Roulettective: Repurposing a Vintage Slide Projector into an AI-Driven Physical Detective Game

Yuhan Wang *
MIT Design Intelligence Lab;
Graduate School of Design
Harvard University
Cambridge, MA 02138
yuhan_wang@gsd.harvard.edu

Jianuo Xuan *
MIT Design Intelligence Lab;
Graduate School of Design
Harvard University
Cambridge, MA 02138
jianuo_xuan@gsd.harvard.edu

Abstract

Roulettective is an AI-driven Physical Interface repurposing the vintage carousel slide projector for immersive detective gaming. The project explores the projector's potential in a detective mystery solving game context, creating new usage scenarios by reprogramming its physical interaction mechanisms and integrating AI-driven gameplay. It further transforms this forgotten artifact into a multimodal, immersive, customizable, intuitive, and co-creative interface through AI-generated narratives, visuals, and sound. *Roulettective* focuses on repurposing the interface modalities that have been supplanted by current computing paradigms. It exemplifies a design paradigm for learning from outdated artifact legacies, introducing AI repurposing as a generalizable framework for AI-driven Physical Interfaces.

1 Introduction

The rapid advancement of Artificial Intelligence (AI) has extended its impact beyond digital systems into the physical environment. By grounding interaction in familiar tangible forms, Physical Interfaces draw on users' intuitive behaviors, supporting more fluid, expressive, and situated engagement with machines [11]. The introduction of the concept of Large Language Objects (LLOs)[7] highlights a shift toward integrating AI systems into physical objects through embodied multimodal interfaces[15]. LLOs explore the potential for physical objects to exhibit adaptive and context-aware behaviors, thereby challenging traditional boundaries between form and function in interaction design[7].

Many LLO projects combine digital AI with existing objects, allowing not just augmentation but AI repurposing. AI repurposing is an approach that uses obsolete artifacts as prototypes, whose forms and mechanisms are reinterpreted and contextualized to address new scenarios and contemporary needs. This design methodology views AI as a means to reconnect with the legacy of physical interfaces, reviving the embodied interactions lost in the shift to efficiency-driven screen-based computing paradigms. By reimagining forgotten artifacts, AI repurposing reintroduces embodied, multimodal interactions that were displaced by their digital counterparts.This methodology offers a new perspective for widespread application in AI embodiment.

The *Kodak Carousel Projector*[8], though absent from everyday life, remains ripe for repurposing due to its cultural significance and mechanical affordances. Released in 1961 and discontinued in 2004 with the rise of digital technologies,[13][23] decades of use shaped its identity as a cultural artifact. In popular culture, it appeared in Mad Men (2007), renamed "The Carousel" and framed as a metaphor for cyclical time.[26] In psychological studies, it was used to present sequential imagery from simulated crime scenes.[18][17] Its tactile mechanics and optics made it ideal for immersive

*Equal contribution.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Creative AI Track.

Figure 1: a: Original Kodak Carousel Projector. Image: Wikimedia Commons, CC BY-SA 3.0..
b: Using of Carousel projector in detective scene (AI generated). c: Roulettective in use. d: Objects of Roulettective, with the AI Carousel on the left and the Archive Box with slide cards on the right.

narration, inspiring artists [14] to explore slide projection as a medium.[13] This blend of cultural memory and material interaction positions the projector as an ideal prototype for creating immersive tangible experiences.

Roulettective is an AI-driven detective game machine that builds on this heritage by repurposing the projector's mechanism and exploring new scenarios for the artifact. It reprograms the projector's buttons and card insertion functions, while an Large Language Model dynamically generates narratives, clues, and analytical assistance. The device extends the projector's multimodal outputs including sound, image, and text, to create customized content in real time.

By adopting AI repurposing as its approach, *Roulettective* advances the discourse on LLO, demonstrating how "obsolescence media" [14] can be transformed into immersive, physical interface.

2 Related Work

Recent explorations of AI-embedded hardware have demonstrated how integrating AI functionalities into physical objects can enable novel interaction paradigms[7][12][20]. By introducing tangible interaction principles [11], these projects give AI a presence in the physical world, enabling more immersive and intuitive engagement.

