
The Anemoia Device: A Tangible AI System for the Co-creation of Synthetic Memories through Scent

Cyrus Clarke*

Massachusetts Institute of Technology
cyrusc@mit.edu

Nomy Yu

Harvard University
nomy_yu@gsd.harvard.edu

Melo Chen

Harvard University
qiyao_chen@gsd.harvard.edu

Yuen Zou

Harvard University
yuen_zou@gsd.harvard.edu

Hiroshi Ishii

Massachusetts Institute of Technology
ishii@mit.edu



Figure 1: Left: The Anemoia Device. Right: User interaction steps: (1) insert photo, (2) generate caption, (3) prompt via dials, (4) generate story and scent, (5) dispense scent, (6) smell fragrance.

Abstract

1 We present the Anemoia Device, a synthetic memory generator that uses generative
2 AI to provoke nostalgia for a time you have never experienced. The system trans-
3 forms an archival photograph into a multi-sensory artifact through a novel synthesis
4 of vision, language, and olfactory technologies. This multi-modal pipeline uses
5 a tangible, dial-based interface to guide an LLM in performing a cross-modal,
6 semantic-to-olfactory translation. The work is an inquiry into memory malleability
7 in an age of AI, proposing an alternative to conventional screen-based interaction
8 through an intentional, embodied ritual that positions the user as an active co-author,
9 rather than a passive consumer.

*Corresponding author

1 Introduction

Can you feel nostalgia for a past you never lived? What stories might AI be able to tell you about that past, and more uniquely, what might that story smell like? The Anemoia Device explores this possibility by co-authoring synthetic memories with generative AI models. It is built on the proposition that a tangible, multi-sensory experience, through the production of immersive narratives and corresponding scents, can help forge an emotional connection to an unlived past, incorporating an external memory into a user’s embodied reality.

The term “anemoia” is a neologism, meaning ‘nostalgia for a time you never experienced’[1]. This form of nostalgia aligns with contemporary theories of memory that argue our cognitive recall includes the external world, from Cloud infrastructure to the memory of others [14, 33]. This is especially evident in an era where we increasingly externalise our memories in vast digital archives where they are exposed to algorithmic influence [13].

Herein lies a critical issue. As we treat our digital archives not just as files or photographs, but as reliable proxies for our memories, the ever-growing capabilities of AI to author synthetic memories is a growing risk [32]. This technological capability forces us to confront the fact that human memory is not a static archive, but a malleable and reconstructive process susceptible to sensory suggestion[27]. The dominant trajectory of generative AI development often prioritises photorealism and persuasive power, a path that risks exploiting our cognitive vulnerabilities for persuasive ends [31], potentially eroding collective trust and devaluing authentic experience.

In response, we argue for an alternative direction, aligned with poetic exploration and generative empathy. The challenge is not simply to replicate or re-author the past, but to create new co-authored narratives that expand our imaginative and emotional worlds. This necessitates a shift in design values: from speed and scale to slowness and reflection. From black-box persuasion to tangible, user-guided collaboration. From purely visual output to multi-sensory, embodied experiences.

To provide this, we utilise two underexplored modalities in AI interaction: olfaction and tangible interfaces. Olfaction offers a unique and powerful ability to evoke vivid memories [11, 20], yet technical challenges have limited its application in computational systems [36]. However, recent research has demonstrated that Large Language Models (LLMs) can approximate human-scent perceptual judgments from text [24]. Our system builds on this approach to provide a semantic foundation for olfactory experiences, linking the associative space of language with the evocative world of odorants. Understanding that embodied interactions would enhance the subjectivity of the experience [38, 41], our approach grounds the system in physical acts from inserting analogue photographs to physically manipulating dials.

In this manner, the Anemoia Device transforms the open-ended and immaterial act of prompting into a structured, co-creative ritual. The system synthesizes and orchestrates generative models to transmute an impersonal, external memory into a more subjectively significant, co-authored experience. The result is a multi-sensory artifact, a generated story paired with a dedicated fragrance, dispersed through an integrated olfactory system, connecting people to pasts and memories they have never lived.

2 Related works

This work is situated at the intersection of four key domains: contemporary theories of memory; critical work on the role of AI in generating synthetic memories; foundational research on the link between olfaction and memory; and finally principles from tangible and embodied interaction.

2.1 Extended memory

While traditional models frame memory as a brain-bound process based on lived experience [7], our work builds on contemporary theories of extended memory that explain how memory is increasingly stored and accessed via external media such as digital archives [14, 38, 41]. This notion of extended memory is not limited to technological artifacts, demonstrated by the existence of vicarious memories, where second-hand accounts start to share phenomenological and functional properties with personal memories [33, 39]. Our work is situated in this conceptual space, exploring how anemoia can be understood as a technologically mediated vicarious memory, co-authored with AI.

61 2.2 Synthetic memory and malleability

62 Contrary to the idea of memory as a perfect recording of the past, research shows it is a reconstructive
63 process [37], where misinformation and falsified content can lead people to form rich false memories
64 [27]. As our cognitive processes actively rebuild memories rather than simply replaying them, our
65 cognitive processes are malleable and susceptible to external influence [28]. It is this malleability
66 that makes synthetic memories, recollections of events that either never occurred or are dramatically
67 distorted, a psychological reality [22]. This vulnerability has been demonstrated to significantly
68 increase with exposure to sensory media [30].

69 2.3 Olfaction: affective power and computational challenges

70 Olfaction is known to evoke vivid memories, this phenomenon famously known as the Proust Effect
71 [35], stems from the direct pathways between the olfactory system and the neurological centres for
72 emotion and memory [12, 21, 25]. Leveraging this power by integrating scent into computational and
73 AI systems presents significant challenges due to difficulties with digitisation and high interpersonal
74 variability in perception [36, 40]. Machine learning research in this domain has primarily focused
75 on predicting odour properties from molecular structures [2–5, 17, 26] with notable work from
76 Pyrume to build a comprehensive foundational data infrastructure [19], while recent research has
77 demonstrated that LLMs can represent human olfactory information from textual data alone [24].
78 Interactive systems meanwhile have largely been limited to pre-defined mappings, where a specific
79 input triggers a single, pre-loaded scent [6, 10, 15]. Our project addresses this gap, exploring dynamic
80 scent mappings that are conceptually linked to AI-generated narratives.

