arXiv:2309.13702v2 [cs.CL] 30 Sep 2023

Skill Check: Some Considerations on the Evaluation of Gamemastering Models for Role-playing Games

Santiago Góngora[1[0000−0002−6256−7577]], Luis Chiruzzo[1[0000−0002−1697−4614]], Gonzalo Méndez[2[0000−0001−7659−1482]], and Pablo Gervás[2[0000−0003−4906−9837]]

[1] Universidad de la República, Uruguay {sgongora,luischir}@fing.edu.uy
[2] Universidad Complutense de Madrid, Spain {gmendez,pgervas}@ucm.es

Abstract. In role-playing games a Game Master (GM) is the player in charge of the game, who must design the challenges the players face and narrate the outcomes of their actions. In this work we discuss some challenges to model GMs from an Interactive Storytelling and Natural Language Processing perspective. Following those challenges we propose three test categories to evaluate such dialogue systems, and we use them to test ChatGPT, Bard and OpenAssistant as out-of-the-box GMs.

Keywords: Role-playing Games · Natural Language Processing · Interactive Storytelling · Computational Creativity · Dialogue Systems

1 Introduction

Probably no one wants to hear somebody say "*Watch out! Behind that door there's a giant monster!*"; except if they are playing a role-playing game (RPG), using their imagination to visit endless worlds and having lots of fun.

Tabletop role-playing games (TTRPGs) consist of two or more players that collaborate in order to create a story, while acting as characters. One of these players is the Game Master (GM), who is the one in charge of creating the world where the narrated events take place, describing the non-playable characters the human players meet and the situations they face. Having a player acting as the GM is one of the characteristics that most TTRPGs share [12].

Capturing the essence of RPGs has long been one of the goals of Interactive Storytelling (IS) research [21,25]. However, through the years only limited solutions have been found, typically by having a lot of premade scenes that can be mixed to generate other narrative structures[3], but pushing the player's freedom aside[4] [25]. To automate a GM is a big challenge for Natural Language Processing (NLP) and Artificial Intelligence, due to its complexity on dialogue and creativity [8].

[3] For example "Call of Cthulhu: The Official Video Game", an adaptation of the RPG.
[4] An interesting example of this is "The Stanley Parable", a novel videogame that makes the players think about free will and the impact of their actions.

Our long-term goal is to model the diverse set of skills that a GM needs to play RPGs. This long path must lead to an explainable, grounded and controllable model, so human-in-the-loop features should be taken into consideration to meet the needs reported by [1] and [30].

In this paper we will take a first step by proposing, inspired in core aspects of RPGs, a set of unit test categories to evaluate such GM models. We also use these brand new tests to evaluate ChatGPT[5], Bard[6] and OpenAssistant[7] [16] as out-of-the-box automated GMs, both for Spanish and English.

2 Previous work

The study of the role of GMs in narrative is nothing new and also some efforts have been made to explicitly model their capabilities. [27] reflects on the core concepts of RPGs and gamemastering. [21,22] consider GMs and RPGs as a framework to build IS systems, in order to tackle the *Interactive Dilemma*, the conflict and balance between the player's will and the designer's choices. Closely related to this concept is the GM's skill to improvise some aspects of a scene due to unexpected players' actions, and [18,19] discuss concepts, approaches and architectures taking that into account. As they present insightful discussions about modelling narrative improvisation, they are some of the most clarifying works for us to date.

Most of the latest works pursue the modelling of GMs for *Dungeons & Dragons* (D&D), called *Dungeon Masters* (DMs), since it is the most popular RPG and finding data is easier than for other games. For example [8] describes the complexity of modelling the (D&D) game, performing experiments with neural models and using *control features* to guide their outputs. They also describe a *gameplay dataset* in English used for training. [28] tries to create a DM model with the ability to predict player's actions, modelling the *Dungeon Master-Player* interactions using Theory-of-Mind and Reinforcement Learning. The recent published datasets are also centered on modelling the D&D game. [24] presents one of the most complete datasets to study D&D interactions, consisting of transcriptions of the popular *Critical Role* web show. [17,29] also present datasets of D&D players' interactions from online text-based playing sessions.

