US, 2009.

[47] Qifang Liu et al. “Framework to optimise two-dimensional DIC measurements at different orders of accuracy for concrete structures”. In: *Structures* 28 (2020), pp. 93–105. DOI: 10.1016/j.istruc.2020.08.056.

[48] Hubert W Schreier, Joachim R Braasch, and Michael A Sutton. “Systematic errors in digital image correlation caused by intensity interpolation”. In: *Optical engineering* 39.11 (2000), pp. 2915–2921. DOI: 10.1117/1.1314593.

[49] David Lecompte et al. “Quality assessment of speckle patterns for digital image correlation”. In: *Optics and Lasers in Engineering* 44.11 (2006), pp. 1132–1145. DOI: 10.1016/j.optlaseng.2005.10.004.

[50] John E Hockett and Peter P Gillis. “Mechanical testing machine stiffness: Part I—Theory and calculations”. In: *International Journal of Mechanical Sciences* 13.3 (1971), pp. 251–264. DOI: 10.1016/0020-7403(71)90007-5.

[51] Kuk-Cheol Kim et al. “Improved method for fracture toughness testing using compliance method for high toughness materials”. In: *Journal of Pressure Vessel Technology* 145.1 (2023), p. 011501. DOI: 10.1115/1.4055089.

[52] DINEN ISO et al. *Metallic materials — Calibration and verification of static uniaxial testing machines — Part 1: Tension/compression testing machines — Calibration and verification of the force-measuring system*. Tech. rep. 2018, pp. 1–22.

[53] Liberato Ferrara and Marco di Prisco. “Mode I fracture behavior in concrete: Nonlocal damage modeling”. In: *Journal of Engineering Mechanics* 127.7 (2001), pp. 678–692. DOI: 10.1061/(ASCE)0733-9399(2001)127:7(678).

[54] Tarek Almusallam et al. “Analytical and experimental investigations on the fracture behavior of hybrid fiber reinforced concrete”. In: *Cement and Concrete Composites* 74 (2016), pp. 201–217. DOI: 10.1016/j.cemconcomp.2016.10.002.

[55] Samuel Sanford Shapiro and Martin B Wilk. “An analysis of variance test for normality (complete samples)”. In: *Biometrika* 52.3–4 (1965), pp. 591–611. DOI: 10.1093/biomet/52.3-4.591.

[56] Douglas C Montgomery. *Design and Analysis of Experiments*. Hoboken, NJ: John Wiley & Sons, 2017.

[57] Muhammad Aftab Khan and Qaiser uz Zaman Khan. “Optimization of mechanical properties of steel and polypropylene fiber composites using RSM”. In: *Materials Research Express* 12.2 (2025), p. 025304. DOI: 10.1088/2053-1591/adb37d.

[58] International Organization for Standardization. *Environmental management — Life cycle assessment — Principles and framework*. Tech. rep. ISO 14040:2006. ISO, 2006.

[59] International Organization for Standardization. *Environmental management — Life cycle assessment — Requirements and guidelines*. Tech. rep. ISO 14044:2006. ISO, 2006.

[60] Bruno L Damineli et al. “Measuring the eco-efficiency of cement use”. In: *Cement and Concrete Composites* 32.8 (2010), pp. 555–562. DOI: 10.1016/j.cemconcomp.2010.07.009.

[61] Taehyoung Kim, Sungho Tae, and Seungjun Roh. “Assessment of the CO2 emission and cost reduction performance of a low-carbon-emission concrete mix design using an optimal mix design system”. In: *Renewable and Sustainable Energy Reviews* 25 (2013), pp. 729–741. DOI: 10.1016/j.rser.2013.05.013.

[62] Luis MP Matos et al. “A new inverse analysis approach for predicting the fracture mode I parameters of fibre reinforced concrete”. In: *Engineering Fracture Mechanics* 246 (2021), p. 107613. DOI: 10.1016/j.engfracmech.2021.107613.

