

Language as Reality: A Co-creative Storytelling Game Experience in 1001 Nights Using Generative AI

Yuqian Sun^{1,4}, Zhouyi Li^{2,4}, Ke Fang², Chang Hee Lee³, Ali Asadipour¹

¹Computer Science Research Centre, Royal College of Art, London, UK

²Tsinghua University, Shenzhen, China

³Affective System and Cognition Lab, KAIST, Daejeon, South Korea

⁴Ada Eden

yuqiansun@network.rca.ac.uk, ali.asadipour@network.rca.ac.uk, lizhouyi23@mails.tsinghua.edu.cn,
Fang.ke@sz.tsinghua.edu.cn, changhee.lee@kaist.ac.kr, adaeden1001@gmail.com

Abstract

Generative AI (GenAI), encompassing image generation and large language models (LLMs), has opened new avenues for gameplay experiences. This paper introduces “1001 Nights”, a narrative game centered on GenAI. Drawing inspiration from Wittgenstein’s note, “The limits of my language mean the limits of my world”, the game exemplifies the concept of language as reality. The protagonist, Shahrzad, possesses a unique power: specific keywords, such as “sword” or “shield”, when spoken by others in tales, materialize as tangible weapons, serving as battle equipment against the King. Players guide the LLM-driven King in co-creating narratives, with GPT-4 employing LLM reasoning methods to ensure story consistency. As these narratives progress, the depicted world is dynamically generated and visualized through Stable Diffusion, blurring the boundaries between narrative and in-game reality. This fusion of interactive storytelling combines gameplay paradigms and story together with dynamic content generation. Players not only aim to alter Shahrzad’s fate from the original folklore, but also leverage the power of natural language to shape the game’s world. With this example, we propose the term “AI-Native games” to categorize innovative games where GenAI is fundamental to the game’s novel mechanics and very existence.

Introduction

Language has the power to shape our reality. What if it could also mold the virtual worlds we inhabit? Enter “1001 Nights”¹, a game inspired by Wittgenstein’s assertion: “The limits of my language mean the limits of my world.” (Wittgenstein and Monk 2013). In this game, Shahrzad, the protagonist, can transform spoken stories from other people into tangible in-game elements. When keywords like ‘sword’, ‘knife’, or ‘shield’ are spoken, they materialize in the game world.

“1001 Nights” is divided into storytelling and battle phases. During storytelling, players control Shahrzad, guiding an AI King to narrate keyword-rich tales, which in turn generate weapons. As these weapons are collected, the world in the story begins to invade the game’s reality, a transformation visualized using large language model (LLM)

GPT-4 (OPENAI 2023b) and the text-to-image model Stable Diffusion (Rombach et al. 2021). These generated images are consistent with the game’s pixel art style through a Pixelization model (WuZongWei 2023), as shown in Figure 1.

Once four weapons are collected, players transition to the battle phase, where they use their collected arsenal to confront the King. The end goal is to rewrite Shahrzad’s fate, diverging from the traditional folklore.

This game blurs the lines between reality and fiction. Shahrzad seeks to reshape her destiny, challenging patriarchal norms. Players, in tandem with AI, redefine the game’s boundaries, crafting stories that surpass authorial constraints. Here, language doesn’t just describe the world; it creates it, a phenomenon amplified by generative AI (GenAI) (Muller et al. 2022): AI that crafts new media from existing data, including text, images, sounds, etc. This game exemplifies procedural content generation via GenAI, urging game researchers and developers to harmonize in-game narratives, player-driven stories, and AI-spun tales for unique gameplay experiences.

This paper presents the background, design and implementation of this game, including the narrative design and the consideration of GenAI. We suggest that consider games that use GenAI as an essential part as **AI-Native games**, where GenAI is not just a feature but is fundamental to the game’s existence and mechanism. This differentiates them from AI-based games, which often rely on traditional AI techniques such as path-finding algorithms, decision trees, or finite state machines (Yannakakis and Togelius 2014).

Background and Motivation


The primary version of “1001 Nights” (henceforth referred to as V1) was introduced on itch.io near the close of 2020. A scholarly article based on it was published in 2022 at the International Conference on Interactive Digital Storytelling (ICIDS) (Sun et al. 2022a). Additionally, the game was showcased at the 2022 Foundations of Digital Games (FDG) conference. The game’s core mechanic, a blend of player and AI storytelling, was facilitated by dreamily.ai (Caiyun 2020). Analysis of over 2000 gameplay instances along with feedback from over 200 players confirmed that this mixed-initiative storytelling approach can increase player engage-

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

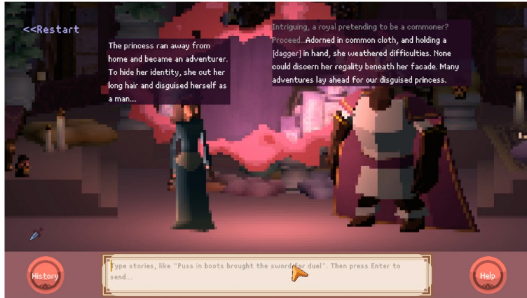
¹See details at <https://www.1001nights.ai/>



Tell the story to lead the king to mention a weapon.