81 2.4 Tangible and embodied interaction

82 It is well established that tangible interfaces can enhance cognition and create more meaningful
83 experiences compared with abstract, detached information processing [16, 23]. Contemporary
84 theories around extended cognition help to explain this, proposing that external information requires
85 a process of embodiment and incorporation into the body (*in corpore*) to become a meaningful part
86 of our cognitive experience [14, 38, 41]. This suggests an interface designed for memory-making
87 would benefit from a physical and sensory led approach. Our tangible interface is guided by these
88 principles, with the olfactory output completing the process of bodily incorporation.

89 3 Anemoia Device

90 The Anemoia Device is a tangible system that implements a novel multi-modal pipeline for generating
91 olfactory experiences. The architecture is designed to support generative world-building, where a
92 user transforms an archival photograph into a multi-sensory artifact. This is achieved by integrating
93 a Vision-Language Model (VLM), a tangible interface for user guidance, and LLMs for narrative
94 generation and cross-modal translation from semantics to a scent-vector space.

95 3.1 System architecture

96 The full pipeline as illustrated in Figure 2, proceeds in six key stages: (1) VLM-based image
97 captioning of a photographic input. (2) LLM-based subject classification based on the caption. (3)
98 The user’s low-dimensional inputs via three dials are processed. (4) A final narrative is synthesized by
99 an LLM from the user inputs. (5) A scent formula is generated via a few-shot, cross-modal translation
100 from the narrative. (6) Formula quantities are calculated into pump timings and the final scent output
101 is rendered by the integrated olfactory display.

102 3.2 User interaction

103 The physical form of the device is based on the metaphor of distillation, organized vertically from
104 top to bottom (See Figure 3) where a memory is gradually processed and concentrated into a scent
105 through a multi-sensory ritual. User interaction is anchored in the physical world, beginning with the
106 insertion of a physical photograph.

107 The core of the interaction is the user-guided narrative construction using three physical rotary dials.
108 This tangible interface provides a constrained, structural and exploratory experience, compared with

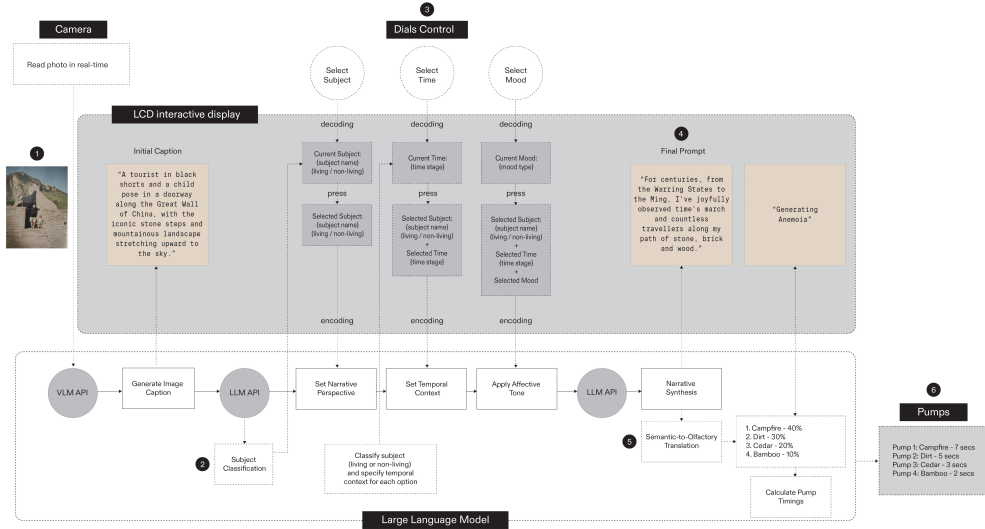


Figure 2: Overview of system architecture from image captioning to scent dispersal

a typical open-ended prompting process. The three dials are designed to align with the concept of anemoia by allowing the user to define the parameters of an imagined past:

- Dial 1 - Subject (perspective): enables the user to inhabit an unfamiliar perspective within the scene by selecting the point-of-view of a subject (e.g. "old man," "tree," "bicycle") based on the VLM-LLM subject classification from the photograph. As part of this initial classification, the LLM determines whether the subject is living or non-living, which directly informs the options available on the next dial.
- Dial 2 - Time (temporal context): situates the selected perspective within a timeline. Options are dynamically filtered. If a living subject is selected, the options are ["childhood", "youth", "adulthood", and "elderly"]. If a non-living subject is chosen the options are ["raw material", "manufacture", "in-use", and "decay"]. This ensures conceptual coherence.
- Dial 3 - Mood (affective tone): allows the user to infuse the synthetic memory with a specific emotional tone from a fixed set of options: "happy," "sad", "calm", "angry".

3.3 LLM semantic-to-olfactory translation

As outlined in Section 3.1, the VLM image caption and the user's three conceptual parameters are synthesized into a structured prompt for an LLM (GPT-4o). The LLM's first task is to generate a short, poetic narrative that integrates these multi-modal inputs. This step transforms the sparse user input into a rich, high-dimensional semantic context.

The LLM's second, more novel task is to perform a few-shot, cross-modal translation from the generated narrative to an olfactory output. The foundation of this translation process is a scent library that defines the target olfactory space for the model. This library consists of our ($N = 39$) available pre-blended fragrance oils. Rather than relying on simple one-word labels, each scent is annotated with a set of semantic descriptors. These descriptors add details such as primary olfactory notes, associated concepts, and emotions to help translate fragrances into a machine-readable semantic space (see Appendix A.4 for details).

To enhance this process, we constructed an Olfactory Knowledge Base containing example pairings of narratives, their ideal scent vectors based on scent profiles, and a textual rationale explaining the semantic mapping. For each generation, we use in-context learning: 2-3 of these examples are randomly selected and prepended to the prompt, effectively teaching the model the desired reasoning process in real-time. The model is then prompted to analyse its own newly generated narrative and following the provided examples, output a proportional recipe of up to four scents that best represents




Figure 3: The Anemoia Device is vertically structured with three functional zones. The left schematic highlights the internal structure and modular segmentation, the right image shows the final prototype.

the story’s emotional and thematic content. An example is provided in Appendix A.5. This output can be considered a sparse vector in the N-dimensional space of our available scents, and this few-shot approach yields more consistent and conceptually grounded results than a purely zero-shot method.

