An Interaction Framework for Studying Co-Creative AI

Matthew Guzdial
mguzdial3@gatech.edu
Georgia Institute of Technology
Atlanta, Georgia

Mark Riedl
riedl@cc.gatech.edu
Georgia Institute of Technology
Atlanta, Georgia

ABSTRACT

Machine learning has been applied to a number of creative, design-oriented tasks. However, it remains unclear how to best empower human users with these machine learning approaches, particularly those users without technical expertise. In this paper we propose a general framework for turn-based interaction between human users and AI agents designed to support human creativity, called co-creative systems. The framework can be used to better understand the space of possible designs of co-creative systems and reveal future research directions. We demonstrate how to apply this framework in conjunction with a pair of recent human subject studies, comparing between the four human-AI systems employed in these studies and generating hypotheses towards future studies.

CCS CONCEPTS

• **Human-centered computing** → **Interaction design theory, concepts and paradigms**.

KEYWORDS

co-creation, machine learning, frameworks

INTRODUCTION

In recent years there have been many machine learning (ML) systems that demonstrate the ability to create in domains like storytelling, music, visual art, and games [5, 16–18]. Modern ML systems have a high technical knowledge requirement. Making creative artifacts with ML often requires directly tweaking a system's parameters, which does not fit naturally into creative workflows [10]. *Co-creative systems*, also called mixed-initiative [1] or centaur [6] systems, serve as an alternative paradigm to fully autonomous systems by allowing a human creator to work alongside an AI algorithm toward a joint goal. Co-creative systems can help to democratize creative ML algorithms and support human creativity. However, the question of how to design the interaction between ML algorithm and human creator remains under-explored.

There have been several prior attempts to classify the kinds of interaction paradigms computers and humans might engage in. Three common classes are: (1) creativity support tool, (2) autonomous, generative system, and (3) co-creative system [3, 12]. Lubart [15] suggested that computers can play four roles in creativity support: *nanny*, *pen-pal*, *coach*, and *colleague*. These prior frameworks or classifications are under-specified, making it difficult to apply them to meaningfully compare between systems for the purpose of evaluation or hypothesis generation.

In this paper we present a framework for representing turn-based, co-creative systems. In particular, we present structure and language for identifying and discussing the common components of these systems across domains. we demonstrate how this framework can be used to represent Lubart's four roles, and the three common classes described above. We further present the utility of this framework for eliciting new research hypotheses about co-creative systems in a case study based on prior human subject studies and simulated tests. This framework is designed for comparing existing systems, as a framing device during initial design of novel systems, and as a tool for generating hypotheses.

FRAMEWORK

We present our framework in Figure 1. It describes systems with a human **user** and **AI** agent collaboratively working to design some **artifact** in a turn-based fashion. We employ a turn-based framework because it is a common way of organizing co-creative interactions [3, 12, 13] and because it suits evolutionary and reinforcement-learning approaches that require discrete steps [2, 7, 8, 14].

Each component of this framework is meant as both a representation of some part of a system and a means of prompting design reflection:

(1) **Start**: How does the interaction start? Who has the first turn?
(2) **Actions**: What actions is each partner capable of based on implementation and interface design.
  (a) **Artifact Actions**: During each turn, what are the actions each partner can take during their turn that directly impact the artifact?
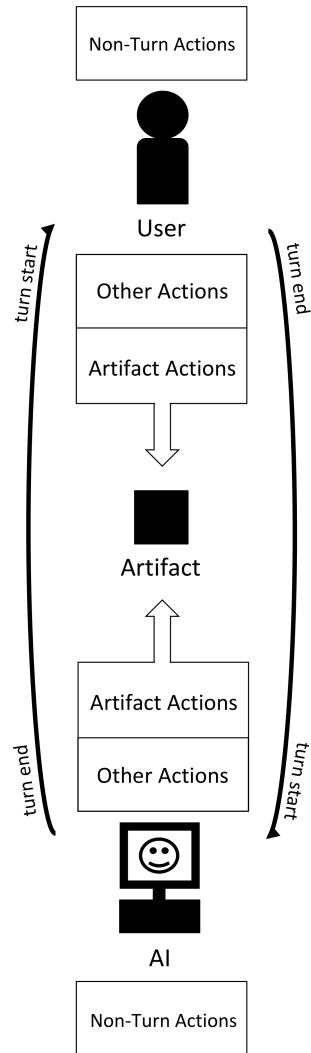
Figure 1: Illustration of our framework.

