




# Human agents, generative AI, and innovation: A formal model of hybrid creative process

Mattia Pedota<sup>a,\*</sup> , Francesco Cicala<sup>b</sup> , Alessio Basti<sup>c</sup> 

<sup>a</sup> Department of Management, Economics, and Industrial Engineering, Politecnico di Milano, Italy

<sup>b</sup> Google Zürich, Switzerland

<sup>c</sup> Department of Engineering and Geology, University of Chieti-Pescara, Italy

## ARTICLE INFO

### Keywords:

Artificial intelligence  
Augmentation  
Creativity  
Innovation  
Knowledge  
Search

## ABSTRACT

Generative AI (GenAI) has rapidly emerged as a revolutionary technology that enables new ways to generate and recombine knowledge. Despite its significant potential, research on GenAI's role in enhancing creativity and innovation is still in its early stages. The present work advances this emerging field by focusing on the human-GenAI dyad. Specifically, we propose a formal model of hybrid creative process designed to maximize the synergistic potential of the human-GenAI interaction. Drawing on machine learning literature, we conceptualize GenAI as a superposition of latent entities. Through formal argumentation, we demonstrate that optimal creative outcomes arise when human agents actively select the most appropriate entity from the complete spectrum of potential alternatives for the problem at hand. Finally, we outline the ideal iterative process required to asymptotically converge toward these optimal entities. Beyond its practical utility for managers, our model provides new insights into human-GenAI mutual augmentation, the nature of creativity, and the skills and cognitive properties involved.

## 1. Introduction

Generative AI (GenAI) is rapidly emerging as one of the most groundbreaking technologies of the century. Its revolutionary aspect lies in the ability to recombine knowledge by uncovering patterns within text, images, and sound with unprecedented effectiveness. This enables GenAI algorithms to create new knowledge based on the patterns inferred in the training data, a property that had always been exclusive to human agents (Mariani and Dwivedi, 2024; Raisch and Fomina, 2025). GenAI is already having a remarkable impact on industries like healthcare, finance, education, and entertainment, and its performance is predictably increasing at a startling rate according to known scaling laws (Kaplan et al., 2020).

Firms operate in a constantly evolving environment, which requires them to innovate and reconfigure their resources and capabilities to improve their evolutionary fit (Nelson and Winter 1973; Teece et al., 1997; Teece, 2007). Radical changes within and/or outside the firm often pose problems without clear pathways to solutions. Possible problems range from developing a new product to devising a strategy to react to a lateral threat. Finding solutions to such problems often requires navigating a virtually unbounded space. This spawned a

long-standing stream of research investigating the drivers, impediments, and outcomes of organizational search (Fleming and Sorenson, 2004; Katila and Ahuja, 2002; March, 1991).

Due to bounded rationality and cognitive inertia, firms rely on heuristics and fall prey to biases in their search endeavors, which result in local search and thinking within existing paradigms, often with detrimental consequences on their survival (Gavetti et al., 2012; Levinthal and March 1993; Simon, 1955; Tripsas, 2009; Tripsas and Gavetti, 2000). Finding solutions to novel problems is likely to require escaping extant paradigms, often through groundbreaking ideas (Girotra et al., 2010; Kornish and Ulrich, 2011). Idea generation and selection, along with their underlying driver (i.e. creativity), have been the object of numerous studies within the innovation domain. Specifically, scholars have investigated the determinants of creativity, the characteristics of idea generation and selection processes, and the creative dynamics stemming from interactions in large groups (Amabile, 1983, 1988, 2020; Amabile and Pratt, 2016; Deichmann and Jensen, 2018; Harvey, 2014; Mannucci and Perry-Smith, 2022; Perry-Smith and Mannucci, 2017; Simonton, 1999, 2023; Woodman et al., 1993).

GenAI holds great promise for enhancing the ability of firms to generate creative ideas and solutions. However, research on this topic is

\* Corresponding author.

E-mail address: [mattia.pedota@polimi.it](mailto:mattia.pedota@polimi.it) (M. Pedota).

<https://doi.org/10.1016/j.technovation.2025.103323>

Received 19 November 2024; Received in revised form 6 June 2025; Accepted 19 July 2025

Available online 7 August 2025

0166-4972/© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

still in its early stages (Akter et al., 2023; Chiarello et al., 2024; Mariani and Dwivedi, 2024; Singh et al., 2024). Bouschery et al. (2023) theorize some mechanisms whereby GenAI may augment innovative efforts (particularly in new product development), by aiding human agents in exploring both the problem and the solution space through text summarization, sentiment analysis, and idea generation. Relatedly, Bilgram and Laarmann (2023) provide a series of use-cases illustrating how GenAI may augment exploration, ideation, and prototyping. Raisch and Fomina, (2025) create a taxonomy of human-(Gen)AI interactions in solution search, theorizing that autonomous search leads to more distant solutions, sequential search to more local solutions, and interactive search to more recombinative solutions. Dell'Acqua and colleagues (2023) document the emergence of two archetypes of human-GenAI interactions: centaurs (who selectively delegate subtasks to GenAI) and cyborgs (who collaborate closely with GenAI). At the same time, they note that “the best approaches to using AI are not fully understood and need to be deeply examined by scholars and practitioners” (Dell'Acqua et al., 2023: p.17).

This remark highlights an important gap: despite the emerging body of research, little is known on how to optimize human-GenAI interactions for creativity. While it is clear that human-GenAI interactions have a strong value-generating potential in organizations (especially in terms of innovative capabilities), extant literature on the relationship between AI and innovation (and AI and creativity in particular) is still silent on how to actively maximize this potential. Experiments have been conducted in different settings, showing that human-GenAI and human-only creative outcomes are different in terms of appropriateness and innovativeness, with the former usually winning in the appropriateness dimension (Boussieux et al., 2024; Dell'Acqua et al., 2023; Doshi and Hauser, 2024; Girotra et al., 2023; Koivisto and Grassini, 2023; McGuire et al., 2024). However, the focus of these studies is mostly on the outcome, leaving the process dimension in the background. This lacuna prompts us to formulate the following research question: *How can human-GenAI interactions be structured to maximize creative potential in performing creative tasks?*

We address this question by proposing a model of hybrid creative process that considers the peculiarities and relative advantages of each component of the human-GenAI dyad. Drawing on machine learning literature, we proceed from a detailed account of the cognitive architecture of GenAI to define a process that maximizes the creative potential of the human-GenAI interaction. At the core of the model lies the conceptualization of GenAI as a superposition of latent states, which we term entities (Shanahan et al., 2023). Given this premise, we show that maximizing the creative potential of the human-GenAI interaction primarily depends on the effective elicitation of simulated entities.

This work contributes to management theory and practice in three important ways. From a theoretical standpoint, we extend the rich research tradition modeling the (human-only) creative process (Amabile, 1983, 1988; Amabile and Pratt, 2016; Kozbelt et al., 2010; Simonton, 1999, 2023). This work opens a new research direction by formalizing a new model of hybrid creative process involving human-GenAI interaction. In so doing, it also contributes to the more recent literature on AI-human ensembles, hybrid intelligence, and AI-enabled automation-augmentation (Anthony et al., 2023; Choudhary et al., 2023; Raisch and Fomina, 2025; Raisch and Krakowski, 2021).

Furthermore, our model implies a redistribution of the importance of creativity-relevant skills and the underlying cognitive properties, by advancing the role of human agents as selectors and orchestrators of GenAI entities. Creativity is often associated with extraordinary cognitive features and skills (Barron and Harrington, 1981; Guilford, 1984; Singh, 1986). In contrast, our work foregrounds the value of a mix of more ordinary skills (e.g. adaptive learning, rational decision making, contextual understanding) for unleashing the creative potential of the human-GenAI dyad, with notable implications for AI-based competitive advantage (Kemp, 2024; Krakowski et al., 2023).

From a practical standpoint, we provide firms with a powerful tool to

enhance their performance in creative tasks, including R&D, organizational search, resource reconfiguration, problem-solving, and new product development. As managers and employees have just started to experiment with the inclusion of GenAI in such tasks, they are often ill-equipped to exploit its full potential. By showing how to optimally structure the interaction between human agents and GenAI, our model also has a valuable prescriptive purpose.

The remainder of the paper is organized as follows. Section 2 lays the theoretical framework for our contribution based on a corpus of machine learning literature. Section 3 develops our model of hybrid creative process. Section 4 concludes by discussing implications for management theory and practice.

## 2. The cognitive architecture of GenAI models

The last two years have witnessed remarkable strides and increasing competition among top contenders for the development of novel AI models, particularly generative AI (GenAI) models. We define GenAI as a computational technique aimed at producing content, such as text, images, or music, in response to (primarily text-based) user prompts, through patterns and structures learned from training data (Feuerriegel et al., 2024). Currently, the most popular GenAI model is the Large Language Model (LLM), capable of navigating the complex patterns of language to generate human-like text. OpenAI's pioneering LLM GPT-3 already excelled in a wide array of tasks through few-shot learning (Brown et al., 2020). Google AI's PaLM further extended capabilities in multilingual and reasoning tasks. OpenAI's subsequent iteration, GPT-4, showcased notable developments in generating human-like text and introduced advanced abilities such as zero-shot and in-context learning (Achiam et al., 2023). Similarly, Google DeepMind's Gemini utilized reinforcement learning to enhance proficiency in language understanding and generation (Hoffmann et al., 2022). Behind these advancements lie extensive training data and sophisticated modeling and optimization strategies, such as multimodal modeling and parameter scaling (Kaplan et al., 2020; Li et al., 2021). Significant advancements have been made in multimodal GenAI, showcasing GenAI's potential in various creative domains. SunoAI's music creation model exemplifies the capabilities of GenAI in creating music, generating original compositions that can mimic different genres and styles in a human-like fashion. OpenAI's DALL-E stands out for its ability to generate GPT-powered high-quality images from textual descriptions. Other significant text-to-image models include MidJourney and Stable Diffusion (Rombach et al., 2022).