3.4 Olfactory interface

The digital scent vector is rendered into a physical fragrance by a custom olfactory display. This subsystem is controlled by a XIAO ESP32S3 microcontroller, which receives commands from the host system. Upon reception of a scent formula, the host application calculates the precise duration each pump must be activated to dispense the appropriate proportion of fragrance oil. These timing commands are sent to the microcontroller and relays to control four 12V DC peristaltic pumps, capable of dispensing fragrance oils with millilitre accuracy. The four pumps draw from separate glass vials and dispense their liquids through silicone tubing into a shared blending vessel. This vessel, containing the final composed scent, can then be removed by the user, completing the pipeline from abstract prompt to physical, olfactory experience.

Additional details of hardware and software implementation as well as notes on fabrication are provided in the Appendix.

4 Discussion

4.1 Co-authoring synthetic memories

By co-authoring synthetic memories through a tangible, olfactory driven process, the Anemoia Device moves beyond conventional screen-based interactions to explore a more embodied and poetic form of human-AI collaboration. This work proposes a fundamentally different archetype for our relationship with generative machines that prioritises intentionality, embodiment and co-creation.

This stands in contrast to the dominant trajectory of AI applications, which often centre on the creation of photorealistic fictions designed to persuade. This approach, reinforced by frictionless, screen-based interfaces conducive to passive consumption, places the user in a position where their cognitive vulnerabilities can be exploited. While the device leverages similar ‘black box’ language models to everyday consumer products, it wraps them in an overtly mechanical object designed for transparency and intentionality. The physical nature of interacting with the device takes the user from

167 a passive consumer to an active co-author. The resulting scent and story are not a deception to be
168 believed, but the poetic outcome of a co-creative act. This approach suggests a more humanistic path
169 for synthetic media, where rather than attempt to blur the line with reality, AI models are used to
170 create beautiful, honest fictions that we can hold, smell, and reflect upon.

171 4.2 Limitations and future directions

172 As a speculative research artifact, this work has several limitations that provide opportunities for
173 future inquiry. This paper presents the system and its conceptual framework without a formal user
174 study to validate its capacity to provoke anemoia. The olfactory display uses a curated array of 39
175 pre-blended fragrance oils with a four-scent output array. While designed to help reduce perceptual
176 blending issues [8], this limits the breadth and precision of olfactory experiences the device can offer.
177 Finally our current semantic-to-olfactory translation uses a one-directional semantic approach which
178 has no model of underlying chemical structures or perceptual interactions between odorants. This
179 model in turn cannot benefit from user-feedback. Future work will seek to address these limitations.
180 Firstly, conducting empirical validation of the system to gather data on user experience. Recognising
181 the highly subjective nature of scent perception, a key goal is to create a more dynamic user-adaptive
182 system that incorporates feedback into a reinforcement learning model, helping to fine-tune the
183 LLM’s semantic mapping. This combined with a more comprehensive fragrance library would help
184 to create a system that is a powerful, personalisable scent engine capable of evoking the emotional
185 texture of any narrative.

186 5 Conclusion

187 By designing a system that co-authors synthetic, olfactory memories, this project demonstrates a
188 novel pathway for human-AI interaction that is intentional, embodied, and poetic. We have presented
189 a multi-modal pipeline that translates an external visual memory into a multi-sensory artifact, using a
190 tangible interface for user guidance and LLMs for cross-modal generation from narrative semantics
191 to a physical scent. A key contribution of this work lies in its novel synthesis of vision, language,
192 and olfactory technologies to create a generative system. This demonstrates how AI tools can be
193 orchestrated to produce rich, affective experiences that engage the challenging domain of olfaction
194 and points towards a future where we design not just for individual AI modalities, but for the complex,
195 multi-sensory systems of our embodied humanity.

196 Beyond a technical system, this work proposes an alternative archetype for our relationship with
197 generative machines. In a landscape often focused on the speed, scale, and persuasive power of
198 AI, our approach promotes a different set of values: transparency, co-creation, and reflection. This
199 suggests a trajectory for creative AI where technology is used not to replicate or deceive, but to
200 expand our capacity for imagination and empathy. Ultimately, the Anemoia Device is a provocation
201 that asks what it means to remember when memory itself can be generated, what it means to feel
202 when that feeling is co-authored with a machine, and what it means to be human when we can craft
203 beautiful, fragrant fictions of pasts we have never lived.

204 6 Acknowledgments

205 This project was initially developed during the How to Make (Almost) Anything class taught by
206 Professor Neil Gershenfeld and the Tangible Interfaces class taught by Professor Hiroshi Ishii at MIT.
207 We also thank Isabela Sanchez Taípe for her contributions to the early research and development of
208 this work, and Dr. Ayan Abukar for her analytical and creative support.

209 References

- 210 [1] Anemoia | The Dictionary of Obscure Sorrows. <https://www.thedictionaryofobscuresorrows.com/concept/anemoia>.
- 211 [2] Harnessing graph neural networks to craft fragrances based on consumer feedback. *Computers &*
212 *Chemical Engineering*, 185:108674, June 2024. ISSN 0098-1354. doi: 10.1016/j.compchemeng.
213 2024.108674.