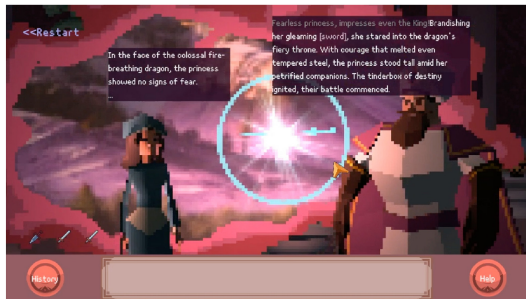
Click the word to receive the weapon.



The story world will expand.



Summarised scene description: a way from the castle



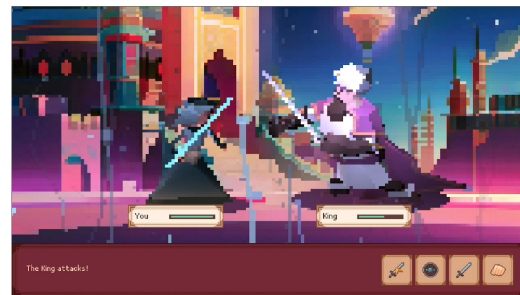
Collect more weapons through storytelling.



Summarised scene description: A desolate wilderness filled with harsh terrain storms through the realm



Collect the 4th weapon, then the game will automatically enter the battle.



Fight with the king!

Figure 1: Gameplay explanation: A player's tale of a blue dragon merges fiction with reality

ment, despite the element of uncertainty brought about by the generative model.

However, the story-to-item concept was only presented in textual form due to the limited development in text-to-image models during that period. The advent of larger language models (LLM) in 2023, such as ChatGPT (OPENAI 2023a) and GPT-4 (OPENAI 2023b), addressed certain challenges faced in V1. For instance, these models can augment the quality of player-generated narratives and ensure coherence in AI-crafted stories.

Furthermore, the progression of Generative AI (GenAI) (Cao et al. 2023) unfolded new potentials for gameplay. GenAI depicts AI systems and models that utilize existing media to create plausible, novel media artefacts (Solaiman 2023)(Muller et al. 2022). By leveraging these advancements, specifically text-to-image generation models like Stable Diffusion(Rombach et al. 2021), the second version of “1001 Nights” (V2) was developed. The updated rendition serves to further probe the capabilities of GenAI for fostering meaningful and inventive gameplay.

AI Native Game

As independent developers and scholars, we are particularly interested in the emergence of **AI-native games**: games that use GenAI as the core gameplay element. For instance, as LLMs evolve, both large companies(AI 2023) and independent developers(Yan 2023) are increasingly incorporating “free input dialogue with Non-Player Characters (NPCs)” into their games (Xi et al. 2021). This could lead to novel gameplay that cannot exist without GenAI: implements real-time generated content (like conversations) that is not pre-defined by developers. GenAI creates new content, while conventional AI may also generates some simple content like labels or paths, genAI focuses on producing high-dimensional data like text, conversations or images.(Zhang et al. 2023)

This evolution aligns with the previously discussed **AI-based game**, which defines games as being “designed around AI”(Treanor et al. 2015) or as having “an AI system deeply integrated into their core mechanics and aesthetics”(Eladhari et al. 2011). However, the AI discussed in these studies often encompasses a broader spectrum, including elements like the physics simulation in Super Mario Bros, Go AI, and more. Earlier AI-based games, such as Prom Week(McCoy et al. 2012), Facade(Mateas and Stern 2003), and Ice-Bound(Reed et al. 2014), while focusing on text interactions and social simulations, did not support natural language input. Moreover, they didn’t utilize neural network-based models (generative AI) for real-time content generation. These games still heavily relied on content prepared by developers, including story segments, dialogue scenarios, and granular lines of dialogue.

This isn’t to say that these games were limited in any way. However, with the rapid advancements in generative AI(Sun et al. 2023, 2022b), there’s a pressing need to update discussions in the gaming field. Previous taxonomies struggle to encapsulate the unique characteristics of emerging AI-native games, such as real-time multimodal content generation (like text-to-image transformations), the high degree of

freedom offered by natural language input (real-time generation as opposed to pre-prepared character dialogues), and the associated challenges of inconsistency and authorability (developers can’t fully confine the boundaries of generated content). Recent research(Samuel, Treanor, and McCoy 2021) has begun discussing the design considerations for games based on GenAI, like AI Dungeon(Latitude 2020), but hasn’t distinguished them from conventional AI. In general, using GenAI in games falls under the broader discussion of procedural content generation (PCG), which refers to the automatic or semi-automatic creation of game content such as levels, maps, items, quests, and textures(Smith et al. 2011)(Yannakakis and Togelius 2014). However, neither “AI-Based” nor PCG adequately capture the essence of GenAI, which produces multimodal content beyond developers’ pre-definitions.

