



EmoNest: Crafting Emotionally-Adaptive Experiences with Generative AI

Yue Huang

University of Notre Dame

Siyuan Wu

Independent Developer

Abstract

We present EmoNest, a generative AI framework for creating interactive, emotionally adaptive storytelling experiences. By integrating advanced language and vision models with user profiling and real-time narrative adaptation, EmoNest generates personalized stories that reflect each user’s emotional state and personal background. Our approach empowers users to co-create immersive narratives that promote engagement and emotional resonance. We demonstrate EmoNest’s potential for delivering personalized artistic experiences and discuss its implications for the development of emotionally intelligent AI.

1 Introduction

The recent surge in generative AI research has fundamentally transformed how content is created, enabling new applications that span image and video synthesis, creative writing, code generation, and beyond (Cao et al., 2025). As these technologies mature, there is a growing shift from simply producing content to crafting engaging, interactive, and emotionally meaningful experiences. In particular, generative AI has emerged as a promising tool for dynamic narrative generation (Ran et al., 2025), visual storytelling (Kou et al., 2024), and human-AI co-creation (Sun et al., 2022), with increasing attention paid to its potential in supporting emotional and psychological well-being (Sabour et al., 2024; Xie and et al., 2022). Beyond creative tasks, the capacity of generative models to understand, reflect, and adapt to human emotions has opened up new frontiers for applications in mental health, emotional expression, and personalized therapeutic experiences. Recent advances in affective computing, empathetic dialogue modeling, and emotion-conditioned generation have demonstrated the value of tailoring AI-driven interactions to the unique emotional states and personal backgrounds of users (Liu et al., 2024; Fu et al., 2025a).

However, most existing approaches focus either on open-ended creativity or on emotionally adaptive dialogue, rarely combining both to deliver highly personalized and immersive artistic experiences. This gap motivates our exploration of a fundamental question: **Can generative AI be effectively harnessed to create emotionally valuable, story-driven experiences that are both interactive and tailored to individual users’ emotional contexts?**

To address this, we propose **EmoNest**—a novel, modular framework that leverages state-of-the-art language and vision models to generate personalized, emotionally resonant stories. EmoNest dynamically integrates user profiling, multi-turn interaction, and adaptive narrative construction, combining large language models with multi-modal content generation. Through an interactive storytelling pipeline, EmoNest actively incorporates users’ personal context, emotional state, and creative intent, empowering users not only as passive recipients but as co-creators of their own narrative journeys.

Our objectives in this work are twofold: (1) to introduce the design and implementation of the EmoNest framework for emotionally adaptive story-based experience generation; and (2) to investigate how generative AI can provide genuine emotional comfort, engagement, and insight through personalized




Figure 1: The interface of questionnaire, tagging, and controllable inputs.

artistic expression; By bridging the fields of generative art, affective computing, and interactive storytelling, we aim to lay the foundation for a new class of creative AI applications that are not only intelligent, but truly emotionally aware and responsive.

2 EmoNest

We present **EmoNest**, a modular generative system designed to deliver personalized, emotionally resonant interactive story experiences. EmoNest leverages multiple types of generative models, including LLMs, text-to-image models, and video generative models, to capture user profiles and drive multi-turn, adaptive narrative generation.

2.1 Before Generation: User Profiling & Story Preparation

The first stage of EmoNest is dedicated to capturing the user’s personal context and creative intent. Typically, the user initiates interaction by submitting a brief, personally meaningful **seed experience**, which serves as the thematic anchor for subsequent narrative generation. To make the generated story more vivid, immersive, and resonant with the user’s emotional state, we conduct a comprehensive user profiling process after the submission of this seed experience. This involves enabling the LLM to actively elicit clarifying information via questionnaires, analyze user-defined tags for deeper personalization, and support user-specified keywords or guiding sentences to ensure controllability and alignment with the desired story style and elements.

a) Adaptive Questionnaire and Structured Response Capture. To enrich the user profile, the system employs an LLM to dynamically generate 5–10 clarifying questions tailored to the submitted seed experience. These questions are multi-modal (including multiple-choice and free-text formats) and cover aspects such as emotional state, background, coping style, and interests. User responses are subsequently ingested, normalized, and stored in a structured format, supporting both constrained and open-ended answers. This enables fine-grained downstream adaptation of the narrative to the user’s specific context.