It is important to note that all of these recent works are about D&D, while our main objective is to work on the general aspects of a GM, regardless of the specific game or theme. Additionally, all of them rely on English resources.

3 A list of gamemastering challenges

Most of the works mentioned in the previous section discuss difficulties faced while modelling some aspects of RPGs. However, as a way of introducing some

[5] https://openai.com/blog/chatgpt
[6] https://bard.google.com
[7] https://huggingface.co/chat/

details that guide our long-term goal and justify the test categories we propose, we would like to convey our thoughts on some challenges that a GM must face while running an RPG session. This list is not exhaustive and there may be other challenges that are not described here.

I. World and story design. As storytellers, GMs must create and manage a rich and coherent world, populated with diverse forms of life (e.g. plants and animals) and characters. In this fictional world is where the players' characters will live and act. They also need to create some interesting places (e.g. an old library) and challenges for the players, which can be logic puzzles, tactic battles, complex dialogues with characters, or other challenges (e.g. the library has hidden rooms). Usually these situations are intended to be solved by teaming up with other characters, collaborating and using the different skills that they may master. It is useful if a GM can also measure how interesting these challenges are for the players, and how meaningful they are for the development of their characters or other characters that live in the fictional world. That is the reason why it is important that such a model can take *creative responsibility* [7,10] while being able to explain what the plan and objectives of each utterance are.

II. Extract player's actions from input. Since TTRPGs are played through a discussion between the players, these games have an inherent *conversational nature*. Therefore classic research problems related to dialogue systems [9] are fundamental to model GMs. More specifically, in order to *understand* (i.e. semantically represent) the actions taken by the players, decide if they are possible in the fictional world and then determine the outcomes, the GM model should have the ability to semantically analyze their inputs.

III. Commonsense reasoning. Commonsense reasoning is an important research area within NLP [26], and despite the great advances made in the area it remains as one of the hardest tasks [11], even for the recent Large Language Models (LLMs) like ChatGPT [23]. The relation between this classic task and the challenges for a GM is direct: since commonsense is an inherent part of our human identity, it naturally arises when playing RPGs. It is important to note that this challenge is related but different to the previous challenge: a model can semantically represent what a human is saying, but maybe the action does not make sense in some context. For instance, sometimes players may want to do actions that are possible in the real world but not in the fictional world (e.g. a character wants to play basketball but there is no gravity in her world).

IV. Track the game state. One of the core aspects of RPGs is to let the players act as they wish, what in IS is usually called *user agency* [25]. Making the players feel this way while thoroughly tracking the state of items (i.e. objects) and characters is one of the greatest problems for IS [6,18]. To track some component of the game is to know where it is, how hurt (in case of a living being) or damaged it is (in case of an object), and other properties that it may have (e.g. intensity of the magic property of a sacred object). This game state must be constantly updated as the world changes and the story moves forward.

Finally, we would like to mention other relevant aspects for this long path. In first place, we think it is crucial that the narrative structure and the game state

may be represented using a human-readable format. Since RPG games are used in educational [13] and therapy environments [2], such GM models could be used to create serious games with a wide range of objectives. Having the possibility of visualizing and customizing the boundaries of an RPG session is extremely crucial for that kind of applications.

In second place, it is fundamental that these models generate respectful and *ethical* outputs, to make the players feel safe and included. In modern RPGs like "Alice is Missing"[8] there are mechanics to silently communicate the rest of players that something recently said was hurtful or uncomfortable. This is crucial when working with neural systems or LLMs, which are known to *hallucinate* [14] and generate offensive outputs [4].

Last but not least, we have to keep in mind that GMs are constantly adapting the game to fit the players' choices, so they have the additional requirement of facing every challenge described here on the fly. That big challenge is related to what [18] previously described as *open-world improvisational storytelling*.