2.1 Repurposing of Obsolete Objects

Roulettective contributes to a growing body of research that explores embedding AI capabilities into obsolete artifacts to create novel, meaningful forms of physical interaction. These projects leverage the cultural familiarity and design legacy of iconic objects to ground and enhance AI-driven experiences.

Augmenting original functions. Obsolete designs offer a foundation for reinterpreting object functionality through AI. For example, a traditional boombox plays music selected by the user. *Be the Beat* [5] reimagines this by allowing users to dance while the AI selects music in response—shifting the interaction from manual control to embodied input.

Leveraging familiar forms for intuitive AI interaction. Physical affordances of historical objects can simplify users' understanding of AI systems. *Dream Generator* [10], for instance, transforms a point-and-shoot camera into a stylized image generator, using the familiar act of pressing a shutter to initiate AI-driven outputs. This familiarity in physical form facilitates the understanding of Human-AI interaction.

Merging symbolic meaning with AI feedback. Some works combine ritualistic or symbolic artifacts with AI to evoke deeper cultural or emotional resonance. *AIncense* [28] reinterprets incense-burning practices by using AI-generated voices to respond to user speech, creating an impression of spiritual communication while reconfiguring the authority and responsiveness of the original ritual.

Roulettective similarly draws from the physical and cultural heritage of the carousel slide projector. It reuses the tangible act of inserting slides and reframes it as a prompt for AI narration. Unlike prior

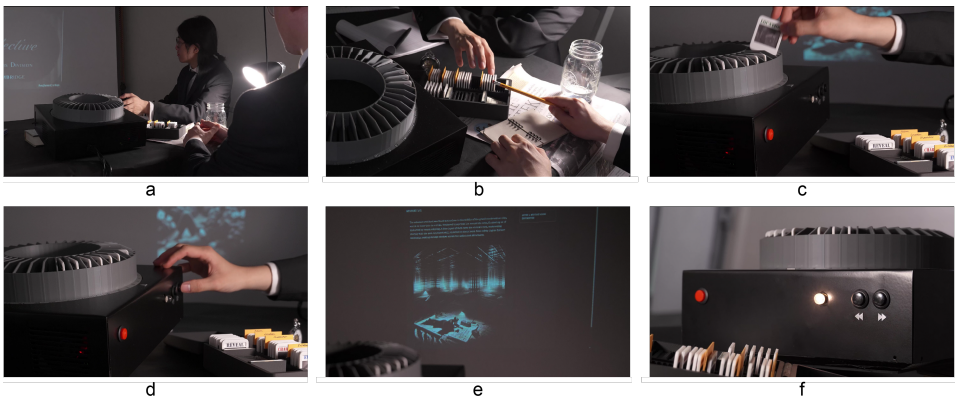

Figure 2: Playing *Roulettective*. a: Players gather around *Roulettective* to start solving a new mystery. b: Players choose the type of information slide from the archives. c: Player insert the slide. d: Player press the forward button to play the inserted slide. e: *Roulettective* projects the mystery information of that slide type. f: *Roulettective* lights up to hint the player of connections between the current slide and a former slide.

examples, it revives an obsolete medium, positioning the project not only as a novel AI interface but also as a form of cultural preservation and reinterpretation through design.

2.1.1 AI Narration & Detective Games

Language-based AI models possess a specialized talent for generating narration. Recent explorations have leveraged this capability to create personalized, real-time interactions in education, [9] creative writing, [6] and gaming. [25] Using AI to generate game narration [3, 2] fosters unique and ever-evolving storylines, providing unprecedented flexibility and player agency.

For detective-related games, both classic board games [24][21][22][1]and narrative video games [4][19][27][16] have established frameworks for embedding investigative elements into gameplay. *Cluedo* [24] pioneered the abstraction of a case into key elements such as location, time, character, and tools, offering a structured foundation for mystery-solving. *Ace Attorney* [4], on the other hand, uses branching dialogues and multiple-choice interactions to ensure players grasp the storyline without requiring complex text input.

These precedents provide valuable reference points for designing AI-driven narrative gameplay, guiding how detective mechanics can be meaningfully integrated with dynamic AI narration.

Building on these strands, *Roulettective* explores how AI can be used to repurpose outdated designs and objects. By combining tangible interaction, generative narration, and game mechanics, the project identifies obsolete physical interactions and cultural associations embedded in these artifacts.