GenAI demonstrates notable strengths over human capabilities in certain domains, mainly due to its underlying cognitive architecture and training methodology. Its efficiency in processing and analyzing vast corpora of data, coupled with its ability to identify complex patterns, enables it to generate coherent output across a wide range of applications (Nijkamp et al., 2022; Brown et al., 2020; Radford et al., 2019).

The scalability and processing speed of GenAI constitute significant advantages, allowing it to analyze extensive datasets and generate insights at a rate far surpassing human cognitive abilities (Brown et al., 2020). These advantages stem largely from GenAI's associative memory, namely its ability to establish connections between diverse pieces of information based on learned patterns. This capability is predominantly horizontal in nature, enabling GenAI to create shallow associations across a broad spectrum of topics. Nevertheless, while this feature allows e.g. for the generation of diverse text and the emulation of human ingenuity, it often results in responses that are contextually superficial (Bender et al., 2021).

GenAI's reliance on statistical patterns identified within training data, rather than explicit symbolic or causal understanding, constrains the development of a robust internal world model. This limitation affects performance in tasks requiring deep hierarchical reasoning, including mathematical computations or rigorous logical deductions (Hendrycks et al., 2021). Consequently, GenAI outputs, while often superficially

plausible, can lack consistency and adherence to logical structures (Bender et al., 2021; Bubeck et al., 2023).

Enhancing GenAI's deeper reasoning capabilities remains an active research frontier. Current approaches explore various promising directions, including the integration of symbolic reasoning capabilities, the development of neuro-symbolic hybrid models (Lamb et al., 2020), and architectures like JEPA (Joint-Embedding Predictive Architecture) that aim to incorporate predictive world models (LeCun, 2022; Bardes et al., 2024). Additionally, prompting techniques such as chain-of-thought (Wei et al., 2022 a,b) and tree-of-thought (Yao et al., 2023) have demonstrated improvements in complex reasoning tasks. Simultaneously, challenging benchmarks like MATH (Hendrycks et al., 2021) and BIG-Bench (Srivastava et al., 2022) are being developed to rigorously evaluate progress in this domain.

### 3. A model of hybrid creative process

In this section, we describe the theoretical approach for modeling and optimizing the hybrid creative process. Methodologically, we rely on mathematical deduction and formal argumentation. In our framework, GenAI is represented as a dynamic mixture of latent states, referred to as "entities", each contributing with a certain weight. This representation parallels the principles of Bayesian updating. By relying on mathematical models that quantify the creativity of an output, we analytically demonstrate that, for any creative task, there is always at least one entity whose activation can maximize the creative outcome. In formal terms, these findings imply that the design of prompts may function as a control mechanism that reweights the probability distribution over the latent entities, suppressing the activation of suboptimal states while increasing the likelihood of selecting the entity that maximizes the creative outcome. Furthermore, we describe a heuristic strategy inspired by Bayesian optimization to navigate toward entity configurations that yield higher creative performance. Finally, we discuss the cognitive skills underlying this process and the importance of decomposing problems into subcomponents while effectively integrating the solutions provided by each entity. Table 1 below summarizes the key terms and formulas underlying the model.

#### 3.1. GenAI as a superposition of latent entities

We conceptualize a GenAI as a superposition of latent entities, each of which can be elicited through suitable prompts (Shanahan et al., 2023). The notion of superposition in this context refers to GenAI being a dynamic amalgamation of multiple states. This phenomenon can be regarded as a specific manifestation of in-context learning (ICL). ICL emerges as a powerful paradigm where GenAI models interpret and generalize from a few contextually provided examples without parameter updates (Brown et al., 2020). Theoretical interpretations suggest that ICL mimics implicit Bayesian inference, with the models performing probabilistic inferences based on the examples provided. This is akin to updating beliefs with new evidence, allowing models to predict outcomes by identifying shared latent concepts among examples (Zhao et al., 2023). The process inherently involves activating various latent entities within the model based on the given prompts, where each latent entity represents a distinct mode of expression, perspective, or knowledge domain, dynamically contributing to the model's response. The capability of ICL is attributed to extensive and diverse pretraining, enabling GenAI to internalize patterns and generalize across tasks effectively (Wei et al., 2022b; Bubeck et al., 2023). Additionally, meta-optimization suggests that the inference process during ICL resembles gradient descent optimization, where the model's internal dynamics adjust representations similarly to parameter updates in gradient-based learning (von Oswald et al., 2023).

The adaptive mechanism of entity superposition can be formalized by a probabilistic model, where each entity is associated with a probability indicative of its activation likelihood. As the response is associated

**Table 1**

Summary of the primary elements of our hybrid creative process model in terms of key concepts and methodological components.

LLM	Large Language Models, like GPT, Gemini, and PaLM
GenAI	Generative Artificial Intelligence, including LLMs, as well as other generative models (e.g., MidJourney, DALL-E, SunoAI)
ICL	In-Context Learning, a paradigm where GenAI models learn from a few context-provided examples without updating their parameters
Simulated Entities $\{\theta_i\}$	Latent states or configurations within a GenAI model that can be "activated" via specific prompts, representing different expressive modes or knowledge domains
$\{\alpha_i\}_{i=1,\dots,N}$	Normalized positive coefficients indicating the probability that the model exactly embodies the entity $\{\theta_i\}$
$ \psi\rangle = \sum_i \alpha_i  \theta_i\rangle$	The state of a GenAI can be seen as a superposition of entities
$ \psi'\rangle = \mathcal{O}( \psi\rangle) = \sum_i \alpha'_i  \theta_i\rangle$	When the GenAI interacts with an input, the state $ \psi\rangle$ evolves according to an operator $\mathcal{O}$ , which modifies the probability distribution over the entity states
<b>Problem Decomposition</b>	The strategy of breaking a complex problem into smaller, manageable sub-problems/tasks $T$ to simplify analysis and resolution
$c = u \cdot l$	The creativity of an outcome is assumed as the product between two factors. E.g., utility $u$ and learning $l$ (Tsao et al., 2019)
$l_{ \psi\rangle, T} = l_{\theta_i, T}$ where $Z \sim \text{Categorical}(\alpha_1, \dots, \alpha_N)$	The learning value of the solutions for task $T$ generated by the GenAI is modelled as that of the single elicited entity
$P(l_{ \psi\rangle, T} > k) \leq P(l_{\theta_M, T} > k)$ , where $M = \text{argmax}_i \{P(l_{\theta_i, T} > k)\}$	There always exists a $T$ -specific entity that is equally or more likely to yield a solution with a high value of learning (and ultimately of $c$ )
$\mathbf{P}_\omega(\omega_{k+1}) := \mathbf{P}(\omega_{k+1}   \omega)$ $\propto \sum_{\theta_i \in L_D} \mathbf{P}_{\theta_i, \omega}(\omega_{k+1}) \mathbf{P}_{\theta_i}(\omega   \theta_i)$	By using suitable prompts $\omega$ , the generation process of e.g. an LLM can be guided so that essentially only the desired entities (those in $L_D$ ) contribute to it
<b>Bayesian Optimization</b>	A probabilistic technique can be heuristically employed to iteratively refine the selection of entities by updating beliefs based on observed performance
<b>Output Integration</b>	The phase where results from various sub-problems are collected and combined to form a coherent, complete solution

with the activated entity, proper prompting enables more coherent and contextually appropriate output across diverse scenarios.

Entities are divided into agentic and non-agentic (Achiam et al., 2023). Agentic entities exhibit behaviors and responses that resemble those of intentional agents, despite the underlying GenAI model lacking genuine agency or intentionality (Rabinowitz et al., 2018). In contrast, non-agentic entities provide descriptive or informational content without implying an underlying goal or intention. In our model, the activated entity can be either agentic or non-agentic, as long as it embeds

relevant disposition and expertise relative to the focal problem.

An entity may coincide with a deeply specialized sub-model (Hajikhani and Cole, 2024), which focuses on a narrow domain of expertise. Alternatively, it might embody a psychological entity composed of specific crafted traits (Wang et al., 2025), allowing the model to emulate nuanced thought patterns and human behavior. In an LLM-based GenAI, the activated entity might encode different writing styles, each offering unique features such as deciding whether to use figures of speech (e.g., metaphors, onomatopoeia, hyperboles), selecting suitable registers (e.g., formal, slang, technical), and choosing different sentence structures. For text-to-image GenAIs, the activated entity might encode different styles for image advertising to promote products, such as hyperrealistic, impressionistic, or comic art (Ramesh et al., 2021). Similar arguments can be made for other types of GenAI, where the activated entity would vary based on the specific type of generated content.

