Psychologically-inspired generative AI videos for supporting creativity

Shabnam Hakimi, Monica Van, Matthew K. Hong, Kalani Murakami,
Pablo E. Paredes, and **Matthew Klenk**
Human-Centered AI, Toyota Research Institute, Los Altos, CA USA
{shabnam.hakimi, monica.van, matthew.hong, kalani.murakami, pablo.paredes, matthew.klenk}@tri.global

Abstract

Generative AI has become a powerful tool for supporting creativity across a breadth of disciplines. Nonetheless, applications are generally domain- and/or task-specific, providing targeted support for creative tasks like design ideation. Here, we present a domain-general approach to using generative AI to amplify human creativity by targeting its underlying psychological processes at high temporal resolution. We describe a generative AI workflow for creating two-minute video interventions designed to increase open-minded, flexible thinking, a central domain-general component of creative processes. We also present human participants' ratings of and reactions to these videos, detailing implications for future work.

Introduction

> If the mind creates a world that is limited, where we think we don't have enough worthwhile ideas or material, we will not see the inspiration the universe is providing.
>
> – Rick Rubin, *The Creative Act: A Way of Being* (p. 202)

The mind is the filter through which humans make sense of and shape the world. The success of a human's creative process, whether measured by the quality of the experience or the output, depends on the mind's capability to represent ideas in new, surprising, and valuable ways (Boden 2004). Across the many theoretical accounts of creativity, including Boden's (2004) influential work, a common set of psychological substrates scaffold creative processes, with **psychological flexibility** emerging as the critical connective tissue that enables what we consider *creativity* (Benedek et al. 2014). Psychological flexibility describes a set of dynamic processes that support openness and adaptation to changing situational demands; shifting of mindsets or behavioral repertoires; and balancing competing desires and needs, while considering how they relate to values (Kashdan and Rottenberg 2010). Further, psychological flexibility supports the executive functions considered necessary for creativity (Nijstad et al. 2010; Kashdan and Rottenberg 2010; Hayes et al. 2020; Zabelina, Friedman, and Andrews-Hanna 2019) by modulating an array of self-regulatory processes that enable higher-order thought and behavior(Zabelina and Robinson 2010).

Psychological flexibility's centrality in the creative process highlights its potential as a target for intervention. Indeed, prior research has leveraged a variety of psychological interventions for enhancing psychological flexibility, with varying degrees of success (e.g., (Kashdan and Rottenberg 2010; Uddin 2021)). Both between- and within- individual variability at the trait and state levels may be one cause for mixed results, potentially limiting the efficacy of existing approaches. Further, prior psychological creativity support has been largely domain- and task- specific, with interventions focused on narrow tasks such as increasing flexibility when generating semantic concepts during ideation (e.g., in a Remote Associates Test (Wu et al. 2020)). Advances in artificial intelligence (AI), particularly in generative AI models, suggest new opportunities for intervening on psychological flexibility to amplify human creativity because of their immense capacity for personalization and generating variability at scale.

Can we leverage both psychological principles and generative AI to aid the mind in creating Rubin's internal world with fewer limitations? Here, we present a human-centered method for automated generation of video interventions for creativity designed to increase psychological flexibility across domains. We describe the system and its approach to targeting psychological flexibility at the fine temporal resolution of the process level with reduced dependence on the specific task or domain. We also describe the results of a study investigating human viewers' affective responses to the generative AI videos. Then, we discuss the implications of the system for impacting creative processes, reviewing opportunities for further research.

Background

In psychology, creativity is often conceptualized as a durable personality characteristic, or trait, that correlates strongly with intelligence and openness (Zhang, Wang, and Zhang 2020; Benedek et al. 2014; McCrae 1987). Creative individuals also demonstrate greater psychological flexibility (Kenett et al. 2018; Smith and Ward 2012). At the trait level, creativity and *creative mindset*, which reflects metacognitive knowledge about creativity and beliefs about its malleability with time and effort, are associated with higher creative performance and creative self-efficacy(Royston and Reiter-Palmon 2019; O'Connor, Nemeth, and Akutsu 2013; Karwowski 2014). Despite data suggesting that some traits, including openness, may be malleable over long periods of time (Roberts et al. 2017; Dolbier, Dieffenbach, and Lieberman 2025), focusing on traits not only reinforces the potentially detrimental belief that creativity is fixed (Karwowski 2014), but may also lead to missed opportunities for targeted support in the immediate to short-term.

Creative State Interventions

Creativity can also be considered a transient state, although this construction does not imply a specific temporal resolution. Recent work suggests that the creative mindset may be affected in a more state-like manner through a perspective-taking instructional prime, leading to increased originality of ideas (Chou and Tversky 2020). Moving to an even shorter timescale (minutes to hours), prior work has shown that expectancy violations increase creative cognition by increasing flexibility (Ritter et al. 2012), with positive affect mediating this relationship (Rowe, Hirsh, and Anderson 2007; Dolbier, Dieffenbach, and Lieberman 2025). Related work has also demonstrated the importance of affect for state creativity interventions, showing that awe, which is characterized in part by complex affective experience, can enhance creative cognition and everyday creativity (Chirico et al. 2018; Zhang et al. 2024). Further evidence suggests that both pharmacological and even digital psychedelic experiences, which have long been considered powerful modulators of creative flexibility, also operate at a similar timescale (Gandy et al. 2022; Greco et al. 2025). Overall, such work complements a history of curricular creativity training interventions targeting not only flexibility, but also fluency and originality over a period of weeks to months (Bi et al. 2020). Interventions designed to impact even shorter timescales (seconds to minutes) have also shown some success in enhancing performance in creative contexts (Albulescu et al. 2022). Despite some evidence of efficacy, these approaches lack specificity both at the target level (i.e., the process being affected is unclear) and at the outcome level (i.e., the relationship between creativity and 'performance' is unclear).

