

# PAYADOR: A Minimalist Approach to Grounding Language Models on Structured Data for Interactive Storytelling and Role-playing Games

Santiago Góngora<sup>1</sup>   Luis Chiruzzo<sup>1</sup>   Gonzalo Méndez<sup>2</sup>   Pablo Gervás<sup>2</sup>

<sup>1</sup> Instituto de Computación, Facultad de Ingeniería, Universidad de la República, Uruguay

<sup>2</sup> Facultad de Informática, Universidad Complutense de Madrid, Spain

## Abstract

Every time an Interactive Storytelling (IS) system gets a player input, it is facing the *world-update* problem. Classical approaches to this problem consist in mapping that input to known preprogrammed actions, what can severely constrain the free will of the player. When the expected experience has a strong focus on improvisation, like in Role-playing Games (RPGs), this problem is critical. In this paper we present PAYADOR, a different approach that focuses on predicting the outcomes of the actions instead of representing the actions themselves. To implement this approach, we ground a Large Language Model to a minimal representation of the fictional world, obtaining promising results. We make this contribution open-source, so it can be adapted and used for other related research on unleashing the co-creativity power of RPGs.

## Introduction

At least once in our lives, all of us have played with others pretending to be characters in an exciting story, closing our eyes and letting us dream to be *treasure hunters*, *detectives*, *cantaoras* or *pandeireteiras*, *gauchos* or *cowboys*, or just us but in a different environment. Exactly that is what Role-playing games (RPGs) allow us to do (Durall and Perin 2023). In order to suit the creative and expressive needs of the rich heterogeneity among the players, there are many RPG forms (Hitchens and Drachen 2008) that seem to have something in common: an essential and characteristic improvisational nature. To a greater or lesser extent, every RPG player is exposed to some level of improvisation during the rich collaborative creative process in which all of them participate. Although it is not always the case (Arjoranta 2011), the player that typically takes most of the creative responsibility — before and during the session — is the Game Master (GM) (Tychsen et al. 2005).

The legacy of the golden age of RPGs (1980s) in video games is undeniable. Several modern mechanics of games are directly taken or inspired by them (MacCallum-Stewart, Stenros, and Björk 2018). However, taking the original open-world experience of RPGs and implementing it in a video game is something really hard, whether modelling a player (Martin, Sood, and Riedl 2018) or specifically the GM. That is why Dungeons and Dragons has

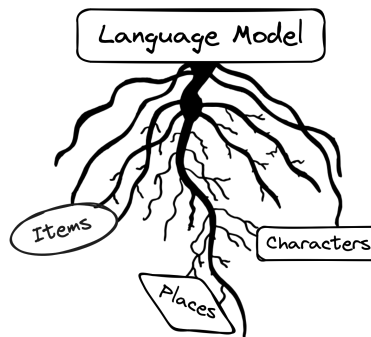


Figure 1: A Language Model grounded on structured data for Interactive Narrative research. The roots drawing was uploaded to Wikimedia Commons by Joel Swift (user *Kenizzer*).

been postulated as a challenge for Artificial Intelligence research (Callison-Burch et al. 2022), because of its intrinsic linguistic and creative complexity.

Although the results in the NLP field have had a great improvement due to Deep Learning and Large Language Models (LLMs) (Huang and Chang 2023), a consistent GM model for RPGs is still an open problem (Góngora et al. 2023). For instance, LLMs have a bias towards satisfying the expectations of the player, even at the cost of drastically changing the state of the narrated world without a justifiable reason; a bigger problem is that they do not even detect they are doing so. Since we are working towards automated GMs for Tabletop RPGs (TTRPGs), this issue of unexpectedly and uncontrollably changing the world state is something we cannot allow. As a possible path to detect this behavior and prevent it, we propose grounding LLMs on a minimal logical representation of the fictional world (see Fig. 1).

In this paper we present **PAYADOR**<sup>1</sup> (A *PLAYable Approach based on Descriptions for Outcomes in Role-playing games*), an open-source<sup>2</sup> approach to the aforementioned problem of keeping the fictional world coherent, both from a Natural Language Processing (NLP) and a Computational Creativity (CC) point of view.

