

# What if I Were?: Synthetic Lived Experience based on Generative Models

Sun-Kyung Lee\*

Digital Convergence Research Lab  
Electronics and Telecommunications  
Research Institute  
Daejeon, Republic of Korea  
sklee2014@etri.re.kr

Hyuntae Jeong

Digital Convergence Research Lab  
Electronics and Telecommunications  
Research Institute  
Daejeon, Republic of Korea  
htjeong@etri.re.kr

Se Won Oh

Digital Convergence Research Lab  
Electronics and Telecommunications  
Research Institute  
Daejeon, Republic of Korea  
sewonoh@etri.re.kr

Jeong Mook Lim

Digital Convergence Research Lab  
Electronics and Telecommunications  
Research Institute  
Daejeon, Republic of Korea  
jmlim21@etri.re.kr

Kyoung Ju Noh

Digital Convergence Research Lab  
Electronics and Telecommunications  
Research Institute  
Daejeon, Republic of Korea  
kjnoh@etri.re.kr

Seungeun Chung

Digital Convergence Research Lab  
Electronics and Telecommunications  
Research Institute  
Daejeon, Republic of Korea  
schung@etri.re.kr

Gyuwon Jung

Digital Convergence Research Lab  
Electronics and Telecommunications  
Research Institute  
Daejeon, Republic of Korea  
gwjung@etri.re.kr

## Abstract

People often wonder what it would be like to be in a different situation: how others live, and how their own life might unfold under different circumstances. In this work, we propose Synthetic Lived Experience (SLE) as computationally generated, first-person accounts that approximate “what I might do and feel if I were this person, in this situation” or “how my life might look if my context changed.” Using a generative model-based pipeline, we specify personas and contexts, and then synthesize episodic narratives and corresponding visualizations that depict actions, environments, and emotions. We present initial qualitative examples of SLEs that illustrate how generated first-person narratives can surface everyday lived experience under specified personas and contexts. We also discuss key limitations and outline opportunities for future work.

## CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; • **Applied computing** → **Sociology**.

## Keywords

Large Language Models, Generative Models, Lived Experience, Synthetic, Experience Augmentation

\*Corresponding author



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI EA '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2281-3/26/04  
<https://doi.org/10.1145/3772363.3798417>

## ACM Reference Format:

Sun-Kyung Lee, Hyuntae Jeong, Se Won Oh, Jeong Mook Lim, Kyoung Ju Noh, Seungeun Chung, and Gyuwon Jung. 2026. What if I Were?: Synthetic Lived Experience based on Generative Models. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3772363.3798417>

## 1 Introduction

People are constantly curious about other lives. We imagine what it would be like to have a different job, to have grown up in another city, or to have been born into a completely different background. According to Markus’s theory of possible selves [12], people gain motivation, adjust their behavior, and set long-term directions for change through imagining such self-relevant futures and alternatives. Through narrative transportation [5] and experience-taking [7], people also come to empathize with the actions and emotions of others in stories, and sometimes change their own beliefs and attitudes as a result.

The rise of social media has brought these other lives even closer. Through constantly updating feeds and stories, we develop a kind of background awareness of what others are doing, even without interacting with them [10]. Yet it remains difficult to understand, in a truly lived sense, what it is like to be in their exact situation: what they experience, what they see and hear, and how they feel in that moment. Differences in time, place, and social and cultural background still stand in the way of deep empathy and experiential understanding, and can even trigger negative feelings of envy, inadequacy, and anxiety [2, 14]. This raises a central question for our work: “*Can we systematically simulate lives different from our own in order to augment our experience sustainably?*”

