Prototyping Slice of Life: Social Physics with Symbolically Grounded LLM-based Generative Dialogue

Mike Treanor
treanor@american.edu
American University
Washington, DC, USA

Ben Samuel
bsamuel@cs.uno.edu
University of New Orleans
New Orleans, LA, USA

Mark J. Nelson
mnelson@american.edu
American University
Washington, DC, USA

ABSTRACT

This paper describes a prototype for the social physics game *Slice of Life*, and how it makes use of the underlying social simulation system's sophisticated state to generate symbolically grounded prompts for a large language model (LLM) in order to generate context appropriate character dialogue. Rather than using LLMs for novelty or economic reasons, the underlying social simulation technology, we argue, necessitates this approach in order to make it feasible to have nuanced dialogue that reflects the many ways the characters could have gotten themselves into particular social situations. The primary goal of this paper is to illustrate how the generative possibilities of LLMs can be uniquely useful when applied as a controlled natural language generation (NLG) system, without giving up authorial control of the gameplay or story, or sacrificing accepted good game design practice.

CCS CONCEPTS

• **General and reference** → **Design**; • **Applied computing** → **Media arts**; • **Computing methodologies** → **Natural language generation**.

KEYWORDS

Social Simulation, Large Language Models, Playable Experiences

ACM Reference Format:
Mike Treanor, Ben Samuel, and Mark J. Nelson. 2024. Prototyping Slice of Life: Social Physics with Symbolically Grounded LLM-based Generative Dialogue . In *Proceedings of the 19th International Conference on the Foundations of Digital Games (FDG 2024), May 21–24, 2024, Worcester, MA, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3649921.3656988

1 INTRODUCTION

Large language models (LLMs) have shown themselves to be powerful tools, capable of producing vast tracts of text and images, and of solving otherwise intractable problems (not to mention introducing new, as-of-yet unresolved issues). Though the technology underlying LLMs is sophisticated, they are simple to use: one need simply supply a prompt written in natural language.

As powerful as they are, LLMs struggle with a number of issues, including the problem of *hallucination*, the term given to LLMs presenting invented information as fact. Hallucinations can be dangerous, as they are largely indistinguishable from logically sound generated content, and the onus of separating fact from fiction rests upon the reader. Hallucinations can be lessened through *prompt engineering*, the term used to describe crafting prompts in such a way to produce desired output from an LLM.

This paper presents our approach to hallucination-resistant use of an LLM for game dialogue, prototyped in a new social simulation game. *Slice of Life* relies on generative text to realize contextual dialogue driven by a procedural simulation system, but gameplay is controlled by a separate symbolically driven AI system based on Comme il Faut (CiF) [8]. We have extended CiF to also perform the prompt engineering for LLM-generated dialog. We argue these two systems compliment each other well by shoring up the others' weaknesses: generating dialogue in CiF-driven systems has been laborious, and near impossible in some situations. Handing off that responsibility to an LLM significantly decreases authorial burden. Conversely, CiF excels at maintaining and reasoning over complex game state, which changes constantly based on player actions; in other words, CiF never hallucinates.

By incorporating game state, as maintained by CiF, into prompts fed into an LLM, the *Slice of Life* prototype showcases an LLM that is capable of generating character dialogue that is grounded in state and reality. Moreover, the tone and content of the generated dialogue adapts to changes in game state as driven by player action.

2 RELATED WORK

Slice of Life is in the genre of "social physics" games. In the same way physics games give players goals within a simulation of gravity, space, collisions, etc., and players discover emergent solutions, a social physics game gives players a social state to achieve, and sets them loose interacting with the social simulation [5, 7]. This particular project has roots in Comme il Faut and its successor, Ensemble [8, 13] (see Section 4.1 for how we build on those systems). As this paper is partly about alleviating authorial burden, particularly relevant is DeKerlegand et al. [2] on authoring challenges.

Natural language generation (NLG) is one way to address the authoring problem in procedural narrative. Standard game dialogue writing methods run into the problem that it is infeasible to write dialogue ahead of time for each combinatorial variation of situations [4, 19]. One could avoid text entirely, through *Sims*-style visual communication of game state. But if we do want dialogue, some form of NLG is necessary. Examples go back a number of years and have used many approaches: templates, generative grammars, storylet systems, etc. [6, 9, 12, 15].