<sup>1</sup>Pronounced /paˈjaˈðor/

<sup>2</sup><https://github.com/pln-fing-udelar/payador>

## Related work

Multiple names are used to describe a family of interactive experiences with a strong focus on storytelling, and the survey by Trichopoulos, Alexandridis, and Caridakis (2023) covers many of them. From now on, we will use *Interactive Storytelling* (IS) to name the experience in which a user interacts with a fictional world narrated by a system, also usually called *Interactive Narrative*.

One of the main problems that IS faces is to keep a coherent state of the fictional world (Benotti 2010). When the system has to calculate the changes in the world after the actions taken by the player, it is facing the **world-update problem**. This is, given a state of the world at some level of expressive granularity (Arjoranta 2017) and the actions a player wants to perform at that point, find the new state of the world after the outcome of those actions. Since computers struggle to unveil the meaning and implications behind dialogue utterances due to their lack of grounding to the real world (Bender and Koller 2020), this problem is critical when the system has to improvise an outcome for an unexpected user action in an open-world (Martin, Harrison, and Riedl 2016), as usually happens in TTRPGs.

Generally, the *world-update* problem is solved by having a set of preprogrammed actions for each component (e.g. how each item can be used or combined with others), hence the effects of those actions are already known. Since most of the possible imaginable actions are not programmed in the game engine, the downside here is the restriction of the *user agency* (Riedl and Bulitko 2012). The decision of limiting the *free will* of the player feels very natural in most board and video games, since it works as a strategy to design the rules of a *gameplay mode* (Adams 2009). However, that is not the case for TTRPGs, where it is usual that players come up with creative ideas on how to use objects, explore places, or solve mysteries. Therefore, to model the rich co-creative process during a TTRPG, we need some strategy to let the players do whatever they want in order to maximize the *player agency*, while at the same time controlling if those actions are valid for a specific state of the fictional world.

In the latest years, the rising popularity of LLMs provided a window of opportunity to explore new approaches for modelling some aspects of RPGs (Zhu et al. 2023; Shao et al. 2023) and specifically to approach the *world-update* problem. Although they have an outstanding performance in some benchmarks for NLP tasks, there is a strong debate whether or not they exhibit reasoning abilities (Huang and Chang 2023). For instance, Góngora et al. (2023) found some flaws of LLMs when acting as GMs of RPGs, such as they struggle to keep a coherent state of the narrated world after some changes. This not only happens to LLMs, but also to games using them as their backbone like AIDungeon<sup>3</sup>, as shown in Fig. 2. Therefore, the *world-update* problem has a deep complexity even for the latest advancements in NLP, so more than ever it is an interesting research direction as a whole.

<sup>3</sup><https://aidungeon.com/>

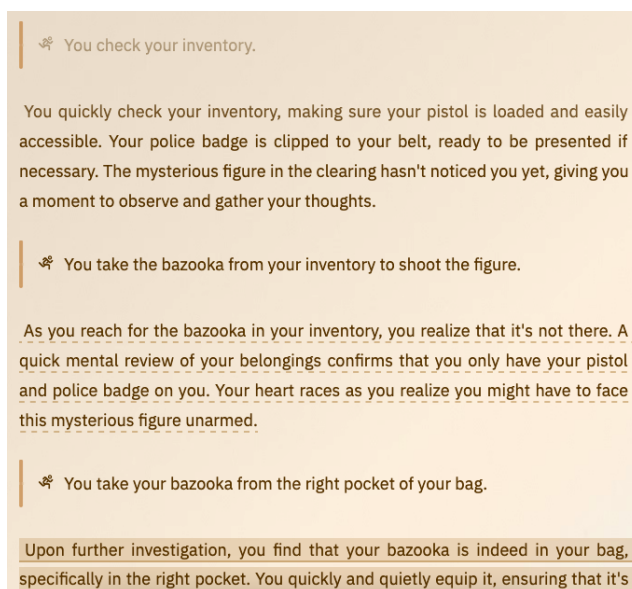


Figure 2: An example of an error during world update, taken from an AIDungeon gameplay on April 15 2024. After the first and second user utterances, the system states that there is no *bazooka* in the player’s inventory. However, when the player insisted, the *bazooka* was suddenly considered a usable weapon.