Within this framework, each entity represents a configuration within the GenAI's latent space  $L$ . The state of the GenAI, denoted by  $|\psi\rangle$ , is a superposition of these individual entity states  $\{|\theta_i\rangle\}_{i=1,\dots,N}$ , represented as:

$$|\psi\rangle = \sum_i \alpha_i |\theta_i\rangle,$$

where  $\alpha_i$  are positive coefficients indicating the degree to which the model embodies each entity state  $|\theta_i\rangle$ . These coefficients are subject to the normalization condition  $\sum_i \alpha_i = 1$ .

When the GenAI interacts with an input prompt, the state  $|\psi\rangle$  evolves according to an operator  $O$ , which modifies the probability distribution over the entity states. The new state  $|\psi'\rangle$  is given by  $|\psi'\rangle = O(|\psi\rangle) = O(\sum_i \alpha_i |\theta_i\rangle) = \sum_i \alpha'_i |\theta_i\rangle$ , where  $\alpha'_i$  are the updated probabilities after the application of the operator  $O$ . Thereafter, the GenAI's latent superposition collapses, probabilistically selecting a single entity to generate the final output.

This framework provides a foundation for optimizing entity activation within the GenAI, also facilitating the decomposition of complex tasks into manageable subproblems. By strategically leveraging targeted prompts, it is possible to fine-tune activation probabilities to enhance coherence and performance. As both human agents and GenAI face processing capacity challenges, employing specialized entities for each subproblem not only alleviates this load but also maximizes creative output. Thus, understanding the dynamic interplay of entities within GenAI sets the stage for more effective interaction strategies and a structured problem-solving process, as detailed in the following subsections.

### 3.2. Simulated entity elicitation

As discussed previously, GenAI can be seen as a superposition of latent entities  $|\theta\rangle \in L$  (with  $L$  being the latent space) that evaluates inputs to generate prompt-specific outputs. For instance, in the case of an LLM, the model evaluates  $P_\omega(\omega_{k+1})$ , i.e. the approximation of the conditional probabilities of the token  $\omega_{k+1}$  appearing as the next term in the token sequence  $\omega := (\omega_0, \dots, \omega_k)$ . By using appropriate prompts  $\omega$ , it is possible to guide the generation process (Nardo, 2023). Indeed,

$$P_\omega(\omega_{k+1}) = \sum_{|\theta\rangle \in L} P_\omega(\omega_{k+1}, |\theta\rangle) = \sum_{|\theta\rangle \in L} P_{|\theta\rangle, \omega}(\omega_{k+1}) P_\omega(|\theta\rangle) \propto \sum_{|\theta\rangle \in L} P_{|\theta\rangle, \omega}(\omega_{k+1}) P_{|\theta\rangle}(\omega) P(|\theta\rangle),$$

where the subscripts denote the conditioning, and  $L_D$  only contains desired entities, e.g. those that are more appropriate for solving the given task; those in  $L_U$ , such that  $L_U = L \setminus L_D$  have negligible  $P_\omega(|\theta\rangle)$  (if  $\omega$  is well-suited for inhibiting the undesired entities). This means that primarily entities in  $L_D$  contribute to creating the sought outputs. By

observing the response  $P_\omega(\omega_{k+1})$  it is possible to progressively refine  $\omega$  so that the probabilities associated with undesired entities become negligible, allowing the model to generate the desired outputs by focusing on the relevant entities. Models other than LLMs follow analogous (GenAI-specific) processes.

We now explore why using suitable simulated entities (those in  $L_D$ ) can be advantageous for the creative generation process using multiplicative models of creative variations. Several models have been introduced to quantify the creativity of a generated output. Simonton (2015, 2023) proposes that the creative outcome  $c$  of an idea can be subjectively quantified as  $c = u(1-p)(1-v)$ , where  $u$  represents the eventual utility of the idea,  $p$  indicates the initial probability that the idea would have been generated, and  $v$  signifies the degree of certainty regarding the eventual utility of the idea. All the parameters range in  $[0, 1]$ , but simplifying them to the extremes of this interval (i.e., 0 or 1) results in a model with 8 distinct pure types, forming  $2^3$  combinations. Among these, one type  $(u, p, v) = (1, 0, 0)$  embodies an optimally creative outcome, while the remaining seven pure types represent uncreative outputs. As both factors  $(1-p)$ , related to originality, and  $(1-v)$ , related to Campbell's (1960) and Simonton's (2010) blindness, can be considered indicators of learning, the formula above can be modified with the insertion of a learning term (Tsao et al., 2019).

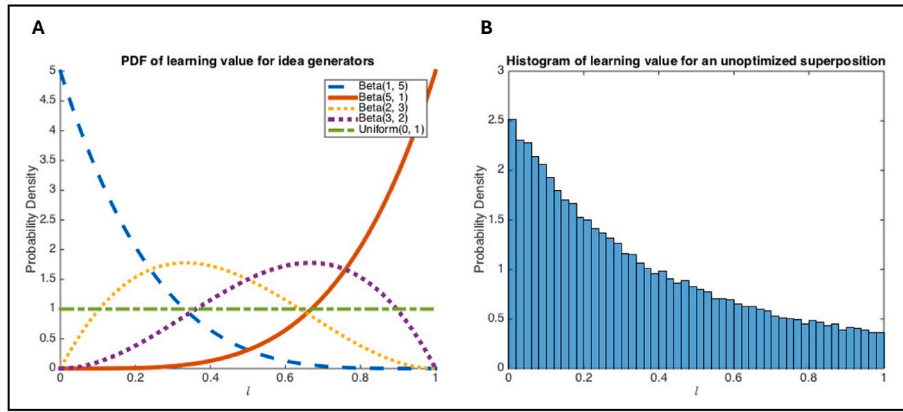
In their model, Tsao et al. (2019) defined  $c$  as the product between  $u$  and  $l$ , where  $u$  remains as described before, and  $l$  characterizes learning. Thus, the formula characterizes a creative outcome as the product of utility and learning, serving as a measure of "useful learning". If either utility or learning is low, the creative outcome is also low; for a high creative outcome, both utility and learning must be high.

For the sake of clarity, let us suppose now that, regarding a creative task  $T$ , each point in  $[0, 1]$  corresponds to the learning value  $l$  of a solution of the task, and assume that every generated output has a utility value equal to 1 if it is a solution, and 0 otherwise. Note: a completely analogous discussion could be conducted assuming that the value in  $[0, 1]$  encodes, e.g., the novelty associated with having cast that idea (Runco and Jaeger, 2012), or any other appropriate measure of creativity.

Thus, the learning parameters of the solutions for a task obtained by an idea generator, whether this generator is a human agent, a GenAI, or a combination of both, can be regarded as realizations of a random variable over the interval  $[0, 1]$ . The probability distribution may significantly vary among generators depending on the features of the generators themselves. For instance, one might be associated with a Beta distribution with parameters  $(1, 5)$ , indicating a low likelihood of producing high learning value outputs, a Beta distribution with  $(5, 1)$ , suggesting a high ability to generate high learning value outcomes, or a uniform distribution, implying an equal probability of producing outputs across the entire range of learning values (see Panel A of Fig. 1 for a graphical representation of the probability density functions of different random variables).

Now, let us take  $N$  simulated entities  $\{|\theta_i\rangle\}_{i=1,\dots,N}$  such that the probability for the individual state to generate a solution with a learning value  $l$  for the task  $T$  follows a distribution in  $[0, 1]$ . Note that we did not specify the random variables associated with the generators; they can thus follow any probability distribution within the range. In this context, framing the GenAI model as a superposition  $|\psi\rangle$  of entities  $\{|\theta_i\rangle\}_{i=1,\dots,N}$  implies that its ability to generate solutions with a specific learning value distribution reflects at any given time the capability of one, and only one of the multiple entities, depending on which entity is active. Hence, being  $l_{|\psi\rangle, T}$  and  $\{l_{|\theta_i\rangle, T}\}_{i=1,\dots,N}$  the random variables in  $[0, 1]$  regulating the learning value of the solutions for the task  $T$  generated, respectively, by the GenAI model and its components, we have that  $l_{|\psi\rangle, T} = l_{|\theta_z\rangle, T}$  where  $Z \sim \text{Categorical}\{\alpha_1, \dots, \alpha_N\}$ ,  $P(Z=i) = \alpha_i$ , (assumed as independent on  $\{l_{|\theta_i\rangle, T}\}_{i=1,\dots,N}$ ).

