

D.M.A.I.H.: Deepfake-Inspired Few-Shot Learning Approach with Stable Diffusion for Digital Mourning

Btissam Acim, Hamid Ouhnni, Nassim Kharmoum, Soumia Ziti

Computer Science Department-Intelligent Processing and Security of Systems (IPSS) Laboratory-Faculty of Sciences,
Mohammed V University in Rabat, Morocco

Abstract—Digital mourning (*deuil numérique*) is the use of digital and AI-based technologies to preserve, recontextualize, and extend the memory of deceased loved ones through personalized and meaningful virtual representations. The digital mourning process requires innovative technologies capable of preserving the memory of deceased loved ones in meaningful and humanized ways. This paper proposes a novel generative approach, D.M.A.I.H. (Digital Mourning with Artificial Intelligence for Healing), for digital grief, with a focus on moral support and the mental health of bereaved relatives, using Stable Diffusion with a few-shot learning adaptation mechanism. The system takes as input a small set of personal references (e.g. a portrait, contextual images such as the person's home, and a short descriptive script) and outputs high-quality, photorealistic images of the deceased in different meaningful contexts, process closely related to deepfake generation but redirected here toward therapeutic and commemorative purposes. Unlike traditional generative models requiring large datasets, Few-shot personalization is leveraged to adapt Stable Diffusion to each individual with minimal data, enabling the generation of personalized digital albums. Experimental results show that the model consistently preserves identity in the images it produces, and contextual control ensures emotional resonance. In particular, identity similarity scores for the generated images ranged from 0.88 to 0.93, with an average score of 0.91, testifying to strong identity preservation across all outputs. The innovation of this study is a foundation for AI-based memorialization, balancing technological innovation with concerns over privacy, authenticity and cultural sensitivity, and psychological comfort.

Keywords—*Stable diffusion; few-shot learning; deepfake; Artificial Intelligence (AI); generative AI; digital mourning*

I. INTRODUCTION

Generative Artificial Intelligence (AI) is developing at a fast rate and transforming various industries such as art, education, healthcare [1], [2] and entertainment. Here, one of the newly emerging and yet to be discovered uses is digital mourning, where AI can assist in keeping the memory of deceased loved ones alive and providing psychological and ethical assistance to bereaved families [3], [4]. Traditional modes of commemoration, photographs, videos, or symbolic objects remain crucial but static and limited in their ability to capture the emotional richness and contextual diversity of an individual's life. Generative AI, and in particular the proposed approach, D.M.A.I.H. (Digital Mourning with Artificial Intelligence for Healing), enables a dynamic and customized recontextualization of memory, generating visual forms that can enrich the affective bond and accompany mourning [5].

This area, however, presents several difficulties. Ensuring identity consistency in the generated images is the primary challenge, considering that it is essential to accurately preserve the individual's distinctive characteristics. Limited data availability is also a primary constraint, since there are seldom several personal photos. Moreover, there are psychological and ethical issues to be addressed with sensitivity: while this technology has the potential to boost mental health and well-being through moral and empathic remembering, as substantiated by empirical studies on chatbots such as Woebot [6] and longitudinal human-computer interaction for affective care [7], it must also safeguard dignity and authenticity in commemoration [8].

The proposed approach introduces D.M.A.I.H., an innovative method that leverages enhanced Stable Diffusion with few-shot learning to generate personalized memorial albums from a small number of reference images, contextual backgrounds, and descriptive scripts. Unlike conventional generative models that require large datasets, D.M.A.I.H. adapts efficiently to each individual, ensuring both identity preservation and faithful contextualization [9].

As opposed to existing digital commemoration sites such as Facebook Memorialized Accounts [10], Instagram memorial accounts [11], or websites such as Forever Missed, which allow for mostly static or symbolic remembering, our approach offers a generative pipeline that can generate personalized, contextualized, and photorealistic memorial albums. This more appropriately situates the contribution in the field and highlights its novelty without overstating it.

First, D.M.A.I.H. is introduced as a generative approach inspired by recent advances in deepfake synthesis, but oriented toward ethically grounded memorialization, specifically designed for digital mourning and capable of producing personalized albums with high emotional value. Second, it is demonstrated that the integration of few-shot learning into Stable Diffusion effectively overcomes the limitation of scarce data while maintaining identity consistency. Finally, the D.M.A.I.H. situated within a broader perspective of psychological and moral support, emphasizing the potential of this technology to assist the mental health of grieving families while opening an ethical debate on its responsible use [12].