- 214 [3] Rayane Achebouché, Anne Tromelin, Karine Audouze, and Olivier Taboureau. Application
215 of artificial intelligence to decode the relationships between smell, olfactory receptors and
216 small molecules. *Scientific Reports*, 12(1):18817, November 2022. ISSN 2045-2322. doi:
217 10.1038/s41598-022-23176-y.
- 218 [4] Manuel Aleixandre, Dani Prasetyawan, and Takamichi Nakamoto. Automatic scent creation by
219 cheminformatics method. *Scientific Reports*, 14(1):31284, December 2024. ISSN 2045-2322.
220 doi: 10.1038/s41598-024-82654-7.
- 221 [5] Manuel Aleixandre, Dani Prasetyawan, and Takamichi Nakamoto. Generative Diffusion Net-
222 work for Creating Scents. *IEEE Access*, 13:57311–57321, 2025. ISSN 2169-3536. doi:
223 10.1109/ACCESS.2025.3555273.
- 224 [6] Judith Amores and Pattie Maes. Essence: Olfactory Interfaces for Unconscious Influence of
225 Mood and Cognitive Performance. In *Proceedings of the 2017 CHI Conference on Human*
226 *Factors in Computing Systems*, pages 28–34, Denver Colorado USA, May 2017. ACM. ISBN
227 978-1-4503-4655-9. doi: 10.1145/3025453.3026004.
- 228 [7] R. C. Atkinson and R. M. Shiffrin. Human Memory: A Proposed System and its Control
229 Processes1. In Kenneth W. Spence and Janet Taylor Spence, editors, *Psychology of Learning*
230 *and Motivation*, volume 2, pages 89–195. Academic Press, January 1968. doi: 10.1016/
231 S0079-7421(08)60422-3.
- 232 [8] S. Barkat, E. Le Berre, G. Coureaud, G. Sicard, and T. Thomas-Danguin. Perceptual blending
233 in odor mixtures depends on the nature of odorants and human olfactory expertise. *Chemical*
234 *Senses*, 37(2):159–166, February 2012. ISSN 1464-3553. doi: 10.1093/chemse/bjr086.
- 235 [9] Quentin Bolsee. gerber2img, n.d. URL [https://quentinbolsee.pages.cba.mit.edu/](https://quentinbolsee.pages.cba.mit.edu/gerber2img/)
236 [gerber2img/](https://quentinbolsee.pages.cba.mit.edu/gerber2img/).
- 237 [10] Stephen Brewster, David McGookin, and Christopher Miller. Olfoto: Designing a smell-based
238 interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing*
239 *Systems*, CHI '06, pages 653–662, New York, NY, USA, April 2006. Association for Computing
240 Machinery. ISBN 978-1-59593-372-0. doi: 10.1145/1124772.1124869.
- 241 [11] Wei-Yu Chen, Chu-Ching Huang, Tzu-Min Lo, Wei-Chun Yuan, Hsin-Yu Tsai, Ting-Fang
242 Cheng, and Fang-Ling Chang. The Research on Scent and Fragrance in Memory with Machine
243 Learning. In *2020 IEEE International Conference on Consumer Electronics - Taiwan (ICCE-*
244 *Taiwan)*, pages 1–2, September 2020. doi: 10.1109/ICCE-Taiwan49838.2020.9257990.
- 245 [12] Simon Chu and John Joseph Downes. Long live Proust: The odour-cued autobiographical
246 memory bump. *Cognition*, 75(2):B41–B50, May 2000. ISSN 0010-0277. doi: 10.1016/
247 S0010-0277(00)00065-2.
- 248 [13] Andy Clark. Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence.
249 January 2004.
- 250 [14] Andy Clark and David Chalmers. The Extended Mind. *Analysis*, 58(1):7–19, 1998. ISSN
251 0003-2638.
- 252 [15] David Dobbelsstein, Steffen Herrdum, and Enrico Rukzio. inScent: A wearable olfactory
253 display as an amplification for mobile notifications. In *Proceedings of the 2017 ACM Inter-*
254 *national Symposium on Wearable Computers*, ISWC '17, pages 130–137, New York, NY,
255 USA, September 2017. Association for Computing Machinery. ISBN 978-1-4503-5188-1. doi:
256 10.1145/3123021.3123035.
- 257 [16] Paul Dourish. *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press, 2001.
258 ISBN 978-0-262-54178-7.
- 259 [17] Dewei Feng, Carol Li, Wei Dai, and Paul Pu Liang. SMELLNET: A Large-scale Dataset for
260 Real-world Smell Recognition, May 2025.
- 261 [18] Neil Gershenfeld. mods CE. URL <https://modsproject.org>.

- [19] Elizabeth A. Hamel, Jason B. Castro, Travis J. Gould, Robert Pellegrino, Zhiwei Liang, Liyah A. Coleman, Famesh Patel, Derek S. Wallace, Tanushri Bhatnagar, Joel D. Mainland, and Richard C. Gerkin. Pyrfume: A window to the world's olfactory data. *Scientific Data*, 11(1):1220, November 2024. ISSN 2052-4463. doi: 10.1038/s41597-024-04051-z.
- [20] Rachel S. Herz and Jonathan W. Schooler. A naturalistic study of autobiographical memories evoked by olfactory and visual cues: Testing the Proustian hypothesis. *The American Journal of Psychology*, 115(1):21–32, 2002. ISSN 0002-9556.
- [21] Rachel S Herz, James Eliassen, Sophia Beland, and Timothy Souza. Neuroimaging evidence for the emotional potency of odor-evoked memory. *Neuropsychologia*, 42(3):371–378, January 2004. ISSN 0028-3932. doi: 10.1016/j.neuropsychologia.2003.08.009.
- [22] Tetsuya Iidaka, Tokiko Harada, and Norihiro Sadato. False memory for face in short-term memory and neural activity in human amygdala. *Brain Research*, 1591:74–85, December 2014. ISSN 0006-8993. doi: 10.1016/j.brainres.2014.10.003.
- [23] Hiroshi Ishii and Brygg Ullmer. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, CHI '97, pages 234–241, New York, NY, USA, March 1997. Association for Computing Machinery. ISBN 978-0-89791-802-2. doi: 10.1145/258549.258715.
- [24] Murathan Kurfalı, Pawel Herman, Stephen Pierzchajło, Jonas Olofsson, and Thomas Hörberg. Representations of smells: The next frontier for language models? *Cognition*, 264:106243, November 2025. ISSN 0010-0277. doi: 10.1016/j.cognition.2025.106243.
- [25] Maria Larsson, Johan Willander, Kristina Karlsson, and Artin Arshamian. Olfactory LOVER: Behavioral and neural correlates of autobiographical odor memory. *Frontiers in Psychology*, 5, April 2014. ISSN 1664-1078. doi: 10.3389/fpsyg.2014.00312.
- [26] Brian K. Lee, Emily J. Mayhew, Benjamin Sanchez-Lengeling, Jennifer N. Wei, Wesley W. Qian, Kelsie A. Little, Matthew Andres, Britney B. Nguyen, Theresa Moloy, Jacob Yasonik, Jane K. Parker, Richard C. Gerkin, Joel D. Mainland, and Alexander B. Wiltschko. A principal odor map unifies diverse tasks in olfactory perception. *Science*, 381(6661):999–1006, September 2023. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.ade4401.
- [27] Elizabeth F. Loftus. Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory*, 12(4):361–366, January 2005. ISSN 1072-0502-01, 1549-5485. doi: 10.1101/lm.94705.
- [28] Elizabeth F. Loftus and Jacqueline E. Pickrell. The Formation of False Memories. *Psychiatric Annals*, 25(12):720–725, December 1995. doi: 10.3928/0048-5713-19951201-07.
- [29] Odeuropa. Odeuropa smell explorer. <https://explorer.odeuropa.eu/>.
- [30] Henry Otgaar, Mark L. Howe, and Lawrence Patihis. What science tells us about false and repressed memories. *Memory*, 30(1):16–21, January 2022. ISSN 0965-8211. doi: 10.1080/09658211.2020.1870699.
- [31] Pat Pataranutaporn, Chayapatr Archiwaranguprok, Samantha W. T. Chan, Elizabeth Loftus, and Pattie Maes. Synthetic Human Memories: AI-Edited Images and Videos Can Implant False Memories and Distort Recollection, September 2024.
- [32] Elizabeth A. Phelps and Stefan G. Hofmann. Memory editing from science fiction to clinical practice. *Nature*, 572(7767):43–50, August 2019. ISSN 1476-4687. doi: 10.1038/s41586-019-1433-7.
- [33] David B. Pillemer, Kristina L. Steiner, Kie J. Kuwabara, Dorthe Kirkegaard Thomsen, and Connie Svob. Vicarious memories. *Consciousness and Cognition*, 36:233–245, November 2015. ISSN 1053-8100. doi: 10.1016/j.concog.2015.06.010.
- [34] P&J Trading. Scented oils collection. <https://www.pandjtrading.com/>.
- [35] Marcel Proust. *Du côté de chez Swann*. Gaillimard, Paris, 1919.