Despite limited evidence for the effectiveness of creative state interventions at shorter timescales, there are opportunities for impact at a resolution closer to that of creative cognition and affect (milliseconds to seconds). The temporal dynamics of expectancy violations and their connection to flexibility and open-mindedness provide a useful lens through which to interrogate the malleability of creative processes at a finer temporal resolution. Kumar and colleagues (2023) examined how Bayesian surprise, a measure of how new observations affect the difference between prior and posterior beliefs, impacted memory encoding of text narratives. By computing a measure of the degree to which each advancing word shifted the predictive distribution of a large language model (LLM), they were able to explain how transient prediction errors affected the evolving mental representation of the narratives, particularly how the contents were segmented into discrete events (Kumar et al. 2023; Baldi and Itti 2010). Their results support dynamic accounts of perception, learning, and memory that suggest that people maintain a multimodal working model of an event that is iteratively checked against new information; the model is stable and is only updated when prediction errors increase, leading to event boundaries (Zacks, Speer, and Reynolds 2009). This view is in line with predictive processing accounts (Clark 2013; Friston et al. 2021) as well as Wiggins's Information Dynamics of Thinking model (Wiggins 2020). Together, these converging lines of evidence suggest that cognitive states may be 1) operationalized using information-theoretic or energetic measures and 2) can be targeted on a moment-to-moment basis to support complex states like psychological flexibility and creativity. Analysis at this level allows for interventional approaches that are less domain- and/or task-specific, focusing instead on the dynamics of the cognitive system.

Creativity Support and Generative AI

Although they are not typically framed as interventions, systems and tools designed to support human creativity likely affect creativity at a state level, since they aim to bolster creative processes *in situ*. Historically, such support happened outside the digital realm, with individuals employing tools for creative cognition (e.g., analogical thinking, mind-mapping) or shifting context to inspire and clear the mind (e.g., taking a walk, visiting a museum). As creative activities have become more digital, creativity support tools have followed suit, growing dramatically in number while also becoming more fully integrated in creative processes (Frich et al. 2019). Many of these approaches are grounded in the field of Human-Computer Interaction. Across a wide array of creative processes and domains, such tools are more likely to be useful when well-aligned with underlying psychological processes (Amitani and Hori 2002). A growing body of work has further demonstrated how psychological principles may be used to guide the development of creative systems (e.g., (DiPaola, Gabora, and McCaig 2018; Tan, Antony, and Kong 2020; Thölke et al. 2024)).

Advances in generative AI have shifted the landscape of creativity support. While much of the discourse focuses on critical topics such as characterizing machine creativity (e.g., (Peeperkorn, Brown, and Jordanous 2023)) and articulating roles in co-creation (e.g., (Chung, He, and Adar 2021)), generative AI's capabilities for amplifying human creative processes should also be considered. Current research has maintained focus on how generative AI may support the psychological factors needed to complete specific creative tasks within a particular domain (e.g., visual design, creative writing). Artists' adoption of text-to-image generative AI was associated with higher productivity, novelty, and community value (Zhou and Lee 2024). A similar effect was reported using text-to-text generation for story composition, showing that using generative AI for ideation can increase individual creativity, although the overall novelty of stories is reduced (Doshi and Hauser 2024). Interestingly, effects across these and related studies are moderated by prior use of generative AI tools, with gains seen primarily in new users; experienced AI users gained little or no benefit (Fu et al. 2024; Ashkinaze et al. 2024). These findings point to the many nuances associated with the use of generative AI in creativity support, especially the risks of null or negative effects for individuals with greater exposure to and experience with AI.

Considering the large body of work on creativity support with and without generative AI, research has largely focused on ways to support specific domains and/or tasks. As with creative state interventions that emerged from the field of psychology, creativity support tools have not yet explicitly addressed the momentary dynamics of the cognitive and affective processes underlying creativity. Extending on the work of (Kumar et al. 2023), generative AI may provide a new approach for re-aligning interventions with their underlying psychological processes at the appropriate temporal resolution. Further, generative AI agents can act as co-creators of interventions (e.g., (Demszky et al. 2023)), vastly increasing the diversity of interventions while avoiding exposure and expertise effects of direct interaction with the human recipients of creativity interventions.

Creating Psychologically-Inspired Generative AI Video Interventions

We therefore aimed to develop a novel approach to enhancing psychological flexibility that could address gaps evident in currently available creative state interventions and creativity support tools. The system desiderata were as follows:

- Ability to generate varied interventions with minimal instruction so as to retain desirable generative AI capabilities (e.g., novelty)
- User control over generative model representation of principles used to intervene on psychological flexibility and their temporal characteristics (e.g., level of Bayesian surprise over time)
- Capacity to generate a high degree of, or even infinite, variability in the narratives and visual features of the resulting videos in order to facilitate personalization and reduce habituation

The video generation process included two stages: text-to-text narrative expansion and text-to-image video generation. The general approach is shown in Figure 1a; an example pair of videos generated using this approach is shown below in Figure 1b.

Text-to-Text Narrative Expansion

In the first stage of the process, a set of user-generated starting prompts is expanded into narratives, or generative text prompts that will later become videos. We developed a set of initial prompts (`STARTING_PROMPT`) to seed the narratives that would eventually become videos. 13 prompts represented scenes that a viewer might see in daily life:

- *A boot hangs on a telephone line*
- *A door opens into a room*
- *A lamp lights the street below*
- *A mug waits by a coffee machine*
- *A novel sits on a bedside table*
- *A painting hangs on a museum wall*
- *A paperclip sits atop a drafting table*
- *A phone rings in another room*
- *A plane flies across a blue sky*
- *A raindrop falls from the sky*
- *A seedling grows from a sidewalk crack*
- *A speck of dust lands on a drawing pad*
- *A stereo plays music*

Each `STARTING_PROMPT` was passed to GPT-4 (OpenAI 2024) using the syntax below with a temperature of 0.65 and a maximum of 1024 tokens. The goal was to generate sets of prompts that would establish a prior on the context of the narrative and then progressively challenge that belief with each additional narrative step. The difference between the prior from the previous step and the beliefs at the current step is also referred to as Bayesian surprise (Itti and Baldi 2009), which was operationalized using semantic distance (similar to (Kumar et al. 2023)). We expected that more surprise would lead to more updating, and thus a higher likelihood of inducing a psychological state conducive to openness and flexibility; induction of this state would be mediated, in part, by more diverse affect. Further, to generate prompts (and then videos) that would allow us to test how the degree of surprisal impacted the viewer's experience, we instructed the LLM to generate two sets of prompts for each starting prompt, one spanning a relatively *near* and the other a very *far* semantic distance. The semantic distance between each pair of generative text prompts was computed *post hoc* using cosine similarity over the sentence embedding space[1].

The system prompt was designed to provide the LLM sufficient instructions to reliably generate sets of additional prompts that represented a narrative continuation of the `STARTING_PROMPT` given the other parameters:

- `STARTING_PROMPT`
- `PROMPT_REFRESH`, the set of prompts $[p_1, ..., p_n]$ corresponding to frames to be generated, $[f_{0,0}, ..., f_{n,0}]$
- `END_POINT`, the desired video duration (s)
- `NUM_STEPS`, the desired frame rate, x (in frames per second, fps)

As noted, the arguments passed to the LLM correspond to the variables denoted in Figure 1a.