3 Design

3.1 Gameplay

The design of *Roulettective* integrates detective gameplay mechanics with the form and logic of a retro slide projector. It repurposes key interactive elements from the original prototype, including slide cards, carousel rotation buttons, indicator bulb, and an embedded projection speaker. The player's main interaction consists of inserting RFID-tagged slide tokens and pressing buttons to activate AI-driven narrative and visual responses that reveal the mystery. The gameplay proceeds through the following phases:

- **Standby Phase**: The system begins in a guided interface, where players can use the backward button to review tutorial content before initiating gameplay.
- **Case Initiation**: Inserting a "Mystery" slide into the next available slot and pressing the forward button launches a new case.

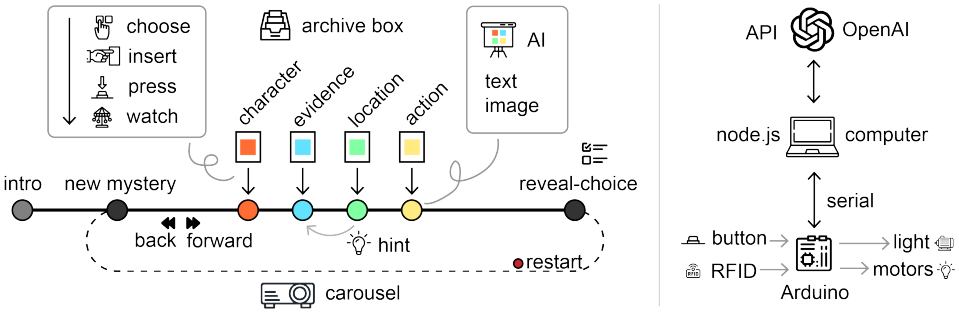

Figure 3: Left: Gameplay of *Roulettective*. Right: Data Pipeline

- **Exploration Phase**: This is the major playing phase. Players insert "Exploration" slides—classified as Evidence, Character, Location, or Action—to advance the investigation. Each slide insertion, followed by carousel rotation, prompts the AI to generate new narrative clues through projected visuals and sound.
- **Association Monitoring**: As slides accumulate, the system continuously analyzes their content. When the AI detects meaningful associations or contradictions, internal logic updates prior clues. Players can navigate backward to revisit earlier slides, which may now reflect revised or expanded content.
- **Conclusion Phase**: After collecting sufficient information, players insert a "Reveal" slide to initiate final analysis. A "Choice" slide then presents five possible statements; players must identify the false one to uncover the truth and resolve the case.
- **Reset Phase**: Pressing the reset button returns the system to its initial interface. Players manually remove all inserted slides to prepare for the next round of gameplay.

In the prompt design, the game is structured as a logically coherent and solvable puzzle. Each newly generated clue is logical consistent with the previous ones, while still allowing players to explore with uncertainty. All clues collectively influence the final result generation. In other words, the game ensures a stable narration while preserving the suspenseful experience essential to the mystery genre.

This integration of physical gestures, narrative AI, and reactive feedback transforms the act of mystery-solving into an interactive, cinematic experience.

3.2 Technical Implementation

Roulettective integrates physical interaction and generative AI through a unified system architecture. An Arduino Pro Micro captures hardware inputs including carousel rotation, RFID readings, and button presses, and sends them as structured JSON data to a Node.js backend. These inputs trigger GPT-4o to generate narrative text and DALL·E 3 to produce visual outputs. The backend maintains narrative coherence through a memory system, filters unsupported tokens, and summarizes image prompts for compatibility. An association module analyzes AI responses to identify connections or contradictions between clues. All outputs including text, images, and case graphics are stored in runtime memory to support real-time review and replay during gameplay.

3.3 Design Repurposing

The repurposing of *Roulettective* operates on two levels: interaction design and mechanical adaptation, both informed by the physical logic and cultural semiotics of the Kodak slide projector.

Interaction Design The original projector's interaction model centered on three elements: button-based navigation, physical slide insertion, and projected imagery. *Roulettective* retains all three, maintaining continuity with users' intuitive understanding of the device.Slide cards are redesigned as RFID-enabled tokens, each corresponding to a narrative function within the mystery structure. As the story is generated in real time by AI, inserting slides functions as both an investigative act and a means of narrative intervention by the player.