- [36] Charles Spence, Marianna Obrist, Carlos Velasco, and Nimesha Ranasinghe. Digitizing the chemical senses: Possibilities & pitfalls. *International Journal of Human-Computer Studies*, 107:62–74, November 2017. ISSN 1071-5819. doi: 10.1016/j.ijhcs.2017.06.003.
- [37] Eleanor Spens and Neil Burgess. A generative model of memory construction and consolidation. *Nature Human Behaviour*, 8(3):526–543, March 2024. ISSN 2397-3374. doi: 10.1038/s41562-023-01799-z.
- [38] John Sutton. Exograms and Interdisciplinarity: History, the Extended Mind, and the Civilizing Process. In *The Extended Mind*, pages 189–225. The MIT Press, May 2010. doi: 10.7551/mitpress/9780262014038.003.0009.
- [39] Zeno Vendler. Vicarious Experience. *Revue de Métaphysique et de Morale*, 84(2):161–173, 1979. ISSN 0035-1571.
- [40] Yu Zhang, Peizhong Gao, Fangzhou Kang, Jiaxiang Li, Jiacheng Liu, Qi Lu, and Yingqing Xu. OdorAgent: Generate Odor Sequences for Movies Based on Large Language Model. In *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, pages 105–114, March 2024. doi: 10.1109/VR58804.2024.00034.
- [41] Gregorio Zlotnik and Aaron Vansintjan. Memory: An Extended Definition. *Frontiers in Psychology*, 10:2523, November 2019. ISSN 1664-1078. doi: 10.3389/fpsyg.2019.02523.

A Technical Appendices and Supplementary Material

A.1 Software implementation

The software architecture includes firmware on the Arduino Uno and XIAO ESP32S3, the central Python host application, and external cloud AI services. A central host application (Python 3.10+) orchestrates the workflow, managing separate serial connections (`pyserial`) to the Uno (input/feedback) and XIAO (output commands), handles image encoding, makes API calls (`requests` to Anthropic/OpenAI), manages state, generates dynamic options, and performs deterministic pump timing calculations (ref. Algorithm 2).

Input controller firmware (Arduino C++ on Uno) reads encoders/buttons, updates the LCD (`LiquidCrystal_I2C` library), sends confirmed selections, and displays received dynamic options (ref. Algorithm 1). Output controller firmware (Arduino C++ on Xiao ESP32S3) is dedicated to receiving pump commands via serial and precisely controlling the relay modules using non-blocking timers (`millis()`), implementing the execution phase of Algorithm 2.

VLM/LLM interaction and prompting use a structured, multi-stage prompting strategy. First, for initial image captioning, Anthropic’s API (Claude Sonnet 3.7) is invoked to generate a concise textual description (max. 145 characters) from the webcam’s base64-encoded image. All subsequent generative tasks leverage OpenAI’s GPT API (model GPT-4o). Firstly, the LLM synthesizes inputs based on the user’s dial selections to generate the final narrative. This serves as the prompt for semantic-to-olfactory translation performed using a few-shot, in-context learning approach. This involves dynamically pre-pending the prompt with examples drawn from the curated olfactory knowledge base, which contains narrative-to-scent vector pairings and rationales. This method guides the LLM to map its own generated narrative to a proportional scent formula, utilizing format constraints and embedded knowledge.

Tangible prompting logic is described in Algorithm 1, detailing the user interaction flow for parameter configuration and the dynamic feedback loop between the user, input controller (Arduino) and host application. Narrative-to-scent logic, shown in Algorithm 2, outlines the end-to-end process from narrative input to scent output, combining the generation of the LLM formula, the calculation of the host time and the execution of the output controller (XIAO / relay). Example outputs are illustrated in Figure 4.