In brief, the D.M.A.I.H. approach was adopted since it is an approach that addresses well the core issues of digital bereavement, i.e., sparsity of data, identity integrity, and emotional relevance. The existing approaches, though promising, are afflicted by either excessive computational cost,

low identity integrity, or limited contextual adaptability and thus are less suitable for the sensitive nature of loss. On the other hand, D.M.A.I.H. integrates Stable Diffusion and few-shot learning to bypass these limitations to develop a robust generative framework for personalized memorial albums.

The remaining part of this paper is organized as follows: Section II gives an overview of related work, Section III presents the proposed model, Section IV provides experimental results, Section V discusses the findings in the context of previous studies, and Section VI concludes with future directions and ethical implications.

II. RELATED WORK

Generative AI has been marked by an ongoing enhancement of Generative Adversarial Networks (GANs) [13] and diffusion models. *Deepfake* techniques, originally popularized through GAN-based face synthesis, illustrate both the power and the risks of these generative systems. StyleGAN, for example, has performed incredibly well in face generation [14], while Stable Diffusion, a diffusion model, has transformed text-to-image synthesis through its controllability and replication of visuals [15]. In order to fulfill personalization needs with limited data, various approaches have emerged, such as DreamBooth [16], LoRA (Low-Rank Adaptation) [17], and DreamArtist [18]. These approaches facilitate identity preservation with as few as 5 to 15 images, and thus are suitable in situations where the data is not abundant. However, their applications have largely been limited to designing avatars, artwork, or entertainment, and have been investigated extremely sparingly in delicate topics such as loss.

Digital technology has revolutionized commemoration practices, to the extent of producing what has been referred to as digital mourning. Social media such as Facebook Memorialized Accounts [19], or Instagram Memorial Accounts [20], allow surviving relatives to keep deceased individuals' pages and post memories day and night. Memory-keeping websites such as Forever Missed and Legacy.com also provide virtual space for memory preservation.

Alongside these static memorials, AI-driven initiatives have emerged. HereAfter, AI records a person's voice and stories during their lifetime, enabling interactive experiences after death. StoryFile offers similar interactive video testimonies. Deep Nostalgia (MyHeritage) animates old photographs using AI, while projects such as Eternime envisioned the creation of persistent digital avatars. Research prototypes and patents, such as Microsoft's so-called "Deadbots", have also suggested the use of digital traces to simulate posthumous conversations. While these solutions illustrate the growing interest in digital commemoration, they remain limited to static memorials, conversational agents, or simple animations, without providing truly personalized visual albums.

In addition to commemoration, AI is more and more studied for its therapeutic applications in mental health treatment. Therapeutic robots such as Replika are used for emotional care, and virtual avatars have been explored for therapy. Existing studies highlight the ability of AI to foster well-being, empathy, and psychological resilience [21], [22]. These approaches often

favor verbal or dialogic interaction, while the visual and contextual dimensions of loss support are largely untapped.

This review demonstrates that while generative personalization methods and digital mourning platforms exist separately, no study to date has proposed a unified pipeline for generating memorial albums using Stable Diffusion with few-shot learning. Current platforms are limited to conversation or static commemoration, whereas AI-based personalization methods remain confined to art and entertainment. The proposed approach addresses this gap by introducing D.M.A.I.H., an innovative generative framework capable of producing photorealistic and contextualized albums of deceased individuals, thereby bridging generative AI research with the psychological needs of digital mourning.

III. PROPOSED MODEL

The proposed D.M.A.I.H. approach leverages the generative capacity of Stable Diffusion, enhanced with few-shot learning, to create personalized memorial albums. Stable Diffusion was selected as the backbone of the proposed approach because it combines photorealistic image generation with efficient conditioning mechanisms, enabling both identity preservation and contextual control. Compared to GAN-based models, which must operate on vast datasets and are often beset with instability (e.g. mode collapse), diffusion models and Stable Diffusion are more stable, more robust, and more scalable for supervised image creation. While deepfake techniques have traditionally relied on GAN-based architectures, the D.M.A.I.H. approach demonstrates how diffusion models can achieve similar identity-focused synthesis while being redirected toward therapeutic and commemorative purposes.

The whole pipeline of the D.M.A.I.H. approach is illustrated in Fig. 1, where identity extraction, contextual conditioning, and text guidance are integrated into the U-Net backbone of Stable Diffusion.