Text-to-Image Video Generation

For each starting prompt and its associated sets of generated prompts, we created a two-minute, 12-fps video using Deforum[2], which uses Stable Diffusion v1.5 (Rombach et al. 2022), and a custom version of the AUTOMATIC1111 web interface[3]. Briefly, each of the prompts p_n was used to generate a static image, $f_{n,0}$. In order to create the illusion of movement in time, 50 interpolated images ($[f_{i,0}, ..., f_{i,j}]$) were generated between each pair of frames, $f_{n,0}$ and $f_{n+1,0}$.

[1]https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2
[2]https://deforum.art/
[3]https://github.com/AUTOMATIC1111/stable-diffusion-webui/

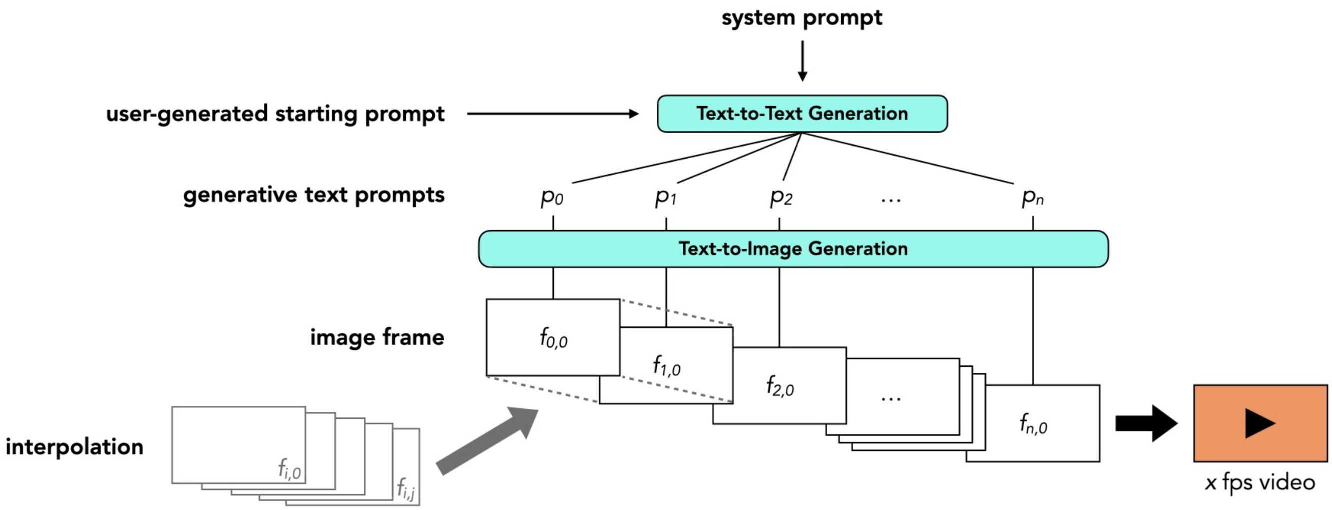

(a) System for automated generation of videos designed to enhance psychological flexibility. To employ the psychologically-inspired generative AI method, the user prompts the system with a set of specific parameters, including a starting prompt that will seed the expansion of a narrative through a text-to-text generative model (i.e., LLM), resulting in a set of generative text prompts. The generative text prompts are then passed to a text-to-image model with parameters for visual features like style. Generated image frames are then interpolated together, leading to a video with features optimized to target the dynamic features of the underlying creative process, psychological flexibility.

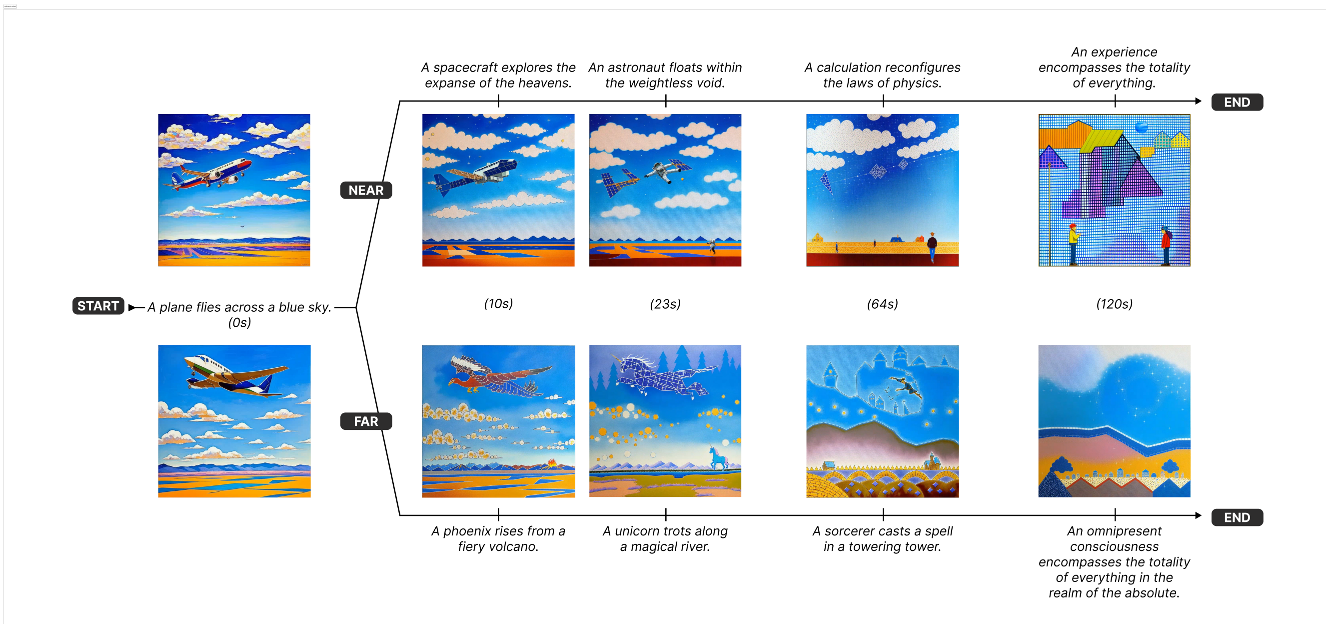

(b) Sample pair of generative videos created using system. Beginning with the user-generated starting prompt, "A plane flies across a blue sky," the system generated two sets of narrative prompts that differ in the degree of Bayesian surprise between sequential prompts. Bayesian surprise was indexed using semantic distance, leading to two set of prompts: one covering a relatively *near* and one covering a relatively *far* semantic distance. All other video generation parameters were the same.