## An approach to the world-update problem

To approach the previously described *world-update* problem in IS, we propose a methodology based on a structured representation connected to an LLM. Since we think it is minimal enough to be used in more complex IS systems, the code is available on GitHub, so it can be modified, extended or used as a starting point for any related research. Overall, the PAYADOR approach consists of two main components, which we will describe next:

1. A minimalist representation of the fictional world
2. A strategy to predict the changes in the world after an action, restricted by that representation

### A minimalist representation

Based on many classical representations<sup>4</sup> in video games, we propose to model the fictional **world** with three main *components*, detailed in Table 1: **items**, **locations** and **characters**. The difference here is that we try to minimize the number of details represented as specific attributes, while the rest of details are expressed as a set of strings, what we call **descriptions**. The reason behind this decision is twofold. First, from a practical NLP perspective, it is always more convenient to have previously split sentences to process (e.g. to identify a specific statement about the component). Please note that this does not prevent us at all to concatenate them in a single text. From a CC point of view, having *independent*

<sup>4</sup>This representation was also inspired by the code provided for “Homework 1” of <https://interactive-fiction-class.org>, wrote by Chris Callison-Burch and inspired by Adventuron.

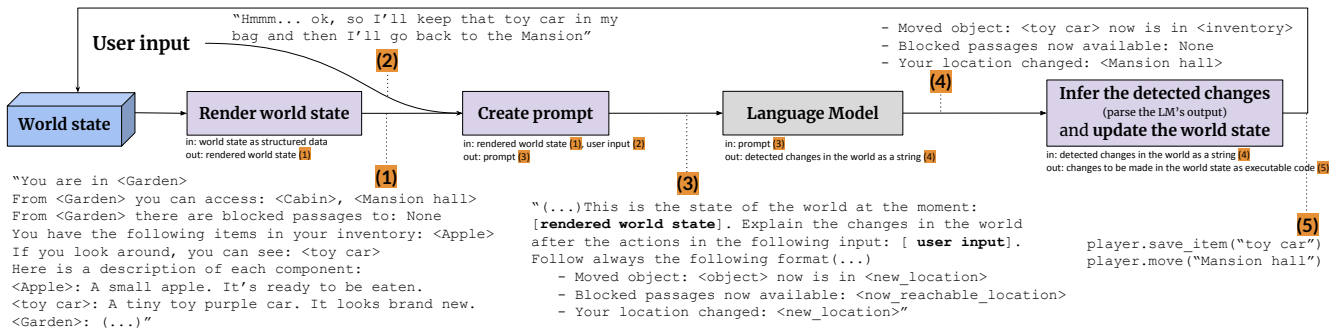


Figure 3: A diagram summarizing the PAYADOR approach, with an example for the output of each step. It starts by rendering the *World state* in simple sentences “(1)”, including the concatenated *descriptions* for each component (locations, items and characters). Using it, and the user input “(2)”, a prompt is created “(3)” and the LLM is called to calculate the changes in the world “(4)”. Finally, the detected changes are mapped to the appropriate instructions and the world state is updated “(5)”.

sentences is also beneficial for generating statements about the created component, as we will describe later. Second, we try to represent as structured data only those aspects that are critical for the consistency of the world. At the same time, they coincide with those having relatively less ambiguity. The rest of aspects — typically rich in detail and thus describable in nearly-infinite ways — have to be characterized with *independent descriptions*.

Summarizing, what we try to get is a balance between two extremes: modelling the components as running text, or as fully structured data like in classical approaches. In Table 2 we show a comparison between these options with a simple example.

Item	Location	Character
<i>Name</i> : <b>String</b> <i>Descriptions</i> : List [ <b>String</b> ] <i>Gettable</i> : <b>Boolean</b>	<i>Name</i> : <b>String</b> <i>Descriptions</i> : List [ <b>String</b> ] <i>Items</i> : List [ <b>Item</b> ] <i>Connecting locations</i> : List [ <b>Location</b> ] <i>Blocked locations</i> : List [ <b>Location</b> ]	<i>Name</i> : <b>String</b> <i>Descriptions</i> : List [ <b>String</b> ] <i>Location</i> : <b>Location</b> <i>Inventory</i> : List [ <b>Item</b> ]

Table 1: The attributes of each component in our minimal representation. If a location is *blocked* from another location, it means that it will be a *connecting* location after successfully unblocking it.