A given prompt or history of prompts determines a set of weights  $\{\alpha_i\}_{i=1,\dots,N}$  that prioritizes some entities over others. Crucially, a care-



**Fig. 1.** Panel A illustrates examples of probability density functions representing the ability of five generators (humans, simulated entities  $\{|\theta_i\rangle\}_{i=1 \dots 5}$ , etc) of providing solutions with learning value in the range  $[0, 1]$  for a given task  $T$ . Each curve represents a specific process. Panel B shows the probability density function as a histogram obtained by simulating (with  $10^5$  realizations) learning values associated with a superposition  $|\psi\rangle$  of the five solution generator of Panel A, with specified weights  $\alpha_i\}_{i=1 \dots 5}$  (taking the variables in the order shown in the legend of Panel A, respectively:  $\alpha_1 = 1/2$ ;  $\alpha_2 = 1/16$ ;  $\alpha_3 = 1/8$ ;  $\alpha_4 = 1/4$ ;  $\alpha_5 = 1/16$ ). It is evident that an unoptimized superposition, i.e. a GenAI model with a naive interaction in this context, exhibits a lower expected value of generating ideas with high learning value compared to utilizing some of the individual underlying processes (i.e. simulated entities) that compose the entire process (e.g. Beta(5,1)).

fully crafted sequence of prompts can further reduce the weights of entities that are useless or even detrimental to the purpose, thus optimizing the generative process. Indeed, we have that  $P(l_{|\psi\rangle,T} > k) \leq P(l_{|\theta_M\rangle,T} > k)$ , where  $M = \operatorname{argmax}_i \{P(l_{|\theta_i\rangle,T} > k)\}$ , i.e. there always exists a task-specific simulated entity of the model that is equally or more likely to yield a solution with a high value of learning (and ultimately of the creativity  $c$ ) for the task  $T$ . Furthermore, there also always exists an entity for which the expected learning value is higher than, or equal to, the one associated with any  $|\psi\rangle$  (for a computational example, see Panel A, in particular Beta(5,1), and Panel B of Fig.1). Therefore, the possibility of dynamically modifying the weights enables increasing e.g. the probability  $\alpha_M$  of eliciting the desired entity  $|\theta_M\rangle$ , maximizing the probability of generating a highly creative solution. Clearly, there also exist entities with expectations lower than or equal to that of  $l_{|\psi\rangle,T}$  (see Panel A, in particular Beta(1,5), and Panel B of Fig.1). Therefore, the benefits of entity elicitation hinge upon the ability to search for and elicit an appropriate entity.

While we have shown that an optimal entity maximizing the expected creativity of its output for a given task always exists, finding it is unlikely due to bounded rationality (Simon, 1955). The main implication of the existence of an optimal entity is the prescription of a process to heuristically converge toward it. Analytically, the process resembles a Bayesian optimization approach, i.e. a probabilistic model-based technique for finding the maximum/minimum of an objective function that is expensive to evaluate, by iteratively building a surrogate model to make informed decisions about where to sample next (Shahriari et al., 2015). Specifically, the goal is to maximize a function  $C(|\theta\rangle)$ , which can be interpreted as an estimate of the probability that the creativity of the outcomes (e.g. modeled as Tsao et al., 2019) may be equal to 1 when using  $|\theta\rangle$  as the entity. The argument that achieves its maximum,  $|\theta_M\rangle = \operatorname{argmax}_i \{C(|\theta_i\rangle)\}$ , represents the optimal entity to converge toward<sup>1</sup>.

For converging to  $|\theta_M\rangle$ , the human agent must balance between trying new, untested regions of the latent space and focusing on areas that have shown high performance, e.g. guided by an heuristic version of

the expected improvement (EI; Gelbart et al., 2014). In this framework, EI measures the expected gain from sampling a particular entity, with the aim of improving upon the best response observed so far. EI may help the human agent decide which point in the latent space to test next by evaluating how much improvement one might expect by trying a novel entity compared to the best one observed  $|\theta_+\rangle$ . Mathematically, the EI at a state  $|\theta\rangle$  is defined as the expected value of the improvement:  $A(|\theta\rangle) = E(\max\{0, C(|\theta\rangle) - C(|\theta_+\rangle)\})$ , where  $E$  denotes the expected value of a random variable. Under Gaussian assumptions,  $A(|\theta\rangle)$  depends on two key parameters associated with  $|\theta\rangle$ : the predicted mean and variance. A higher mean indicates that the entity configuration is expected to provide high-quality, creative responses based on current knowledge, while a higher variance indicates greater uncertainty about the predicted quality, suggesting the potential for discovering unexpectedly high-performing configurations.

At each iteration, the human agent selects an entity to test based on past observations, identifying a new point in the latent space that seems promising. Eliciting the selected entity through prompts, the human agent evaluates the quality of the responses and incorporates this new information to refine its prior on  $C(|\theta\rangle)$ , akin to updating a mental model of the latent space. This process is then repeated iteratively, continuously adjusting the representation of the potential of the points in the latent space. If significant improvements are no longer observed, the process can be considered converged.

From the perspective of the human agent, three key skills underlie the success of the search process. The first is the capability of updating the (mental) model of  $C(|\theta\rangle)$  as a function over the latent space based on new observations. This involves learning the relevance of the entities for the problem at hand, by intuitively refining estimates of both the average response and responses' variability for each entity configuration, leading to an iterative process that enables the human agent to form a progressively clearer picture of which entities are likely to perform well and which hold potential but are indeterminate. Given the vast number of potential entities, it is impossible for the human agent to ascertain the value of all configurations. Therefore, this process relies heavily on an intuitive understanding of the entities' relevance in the latent space.

The second is the human agent's decision-making skill regarding subsequent search steps, where the options to choose from depend on what has been learned about the latent space. This skill refers to the human agent's accuracy in determining the maximum expected improvement  $A(|\theta\rangle)$  for the next step. Unlike the first skill, which is about updating the model based on past observations, this one

<sup>1</sup> As discussed above, the impossibility of analyzing all the entities (e.g., due to the high-dimensionality of the latent space and the bounded rationality of the human agent) makes it utopian to reach the truly optimal one. Nevertheless, for the sake of simplicity, we will continue using the term "optimal entity" instead of a more elaborate description that highlights subjective considerations, such as the "entity that the human agent believes to be optimal" for solving the subproblem.

emphasizes strategic decision-making for future steps. As the human agent selects the next entity to explore based on extant knowledge, this skill requires synthesizing information efficiently and effectively to decide on the next action. A crucial aspect of it is the avoidance of local search biases (Levinthal and March, 1993; March, 1991 Tripsas and Gavetti, 2000). While many search processes are constrained by path dependence and resource scarcity (Levinthal, 1997), this one places minimal limitations on the direction and extent of exploration. In particular, distant exploration incurs very little cost and could potentially reveal highly promising areas (e.g. by visiting entities with great subjective uncertainty) of the latent space that local moves might miss. While moving close to previously high-performing entities ensures finding a local maximum, it does not guarantee discovering the global maximum. Still, human agents may dwell on neighboring areas in the latent space (e.g. by changing only a few attributes per iteration), due to anchoring effects and cognitive inertia (Tripsas and Gavetti, 2000; Tversky and Kahneman, 1974). Thus, rationality is required in balancing the trade-off between finding the region with the highest potential for a global maximum and exploring it effectively to move as close as possible to its (local) maximum.

The third is the technical skill to design effective prompts that elicit the desired entity  $|\theta\rangle$  while inhibiting undesired processes. This requires a deep understanding of the functioning of the GenAI being used, including an intuitive grasp of how different prompts influence the responses and allow for reliably bringing out the specific traits of the entity being tested. This nuance of prompt engineering increases the likelihood of each interaction yielding valuable responses, facilitating a more efficient and precise exploration process. Following the receipt of responses, the iterative process restarts, reiterating the application of the first skill to evaluate and refine the model  $C(|\theta\rangle)$  of the entities' relevance.

### 3.3. Decomposing problems and integrating entity solutions

Modularization is a cornerstone of management and organization theory (Eppinger and Browning, 2012; Ethiraj and Levinthal, 2004; Ulrich and Eppinger, 1995). In the realm of problem-solving, breaking down a problem in parts, solving them separately, and then integrating the resulting solutions is a well-known design thinking practice (Liedtka, 2015; Verganti et al., 2021). In this subsection, we explain how modularization can also be exploited to optimize the effectiveness of simulated entity elicitation for creative purposes. While some of the underlying principles partly overlap with well-known benefits of problem modularization (e.g. mitigation of cognitive overload), others stem from specific peculiarities of this instance (e.g. GenAI features and human-GenAI relative advantages).

Human working memory has limited processing capacity, which can be overwhelmed by large amounts of data to handle (Richardson et al., 1996). This limitation means that when faced with excessive information, the ability to process and retain details is significantly compromised, leading to cognitive overload. Interestingly, challenges with processing capacity aren't exclusive to human agents; they extend to interactions with GenAI as well (Achiam et al., 2023). While technical constraints (e.g. reaching token limits) do play a role, the primary challenge stems from the inherent limitations in terms of fine attention span. This limitation leads GenAI to inadvertently favor certain subsets of information over others, ultimately resulting in forgetfulness (Luo, et al., 2023). Moreover, context switching between varied tasks can further impede the model's ability to maintain coherence. Segmentation strategies, which involve breaking down a problem into smaller, more manageable subproblems (Halford et al., 1998), can thereby prove advantageous. Implementing such strategies not only mitigates overload, but also facilitates the attainment of creative output generated by different entities.

This implies structuring the problem domain in a way that allows for modularization and coordination of subproblems. Independent

components can be addressed concurrently, with each component assigned to a different entity for parallel processing (within different GenAI-instances/conversations). On the other hand, components that exhibit unidirectional dependencies may need to be handled sequentially (i.e. within the same GenAI-instance/conversation), with the model roleplaying differently for each task to ensure the coherent integration of outputs<sup>2</sup>.

Once a satisfactory entity has been identified for each subproblem, the human agent elicits it an arbitrary number of times to generate multiple ideas. The higher the variance associated to the selected entity, the higher the number of times it needs to be elicited to capture the full spectrum of ideas it may generate. Furthermore, even apparently simple subproblems are still likely to allow for recombination of ideas (Fleming and Sorenson, 2004). Thus, ideas generated by each entity within each subproblem may be taken as they are, fine-tuned by the human agent, or sequentially recombined to generate further ideas. In this phase, it may still be convenient to retain many promising ideas rather than converge on a single best idea for each subproblem. The reason lies in the interdependencies between subproblems: as ideas for each subproblem are to be combined toward solving the overarching problem, the quality of ideas for each subproblem partly depends on the pool of ideas of the other subproblems.