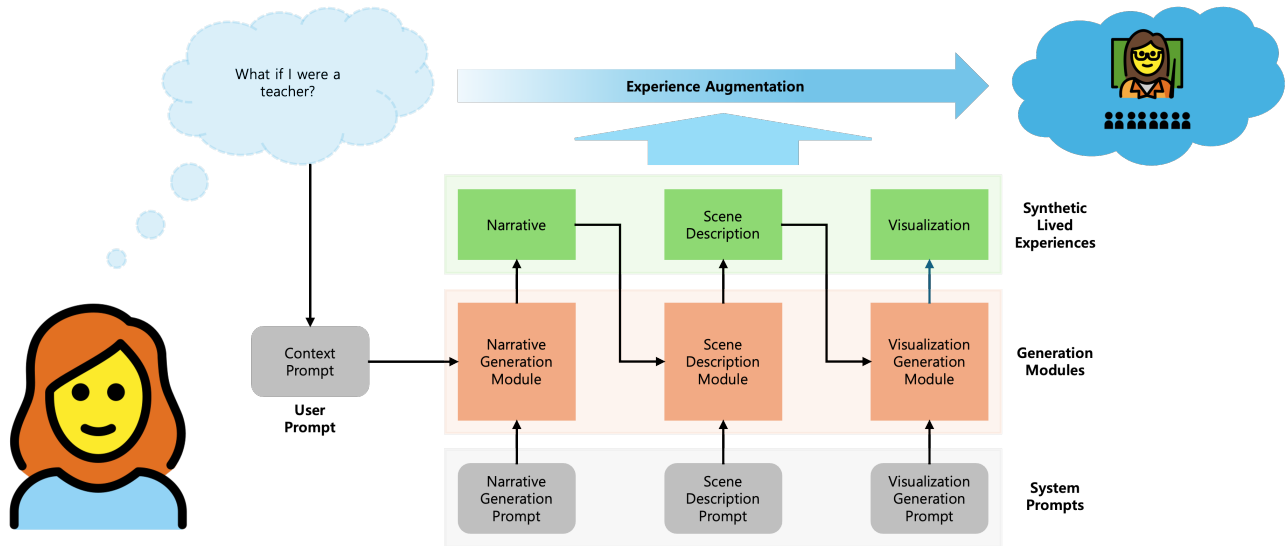


Figure 1: Overview of the proposed Synthetic Lived Experience

Recent advances in generative models offer a new way to approach this curiosity and question. Large language models (LLMs) [1, 19], drawing on vast general knowledge, can now produce coherent narratives about everyday situations from relatively compact prompts. Image generation models [17] have also reached a level where they can create realistic scenes comparable to photographs. When combined, these two developments suggest that generative models need not remain mere tools for “making stories and pictures.” Instead, they can become a medium for constructing virtual approximations of lived experience for diverse personas and what-if scenarios.

Several lines of prior work are related to this idea, including (i) LLM-based systems [3, 6] for conveying emotions and fostering empathy, (ii) persona-based role-playing [15, 20, 23], and (iii) perspective-taking interventions [24]. However, prior work has predominantly emphasized interactive dialogue, role-play, or decision-centric exchanges, where the system’s unit of generation is a conversational turn or a discrete choice. What remains less explored is a narrative-centered representation of lived experience: a coherent first-person episode that sustains context over time and foregrounds what the person notices, feels, and reflects on as ordinary events unfold.

In this paper, building on these threads, we propose the concept of **Synthetic Lived Experience (SLE)**. We define SLE as a **virtual lived experience, constructed in the first-person by generative models**. In other words, even without actually having lived that life, a user can vicariously live through another life situation or what-if scenario as if they themselves were inside it. Whereas prior work primarily simulates dialogue or one-off decisions, SLE emphasizes holistic lived experience, expressed through narratives and visualizations. More concretely, SLE focuses not only on what happened, but also on what the person notices and perceives, what emotions and thoughts arise, and what the surrounding atmosphere

and context are like, all written in a first-person voice. On top of this, SLE adds visualizations grounded on the narrative, enabling a more immersive experiential engagement.

The goal of this paper is to propose the concept of SLE and present a basic framework along with initial results. We present SLE as a method for generating plausible, inspectable, and revisable approximations of other lives and other possible versions of myself. These approximations enable lightweight tests of assumptions, help users externalize and elaborate what they imagine, provoke reflection, and support early-stage ideation. Ultimately, SLE aims to move beyond imagining others in the abstract by fostering a deeper understanding of others’ lived experiences.