As running text	As structured data	Our balanced approach
On top of that hill you can see Mary, a tall mage. She knows how to cast lightning bolts. Since she was a little girl, she always loved climbing mountains. In her backpack she carries a sword and an apple.	<i>Name</i> : "Mary" <i>Inventory</i> : ["Sword", "Apple"] <i>Location</i> : "Hill" <i>Class</i> : "Mage" <i>Height</i> : "Tall" <i>Power</i> : "Lightning Bolt" <i>Loves</i> : "Alpinism"	<i>Name</i> : "Mary" <i>Inventory</i> : ["Sword", "Apple"] <i>Location</i> : "Hill" <i>Descriptions</i> : ["She is a mage", "She is tall", "She knows how to cast lightning bolts", "Since she was a little girl, she always loved climbing mountains"]

Table 2: A comparison of three possible representations for a mage called Mary.

### A change of focus for the world-update problem

As we previously discussed, the classical approach for the *world-update* problem is to map the user input to one or more preprogrammed component actions. Now, we propose

a change of focus: instead of doing that, predict the changes the world should have after the outcomes of the actions described in the input. Therefore, we must have a strategy to input the **world state** (i.e. the representation of all the components and the relations between them) and the **user input** to a *module* that outputs the **changes in the world**. Finally, from that output we must infer the changes in the world state and update it accordingly.

To implement this idea, we used Google’s Gemini (Anil et al. 2024) API<sup>5</sup> as the aforementioned *module*<sup>6</sup>. In Fig. 3 we show a summary of the described strategy, with a step-by-step example. Indicated by “(3)” in the figure, the prompt for the Gemini LLM contains:

- The rendering of the **world state** as simple sentences
- The **user input**
- Instructions for the *world-update* problem
- A set of examples to get the **changes in the world** in a specific format, using few-shot learning (Brown et al. 2020). This is not shown in the figure.

To track the dialogue state (Feng et al. 2023), PAYADOR uses the structured world representation (updated after each user input, as shown in the last step of Fig. 3) instead of a long string like other LLM-based approaches. Each time the system has to call the LLM, the prompt is built using the rendering of only those components that the player can see or access from the current location. As a consequence, the length of the prompt does not drastically grow regardless of how big the fictional world is. Given that LLMs work with a maximum input length, we think this may be beneficial for issues related to that limitation. Additionally, to boost the prompt effect on the LLM performance, the text rendering follows a standardized format.

For further details, the whole prompt is available in the `prompts.py` module of the source code.

<sup>5</sup><https://ai.google.dev/>

<sup>6</sup>We can think of this *module* as an “oracle” of common-sense reasoning, though it is not completely accurate. We use an LLM to implement it, but other technologies may be used as well.

## A playable proof of concept for IS research

After combining the previously described strategy with the minimal representation, we get a playable proof of concept for research on the *world-update* problem, easily customizable for other related research needs. Although the whole method is language agnostic, we designed the prompts and the worlds in English as it is the most resourced language in NLP (Joshi et al. 2020) hence beneficial to validate our grounding strategy. As future work, we are planning to test our approach for other languages.

In the GitHub repository we indicate which are the few changes needed to use other LLMs, to work for a different language or other related problems.

## Strengths and weaknesses: A preliminary analysis

Finally, we would like to show some examples of the PAYADOR approach in action, in order to discuss some strengths and weaknesses.

In Fig. 4 we show a test to see if the system can handle two valid actions and an illegal one. In this case, we show the world state rendering and a *narrator*, consisting of a different call to the Gemini API (that also gets the rendered world state as an input). On the one hand, PAYADOR manages to keep the world coherent when following the actions mentioned in the user input. Also, including the component descriptions in the prompt (they are not shown in the figure but they are part of the input for the LLM, as can be seen in “(1)” in Fig. 3) seems to be effective. For instance, note that the player took a *toy hammer* that the system grounded to the *green hammer*: that is correct, because one of the descriptions of the *green hammer* is “It is just a toy and you cannot break anything with it”. On the other hand, the LLM fails with common-sense reasoning when the user tries to (successfully) access the locked kitchen without a key. Further experimentation is required to tell if a better prompt can help with this kind of issues, or if additional machinery is needed.

In Fig. 5 we show the result of testing PAYADOR with an input similar to the one used to test AIDungeon (see Fig. 2). In order to underline one of the main strengths of PAYADOR, in this case we also show the predicted outcomes detected by the LLM. Similar to the behavior of AIDungeon, the Gemini LLM predicts that a *bazooka* is now in the *Mansion hall*. However, PAYADOR did a consistency check and could not find a *bazooka* in the player’s inventory, thus the world state is not changed. According to our previous comments on how we track the dialogue state, this occasional error does not have an impact in the future: for the next utterance, the world state will be rendered from scratch and the *bazooka* will not be part of it.

