

Evaluating a Casual Procedural Generation Tool for Tabletop Role-Playing Game Maps

Henry Crain

Department of Computer Science
North Carolina State University
Raleigh, NC, United States
hrcrain@ncsu.edu

Dan Carpenter

Department of Computer Science
North Carolina State University
Raleigh, NC, United States
dcarpen2@ncsu.edu

Chris Martens

Department of Computer Science
North Carolina State University
Raleigh, NC, United States
martens@csc.ncsu.edu

Abstract—Tile maps are useful for a wide variety of games, particularly tabletop role-playing games. However, existing tools for creating them usually either require the map to be created entirely by hand, or procedurally generate the entire map with only a few parameters being controlled by the user. In previous work, we presented a mixed-initiative tool for generating tile maps that gives the user precise control over the generated map while still allowing them to take advantage of procedural generation. In this work, we empirically evaluate this tool by asking five users to complete three thinkaloud tasks and a post-task survey. We found that users found the tool fun to use and that it made map design easier. We discuss takeaways and how this tool and similar tools can be made better.

Index Terms—answer set programming, casual creator, games, mixed-initiative co-creativity, procedural content generation, user experience, user study

I. INTRODUCTION AND RELATED WORK

Tile maps, or 2D images assembled from smaller images of a regular size and shape [1], are commonly used in digital and tabletop games to represent virtual environments. The smaller images, called tiles, represent physical features or parts of features in the environment. An example of a tile map is shown in Fig. 1. Tile maps are used in many genres of games; the focus of this project is tabletop role-playing games (TTRPGs), analog games in which players roleplay their characters' actions, and those actions are mediated by a ruleset and a *game master*, who also designs and describes the scenario for the game.

The goal of our research is to help game masters across a range of experience levels make tile maps more easily. There exist many powerful editors for manually creating tile maps, such as Tiled, [3], but these require the user to create the entire map by hand, which can be (as one of our study participants described it) “time-consuming and tedious.” This difficulty could be alleviated by a procedural generation system., such as donjon’s Fantasy Town Generator [4]. However, donjon only gives the user control over high-level parameters such as size and environment.

In our previous work [5], we presented the design of a tool that combines the benefits of manual editors and procedural generators by following *mixed-initiative co-creativity* (MI-CC) principles [6]. MI-CC is defined as “the task of creating




Fig. 1. An example of a tile map and the tiles that comprise it [2]

artifacts via the interaction of a human initiative and a computational initiative.” Previous evaluations of mixed-initiative map generation tools suggest that users generally have positive experiences with them [7], [8] but a common criticism is that users want more control [7], [9].

Our MI-CC tool, Envoi, allows a user to place tiles and specify general rules for the map and use answer set programming (ASP) to complete the map. ASP is a form of logic programming that is an established method of generating game content [10]–[12]. While traditional logic programming, like in Prolog [13], allows a user to state facts about how the world *is*, ASP allows a user to express a *set* of possible worlds [10].

When designing Envoi, we aimed for a balance of user-friendliness and expressivity. It thus aligns with Compton and

Mateas’ concept of a *casual creator* [14], a tool that supports intrinsically motivated creativity and is accessible to users with no background knowledge or experience. In the previous paper [5], we assessed Envoi subjectively along each of Compton and Mateas’ design patterns for *casual creators*, which provide a compelling target for this work. While some of the casual creator patterns are possible to assess by observation, a few of them require empirical evaluation. Evaluation of these criteria, along with gathering baseline data from user interaction with Envoi, is thus the focus of this paper.

II. STUDY DESIGN

Our study was driven by the following research questions:

- RQ1. How easy is Envoi to use?
- RQ2. How fun is Envoi to use?
- RQ3. How does Envoi aid the process of map design?
- RQ4. How well does Envoi implement casual creator design patterns?

Our study consisted of a thinkaloud session in which participants used Envoi while a researcher listened to and logged their comments, followed by a post-study survey that was sent to participants after the session. We sought participants from gaming clubs and game design courses at North Carolina State University, and participants self-selected to participate based on their interest in the subject matter. Participants had to have at least some experience with tabletop role-playing games, so that they could evaluate the usefulness of this tool for that purpose.