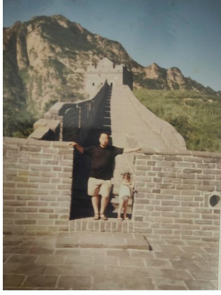
	Objective Caption	Extracted Subjects	Selected Subject	Select Time	Select Mood	Final Prompt	Scent Formula
	'A tourist in black shorts and a child pose in a doorway along the Great Wall of China, with the iconic stone steps and mountainous landscape stretching upward to the sky.'	0: Tourist (living) 1: Dog (living) 2: Great Wall of China (non-living) 3: Mountains (non-living)	Great Wall of China (non-living)	Raw-material Manufacture In-use Decay	Happy Sad Angry Calm	For centuries, from the Warring States to the Ming, I've joyfully observed time's march and countless travellers along my path of stone, brick, and wood.	Campfire - 40% Dirt - 30% Cedar - 20% Bamboo - 10%
			Tourist (living)	Childhood Youth Adulthood Elderly	Happy Sad Angry Calm	Where ancient guardians once watched, I feel time stretch like the Wall itself. My childish wonder mirrors my past, our souls becoming part of this stone tapestry.	Moss - 35% Night Air - 30% Cedar - 25% Green Tea - 10%

Figure 4: Example outputs demonstrating narrative and scent generation for different user inputs (Subject, Time, Mood) based on a sample image caption.

Algorithm 1 Tangible prompting interaction logic

```

Initialize Phase
1: Host: Get VLM caption & LLM subject classifications.
2: Host → Arduino Uno (Serial): Send initial Subject options (OPTS:SUBJ:...).
3: Arduino Uno: Display Subject options on LCD.
4: activeParam ← SUBJECT.
Interaction Loop Phase
5: loop
6:   Arduino Uno: Monitor Encoders/Buttons.
7:   if Rotation detected for activeParam then
8:     Arduino Uno: Update internal selection focus & LCD display.
9:   else if Button Press detected for activeParam then
10:    selectedValue ← Focused option on Arduino Uno.
11:    Arduino Uno → Host (Serial): Send confirmation (CONF:PARAM:Value).
12:    Host Processing: Parse confirmation, store selectedValue.
13:    if activeParam = SUBJECT then
14:      Host: Determine dynamic Period options based on tag.
15:      Host → Arduino Uno (Serial): Send Period options (OPTS:PERD:...).
16:      activeParam ← PERIOD.
17:      Arduino Uno: Update internal list for Period dial.
18:    else if activeParam = PERIOD then
19:      Host: Determine fixed Mood options.
20:      Host → Arduino Uno (Serial): Send Mood options (OPTS:MOOD:...).
21:      activeParam ← MOOD.
22:      Arduino Uno: Update internal list for Mood dial.
23:    else if activeParam = MOOD then
24:      Exit Loop.
25:      Arduino Uno: Update LCD for new activeParam.
Post-Loop: Host proceeds to Narrative Synthesis.

```

Algorithm 2 Narrative-to-Scent generation logic

Require: Final narrative_sentence, scent_library with descriptors, olfactory_knowledge_base, pump_assignments dict, MAX_PUMP_DURATION_S.

Ensure: Physical scent emission.

Host Application - Few-Shot Formula Generation

```
1:  $k \leftarrow 3$  ▷ Number of examples for in-context learning
2: selected_examples  $\leftarrow$  RANDOMLYSELECT(olfactory_knowledge_base, k)
3: formatted_examples  $\leftarrow$  FORMATEXAMPLES(selected_examples) ▷ Formats Narrative, Rationale, and Recipe
4: prompt  $\leftarrow$  CONSTRUCTPROMPT(system_instruction, formatted_examples, scent_library, narrative_sentence)
5: formula_string  $\leftarrow$  LLM_API(prompt, model="GPT-4o")
6: formula_dict  $\leftarrow$  VALIDATEANDPARSE(formula_string) ▷ Parse into {name: percentage}
```

Host Application - Timing Calculation

▷ Deterministic

```
7: commands  $\leftarrow$  []
8: if IsEmpty(formula_dict) then return ▷ Handle error

9: proportions  $\leftarrow$  VALUES(formula_dict)
10: max_proportion  $\leftarrow$  MAX(proportions)
11: if max_proportion  $\leq 0$  then return
12: scale_factor  $\leftarrow$  MAX_PUMP_DURATION_S / max_proportion
13: for all name, proportion in formula_dict do
14:   duration  $\leftarrow$  proportion * scale_factor
15:   pump_index  $\leftarrow$  GET(pump_assignments, name, default=-1)
16:   if pump_index  $\neq -1$  then
17:     cmd_string  $\leftarrow$  FORMAT("pump{pump_index}:{duration:.2f}")
18:     ADD(cmd_string, commands)
```

Host Application - Dispensing Trigger

```
19: Wait for User Trigger signal from tangible interface.
20: for all cmd in commands do
21:   Host  $\rightarrow$  Xiao ESP32S3 (Serial): Send cmd + '
22:   n'.
23:   DELAY(50ms) ▷ Short pause between commands
```

Output Controller (Xiao ESP32S3) - Execution

```
24: Loop:
25:   If Serial data ("pumpX:Y.YY") received Then
26:     pumpIndex, duration  $\leftarrow$  PARSECOMMAND(command).
27:     Set target end time: pumpEndMillis[pumpIndex]  $\leftarrow$  millis() + duration*1000.
28:     Activate relay: digitalWrite(relayPins[pumpIndex], HIGH).
29:   End If
30:   CHECKPUMPTIMERS ▷ Deactivates relays when millis()  $\geq$  pumpEndMillis
31: End Loop
```

356 **A.2 Hardware implementation**

357 The hardware system comprises distinct input and output modules connected to a central host
358 computer, alongside the necessary power and fluidic components. Key components and their inter-
359 connections are shown in Figure 5.

360 The input module uses a USB webcam to capture images. Three KY-040 rotary encoders and an I2C
361 20x4 LCD screen form the tangible interface, connected to the Arduino Uno (see Figure 6). The
362 output module consists of four 12V DC peristaltic pumps used to dispense fragrance oils. Each pump
363 is switched by a separate 12V relay module. The logic inputs of the relays are controlled by a Seeed
364 Studio XIAO ESP32S3 (Output Controller), which is mounted on a custom PCB. Figure 7 shows a
365 simplified representation. Power distribution is handled by powering the Arduino and XIAO via USB,




Figure 5: Core hardware components used in the Anemoia Device: (A) peristaltic pumps (x4), (B) 12V relay modules (x4), (C) arduino uno (input controller), (D) LCD screen module, (E) custom PCB with XIAO ESP32S3 (output controller), (F) KY-040 rotary encoder modules (x3), (H) USB webcam.