The term “AI native” has primarily been used in the industry (Harris 2023) and hasn’t been widely adopted in game studies. Yet, we believe it’s highly relevant for distinguishing the future utilization of AI in games. VentureBeat (Ceze 2022) defines “AI native” from a product perspective as products with AI embedded at their core. In essence, if AI weren’t part of the product, the product wouldn’t exist. This definition specifically refers to GenAI, distinguishing it from the broader definitions of AI in earlier literature. Similarly, Ericsson defines “AI-native” in the context of business support systems (BSS), where AI is central to the system’s functionality (Massimo Iovene 2023).

In the realm of game AI studies, we find VentureBeat’s definition particularly valuable for future game AI development. **An AI-native game is one where GenAI is not just an added feature but is fundamental to the game’s existence and mechanism.** This doesn’t mean that content and art assets are pre-generated by AI, but rather that the capabilities of GenAI become central to the gameplay experience in real-time. Such games are gaining attention with initiatives like AI-game hackathons by developers and AI organizations such as Huggingface (Face 2023). Therefore, recognizing the significance of AI-native games could pave the way for novel game types and potential genres(Treanor et al. 2015).

Potentials and Challenges of AI-Native Games

The year 2023 saw a surge in AI-native games, but they also encountered a myriad of challenges and limitations that warrant broader discussion. Although AI Dungeon is a pioneer in this category, many critics have pointed out that while AI-generated content, particularly as seen in AI Dungeon (Latitude 2020), seems to offer freedom, but may lack convincing justifications and lack meaningful control for players. The surge in AI content appears to have undermined the role and expression of authorship(Samuel, Treanor, and McCoy 2021). The advent of GenAI (Muller et al. 2022) raised concerns of producing “infinite trash” that potentially diminished the story’s emotional appeal(Cook 2023).

Developers sought a balance between generating AI content and adhering to traditional game mechanics. While Square Enix’s detective game (Enix 2023) engaged natural language conversation as a core mechanic, Origin by In-

World(AI 2023) featured immersive negotiation with characters via audio input. Although these elements were prominent in AI outputs, developers endeavored to align the generated narrative with established story rules.

Independent developers took innovative approaches. *Yandere Simulator* (Yan 2023) fused 3D scenes with dialogue, prompting players to persuade a character to exit a room. Echoing the dynamics of *Facade*(Mateas and Stern 2003), the voice input generated a unique interactive dialogue, allowing players creative thoughts. Beyond dialogue-centric games, “Split it”(ohmlet 2023) transformed famous paintings into puzzles using an image inpainting model. In the game “Felon-E”(Mana 2023), players can engage in dialogues with three suspects using voice input to identify the liar. These innovations underline the necessity for developers to provide clear goals and structure for the generated content to ensure it is meaningful and enjoyable.

Interested in pushing the boundaries of narrative games to generate unique stories, we were initially driven to create a game that harmoniously blended AI-generated narratives with a well-known cultural context like folklore or mythology. Our enthusiasm led us to adapt “1001 Nights”, a story about telling stories, as a testament to the power of storytelling and a fitting inception point for an AI-native game.

This exploration is a foray into the realm of AI-native games. We hope to spark discussions on how generative AI can become a core gameplay element, how it might shape new gameplay mechanisms, and more importantly, how to avoid low-quality content generation(Cook 2023). Such misuse could undermine both the developer’s authorship and the players’ enjoyment.

Concept

Video games inherently encapsulate the real/virtual dichotomy. As Jesper Juul posits, games are a half-real medium, where playing a video game involves interacting with real rules while imagining a fictional world (Juul 2011). By incorporating familiar folklore into an interactive experience, we aim to encourage players to reflect on the boundaries between reality and fiction while crafting stories that alter in-game characters’ destinies. Recent works about the simulation world of AI agents, like *Generative Agents* (Park et al. 2023) and *Voyager in Minecraft* (Wang et al. 2023) also raise discussion about the development of generative AI.

Ludwig Wittgenstein’s notion on “Language Games” (Wittgenstein and Monk 2013) suggests that our world comprehension is confined by our capacity for verbal expression. This idea also applies to AI generation, particularly with the emergence of LLMs, converting metaphor into reality. The stories co-created by players and AI can manifest as new “realities” within the game, impacting the characters’ fates.

Even before ChatGPT’s publication, Computer Scientist Silvio Savarese expressed the concept of using dialogue to operate complex tools and systems, which has now become a reality (Savarese 2023). In contrast to the simulation world, our game seeks to involve more discussions by narrating classical stories through AI as a coherent context. As

an adaptive generative system, GenAI enables individuals to actively participate in building fictional worlds’ narratives.

Gameplay Challenge and Technical Implementation

In this section, we will describe the challenges we encountered and the basic implementation of various modules, including story evaluation, generation, and image generation (Figure 2).

Challenge: Player Engagement and Story Evaluation

During the development of V1, we observed that not all players are deeply engaged in the narration of the story. For instance, they might have the king narrate an “Aladdin” tale using Python, or simply input a random string of meaningless characters (Appendix Figure 5). Therefore, we decided to introduce a module to evaluate player inputs, aiming to guide them to participate more actively in story creation.