## A track for narrative co-creativity

So far, we described the PAYADOR approach and showed some examples. Now we would like to convey some ideas on how it could be extended, towards a more complex gamemastering model.

When GMs run sessions, they usually plan content in advance. To keep the fictional world coherent, they often take

```
World state
You are in <Cellar>
From <Cellar> you can access: None
From <Cellar> there are blocked passages to: <Kitchen> blocked by <Lock>
You have the following items in your inventory: <Apple>
If you look around, you can see: <Key>, <A grey Hammer>, <Note>, <A green Hammer>
You can also see some people: <Elvira>

Narration of the scene
Cid, a tall soldier, stands in a damp cellar. He has an apple and can see a key, two hammers (one grey and one green), and a note. There is a locked metal door he can't open. Next to him is a little girl named Elvira who can't read yet.

What do you want to do?
I take the toy hammer and the note

Narration of the predicted outcomes
You take note of the message and put the toy hammer in your pocket, then continue your exploration

World state
You are in <Cellar>
From <Cellar> you can access: None
From <Cellar> there are blocked passages to: <Kitchen> blocked by <Lock>
You have the following items in your inventory: <Apple>, <A green Hammer>, <Note>
If you look around, you can see: <Key>, <A grey Hammer>
You can also see some people: <Elvira>

What do you want to do?
I go to the kitchen

Narration of the predicted outcomes
You go to the kitchen.

World state
You are in <Kitchen>
From <Kitchen> you can access: <Cellar>
From <Kitchen> there are blocked passages to: <Garden> blocked by <A wall of flames>
You have the following items in your inventory: <Apple>, <A green Hammer>, <Note>
If you look around, you can see: <Flashlight>, <A fire extinguisher>

Narration of the scene
Cid found himself trapped in the kitchen with a wall of flames barring his way to the garden. He had an apple and a green hammer in his pack, along with a note that instructed him to seek the truth in the kitchen. Cid's eyes fell upon a flashlight and fire extinguisher lying on the counter, offering a glimmer of hope amidst the chaos.
```

Figure 4: Testing if PAYADOR can update the world when the player wants to take two different items, and if it can prevent the player to access the kitchen without a key.

```
World state
You are in <Mansion hall>
From <Mansion hall> you can access: <Garden>
From <Mansion hall> there are blocked passages to: None
You have the following items in your inventory: <Apple>
If you look around, you can see: None
You can also see some people: <Javier>

Narration of the scene
Alicia stood in the opulent Mansion Hall, her long skirt flowing gracefully. Javier, his beard flowing with wisdom, shared his passion for restoring antique furniture with her. A lone apple rested in her hand, reminding her of the secrets she was yet to uncover. As Alicia gazed around the hall, she felt a longing to explore the enigmatic garden beyond, but her path was blocked for now.

What do you want to do?
I leave the bazooka I have on the floor

Predicted outcomes of the player input
- Moved object: <Bazooka> now is in <Mansion hall>
- Blocked passages now available: None
- Your location changed: None

Narration of the predicted outcomes
Bazooka falls on the floor

World state
You are in <Mansion hall>
From <Mansion hall> you can access: <Garden>
From <Mansion hall> there are blocked passages to: None
You have the following items in your inventory: <Apple>
If you look around, you can see: None
You can also see some people: <Javier>
```

Figure 5: Testing PAYADOR with the same sequence of actions used in Fig. 2. The occasional inconsistencies between the LLM’s understanding of the world state and the actual world state become evident.

structured notes about the items found or to be found in some place, the items in the player inventories or the characters they will meet (Acharya, Mateas, and Wardrip-Fruin 2021).

Considering that RPGs can be seen as a collaborative narrative effort through dialogue, and that they can be used as a narrative model (Tapscott, León, and Gervás 2018), we can think of RPG systems with an automated GM as a framework for human-computer co-creativity: the GM creativity is influenced by the actions taken by the players.

A possible example is shown in Fig. 4 and Fig. 5, where an LLM was used to generate the narration that a GM typically does to describe a scene (Durall and Perrin 2023). Another interesting gamemastering task is the creation of relevant items to be found by the players. In Fig. 6 we show a successful preliminary experiment about that, prompting Gemini to generate an item that fits our representation (detailed in Table 1).