A. Thinkaloud

We conducted thinkaloud sessions using the Zoom video conferencing platform. To ensure that participants had a consistent experience, we had them use the remote control feature to control Envoi on the researcher’s computer. During the study, participants were asked to complete three tasks. Task 1 was a tutorial task designed to introduce the features of Envoi. In Task 2, participants were given a description of a city, consisting of a wealthy, well-maintained inner city and a sprawling outer city, and were asked to create a map that fit the description. In Task 3, participants were directed to make any city map they wanted.

B. Survey

Immediately after the session, participants were sent a survey asking about their past experiences with TTRPGs and map-making and their experience with Envoi. The survey started with the following background questions:

- How much experience do you have as a dungeon master or game master for TTRPGs? (multiple choice, shown in Fig. 2)
- How much experience do you have designing maps for TTRPGs? (multiple choice, shown in Fig. 2)
- What tools, if any, have you used to design maps? (free response)
- What parts, if any, do you find difficult about making maps? (free response)

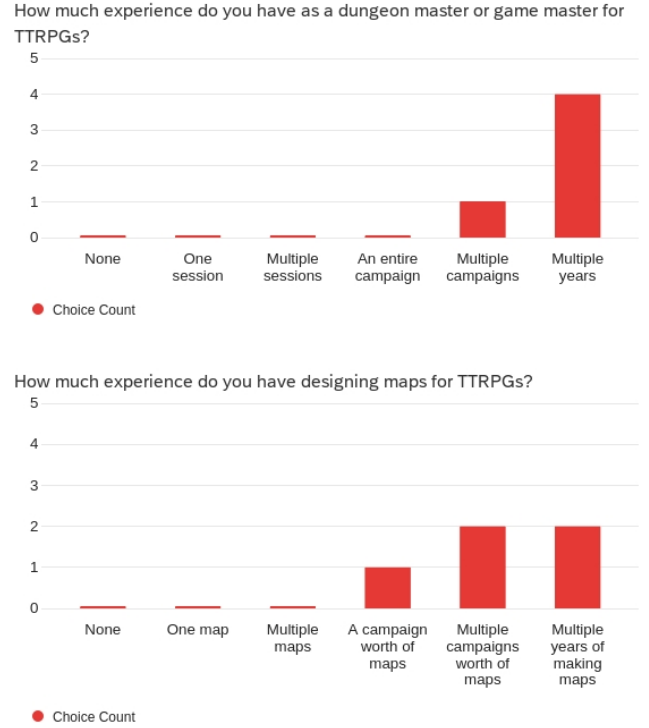


Fig. 2. Demographic results from the post-task survey

Participants were then asked a series of Likert scale questions about the experience of using Envoi and how it facilitated the creative process. These questions can be seen in Fig. 3. We also included a free response field to allow participants to explain their answers if they chose. Finally, we included a few free response questions that could not adequately be answered by a Likert scale. These were:

- What parts of the tool, if any, were confusing?
- What about the tool do you think could be improved?
- What limitations, if any, prevented you from fully expressing your ideas?
- If applicable, does the tool address any parts of map design that you previously found difficult?

III. RESULTS

We had a total of seven participants. However, a bug occurred in the last two sessions that we were not previously aware of that impacted the user’s experience, so we excluded these two participants. The first participant was accidentally shown an outdated version of Envoi, but the differences were minor, and so we do not think our overall conclusion was affected.