FDG 2024, May 21–24, 2024, Worcester, MA, USA

ACM ISBN 979-8-4007-0955-5/24/05
https://doi.org/10.1145/3649921.3656988

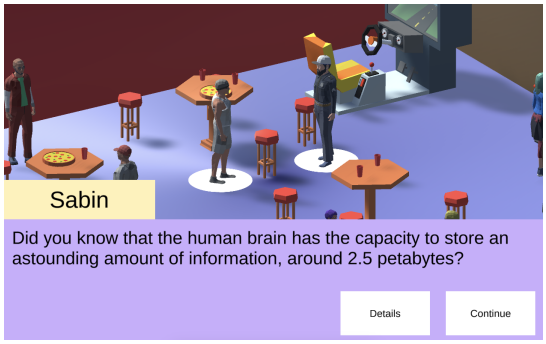

Figure 1: A screenshot showing Sabin making Small Talk with his brother while trying to impress him.

We turn to LLMs to help with this NLG task, so work on LLMs for game dialogue is relevant. Akoury et al. [1] have an LLM generate text in place of masked-out dialogue lines from *Disco Elysium* (although they aim more at using games for LLM evaluation than NLG for procedural narrative). Müller-Brockhausen et al. [10] focus more directly on LLMs in game design, and make a distinction between *chatter* and *dialogue*. They argue that "chatter yields more promise for integration" of LLMs into games, because it sidesteps issues of LLM-generated text going off the rails. Their solution is to use LLMs for non-story-critical filler text, which they call chatter. Our approach aims to avoid the problem they identify with LLM dialogue by keeping all state symbolically grounded, and using the LLM solely to do surface-text realization of that symbolic state.

Finally, we make the opposite representational choice from the otherwise related *Smallville* social simulation system [11]. They use LLM-generated text as the simulation representation itself, storing generated text in various places as the simulation state (for example, agents store generated lines of text in their "memory" to condition later generation). We keep the simulation state entirely as symbolic state, and use the LLM for surface text realization but not simulation progress. The following section motivates that design decision.

3 MOTIVATION: AI-BASED GAMEPLAY

Slice of Life explores AI-based gameplay within a social simulation game. AI-based gameplay foregrounds the operation and structure of the underlying artificial intelligence techniques used to create it. In other words, the AI system is a core part of gameplay, requiring players to engage with and consider it as an end of its own [3, 18].

Samuel et al. [14] argue that games with AI-based gameplay should strive to adhere to the same generally accepted features of "good" game design as other games – interpretability, consistency/-coherence, agency, and authorability. They point out those design values are difficult to maintain with AI based on machine learning. With the largely black box nature of learned models, the underlying system is difficult for players to interpret or understand, so understanding it cannot be made a core gameplay mechanic. Samuel et al. [14] argue that symbolic approaches are better suited to AI-based gameplay. Symbolic models have consistent and interpretable internal representations of the worlds they represent, which makes them ideal for types of AI-based gameplay that depend on players reasoning about the system's underlying state.

When it comes to authoring, the story is more complicated. Careful symbolic knowledge engineering decisions can make authoring somewhat match the how humans think (e.g. the symbolic relationship `family(Edgar,Sabin)`), but it quickly becomes overwhelming to track many of symbolic facts and to communicate them to the player. Our goal is to use large language models to solve that problem, communicating the current game state (which may involve many symbolic facts), but without giving the model any influence over the core game state or progression.

4 SLICE OF LIFE

Our playable prototype, *Slice of Life*, uses CiF for state management and social simulation, and an LLM for dialogue generation. The player plays as a disembodied pizzeria owner who can see the future social struggles of his customers. The goal is to prevent those events from occurring, to achieve customer satisfaction. Like *Prom Week*, this game is a social physics game where the goals are to achieve a symbolically represented social state, and player's actions modify the social simulation system on a deterministic basis.

To play the game, players first select a character, then select a second character to interact with. Each character is either a patron or an employee of the pizza parlor, and game states are largely defined through their relationships with other people (who they are family with, who they are friends with, etc.). Each character has individual characteristics, both permanent (e.g., employee of the restaurant) and temporary (e.g., "hangry"). Under the hood, interactions employ a CiF-based social practice system. From the player's viewpoint, they are shown a window with buttons that correspond to actions that character can take. The available actions and social practices are determined by CiF from the characters' social state and sets of authored social influence rules. An example of a social practice is *Small Talk*, shown in Figure 2.