4 How to evaluate such models?

The procedure to evaluate creative systems (i.e. appropriate experiments and metrics) has long been a subject of debate, and remains one of the main problems of the field [10,15]. Since TTRPGs can be modelled as a series of utterances in a complex dialogue [12], we will assume that a GM model will always have a *conversational nature*, as we mentioned in *challenge II*. This gives us a general guideline: there is always a player who is asking or trying to do something, and another player answering or reacting to it.

The first idea that comes to mind could be to ask humans to play and evaluate the models based on their reaction. Although we consider important to measure how fun it is to play with the models, the humans' judgments can be very subjective, not very specific, and also biased by the fluency of the generated text [3]. This bias can be stronger when working with LLMs, since they are trained to sound very natural to the human reader (exploiting the patterns behind the form of massive amounts of texts [5]), what can lead to distract the evaluators from their goal of judging specific characteristics of the models' output.

Hence, we would like to take an approach on evaluating basic, almost essential, skills that a GM should master. We propose three different test categories related to the previously described challenges: *commonsense reasoning*, the capacity to track *items* in the world and the ability to coherently design *maps*. These categories were designed reflecting on core characteristics of RPGs, so we think they can be used to evaluate any system trying to model a GM, independently of the theme and features of the modelled game, and the technology used to play it. We also hope these categories work as a guide for human evaluators, helping them to judge models while reducing the subjectivity, the mentioned biases and the evaluation noise as possible. We will describe each of them next.

[8] https://www.huntersentertainment.com/alice-is-missing

4.1 GM-P-GM pattern

In *challenges II* and *III* we discussed the importance of pragmatics and commonsense reasoning for a GM model. In order to evaluate the performance on this challenges we propose the GM-P-GM pattern, a formalization of the most elemental interaction between a GM and a player [12]. Specifically, we propose to evaluate the model's ability to judge the feasibility of a player's action:

- GM_1: Narrates a **situation** to solve in some **context**.
- *Player*: Describes the **actions** to overcome that **situation**.
- GM_2: Validates if those **actions** are feasible for that **context**, and next narrates the **outcomes**.

To run this test we give the model the GM_1 and *Player* contradictory utterances and ask it to generate the GM_2 utterance. If the GM model prevents the action and explains why it is an inconsistency, the test is passed. A failure case is shown in table 1.

GM	You enter the building to look for the lost suitcase. However, the door is covered in flames. To your right you see a fire extinguisher, which seems to be ready to use. You can't use anything from your backpack because you can hardly move your arms. What will you do?
Player	I take a bucket of water from my backpack and put out the fire
GM	As the player grabs the bucket of water from their backpack, they realize it's not nearly enough to put out the flames completely. However, they quickly pour the water onto the fire, creating a small opening in the flames (...)
Result	Failed

Table 1. An example of a *GM-P-GM* unit test.

4.2 Item tracking

As we described in *challenge IV*, item tracking is one of the fundamental problems for gamemastering. It is important both to track items that are located in specific places (e.g. on the ground, inside a safe) and to track items that belong to the characters. The latter is usually referred as the *inventory* of each character but, in order to avoid the use of specific RPG words, in this test we will use *bag* or *backpack* as more neutral words for a similar concept. To test if the model can track objects inside or outside the *backpack*, we follow this scheme:

1. *Player*: "What is in my backpack?"
2. *GM*: $Object_1, Object_2, Object_3, Object_4, ...$
3. *Player*: "I take $Object_K$ and leave it on *another location*"
4. *GM*: *Narrates outcome*
5. *Player*: "What is in my backpack?"
6. *GM*: $Object_1, Object_2, Object_3, ...$
7. *Player*: "What is on *another location* ?"
8. *GM*: $Object_A, Object_B, ...$

If the lists given in steps #2 and #6 are the same or contain other unrelated objects, then the test is failed. In case the object is not listed in step #6, the step #8 tries to check if the model updated its new location. A failure case is shown in table 2.