After selecting a series of promising ideas for each subproblem, the human agent is responsible for bringing them together. Here is where human agents bring to bear their relative advantage vis-à-vis GenAI. First, they integrate a wealth of information that goes beyond the creative task (Balasubramanian et al., 2022). Human agents know where, why and how the idea is to be deployed, the cultural and social nuances of the context (e.g. the characteristics of the target niche of consumers), the firm's corporate vision and its strategy (both short-term and long-term), and the explicit and implicit criteria whereby the idea will be evaluated (Csikszentmihalyi, 1996). While a part of this information may be fed to the AI through voluminous ad-hoc training data, a large part of it would be inevitably lost, as it involves a significant degree of tacit knowledge (Nonaka, 1994; Nonaka and Takeuchi, 1995; Nonaka and von Krogh, 2009). Second, human agents can construct semantic relationships between ideas, rather than mere correlational links. They make sense of incoming knowledge by linking it to extant knowledge, i. e. they naturally engage in associative learning (Cohen and Levinthal, 1990; Todorova and Durisin, 2007). Semantic connections are better suited to construct a convincing narrative linking together the ideas associated with each subproblem.

Thus, human agents leverage their tacit knowledge and contextual understanding on the one hand, and the ability to create semantic connections on the other, to recombine the ideas emerging from each subproblem toward solving the overarching problem. Considering an example from the video game industry, a subject may decide to give a character (created by an ad-hoc elicited entity) a dark twist as it is combined with a post-apocalyptic setting (created by a different elicited entity), both for thematic coherence (semantic understanding) and for a better alignment with a new trend started by a recent movie (contextual understanding).

In most business contexts, extremely good ideas are exponentially more valuable than good ideas (Girotra et al., 2010). A correct output integration ensures that the value of extreme ideas generated by the various ad-hoc entities is preserved without sacrificing contextual adherence and sociocultural (as well as strategic) alignment. We regard this approach as superior to keeping the overarching problem as a whole

<sup>2</sup> This approach bears resemblance to the classical "divide-and-conquer" paradigm, wherein a complex problem is decomposed into smaller, more tractable subproblems of the same type, which are solved independently and subsequently recombined to yield a solution to the original problem (Smith, 1985). However, our method allows for subproblems that may differ substantially in nature from the original problem.

and eliciting a single entity to solve it, for multiple reasons. First, keeping the problem as a whole makes the search for an optimal entity more difficult, as the elicitation of an entity for a complex problem requires relatively more prompts. Second, once the entity has been identified, it would be computationally less effective at tackling a complex problem, due to limited processing capacity (Achiam et al., 2023). Third, the chosen entity would be less likely to generate extremely creative ideas, as the higher complexity of the problem would lead it toward generic responses that average out all the inputs. Fourth, the chosen entity would implicitly face the task of connecting the micro-ideas underlying the idea it generates. In this respect, its incapability of performing semantic reasoning would likely lead to an incongruent narrative and/or a comparatively weak contextual fit (Balasubramanian et al., 2022; Bender et al., 2021; Townsend et al., 2024). While such issues may still be mitigated via human fine-tuning, fine-tuning the various dimensions of a complex idea without changing its overall value may be significantly more difficult than fine-tuning micro-ideas as they are recombined.

Fig. 2 provides a graphical overview of our framework.

#### 4. Discussion

This study set out to explore a simple yet foundational question in the emerging literature on human-GenAI collaboration: *How can human-GenAI interactions be structured to maximize creative potential in performing creative tasks?* In response, we propose a new model of hybrid creative process that conceptualizes GenAI as a dynamic superposition of latent entities. Specifically, we demonstrate that optimal creative outcomes can be systematically pursued by searching for an optimal simulated entity, guided by an iterative heuristic process akin to Bayesian optimization.

Extant studies have primarily focused on the emerging outcomes of GenAI-supported creativity and typologies of interaction modes (e.g. Boussioux et al., 2024; Dell’Acqua et al., 2023; Doshi and Hauser, 2024; Girotra et al., 2023; Koivisto and Grassini, 2023; McGuire et al., 2024). In contrast, we leverage a process-oriented view grounded in the cognitive properties of GenAI and human agents to theoretically deepen the optimization of human-GenAI interaction toward creative outcomes. In so doing, we achieve three main results.

First, we reconceptualize GenAI not as a monolithic tool, but as a superposition of latent entities, each representing a unique perspective,

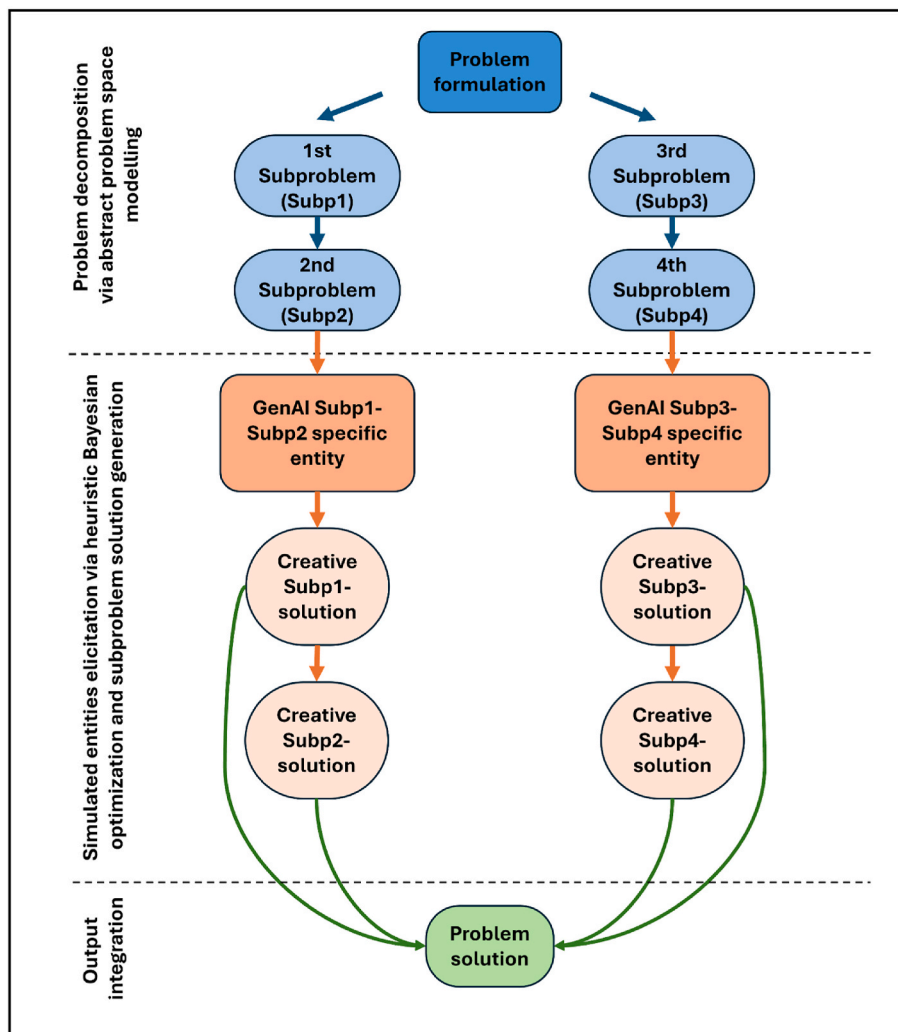


Fig. 2. A graphical representation of the model of human-GenAI hybrid creative process. After formulating the problem to be solved, it is necessary to analyze it in an abstract problem space, allowing for its decomposition into an appropriate number of subproblems (depending on the nature of the problem, the GenAIs used, and other contextual factors). The subproblems can then be addressed sequentially or in parallel, based on their level of interdependency. In the example shown in the figure, the problem has been divided into four subproblems, which are sequentially paired (Subp1 with Subp2, and Subp3 with Subp4). For each of the two pairs, a suitable simulated entity is sought through a heuristic approach resembling the Bayesian optimization technique. Once the four outcomes are generated, they are coherently integrated to form the solution to the initial problem.

mode of reasoning, or domain-specific competence. This framing opens new possibilities for understanding and guiding GenAI behavior in creative contexts and beyond. Second, we provide a formal basis for the idea that human agents can strategically optimize the creative process through the exploration of simulated entities. This represents an advancement with respect to conventional creative processes limited to the exploration of problem and solution spaces. Third, we incorporate principles of problem decomposition and modular recombination to show how human and GenAI capabilities can be integrated in a structured, synergistic process. Taken together, these results contribute to shaping a new way to understand hybrid creativity, with several theoretical and practical implications.