(b) **Other Actions**: During each turn, what are the actions that do not directly impact the artifact? For example, an AI agent explaining its actions or internal state. [9, 19].

(c) **Non-Turn Actions:** What actions can each partner take outside of their turn? For example does the AI attempt to actively learn from or model the user. [7, 13]? Does the user have the ability to observe or explore the artifact?

(3) **AI**: What algorithm or technique is used to train or design the AI agent? Is this agent static, choosing between a series of existing behaviors, or learning during the interaction?

(4) **User**: Who is the imagined or intended user for this system? Prior work has demonstrated the importance of designing these systems towards a particular target user [11].

(5) **Turns**: How do turns begin and end? Which partner has agency over these decisions?

(6) **End**: How does the interaction end? Which partner or conditions must be satisfied?

Comparison to Prior Frameworks

Three common ways of describing interaction between an AI agent and a human user are: (1) creativity support tool, (2) autonomous, generative system, and (3) co-creative system. All three of these can be expressed with our proposed framework. First, a creativity support tool can be represented in this framework with the AI agent having no Artifact Actions, only taking actions when the user is not. Second, an autonomous generative system instead involves the user not having any Artifact Actions, but a single other action that starts the AI process that completes the artifact, thus ending the AI agent's only turn. The third involves a human and an AI agent taking turns and most clearly fits this framework. Still, there are many ways in which humans and AI agents can co-create and our framework allows for a more fine-grained level of description.

Lubart [15] presents an influential framework that classified computer-human interactions into four categories: *nanny*, *pen-pal*, *coach*, and *colleague*. A computer-as-nanny system conducts regular check-ins to monitor, support, or manage the user. Our framework can represent nanny-like interactions as Other Actions and Non-turn Actions taken by an AI agent that has no Artifact Actions. A computer-as-pen-pal system sends and receives messages from the human creator in order to encourage creative reflection. Our framework represents this by removing Artifact Actions from both user and AI agent. A computer-as-coach system teaches the user strategies and skills for creation. Our framework can represent a coach with an empty set of Artifact Actions and using Other Actions and Non-turn Actions that teach and provide constructive feedback. A computer-as-colleague system acts as a peer, taking turns to contribute to the artifact. In our representation, this is the only interaction configuration in which the AI agent can make Artifact Actions. A computer-as-colleague system can also have Other Actions and Non-turn Actions, providing a wide range of possible agent interaction designs.
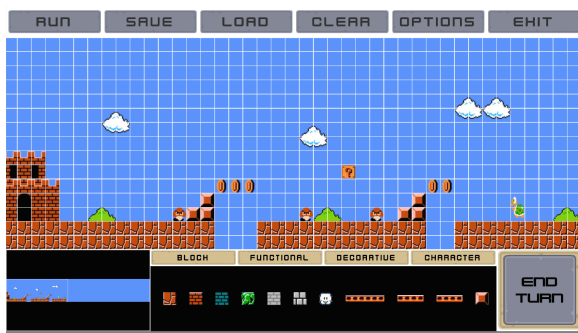
Figure 2: Screenshot of the interface for the systems in the first study.

[1]A video of the Bayes Net agent can be found here: https://youtu.be/6Ska6Y9Wnvo

[2]Video of this agent can be found here: https://youtu.be/UkMeM5Ty1lA

Case Study

Our co-creative agent framework allows us to draw comparisons between any two systems that can be represented within it. In this section we discuss our recent work on studying co-creativity in the context of a *Super Mario Bros.* level design editor. Two studies, described in [7], sought to qualitatively understand how users of co-creative systems thought of the AI agent and compared different ML-based AI partners. The first study interface is shown in Figure 2 and allowed users to interact with one of three distinct ML-based AI partners in the backend: an LSTM, a Markov Chain, and a Bayes Net.[1] Between the studies we developed the interface further, removing an ability for the user to play their current level during their turn, added the ability to ask a study conductor for an explanation of the AI agent's last action (simulating a future explainable AI sub-system), and added a button corresponding to a user action undoing all of the AI agent's last turn. We also replaced the AI agent with an active learning, deep reinforcement agent.[2] We note that this paper's framework is not described in [7].