Once a social practice or action is selected, a prompt is generated based on CiF's current representation of the social state (see Figure 3). This prompt is fed into an LLM, which is ultimately asked to provide a line of dialogue befitting the current situation. It is interesting to note that the prompts do not reference the state directly (e.g., that the two characters are brothers), but rather the social influence rules that were flagged as true because of the state (e.g., family members are kind to one another).

4.1 The Social Simulation System

Slice of Life is the newest game to make use of the CiF family of social simulation systems [8, 17]. In addition to the symbolically-grounded LLM dialogue generation, another contribution of this paper is the first game-based publication to use the Social Practice branch of the CiF family, which derives from Ensemble [13]. The Social Practice system enables turn-based interactions between characters as they navigate something like a dialogue tree, but where each branch is chosen based on a sophisticated social reasoning model.

A social practice is made up of a series of linked stages with two participants taking turns as the active character. Stages are comprised of a set of actions that contain a character performance

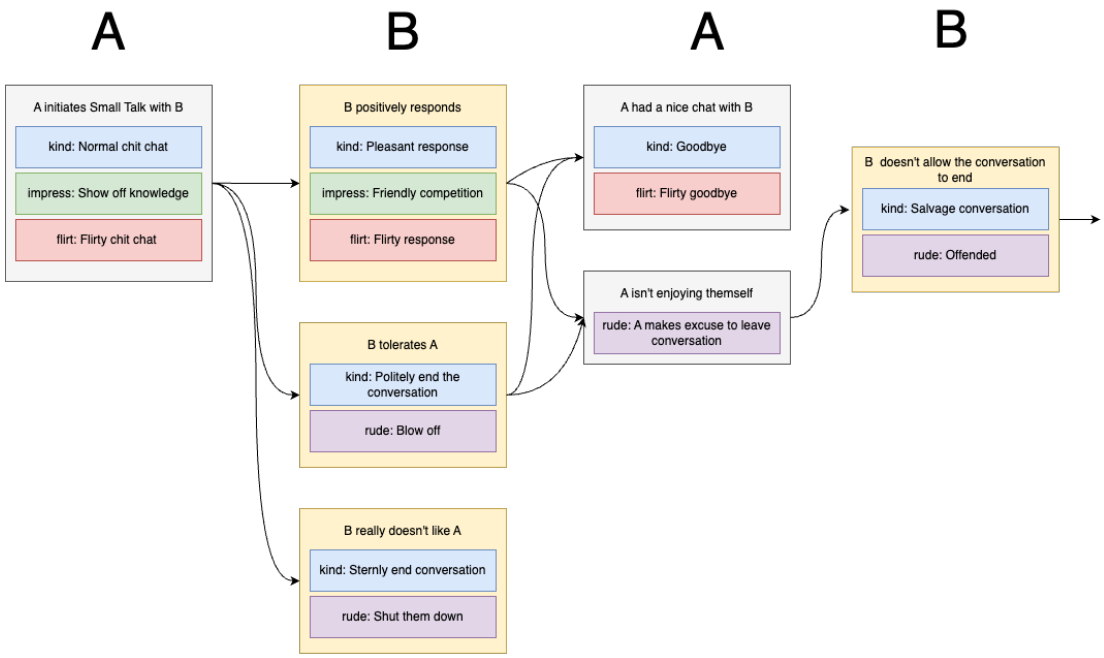

Figure 2: A partial diagram of a social practice showing all of the possible paths for characters A and B to take turns choosing how to respond to each other based on their possible intents: kind, flirt, rude, and impress.

(generally a line of dialogue, animation, or visualization) and updates to the social record called effects. Each action is authored to conceptually represent an aspect of a practice. A stage "A is having a nice chat with B" could contain actions where participants express that nice chat in different ways: 'kindly enjoying the conversation' or 'taking the conversation in a flirty direction'). The stages are linked as a directed acyclic graph. A practice ends when there are no more stages. The set of all possible paths through the actions of a social practice represent the space of all possible ways two characters could engage in that type of social interaction.

When choosing to act, agents choose the *highest scored* action that is available to them. The available actions are determined by the current stage and all of the available actions in the linked stages. Where these scores come from is one of the more interesting parts of this system. Every action is tagged with an *intent* and these intents are given a weight based upon the sum of all true *microtheory rules*. The intents in *Slice of Life* are: kind, rude, flirt, and impress.