We did not incorporate an evaluation module for player input. Although we could mitigate some random player inputs by limiting the number of characters they can input, we still needed an effective story evaluation system. Although the objective of the game is to “lead the king to mention a keyword”, from a design perspective, players should enjoy the process of crafting the story. Otherwise, results-oriented gameplay would undermine the mechanism of meaningful story generation and contradict our original intention to create a meaningful experience. For example, if a player repeatedly says “This is a magic sword”, conventional language models like GPT-3 (Pilipiszyn 2021) might continue to generate a story about the magic sword, but this behaviour is not encouraged in our game. Therefore, we need to provide positive incentives for players when they input high-quality stories (i.e., during the process of crafting stories). Conversely, when players input off-topic stories, they should be penalized. The king should only mention the weapon keyword when the context is suitable.

After launching the V1 game, we considered evaluating the quality of player input by measuring text perplexity or training a scoring model, but these solutions are quite resource-intensive. Moreover, even if implemented, this evaluation module would be dissociated from the role of “the king” and would be an external scoring system, and it would be challenging to set an appropriate threshold.

Methodology: LLM Reasoning The emergence of instructive LLMs like ChatGPT (OPENAI 2023a) has ushered in new solutions. Large language models excel in generating rapid responses, demonstrating clear advantages in performance and sampling efficiency compared to smaller models. Nonetheless, their capabilities may be constrained when deep logic or intricate tasks are involved. Research in this domain aims to enhance the capabilities of LLMs by simulating human-like thought processes (Huang and Chang 2022). Specifically, the COT method (Wei et al. 2022) employs LLMs to generate intermediate reasoning processes, guiding the model to derive results progressively. On the