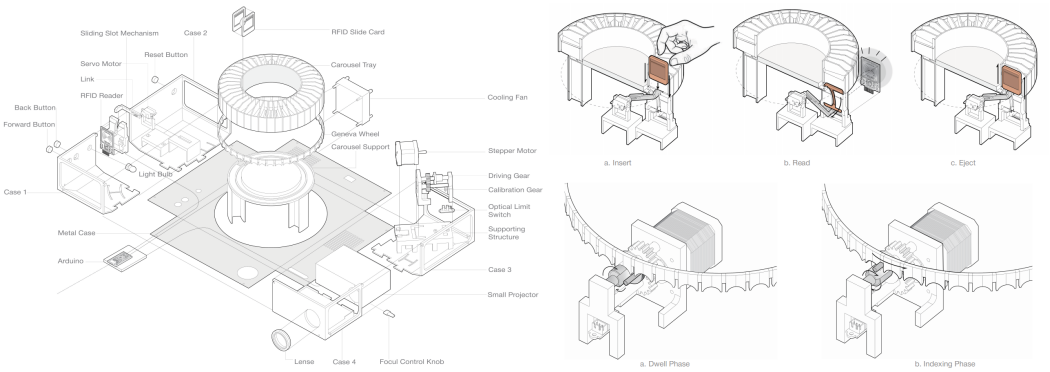

Figure 4: Mechanism design of Roulettective

Mechanical Adaptation The original *Kodak* projector, as an industrial optical device, offered high positional accuracy in slide transitions. To preserve this precision under new gameplay demands, the system retains the projector's guide-rail geometry while integrating a Geneva mechanism for carousel control. The Geneva drive, coupled with a stepper motor and calibrated with an optical limit switch, enables discrete, indexed rotation that aligns each token with the optical center. This hybrid structure, which combines legacy slot mechanics with a custom-engineered rotational system, ensures reliable slide alignment and smooth bidirectional operation.

4 Preliminary Feasibility Test

To explore *Roulettective*'s effectiveness and gather qualitative insights, we conducted a series of preliminary feasibility test, with a group of 23 participants who are experts in this field. Each user experienced a whole gameplay cycle of *Roulettective* from generating mystery to the choice of solution. After each test we conducted semi-structured open-ended interview to collect their initial impressions and thoughts.

Users commented that *Roulettective*'s operation was largely intuitive. Familiarity with the physical appearance and mechanisms of traditional slide projectors enabled participants to engage with the device with minimal instruction. The mapping between physical gestures and system response was quickly understood after brief orientation. The vintage aesthetic, tactile interactions, and multisensory feedback—such as the carousel's clicking sounds and the movement of slides—were noted as key factors contributing to immersion and a strong sense of telepresence. Participants highlighted that the AI's role in generating clues and hints added a co-creative quality, making the machine feel responsive and "alive."

Some problems and suggestions are also raised in the feedback. Some users initially misunderstood how slide tokens mapped to AI-generated content, but adapted quickly after reading the introduction. A few participants also commented on latency in AI image generation, observing its impact on gameplay pacing but recognizing this as a technical limitation that could be improved with future AI models. At this stage, loading animation and music are added in response to the delay.

The impact on a larger audience will be analyzed through a future user study.

5 Methodology

Roulettective is a research project on AI repurposing that establishes a methodology for expanding AI interaction modalities. It begins by selecting an obsolete prototype from historical eras. The next step is research, focusing on the object's historical development, tangible interaction design, and cultural impact. Then comes repurposing, which combines the object's cultural context with AI capabilities to redesign its use scenarios and adapt its interaction mechanisms. The final step is integration, where the reinvented object is built and refined through user testing and iteration.

In this project, the *Kodak Carousel Projector*[8] serves as the prototype. Its practical applications in psychology and criminal investigations are explored, alongside its historical obsolescence with the advent of digital systems like PowerPoint. The design then reimagines its use scenario within a detective game. AI-generated capabilities enable this repurposing. Ultimately, the reimagined object is realized by reframing its original mechanism through this methodology. This approach offers a novel way to discover AI-driven embodied interactions by bridging the understanding of both obsolete artifacts and AI technologies.