For an example of how our framework helps guide the study of co-creative systems, we can describe the impact the AI algorithm (Markov Chain, Bayes Net, and LSTM) had on the user's perceptions of the co-creative system. We found surprisingly few statistically significant differences in user perception across these systems. There was a significant results found that the Bayes Net was more creative in behavior than the LSTM, that the LSTM was more challenging to use than the Bayes Net, and that the Markov Chain was more creative in behavior than the Bayes Net. Represented in the language of our framework we anticipate this was due to the Bayes Net producing less changes per action compared to the LSTM or Markov Chain, while also including actions to add decorative elements that both other agents did not. All three agents had an "Other Action" for observing the current level, which differed in scale. The Markov Chain perceived only small windows of the current level, the Bayes Net a full screen, and the LSTM almost the entire level. This influenced the results we found, with the Markov Chain adding more locally interesting content, leading to the higher perception of its creativity.

Changing the other actions available to the user changed their experience working with the co-creative system. In the second study using a Deep RL agent, we added an undo button. Surprisingly, users kept an average of 15% more of the AI agent's changes compared to the three systems without an undo button and self-reported 15% less frustration over all. While the latter seems to follow from the addition of the undo button, the former goes against expectations. However, this undo button could certainly have helped the Deep RL AI agent's non-turn action of actively learning from the user, giving a very clear signal when its changes weren't appreciated and allowing it to learn faster.

Including active learning as a non-turn action for the Deep RL AI agent improved simulated test performance [8] and human self-reports [7]. In the second study we were surprised to find that despite

designing the Deep RL System with Lubart's computer-as-colleague interaction framework in mind, users perceived the agent as a friend, collaborator, student, and manager [7]. This suggests that even with an otherwise fixed system, variances in users can lead to different perceptions of the AI partner.

Research Hypothesis Generation

The above case study is not exhaustive in terms of research questions about how to design more effective co-creative systems. The framework suggests that there may be human-subject studies involving other parts of the framework that the case study didn't address. For example, the framework tells us that research questions can be crafted around how turns are ended. Anecdotal evidence suggests that users that interacted with the AI more had a more positive experience. Could the frequency of turns be altered to further increase positive perceptions of the creative experience, perhaps through predictive interruptions? Likewise, should the user be able to interrupt the AI agent?

There are numerous types of other actions that can be considered where the user interacts with or explores what they have created without directly changing the artifact. For example, testing or any means of deriving additional observations of the artifact would be considered an other action.

What can and should the user be doing while the AI agent is taking its turn to increase engagement? Can the user request explanations for what the AI agent is doing when changing the artifact? Explanations have been shown to positively affect user perceptions of intelligent systems [4, 9]. Should explanations be offered automatically as an Other Action belonging to the AI agent?

CONCLUSIONS

In this paper we present a general framework for representing turn-based co-creative systems. Such a framework allows for direct comparisons between distinct systems, from which one can generate discussion, hypotheses, and potential generalizations on the impact of these differences. We anticipate that this framework can further serve as a means of eliciting design reflection during the development stages of such a tool. Organizing frameworks like this can serve as a vector for developing general design knowledge for best practices in developing usable machine learning approaches.

ACKNOWLEDGMENTS

REFERENCES

[1] JE Allen, Curry I Guinn, and E Horvtz. 1999. Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications* 14, 5 (1999), 14–23.

[2] Alberto Alvarez, Steve Dahlskog, Jose Font, Johan Holmberg, Chelsi Nolasco, and Axel Österman. 2018. Fostering creativity in the mixed-initiative evolutionary dungeon designer. In *Proceedings of the 13th International Conference on the Foundations of Digital Games*. ACM, 50.