GM	(...) Aside from the mysterious key, your backpack also contains the usual adventurer's equipment: a sturdy rope, a grappling hook, a water flask, some dried rations, a small first-aid kit, a compass, a map of the surrounding area, and a trusty dagger at your side.
Player	I take the rope, the map and the key and I leave them on the floor
GM	You carefully set the rope, map, and key down on the floor, making sure they are within reach in case you need them later. (...)
Player	What's in my backpack?
GM	As you take a closer look at your backpack, you notice that it contains a few essential items that any adventurer would find useful: A sturdy rope: (...) this rope is perfect for climbing or rappelling down steep surfaces. (...) A map of the surrounding area: this map shows the nearby landmarks, towns, and topography of the region, allowing you to plan your route. (...)
Result	Failed

Table 2. An example of an *Item tracking* unit test.

4.3 Map design

When GMs create a fictional world they also have the challenge to design an interesting and coherent concrete map for the players to explore, as we described in *challenge I*. Although we also justified the importance of having interesting situations for the players, for now we do not focus on how attractive, large or complex the maps generated are.

Inspired in the scheme of the *Item tracking* category, we propose to ask the model to describe the accessible locations from the current location. Then the player must choose to move to one of those locations and go back again to the original location, in order to test the map's coherence. If during this process the available locations are coherent, the test is passed. A successful case is shown in table 3.

GM	(...) Current place: You find yourself standing at the entrance of a dark and damp cave. The entrance is surrounded by tall trees, and the sound of running water can be heard in the distance. Accessible places: Inside the cave, Back to the forest, Along the river
Player	I want to explore the cave
GM	(...) Current place: Inside the cave Accessible places: The cave entrance, Deeper into the cave
Player	I want to go outside
GM	(...) Current place: Outside the cave Accessible places: Inside the cave, Along the river, Deeper into the forest
Result	Passed

Table 3. An example of a *Map design* unit test.

5 Large Language Models as out-of-the-box GMs

In order to perform some experiments with our three proposed categories, we wrote five unit tests in Spanish and English for each of them. We then used those unit tests to evaluate the out-of-the-box GM skills of three LLMs: ChatGPT[9], Bard[10] and the LLaMA-based OpenAssistant chatbot[11]. We consider these three LLMs appropriate since they are dialogue systems (i.e. they have a *conversational nature*), a fundamental requirement for a GM model.

[9] "ChatGPT Mar 14 Version. Free Research Preview"

[10] Bard Experimental. Accessed on the 17th of July, 2023.

[11] "Model: OpenAssistant/oasst-sft-6-llama-30b". Accessed on the 14th of July, 2023.

Category	**OA** [ES]	**BARD** [ES]	**CGPT** [ES]	**OA** [EN]	**BARD** [EN]	**CGPT** [EN]
GM-P-GM	0 out of 5	1 out of 5	1 out of 5	1 out of 5	1 out of 5	0 out of 5
Item	0 out of 5	0 out of 5	2 out of 5	0 out of 5	3 out of 5	1 out of 5
Map	0 out of 5	3 out of 5	3 out of 5	0 out of 5	2 out of 5	3 out of 5
Total	0 out of 15	4 out of 15	6 out of 15	1 out of 15	6 out of 15	4 out of 15

Table 4. Number of passed tests for each of the categories described in section 4, testing OpenAssistant (OA), Google's Bard and ChatGPT (CGPT), both for English and Spanish. The last row shows the sum of the passed tests for each model-language pair.

Since these are only preliminary experiments we consider really important to make the experimental logs open, because it can help the reader to critically examine the results reported here and reflect on the real flaws and strengths of both our proposed test categories and the evaluated models. Therefore, the detailed logs of the experiments and their results are accessible on GitHub[12].

We detail the analysis in two subsections, the quantitative results and the qualitative observations.