A. Approach Overview

The proposed D.M.A.I.H. approach is designed to address three key challenges in digital mourning: data scarcity, identity preservation, and emotional relevance. The system takes as input a small set of target images of the deceased, contextual scenes representing meaningful environments (e.g. home, farm, or interior spaces), and a descriptive script provided by relatives.

These heterogeneous inputs are processed by specialized encoders and conditioning modules, then fused within the denoising process of Stable Diffusion's U-Net. The result is the generation of multiple high-quality, photorealistic images which collectively form a personalized memorial album that preserves both identity fidelity and contextual meaning.

B. Detailed Description of Model Blocks

1) *Target (Identity extraction)*: The reference images of the deceased are first segmented to isolate the face. A Face Encoder (ArcFace or InstantID) extracts identity tokens that encode the distinctive visual traits of the individual, ensuring faithful identity preservation [23].

2) *Scene (Contextual integration)*: Contextual images are processed with ControlNet with depth, Canny edges, and pose

constraints. This guarantees realistic integration of the target subject into meaningful environments [24].

3) *Descriptive script (Semantic guidance)*: A descriptive script provided by relatives is embedded using CLIP/BERT. These embeddings introduce semantic and emotional guidance [25], [26], reinforcing contextual fidelity and emotional resonance in the generation process.

4) *Fusion in U-Net (Stable diffusion)*: Identity tokens, contextual constraints, and text embeddings are fused during the denoising process of the U-Net backbone within Stable Diffusion, yielding coherent and photorealistic outputs.

5) *Output (Memorial album)*: The system produces multiple personalized images under different contexts. These are assembled into a digital memorial album, preserving both the identity of the deceased and the emotional significance of the scenes.

ArcFace was selected due to its state-of-the-art performance on face recognition benchmarks such as LWF (>99.8%) [27], ensuring robust identity preservation. CLIP [28] has been demonstrated to possess robust text–image alignment on large-scale datasets, enabling coherent semantic composition. ControlNet [29] has demonstrated improved performance on controllable generation tasks (pose, depth, edge), which is suitable for contextual fidelity in digital mourning.

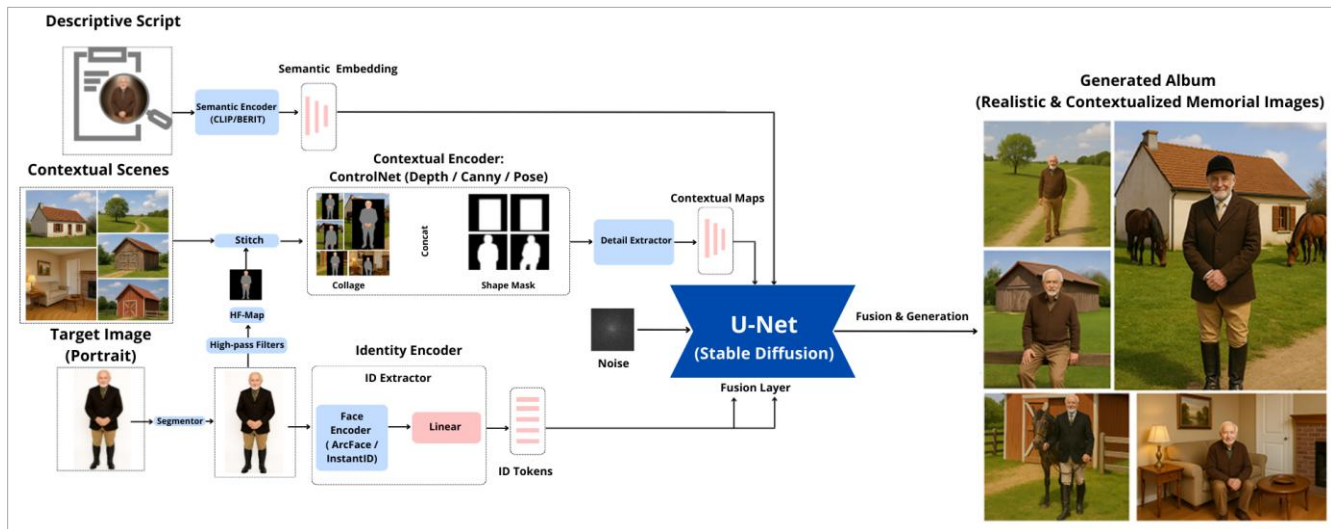


Fig. 1. Architecture of the D.M.A.I.H. approach for personalized digital mourning.

C. Motivation for Choosing Stable Diffusion and U-Net

Several personalization techniques have been explored in the literature:

1) *DreamBooth*: provides strong identity preservation but requires costly fine-tuning for each new subject.