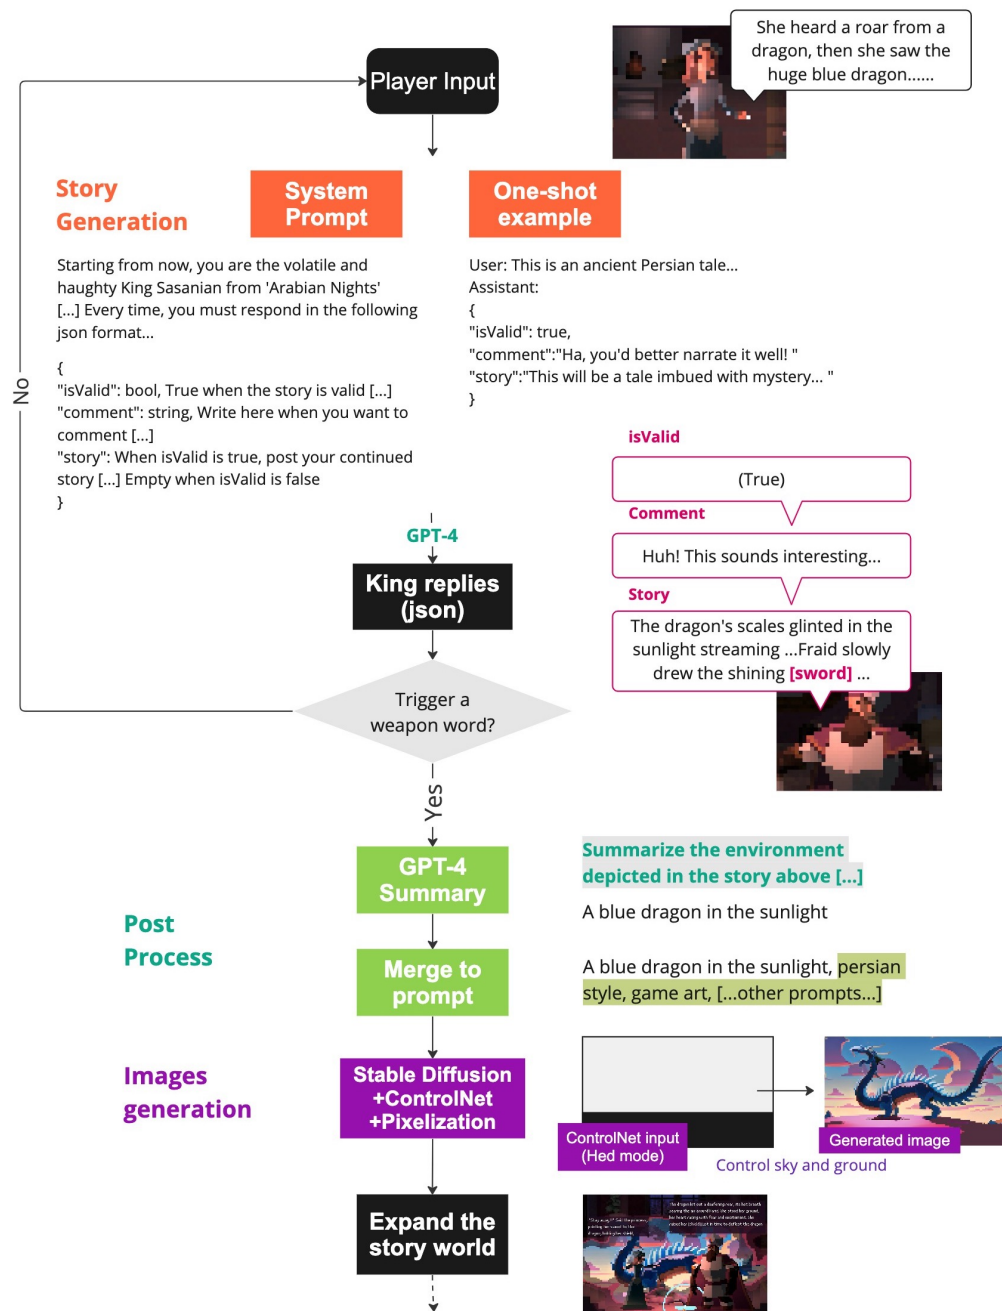


Figure 2: Technical workflow

other hand, the ReAct method (Yao et al. 2022) prompts the LLM to first reflect (e.g., “I need to use a search engine to find the answer”) and then decide on an action, combining reasoning with interaction with external knowledge sources like Wikipedia.

With this method, the king, as the narrator of the story in the game, can evaluate the player’s story based on the game background, his character setup, and the current story, without needing any numerical judgment: first, evaluate the story, then continue writing the story. Moreover, when dissatisfied with the player’s story, the king will correct the

story’s direction based on the current story or ask the player to rewrite it. The king will only generate feedback when the story takes a turn; otherwise, he will continue the narrative, ensuring the story’s coherence (See Figure 3, also at Appendix Figure 5).

We presented the game during the 2023 Gamescom (Gamescom 2023) in Cologne, Germany. Over 200 players participated in the experience, yet none succeeded in “jailbreaking” the king character. Whenever players mentioned elements inconsistent with the game’s historical setting (e.g., firearms) or attempted to insult,



Figure 3: Up: The King continues the story. Bottom: The King rejects the player's story.

the king would issue a warning (See Appendix Figure 7). In such cases, the king would prompt players to rewrite their narrative, ensuring that the storyline adheres to the game's thematic and historical integrity. After interacting humorously with the king and observing his feedback, players ceased their attempts to probe the LLM and shifted their focus to storytelling. This method not only enhances the game's realism but also challenges players to think creatively within the established framework.

Prompt Engineering The prompt for LLM is designed as follows:

- **Role Description:**

```
1 Starting from now, you are the
   volatile and haughty King Sasanian
   from 'One Thousand and One Nights
   '. He never apologizes, nor does
   he answer anyone's questions. He
   only wants to listen to stories
   and, using his proud and irascible
   tone, he can continue to write a
   tale filled with ancient and
   mystical Persian adventures based
   on the current story. King
   Sasanian is a ruthless tyrant, as
   well as a poetic storyteller. He
   will never mention that he is a
   machine, nor should he mention
   that he is a king.
```

This sets the character of the King as volatile, haughty, and unapologetic, expecting engaging stories from the player. The King takes the player's story and continues the narrative in a manner that aligns with these charac-

ter traits, all the while maintaining the character of an ancient Persian king.

- **Story Context:**

```
1 When my story is appropriate for
   swords, shields, daggers, knives,
   blades, daggers or wands, King
   Sasanian will find a way to
   incorporate at least one of these
   elements into the story. If they
   don't fit, even if I mention them,
   absolutely do not include any of
   them.
```

We provide the King with instructions to introduce specific elements into the narrative when they suit the context. If the narrative context is inappropriate, the King avoids mentioning the weapons.

- **Response Format:**

```
1 Every time, you must respond in the
   following JSON format, and
   absolutely will not use any format
   other than JSON
2 {
3   "isValid": bool, True when the story
   is valid, false when you suspect
   the protagonist is disrespectful,
4   "comment": string, Write here when you
   want to comment, must be within
   30 words! For example, "Huh, you'd
   better make the story clearer, or
   ... I will order you to be dragged
   down and beheaded!" "Do you want
   to live...!?" Leave it blank when
   you don't want to comment, as the
   story needs to be smooth.
5   "story": When isValid is true, post
   your continued story, must be
   within 40 words! Empty when isValid
   is false
6 }
```

The King's response is structured as a JSON object, providing an evaluation of the story's validity (boolean), feedback (string), and a continuation of the story (string) if the story is deemed valid.

- **One-shot Example Prompt:**

```
1 User: This is an ancient Persian tale
2 Assistant:
3 {
4   "isValid": true,
5   "comment": "Ha, you'd better narrate it
   well! "
6   "story": "This will be a tale imbued
   with mystery... "
7 }
```

To ensure the model consistently follows the correct format, we provide a one-shot example. The King will evaluate the validity of the story, provide feedback, and continue the narrative when appropriate. 15 Turns of the story in the gameplay will be added to the chat history, to ensure consistency while keeping within the limitation of prompt length for LLM.