b) Tagging and Profile Enrichment. Users may further annotate themselves with self-descriptive tags (e.g., personality traits, interests), which are both suggested by the LLM and user-customizable. All information is consolidated into a `UserProfile` data structure, which can be further abstracted into a concise summary via LLM-based aggregation.

c) Story Preparation and Controllable Inputs. To further empower user agency and ensure controllability over the generated narrative, EmoNest allows users to specify additional *story keywords* or to input guiding sentences that reflect their desired story style, elements, or scenarios. These user-provided signals are explicitly incorporated into the story generation pipeline, enabling control over narrative aspects such as tone, themes, character archetypes, or specific settings. This preparatory phase ensures that both the macro-level style and micro-level elements of the story are aligned with the user’s preferences, facilitating a more engaging and personalized storytelling experience.

2.2 Story Generation Pipeline

Story creation in EmoNest is decomposed into five compositional components, each tightly coupled with the evolving user profile:




Figure 2: EmoNest Workflow

- **Story Foundation:** Generation of the story title, core concept, thematic focus, and emotional tone, grounded in seed experience and user profile.
- **Scene Construction:** Creation of principal settings, temporal context, atmosphere, and unique environmental cues.
- **Character Design:** Instantiation of protagonist, supporting, and antagonist roles; character attributes are explicitly aligned with user-specified traits and goals.
- **Narrative Structuring:** Automated synthesis of multi-act narrative scaffolding to support interaction and non-linear story progression.
- **Opening and Interactive Elements:** Generation of the introductory narrative, dialogue, initial branching points, and hidden options, providing a launchpad for immersive user interaction.

Each stage is realized via successive LLM invocations, conditioned on the cumulative system context and outputs of prior components. Implementation details are encapsulated within dedicated API interfaces and prompt engineering modules.

2.3 Interactive Storytelling and Progression

Dialogue Tracker. The Dialogue Tracker is responsible for persistently storing all relevant information throughout the interactive storytelling session, including user-system dialogue history, evolving plot developments, character attributes, and stateful contextual metadata. This module provides structured access for downstream modules, enabling seamless retrieval and update of narrative elements, character trajectories, and user-driven choices at each turn. The Dialogue Tracker thus forms the foundation for both consistent story continuity and rich, context-aware narrative adaptation.

Meta-Reflector. Interactive storytelling systems face several fundamental challenges: 1) Story progression can be sluggish, with narratives occasionally stalling or lacking timely advancement; it is often unclear when and how to appropriately propel the plot. 2) Effective narrative development requires tight alignment with evolving user input, ensuring the story remains emotionally relevant and responsive to the user’s needs. 3) Purely textual interactions may be insufficient for deep immersion, highlighting the need for richer, multi-modal interactive elements.

To address these challenges, we introduce the Meta-Reflector module, which orchestrates higher-order reasoning and dynamic narrative control through several key mechanisms: 1) To prevent conversational stagnation, the Meta-Reflector performs periodic summarization—after every three dialogue rounds, an LLM-generated summary distills recent developments; after every six rounds, consecutive summaries are compared, and, if narrative stasis is detected, the system forcibly advances the plot. 2) The module continuously monitors the user’s emotional state and pivotal decisions, enabling reflective and emotionally attuned responses that keep the narrative closely aligned with the user’s affective journey. 3) In appropriate contexts, the Meta-Reflector proactively introduces interactive elements




Figure 3: User interface of EmoNest.

powered by generative models—including images, videos, or executable JavaScript code—to enhance engagement and foster a more immersive, multi-modal storytelling experience.

2.4 User Interface

As illustrated in **Figure 3**, the user interface is composed of several distinct components. Component ① displays the story title along with a brief description of the opening scene. Component ② presents the dialogue between characters, featuring the user’s role on the right and non-player characters (NPCs) on the left. Component ③ provides the interactive module (in this example, a JavaScript-based element), where the user can engage by clicking relevant buttons. Component ④ offers narrative feedback following the user’s response, further describing the current stage of the story. Component ⑤ is the input area for conversation, enabling the user to either select recommended responses to NPCs or enter custom input, with the option to specify the replying role (e.g., in this case, the user responds as Lily Woods). Finally, Component ⑥ displays detailed story information, including scene details (such as location, key events, and scene description), character profiles, and journey data (e.g., transitions between different scenes).