2) *LoRA*: reduces training costs but may sacrifice fine-grained identity fidelity.

3) *Textual inversion*: is lightweight but limited in flexibility and often degrades image quality.

Conversely, Stable Diffusion's U-Net backbone supports [23], the straightforward combination of multiple conditioning sources (identity, context, and text) without full retraining. This balance of fidelity, controllability, and efficiency makes it ideally suited to few-shot personalization in digital bereavement, where personal images are minimal.

It was observed that this is the first such use of Stable Diffusion with few-shot learning [30] in the context of digital mourning, attempting to create photorealistic, contextualized, and emotionally evocative memorial albums.

This is also novel in technique, but it also has a profound psychosocial connotation. By producing coherent and

personalized visual representations of the deceased, the system moves beyond image creation to the design of meaningful artifacts that can play a role in the grieving process.

Unlike static images, which are frozen in time, the designed memorial albums create an active and contextualized redesign of memory, allowing families to reopen contact with their loved ones in emotionally significant spaces.

Such recontextualization may have the effect of consolidating the affective bond with the deceased, conferring comfort and continuity upon the process of grieving. It does indeed support the presentation of grief in a visual and narrative form that is consonant with traditional rituals of commemoration. Such personalized albums can thus function as therapeutic agents, facilitating the process of coming to terms and lessening the sense of loss and lack.

From a psychological perspective, this intervention positions generative AI not only as a creative tool but also as a potential aid for mental health and emotional stability. By enabling families to preserve identity, remember the past, and create new symbolic meaning, the proposed D.M.A.I.H. approach can act as a moral and empathic companion bridging technological innovation with fundamental human values of memory, dignity, and resilience.

D. Comparative Analysis of Methods

The points out that while DreamBooth and LoRA give promising results, they are either computationally expensive or less accurate in terms of maintaining identity. The discussed U-Net-based Stable Diffusion pipeline ensures the best trade-off

between computational expense, identity fidelity, and context flexibility, which renders it particularly well-suited for digital grief purposes. To further illustrate this comparison, Table I provides a structured overview of the main methods in terms of data requirements, training cost, identity preservation, context control, and suitability for digital mourning.

TABLE I. COMPARATIVE ANALYSIS OF GENERATIVE PERSONALIZATION METHODS FOR DIGITAL MOURNING

Method	Data Requirement	Training Cost	Identity Preservation	Control Suitability for Digital Mourning
DreamBooth	5–15 images	High (fine-tuning)	Strong	Limited (costly per subject)
LoRA	5–10 images	Moderate	Moderate	Possible but less precise
Textual Inversion	3–5 images	Low	Weak to Moderate	Not optimal (loss of fidelity)
D.M.A.I.H. approach	5–15 images	Low to Moderate	Strong (0.88–0.93, avg. 0.91)	Highly suitable

As shown in Table I and supported by quantitative validation. (0.88–0.93 identity similarity scores), D.M.A.I.H. achieves greater identity preservation and contextual fidelity than previous approaches.

IV. EXPERIMENTAL RESULTS

To validate the effectiveness of the proposed D.M.A.I.H. approach, an experimental use case representative of the digital mourning process was designed. This case relies on three types of inputs:

- Reference portrait of an elderly woman.
- Contextual scenes representing meaningful environments (family house, farm, interior).
- A descriptive script providing semantic and emotional guidance.

These inputs were integrated into the pipeline, which combines Stable Diffusion with few-shot learning and the U-Net backbone, to generate a personalized memorial album.



Fig. 2. Experimental use case of the D.M.A.I.H. approach.

Fig. 2 illustrates this experimental setup. On the left, the inputs are shown (portrait of the woman, contextual scenes, and descriptive script); in the middle, the Stable Diffusion with few-shot learning process is represented; and on the right, the output is displayed as a five-image memorial album, where the deceased is realistically recontextualized into different meaningful contexts.

The generated images demonstrate the system's ability to preserve identity fidelity while ensuring contextual relevance and emotional resonance, confirming the suitability of the proposed approach for digital mourning applications.