[3] Nicholas Davis, Chih-Pin Hsiao, Yanna Popova, and Brian Magerko. 2015. An enactive model of creativity for computational collaboration and co-creation. In *Creativity in the Digital Age*. Springer, 109–133.

[4] Upol Ehsan, Pradyumna Tambwekar, Larry Chan, Brent Harrison, and Mark Riedl. 2019. Automated Rationale Generation: A Technique for Explainable AI and its Effects on Human Perceptions. In *Proceedings of the International Conference on Intelligence User Interfaces*.

[5] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2015. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576* (2015).

[6] Ian M Goldstein, Julie Lawrence, and Adam S Miner. 2017. Human-machine collaboration in cancer and beyond: the centaur care model. *JAMA oncology* 3, 10 (2017), 1303–1304.

[7] Matthew Guzdial, Nicholas Liao, Jonathan Chen, Shao-Yu Chen, Shukan Shah, Vishwa Shah, Joshua Reno, Gillian Smith, and Mark Riedl. 2019. Friend, Collaborator, Student, Manager: How Design of an AI-Driven Game Level Editor Affects Creators. *CHI Conference on Human Factors in Computing Systems* (2019).

[8] Matthew Guzdial, Nicholas Liao, and Mark Riedl. 2018. Co-Creative Level Design via Machine Learning. *arXiv preprint arXiv:1809.09420* (2018).

[9] Matthew Guzdial, Joshua Reno, Jonathan Chen, Gillian Smith, and Mark Riedl. 2018. Explainable PCGML via Game Design Patterns. *arXiv preprint arXiv:1809.09419* (2018).

[10] Bhautik Joshi, Kristen Stewart, and David Shapiro. 2017. Bringing impressionism to life with neural style transfer in come swim. In *Proceedings of the ACM SIGGRAPH Digital Production Symposium*. ACM, 5.

[11] Anna Kantosalo, Jukka M Toivanen, Ping Xiao, and Hannu Toivonen. 2014. From Isolation to Involvement: Adapting Machine Creativity Software to Support Human-Computer Co-Creation.. In *ICCC*. 1–7.

[12] Pegah Karimi, Kazjon Grace, Mary Lou Maher, and Nicholas Davis. 2018. Evaluating Creativity in Computational Co-Creative Systems. *arXiv preprint arXiv:1807.09886* (2018).

[13] Antonios Liapis, Georgios N Yannakakis, and Julian Togelius. 2013. Designer modeling for personalized game content creation tools. In *Proceedings of the AIIDE workshop on artificial intelligence & game aesthetics*.

[14] Antonios Liapis, Georgios N Yannakakis, and Julian Togelius. 2013. Sentient Sketchbook: Computer-aided game level authoring.. In *Proceedings of ACM Conference on Foundations of Digital Games*. FDG, 213–220.

[15] Todd Lubart. 2005. How can computers be partners in the creative process: classification and commentary on the special issue. *International Journal of Human-Computer Studies* 63, 4-5 (2005), 365–369.

[16] Lara J Martin, Prithviraj Ammanabrolu, Xinyu Wang, William Hancock, Shruti Singh, Brent Harrison, and Mark O Riedl. 2018. Event representations for automated story generation with deep neural nets. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

[17] Yuki Saito, Shinnosuke Takamichi, and Hiroshi Saruwatari. 2018. Statistical parametric speech synthesis incorporating generative adversarial networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 26, 1 (2018), 84–96.

[18] Adam Summerville, Sam Snodgrass, Matthew Guzdial, Christoffer Holmgård, Amy K Hoover, Aaron Isaksen, Andy Nealen, and Julian Togelius. 2017. Procedural Content Generation via Machine Learning (PCGML). *arXiv preprint arXiv:1702.00539* (2017).

[19] Jichen Zhu, Antonios Liapis, Sebastian Risi, Rafael Bidarra, and G Michael Youngblood. 2018. Explainable AI for Designers: A Human-Centered Perspective on Mixed-Initiative Co-Creation. *Computational Intelligence in Games* (2018).