In this paper, we make three main contributions:

- **Conceptualizing SLE:** We articulate SLE as a generated approximation of what it might feel like to live other lives in a particular situation.
- **Building SLE Pipeline:** We describe a generative model-based pipeline that takes user-designated persona and context specifications as input and produces episodic first-person narratives and associated visualizations as output.
- **Initial Qualitative Results:** We present examples of generated SLEs, discuss their strengths and limitations, and outline future directions.

## 2 Methods

We implement Synthetic Lived Experience (SLE) generation as a three-stage pipeline: (i) Narrative Generation Module, (ii) Scene Description Module, and (iii) Visualization Generation Module. Each module is controlled by a system-level prompt and consumes the output of the previous stage. Across the pipeline, we use a Context Prompt to specify who the experience is about and in what situation. We refer to the resulting Narrative, Scene Description, and Visualization collectively as the generated SLEs.

Type	Use	Prompt
User Prompt	Context	<i>e.g., A woman in her 50s, a high school teacher, teaching in a classroom, two students making noise</i>
System Prompt	Narrative Generation	Please generate a first person narrative given "[CONTEXT]" 1. Start with a short "Persona Recap" about who I am and my situation 2. Create an "Episode" and write in the first person ("I") describing: - my actions - my perceptions - my emotions and thoughts 3. End with a short "Reflection" on what mattered most to me in this experience
	Scene Description	Please generate a concise visual scene description for image generation given "[NARRATIVE]" 1. Summarize into 1–2 short sentences of vivid scene description 2. Focus on what could be seen in a single image: setting, people, objects, and overall atmosphere
	Visualization Generation	Please generate a realistic image given "[SCENE DESCRIPTION]"

Figure 2: Summary of the four prompts used for SLE generation

## 2.1 Input Prompts

There are four main input prompts, comprising one user prompt and three system prompts. Each prompt is summarized in Fig. 2. The user-driven Context Prompt is the scenario-specific input authored by the user. Several key pieces of information it can encode are as follows:

- **Persona:** Who the SLE is about (e.g., another person or an alternative version of myself), described through relevant attributes such as age, occupation, role, and characteristics.
- **Location:** Where the SLE takes place, including the broader setting (e.g., city, region) and place type (e.g., workplace, home, public transit).
- **Time Frame:** When the SLE occurs and its temporal scope, such as the time of day or duration.
- **Situation:** What is happening in the SLE, including the main activity or event and any relevant social or environmental conditions (e.g., who is present, ongoing pressures)

The first system prompt is a Narrative Generation Prompt, which configures the module to produce first-person, episodic narratives. It specifies the output structure (e.g., persona recap, episodes, reflection), writing style (first person, everyday detail, emotions), and constraints (e.g., "avoid stereotypes", "stay within the given context"). Scene Description Prompt defines how to transform a narrative into concise, visualization-friendly descriptions: one or two sentences per scene, focusing on visible settings, people, objects, and atmosphere. Visualization Generation Prompt encodes fixed preferences such as a neutral artistic style or level of realism (e.g., "keep the main character recognizable," "everyday documentary style"), and how to interpret the Scene Descriptions.

## 2.2 Generation Modules

The aforementioned prompts control each generation module. The Narrative Generation Module takes the Context Prompt and Narrative Generation Prompt and returns a structured, first-person Narrative. Then, the Scene Description Module takes the Narrative and Scene Description Prompt to generate the Scene Description for

the visualization generation process. Finally, the Scene Description and the Visualization Generation Prompt are fed into the Visualization Generation Module to generate a realistic Visualization that depicts key moments.

## 2.3 Synthetic Lived Experiences

By exploiting the prompts across the three modules, the pipeline produces three complementary artifacts that together constitute an SLE. The Narrative is a structured first-person account including persona recap, episodic moments, and reflection. The Scene Description distills each episode into visually grounded Scene Descriptions. Finally, the Visualization provides a concrete view of key moments in the Narrative, offering an additional modality through which users can inspect the experience.

## 3 Results

### 3.1 Implementation Details

As a prototype, we used a frontier LLM (*GPT 5.2 Thinking*) for the whole process. Note that each module can be replaced with other open source models for further usability and variability.