In the experiments, the scalability of the proposed approach was evaluated in generating a large number of personalized images per subject. Specifically, it was possible to produce more than 200 distinct images while maintaining satisfactory identity and contextual fidelity. This result demonstrates the practical scalability of the D.M.A.I.H. approach and its relevance for creating enriched memorial albums. However, this generative capacity remains subject to several technical constraints, including the available GPU memory, the inference time required to produce large batches of images, and the quality of input data, which directly affects the coherence and emotional relevance of the generated outputs. To additionally challenge the

generalizability of D.M.A.I.H., new albums were generated with alternate reference sets, with stable identity preservation and contextual integration. For comparison as a baseline, we applied DreamBooth [31] to the same inputs. While DreamBooth produced visually realistic outputs, it performed worse in terms of contextual fidelity, particularly in grounding descriptive scripts in generated scenes. D.M.A.I.H., by comparison, resulted in higher identity similarity scores (avg. 0.91) and improved contextual coherence.

Our findings are consistent with prior reports of trade-offs between identity fidelity and contextual responsiveness in personalization methods. The results extend these findings by demonstrating that D.M.A.I.H. achieves a more optimal balance, particularly for sensitive applications such as digital mourning.

A. Qualitative Analysis

The output images show that the identity of the subject is preserved across all contexts. The facial features of the individual remain consistent, confirming the effectiveness of the identity extraction module. Correspondingly, the contextual integration is also realistic: the subject is depicted believably in the scenes provided (house, farm, interior), achieving semantic coherence between inputs and outputs. The descriptive script also emphasizes the emotional impact, as it guides the generation to representations that achieve high affective value.

As a result, the album produced is an active visual remembrance rather than static common images. By representing them in several contextualized and emotionally relevant ways, it allows families to symbolically reconnect with deceased loved ones.

B. Quantitative Analysis: Identity Fidelity

In addition to the qualitative analysis, the identity fidelity between the reference image and the generated images of the album was evaluated. For this purpose, A face encoder (ArcFace) was used [24] to extract embeddings and compute the cosine similarity defined as follows:

$$S(I_{ref}, I_{gen}) = \frac{f(I_{ref}) \cdot f(I_{gen})}{\|f(I_{ref})\| \|f(I_{gen})\|}$$

where, $f(I)$ denotes the embedding of an image, I_{ref} is the reference image, and I_{gen} is a generated image. A value close to 1 indicates strong identity preservation.




Fig. 3. Cosine similarity scores of generated images (D.M.A.I.H.).

Fig. 3 reports the scores obtained for the five generated images of the album. Results ranging between 0.88 and 0.93, with an average similarity score of 0.91, further demonstrate consistent identity preservation and confirm that the D.M.A.I.H. approach achieves high identity consistency despite the limited number of input images. These validation measures are necessary to demonstrate robustness. By combining quantitative analysis (identity similarity scores) and qualitative analysis (contextual and semantic coherence), the generated outputs are ensured to be not only visually acceptable but also retain the emotional and memorial value required in digital mourning software.

C. Dual Scope of the D.M.A.I.H. Approach

On the technical level, D.M.A.I.H. introduces, for the first time in the field of digital mourning, the combination of Stable Diffusion with few-shot learning. This synergy enables the creation of photorealistic, contextualized, and highly personalized memorial albums from only a very limited number of reference images. By leveraging generative AI with advanced adaptation techniques, the system overcomes the traditional limitations of data scarcity in sensitive contexts such as mourning.

On the psychosocial level, D.M.A.I.H. transcends the boundaries of image generation to become a medium of moral and psychological support. By offering authentic and emotionally resonant representations, it strengthens the symbolic continuity of grief, consolidates the bond between families and the deceased, and complements rather than replaces traditional mourning rituals.

Overall, this paper presents generative AI both as a driver of technological innovation and as an instrument of psychological resilience, emotional well-being, and mental health. It highlights the possibility of an empathetic and ethically responsible use of AI in supporting families, where advanced technology is balanced with fundamental human values such as dignity, memory, and humanity.

V. DISCUSSION

The results obtained highlight the effectiveness and relevance, both from a technical and a psychosocial perspective.

A. Technical Strengths

The D.M.A.I.H. approach, based on Stable Diffusion with few-shot learning, demonstrated its ability to robustly preserve identity fidelity, even with a very limited number of reference images. Similarity scores (0.88–0.93) confirm high identity consistency across generated outputs. Additionally, contextual scene composition using ControlNet ensures realistic placement of the subject in semantically relevant contexts, and the script description ensures semantic and emotional consistency. Compared to existing personalization methods, such as DreamBooth, LoRA, or Textual Inversion, A more favorable trade-off between visual quality is achieved, flexibility, and computational efficiency, without requiring expensive retraining for each new subject, and illustrates how techniques often linked to deepfake generation can be redirected toward ethical applications.