5.1 Quantitative results

After we ran the tests using the aforementioned models, we carefully examined the outputs and determined the results for each test, shown in table 4. As can be seen, the performance on the *GM-P-GM* category is really low, regardless of the language or model. This result is aligned with those in [23], where commonsense reasoning was one of the remarkable flaws of ChatGPT. However, *Item tracking* and *Map design* tests were quite good both for ChatGPT and Bard. Although these preliminary experiments do not report a big gap in the results for ChatGPT or Bard between languages, they do unveil their strength over OpenAssistant. In most cases OpenAssistant just could not finish the test, due to generating nonsensical outputs that had nothing to do with the narrated events. That problem was even deeper for the tests in Spanish.

5.2 Qualitative observations

The first and most important observation is that ChatGPT and Bard are really good at making the user feel that is playing with a real GM. There is a world to interact with, characters to meet and items to use. Everything seems perfect if the player chooses an action from those suggested by the model, although it is far from perfect when having to improvise new scenes and keep it coherent. OpenAssistant, however, struggles to deliver a minimum interactive experience and the tests had to be repeated several times to obtain a reasonable output. Our evaluation methodology does not distinguish that kind of errors, hence this aspect cannot be inferred by just comparing the quantitative results for each

[12] https://github.com/sgongora27/skill-check-GM-tests

category between models (e.g. Bard failed the Spanish *Item tracking* tests due to wrongly list the available items, while OpenAssistant failed them because could not even give a proper output). However, we think that the quantitative results do represent the strengths and weaknesses of each model (e.g. ChatGPT is better at world coherence than commonsense reasoning) but also the "Total" scores provide an accurate comparison of the experience provided by the different models.

The second observation is about the contents generated by the models when taking the *creative responsibility*. Almost every scene *generated* by the models took place in a medieval-fantasy setting. This relation between RPGs and a medieval setting is aligned with the previous comments in section 2: most of the available data about RPGs is in fact about D&D. As LLMs reproduce the biases in their training data [4], this shows that more work on other RPGs with different themes is needed. There is also a notorious absence of diversity of plots; after playing a few hours the narrated events and the available places start to repeat. Although this is related to the previous comment about the medieval settings and the biased data, it is important to have in mind that a great diversity of plots can be created using a medieval-fantasy setting[13], so they are independent flaws and might be studied separately.

Our third observation is about these models' tendency to constantly adjust the output to the prompt. If the player says or tries to do something then the output will try to adjust the narrative to it, without letting the player to feel any mystery about the plot. This is not a good sign for the skills we described in *challenge I*.

5.3 Limitations

Although we propose the test categories to assist the evaluation of GM models, the human subjectivity is still there. In addition to the difficulties faced to decide whether or not the test was passed, this subjectivity can also be present in the prompts design as well, as in the case of the *GM-P-GM* tests which need a specific human-designed case in order to run (i.e. a situation to solve and a player's solution to it). To perform a deeper evaluation and extract stronger conclusions we would need a diverse team of human evaluators and a bigger number of tests.

It is important to highlight that the difficulties faced when evaluating a creative system, added to the nearly-infinite input space that RPGs offer, make the evaluation even harder. Furthermore, the LLMs show a tendency to irregularly move the story forward: sometimes the model's output narrates a single event happening immediately and sometimes narrates long scenes. Not having a reliable mechanism (e.g. a symbolic representation) to restrict the model makes the execution of these tests more unpredictable, forcing the human evaluator to take unexpected decisions on the fly. For example, it would be positive for the *Map*

[13] This is evidenced by the massive amount of adventures published for RPGs with this theme, such as *Dungeons & Dragons* or *Pathfinder*.

design tests to have some kind of constraints and visualization components, to perform an in-depth analysis of the different reachable places in a given scene but without moving the story forward. Also, these dialogue models compute the utterances each time a new input is sent, what makes the replication experiments harder. Additionally we share the same limitations found by [23] regarding the needed time to run a small set of tests.

6 Conclusions and future work

In this paper we discussed some challenges to face in order to model the skills that GMs need to play RPGs, like creating and managing a fictional world, tracking the game state and understanding the players' actions.