2.5 Story Termination and Summary Output

Upon story conclusion, the system synthesizes a comprehensive summary that integrates the initial seed experience, narrative evolution, and emotional milestones. Following the “End Story” trigger, EmoNest compiles the full interaction log and generates a final output, which is presented to the user along with a reflective message designed to maximize emotional value and promote user insight or healing. This final output primarily consists of a structured summary that covers the following core aspects: Story Arc (a high-level overview of the narrative trajectory), Character Development (the evolution and growth of key characters throughout the story), and Key Insights (personalized reflections and emotional takeaways derived from the user’s journey). An example is shown in **Figure 4**.

3 Case Study

EmoNest supports a wide variety of interactive elements. As shown in **Figure 5**, EmoNest generates various kinds of code-based interactive elements. Specifically, the generated elements




Figure 4: Example of the story termination as well as the summary. From left to right, the blocks represent Story Arc, Character Development, and Key Insights, respectively.




Figure 5: Generated interactive elements through JavaScript code.

encompass multiple interaction types, such as dynamic text displays (e.g., simulating the opening and gradual reveal of a letter), simple puzzle or decryption games (e.g., unlocking a door or cracking a password-protected box), and visually engaging information or Easter egg presentations. These interactive modules enrich the user experience by introducing diverse modalities of engagement within the narrative flow. EmoNest supports a wide variety of interactive elements, including the dynamic generation of diverse and vivid scene images. The system can create richly detailed visual environments that enhance the narrative and provide immersive, context-aware backdrops.

Table 1: An example of character description.

Character	Description
Emily Green (You)	A 28-year-old woman with shoulder-length auburn hair often tucked behind her ear. She wears a worn denim jacket over a flannel shirt, dark skinny jeans, and carries a leather satchel with a notebook and camera peeking out.
Lily Woods	A key figure. A tall, freckled woman with long, curly chestnut hair often tied back with a silk scarf. She favors cozy cardigans in autumnal hues and always carries a sketchbook.
Arthur McNeil	A key figure. A dignified man in his mid-60s with silver hair swept back and wire-rimmed glasses. He wears tweed jackets with elbow patches and carries a pocket watch fused to a brass chain.
Margaret Davis	A key figure. A stern woman in her mid-50s with charcoal-gray hair pulled into a tight bun. She wears tailored blazers and brooches shaped like roses—always immaculately pressed.
The Faded Letter	A key figure. An aged, yellowed envelope with delicate cursive peeking through a cracked seal. When read aloud, it seems to "whisper" under one's breath, and its ink seems to shift in the attic's slanted sunlight.

EmoNest is capable of generating rich and diverse characters. As illustrated in Table 1, each with vivid visual descriptions and distinct personalities. The generated cast spans a wide range of ages, appearances, and personal styles—from a young, adventurous protagonist to dignified elders and




Figure 6: Generated characters' profiles through text-to-image models.

even personified objects such as “The Faded Letter.” This diversity not only enhances narrative depth, but also allows for more immersive and emotionally engaging storytelling.

The stories generated by EmoNest are highly aligned with the user’s personal experiences. In [Appendix C](#), we show some cases about story details and assess how the story aligns with the users’ emotional experiences. The cases demonstrate that the stories generated by EmoNest are highly aligned with users’ personal experiences. For each input—often describing nuanced emotional struggles such as difficulties with identity in close relationships, coping with grief, or anxiety about the future—EmoNest provides a targeted psychological analysis and narrative response that directly addresses the user’s specific emotional concerns. As shown in the “Analysis” and “Assessment” columns, the system consistently tailors its feedback and storylines to the user’s context, ensuring that the generated stories deeply resonate with the user’s real-life emotions and experiences.

4 Conclusion and Future Work

In this work, we presented EmoNest, a novel generative AI framework for emotionally adaptive, interactive storytelling. By integrating advanced language and vision models with real-time user profiling, EmoNest enables highly personalized and emotionally resonant narratives, allowing users to co-create stories that foster engagement, reflection, and connection.

Looking ahead, several important directions remain. First, a large-scale human evaluation has not yet been conducted. Comprehensive user studies with diverse populations are a critical next step to assess EmoNest’s emotional efficacy, personalization, and broader impact. Second, improving the consistency and coherence of generated visual content—especially character appearance, scene attributes, and symbolic motifs—will be crucial for immersive storytelling. Future efforts will explore cross-modal alignment, persistent character embeddings, and context-aware control. Finally, enhancing game fluidity and reducing latency remain challenges. We plan to optimize model efficiency through asynchronous generation, local caching, and lightweight surrogate models to ensure a more seamless user experience.