Each intent has a set of *microtheory rules* that describe reasons an agent could have that intent. These rules are based on the social schema (the set of things that can be true between agents) [13], and have weights. The sum of all true microtheory rules determines the agent's desire to hold that intent. For example: "If you are family with someone, you are much less inclined to be rude to them", or $family(x, y) \implies rude(-3)$. A little more complicated example: "If you're an employee on a break, and a customer comes to talk to you, you will be a little more inclined to be rude to them", or

$$on\_break(x) \wedge employee(x) \wedge \neg employee(y) \implies rude(+2)$$

Our prototype currently has around 100 such rules, but thousands can be supported.

4.2 The LLM Dialogue System

Figure 2 shows the majority of the social practice *Small Talk*. From left to right, agents A and B take turns choosing which action to take from the linked stages. In this example, once A has chosen their action from the "A Initiates Small Talk with B" stage, there are seven possible actions available for B to choose from. The nature of this structure makes it so there are many different ways that can potentially have led to an action being chosen. From an authoring point of view, this makes it very difficult to write coherent dialogue that respects what happened before. For example, in Figure 2, let's say B is choosing to have a "Flirty response". At the time of authoring, we don't know what action preceded it: was it "Normal chit chat", "Show off knowledge" or "Flirty chit chat"? This problem only gets worse the further into a social practice the characters get.

Our solution is to generate contextually relevant dialog for each action by feeding the relevant context (game state and path) to an LLM. We choose relevant features of the symbolically authored social practices and dynamic social state, and retrieve their current values. These are converted to text and used to to construct an LLM prompt, which is sent to an LLM to generate the actual dialogue.

We have experimented with two broad classes of LLMs, each of which requires a different approach to constructing the prompt. *Base* models are trained to model language directly, and work as a kind of fancy autocomplete. The type of text produced can be influenced by a variety of prompting strategies. *Instruction-tuned* models have an additional layer of training, learning to respond to instructions stated in natural language. For the current prototype, we use Google Gemini's API [16], as it has a free tier and produces reasonably good output. It is an instruction-tuned model,

```
Edgar and Sabin are engaged in a conversation generally
referred to as: Small Talk.

So far, this is what they have said:
- Sabin said the following with the intent of impress:
"My friend is a world-renowned artist; her work is so
impressive, it's been featured in museums around the
globe. I'd love to introduce you to her sometime."

Edgar is going to respond with the intent of being kind.
The reasons Edgar wants to be kind are:
- Family members are kind to one another.
- Loyal people are kind.
- Friendly people are kind.

The way Edgar is going to respond can be described as
"Just a normal chit chat response to someone Edgar
likes". Remember that Edgar has the intent of being kind.
Please create a single line of dialogue that Edgar would
say in the situation just described.

Try to represent the reasons they intend to respond in
that way in the line of dialogue as much as possible.
```

Figure 3: Example prompt, including the action's intent and the aspects of social state that led to that intent's selection.

so we generate the prompts by concatenating a series of declarative statements about the game state, followed by imperative requests specifying desired properties of the generated dialogue.

To retrieve relevant game state that should go into the prompt, we determined ahead of time the salient features of the social practice and social state. For now these are not dynamic and the same features are always retrieved: the social practice's name, the history of conversation within the practice, the name of the action that we asking for dialogue from, and probably most important, the intent of the action. To provide context for the intent, we retrieve the specific microtheory rules that influenced that intent being chosen (this part is dynamic, selecting the active microtheory rules). It is especially in these last two features that we view our approach as novel, as the microtheory-based action selection system is quite complex and captures nuance of the highly dynamic social state, which poses challenges for NLG dialogue systems.

Each piece of selected state is translated into English sentences through a simple templated `toString` method, and the prompt is assembled from those sentences. In this respect, the construction of the prompt has similarities to existing templated NLG dialogue systems, with the difference that it is not surface text intended for player consumption – it will be the LLM's job to take this rough templated text dump and turn it into a nice line of dialogue. See Figure 3 for an example of a complete prompt.

An example Small Talk social practice dialogue:

Sabin (intent: impress): "My friend is a world-renowned artist; her work is so impressive, it's been featured in museums around the globe. I'd love to introduce you to her sometime."
Edgar (intent: kind): "Wow, how wonderful! That must make you very proud of her - I would love to meet her if she's ever back in town."
Sabin (intent: rude): "Sure, why not... I guess a meeting would be 'fine,' or whatever."