A. Thinkaloud

Participants did not speak much in task 1, as it was primarily a tutorial. We did observe, however, that parts of the interface were non-intuitive. For example, a common mistake was to try to select tiles on the map while in the mode to place tiles,

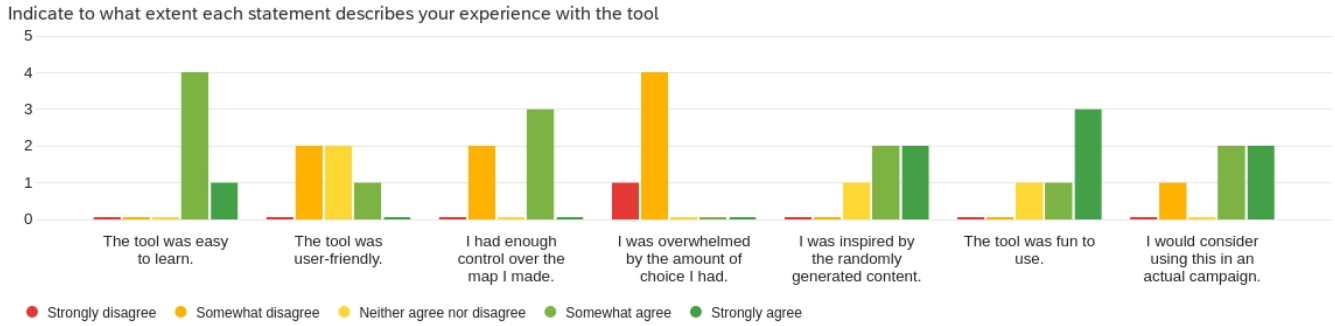


Fig. 3. Results from the Likert scale questions from the post-task survey

resulting in a large group of tiles being placed instead. In the rules interface, it was also non-intuitive that new rules were added to the bottom of the list and the window did not automatically scroll to them.

In task 2, all participants were able to represent the city described to them reasonably well, with participant 4 specifically saying upon finishing the task that they could see themselves using the map in a game session. Interestingly, all participants took the same general approach of manually creating the inner city, which was more organized, and procedurally generating the outer city. An obstacle to this approach was the fact that there was no way to designate specific rules for a specific part of the city.

In task 3, participants made a variety of maps, including an island village, a walled port city, and a city with multiple tiers of walls inspired by *Attack on Titan* [15]. They took a variety of approaches. One participant started by deleting all the rules, seeing what was generated, and then exploring the results of adding different rules. Another started by manually adding a sea, a river, and the important parts of their city and then procedurally generating smaller details.

B. Survey

As shown in Fig. 2, all of our participants had served as a dungeon master or game master for multiple campaigns before, and all but one said they had multiple years of experience in such a role. With regards to specifically making maps, one participant had made maps for a single campaign, two had made maps for multiple campaigns, and two have been making maps for multiple years.

Results of the Likert scale questions are shown in Fig. 3. All participants agreed that Envoi was easy to learn, with one strongly agreeing. Opinions on its overall ease of use were much less favorable, with two participants disagreeing that it was user-friendly, two neutral, and only one agreeing. One participant who agreed that Envoi was easy to learn but did not agree that it was user-friendly added that, “I’ve worked with rule-based interactive systems in the past, and I think that may have contributed to the ease of the tool.”

Despite these difficulties, four out of five participants agreed that Envoi was fun to use, with three strongly agreeing and

none disagreeing. Participants seemed amused by the absurd results the generator sometimes produced (see Fig. 5), and one commented, “I like the idea of random generation with your own rules and such, and the grid layout was easy to work with.” Two participants somewhat agreed and two strongly agreed that they would consider using Envoi in an actual campaign, with one disagreeing. One user did qualify their agreement by saying that if they used it in a campaign they would have to be able to make larger maps in a reasonable amount of time.

As mentioned, participant 1 was accidentally shown an outdated version of Envoi that did not notify the user when the generator was working or when the map was unsatisfiable, did not allow the user to cancel the generation, and had a bug that could cause it to freeze. Perhaps as a result, participant 1 was the only participant who did not agree that Envoi was fun to use, the only participant who said they would not use it in an actual campaign, and the only participant who did not definitely say they think it made map creation easier.

Three out of five participants said they had enough choice over the generated content, but two participants said they did not have enough choice. No participants felt overwhelmed by the amount of choice they had. Two participants somewhat agreed and two strongly agreed that they were inspired by the generated content. Only participant 1 had no opinion. A specific example we observed was when the generator placed a 2x2 square of road tiles near the entrance to participant 3’s city for task 3; they decided to interpret this as a market square and lock it (see Fig. 4).