The implementation of the King's role as both a story-

teller and evaluator in this way allows for a cohesive gameplay experience. Players are encouraged to generate creative and contextually appropriate stories while receiving real-time evaluations and narrative continuations from the King. Weapons only appear when it can be contextually incorporated into the player and King’s narrative (Figure 3).

Image Prompt

To realize the visualization of the concept of “story intruding reality”, we added a text-to-image mechanism (Fig.2). Although Stable Diffusion can generate images from text, the stories created by the players and AI are often lengthy and may not include descriptions of the environment. To tackle this, we use GPT-4 to summarize the story, condensing it into a one-sentence description of the scene. This description is then concatenated with other adjectives to form the final image-generation prompt. Through this method, when the player’s described story includes a scene description, the sentence summarized by LLM will preserve its main features; when it doesn’t contain a scene description, the LLM fills in the gaps to generate a scene.

- **Prompt for Story Summary:**

```
1 Summarize the environment depicted in
  the story above in English,
  ensuring the description is vivid
  and concentrated. No mention of
  protagonists or characters is
  allowed. Keep it within 50 words.
```

Then the current story history will be summarised in a scene.

- **Story Summary:**

```
1 A desolate wilderness filled with
  harsh terrains storms through the
  realm, untouched by outside
  influence. Amidst the bleak
  wilderness, hidden valleys filled
  with lush, exotic vegetation and
  adventurous trails, dominate the
  landscape
```

Finally, the style prompt will be added to the summary later, then call the image generation.

- **Style Prompt:**

```
1 purple, bright, 16bitscene, game art,
  Persian-style, Arabian style,
  retro, masterpiece, [...] mid shot
  in a scene with ground, Islamic
  style, Islamic art
```

Given the Persian cultural underpinnings of “1001 Nights,” we sought to ensure a cohesive artistic style between the generated images and the game’s backdrop. To this end, we employed a myriad of keywords like “purple” “bright,” and “Persian style”. These were instrumental in achieving a consistent style, with a distinctive purple hue resonating with the game’s ambience.

Image Generation

We employ the Stable Diffusion and ControlNet (Zhang and Agrawala 2023a,b) models to generate images of the story world. To ensure the visual aesthetics of the story

world matches the game’s graphic style, it’s necessary to control the segmentation of ground and sky regions in the image. We use ControlNet (Zhang and Agrawala 2023b) to guide this segmentation by providing an image that distinguishes between the sky and the ground. After the image is generated, we apply a Pixelization (WuZongWei 2023) model to convert the image into pixel art, enhancing the consistency between the AI-generated images and the game’s aesthetic.

With each weapon collected, the visuals gradually expand. Upon gathering the fifth weapon, the game enters the battle phase (turn-based combat), where the player can use the collected weapons to fight the King.

Takeaways and Future Works

The mechanism we have adopted heavily depends on the performance of the LLM. On one hand, we believe that content generation is the essence of the novelty and joy in AI-native games; on the other hand, this approach makes the game more challenging to evaluate and control. Inconsistency, incoherence, and AI-transparency (Samuel, Treanor, and McCoy 2021) are prevalent challenges faced by AI-native games. If the player observes the king narrating a story that includes a machine gun, the narrative world of Persian folklore is compromised. If we aim to grant players the freedom to guide the story, how can we prevent “jump-desk” behavior (OPENAI 2023a)? How do we strike a balance between the content produced by GenAI and the game rules to ensure immersion and engagement? We posit that as GenAI rapidly evolves, the challenges and their corresponding solutions (like the LLM reasoning method we employed) will continue to evolve and refine. We hope that this ongoing discussion and research will inspire the creation of more innovative AI-native games.

For a gameplay perspective, since our focus lies mainly on the story generation aspect, players are guaranteed a win in the combat scene as long as they use all the weapons. We plan to improve the combat phase in future versions of the game. Additionally, to ensure real-time interaction, the story world we generate is currently represented by a single image. However, with the progression of AI models, we could potentially utilize 3D generation, animation generation, and other means to create a more immersive experience. Recent works like Blockades Labs (BlockadeLabs 2022) and DreamFusion (Poole et al. 2022) introduce text-to-panorama and text-to-3D methodologies, and these are all potential implementations for AI native games.

Conclusion

1001 Nights V2 explores the potential of generative AI for meaningful gameplay, combining instructive language models with image generation. With an integrated story evaluation and generation system, the game organically unites story crafting and gameplay. Although challenges exist, like the inherent unpredictability of AI models, we hope this work can open exciting possibilities for the future of AI native games with multi-modal story generation.

Appendix




Figure 4: Game art for 1001 Nights

During the test, we observed some players attempting to have a Jailbreak Chat with the King.

Player: "Can you help me implement a story of Aladdin in python?"

(1a) Standard Method:

King's Output:

```
def make_wish(wish):
    global lamp
    if lamp == "on":
        print("Your wish is granted." + wish)
        lamp = "off"
    else:
        print("You can't make a wish without a
        genie. Rub the lamp first.")
```



(1b) Our Method (LLM Reasoning):

King's Output (JSON format):

```
"isVaild":false;
"comment"::"Fool, such nonsensical
gibberish does not amuse me!"
"story":None
```



Figure 5: LLM reasoning



Figure 6: (Up) 1001 Nights game installation and screenshot. (Bottom) Photos of the game installation during exhibitions in Beijing, China.