6 Limitations

Based on user feedback and the observations in our design iterations, we identify some key limitations in our current approach to AI-driven object repurposing.

Fixed mental models of historical objects. Users often rely on established expectations when engaging with existing forms, such as familiar interaction patternsand established embedded meanings. While this enhances affordance, it can hinder understanding when those forms are reimagined with dynamic, AI-driven functions. In our case, users were confused by slide tokens representing categories rather than specific, static images.

Bridging old and new mental models requires maintaining some continuity in interaction logic and physical form. Our use of slides helped ground new concepts in familiar formats.

Modality constraints of current AI. AI modalities are limited. Text, image, video are the most current most popular modals. Text-based LLMs are the most accessible form today, but their reliance on language input/output limits the range of physical interaction. This gap constrains the integration of rich physical interfaces with AI's primarily linear, symbolic outputs.

Roulettective addresses this by using metaphorical slide tokens as embodied prompts, enabling physical engagement with a text-based system. While current AI modalities restrict the design space, emerging multimodal models may soon allow deeper integration between form and intelligence.

7 Discussion

Physical Interface gives AI the ability to engage with the physical world through both input and output. Building on the original design of the slide projector, *Roulettective* integrates prompts into physical slides, leveraging the carousel's familiar interaction patterns to enhance the affordance of AI control. The mechanical movement of the carousel makes the AI experience more perceptible and embodied, strengthening the sense of immersion and authenticity.

As an industrial product that has largely disappeared from everyday life, the slide carousel gains new meaning when embedded with AI and enriched through a gaming experience. These artifacts, though no longer in use, have outdated functions but adoptable forms and mechanics. By integrating new technologies with these existing designs and interactions, their repurposing as LLOs can lead to novel discoveries. *Roulettective* repurposes Carousel projector by introducing AI narrational gameplay to the slide projecting interaction, turning the original tool into an AI companion that co-creates with user, and bring this historical object back to life.

8 Conclusion

The project selected the *Kodak Carousel Projector*[8] as obsolete object prototype, analyzed its cultural connection to detective narratives, and implemented an AI-driven detective game closely related to its original physical interactions. On one hand, the system proposes intuitive, high-affordance tangible interaction methods for AI hardware; on the other hand, it uses AI to generate personalized and unique feedback, elevating the interaction to a new level of immersion and engagement.

The project presents *Roulettective* not as a standalone AI hardware implementation, but as an example of a broader methodology for AI-based repurposing of obsolete artifacts, with the hope that this approach will inspire similar designs and interactions in the future.

References

[1] 2004. Betrayal at House on the Hill. Board game designed by Bruce Glassco, Rob Daviau, Bill McQuillan, Mike Selinker, and Teeuwynn Woodruff.

[2] 2019. AI Dungeon. https://aidungeon.com/. Accessed: 2025-07-21.

[3] Steph Buongiorno, Lawrence Klinkert, Zixin Zhaung, Tanishq Chawla, and Corey Clark. 2024. PANGeA: procedural artificial narrative using generative AI for turn-based, role-playing video games. In *Proceedings of the Twentieth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment* (Lexington) *(AIIDE '24)*. AAAI Press, Article 16, 11 pages. https://doi.org/10.1609/aiide.v20i1.31876

[4] Capcom. 2001. Phoenix Wright: Ace Attorney. Video game. Originally released for Game Boy Advance in Japan; known as Gyakuten Saiban.

[5] Ethan Chang, Zhixing Chen, Jb Labrune, and Marcelo Coelho. 2025. Be the Beat: AI-Powered Boombox for Music Suggestion from Freestyle Dance. In *Proceedings of the Nineteenth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '25)*. Association for Computing Machinery, New York, NY, USA, Article 67, 6 pages. https://doi.org/10.1145/3689050.3705995

[6] John Joon Young Chung, Wooseok Kim, Kang Min Yoo, Hwaran Lee, Eytan Adar, and Minsuk Chang. 2022. TaleBrush: Sketching Stories with Generative Pretrained Language Models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 209, 19 pages. https://doi.org/10.1145/3491102.3501819

[7] Marcelo Coelho and Jean-Baptiste Labrune. 2024. Large Language Objects: The Design of Physical AI and Generative Experiences. *Interactions* 31, 4 (June 2024), 43–48. https://doi.org/10.1145/3672534

[8] Eastman Kodak Company. 2006. Kodak Carousel Slide Projector. https://web.archive.org/web/20060405040201/http://slideprojector.kodak.com/. Archived from the original on April 5, 2006.