### 3.2 Generated SLEs

Since we do not conduct user studies at this stage, our results consist of qualitative examples and artifacts produced by the SLE pipeline. Given the user prompt, the system first concretizes the persona through persona recap. Then, the detailed episode is generated, describing the situation and emotion of the generated self. Finally, the reflection summarizes the first-person overall thoughts on the event. Then, the Scene Description properly summarizes the core moment of the Narrative into short sentences. Finally, generated Visualizations provide illustrative glimpses of the situation. All SLEs may prompt the imagination and help expand the experience. Two examples in Fig. 3 also demonstrate the variability of SLEs under different user-given contexts. These examples illustrate the potential of SLE artifacts as materials for early-stage design, or structured reflection.



<b>User Prompt</b>	A woman in her 50s, a high school teacher, teaching in a classroom, two students making noise	
<b>Narrative</b>	<p><b>Persona Recap</b> I'm a woman in my 50s and a high school teacher. Right now, I'm in my classroom in the middle of a lesson, trying to keep the room focused while two students keep making noise.</p> <p><b>Episode</b> I'm halfway through explaining the day's concept when I hear it again—two voices, low at first, then rising into a rhythm of whispered jokes and chair-scrapes. (...) I don't corner them. I just catch one of them briefly and say, in a low voice, "Tomorrow—start strong. Help me keep the room steady." He doesn't smile, but he nods. It's small. It's enough.</p> <p><b>Reflection</b> What mattered most to me in this experience was protecting the learning space without turning it into a power struggle. I wanted the class to feel steady—like a place where attention and respect aren't optional, and where my calm matters as much as my content. Even on a tired day, I cared most about guiding the room back to focus while still seeing those two students as kids who need direction, not enemies to defeat.</p>	
<b>Scene Description</b>	A sunlit high school classroom mid-lesson: a tired but composed woman teacher in her 50s stands in the aisle with a marker in hand, eyes calmly fixed on two whispering students in the back corner. Other teens sit at desks with notebooks open, some glancing over with half-smiles, while a bright whiteboard behind the teacher glows with partially written notes, creating a tense-but-controlled, quietly authoritative atmosphere.	<b>Visualization</b> 
<b>User Prompt</b>	A man in his 20s, an instagram influencer, working out at a gym, a fan trying to get an autograph	
<b>Narrative</b>	<p><b>Persona Recap</b> I'm a man in my 20s who makes a living as an Instagram fitness influencer. I'm at my usual gym during a midday workout—half training, half filming—trying to stay locked in while still being "on" in public.</p> <p><b>Episode</b> I step off the treadmill and wipe the sweat from my forehead with the hem of my shirt, phone already in my hand. The screen is smudged from fingerprints and chalk dust, but the front camera still makes everything look cleaner than it feels. (...) Then I press stop. I'm not sure what the camera caught—me lifting, or me becoming someone in someone else's story. Both, maybe. I set my grip for the next set, and when I lift again, it feels different. Not heavier. Just... more meaningful.</p> <p><b>Reflection</b> What mattered most to me was realizing my impact wasn't in the perfect posts or the best angles—it was in the quiet, real-life moments where someone told me I helped them change. The autograph wasn't about my name; it was proof to him that his effort was real, and being part of that made me want to carry my influence more carefully—and more sincerely.</p>	
<b>Scene Description</b>	A bright midday gym scene: a sweaty, athletic man in his 20s (Instagram fitness influencer) stands by a squat rack and tripod-mounted phone, towel in hand, while signing a small bent notebook for a nervous younger fan holding his phone nearby. Around them, weight plates and machines blur into the background as a few gym-goers glance over, creating a lively, slightly tense "public spotlight" atmosphere amid clanking metal and workout energy.	<b>Visualization</b> 

Figure 3: Examples of the generated SLEs

### 3.3 Discussion

The initial results suggest that SLE can often produce coherent first-person narratives that appear plausible, and may serve as tentative design artifacts for early-stage ideation. However, our current prototype relies primarily on an LLM's general knowledge, which inherits well-known limitations [9, 13] including hallucination, ultimately constraining realism and reducing the degree of personalization. This also raises ethical concerns about biased or stereotypical depictions when modeling specific identities or social contexts. Despite these limitations, SLE may be useful as a rapid what-if prototyping tool: it enables researchers and designers to quickly explore alternative scenarios and generate early insights, offering a preliminary "glimpse" that can guide subsequent, more rigorous investigation and validation.