In particular, compared to DreamBooth [32], Costly fine-tuning was required, and weaker contextual fidelity was observed in the baseline implementation; higher identity similarity scores (average 0.91) were achieved. Similarly, while LoRA [31] and Textual Inversion [33] provide lightweight alternatives, they often suffer from degraded identity fidelity or reduced image quality. D.M.A.I.H. therefore offers a stronger balance of identity preservation, contextual adaptability, and computational efficiency, which is especially relevant in sensitive contexts such as bereavement. New innovations such as Perfusion [34] have attempted to reduce fine-tuning costs but keep personalization potential high; however, they fall short regarding identity preservation. In contrast, the model is stronger and more contextual in design.

B. Comparison with Existing Approaches

Current online mourning services, such as Facebook Memorialized Accounts, Forever Missed, or Legacy.com, just offer static memorials. AI-based initiatives such as HereAfter AI, StoryFile, or Deep Nostalgia focus mainly on interactive dialogue or simple photo animation [35]. While innovative, these tools do not provide a comprehensive generative framework capable of producing personalized, contextualized visual albums. The D.M.A.I.H. approach fills this gap by introducing a pipeline that combines identity preservation, contextual control, and emotional resonance into a single system. Moreover, by explicitly stating similarity of identity and correspondence of context, the analysis extends beyond qualitative judgment to provide quantifiable proof of resilience. This not only justifies the comparative analysis by virtue of being against other existing personalization methods, but it positions D.M.A.I.H. as not merely a theoretical contribution but as an empirically proven framework aligned with existing research.

C. Psychosocial Impact

Beyond technical innovation, the main strength the D.M.A.I.H. approach lies in its role as a potential tool for moral and psychological support. The generated albums allow families to symbolically reconnect with their deceased loved ones, transcending the limits of traditional static photographs [36]. This recontextualization of memory as a dynamic process can be an additional therapeutic tool supporting the mourning process, reinforcing emotional stability, and enabling psychological coping in bereaved families.

D. Limitations and Future Perspectives

Despite the positive outcomes of the D.M.A.I.H. approach, it is important to recognize a number of its limitations. Generalizability is limited because the evaluation is still primarily qualitative and based on a single use case. To evaluate the model's robustness, a larger-scale validation process involving user studies in various cultural contexts will be necessary. The integration of multimodal memories (voice, video, and text), the reinforcement of ethical control mechanisms across the pipeline, and the creation of safeguards to reduce risks like emotional dependency are some future research directions. To guarantee greater empirical validation and ethical robustness, the assessment should also be expanded to larger datasets and carried out in conjunction with psychologists and ethicists.

Similar scalability and identity drift issues also have been encountered in recent personalization studies, validating that the proposed extensions support multimodal integration, adversarial robustness, and stronger ethical guardrails are timely and valid directions.

E. Ethical and Legal Considerations

The use of the D.M.A.I.H. approach in the context of bereavement also raises serious ethical and legal concerns. The issue of consent post-mortem remains on everyone's minds, as current legal regulations provide minimal guarantees. The GDPR and upcoming AI Act for Europe impose transparency and traceability obligations on AI systems but only apply partially to posthumous data across jurisdictions. Such risks include the possibility of inducing false memories, emotional manipulation, and psychological dependence among families. In response to these challenges, the D.M.A.I.H. approach is grounded in an ethical responsibility model, viewing the suggested system as a supportive tool supplementary to the natural grieving process rather than a substitute for it. Future efforts must focus on establishing ethical norms and regulatory frameworks to ensure that generative AI is applied in a thoughtful, empathetic, and culturally sensitive manner.

VI. CONCLUSION

In this paper, the D.M.A.I.H. approach was introduced, a novel pipeline for digital mourning based on Stable Diffusion with few-shot learning, capable of generating personalized, photorealistic, and contextualized memorial albums from a very limited number of reference images. The D.M.A.I.H. approach meets three of the most significant challenges in this area: identity preservation, contextual fidelity, and emotional relevance. Experimental outcomes show that the system proposed herein effectively preserves identity through a change of context, and descriptive scripts ensure semantically and emotionally relevant outputs, illustrating how techniques commonly linked to deepfake generation can be ethically redirected.

Aside from the technical relevance, this study emphasizes the psychosocial significance of the D.M.A.I.H. approach as a tool for moral and psychological attention. Through the dynamic recontextualization of remembrance, the system offers relatives an opportunity to symbolically reunite with their deceased kin, thus complementing everyday acts of memory and encouraging mental health and emotional balance. At the same time, this study identifies important ethical and legal issues, particularly regarding consent, authenticity, and cultural sensitivity. Their resolution will be needed to assure proper use of these technologies.