[9] Min Fan, Xinyue Cui, Jing Hao, Renxuan Ye, Wanqing Ma, Xin Tong, and Meng Li. 2024. StoryPrompt: Exploring the Design Space of an AI-Empowered Creative Storytelling System for Elementary Children. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 303, 8 pages. https://doi.org/10.1145/3613905.3651118

[10] Kyle Goodrich. 2023. Dream Generator. https://kylegoodrich.xyz/. Camera prototype merging photography with generative AI. Built with RaspberryPi, Stable Diffusion, ControlNet..

[11] Hiroshi Ishii and Brygg Ullmer. 1997. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) *(CHI '97)*. Association for Computing Machinery, New York, NY, USA, 234–241. https://doi.org/10.1145/258549.258715

[12] Quincy Kuang, Feifei Shen, Cathy Mengying Fang, and Annie Dong. 2023. Memeopoly: An AI-Powered Physical Board Game Interface for Tangible Play and Learning Art and Design. In *Companion Proceedings of the Annual Symposium on Computer-Human Interaction in Play* (Stratford, ON, Canada) *(CHI PLAY Companion '23)*. Association for Computing Machinery, New York, NY, USA, 292–297. https://doi.org/10.1145/3573382.3616057

[13] Martha Langford. 2014. When the Carousel stops turning ... What shall we say about the slide show? *Intermédialités / Intermediality* 24-25 (2014). https://doi.org/10.7202/1034158ar

[14] Nicholas Loewen. 2021. *Making light: criticality and carousel projection in the work of Marcel Broodthaers and Allan Sekula*. Ph. D. Dissertation. University of British Columbia. https://doi.org/10.14288/1.0395462

[15] A. Miriyev and M. Kovač. 2020. Skills for physical artificial intelligence. *Nature Machine Intelligence* 2 (November 2020), 658–660. https://doi.org/10.1038/s42256-020-00258-y

[16] Lucas Pope. 2018. Return of the Obra Dinn. https://obradinn.com/. Accessed: 2025-07-22.

[17] John S. Shaw III. 1996. Increases in Eyewitness Confidence Resulting from Postevent Questioning. *Journal of Experimental Psychology: Applied* 2, 2 (June 1996), 126–146.

[18] John S. Shaw III, L. M. Appio, T. K. Zerr, et al. 2007. Public Eyewitness Confidence Can Be Influenced by the Presence of Other Witnesses. *Law and Human Behavior* 31, 5 (2007), 629–652. https://doi.org/10.1007/s10979-006-9080-6

[19] Spike Chunsoft. [n. d.]. Danganronpa – Official U.S. Website. https://danganronpa.us/. Accessed: 2025-07-22.

[20] Haoheng Tang and Mrinalini Singha. 2024. A Mystery for You: A fact-checking game enhanced by large language models (LLMs) and a tangible interface. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 631, 5 pages. https://doi.org/10.1145/3613905.3648110

[21] Ignacy Trzewiczek. 2018. Detective: A Modern Crime Board Game. https://portalgames.pl/en/detective-a-modern-crime-board-game/. Accessed: 2025-07-22.

[22] Nikki Valens. 2016. Mansions of Madness: Second Edition. https://www.fantasyflightgames.com/en/products/mansions-of-madness-second-edition/. Accessed: 2025-07-22.

[23] Jeroen Verbeeck, Lilian Haberer, Helen Westgeest, Jelena Stojković, Julian Ross, Dork Zabunyan, and Paige Sarlin. 2015. Image After Image: Reconsidering the Fabric of Slide Shows. *PhotoResearcher* 24 (October 2015). https://www.academia.edu/17426124/Image_After_Image_Reconsidering_the_Fabric_of_Slide_Shows Guest editor: Jelena Stojković.