Together, these directions aim to strengthen EmoNest’s potential in emotionally intelligent, multi-modal human-AI interaction, with wide-reaching implications for creative expression and emotional well-being.

References

- Yihan Cao, Siyu Li, Yixin Liu, Zhiling Yan, Yutong Dai, Philip Yu, and Lichao Sun. A Survey of AI-Generated Content (AIGC). *ACM Comput. Surv.*, 57(5), January 2025. ISSN 0360-0300. doi: 10.1145/3704262. URL <https://doi.org/10.1145/3704262>.
- Yiting Ran, Xintao Wang, Tian Qiu, Jiaqing Liang, Yanghua Xiao, and Deqing Yang. Book-World: From Novels to Interactive Agent Societies for Creative Story Generation. *arXiv preprint arXiv:2504.14538*, 2025.
- Ziyi Kou, Shichao Pei, and Xiangliang Zhang. LeMon: Automating Portrait Generation for Zero-Shot Story Visualization with Multi-Character Interactions. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '24, page 1418–1427, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704901. doi: 10.1145/3637528.3671850. URL <https://doi.org/10.1145/3637528.3671850>.
- Yuqian Sun, Xuran Ni, Haozhen Feng, Ray LC, Chang Hee Lee, and Ali Asadipour. Bringing stories to life in 1001 nights: A co-creative text adventure game using a story generation model. In *International Conference on Interactive Digital Storytelling*, pages 651–672. Springer, 2022.
- Sahand Sabour, Siyang Liu, Zheyuan Zhang, June M Liu, Jinfeng Zhou, Alvionna S Sunaryo, Juanzi Li, Tatia Lee, Rada Mihalcea, and Minlie Huang. Emobench: Evaluating the emotional intelligence of large language models. *arXiv preprint arXiv:2402.12071*, 2024.
- Weidi Xie and et al. Psychology-guided controllable story generation. In *COLING*, 2022.
- Zhiwei Liu, Kailai Yang, Qianqian Xie, Tianlin Zhang, and Sophia Ananiadou. Emollms: A series of emotional large language models and annotation tools for comprehensive affective analysis. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5487–5496, 2024.
- Yumeng Fu, Junjie Wu, Zhongjie Wang, Meishan Zhang, Lili Shan, Yulin Wu, and Bingquan Liu. LaERC-S: Improving LLM-based Emotion Recognition in Conversation with Speaker Characteristics. In Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert, editors, *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6748–6761, Abu Dhabi, UAE, January 2025a. Association for Computational Linguistics. URL <https://aclanthology.org/2025.coling-main.451/>.
- Shuai Yang, Yuying Ge, Yang Li, Yukang Chen, Yixiao Ge, Ying Shan, and Yingcong Chen. Seed-story: Multimodal long story generation with large language model. *arXiv preprint arXiv:2407.08683*, 2024a.
- Phoebe J Wang and Max Kreminski. Guiding and diversifying LLM-based story generation via answer set programming. *arXiv preprint arXiv:2406.00554*, 2024.
- Shayan Talaei, Meijin Li, Kanu Grover, James Kent Hippler, Diyi Yang, and Amin Saberi. StorySage: Conversational Autobiography Writing Powered by a Multi-Agent Framework. *arXiv preprint arXiv:2506.14159*, 2025.
- Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation. In *ACL*, 2018.
- Hannah Rashkin, Antoine Bosselut, Maarten Sap, and Yejin Choi. PlotMachines: Outline-conditioned generation with dynamic plot state tracking. In *EMNLP*, 2020.
- Haofeng Wen and et al. GROVE: A retrieval-augmented complex story generation framework with a forest of evidence. *arXiv preprint arXiv:2306.15625*, 2023.
- Tarik Belouadi and Steffen Eger. ByGPT5: End-to-End Style-conditioned Poetry Generation with Token-free Language Models. In *ACL*, 2023.
- Andy Coenen and et al. Wordcraft: a human-ai collaborative editor for story writing. In *EACL 2021 Workshop on Human-AI Co-Writing*, 2021.
- Ahmed Radwan and et al. SARD: A human-ai collaborative story generation. *arXiv preprint arXiv:2403.01575*, 2024.