#### 4.1. Theoretical implications

Traditional (human-only) creative processes in management emphasize the set of skills and cognitive features allowing human agents to recombine knowledge effectively toward creative outcomes, including mastery over the focal domain (e.g. knowledge, expertise), creativity-relevant traits (e.g. ad-hoc heuristics, thinking outside the box), and psychological motives (e.g. intrinsic motivation), subject to organizational enablers and pressures (Amabile, 1983, 1988; Amabile and Pratt, 2016; Csikszentmihalyi, 1996; Woodman et al., 1993). Drawing on these underlying factors, human agents (whether individually or collectively) generate a large pool of ideas and subsequently converge on the ideas deemed most promising (Cropley, 2006; Simon, 1999, 2023). Computerized technological artifacts like search engines, word-processing software, and computer-aided design have long been acknowledged to support the creative process by helping human agents retrieve, visualize, and elaborate knowledge (Lubart, 2005; Shneiderman, 2007). However, GenAI is unique in going well beyond a mere supporting role, by potentially taking over the entire task of idea generation. Relative to this stream of literature, we open a new research direction based on the concept of hybrid human-GenAI creative process.

Confronting environmental turbulence, strategic adaptation, and technological evolution, firms constantly navigate problem spaces and draw upon their resources and creativity in search for solutions (Fleming and Sorenson, 2004; Katila and Ahuja, 2002; March, 1991), often in the form of groundbreaking ideas (Girotra et al., 2010; Kornish and Ulrich, 2011). Recent research recognizes the potential of GenAI in supporting the search for solutions (Bilgram and Laarmann, 2023; Bouschery et al., 2023; Mariani and Dwivedi, 2024) and theorizes on the outcomes that different human-GenAI search configurations are likely to yield (Dell'Acqua et al., 2023; Raisch and Fomina, 2025). We complement these studies by unpacking the peculiarities of the cognitive architecture of GenAI and theorizing a hybrid process that maximizes the creative potential of the human-GenAI interaction. Our model underscores the necessity to steer clear of a naïve implementation of GenAI within traditional creative processes and define a new creative process that blends the relative strengths of the cognitive architectures of human agents and GenAI.

By conceptualizing GenAI as a superposition of latent entities, we mathematically show that maximizing the creative potential of solutions generated by GenAI requires human agents to focus on eliciting the optimal entity for the problem at hand. Converging toward the optimal entity requires the interacting human agent to navigate the space of potential entities through an iterative process that resembles a Bayesian optimization approach. For the threefold purpose of facilitating the search for the optimal entity, mitigating GenAI cognitive overload problems, and leveraging human relative advantages in contextual and semantic understanding (Balasubramanian et al., 2022; Townsend et al., 2024), we prescribe decomposing the problem beforehand, finding the optimal entity for each subproblem, and subsequently integrating each entity's output.

Decomposing the problem and eliciting suitable entities rely on skills typically associated with analytical tasks. A suitable decomposition

presupposes identifying its core elements and understanding their relationships and interactions, through logical reasoning and systems thinking. Converging toward the optimal entity presupposes iteratively updating the human agent's mental model of which entities work best, making rational decisions regarding the next entities to test, and eliciting them through suitable sequences of prompts. These skills depend on the human agent's ability to synthesize information, learn adaptively through retrospective analysis, and perform rational decisions, including the elusion of biases like anchoring (Cabantous and Gond, 2011; March, 1991; Tyre and Von Hippel, 1997; Tversky and Kahneman, 1974). On the other hand, the integration of the entities' outputs relies on contextual and semantic understanding, which presuppose social, emotional, and verbal skills (Balasubramanian et al., 2022; Huy, 1999; Kemp, 2024; Townsend et al., 2024; Wright, 2002).

While domain-relevant expertise and analytical thinking have always played a role in creative settings (especially technical problem solving) (Amabile, 1983, 1988; Amabile and Pratt, 2016; Grilli and Pedota, 2024), they acquire much greater importance in interaction with GenAI, even within purely creative tasks. This is because analytical skills now matter not only as a means to navigate problem and solution spaces, but also as a means to structure the interaction with GenAI and navigate the space of potential entities. At the same time, softer skills underlying contextual and semantic understanding also become more relevant, for the mitigation of GenAI's shortcoming and the creative recombination of knowledge modules embedded in subproblems. These insights contribute significantly to the contemporary debate on the upskilling vs deskilling effect of AI (Acemoglu and Restrepo, 2018; Jia et al., 2024; Raisch and Krakowski, 2021; Xue et al., 2022; Zirar et al., 2023). Our model implies that only a naïve implementation of GenAI may result in a total substitution effect, with GenAI taking over the entire idea generation task. In contrast, maximizing GenAI's creative potential does require keeping human agents in the loop, potentially placing an even higher demand on the breadth and depth of the skills required.

#### 4.2. Practical implications

The present work offers valuable prescriptions for firms of any sector and dimension, especially those that frequently engage in creative endeavors and/or problem solving. The implementation of the proposed model may help firms generate more creative solutions as a result of the interaction between employees and GenAI. Our work also entails that the aforementioned mix of hard and soft skills becomes more valuable, with significant implications for human resource management and the job market. In this respect, we suggest that the introduction of GenAI as an integral part of the creative process may reduce the weight of innate extraordinary characteristics (Barron and Harrington, 1981; Guilford, 1984; Singh, 1986). However, this is far from implying that all individuals and organizations will have similar innovation capabilities by merely implementing the suggested process. In contrast, it implies that innovation capabilities will increasingly depend on the ability to acquire and nurture the mix of hard and soft skills underlying the suggested process, foregrounding another facet of AI-based competitive advantage (Kemp, 2024; Krakowski et al., 2023).

Our model also prescribes actively organizing toward a facilitation of optimal human-GenAI interactions. Even in creative industries, where pure analytical skills have traditionally played a less relevant role, firms now need a workforce that can couple soft skills and creative thinking with logics, rationality, and hierarchical reasoning. The latter skillset becomes just as instrumental to creativity as the former, as it represents the bridge that allows employees to leverage the creativity of GenAI. This implies that human resource managers should increasingly hire more balanced profiles, even for purely creative roles.

The same reasoning extends to employee training. Firms should develop targeted learning programs that not only enhance traditional creative abilities but also foster computational thinking, structured

problem-solving, and adaptability in working with AI-generated content. This involves training employees to navigate the space of potential entities and iteratively evaluate them, as well as critically assess, refine, and integrate GenAI output with their own, in accordance with our hybrid model. Relatedly, when individual profiles lack balance, managers may consider forming ad-hoc cross-functional teams to ensure the adequate mix of skills that maximizes human-GenAI creativity.

#### 4.3. Limitations

Despite these contributions, our work has intrinsic limitations. First, a part of our arguments does not descend directly from mathematical deduction, but from standard theory building. While grounded in extant literature and logical reasoning, as in all conceptual papers, our arguments need empirical deepening and testing. In particular, as we describe the steps and skills needed to converge toward the optimal entity, we are largely silent on the difficulties that human agents may encounter in practice. We are also silent on the heuristics that human agents may use in lieu of the proposed steps and how they compare in efficiency and effectiveness to our idealized process. The theoretical nature of our endeavor also prevents us from considering contextual variables (e.g. characteristics of the organizational environment).

Another limitation regards the absence of the team level of analysis in the human component of the dyad. Beyond the (purely deductive) prescription on the choice of the optimal entity, our arguments on problem decomposition and re-integration do not delve into team-level considerations. In contrast, they focus on the comparative advantage of the human cognition in itself (independently of how “distributed” it may be in a team). While our theory can reasonably be extended to small teams as it is (as in other scholarly work on creativity, see e.g. [Amabile, 1988](#)), it leaves interesting questions open on how the various steps may be organized in a large team.

#### 4.4. Future research directions

From a theoretical perspective, future research could join our endeavor by adopting the team or even the organization as the focal level of analysis. Creativity is a multilevel construct ([Woodman et al., 1993](#)), and the impact of GenAI on it should be explored accordingly ([Grilli and Pedota, 2024](#)). Therefore, studies at the organizational and (large) team levels are as important as those at the individual (or small team) levels. Specifically, our contribution spawns a series of intriguing questions on the potential decoupling, modularization, and delegation of each step of the process, as well as the distribution of the skills required within and across teams, and the role of hierarchies, control mechanisms, and organizational enablers. We still know relatively little on the impact of AI on organizational creativity, and our perspective on human-GenAI interaction may provide a new lens for future theory-building efforts.

Empirically, experimental research could complement our model by measuring the gain in creative performance stemming from its implementation. Experiments would also be instrumental in understanding how relevant hard and soft skills are in the proposed creative process, and what is the impact of different contextual variables (e.g. innovation-friendly organizational environments, supportive leadership, and time constraints). As GenAI-driven team-level interactions are difficult to model theoretically, experiments on teams with varying skill distributions in different types of creative tasks may be needed for solid theory building at the team level of analysis.

We also encourage qualitative empirical research to explore the spontaneous unfolding of our model in practice. As our model depicts an ideal process, part of its usefulness stems from comparing it to how individuals actually implement it in real-world contexts, akin to comparing actual solutions to game-theory optimal ones in strategic interactions. Qualitative research may reveal difficulties, biases, and heuristics that constrain and guide implementation in different contexts.

This may not only provide a richer theoretical understanding, but also more specific practical prescriptions.

## Conclusion

This study highlights the transformative role of GenAI in creative processes, moving beyond traditional support functions to a more integrated, interactive dynamic with human cognition. Through a model centered around the innovative notion of simulated entity elicitation, we emphasize the evolving drivers of creativity in the age of AI. Rather than marginalizing the human role, GenAI necessitates a refined blend of analytical and contextual abilities. While our contribution offers a foundation for optimizing human-GenAI interaction, empirical research is needed to assess its practical implementation and the full extent of its organizational implications. Ultimately, we hope that our work may stimulate further inquiry into the creative synergy between human agents and GenAI, which we consider a defining challenge of the AI era.