## 4 Future Work

Our proposed SLE framework is designed to be extensible, and several promising directions can broaden its practicality and impact.

- **System Usability:** In order to support smoother end-to-end interaction, interface designs that help users specify personas, contexts and desired experience qualities without requiring prompt-engineering expertise are required. In addition, interactive controls for editing narrative constraints, interpretability on why particular events were generated could make SLE more accessible for everyday users.
- **System Composition:** To better personalize experiences, SLE could incorporate a dedicated memory module that ingests real-world data sources such as diaries, notes, or

lifelog [4] streams. A practical research direction is to develop pipelines that (i) compress and summarize noisy personal data into structured representations, (ii) retrieve relevant memories given the current narrative context, and (iii) integrate retrieved content into narrative generation in a controllable way.

- **System Factuality:** Relying solely on generative models can make it difficult to maintain long-horizon coherence and realistic causal dependencies. Building on the memory module above, future work can introduce a RAG [11] based retrieval or causal reasoning [8, 16] layer that explicitly models “if-then” relationships (e.g., routines, social dynamics, constraints, and consequences) and uses them to guide narrative progression. Also, in order to enhance visual artifacts, models such as image-text-to-text [25] can be used to reflect a realistic self in the image. Such mechanisms could improve consistency across episodes, reduce implausible jumps, and better preserve the logic of counterfactual life trajectories.
- **System Richness:** Future systems could generate complementary modalities such as short-form video, audio (voiceover, ambient sound), or storyboard-like sequences to increase immersion and affective engagement. Recent development of text-to-video [21, 26] and any-to-any [18, 22] models can be used.
- **System Validation:** Finally, rigorous evaluation with real users is essential. Future work should conduct controlled and in-the-wild user studies to assess experiential outcomes (e.g., presence, empathy, self-reflection, learning), usability, and potential risks (e.g., emotional discomfort, over-identification). Longitudinal studies may be particularly valuable for understanding whether repeated SLE use produces meaningful changes in perspective-taking or decision-making.

## 5 Conclusion

We proposed the concept of Synthetic Lived Experience (SLE) and presented an initial pipeline and example artifacts to illustrate how first-person alternative-life experiences can be generated, followed by a discussion of key implications, limitations, and future research directions. While SLE is not a substitute for real experience or factual evidence, we suggest that if designed and used responsibly, it can serve as a reflective medium for perspective-taking and may help people better understand one another in an increasingly fragmented world.

## Acknowledgments

This work was supported by Electronics and Telecommunications Research Institute(ETRI) grant funded by the Korean government. [26ZR1100, Research on Intelligent Industrial Convergence].