- Weize Mao and et al. Procedural Content Generation via Generative Artificial Intelligence. *arXiv preprint arXiv:2403.03065*, 2024.
- Chengpeng Hu, Yunlong Zhao, and Jialin Liu. Game generation via large language models. In *2024 IEEE Conference on Games (CoG)*, pages 1–4. IEEE, 2024.
- Mingyu Yang, Junyou Li, Zhongbin Fang, Sheng Chen, Yangbin Yu, Qiang Fu, Wei Yang, and Deheng Ye. Playable game generation. *arXiv preprint arXiv:2412.00887*, 2024b.
- Haoxuan Che, Xuanhua He, Quande Liu, Cheng Jin, and Hao Chen. Gamegen-x: Interactive open-world game video generation. *arXiv preprint arXiv:2411.00769*, 2024.
- Jiaying Fu, Xiruo Wang, Kate Vi, Zhouyi Li, Chuyan Xu, and Yuqian Sun. "I Like Your Story!": A Co-Creative Story-Crafting Game with a Persona-Driven Character Based on Generative AI. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pages 1–5, 2025b.
- Yongkui Wang and et al. Open-world Story Generation with Structured Knowledge Enhancement: A Comprehensive Survey. *Neurocomputing*, 2023.
- Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. Towards empathetic open-domain conversation models: a new benchmark and dataset. In *ACL*, 2019.
- Hao Zhou and et al. Emotional chatting machine: Emotional conversation generation with internal and external memory. In *AAAI*, 2018.
- Bodhisattwa Prasad Majumder and et al. Mime: Mimicking emotions for empathetic response generation. In *EMNLP*, 2020.
- Yuxuan Lei, Dingkan Yang, Zhaoyu Chen, Jiawei Chen, Peng Zhai, and Lihua Zhang. Large vision-language models as emotion recognizers in context awareness. *arXiv preprint arXiv:2407.11300*, 2024.
- Dawei Huang, Qing Li, Chuan Yan, Zebang Cheng, Yurong Huang, Xiang Li, Bin Li, Xiaohui Wang, Zheng Lian, and Xiaojiang Peng. Emotion-Qwen: Training Hybrid Experts for Unified Emotion and General Vision-Language Understanding. *arXiv preprint arXiv:2505.06685*, 2025.
- Saizheng Zhang and et al. Personalizing dialogue agents: I have a dog, do you have pets too? In *ACL*, 2018.
- Stephen Roller and et al. Recipes for building an open-domain chatbot. *arXiv preprint arXiv:2004.13637*, 2020.
- Li Li, Peilin Cai, Ryan A Rossi, Franck Dernoncourt, Branislav Kveton, Junda Wu, Tong Yu, Linxin Song, Tiankai Yang, Yuehan Qin, et al. A Personalized Conversational Benchmark: Towards Simulating Personalized Conversations. *arXiv preprint arXiv:2505.14106*, 2025.
- Sandra C Matz, Jacob D Teeny, Sumer S Vaid, Heinrich Peters, Gabriella M Harari, and Moran Cerf. The potential of generative AI for personalized persuasion at scale. *Scientific Reports*, 14(1):4692, 2024.

A Background

Generative AI for Artistic Creation. Automatic story generation and other creative writing tasks have seen rapid progress due to large-scale pre-trained language models and advances in controllable text generation (Yang et al., 2024a; Wang and Kreminski, 2024; Talaei et al., 2025). Early works such as Fan et al. (2018) introduced hierarchical frameworks to improve narrative coherence, while Rashkin et al. (2020) proposed dynamic plot state tracking to better condition generated stories on outlines. More recent approaches emphasize structured control and knowledge injection, e.g., PICS (Xie and et al., 2022) leverages psychological state modeling for protagonist-driven story arcs, and GROVE (Wen and et al., 2023) employs retrieval-augmented prompting for complex plot development. BookWorld (Ran et al., 2025) constructs a multi-agent simulation grounded in existing novels to generate interactive, original stories.

Creativity beyond stories, such as poetry, script, and game content, has also benefited from generative models. ByGPT5 (Belouadi and Eger, 2023) explores style-conditioned poetry generation, demonstrating the potential of end-to-end token-free Transformers. For collaborative and interactive writing, tools like Wordcraft (Coenen and et al., 2021) and SARD (Radwan and et al., 2024) enable mixed-initiative human-AI co-creation through dialogue or node-based interfaces. In game design and procedural content generation, generative AI has been employed for levels, quests, and creative visuals (Mao and et al., 2024; Hu et al., 2024; Yang et al., 2024b; Che et al., 2024; Fu et al., 2025b). Despite these advances, challenges remain in controllability, diversity, and human-AI collaboration for open-ended creative tasks (Wang and et al., 2023; Kou et al., 2024).