## CRedit authorship contribution statement

**Mattia Pedota:** Writing – review & editing, Writing – original draft, Validation, Supervision, Conceptualization. **Francesco Cicala:** Writing – review & editing, Methodology, Formal analysis. **Alessio Basti:** Writing – review & editing, Methodology, Formal analysis, Conceptualization.

## Conflict of interest statement

The authors confirm that they have no conflict of interest to disclose.

## Acknowledgements

This work is part of the HumanTech Project, funded by the Italian Ministry of University and Research (MUR) for the period 2023–2027 under the “Departments of Excellence” programme (L. 232/2016).

## Data availability

No data was used for the research described in the article.

## References

- Acemoglu, D., Restrepo, P., 2018. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review* 108 (6), 1488–1542.
- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., McGrew, B., 2023. Gpt-4 technical report. *arXiv*, 2303.08774.
- Akter, S., Hossain, M.A., Sajib, S., Sultana, S., Rahman, M., Vrontis, D., McCarthy, G., 2023. A framework for AI-powered service innovation capability: Review and agenda for future research. *Technovation* 125, 102768.
- Amabile, T.M., Pratt, M.G., 2016. The dynamic componential model of creativity and innovation in organizations: Making progress, making meaning. *Research in Organizational Behavior* 36, 157–183.
- Amabile, T.M., 1983. The social psychology of creativity: A componential conceptualisation. *Journal of Personality and Social Psychology* 45 (2), 357.
- Amabile, T.M., 1988. A model of creativity and innovation in organizations. *Research in Organizational Behavior* 10 (1), 123–167.
- Amabile, T.M., 2020. Creativity, artificial intelligence, and a world of surprises. *Academy of Management Discoveries* 6 (3), 351–354.
- Anthony, C., Bechky, B.A., Fayard, A.L., 2023. “Collaborating” with AI: Taking a system view to explore the future of work. *Organization Science* 34 (5), 1672–1694.
- Balasubramanian, N., Ye, Y., Xu, M., 2022. Substituting human decision-making with machine learning: Implications for organizational learning. *Academy of Management Review* 47 (3), 448–465.
- Bardes, A., Garrido, Q., Ponce, J., Chen, X., Rabbat, M., LeCun, Y., Ballas, N., 2024. Revisiting feature prediction for learning visual representations from video. *Meta AI Blog*.
- Barron, F., Harrington, D.M., 1981. Creativity, intelligence, and personality. *Annual Review of Psychology* 32 (1), 439–447.
- Bender, E.M., Gebru, T., McMillan-Major, A., Shmitchell, S., 2021. On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency, FAccT '21*, 610–623.

- Bilgram, V., Laarmann, F., 2023. Accelerating innovation with GenAI: AI-augmented digital prototyping and innovation methods. *IEEE Engineering Management Review* 51 (2), 18–25.
- Bouschery, S.G., Blazevic, V., Piller, F.T., 2023. Augmenting human innovation teams with artificial intelligence: Exploring transformer-based language models. *Journal of Product Innovation Management* 40 (2), 139–153.
- Boussieux, L., Lane, J.N., Zhang, M., Jacimovic, V., Lakhani, K.R., 2024. The crowdless future? Generative AI and creative problem-solving. *Organization Science* 35 (5), 1589–1607.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Amodei, D., 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems* 33, 1877–1901.
- Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Zhang, Y., 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv* 2303, 12712.
- Cabantous, L., Gond, J.P., 2011. Rational decision making as performative praxis: Explaining rationality's Éternel Retour. *Organization Science* 22 (3), 573–586.
- Campbell, D.T., 1960. Blind variation and selective retentions in creative thought as in other knowledge processes. *Psychological Review* 67 (6), 380.
- Chiarello, F., Giordano, V., Spada, I., Barandoni, S., Fantoni, G., 2024. Future applications of generative large language models: A data-driven case study on ChatGPT. *Technovation* 133, 103002.
- Choudhary, V., Marchetti, A., Shrestha, Y.R., Puranam, P., 2023. Human-AI ensembles: When can they work? *Journal of Management*, 01492063231194968.
- Cohen, W.M., Levinthal, D.A., 1990. Absorptive capacity: A new perspective on learning and innovation. *Administrative Science Quarterly* 35 (1), 128–152.
- Cropley, A., 2006. In praise of convergent thinking. *Creativity Research Journal* 18 (3), 391–404.
- Csikszentmihalyi, M., 1996. *Flow and the psychology of discovery and invention*. Harper Collins, New York, NY.
- Deichmann, D., Jensen, M., 2018. I can do that alone... or not? How idea generators juggle between the pros and cons of teamwork. *Strategic Management Journal* 39 (2), 458–475.
- Dell'Acqua, F., McFowland III, E., Mollick, E.R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Lakhani, K.R., 2023. Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper* 24–2013.
- Doshi, A.R., Hauser, O.P., 2024. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances* 10 (28), eadn5290.
- Eppinger, S.D., Browning, T.R., 2012. *Design Structure Matrix Methods and Applications*. MIT Press, Cambridge, MA.
- Ethiraj, S.K., Levinthal, D., 2004. Modularity and innovation in complex systems. *Management Science* 50 (2), 159–173.
- Feuerriegel, S., Hartmann, J., Janiesch, C., Zschech, P., 2024. Generative AI. *Business & Information Systems Engineering* 66 (1), 111–126.
- Fleming, L., Sorenson, O., 2004. Science as a map in technological search. *Strategic Management Journal* 25 (8–9), 909–928.
- Gavetti, G., Greve, H.R., Levinthal, D.A., Ocasio, W., 2012. The behavioral theory of the firm: Assessment and prospects. *Academy of Management Annals* 6 (1), 1–40.
- Gelbart, M.A., Snoek, J., Adams, R.P., 2014. Bayesian optimization with unknown constraints. *arXiv* 1403, 5607.
- Girotra, K., Meinck, L., Terwiesch, C., Ulrich, K.T., 2023. Ideas are dimes a dozen: Large language models for idea generation in innovation. SSRN, 4526071.
- Girotra, K., Terwiesch, C., Ulrich, K.T., 2010. Idea generation and the quality of the best idea. *Management Science* 56 (4), 591–605.
- Grilli, L., Pedota, M., 2024. Creativity and artificial intelligence: A multilevel perspective. *Creativity and Innovation Management* 33 (2), 234–247.
- Guilford, J.P., 1984. Varieties of divergent production. *The Journal of Creative Behavior* 18 (1), 1–10.
- Hajikhani, A., Cole, C., 2024. A critical review of large language models: Sensitivity, bias, and the path toward specialized ai. *Quantitative Science Studies* 1–22.
- Halford, G.S., Wilson, W.H., Phillips, S., 1998. Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behavioral and Brain Sciences* 21 (6), 803–831.
- Harvey, S., 2014. Creative synthesis: Exploring the process of extraordinary group creativity. *Academy of Management Review* 39 (3), 324–343.
- Hendrycks, D., Burns, C., Kadavath, S., Arora, A., Basart, S., Tang, E., Steinhardt, J., 2021. Measuring mathematical problem solving with the math dataset. *Advances in Neural Information Processing Systems* 35.
- Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., Sifre, L., 2022. Training compute-optimal large language models. *arXiv* 2203, 15556.
- Huy, Q.N., 1999. Emotional capability, emotional intelligence, and radical change. *Academy of Management Review* 24 (2), 325–345.
- Jia, N., Luo, X., Fang, Z., Liao, C., 2024. When and how artificial intelligence augments employee creativity. *Academy of Management Journal* 67 (1), 5–32.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T.B., Chess, B., Child, R., Amodei, D., 2020. Scaling laws for neural language models, 08361. *arXiv*, 2001.
- Katila, R., Ahuja, G., 2002. Something old, something new: A longitudinal study of search behavior and new product introduction. *Academy of Management Journal* 45 (6), 1183–1194.
- Kemp, A., 2024. Competitive advantage through artificial intelligence: Toward a theory of situated AI. *Academy of Management Review* 49 (3), 618–635.
- Koivisto, M., Grassini, S., 2023. Best humans still outperform artificial intelligence in a creative divergent thinking task. *Scientific Reports* 13 (1), 13601.
- Kornish, L.J., Ulrich, K.T., 2011. Opportunity spaces in innovation: Empirical analysis of large samples of ideas. *Management Science* 57 (1), 107–128.
- Kornish, L.J., Ulrich, K.T., 2011. Opportunity spaces in innovation: Empirical analysis of large samples of ideas. *Management Science* 57 (1), 107–128.
- Kozbelt, A.K., Beghetto, R.A., Runco, M.A., 2010. Theories of Creativity. In: Kaufman, J. C., Sternberg, R.J. (Eds.), *The Cambridge Handbook of Creativity*. Cambridge University Press, New York, NY, pp. 20–47.
- Krakowski, S., Luger, J., Raisch, S., 2023. Artificial intelligence and the changing sources of competitive advantage. *Strategic Management Journal* 44 (6), 1425–1452.
- Lamb, L.C., D'Avila Garcez, A., Gori, M., Prates, M., Avelar, P., Vardi, M., 2020. Graph neural networks meet neural-symbolic computing: A survey and perspective. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence* 29, 4877–4884.
- LeCun, Y., 2022. A Path Towards Autonomous Machine Intelligence, 08239. *arXiv*, 2201.
- Levinthal, D.A., 1997. Adaptation on rugged landscapes. *Management Science* 43 (7), 934–950.
- Levinthal, D.A., March, J.G., 1993. The myopia of learning. *Strategic Management Journal* 14 (S2), 95–112.
- Li, Y., Zhang, K., Wang, J., Gao, X., 2021. A cognitive brain model for multimodal sentiment analysis based on attention neural networks. *Neurocomputing* 430, 159–173.
- Liedtka, J., 2015. Linking design thinking with innovation outcomes through cognitive bias reduction. *Journal of Product Innovation Management* 32 (6), 925–938.
- Lubart, T., 2005. How can computers be partners in the creative process: classification and commentary on the special issue. *International Journal of Human-Computer Studies* 63 (4–5), 365–369.
- Luo, Y., Yang, Z., Meng, F., Li, Y., Zhou, J., Zhang, Y., 2023. An empirical study of catastrophic forgetting in large language models during continual fine-tuning. *arXiv* 2308, 08747.
- Mannucci, P.V., Perry-Smith, J.E., 2022. “Who are you going to call?” Network activation in creative idea generation and elaboration. *Academy of Management Journal* 65 (4), 1192–1217.
- March, J.G., 1991. Exploration and exploitation in organizational learning. *Organization Science* 2 (1), 71–87.
- Mariani, M., Dwivedi, Y.K., 2024. Generative artificial intelligence in innovation management: A preview of future research developments. *Journal of Business Research* 175, 114542.
- McGuire, J., De Gremer, D., Van de Cruys, T., 2024. Establishing the importance of co-creation and self-efficacy in creative collaboration with artificial intelligence. *Scientific Reports* 14 (1), 18525.
- Nardo, C., 2023. The Waluigi Effect. Retrieved July 17th, 2024 from. <https://lesswrong.com/posts/D7PumeYTDPBTP3i7/toward-the-waluigi-effect-mega-post>.
- Nelson, R.R., Winter, S.G., 1973. Toward an evolutionary theory of economic capabilities. *The American Economic Review* 63 (2), 440–449.
- Nonaka, I., 1994. A dynamic theory of organizational knowledge creation. *Organization Science* 5 (1), 14–37.
- Nonaka, I., Takeuchi, H., 1995. *The Knowledge Creating Company: How Japanese Companies Create the Dynamics of Innovation*. Oxford University Press, New York.
- Nonaka, I., von Krogh, G., 2009. Tacit Knowledge and Knowledge Conversion: Controversy and Advancement in Organizational Knowledge Creation Theory. *Organization Science* 20 (3), 635–652.
- Nijkamp, E., Pang, B., Hayashi, H., Tu, L., Wang, H., Zhou, Y., Xiong, C., 2022. Codegen: An open large language model for code with multi-turn program synthesis. *arXiv* 2203, 13474.
- Perry-Smith, J.E., Mannucci, P.V., 2017. From creativity to innovation: The social network drivers of the four phases of the idea journey. *Academy of Management Review* 42 (1), 53–79.
- Rabinowitz, N., Perbet, F., Song, F., Zhang, C., Eslami, S.A., Botvinick, M., 2018. Machine theory of mind. In: *Proceedings of the 35th International Conference on Machine Learning*, 80. PMLR, pp. 4218–4227.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., 2019. *Language models are unsupervised multitask learners*. OpenAI Blog. <https://cdn.openai.com/better-language-models/language-models-are-unsupervised-multitask-learners.pdf>.
- Raisch, S., Fomina, K., 2025. Combining human and artificial intelligence: Hybrid problem-solving in organizations. *Academy of Management Review* 50 (2), 441–464.
- Raisch, S., Krakowski, S., 2021. Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review* 46 (1), 192–210.
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Sutskever, I., 2021. Zero-shot text-to-image generation. In: *Proceedings of the 38th International Conference on Machine Learning*, 139. PMLR, pp. 8821–8831.
- Richardson, J.T., Engle, R.W., Hasher, L., Logie, R.H., Stoltzfus, E.R., Zacks, R.T., 1996. Working memory and human cognition. Oxford University Press, New York.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B., 2022. High resolution image synthesis with latent diffusion models. In: *IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695.
- Runco, M.A., Jaeger, G.J., 2012. The standard definition of creativity. *Creativity Research Journal* 24 (1), 92–96.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., De Freitas, N., 2015. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE* 104 (1), 148–175.
- Shanahan, M., McDonnell, K., Reynolds, L., 2023. Role play with large language models. *Nature* 623 (7987), 493–498.
- Shneiderman, B., 2007. Creativity support tools: accelerating discovery and innovation. *Communications of the ACM* 50 (12), 20–32.
- Simon, H.A., 1955. A behavioral model of rational choice. *The Quarterly Journal of Economics* 69 (1), 99–118.