Unlike existing digital commemoration services that are static or symbolic, D.M.A.I.H. introduces a novel generative pipeline that integrates identity preservation, contextual fidelity, and semantic resonance. This renders the approach a unique contribution to the emerging field of AI-based memorialization. From both an ethical and legal standpoint, further roll-out must be managed by healthcare and psychology ethics committees, based on guidelines such as the GDPR and the EU AI Act. D.M.A.I.H. presents a novel generative pipeline that combines identity preservation, contextual fidelity, and semantic

resonance, in contrast to current digital commemoration services that are static or symbolic. Because of this, the proposed approach offers something special to the developing field of AI-based memorialization. Based on rules like the GDPR and the EU AI Act, ethics committees in the fields of psychology and healthcare must oversee any future roll-out from a moral and legal perspective.

REFERENCES

- [1] B. Acim, H. Ouhnni, N. Kharmoum, S. Ziti, and A. Ouajdouni, "The era of blockchains, deepfakes, and NFTs in therapeutic treatments: A comparative study," in *Non-Fungible Tokens (NFTs) in Smart Cities: Advancements and Security Challenges*, U. Khalil, M. Uddin, O. Malik, and A. Dandoush, Eds. Hershey, PA, USA: IGI Global Scientific Publishing, 2026, pp. 119–152.
- [2] W. Abbaoui, S. Retal, N. Kharmoum, and S. Ziti, "Artificial intelligence at the service of precision medicine," in *Proc. Int. Conf. Advanced Intelligent Systems for Sustainable Development (AI2SD)*, Lecture Notes in Networks and Systems, vol. 713, J. Kacprzyk, M. Ezziyyani, and V. E. Balas, Eds. Cham, Switzerland: Springer, 2023.
- [3] H. Ouhnni, B. Acim, M. Belhiah, K. El Bouchti, Y. Z. Seghroucheni, S. N. Lagmiri, R. Benachir, and S. Ziti, "The evolution of virtual identity: A systematic review of avatar customization technologies and their behavioral effects in VR environments," *Front. Virtual Real.*, vol. 6, 2025.
- [4] B. Acim, N. Kharmoum, M. Ezziyyani, and S. Ziti, "Mental health therapy: A comparative study of generative AI and deepfake technology," in *Proc. Int. Conf. Adv. Intell. Syst. Sustain. Dev. (AI2SD 2024)*, vol. 1403, Lecture Notes in Networks and Systems, M. Ezziyyani, J. Kacprzyk, and V. E. Balas, Eds. Cham, Switzerland: Springer, 2025, pp. 351–357.
- [5] B. Acim, N. Kharmoum, and S. Ziti, "Generative AI in medical education: Personalized learning with deepfakes, LLMs, and RAG," in *The Second International Symposium on Generative AI and Education (ISGAIE'2025)*, R. González Vallejo, G. Moukhliiss, E. Schaeffer, and V. Paliktzoglou, Eds., Lecture Notes on Data Engineering and Communications Technologies, vol. 262. Cham: Springer, 2025, pp. 175–188.
- [6] K. K. Fitzpatrick, A. Darcy, and M. Vierhile, "Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial," *JMIR Mental Health*, vol. 4, no. 2, 2017.
- [7] T. W. Bickmore and R. W. Picard, "Constructing and maintaining long-term human–computer relationships," *ACM Transactions on Computer-Human Interaction*, vol. 12, no. 2, pp. 293–327, 2005.
- [8] W. Abbaoui, S. Retal, B. El Bhiri, N. Kharmoum, and S. Ziti, "Towards revolutionizing precision healthcare: A systematic literature review of artificial intelligence methods in precision medicine," *Informatics in Medicine Unlocked*, vol. 46, 101475, 2024.
- [9] H. Ouhnni, M. Belhiah, S. Ziti, K. El Bouchti, and S. Lagmiri, "Digital afterlife: Challenges and technological innovations in pursuit of immortality," *Indonesian Journal of Electrical Engineering and Computer Science*, vol. 37, no. 2, Feb. 2025.
- [10] J. R. Brubaker, G. R. Hayes, and P. Dourish, "Beyond the grave: Facebook as a site for the expansion of death and mourning," *The Information Society*, vol. 29, no. 3, pp. 152–163, 2013.
- [11] M. Gibbs, J. Meese, M. Arnold, B. Nansen, and M. Carter, "#Funeral and Instagram: Death, social media, and platform vernacular," *Information, Communication & Society*, vol. 18, no. 3, pp. 255–268, 2015.