Figure 1: Psychologically-inspired video generation.

In order to create videos with dynamic visual features conducive to enhancing psychological flexibility, we adjusted a set of parameters: diffusion cadence = 6, angle = "0: (0)", zoom = "0:(1.00)", x translation: "0: (0)", y translation: "0: (0)", z translation: "0: (0)", anti-blur kernel_schedule: "0: (5)", anti-blur sigma_schedule": "0: (4.0)", anti-blur amount_schedule: "0: (0.2)", and anti-blur threshold_schedule": "0: (0.0)". Video frame size was 768 x 768 pixels. The seed was randomized for each generation instance and the remaining parameters took default val-

```
Given the [STARTING_PROMPT], consider the
features of the world (including its
observable properties, possible actions, and
 overall context) in which this statement
exists. This is the prior belief. Now,
consider the [END_POINT] as the maximum
value for surprise. We will define surprise
as the degree to which a new observation
challenges or changes this prior belief,
resulting in an updated belief about the
world in which the statement exists.
Generate two sets of prompts that represent
[PROMPT_REFRESH] by incrementally updating
the [STARTING_PROMPT]. Each set should have
a total of [NUM_STEPS] prompts.

The two sets of prompts should differ in how
 surprising the final prompt is relative to
the [STARTING_PROMPT]. Use semantic distance
, or the degree of similarity between
sequential prompts, as a metric for surprise
. The size of the semantic distance between
each pair of prompts should be consistent.
The final prompt should be maximally
different from the [STARTING_PROMPT]. The
two sets of prompts should differ in how
dissimilar the final prompt is from the
original one. One should have a final prompt
 that is far in semantic similarity while
the other should have a final prompt that is
 very far in semantic similarity.
```

Figure 2: System prompt to initiate generative process. The system prompt instructs the generative AI model in how to generate pairs of prompt narratives differing in the progression of Bayesian surprise. Surprise is operationalized using semantic similarity.

ues. Positive prompts were used to constrain the visual style while negative prompts were used to minimize the likelihood of text, copyrighted material, or inappropriate content. After videos were generated, a title card was added to each video to provide basic narrative context to human viewers (e.g., ”A flying plane”).

Evaluating Human Responses to Generative AI Videos

Because prior work has highlighted the importance of rich, complex affective experiences for effective creativity interventions (Rowe, Hirsh, and Anderson 2007; Dolbier, Dieffenbach, and Lieberman 2025; Chirico et al. 2018), we measured human viewers' affective responses to viewing the generative AI videos. We expected that the videos created using our system would be associated with a wide range of affective responses, reflecting not only the complex affective experiences of the viewers but also individual differences in which videos were most compelling to each rater. We further expected to see differences in affect by semantic distance, with far videos leading to more varied affect than near videos.

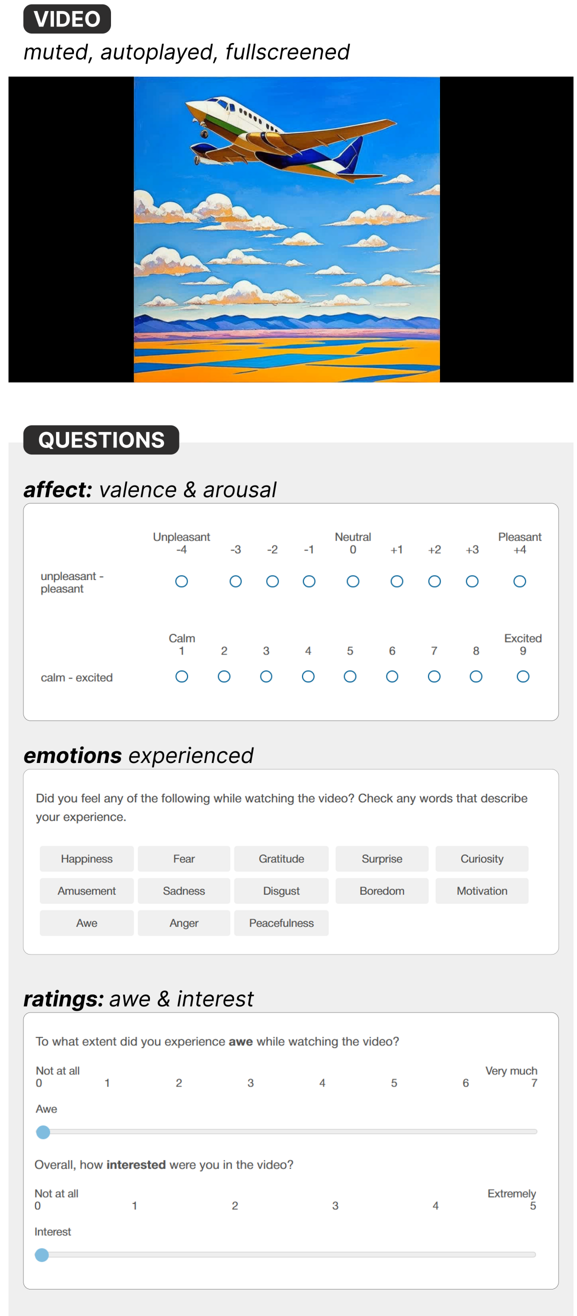

Figure 3: Video Evaluation Task. For each video, participants rated affect, awe, and level of interest; they also selected emotions they experienced while watching the video.

Participants

A total of 70 U.S.-based participants (33 females) evaluated the videos. Participant age ranged from 21 to 74 (median

= 36) and a range of ethnicities were represented (70.0% White, 12.9% Black, 10.0% Asian, 4.29% Mixed, 2.86% Other). Participants were recruited through Prolific[4] and completed the task using Qualtrics[5]. In addition to completing the video evaluation task described below, participants completed questionnaires to assess personality traits and demographics. Participants were compensated a flat rate of $8.50 for their time (mean = 25.5 min., std. = 12.3 min.). All participants provided informed consent and the study was approved by the WCG Institutional Review Board.

Video Evaluation Task Design

Each of 26 videos (one *near* and one *far* semantic distance for each starting prompt) was presented to 10-12 participants. Each participant viewed and rated videos generated from four distinct starting prompts; participants saw two *near* and two *far* videos. Video presentation order was randomized.