- Simonton, D.K., 1999. *Origins of genius: Darwinian perspectives on creativity*. Oxford University Press, New York.
- Simonton, D.K., 2010. Creative thought as blind-variation and selective-retention: Combinatorial models of exceptional creativity. *Physics of Life Reviews* 7 (2), 156–179.
- Simonton, D.K., 2015. On praising convergent thinking: Creativity as blind variation and selective retention. *Creativity Research Journal* 27 (3), 262–270.
- Simonton, D.K., 2023. The blind-variation and selective-retention theory of creativity: Recent developments and current status of BVSR. *Creativity Research Journal* 35 (3), 304–323.
- Singh, B., 1986. Role of personality versus biographical factors in creativity. *Psychological Studies* 31, 90–92.
- Singh, K., Chatterjee, S., Mariani, M., 2024. Applications of GenAI and future organizational performance: The mediating role of explorative and exploitative innovation and the moderating role of ethical dilemmas and environmental dynamism. *Technovation* 133, 103021.
- Srivastava, A., Rastogi, A., Rao, A., Shoeb, A.A.M., Abid, A., Fisch, A., Wang, G., 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv 2206.04615*.
- Teece, D.J., Pisano, G., Shuen, A., 1997. Dynamic capabilities and strategic management. *Strategic Management Journal* 18 (7), 509–533.
- Teece, D.J., 2007. Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal* 28 (13), 1319–1350.
- Todorova, G., Durisin, B., 2007. Absorptive capacity: Valuing a reconceptualization. *Academy of Management Review* 32 (3), 774–786.
- Townsend, D.M., Hunt, R.A., Rady, J., Manocha, P., Jin, J.H., 2024. Are the futures computable? Knightian uncertainty and Artificial Intelligence. *Academy of Management Review* forthcoming.
- Tripsas, M., Gavetti, G., 2000. Cognition, capabilities and inertia: evidence from digital imaging. *Strategic Management Journal* 21 (10/11), 1147–1161.
- Tripsas, M., 2009. Technology, identity, and inertia through the lens of “The Digital Photography Company”. *Organization Science* 20 (2), 441–460.
- Tsao, J.Y., Ting, C.L., Johnson, C.M., 2019. Creative outcome as implausible utility. *Review of General Psychology* 23 (3), 279–292.
- Tversky, A., Kahneman, D., 1974. Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *Science* 185 (4157), 1124–1131.
- Tyre, M.J., Von Hippel, E., 1997. The situated nature of adaptive learning in organizations. *Organization Science* 8 (1), 71–83.
- Ulrich, K.T., Eppinger, S.D., 1995. *Product design and development*. Mc Graw-Hill, New York.
- Verganti, R., Dell’Era, C., Swan, K.S., 2021. Design thinking: Critical analysis and future evolution. *Journal of Product Innovation Management* 38 (6), 603–622.
- Von Oswald, J., Niklasson, E., Randazzo, E., Sacramento, J., Mordvintsev, A., Zhmoginov, A., Vladymyrov, M., 2023. Transformers learn in-context by gradient descent. *Proceedings of the 40th International Conference on Machine Learning 202*, 35151–35174. PMLR.
- Wang, Y., Zhao, J., Ones, D.S., He, L., Xu, X., 2025. Evaluating the ability of large language models to emulate personality. *Scientific Reports* 15 (1), 519.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Zhou, D., 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *Advances in Neural Information Processing Systems* 35, 24824–24837.
- Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Fedus, W., 2022. Emergent abilities of large language models. *arXiv 2206.07682*.
- Woodman, R.W., Sawyer, J.E., Griffin, R.W., 1993. Toward a theory of organizational creativity. *Academy of Management Review* 18 (2), 293–321.
- Wright, P., 2002. Marketplace metacognition and social intelligence. *Journal of Consumer Research* 28 (4), 677–682.
- Xue, M., Cao, X., Feng, X., Gu, B., Zhang, Y., 2022. Is college education less necessary with AI? Evidence from firm-level labor structure changes. *Journal of Management Information Systems* 39 (3), 865–905.
- Yao, S., Yu, D., Zhao, J., Shafraan, I., Griffiths, T., Cao, Y., Narasimhan, K., 2023. Tree of Thoughts: Deliberate Problem Solving with Large Language Models. *Advances in Neural Information Processing Systems* 36, 11809–11822.
- Zhao, W.X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Wen, J.R., 2023. A survey of large language models. *arXiv 2303.18223*.
- Zirar, A., Ali, S.I., Islam, N., 2023. Worker and workplace Artificial Intelligence (AI) coexistence: Emerging themes and research agenda. *Technovation* 124, 102747.