Participant 1 was not sure if Envoi made map design easier. Participant 2 said it did, saying “Yes, it addresses the problem I have of making maps unique but not too unique wonderfully, by combining procedural generation with human intervention and sculpting!” Participant 5 had a similar response. Participant 3 found the rules system interesting and said that the generator “...provid[ed] inspiration that I wouldn’t have otherwise found through manual decision-making.” Participant 4 said Envoi reduced the tedium of placing large numbers of minor buildings and of creating connected roads and rivers.



Fig. 4. A square that was generated near the entrance to participant 3’s city. They decided to interpret this as a market square and lock it

IV. ANALYSIS

A. User Interface

All participants were confused by at least one part of the user interface. Participants 1 through 4 all expressed confusion with the controls for selecting and placing tiles. Participant 5 commented, “One issue with control is that I wish it was easier to handle multiple map elements at once.” There was also no feedback to the user when a set of tile placements and rules were unsatisfiable. In the future, it will be important to add some sort of debugging functionality.

B. Rules

Participants 1, 2, 4, and 5 were all confused by the fact that the rules applied to tiles they manually placed and not just tiles that were generated, especially since they did not always remember every single tile they placed or locked. This frequently became a limitation in task 2; because of the layout of the city that was described to the participants, it would have been useful to have general rules for the whole city but have a specific district that was an exception to those rules, and Envoi currently does not allow such exceptions.

Participants 1, 3, 4, and 5 all expressed some amount of confusion regarding the rule templates themselves, with participant 5 remarking, “The names for each rule didn’t do an amazing job at explaining how they function.” Even when the literal meaning of a rule template’s text was known, it was sometimes hard to predict its effect on the generated content. For example, at one point participant 4 added a Connection rule that all buildings must be connected by roads, but did not require that there actually be any buildings. The result was that the generator tiled almost the entire map with roads, making the few buildings there were trivially connected (see Fig. 5).

One rule template that users did find easy to use was the “Connection” template, which said “All [Tile A] are connected to all [Tile B] by [Tile C]”. This was used in the base ruleset to make sure that all buildings were connected to all other buildings by roads. Participant 2 described this rule as “...more intuitive than the other ones...” and participant 4 said it solved the problem of making connected networks of roads and rivers,

which they identified as one of the more tedious parts of making maps.

C. The Generator

The generator often behaved in non-intuitive ways, with participant 2 saying, “I will say it is fun, but figuring out how to make it do what you want though is a bit aggravating.” We observed that the generator tends to find “lazy” solutions”, as seen in Fig. 5. This quirk of the generator led to predictable patterns in the generated maps. The predictability limited the exploration potential of Envoi, which is an important characteristic of casual creators [14]. We ran Clingo with the `--rand-freq=0.05` option, leading it to make occasional random decisions. We believe this is the reason behind the few non-road tiles (that were not manually placed) in Fig. 5. We may want to increase the random frequency to make more diverse maps.

Another factor limiting the Envoi’s exploration potential was the time it took to generate large or complex maps. Participants were most successful with map sizes between 12x12 (the size of the map in task 1) and 20x20. Participants 1 and 2 tried to make 30x30 maps, but gave up when the generator was still running after 3 and 2 minutes, respectively.

D. Control

In section I we mentioned that a common critique of MI-CC tools was that users desired more control over the generated suggestions. We seem to have somewhat addressed this critique as a slight majority of participants thought they had enough choice. However, two participants said they did not have enough choice, and all participants mentioned in the free response section at least one limitation that prevented them from fully expressing their ideas. These included the inability to create sections of the map with different rules, the fact that rules applied to tiles the users manually placed, the limited selection of tiles and rule templates, the inability to use a hexagonal grid, and performance constraints preventing large maps from being created.