[63] Ana Blanco et al. “Application of constitutive models in European codes to RC–FRC”. In: *Construction and Building Materials* 40 (2013), pp. 246–259. DOI: 10.1016/j.conbuildmat.2012.09.096.

[64] Zdeněk P Bažant and Byung H Oh. “Crack band theory for fracture of concrete”. In: *Matériaux et construction* 16.3 (1983), pp. 155–177. DOI: 10.1007/BF02486267.

[65] Rafic G El-Helou et al. “Triaxial constitutive law for ultra-high-performance concrete and other fiber-reinforced cementitious materials”. In: *Journal of Engineering Mechanics* 146.7 (2020), p. 04020062. DOI: 10.1061/(ASCE)EM.1943-7889.000177.

[66] Matteo Colombo et al. “The basis for ductility evaluation in SFRC structures in MC2020: An investigation on slabs and shallow beams”. In: *Structural Concrete* 24.4 (2023), pp. 4406–4423. DOI: 10.1002/suco.202300114.

[67] Alejandro Enfedaque et al. “Suitability of Constitutive Models of the Structural Concrete Codes When Applied to Polyolefin Fibre Reinforced Concrete”. In: *Materials* 15.6 (2022), p. 2323. DOI: 10.3390/ma15062323.

[68] SM Iqbal S Zainal, Farzad Hejazi, and Ala MA Mafaileh. “Strengthening of Reinforced Concrete slabs using macro and micro synthetic fibers”. In: *Structures* 51 (2023), pp. 1579–1590. DOI: 10.1016/j.istruc.2023.03.120.

[69] ASTM International. *ASTM C496/C496M-17. Standard Test Method for Splitting Tensile Strength of Cylindrical Concrete Specimens*. Tech. rep. West Conshohocken, PA: ASTM International, 2017. DOI: 10.1520/C0496_C0496M-17.

[70] Norman R. Draper George E. P. Box. *Response surfaces, mixtures, and ridge analyses*. 2006. DOI: 10.1002/0470072768.

[71] Raymond H Myers, Douglas C Montgomery, and Christine M Anderson-Cook. *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*. 4th. Hoboken, NJ: John Wiley & Sons, 2016.

[72] H. B. Mann and D. R. Whitney. "On a test of whether one of two random variables is stochastically larger than the other". In: *Annals of Mathematical Statistics* 18.1 (1947), pp. 50–60. DOI: 10.1214/aoms/1177730491.

[73] William H Kruskal and W Allen Wallis. "Use of ranks in one-criterion variance analysis". In: *Journal of the American Statistical Association* 47.260 (1952), pp. 583–621. DOI: 10.1080/01621459.1952.10483441.

[74] Charles Spearman. "The proof and measurement of association between two things". In: *The American Journal of Psychology* (1961), pp. 45–58. DOI: 10.1037/11491-005.

[75] European Committee for Standardization (CEN). *Sustainability of construction works—Environmental product declarations—Core rules for the product category of construction products*. Tech. rep. EN 15804:2012+A2:2019/AC:2021. CEN, 2021.

[76] Valérie Masson-Delmotte et al. *Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Tech. rep. 1. 2021, p. 2391.

[77] ecoinvent Association. *ecoinvent database v3.9 – cut-off, allocation at the point of substitution*. Zurich, Switzerland. Background life-cycle inventory database used for cradle-to-gate emission factors. 2022.

[78] Ente Italiano di Normazione (UNI). *UNI 11039-2:2003. Steel fibre reinforced concrete – Part 2: Test method for first cracking strength and ductility indexes*. Standard. 2003.

[79] Genjin Liu and Dongsheng Li. "Analysis of synergistic effect of bending energy absorption capacity of fiber reinforced concrete based on fiber distribution characteristic". In: *Construction and Building Materials* 447 (2024), p. 138098. DOI: 10.1016/j.conbuildmat.2024.138098.

[80] Julia Blazy and Rafał Blazy. "Polypropylene fiber reinforced concrete and its application in creating architectural forms of public spaces". In: *Case Studies in Construction Materials* 14 (2021), e00549. DOI: 10.1016/j.cscm.2021.e00549.