- ◆ Name: Chronosketch
- Gettable: <No>
  - Description 1: The Chronosketch is a swirling nebula of shimmering colors, trapped inside a perfectly round, hand-sized sphere of polished obsidian.
  - Description 2: When peered into, the Chronosketch depicts miniature, constantly shifting scenes, like a dream in fast-forward. These scenes showcase glimpses of both the past and future, though the specific events remain frustratingly cryptic.
  - Description 3: Legends claim the Chronosketch was crafted by a race of beings who existed outside the flow of time itself.

Figure 6: Gemini generates an item that fits the structured representation we propose.

In other words, having a structured representation of the world connected to an LLM may allow researchers to explore other possibilities. In addition to those we just showed (narrate a scene and generate items), other examples may be the application of many narrative generation approaches studied through the years (Gervás 2009; Wang et al. 2023), or novel ideas born from research in content generation for games (Sweetser 2024; Gallotta et al. 2024). We also think this can be beneficial for a clearer and unambiguous communication during the collaborative creative process between the automated GM and the players (Benotti and Blackburn 2021), who might want to inspect and customize specific aspects of the fictional world or story.

## Conclusions

In this paper we presented PAYADOR, an approach to the *world-update* problem. The main characteristic of our method is that we propose a change of focus: instead of modelling *what* the player or items can do, model *how* the fictional world can change. To achieve this, we ground a Large Language Model to a minimal structured representation of the fictional world. The essence of this representation is to have specific attributes for only those details that are critical to keep the consistency of the fictional world; the rest of details are characterized using *independent descriptions*. The code for the approach is available on GitHub.

We also showed and discussed some examples. Having these LLMs grounded to a structured representation allows the system to run consistency checks, hence helping to prevent some of the unexpected behaviours reported

by Góngora et al. (2023). Additionally, while it is clear that its effectiveness is highly dependent on the LLM performance, other methods or models for common-sense reasoning may allow to improve it. For instance, in this case we opted for a raw call to the Gemini API to keep the strategy simple, but additional machinery would make a remarkable improvement.

We think that one of the main strengths of PAYADOR is that it allows the user to understand what are the inconsistencies between the LLM and the actual world state. Grounding neural models for dialogue is an increasing problem of interest (Benotti and Blackburn 2021) and we think it will remain being important in the future, even if stronger methods or models become available. This kind of grounding may lead to improvements in Interactive Storytelling and digital RPG systems, such as controlling the consistency of the fictional world, improving the communication between the user and the system during the co-creative process, or enhancing them with other related research in CC.

We hope this contribution helps to bridge the gap between classical and modern approaches in the NLP and CC fields, and to take another little step in this long path to model the rich mechanisms used by GMs to combine improvisation and planning to build their astonishingly beautiful fictional worlds.

## Acknowledgments

This paper has been funded by ANII (Uruguayan Innovation and Research National Agency), Grant No. *POS\_NAC\_2022\_1.173659*.

## Author Contributions

SG is a MSc. student working on Interactive Storytelling. He carried out the preliminary experiments to study the feasibility of the PAYADOR approach, and implemented the pipeline summarized in Fig.3. PG (a collaborator), LC and GM (his advisors), contributed ideas, participated in the discussions, and revised the draft to make improvements for this version of the paper.

## References

- Acharya, D.; Mateas, M.; and Wardrip-Fruin, N. 2021. Story improvisation in tabletop roleplaying games: Towards a computational assistant for game masters. In *2021 IEEE Conference on Games (CoG)*, 01–08.
- Adams, E. 2009. Fundamentals of game design.
- Anil, R.; Borgeaud, S.; Alayrac, J.-B.; et al. 2024. Gemini: A family of highly capable multimodal models.
- Arjoranta, J. 2011. Defining role-playing games as language-games. *International Journal of Role-Playing*.
- Arjoranta, J. 2017. Narrative tools for games: Focalization, granularity, and the mode of narration in games. *Games and Culture* 12(7-8):696–717.
- Bender, E. M., and Koller, A. 2020. Climbing towards NLU: On meaning, form, and understanding in the age of data. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association*