On each video evaluation trial, participants were presented with a two-minute, muted video in full screen. Participants could not proceed unless the complete video was viewed in fullscreen. Following the video, participants used anchored, Likert-style scales to rate state affect (valence: *Calm* = 1, *Excited* = 9; arousal: *Unpleasant* = -4, *Neutral* = 0, *Pleasant* = +4). They then reported any emotions they felt while watching the video by checking options from a randomly-ordered list (*Amusement*, *Anger*, *Awe*, *Boredom*, *Curiosity*, *Disgust*, *Fear*, *Gratitude*, *Happiness*, *Peacefulness*, *Motivation*, *Sadness*, *Surprise*). Finally, they rated their level of awe (*Not at all* = 0 to *Very much* = 7) and interest (*Not at all* = 0 to *Extremely* = 5) in response to the video using anchored scales. The design of the ratings questions was adapted from prior work assessing affective responses to awe-inducing videos used in interventions for creativity and pro-sociality (Chirico et al. 2018; Piff et al. 2015). A single trial is depicted in Figure 3.

Analytical Approach

Participant responses were preprocessed using `pandas` v2.2.3[6] in Python v3.1.2[7]. Statistical modeling was done with `statsmodels` v0.14.4[8]. To test the effect of video identity and level of semantic distance on video ratings, we fit linear regressions for each affective rating. To test whether prompt generation achieved the desired level of control over the level of Bayesian surprise between sequential prompts, we computed the absolute value of the difference between each pair of prompts and entered it into an ANOVA. Data visualizations were created using `seaborn` v0.13.2[9].

[4]https://www.prolific.com/
[5]https://www.qualtrics.com/
[6]https://pandas.pydata.org/
[7]https://www.python.org/downloads/release/python-312/
[8]https://www.statsmodels.org/stable/index.html
[9]https://seaborn.pydata.org/

Results

When selecting from 13 available emotions to describe how they felt while watching each video, the modal number of emotions selected was 3 (min = 1, max = 7). Across all videos, *curiosity* was selected most frequently, occurring in 22.4% of responses; *anger* was selected least frequently, appearing in only 0.51% of responses. Overall, the distribution of endorsed emotions was statistically equivalent for near and far semantic distances (χ^2(12, N=764) = 5.18, p=0.95). These data suggest that while semantic distance did not impact the incidence of discrete emotional states, viewing of generative AI videos elicited complex and sometimes conflicting emotions in participants.

Analysis of ratings data reveals variability in line with a dimensional account of emotion and also reflects a wide range of affective experience in responses to the videos. Figure 4 shows how ratings varied by both video and level of semantic distance for valence, arousal, awe, and interest. These data did not, however, fully confirm the hypothesis that affect would be modulated by the level of semantic distance. Video and level of semantic distance explained significant variance in valence ratings (F(13,265) = 2.259, p=0.008), but not others, although trend-level effects were observed for awe (p=0.054) and interest (p=0.081). While each rating showed a marginal effect in the expected direction—that mean ratings would be greater for far compared to near prompts—there was high variance across both videos and participants. Further investigation examined whether the *near* and *far* sets of generative prompts indeed differed in the degree of Bayesian surprise between sequential prompts, since a lack of significant difference might have impacted the analyses of affective ratings. While there was a marginal effect of video, the mean absolute distance between prompts did not significantly differ by semantic distance ($F(1, 24)$ = 0.002, p=0.967).

Discussion

In designing a system for the creation of generative AI video interventions for creativity, we prioritized the ability to target underlying processes associated with psychological flexibility. These core components of creative cognition evolve dynamically over time, with meaningful features emerging at timescales finer than those typically targeted by creativity support (i.e., sub-seconds to seconds rather than minutes or longer). Here, we aimed to bring the dynamics of the intervention more in line with those of the underlying process of interest, *psychological flexibility*. We used generative AI to manipulate, on the order of seconds, the dynamics of surprise in video narratives, thus inducing the viewer to update their model of the world. Through this perturbation of mental dynamics, we aimed to encourage flexibility and openmindedness that could support creativity without regard to specific task or domain.

We demonstrate initial evidence for the success of this approach to intervention generation, demonstrating the effect of the generative AI video interventions on the affective state of human viewers. Because prior research on creative state interventions has noted the importance of affect—especially

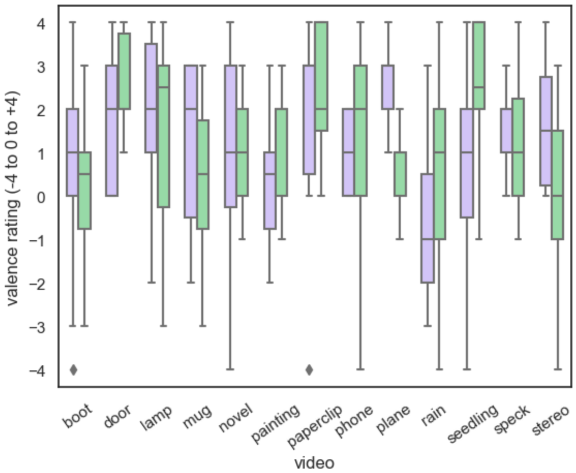

(a) Valence ratings by video.

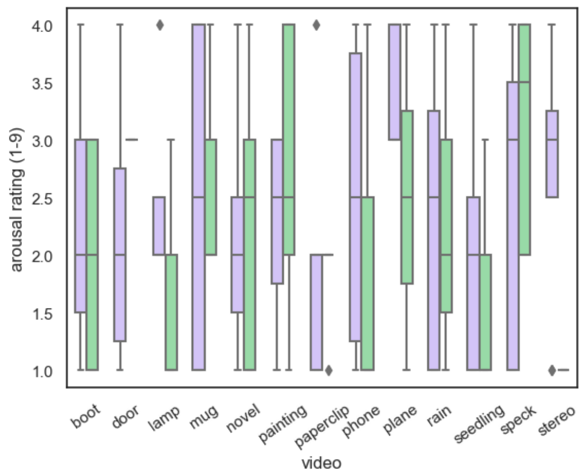

(b) Arousal ratings by video.

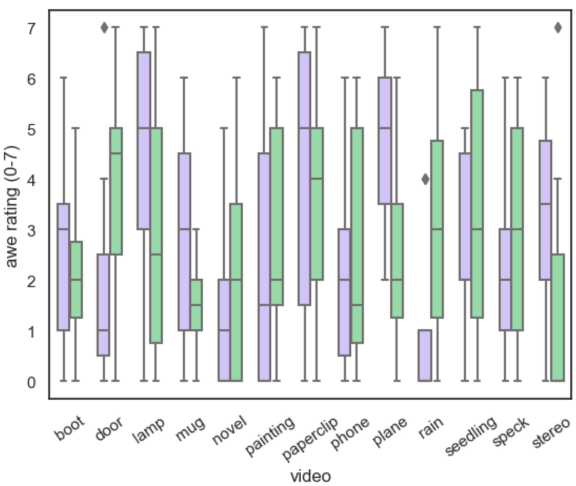

(c) Awe ratings by video.