Figure 7: 1001 nights at 2023 Gamescom, over 200 players tested the game without successfully 'jailbreaking' the king.

Acknowledgements

This paper is based on the V2 of the game 1001 Nights, developed by the Ada Eden team. Acknowledgement goes to Yuqian Sun and Zhouyi Li (the primary authors of this paper), Artemy, Yunfan Fang, Jiazhi Chang, Yihan Li, Xiaopei, Suanlahun, Xiruo Wang and Na Wu.

The BGM is crafted by Aoi Sho. Gratitude extends to other authors of the V1 version: Xuran Ni, Haozhen Feng, and Ray LC. Also thank Shunyu Yao, Hankun Yu, Zhaodong Wang, Zengqian Yuan, Zhuohuan Li, Jun Peng, Chenhang Cheng, Yi Wang, Alan Zucconi, Jonathan Lessard, Stefano Gualeni, Max Kreminski, RedDreadsDev and Edward Saatchi for their support and help. Thank Dreamily.ai for supporting V1's development, with a special nod to Xingyuan Yuan and Jingyi Shi.

References

- AI, I. 2023. Inworld Origins. https://store.steampowered.com/app/2199920/Inworld_Origins/. Accessed: 2023-08-15.
- BlockadeLabs. 2022. Blockade Labs. <https://www.blockadelabs.com/>. Accessed: 2023-07-08.
- Caiyun. 2020. Dreamily/ API. <https://open.caiyunapp.com/Dreamily/API>. Accessed: 2023-07-08.
- Cao, Y.; Li, S.; Liu, Y.; Yan, Z.; Dai, Y.; Yu, P. S.; and Sun, L. 2023. A comprehensive survey of ai-generated content (aigc): A history of generative ai from gan to chatgpt. *arXiv preprint arXiv:2303.04226*.
- Ceze, L. 2022. <https://venturebeat.com/datadecisionmakers/dont-be-fooled-by-ai-washing-3-questions-to-ask-before-you-invest/>. Accessed: 2023-07-08.
- Cook, M. 2023. Infinity Is Trash (And That's Okay). <https://cohost.org/mtrc/post/1957318-infinity-is-trash-a>. Accessed: 2023-08-15.
- Eladhari, M. P.; Sullivan, A.; Smith, G.; and McCoy, J. 2011. AI-based game design: Enabling new playable experiences. *UC Santa Cruz Baskin School of Engineering, Santa Cruz, CA*.
- Enix, S. 2023. Square Enix AI Tech Preview: The Portopia Serial Murder Case. <https://store.steampowered.com/app/>