E. Workflows

A common workflow, which we observed from all participants in task 2 and participant 3 in task 3, was to manually place the important parts of the city and generate the other parts. Participant 3 said, “I think my mindset here is to generate the outer walls myself and let the generator fill in the insides with randomness. I know what I want the shape to be, but the randomness is useful for just simulating culture in random spots.” Features they manually placed included the noble district, river, and sea. They added the major roads manually as well, as they said the generated roads were too “chaotic”.

Interestingly, participant 2 took the opposite approach to task 3. They started by removing all rules and observing the result. They then began to iteratively add and refine rules to create the city they wanted. At the end, they manually added watchtowers, a sea, and a road leading out of the city. They




Fig. 5. The user specified that all buildings must be connected by roads, but did not require a minimum number of buildings. As a result, the generator covered almost the entire map with roads, making the buildings trivially connected.

explained, “For stuff that’s very small, the watchtowers for example, I think it honestly takes more time to add a rule to make a few and also make them more spaced out [than to make them manually], versus with road, it’s nice to have the rules to make it procedurally generate because it makes these interesting weird shapes.” Participant 1 took a similar approach, although they added some rules before generating their first map. Participant 5 used a similar workflow but did not clear the rules; they generated a map from the base ruleset provided and added multiple tiers of inner walls.

F. Evaluation as Casual Creator

In previous work [5], we describe how Envoy implements six casual creator design patterns [14]: Instant Feedback, No Blank Canvas, Limiting Actions to Encourage Exploration, Mutant Shopping, Modifying the Meaningful, and Saving and Sharing. This study provides additional insight into some of them.

1) *Instant Feedback*: Previously, we said that Envoy usually returned a result instantly but remarked that, “...it is conceivable that some rule sets could cause the generator to slow down.” We observed in this study that generating large maps, or small maps with certain combinations of rules, can take an excessive amount of time. Since this became a limitation a couple of times in the study, Envoy will have to be optimized before it truly implements this pattern.

2) *No Blank Canvas*: Users in this study were provided with a basic set of rules and it was suggested, but not required, that they could use it as a starting point for tasks 2 and 3. We

found that this was successful at avoiding the “blank canvas” effect in practice. While not all users used the base ruleset, we did not observe any users struggling to get started with the tasks.

3) *Limiting Actions to Encourage Exploration*: No participants said they felt overwhelmed by the amount of choice they had. This suggests that we have achieved the goal of this design pattern.

4) *Modifying the Meaningful*: Previously we described Envoy as implementing this pattern because the internals are hidden from the user and they only interact with the high-level rules. While this remains the case, in section IV-B we described ways the rules might not have a straightforward effect on the map.

V. CONCLUSION AND FUTURE WORK

To the extent of our knowledge, this is one of the first user studies of an ASP-based tool. We found that Envoy is not very easy to use (RQ1). Despite this, participants broadly found it fun to use (RQ2) and that it addressed parts of map-making that they found difficult (RQ3). Our findings reinforce our belief that Envoy implements a total of five of the eleven casual creator design patterns [14] (RQ4). There is one pattern that we found will require more work (Instant Feedback) and a number that we have already acknowledged as needing more work [5]. Thus, while Envoy can be improved, we have demonstrated the feasibility of ASP-based tools for creative tasks, particularly casual creation.

One participant said that, “The most difficult part of making maps is handling multiple overlaying or adjacent map elements. Such as overlapping buildings or geographic details, or adjacent clustered structures like housing districts.” Multi-layered maps are supported by Tiled [3] and other tools, and it would be a useful feature for future versions of Envoy to support as well.

There is more work that can be done to improve the flexibility of Envoy. Participants 1 and 2 mentioned the limited size of maps the generator could handle as a limitation; it will be important to optimize the generator so that larger maps can be efficiently generated. Participants 1 and 4 mentioned having a limited selection of tiles as a limitation; might be useful to add the ability for the user to define custom tiles. This would not be a difficult feature to support and would greatly expand the range of maps the user can create. Participants 2 and 3 said they felt constrained by the square grid and that a hexagonal grid would allow them to make more organic-looking maps. Since hexagonal maps are common in gaming, it would make sense for Envoy to support them.