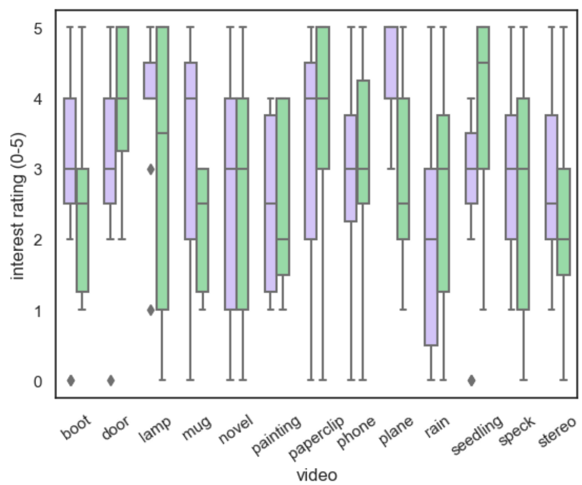

(d) Interest ratings by video.

Figure 4: Generative AI video ratings by user-generated starting prompt. For valence (a), arousal (b), awe (c), and interest (d), mean ratings and interquartile range are depicted for each video. *Near* semantic distance is in green while *far* semantic distance is in purple.

rich, complex affective experiences, like those associated with awe—we focused first on testing whether videos created using our system would be associated with viewing. On the whole, the results of our video evaluation study show increases in affect consistent with that previously seen with awe inductions (e.g., (Piff et al. 2015)), suggesting that the videos have characteristics sufficient for impacting affective states relevant to creativity. The efficacy of these generative AI videos as interventions for creativity remains to be tested directly.

Our approach to creating generative AI video interventions diverges from conventional human-AI co-creative applications of AI-driven storytelling, which increasingly leverage end-to-end models that map textual descriptions directly to video sequences. Large vision models like Runway's Gen-2[10] system and OpenAI's Sora[11] have demonstrated remarkable capability for high-quality video generation (Liu et al. 2024). These models often employ advanced prompt optimization and preference alignment strategies to produce aesthetically enhanced imagery, with the goal of mimicking the data distribution of videos that are available in the public domain.

Our priority was to develop a system for video intervention generation that would provide users with a high level of control, flexibility, interpretability, and reproducibility in all aspects of the generative process. This modular approach enabled us to exploit the latent capabilities of the different generative models in the least restrictive manner. Take, for

[10] `https://runwayml.com/research/gen-2`

[11] `https://openai.com/sora/`

instance, the text-to-text generative pipeline we used to generate the scene descriptions. While strategies exist for refining text-to-image prompts for computational creativity applications (Ibarrola and Grace 2023), we wanted to provide the minimum necessary information to preserve the generative diversity of the LLM, avoiding over-specification while still being able to guide the model toward coherence.

To generate the individual scene descriptions, we leveraged OpenAI's text-to-text generative model for its zero-shot reasoning capabilities (Espejel et al. 2023), semantic processing (Le Mens et al. 2023), and demonstrated aptitude for a range of affective computing tasks (Amin et al. 2024). The generated scene descriptions then informed a text-to-image model (i.e., Stable Diffusion), which created image frames, followed by a diffusion-based interpolation model that synthesized smooth transitions between the image frames. It is important to note that the visual feel of generative videos can significantly impact the viewer's experience. While our framework enables fine-grained adjustments to key parameters like visual style, frame rate, and motion dynamics, we developed heuristics to reduce the overwhelming parameter space for features.

Importantly, many of these heuristics were informed by an explicit choice to eschew realism in the generative AI videos. By selecting starting prompts that reflected likely scenes from daily life and systematically increasing Bayesian surprise over narrative progression, we aimed to control the properties of the distribution of information presented to the viewer. We expected that the aesthetics and production quality could risk this control, especially if an expectation of realism was set and then subsequently violated. Therefore, we chose parameters that would produce videos of acceptable production value that were *conceptually* grounded in viewers' experiences, but unmoored from high expectations about video quality. The unique visuals of the resulting generative AI videos helped establish an event boundary (Zacks and Swallow 2007) that could boost the potential efficacy of the intervention *without* also inducing the strong negative affect associated with narrow memory encoding (Gable and Harmon-Jones 2010). Narrow rather than broad encoding would inhibit the desired outcome of openness and flexibility.

Considerations and Future Directions

While we were able to generate interesting, narratively diverse videos with a range of visual styles, we did not achieve the desired level of control over the LLM's search over the semantic space, reducing the precision of the Bayesian surprise manipulation. The effect of the degree of surprise between sequential prompts was not as expected, perhaps because text-to-text narrative expansion did not reliably result in a consistent semantic distance between prompts where the distance was, on average, smaller for *near* relative to *far*. Future research should explore ways to increase control without overly specifying the system prompt.

We observed a great deal of variability in individuals' affective responses to the generative AI videos. Although such variability was expected and even typical for larger studies of affect in response to viewing videos (e.g., (Samide, Cooper, and Ritchey 2020)), it can limit inference on which videos induce affective states conducive to the desired effects on psychological flexibility. The immense size of the parameter space suggests a need to develop strategies that will allow future work to harness the options for features like narrative topic, aesthetic style, and frame rate to match the preferences of the viewer. Maintaining the appropriate level of meaningful variability is critical for the success of future interventions, since personal salience of stimuli can boost learning and memory (Peters 2006) as well as the efficacy of behavior change interventions (Alslaity, Chan, and Orji 2023). Increasing the level of immersiveness may also facilitate stronger resonance between a video and a human viewer. For example, prior work has shown that immersion can enhance the effect of awe (Chirico et al. 2018) as well as hallucinatory generative AI interventions (Greco et al. 2025) on creative thinking. Increasing immersiveness by including additional sensory modalities or through the use of techniques like virtual reality may allow for a higher likelihood of intervention success in future trials.

Finally, while our explicit focus has been on the potential for enhancing human creativity, our approach may also have implications for machine creativity. Pezzulo and colleagues (2023) argue that shifting generative AI models from passive to active inference may lead to systems more capable of human-like creativity (Pezzulo et al. 2023). Our earlier motivating example describing the formation of event boundaries as a result of active inference on the world (Zacks, Speer, and Reynolds 2009) is just the sort of process that may enable such advances. Future research may find value in exploring the extent to which purposefully shifting a generative AI agent's representation of the world may increase its flexibility.