## References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Raphaël Aubry, Alain Quiamzade, and Laurenz L Meier. 2024. Depressive symptoms and upward social comparisons during Instagram use: A vicious circle. *Personality and Individual Differences* 217 (2024), 112458.
- [3] Ananya Bhattacharjee, Sarah Yi Xu, Pranav Rao, Yuchen Zeng, Jonah Meyerhoff, Syed Ishtiaque Ahmed, David C Mohr, Michael Liut, Alex Mariakakis, Rachel Kornfield, et al. 2025. Perfectly to a Tee: Understanding User Perceptions of Personalized LLM-Enhanced Narrative Interventions. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*. 1387–1416.
- [4] Seungeun Chung, Chi Yoon Jeong, Jeong Mook Lim, Jiyou Lim, Kyoung Ju Noh, Gague Kim, and Hyuntae Jeong. 2022. Real-world multimodal lifelog dataset for human behavior study. *ETRI Journal* 44, 3 (2022), 426–437.
- [5] Melanie C Green and Timothy C Brock. 2000. The role of transportation in the persuasiveness of public narratives. *Journal of personality and social psychology* 79, 5 (2000), 701.
- [6] Riko Horikawa and Tatsuo Nakajima. 2025. BioCanvas: Dynamic Health Visualization through Daily AI-Generated Image and Poem. In *Companion of the 2025 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. 1576–1582.
- [7] Geoff F Kaufman and Lisa K Libby. 2012. Changing beliefs and behavior through experience-taking. *Journal of personality and social psychology* 103, 1 (2012), 1.
- [8] Emre Kiciman, Robert Ness, Amit Sharma, and Chenhao Tan. 2023. Causal reasoning and large language models: Opening a new frontier for causality. *Transactions on Machine Learning Research* (2023).
- [9] Md Tahmid Rahman Laskar, Sawsan Alqahtani, M Saiful Bari, Mizanur Rahman, Mohammad Abdullah Matin Khan, Haidar Khan, Israt Jahan, Amran Bhuiyan, Chee Wei Tan, Md Rizwan Parvez, et al. 2024. A systematic survey and critical review on evaluating large language models: Challenges, limitations, and recommendations. *arXiv preprint arXiv:2407.04069* (2024).
- [10] Ana Levordashka and Sonja Utz. 2016. Ambient awareness: From random noise to digital closeness in online social networks. *Computers in human behavior* 60 (2016), 147–154.
- [11] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems* 33 (2020), 9459–9474.
- [12] Hazel Markus and Paula Nurius. 1986. Possible selves. *American psychologist* 41, 9 (1986), 954.
- [13] Andrea Matarazzo and Riccardo Torlone. 2025. A survey on large language models with some insights on their capabilities and limitations. *arXiv preprint arXiv:2501.04040* (2025).
- [14] Adrian Meier and Benjamin K Johnson. 2022. Social comparison and envy on social media: A critical review. *Current opinion in psychology* 45 (2022), 101302.
- [15] Joon Sung Park, Carolyn Q Zou, Aaron Shaw, Benjamin Mako Hill, Carrie Cai, Meredith Ringel Morris, Robb Willer, Percy Liang, and Michael S Bernstein. 2024. Generative agent simulations of 1,000 people. *arXiv preprint arXiv:2411.10109* (2024).
- [16] Judea Pearl. 2009. *Causality*. Cambridge university press.
- [17] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10684–10695.
- [18] Zineng Tang, Ziyi Yang, Chenguang Zhu, Michael Zeng, and Mohit Bansal. 2023. Any-to-any generation via composable diffusion. *Advances in Neural Information Processing Systems* 36 (2023), 16083–16099.
- [19] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [20] Yu-Min Tseng, Yu-Chao Huang, Teng-Yun Hsiao, Wei-Lin Chen, Chao-Wei Huang, Yu Meng, and Yun-Nung Chen. 2024. Two tales of persona in llms: A survey of role-playing and personalization. *arXiv preprint arXiv:2406.01171* (2024).
- [21] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. 2025. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314* (2025).
- [22] Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. 2024. Next-gpt: Any-to-any multimodal llm. In *Forty-first International Conference on Machine Learning*.
- [23] Quejie Xie, Qiming Feng, Tianqi Zhang, Qingqiu Li, Linyi Yang, Yuejie Zhang, Rui Feng, Liang He, Shang Gao, and Yue Zhang. 2024. Human simulacra: Benchmarking the personification of large language models. *arXiv preprint arXiv:2402.18180* (2024).
- [24] Xiangzhe Yuan, Jiajun Wang, Qian Wan, and Siying Hu. 2025. A Day in Their Shoes: Using LLM-Based Perspective-Taking Interactive Fiction to Reduce Stigma Toward Dirty Work. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. 1341–1359.
- [25] Mao Zheng, Zheng Li, Tao Chen, Mingyang Song, and Di Wang. 2025. HY-MT1.5 Technical Report. *arXiv:2512.24092* [cs.CL]. <https://arxiv.org/abs/2512.24092>
- [26] Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. 2024. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404* (2024).