Generative AI in Emotion Application & Personalization. Emotional and personalized interaction has emerged as a critical research direction in generative dialogue and affective computing. Empathetic dialogue modeling was pioneered by Rashkin et al. (2019), who released a large dataset for training and evaluating empathy in open-domain conversations. Techniques for emotion conditioning in neural generation include the emotional chat machine (Zhou and et al., 2018) and MIME (Majumder and et al., 2020), which dynamically mirrors user emotions for more empathetic responses. Recently, Large Language Model (LLM) and Large Vision-Language Model (LVLM) are widely used for emotional applications, including EmoLLMs (Liu et al., 2024) and LVLMs as emotion recognizers (Lei et al., 2024; Huang et al., 2025).

Personalization is commonly addressed by conditioning on user profiles or personas, as in PersonaChat (Zhang and et al., 2018) and the BlenderBot architecture (Roller and et al., 2020), both of which combine personality with emotional and factual skills in multi-turn conversation. Recently, PersonaConvBench (Li et al., 2025) introduces a large-scale benchmark combining personalized reasoning and conversational structure across multi-turn dialogues. Matz et al. (2024) demonstrate that LLMs can automatically generate personalized messages that are much more persuasive than generic ones.

B Implement Details

The front end of EmoNest is developed using the React framework with TypeScript for type-safe programming and is initialized and built with the Create React App tool. The user interface leverages the Material-UI (MUI) component library to achieve a modern and responsive design. Dependency management is handled by npm, with key dependencies including axios (for HTTP communication with the backend API), react-router-dom (for routing management), as well as a range of type declaration packages to enhance the development experience and code reliability. The front-end project supports hot reloading and unit testing.

C Story Cases

Story Case I

User Input: I struggle to maintain my identity in close relationships, often prioritizing others' needs over my own, which leads to a sense of self-loss and emotional conflict.

Story Analysis:

- **Story Arc Summary:** In *Shadows on the Horizon*, Elena Montrose embarks on a transformative journey marked by four distinct acts—from the murmur of distant memories in *Whispers of the Past* to the reflective renewal in *Dawn of Reflection*. Throughout her passage, she navigates the delicate balance between the heaviness of old wounds and the cathartic power of embracing change, symbolized by her contemplative stroll down a cobblestone path toward a storm-lashed ocean.
- **Character Development:** Elena, our central figure, evolves from a state of inner turmoil and latent grief into a more aware and self-accepting individual. At the beginning of her journey, her eyes mirror unresolved pain and a hesitant acknowledgment of her past. Through pivotal moments—like the recall flashback that rekindles cherished memories and her encounter with the wise Mrs. Walker—Elena learns to integrate her personal history with her present, ultimately fostering a renewed sense of purpose. Even secondary elements, such as the symbolic folding of the street lamp's glow and the hidden carving reading “remember”, underscore her gradual transition from isolation to enlightened connection.
- **Emotional Journey:** Elena's emotional voyage is carefully mapped out by the metaphor of a storm: beginning with the quiet darkness of suppressed memories, escalating into the turbulent surge of the present, and finally easing into moments of calm introspection. This progression mirrors the internal shift from unprocessed pain to gentle acceptance and self-forgiveness. Her steps on the narrow cobblestone path reflect both her hesitancy and her determination to confront the past, culminating in a moment of solace and inner healing.
- **Key Insights:**
 - True healing requires confronting painful memories rather than repressing them, allowing the shadows of the past to be acknowledged and integrated.
 - Mentorship and support—from figures like Mrs. Walker—can provide essential guidance when navigating inner turmoil.
 - Symbolic language and imagery, such as the storm and the carving that reads “remember”, serve as powerful metaphors for resilience and the persistence of memory.
 - Moments of quiet reflection are as crucial as the active confrontation of painful experiences, emphasizing that recovery is a blend of emotional release and calm acceptance.
- **Therapeutic Value:** The story serves as a metaphorical map for anyone facing their own emotional storms. By illustrating Elena's progression from brokenness to wholeness, it validates the therapeutic process of revisiting and reframing past experiences. The narrative encourages readers to see their struggles not as permanent scars but as integral parts of a larger story—one that can lead to healing and enlightenment when met with courage, introspection, and support.