Conclusion

Advances in generative AI suggest opportunities for new automated, scalable approaches to amplifying human creativity that can provide support across tasks and domains. We present a novel method for creating generative AI video interventions designed to enhance psychological flexibility, a central component of human creativity. We provide an initial demonstration of the system's capabilities and present opportunities for future research, including the development of personalized interventions for creativity.

Author Contributions

S.H. and M.K.H. conceptualized the research. S.H., M.V., M.K.H., and K.M. developed the video generation procedure, including associated software. S.H. and M.V. designed and conducted the human subjects study; S.H. analyzed and interpreted data. S.H., M.V., M.K.H., P.E.P, and M.K. drafted the manuscript.

References

Albulescu, P.; Macsinga, I.; Rusu, A.; Sulea, C.; Bodnaru, A.; and Tulbure, B. T. 2022. "Give me a break!" A systematic review and meta-analysis on the efficacy of microbreaks for increasing well-being and performance. *PLOS ONE* 17(8):1–27.

Alslaity, A.; Chan, G.; and Orji, R. 2023. A panoramic view of personalization based on individual differences in persuasive and behavior change interventions. *Frontiers in Artificial Intelligence* Volume 6 - 2023.

Amin, M. M.; Mao, R.; Cambria, E.; and Schuller, B. W. 2024. A wide evaluation of chatgpt on affective computing tasks. *IEEE Transactions on Affective Computing*.

Amitani, S., and Hori, K. 2002. Supporting musical composition by externalizing the composer's mental space. In *Proceedings of the 4th Conference on Creativity & Cognition*, CC '02, 165–172. New York, NY, USA: Association for Computing Machinery.

Ashkinaze, J.; Mendelsohn, J.; Qiwei, L.; Budak, C.; and Gilbert, E. 2024. How AI ideas affect the creativity, diversity, and evolution of human ideas: evidence from a large, dynamic experiment. *arXiv preprint arXiv:2401.13481*.

Baldi, P., and Itti, L. 2010. Of bits and wows: A bayesian theory of surprise with applications to attention. *Neural Networks* 23(5):649–666.

Benedek, M.; Jauk, E.; Sommer, M.; Arendasy, M.; and Neubauer, A. C. 2014. Intelligence, creativity, and cognitive control: The common and differential involvement of executive functions in intelligence and creativity. *Intelligence* 46:73–83.

Bi, H.; Mi, S.; Lu, S.; and Hu, X. 2020. Meta-analysis of interventions and their effectiveness in students' scientific creativity. *Thinking Skills and Creativity* 38:100750.

Boden, M. 2004. *The Creative Mind: Myths and Mechanisms*. Taylor & Francis.

Chirico, A.; Glaveanu, V. P.; Cipresso, P.; Riva, G.; and Gaggioli, A. 2018. Awe enhances creative thinking: An experimental study. *Creativity Research Journal* 30(2):123–131.

Chou, Y.-Y. J., and Tversky, B. 2020. Changing perspective: Building creative mindsets. *Cognitive Science* 44(4):e12820.

Chung, J. J. Y.; He, S.; and Adar, E. 2021. The intersection of users, roles, interactions, and technologies in creativity support tools. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*, DIS '21, 1817–1833. New York, NY, USA: Association for Computing Machinery.

Clark, A. 2013. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences* 36(3):181–204.

Demszky, D.; Yang, D.; Yeager, D. S.; Bryan, C. J.; Clapper, M.; Chandhok, S.; Eichstaedt, J. C.; Hecht, C.; Jamieson, J.; Johnson, M.; et al. 2023. Using large language models in psychology. *Nature Reviews Psychology* 2(11):688–701.

DiPaola, S.; Gabora, L.; and McCaig, G. 2018. Informing artificial intelligence generative techniques using cognitive theories of human creativity. *Procedia computer science* 145:158–168.

Dolbier, S. Y.; Dieffenbach, M. C.; and Lieberman, M. D. 2025. Open-mindedness: An integrative review of interventions. *Psychol. Rev.* 132(1):204–238.

Doshi, A. R., and Hauser, O. P. 2024. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances* 10(28):eadn5290.

Espejel, J. L.; Ettifouri, E. H.; Alassan, M. S. Y.; Chouham, E. M.; and Dahhane, W. 2023. GPT-3.5, GPT-4, or BARD? Evaluating LLMs reasoning ability in zero-shot setting and performance boosting through prompts. *Natural Language Processing Journal* 5:100032.

Frich, J.; MacDonald Vermeulen, L.; Remy, C.; Biskjaer, M. M.; and Dalsgaard, P. 2019. Mapping the landscape of creativity support tools in HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, CHI '19, 1–18. New York, NY, USA: Association for Computing Machinery.

Friston, K.; Moran, R. J.; Nagai, Y.; Taniguchi, T.; Gomi, H.; and Tenenbaum, J. 2021. World model learning and inference. *Neural Networks* 144:573–590.

Fu, Y.; Bin, H.; Zhou, T.; Wang, M.; Chen, Y.; Lai, Z. G. D. C.; Wobbrock, J. O.; and Hiniker, A. 2024. Creativity in the age of AI: Evaluating the impact of generative AI on design outputs and designers' creative thinking. *arXiv preprint arXiv:2411.00168*.

Gable, P., and Harmon-Jones, E. 2010. The blues broaden, but the nasty narrows: Attentional consequences of negative affects low and high in motivational intensity. *Psychological Science* 21(2):211–215. PMID: 20424047.

Gandy, S.; Bonnelle, V.; Jacobs, E.; and Luke, D. 2022. Psychedelics as potential catalysts of scientific creativity and insight. *Drug Science, Policy and Law* 8:20503245221097649.

Greco, A.; Rastelli, C.; Ubaldi, A.; and Riva, G. 2025. Immersive exposure to simulated visual hallucinations modulates high-level human cognition. *Consciousness and Cognition* 128:103808.

Hayes, S. C.; Law, S.; Malady, M.; Zhu, Z.; Bai, X.; et al. 2020. The centrality of sense of self in psychological flexibility processes: What the neurobiological and psychological correlates of psychedelics suggest. *Journal of Contextual Behavioral Science* 15:30–38.

Ibarrola, F. J., and Grace, K. 2023. Prompt diversification for iterating with text-to-image models. In *International Conference on Innovative Computing and Cloud Computing*.