366 while an external 12V DC power supply powers the pumps, switched via the relay contacts. The relay
 367 modules' logic coils (VCC) are powered by the XIAO's 5V output. Fluidics are implemented using
 368 glass reservoirs, connected via silicone tubing through the pumps to a passive mixing point.

369 A.3 Fabrication

370 The device enclosure was constructed using a modular aluminium rod frame with laser-cut frosted
 371 acrylic panels. Key structural joints and component mounts were 3D printed (PLA), creating a layered
 372 internal layout. A custom PCB, designed in Fusion 360 Electronics, processed using gerber2img[9]
 373 and mods [18], and CNC-milled on a Carvera Desktop CNC machine, integrates the Seeed Studio
 374 XIAO ESP32S3 with connectors for four relay modules. This compact board simplifies wiring,
 375 improves reliability, and provides centralised control over the peristaltic pumps. To minimise cross-
 376 interference, the internal layout physically separates logical subsystems (e.g. microcontrollers, relays)
 377 from fluid-handling components (pumps, tubing, reservoirs), each mounted on distinct internal panels
 378 within the frame.




Figure 6: Input module schematic: shows connections between Arduino Uno microcontroller, three rotary encoders and LCD Screen

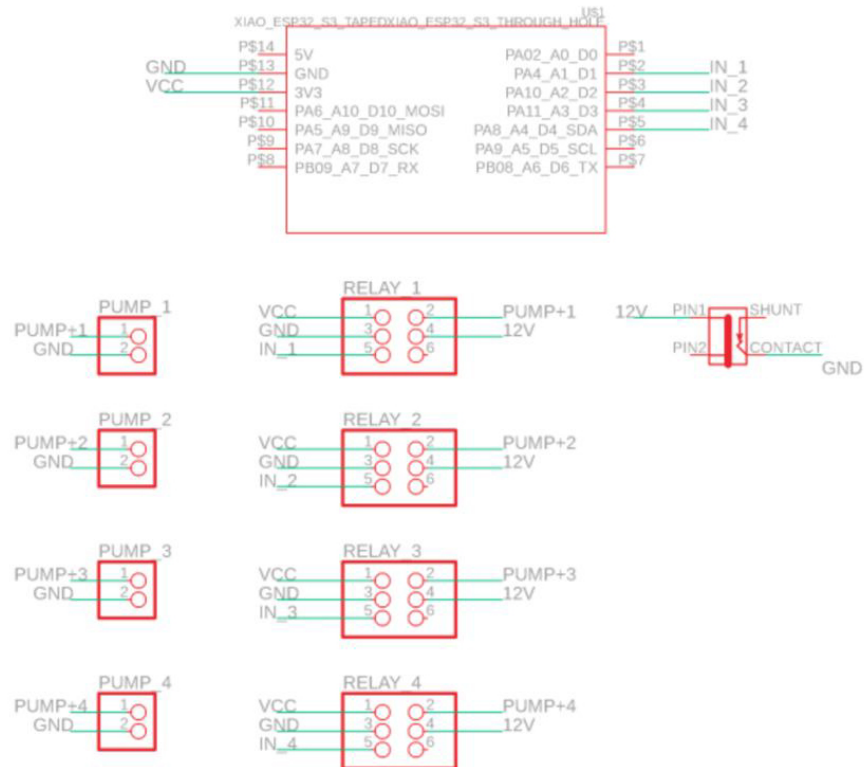


Figure 7: Output module: simplified schematic shows XIAO ESP32S3 controlling the logic inputs of four relay modules. Each relay switches the external 12V supply to its respective pump.

379 A.4 Curated fragrance library

380 This library was prepared via a combination of manufacturer descriptions [34], olfactory notes from
381 Odeuropa - a large scale study of scent notes [29], and the authors' own sensory evaluation.

Table 1: Fragrance library with semantic descriptors

#	Fragrance name	Semantic descriptors / keywords
1	Night Air	cool, crisp, ozonic, dark, mysterious, midnight, dew, subtle, empty street
2	Dirt	earthy, grounding, rich, soil, garden, damp, mineral, fertile, forest floor
3	Ocean Breeze	salty, aquatic, marine, invigorating, fresh, coastal, wind, mineral, expansive
4	Rain	ozonic, petrichor, wet pavement, fresh, rejuvenating, melancholy, atmospheric
5	Fresh Cut Grass	green, sharp, vibrant, summer, lawn, childhood, outdoors, fresh, uplifting
6	Campfire	smoky, woody, warm, nostalgic, embers, resinous, comforting, autumn
7	Sandalwood	woody, creamy, sweet, warm, smooth, meditative, sacred, balsamic, rich
8	Bamboo	green, crisp, clean, watery, airy, light, spa, fresh, zen
9	Forest Pine	coniferous, fresh, invigorating, pine needles, mountain, crisp, woody
10	Peaches & Cream	fruity, sweet, gourmand, comfort, creamy, soft, innocent, summer, dessert
11	Vanilla Bean	gourmand, creamy, baking, warm, sweet, rich, comforting, familiar, edible
12	Baby Fresh	soft, clean, powdery, light, innocent, delicate, talc, comforting, gentle
13	Green Tea	fresh, herbal, calming, zen, bitter, clean, spa, meditative, light
14	Magnolia	floral, creamy, lemony, delicate, southern, elegant, fresh, romantic, rich
15	Honeydew	juicy, fresh, sweet, aquatic, mellow, ripe, summer, light, fruity
16	Amber	rich, resinous, warm, golden, ancient, comforting, oriental, nostalgic
17	Sweet Tobacco	warm, smoky, sweet, leathery, pipe, grandfather, study, rich, vanilla, spicy
18	Cedar	woody, crisp, dry, clean, aromatic, closet, sharp, pencil shavings, grounding
19	Old Books	paper, dusty, woody, library, quiet, intellectual, history, leather-bound
20	Coffee	rich, roasted, bitter, energizing, morning, warm, comforting, dark, cafe, gourmand
21	Leather	earthy, robust, animalic, warm, smoky, worn, library, comforting, masculine
22	Sand	dry, warm, mineral, subtle, beach, sun-baked, grainy, neutral
23	Salty Sea	bracing, marine, fresh, mineral, aquatic, sharp, invigorating, ozonic, seaweed
24	White Tea	clean, delicate, subtle, elegant, fresh, light, spa, herbal, refined
25	Verbena	citrusy, fresh, green, lemony, invigorating, bright, herbal, sharp, clean
26	Seashore	marine, salty, breezy, aquatic, sand, sun, coastal, open-air, driftwood
27	Moss	earthy, damp, green, forest floor, ancient, cool, loamy, natural, shaded
28	Lilac	floral, sweet, fresh, spring, powdery, romantic, nostalgic, delicate
29	Cucumber Melon	fresh, green, juicy, sweet, aquatic, clean, young, light, refreshing
30	Coconut	creamy, tropical, sweet, gourmand, beach, sun-tan lotion, exotic, warm
31	Cinnamon	spicy, warm, sweet, gourmand, exotic, baking, comforting, woody, aromatic
32	Fresh Cotton	clean, soft, airy, laundry, fresh, comforting, light, powdery, simple
33	Violet	powdery, floral, delicate, sweet, earthy, slightly woody, nostalgic, classic
34	Sugar Cookie	sweet, gourmand, comforting, buttery, baking, warm, childhood, school
35	Vanilla	perfumey, sweet, amber, musk, complex, elegant, sophisticated, oriental
36	Pomegranate	fruity, tart, vibrant, juicy, sweet, exotic, sharp, bright, refreshing
37	Lavender	floral, herbal, calming, clean, aromatic, relaxing, sleep, classic, medicinal
38	Pear	juicy, crisp, fruity, fresh, sweet, green, light, delicate, orchard
39	Peppermint	cool, invigorating, crisp, minty, sharp, refreshing, clean, energetic