Following those challenges we proposed three test categories to evaluate any kind of GM model. Although these tests are domain specific, we think they can inspire other evaluation methodologies for dialogue systems.

We also used those test categories to perform preliminary experiments with ChatGPT, Bard and OpenAssistant. We found that ChatGPT and Bard can provide a satisfying gaming experience, but also they struggle when dealing with commonsense reasoning. OpenAssistant was unable to maintain the GM role during most of the tests. All 90 unit tests are available on GitHub.

The difficulties faced to control the models' outputs while running the tests make us think that in the future more *neuro-symbolic* approaches should be explored. We think that would help to keep the test phase more controllable, and also allow the players to examine the narrative details, avoid some scenes that they do not want to play and add another elements that they do.

In the future we would like to improve these test categories and design more to test other gamemastering skills (e.g. model the emotional variation of a character during an interaction with other character [20]).

7 Acknowledgements

This paper has been partially funded by ANII (Uruguayan Innovation and Research National Agency), Grant No. $POS_NAC_2022_1_173659$ and by the project CANTOR: Automated Composition of Personal Narratives as an aid for Occupational Therapy based on Reminescence, Grant No. $PID2019{-}108927RB{-}I00$ (Spanish Ministry of Science and Innovation).

References

1. Acharya, D., Mateas, M., Wardrip-Fruin, N.: Story improvisation in tabletop roleplaying games: Towards a computational assistant for game masters. In: 2021 IEEE Conference on Games (CoG). pp. 01–08 (2021). https://doi.org/10.1109/CoG52621.2021.9619006

2. Arenas, D.L., Viduani, A., Araujo, R.B.: Therapeutic use of role-playing game (rpg) in mental health: A scoping review. Simulation & Gaming **53**(3), 285–311 (2022)
3. Behnke, H., Fomicheva, M., Specia, L.: Bias mitigation in machine translation quality estimation. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 1475–1487. Association for Computational Linguistics, Dublin, Ireland (May 2022). https://doi.org/10.18653/v1/2022.acl-long.104
4. Bender, E.M., Gebru, T., McMillan-Major, A., Shmitchell, S.: On the dangers of stochastic parrots: Can language models be too big? In: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency. p. 610–623. FAccT '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3442188.3445922
5. Bender, E.M., Koller, A.: Climbing towards NLU: On meaning, form, and understanding in the age of data. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 5185–5198. Association for Computational Linguistics, Online (Jul 2020). https://doi.org/10.18653/v1/2020.acl-main.463
6. Benotti, L.: Implicature as an Interactive Process. (L'implicature comme un Processus Interactif). Ph.D. thesis, Henri Poincaré University, Nancy, France (2010)
7. Botelho, L.M.: A guided journey through non-interactive automatic story generation (2021)
8. Callison-Burch, C., Tomar, G.S., Martin, L., Ippolito, D., Bailis, S., Reitter, D.: Dungeons and dragons as a dialog challenge for artificial intelligence. In: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing. pp. 9379–9393. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Dec 2022)
9. Chen, H., Liu, X., Yin, D., Tang, J.: A survey on dialogue systems: Recent advances and new frontiers. SIGKDD Explor. Newsl. **19**(2), 25–35 (nov 2017). https://doi.org/10.1145/3166054.3166058
10. Colton, S., Wiggins, G.A.: Computational creativity: The final frontier? In: European Conference on Artificial Intelligence (2012)
11. Davis, E., Marcus, G.: Commonsense reasoning and commonsense knowledge in artificial intelligence. Communications of the ACM **58**(9), 92–103 (2015)
12. Durall, J., Perrin, S.: Basic Roleplaying: Universal Game Engine. Chaosium Inc. (2023)
13. Gatsakou, C., Bardis, N., Drigas, A.: Role playing vs rpgs as teaching strategies in educational procedure. Technium Soc. Sci. J. **26**, 186 (2021)
14. Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A., Fung, P.: Survey of hallucination in natural language generation. ACM Comput. Surv. **55**(12) (mar 2023). https://doi.org/10.1145/3571730
15. Jordanous, A.: A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. Cogn. Comput. **4**(3), 246–279 (2012). https://doi.org/10.1007/s12559-012-9156-1
16. Köpf, A., Kilcher, Y., von Rütte, D., Anagnostidis, S., Tam, Z.R., Stevens, K., Barhoum, A., Duc, N.M., Stanley, O., Nagyfi, R., ES, S., Suri, S., Glushkov, D., Dantuluri, A., Maguire, A., Schuhmann, C., Nguyen, H., Mattick, A.: Openassistant conversations – democratizing large language model alignment (2023)
17. Louis, A., Sutton, C.: Deep dungeons and dragons: Learning character-action interactions from role-playing game transcripts. In: Proceedings of the 2018 Confer-