Itti, L., and Baldi, P. 2009. Bayesian surprise attracts human attention. *Vision research* 49(10):1295–1306.

Karwowski, M. 2014. Creative mindsets: Measurement, correlates, consequences. *Psychology of Aesthetics, Creativity, and the Arts* 8(1):62.

Kashdan, T. B., and Rottenberg, J. 2010. Psychological flexibility as a fundamental aspect of health. *Clinical psychology review* 30(7):865–878.

Kenett, Y. N.; Levy, O.; Kenett, D. Y.; Stanley, H. E.; Faust, M.; and Havlin, S. 2018. Flexibility of thought in high creative individuals represented by percolation analysis. *Proceedings of the National Academy of Sciences* 115(5):867–872.

Kumar, M.; Goldstein, A.; Michelmann, S.; Zacks, J. M.; Hasson, U.; and Norman, K. A. 2023. Bayesian surprise predicts human event segmentation in story listening. *Cognitive science* 47(10):e13343.

Le Mens, G.; Kovács, B.; Hannan, M. T.; and Pros, G. 2023. Uncovering the semantics of concepts using GPT-4. *Proceedings of the National Academy of Sciences* 120(49):e2309350120.

Liu, Y.; Zhang, K.; Li, Y.; Yan, Z.; Gao, C.; Chen, R.; Yuan, Z.; Huang, Y.; Sun, H.; Gao, J.; He, L.; and Sun, L. 2024. Sora: A review on background, technology, limitations, and opportunities of large vision models.

McCrae, R. R. 1987. Creativity, divergent thinking, and openness to experience. *J. Pers. Soc. Psychol.* 52(6):1258–1265.

Nijstad, B. A.; De Dreu, C. K.; Rietzschel, E. F.; and Baas, M. 2010. The dual pathway to creativity model: Creative ideation as a function of flexibility and persistence. *European review of social psychology* 21(1):34–77.

O'Connor, A. J.; Nemeth, C. J.; and Akutsu, S. 2013. Consequences of beliefs about the malleability of creativity. *Creativity Research Journal* 25(2):155–162.

OpenAI. 2024. GPT-4 Technical Report.

Peeperkorn, M.; Brown, D.; and Jordanous, A. 2023. On characterizations of large language models and creativity evaluation. In *International Conference on Computational Creativity*.

Peters, E. 2006. *The Functions of Affect in the Construction of Preferences*. Cambridge University Press. 454–463.

Pezzulo, G.; Parr, T.; Cisek, P.; Clark, A.; and Friston, K. 2023. Generating meaning: Active inference and generative ai. *PsyArXiv. June* 8.

Piff, P. K.; Dietze, P.; Feinberg, M.; Stancato, D. M.; and Keltner, D. 2015. Awe, the small self, and prosocial behavior. *Journal of personality and social psychology* 108(6):883.

Ritter, S. M.; Damian, R. I.; Simonton, D. K.; van Baaren, R. B.; Strick, M.; Derks, J.; and Dijksterhuis, A. 2012. Diversifying experiences enhance cognitive flexibility. *Journal of experimental social psychology* 48(4):961–964.

Roberts, B. W.; Luo, J.; Briley, D. A.; Chow, P. I.; Su, R.; and Hill, P. L. 2017. A systematic review of personality trait change through intervention. *Psychological bulletin* 143(2):117.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695.

Rowe, G.; Hirsh, J. B.; and Anderson, A. K. 2007. Positive affect increases the breadth of attentional selection. *Proceedings of the National Academy of Sciences* 104(1):383–388.

Royston, R., and Reiter-Palmon, R. 2019. Creative self-efficacy as mediator between creative mindsets and creative problem-solving. *The Journal of Creative Behavior* 53(4):472–481.

Samide, R.; Cooper, R. A.; and Ritchey, M. 2020. A database of news videos for investigating the dynamics of emotion and memory. *Behavior Research Methods* 52:1469–1479.

Smith, S. M., and Ward, T. B. 2012. *The Oxford handbook of thinking and reasoning*. Oxford University Press. chapter Cognition and the Creation of Ideas, 456–474.

Tan, X.; Antony, M.; and Kong, H. 2020. Automated music generation for visual art through emotion. In *International Conference on Computational Creativity*, 247–250.

Thölke, P.; Bellemare-Pepin, A.; Harel, Y.; Lespinasse, F.; and Jerbi, K. 2024. Bio-Mechanical Poet: An Immersive Audiovisual Playground for Brain Signals and Generative AI. In *Proceedings of the 15th International Conference on Computational Creativity*.

Uddin, L. Q. 2021. Cognitive and behavioural flexibility: neural mechanisms and clinical considerations. *Nat. Rev. Neurosci.* 22(3):167–179.

Wiggins, G. A. 2020. Creativity, information, and consciousness: the information dynamics of thinking. *Physics of life reviews* 34:1–39.

Wu, C.-L.; Huang, S.-Y.; Chen, P.-Z.; and Chen, H.-C. 2020. A systematic review of creativity-related studies applying the remote associates test from 2000 to 2019. *Frontiers in psychology* 11:573432.

Zabelina, D. L., and Robinson, M. D. 2010. Creativity as flexible cognitive control. *Psychology of Aesthetics, Creativity, and the Arts* 4(3):136.

Zabelina, D. L.; Friedman, N. P.; and Andrews-Hanna, J. 2019. Unity and diversity of executive functions in creativity. *Consciousness and cognition* 68:47–56.

Zacks, J. M., and Swallow, K. M. 2007. Event segmentation. *Current Directions in Psychological Science* 16(2):80–84. PMID: 22468032.

Zacks, J. M.; Speer, N. K.; and Reynolds, J. R. 2009. Segmentation in reading and film comprehension. *Journal of Experimental Psychology: General* 138(2):307.

Zhang, J. W.; Howell, R. T.; Razavi, P.; Shaban-Azad, H.; Chai, W. J.; Ramis, T.; Mello, Z.; Anderson, C. L.; Monroy, M.; and Keltner, D. 2024. Awe is associated with creative personality, convergent creativity, and everyday creativity. *Psychology of Aesthetics, Creativity, and the Arts* 18(2):209.

Zhang, M.; Wang, F.; and Zhang, D. 2020. Individual differences in trait creativity moderate the state-level mood-creativity relationship. *PloS one* 15(8):e0236987.

Zhou, E., and Lee, D. 2024. Generative artificial intelligence, human creativity, and art. *PNAS nexus* 3(3):pgae052.