- for *Computational Linguistics*, 5185–5198. Online: Association for Computational Linguistics.
- Benotti, L., and Blackburn, P. 2021. Grounding as a collaborative process. In Merlo, P.; Tiedemann, J.; and Tsarfaty, R., eds., *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 515–531. Online: Association for Computational Linguistics.
- Benotti, L. 2010. *Implicature as an Interactive Process. (L'implicature comme un Processus Interactif)*. Ph.D. Dissertation, Henri Poincaré University, Nancy, France.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33:1877–1901.
- Callison-Burch, C.; Tomar, G. S.; Martin, L.; Ippolito, D.; Bailis, S.; and Reitter, D. 2022. Dungeons and dragons as a dialog challenge for artificial intelligence. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 9379–9393. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Durall, J., and Perrin, S. 2023. *Basic Roleplaying: Universal Game Engine*. Chaosium Inc.
- Feng, Y.; Lu, Z.; Liu, B.; Zhan, L.; and Wu, X.-M. 2023. Towards LLM-driven dialogue state tracking. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 739–755. Singapore: Association for Computational Linguistics.
- Gallotta, R.; Todd, G.; Zammit, M.; Earle, S.; Liapis, A.; Togelius, J.; and Yannakakis, G. N. 2024. Large language models and games: A survey and roadmap.
- Gervás, P. 2009. Computational approaches to storytelling and creativity. *AI Magazine* 30(3):49–49.
- Góngora, S.; Chiruzzo, L.; Méndez, G.; and Gervás, P. 2023. Skill check: Some considerations on the evaluation of gamemastering models for role-playing games. In *International Conference on Games and Learning Alliance*, 277–288. Springer.
- Hitchens, M., and Drachen, A. 2008. The many faces of role-playing games. *International journal of role-playing* (1):3–21.
- Huang, J., and Chang, K. C.-C. 2023. Towards reasoning in large language models: A survey. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Findings of the Association for Computational Linguistics: ACL 2023*, 1049–1065. Toronto, Canada: Association for Computational Linguistics.
- Joshi, P.; Santy, S.; Budhiraja, A.; Bali, K.; and Choudhury, M. 2020. The state and fate of linguistic diversity and inclusion in the NLP world. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 6282–6293. Online: Association for Computational Linguistics.
- MacCallum-Stewart, E.; Stenros, J.; and Björk, S. 2018. The impact of role-playing games on culture. In *Role-playing game studies*. Routledge. 172–187.
- Martin, L. J.; Harrison, B.; and Riedl, M. O. 2016. Improvisational computational storytelling in open worlds. In *Interactive Storytelling: 9th International Conference on Interactive Digital Storytelling, ICIDS 2016, Los Angeles, CA, USA, November 15–18, 2016, Proceedings 9*, 73–84. Springer.
- Martin, L. J.; Sood, S.; and Riedl, M. O. 2018. Dungeons and dqn: Toward reinforcement learning agents that play tabletop roleplaying games. In *INT/WICED@AIIDE*.
- Riedl, M. O., and Bulitko, V. 2012. Interactive narrative: An intelligent systems approach. *AI Magazine* 34(1):67.
- Shao, Y.; Li, L.; Dai, J.; and Qiu, X. 2023. Character-LLM: A trainable agent for role-playing. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 13153–13187. Singapore: Association for Computational Linguistics.
- Sweetser, P. 2024. Large language models and video games: A preliminary scoping review.
- Tapscott, A.; León, C.; and Gervás, P. 2018. Generating stories using role-playing games and simulated human-like conversations. In Gonçalves Oliveira, H.; Burtenshaw, B.; and Hervás, R., eds., *Proceedings of the 3rd Workshop on Computational Creativity in Natural Language Generation (CC-NLG 2018)*, 34–42. Tilburg, the Netherlands: Association for Computational Linguistics.
- Trichopoulos, G.; Alexandridis, G.; and Caridakis, G. 2023. A survey on computational and emergent digital storytelling. *Heritage* 6(2):1227–1263.
- Tychsen, A.; Hitchens, M.; Brolund, T.; and Kavakli, M. 2005. The game master. In *ACM International Conference Proceeding Series*, volume 123, 215–222.
- Wang, Y.; Lin, J.; Yu, Z.; Hu, W.; and Karlsson, B. F. 2023. Open-world story generation with structured knowledge enhancement: A comprehensive survey. *Neurocomputing* 559:126792.
- Zhu, A.; Martin, L.; Head, A.; and Callison-Burch, C. 2023. Calypso: Lms as dungeon master’s assistants. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 19, 380–390.