Assessment: The input reveals difficulties in maintaining self-identity within intimate relationships and a tendency to prioritize others' needs excessively, leading to emotional exhaustion and resentment. The analysis highlights that “Vulnerability is a gateway to deeper connections and self-understanding,” encouraging the user to embrace authenticity and express personal needs within close relationships. Additionally, the insight that “Healing often begins in everyday moments and quiet spaces, not grand gestures” suggests that emotional improvement can arise from small, daily experiences and introspection. Overall, the analysis emphasizes self-reflection, patience, and acceptance of subtle life transitions—factors that are instrumental in resolving the user's relational struggles.

Story Case II

User Input: I constantly worry about the future and replay past mistakes, struggling to live in the present or make decisions.

System Analysis:

- **Story Arc Summary:** The narrative follows Morgan Reed's courageous journey from isolation as he confronts the haunting echoes of a hidden past. Beginning in the reflective corridors of memory, Morgan is drawn inward to confront personal shadows and unresolved legacies. Each act—the initial stirrings of reminiscence, the deep dive into memory's reverberations, the courageous confrontation of inner demons, and finally, the rejuvenating renewal in a public, life-affirming space—marks a stage in his quest for healing and reconnection.
- **Character Development:** Morgan Reed evolves from a withdrawn, haunted individual into a person who finds strength in vulnerability. Initially isolated by his past, Morgan learns to acknowledge and embrace painful memories, using moments of internal reflection and external guidance (as exemplified by the supportive entrance of Antonio) to slowly reclaim his power. Key characters, though not extensively detailed, serve as symbolic guides that help him navigate the labyrinth of recollections, emphasizing the importance of both solitary introspection and shared human connection.
- **Emotional Journey:** The emotional arc is a poignant transition from loneliness and uncertainty to empowerment and hope. Morgan's journey is marked by moments of hesitation, deep introspection, and gradual acceptance of his past. The narrative vividly portrays how embracing even the darkest memories can pave the way for renewal, reaffirming that confronting internal shadows can lead to a more connected and joyful existence.
- **Key Insights:**
 - Healing begins with confronting, rather than avoiding, the hidden parts of our past.
 - Vulnerability is a source of strength, not weakness, and it is through acknowledging our true selves that we forge paths to recovery.
 - Internal reflection, coupled with the support of trusted mentors, can transform isolation into a space for growth.
 - The journey toward reconnection is non-linear, marked by pauses and moments of introspection that are essential in the healing process.
 - Even in the smallest, nearly imperceptible details of our lives, there are messages—easter eggs of wisdom—that hint at the possibility of renewal.
- **Therapeutic Value:** The story embodies a therapeutic message by illustrating that the path to healing is deeply personal and requires confronting one's past. Through Morgan's internal struggles and eventual reconnection with supportive figures, the narrative offers a metaphor for reclaiming one's identity after prolonged isolation. This deeply resonant journey invites readers to reflect on their own life stories, recognize the value in embracing both light and shadow, and ultimately, find hope in the promise of renewal.

Assessment: The input mentions that "I" am always worried about the future and repeatedly dwell on past mistakes, making it difficult to live in the present and enjoy the beautiful moments in life. The analysis emphasizes facing the hidden parts of the past, recognizing that vulnerability is a source of strength, achieving growth through self-reflection and support from others, and understanding that the path to healing is nonlinear. All of these can directly or indirectly help to resolve the emotional dilemmas caused by my past mistakes and worries about the future, guiding "me" to treat the past correctly and accept myself, thus finding a way out of the current emotional difficulties.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading “NeurIPS Paper Checklist”,**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: We directly show our research aim and contributions in the abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have discussed the limitation of this paper in [section 4](#).

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This paper does not need theoretical proofs and assumptions.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: This paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.

- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We are not currently releasing the source code for our framework. The paper provides detailed methodology descriptions, but we plan to make the code publicly available in the future after addressing proprietary components and completing additional validation work.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.

- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: This paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: This paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: This paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.

- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research adheres to the NeurIPS Code of Ethics by ensuring transparency, documentation, and measures against potential societal and environmental impacts, as detailed in our methodologies and data handling practices.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss the broader impacts in [section 4](#)

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The dataset and models mentioned in this work do not involve a high risk of misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: In this work, we used pre-trained models following the licenses and terms specified by the creator, and strictly adhered to the licenses for existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: This paper describes the usage of LLMs as they constitute an important component of our core research methodology.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.