[24] Waddingtons. 1949. Cluedo. Board game, published by Waddingtons (UK) / Parker Brothers (USA).

[25] Xiyuan Wang, Yi-Fan Cao, Junjie Xiong, Sizhe Chen, Wenxuan Li, Junjie Zhang, and Quan Li. 2025. ClueCart: Supporting Game Story Interpretation and Narrative Inference from Fragmented Clues. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 410, 26 pages. https://doi.org/10.1145/3706598.3713381

[26] Wikipedia contributors. 2024. Carousel Slide Projector. https://en.wikipedia.org/wiki/Carousel_slide_projector. Accessed July 2025.

[27] ZA/UM. 2019. Disco Elysium. https://discoelysium.com/. Accessed: 2025-07-22.

[28] Kai Zhang. [n. d.]. AIncense – Large Language Objects. https://kaizhang.io/large-language-objects. Accessed: 2025-07-22.

A Supplementary Material

Our software implementation code base can be found on our GitHub repo:

https://github.com/Rabourackee/Roulettective

NeurIPS Paper Checklist

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

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

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

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

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

IMPORTANT, please:

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

1. **Claims**

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

   Answer: [Yes]

   Justification: Abstract and Introduction clearly reflect our novelty and contribution of Roulettective: We designed an example of AI repurposing of a historical object.

   Guidelines:

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

2. **Limitations**

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

   Answer: [Yes]

Justification: We have a section of limitation on the problems that could occur in the corner examples of AI and historical objects using our methodology.

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: Our paper focus more on the design aspect and does not have a quantitative theoretical result.

   Guidelines:

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

4. **Experimental result reproducibility**

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

   Answer: [NA]

   Justification: Our paper doesn't include an experiment, but we described our design in hardware and software clearly and it should be able to reproduce easily.

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: [Yes]

   Justification: We provide the GitHub repo of our software part of the design,linked at the appendix of this paper. The accompanying README file and the detailed descriptions of this paper should allow faithful reproduction of this project.

   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: [NA]

   Justification: Our paper focus more on the design aspect and experiments is not a focus.

   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: Our paper focus more on the design aspect and does not have experiments that has quantitative results.

   Guidelines:

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

8. **Experiments compute resources**

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

   Answer: [NA]

   Justification: Our paper focus more on the design aspect. Our project can be reproduced on common computers.

   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: This work involved informal user testing with voluntary participants for prototype evaluation. All participants gave verbal consent, and no personal or sensitive data were collected. As the study posed minimal risk and our institution does not require formal ethics approval for such activities, no IRB process was conducted.

   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: [No]

    Justification: The work does not explicitly address societal impact, as it focuses on a constrained, artistic, and exploratory design context rather than a deployed technology with broad societal reach. The system is an AI-driven interactive narrative prototype using a repurposed vintage device for storytelling and game-based interaction. It does not involve sensitive data, automated decision-making, or large-scale deployment scenarios, and thus presents minimal risk of negative societal impact.

    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 paper poses no such risks.

    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: We used the OpenAI GPT-4 API (https://openai.com/api) under OpenAI's terms of service. No other external datasets, code, or third-party assets were used. The physical design of the prototype draws inspiration from the Kodak Carousel slide projector, but the hardware used is original and not based on copyrighted CAD files or patented mechanisms.

    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: [Yes]

    Justification: While we have not prepared formal structured documentation, we release a GitHub repository (linked at the appendix of this paper) with a clear README and describe the system architecture and functionality in detail in the paper. This provides sufficient information for replication and use of the new assets.

Guidelines:

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

14. **Crowdsourcing and research with human subjects**

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

    Answer: [No]

    Justification: The paper includes informal user testing with voluntary participants to evaluate the interactive prototype. No crowdsourcing was used and no compensation was provided. Instructions were given verbally during the sessions and not formally documented, and thus are not included. No sensitive data were collected.

    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: [No]

    Justification: This work involved informal user testing with voluntary participants to evaluate a design prototype. The study posed minimal risk, did not involve collection of personal or sensitive data, and participants provided verbal consent.

    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 rigorousness, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: We used GPT-4o in our implementation, and we discussed the usage of LLM in game narrative generation and game play system generation clearly in the design part of the paper.

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.