All participants in our study were experienced game masters and map makers. Since part of the purpose of Envoy is to be accessible to beginners, it would be important to evaluate it with a less experienced audience. Additionally, all participants in the study were students at an American university, and Envoy was designed with Western fantasy in mind. It would be interesting to see how Envoy is used by users who are familiar with other genres of fantasy.

ACKNOWLEDGMENTS

We would like to thank John Thomas Bacher (jtbacher@ncsu.edu) for his contributions to the design and development of Envoi and the POEM lab at North Carolina State University for their feedback.

REFERENCES

- [1] MDN Contributors. (2022, Feb.) Tiles and tilemaps overview. Mozilla. [Online]. Available: <https://developer.mozilla.org/en-US/docs/Games/Techniques/Tilemaps>
- [2] GameFromScratch.com. (2015, Nov.) Tiled map editor tutorial series. [Online]. Available: <https://gamefromscratch.com/tiled-map-editor-tutorial-series/>
- [3] T. Lindeijer. (2022, Feb.) Tiled. [Online]. Available: <https://www.mapeditor.org/>
- [4] drow. (2021) Fantasy town generator. [Online]. Available: <https://donjon.bin.sh/fantasy/town/>
- [5] D. Carpenter, J. T. Bacher, H. Crain, and C. Martens, “Casual creation of tile maps via authorable constraint-based generators,” presented at 1st Workshop on Programming Languages and Interactive Entertainment, Online, Oct. 2021.
- [6] G. N. Yannakakis, A. Liapis, and C. Alexopoulos, “Mixed-initiative co-creativity,” in *Proc. 9th International Conference on the Foundations of Digital Games (FDG’14)*, Fort Lauderdale, FL, United States, Apr. 2014.
- [7] A. Alvarez, S. Dahlskog, J. Font, J. Holmberg, C. Nolasco, and A. Osterman, “Fostering creativity in the mixed-initiative Evolutionary Dungeon Designer,” in *Proc. 13th International Conference on the Foundations of Digital Games (FDG’18)*, Malmo, Sweden, Aug. 2018, pp. 1–8.
- [8] A. Liapis, G. N. Yannakakis, and J. Togelius, “Sentient Sketchbook: Computer-aided game level authoring,” in *Proc. 8th International Conference on the Foundations of Digital Games (FDG’13)*, Chania, Crete, Greece, May 2013, pp. 213–220.
- [9] N. Davis, C.-P. Hsiao, K. Y. Singh, and B. Magerko, “Co-creative drawing agent with object recognition,” in *Proc. 12th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AAAI’16)*, Burlingame, CA, United States, Oct. 2016, pp. 9–15.
- [10] A. M. Smith and M. Mateas, “Answer set programming for procedural content generation: A design space approach,” *IEEE Trans. Comput. Intell. and AI in Games*, vol. 3, no. 3, pp. 187–200, Sep. 2011.
- [11] M. Kreminski, M. Dickinson, J. C. Osborn, A. Summerville, M. Mateas, and N. Wardrip-Fruin, “Germinate: A mixed-initiative casual creator for rhetorical games,” in *Proc. 16th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AAAI’20)*, Online, Oct. 2020, pp. 102–108.
- [12] X. Neufeld, S. Mostaghim, and D. Perez-Liebana, “Procedural level generation with answer set programming for general video game playing,” in *Proc. 7th Computer Science and Electronic Engineering Conference (CEECE ’15)*, Colchester, United Kingdom, Sep. 2015, pp. 207–212.
- [13] *Information technology — Programming languages — Prolog — Part 1: General core*, ISO/IEC Std. 13211-1:1995, 2008.
- [14] K. Compton and M. Mateas, “Casual creators,” in *Proc. 6th International Conference on Computational Creativity (ICCC’15)*, Park City, UT, United States, Jun. 2015, pp. 228–235.
- [15] H. Isayama, *Attack on Titan*. Tokyo, Japan: Kodansha, 2009.