- [12] H. Ouhnni, B. Acim, M. Belhiah, R. Benachir, and S. Ziti, "A rule-based framework for clothing fit recommendation from 3D body 11 | P a g e www.thesai.org (IJACSA) International Journal of Advanced Computer Science and Applications Vol. XXX, No. XXX, 2024 reconstruction," *Int. J. Adv. Comput. Sci. Appl. (IJACSA)*, vol. 16, no. 6, 2025.
- [13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, et al., "Generative adversarial Networks," *Advances in Neural Information Processing Systems (NeurIPS)*, Montreal, Canada, 2014.
- [14] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, USA, 2019.
- [15] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, USA, 2022.
- [16] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, and K. Aberman, "DreamBooth: Fine-tuning text-to-image diffusion models for subject-driven generation," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [17] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," *arXiv:2106.09685*, 2021.
- [18] J. Ma, J. Gu, Y. Jiang, Y. Zhang, and J. Zhu, "DreamArtist: Towards controllable one-shot text-to-image generation via contrastive prompt-tuning," *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2299–2310, 2023.
- [19] S. Walter, "Digital death: Mortality and beyond in the online age," *Mortality*, vol. 20, no. 3, pp. 259–277, 2015.
- [20] M. Gibbs, J. Meese, M. Arnold, B. Nansen, and M. Carter, "#Funeral and Instagram: Death, social media, and platform vernacular," *Information, Communication & Society*, vol. 18, no. 3, pp. 255–268, 2015.
- [21] K. K. Fitzpatrick, A. Darcy, and M. Vierhile, "Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial," *JMIR Mental Health*, vol. 4, no. 2, 2017.
- [22] T. W. Bickmore and R. W. Picard, "Establishing and maintaining long-term human–computer relationships," *ACM Trans. Comput.-Hum. Interact.*, vol. 12, no. 2, pp. 293–327, 2005.
- [23] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, USA, 2019.
- [24] L. Zhang, A. Rao, and M. Agrawal, "Adding conditional control to text-to-image diffusion models," *arXiv:2302.05543*, 2023.
- [25] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," *arXiv:2103.00020*, 2021.
- [26] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, Minneapolis, MN, USA, 2019.
- [27] A. S. Luccioni, C. Akiki, M. Mitchell et al., "Stable Bias: Analyzing Societal Representations in Diffusion Models," *arXiv:2303.11408*, 2023.
- [28] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu, "CosFace: Large margin cosine loss for deep face recognition," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5265–5274, 2018.
- [29] N. Kumari, B. Zhang, E. Shechtman, and J. Y. Zhu, "Multi-concept customization of text-to-image diffusion," *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1931–1941, 2023.
- [30] L. Xu, H. Xie, S.-Z. J. Qin, X. Tao, and F. L. Wang, "Parameter-Efficient Fine-Tuning Methods for Pretrained Language Models: A Critical Review and Assessment," *arXiv:2312.12148*, 2023.
- [31] R. Gal, Y. Alaluf, Y. Atzmon, O. Patashnik, A. H. Bermano, G. Chechik, and D. Cohen-Or, "An image is worth one word: Personalizing text-to-image generation using textual inversion," *ACM Transactions on Graphics (TOG)*, *Proc. of SIGGRAPH 2023*, vol. 42, no. 4, pp. 1–11, 2023.
- [32] Y. Tewel, R. Gal, G. Chechik, et al., "Key-Locked Rank One Editing for Text-to-Image Personalization," *Proceedings of SIGGRAPH 2023*.

- [33] T. Cilloni, C. Fleming et al., "Privacy Threats in Stable Diffusion Models," arXiv:2311.09355, 2023.
- [34] K. K. Fitzpatrick, A. Darcy, and M. Vierhile, "Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): A randomized controlled trial," *JMIR Mental Health*, vol. 4, no. 2, 2017
- [35] T. W. Bickmore and R. W. Picard, "Establishing and maintaining long-term human-computer relationships," *ACM Transactions on Computer-Human Interaction*, vol. 12, no. 2, pp. 293–327, 2005.
- [36] J. R. Brubaker, G. R. Hayes, and P. Dourish, "Beyond the grave: Facebook as a site for the expansion of death and mourning," *The Information Society*, vol. 29, no. 3, pp. 152–163, 2013