The prompt shown in Figure 3 was used to generate the second line of dialogue (what Edgar says). The LLM is asked to generate a line of dialogue that can be described as "Just normal chit chat...", with the intent of being kind, and for the reasons that he is family with Sabin, he's loyal, and he is friendly. The list of reasons can become arbitrarily large, and would be difficult for a human author to account for. Based on our experiments, the LLM performs very well, and rarely makes statements that are unacceptable.

5 CONCLUSION AND FUTURE WORK

This paper presents the prototype for *Slice of Life*, and how it applies the CiF-based social practice system, and symbolically grounded prompts for a large language model (currently Google Gemini), to generate context appropriate dialogue in a very dynamic social state. In addition to the symbolically grounded prompt generation, this paper is the first publication about the application of the social practice system to a game. Future work will involve lots of game design and development to make the game fully playable. On the other hand, we will also be exploring other LLMs and prompt generation approaches.

REFERENCES

[1] Nader Akoury, Qian Yang, and Mohit Iyyer. 2023. A Framework for Exploring Player Perceptions of LLM-Generated Dialogue in Commercial Video Games. In *Findings of EMNLP*.
[2] Daniel DeKerlegand, Ben Samuel, and Mike Treanor. 2021. Pedagogical Challenges in Social Physics Authoring. In *Proc. ICIDS*.
[3] Mirjam P Eladhari, Anne Sullivan, Gillian Smith, and Josh McCoy. 2011. AI-based game design: Enabling new playable experiences. Tech. Rep., UC Santa Cruz.
[4] Charlie Hargood, David E. Millard, Alex Mitchell, and Ulrike Spierling (Eds.). 2022. *The Authoring Problem*. Springer.
[5] Shi Johnson-Bey, Mark J Nelson, and Michael Mateas. 2022. Exploring the Design Space of Social Physics Engines in Games. In *Proc. ICIDS*.
[6] Max Kreminski and Noah Wardrip-Fruin. 2018. Sketching a map of the storylets design space. In *Proc. ICIDS*.
[7] Josh McCoy, Mike Treanor, Ben Samuel, Michael Mateas, and Noah Wardrip-Fruin. 2011. Prom Week: Social physics as gameplay. In *Proc. FDG*.
[8] Joshua McCoy, Mike Treanor, Ben Samuel, Aaron A Reed, Michael Mateas, and Noah Wardrip-Fruin. 2014. Social story worlds with comme il faut. *IEEE Transactions on Computational Intelligence and AI in Games* 6, 2 (2014), 97–112.
[9] Nick Montfort. 2009. Curveship: An interactive fiction system for interactive narrating. In *Proc. Computational Approaches to Linguistic Creativity*.
[10] Matthias Müller-Brockhausen, Giulio Barbero, and Mike Preuss. 2023. Chatter Generation through Language Models. In *Proc. IEEE CoG*.
[11] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proc. UIST*.
[12] James Ryan, Michael Mateas, and Noah Wardrip-Fruin. 2016. Characters who speak their minds: Dialogue generation in Talk of the Town. In *Proc. AIIDE*.
[13] Ben Samuel, Aaron A Reed, Paul Maddaloni, Michael Mateas, and Noah Wardrip-Fruin. 2015. The Ensemble engine: Next-generation social physics. In *Proc. FDG*.
[14] Ben Samuel, Mike Treanor, and Joshua McCoy. 2021. Design Considerations for Creating AI-based Gameplay. In *Proc. AIIDE EXAG Workshop*.
[15] Christina Strong, Manish Mehta, Kinshuk Mishra, Alistair Jones, and Ashwin Ram. 2007. Emotionally driven natural language generation for personality rich characters in interactive games. In *Proc. AIIDE*.
[16] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, et al. 2023. Gemini: a family of highly capable multimodal models. arXiv:2312.11805.
[17] Mike Treanor, Josh McCoy, and Anne Sullivan. 2015. Social play in non-player character dialog. In *Proc. AIIDE*.
[18] Mike Treanor, Alexander Zook, Mirjam P Eladhari, Julian Togelius, Gillian Smith, Michael Cook, Tommy Thompson, Brian Magerko, John Levine, and Adam Smith. 2015. AI-based game design patterns. In *Proc. FDG*.
[19] Robert Zubek. 2024. Narrative. In *Elements of Game Design*. MIT Press. https://mitp-content-server.mit.edu/books/content/sectbyfn/books_pres_0/12040/chapter.pdf Online supplemental material.