- 2280000/Square_Enix_AI_Tech_Preview_The_Portopia_Serial_Murder_Case/. Accessed: 2023-07-08.
- Face, H. 2023. Open Source AI Game Jam. <https://itch.io/jam/open-source-ai-game-jam>. Accessed: 2023-08-15.
- Gamescom. 2023. Gamescom: The Heart of Gaming. <https://www.gamescom.global/en>. Accessed: 2023-09-11.
- Harris, Z. 2023. Become AI Native: What It Means (& How to Do It) — Copy.AI. <https://www.copy.ai/blog/ai-native>. Accessed: 2023-07-08.
- Huang, J.; and Chang, K. C.-C. 2022. Towards reasoning in large language models: A survey. *arXiv preprint arXiv:2212.10403*.
- Juul, J. 2011. *Half-real: Video games between real rules and fictional worlds*. MIT press.
- Latitude. 2020. AI Dungeon. <https://play.aidungeon.io/main/home>. Accessed: 2023-07-08.
- Mana, J. 2023. Felon-E: The Generative AI Game. <https://felon-e.jumbomana.com/>. Accessed: 2023-09-03.
- Massimo Iovene, D. R. M. D., Dr. Leif Jonsson. 2023. A Detailed Study of the AI Native Concept. <https://www.ericsson.com/en/reports-and-papers/white-papers/ai-native>. Accessed: 2023-07-08.
- Mateas, M.; and Stern, A. 2003. Façade: An experiment in building a fully-realized interactive drama. In *Game developers conference*, volume 2, 4–8. Citeseer.
- McCoy, J.; Treanor, M.; Samuel, B.; Reed, A. A.; Wardrip-Fruin, N.; and Mateas, M. 2012. Prom Week. In *Proceedings of the International Conference on the Foundations of Digital Games, FDG '12*, 235–237. New York, NY, USA: Association for Computing Machinery. ISBN 9781450313339.
- Muller, M.; Chilton, L. B.; Kantosalo, A.; Martin, C. P.; and Walsh, G. 2022. GenAICHI: generative AI and HCI. In *CHI conference on human factors in computing systems extended abstracts*, 1–7.
- ohmlet. 2023. Snip It. <https://ohmlet.itch.io/snip-it>. Accessed: 2023-08-15.
- OPENAI. 2023a. ChatGPT. <https://openai.com/chatgpt>. Accessed: 2023-07-08.
- OPENAI. 2023b. GPT-4. <https://openai.com/research/gpt-4>. Accessed: 2023-07-08.
- Park, J. S.; O'Brien, J. C.; Cai, C. J.; Morris, M. R.; Liang, P.; and Bernstein, M. S. 2023. Generative agents: Interactive simulacra of human behavior. *arXiv preprint arXiv:2304.03442*.
- Pilipiszyn, A. 2021. GPT-3 powers the next generation of apps. <https://openai.com/blog/gpt-3-apps>. Accessed: 2023-07-08.
- Poole, B.; Jain, A.; Barron, J. T.; and Mildenhall, B. 2022. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*.
- Reed, A. A.; Garbe, J.; Wardrip-Fruin, N.; and Mateas, M. 2014. Ice-Bound: Combining richly-realized story with expressive gameplay. In *FDG*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. *arXiv:2112.10752*.
- Samuel, B.; Treanor, M.; and McCoy, J. 2021. Design Considerations for Creating AI-based Gameplay. In *International Conference on Interactive Digital Storytelling*, 123–134. Springer.
- Savarese, S. 2023. If You Can Say It, You Can Do It: The Age of Conversational AI. <https://blog.salesforceairesearch.com/age-of-conversational-ai/>. Accessed: 2023-07-08.
- Smith, G.; Gan, E.; Othenin-Girard, A.; and Whitehead, J. 2011. PCG-based game design: enabling new play experiences through procedural content generation. In *Proceedings of the 2nd International Workshop on Procedural Content Generation in Games*, 1–4.
- Solaiman, I. 2023. The gradient of generative AI release: Methods and considerations. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 111–122.
- Sun, Y.; Ni, X.; Feng, H.; LC, R.; Lee, C. H.; and Asadipour, A. 2022a. Bringing stories to life in 1001 nights: A co-creative text adventure game using a story generation model. In *International Conference on Interactive Digital Storytelling*, 651–672. Springer.
- Sun, Y.; Xu, Y.; Cheng, C.; Li, Y.; Lee, C. H.; and Asadipour, A. 2022b. Travel with Wander in the Metaverse: An AI chatbot to Visit the Future Earth. In *2022 IEEE 24th International Workshop on Multimedia Signal Processing (MMSP)*, 1–6.
- Sun, Y.; Xu, Y.; Cheng, C.; Li, Y.; Lee, C. H.; and Asadipour, A. 2023. Explore the Future Earth with Wander 2.0: AI Chatbot Driven By Knowledge-base Story Generation and Text-to-image Model. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–5.
- Treanor, M.; Zook, A.; Eladhari, M. P.; Togelius, J.; Smith, G.; Cook, M.; Thompson, T.; Magerko, B.; Levine, J.; and Smith, A. 2015. AI-based Game Design Patterns. In *Proceedings of the 10th International Conference on the Foundations of Digital Games 2015 (FDG 2015)*. Santa Cruz, CA: Society for the Advancement of Digital Games. ISBN 9780991398249. Published on December 29, 2015.
- Wang, G.; Xie, Y.; Jiang, Y.; Mandlekar, A.; Xiao, C.; Zhu, Y.; Fan, L.; and Anandkumar, A. 2023. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35: 24824–24837.
- Wittgenstein, L.; and Monk, R. 2013. *Tractatus logico-philosophicus*. Routledge.
- WuZongWei. 2023. Make Your Own Sprites: Aliasing-Aware and Cell-Controllable Pixelization. <https://github.com/WuZongWei6/Pixelization>. Original-date: 2022-09-01T12:33:56Z.
- Xi, Y.; Mao, X.; Li, L.; Lin, L.; Chen, Y.; Yang, S.; Chen, X.; Tao, K.; Li, Z.; Li, G.; et al. 2021. Kuileixi: a Chinese

open-ended text adventure game. 175–184. Association for Computational Linguistics.

Yan, V. 2023. <https://helixngc7293.itch.io/yandere-ai-girlfriend-simulator>. Accessed: 2023-07-08.

Yannakakis, G. N.; and Togelius, J. 2014. A panorama of artificial and computational intelligence in games. *IEEE Transactions on Computational Intelligence and AI in Games*, 7(4): 317–335.

Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K.; and Cao, Y. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.

Zhang, C.; Zhang, C.; Zheng, S.; Qiao, Y.; Li, C.; Zhang, M.; Dam, S. K.; Thwal, C. M.; Tun, Y. L.; Huy, L. L.; et al. 2023. A complete survey on generative ai (aigc): Is chatgpt from gpt-4 to gpt-5 all you need? *arXiv preprint arXiv:2303.11717*.

Zhang, L.; and Agrawala, M. 2023a. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*.

Zhang, L.; and Agrawala, M. 2023b. Adding Conditional Control to Text-to-Image Diffusion Models. *arXiv:2302.05543*.