382 A.5 Example from the olfactory knowledge base

383 Below is one of the few-shot examples used to steer the LLM’s semantic-to-olfactory translation.

384 Input narrative

385 “In the fading light of autumn, the leather chair sits quietly by the window, remem-
386 bering the scent of rain on the dusty glass, a comfortable sadness for things long
387 past.”

```
388 {  
389   "Leather": 0.40,  
390   "Rain":   0.30,  
391   "Old Books": 0.20,  
392   "Amber":  0.10  
393 }
```

394

395 **Rationale** The narrative’s primary subject is the *leather chair*, making *Leather* the dominant note.
396 The mention of “rain on the dusty glass” maps directly to the *Rain* scent. The overall mood of
397 “comfortable sadness” evokes nostalgia, expressed with *Old Books* and the slightly somber resinous
398 note of *Amber*.

NeurIPS Paper Checklist

1. Claims

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

Answer: [\[Yes\]](#)

Justification: The abstract and introduction frame the paper's contributions which are conceptual and technical. They clearly state the contribution of the proposal of an alternative archetype for human-AI interaction around synthetic memory. The paper also describes the scope of the technical contribution: a novel multi-modal pipeline that uses an LLM for semantic-to-olfactory translation. The aspirational goal of provoking anemoia is framed as an exploration, which reflects the paper's scope as a speculative design inquiry without a formal user study.

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: In Section 5.2, the paper discusses the absence of a formal user study to empirically validate the experience. It also details the technical constraints of the current prototype. It also includes a critical reflection on the algorithmic approach, noting our current method does not model the underlying chemical or perceptual interactions between scents.

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: The paper presents a speculative design artifact and a conceptual framework. It does not include novel theoretical results in the form of theorems or mathematical proofs.

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: [Yes]

Justification: The paper fully discloses the information needed to understand the main claims, which are conceptual and system-based. We describe the full multi-modal pipeline, including the specific APIs. The methodology for the novel semantic-to-olfactory translation, including the structure of the Scent Library and the use of a few-shot Olfactory Knowledge Base, is described in detail, providing a clear path for other researchers to verify or replicate the approach.

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: At this stage of the research, the full code for the integrated system and the curated Olfactory Knowledge Base are not being publicly released. However, the paper and appendix provides sufficient architectural and methodological detail to allow for the conceptual replication of the system.

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, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The paper specifies the necessary details to understand the system’s design and function. This includes the specific LLM and VLM APIs used, the number of fragrance oils in the library (N=39), the number of scents in a final blend (up to 4), and the parameters of the tangible user interface (Subject, Time, Mood dials). Further details on the specific hardware components and fabrication are provided in the appendix.

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: The paper presents a research artifact and a conceptual framework. It does not include quantitative experiments or user studies that would require statistical significance testing.

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: [Yes]

Justification: The primary computational resources are calls to third-party APIs. The generation of a single narrative-scent pair is computationally inexpensive, completing in under 30 seconds. The on-device computation is handled by a standard host computer (e.g. MacBook Pro with an M1 chip) and a low-power microcontroller (XIAO ESP32S3), requiring minimal memory and processing power. No model training was performed for this project.

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 conducted in the paper conforms with the NeurIPS Code of Ethics. The project is a speculative design inquiry and does not involve human subjects, private data, or high-risk applications.

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: The paper discusses the societal impacts of AI and synthetic memory. The Discussion section contrasts the dominant trajectory of AI development and memory (screen-based, photorealism and persuasion) with the project's proposed alternative (poetic exploration, multi-sensory, and generative empathy). It discusses the potential negative impacts of misusing synthetic media while proposing our system as a model for a more humanistic and ethical path for creative AI.

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 project does not release new models or datasets. The system uses existing, commercially available LLM APIs and a small, curated, non-sensitive knowledge base. The output is a short text and a physical scent, which have limited potential for direct misuse in the manner of large-scale image or video generators.

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: The core of the system relies on commercially available, closed-source assets, and their use conforms to the terms of service provided by Anthropic and OpenAI, respectively. The archival photographs used as inputs for the device are sourced from personal images shared with permission or public domain collections where licenses permit creative reuse.

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: The paper does not release new assets such as datasets, code, or models.

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: The paper does not involve crowdsourcing or 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: The paper does not involve research with human subjects and therefore did not require IRB approval.

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: The paper explicitly describes the multi-stage use of a VLM for image captioning and an LLM for both narrative synthesis and the novel task of few-shot, semantic-to-olfactory translation.

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.