ence of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers). pp. 708–713. Association for Computational Linguistics, New Orleans, Louisiana (Jun 2018). https://doi.org/10.18653/v1/N18-2111
18. Martin, L.J., Harrison, B., Riedl, M.O.: Improvisational computational storytelling in open worlds. In: Nack, F., Gordon, A.S. (eds.) Interactive Storytelling. pp. 73–84. Springer International Publishing, Cham (2016)
19. Martin, L.J., Sood, S., Riedl, M.O.: Dungeons and dqns: Toward reinforcement learning agents that play tabletop roleplaying games. In: INT/WICED@AIIDE (2018)
20. Oñate, A., Méndez, G., Gervás, P.: Emolift: Elevator conversations based on emotions. In: International Conference on Innovative Computing and Cloud Computing (2019)
21. Peinado, F., Gervás, P.: Transferring game mastering laws to interactive digital storytelling. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) Technologies for Interactive Digital Storytelling and Entertainment. pp. 48–54. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
22. Peinado, F., Gervás, P.: Automatic direction of interactive storytelling: Formalizing the game master paradigm. In: Cavazza, M., Donikian, S. (eds.) Virtual Storytelling. Using Virtual Reality Technologies for Storytelling. pp. 196–201. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)
23. Qin, C., Zhang, A., Zhang, Z., Chen, J., Yasunaga, M., Yang, D.: Is chatgpt a general-purpose natural language processing task solver? (2023). https://doi.org/10.48550/ARXIV.2302.06476
24. Rameshkumar, R., Bailey, P.: Storytelling with dialogue: A Critical Role Dungeons and Dragons Dataset. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 5121–5134. Association for Computational Linguistics, Online (Jul 2020). https://doi.org/10.18653/v1/2020.acl-main.459
25. Riedl, M.O., Bulitko, V.: Interactive narrative: An intelligent systems approach. AI Magazine **34**(1), 67 (Dec 2012). https://doi.org/10.1609/aimag.v34i1.2449
26. Sap, M., Shwartz, V., Bosselut, A., Choi, Y., Roth, D.: Commonsense reasoning for natural language processing. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts. pp. 27–33. Association for Computational Linguistics, Online (Jul 2020). https://doi.org/10.18653/v1/2020.acl-tutorials.7
27. Tychsen, A., Hitchens, M.: The many faces of role-playing games. International journal of role-playing **1**(1), 3–21 (2009)
28. Zhou, P., Zhu, A., Hu, J., Pujara, J., Ren, X., Callison-Burch, C., Choi, Y., Ammanabrolu, P.: An ai dungeon master's guide: Learning to converse and guide with intents and theory-of-mind in dungeons and dragons (2022). https://doi.org/10.48550/ARXIV.2212.10060
29. Zhu, A., Aggarwal, K., Feng, A., Martin, L., Callison-Burch, C.: FIREBALL: A dataset of dungeons and dragons actual-play with structured game state information. In: Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 4171–4193. Association for Computational Linguistics, Toronto, Canada (Jul 2023)
30. Zhu, A., Martin, L.J., Head, A., Callison-Burch, C.: Calypso: Llms as dungeon masters' assistants. In: AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE). AAAI, Salt Lake City, UT (10 2023)