Human-AI Interaction Generation: A Connective Lens for Generative AI and Procedural Content Generation

Matthew Guzdial
University of Alberta, Alberta Machine Intelligence Institute (Amii)
guzdial@ualberta.ca

Abstract

Generative AI has recently gained popularity as a paradigm for content generation. In this paper, we link this paradigm to an older one: Procedural Content Generation (PCG). We propose a lens to identify the commonalities between both paradigms that we call human-AI interactive generation. Using this lens, we identify three beneficial attributes then survey recent related work and summarize relevant findings.

1 Introduction

Recent generative AI tools for text and image generation have created a large response in the general public and academic community. These models generally take prompts, often natural language instructions, as input. They then typically output content in whatever media format a tool focuses on in a single-shot manner. While this paradigm for generating content can work well in many instances, it is ill-suited to many others, leading in some part to the wide range of responses to these tools [Vimpari *et al.*, 2023].

Generative AI tools did not invent the concept of generating content with computers. Procedural Content Generation (PCG) dates back to 1983 with the game *Beneath Apple Manor* for the Apple II, and has benefited from decades of research since then [Shaker *et al.*, 2016a]. PCG refers generally to processes (procedures) for generating content with computers, and can be broken into constructive or rules-based methods, search-based methods [Togelius *et al.*, 2011], and machine learning-based methods [Guzdial *et al.*, 2022b], including deep learning methods [Liu *et al.*, 2021]. Recent research has directly investigated applications of Large Language Models (LLMs) to typical PCG problems [Gallotta *et al.*, 2024]. However, we can argue that generative AI tool and PCG are the same process at the highest level, as PCG is broadly defined as procedures for content generation. With this perspective, we can identify how existing PCG research and open research may benefit the future development of these tools.

In this paper, we formally define human-AI interactive generation as a lens to understand the commonalities of PCG and generative AI. We then identify three attributes of this lens that are less present in generative AI tools than existing PCG research. We survey prior work related to these attributes and outline opportunities for future work.

2 Human-AI Interactive Generation

Procedural Content Generation (PCG) research has a number of frameworks or lenses for understanding content generation processes that involve humans [Shaker *et al.*, 2016b; Guzdial *et al.*, 2022a]. However, these frameworks tend to divide generation into autonomous (AI system alone) or mixed-initiative (AI and human together). We argue this is an artificial divide, when even "autonomous" generation processes still require human prompting or to otherwise start the process. We identify human-AI interactive generation as an alternative lens in terms of the below equation, meant to illustrate the commonalities between PCG and Generative AI paradigms. We intentionally leverage the basic function formulation common across many AI paradigms.

$$Y = G(X) \tag{1}$$

In the above equation, Y indicates the set of all sets of possible outputs across all generation processes. The theoretical concept of the set of all possible outputs for a single generation process is historically referred to as a *possibility space* [Rabii and Cook, 2023; Smith, 2017]. A single generation process is similarly referred to as a *generator*. For our lens, we can define the possibility space as Y_g, the set of all possible outputs for a particular generator g. A particular output can then be rendered as y. For example, $Y_{ChatGPT}$ might indicate all possible combinations of strings up to the maximum size of a specific version of ChatGPT's output or $Y_{StableDiffusion}$ might indicate all possible RGB images that Stable Diffusion can output.

G in the above indicates the space of all possible generators. To parallel possibility space we can refer to this as the theoretical *generator space* [Khalifa, 2020]. Though this is a less-investigated concept in PCG literature, it is related to the concept of a hypothesis space of potential generative models. We can then define an individual generator as g or the set of generators of a specific type as G_{Type}, such as G_{LLM} for the set of all the possible generators that involve a Large Language Model. However, a generator g does not need to be a single model or AI approach, such as in the case of ChatGPT which leverages outside services like WolframAlpha as part of its generation process.

X then gives us the set of all possible inputs to all possible generators. A particular input x would then define a specific input, for example a specific prompt and context to an LLM or random noise for a constructive generator [Shaker *et al.*, 2016a]. Our definition of X considers all possible inputs, even if the vast majority would not lead to valid outputs for a generator. We refer to the set of inputs that would lead to a valid output for a generator g as X_g. For example, Trajkova et al.'s "LuminAI" dance partner can be understood as a generator that takes in a recording of human dance x_i and outputs a space of possible dances in the form of a projected character $Y_{LuminAI}$ [Trajkova *et al.*, 2024]. We could not feed this input x_i to ChatGPT and expect it to function, but neither could we feed an appropriate prompt and context x_j for ChatGPT into LuminAI. Thus X_g may change as the generator changes, as with the increasing prevalence of multimodal Generative AI tools.

In the above paragraphs, we have attempted to demonstrate how this simple human-AI interactive generation lens allows us to consider the commonalities between PCG and Generative AI systems. In addition, this lens allows us to identify several desirable attributes of these systems. We will dive further into the commonalities for PCG and generative AI by investigating these attributes.

Human-centered input alignment represents the extent to which X_g for a particular generator g aligns with human expectations for that generator. We contrast this to the general use of the term alignment which largely focuses on the behavior of a generator in terms of its output [Ji *et al.*, 2023; Lopes and Bidarra, 2011]. *Adaptability* is the extent to which the generator g adapts to the user over a sequence of inputs (x_0, x_1, ... x_n) and outputs (y_0, y_1, ... y_n). This in turn is related to continual and lifelong learning, along with player adaptation research, but focuses specifically on generation processes [Parisi *et al.*, 2019]. *Novelty* is the extent to which a generator g can produce outputs within Y_g that go beyond its learned or authored distribution, in other words out-of-distribution generation. This is related to several concepts from core machine learning, including out-of-distribution detection [DeVries and Taylor, 2018] and generalization [Zhang *et al.*, 2021]. We will overview recent work and results related to these attributes in the following sections.

3 Human-centered Input Alignment

The most common valid input for Generative AI tools are prompts. While initially constrained to natural language, Generative AI tools are increasingly focused on increasing the number of valid input modalities, leading to modern multimodal models [Liu, 2023]. In other words, Generative AI tool developers seem to seek to expand X_g to reach X. However, we argue this actually runs counter to successful generation processes in many instances, thus our proposal of the human-centered input alignment attribute. This is part of the broader concept of human-centered design, which argues design processes should be driven by the needs of human users [Cooley, 2000]. This attribute then considers the extent to which the valid inputs for a generator correspond to user preferences over possible valid inputs.

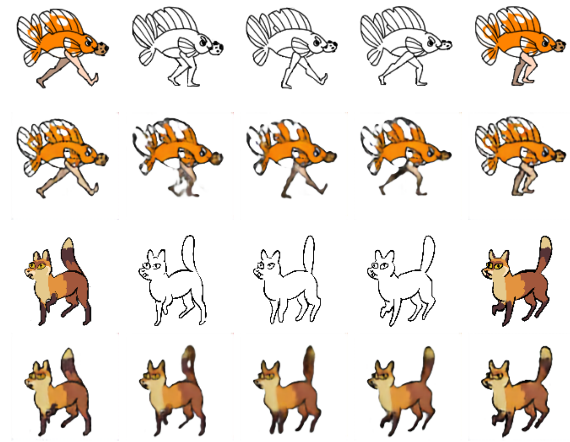

Figure 1: Animations by two artists. For each section the top row is an animation with artist-drawn sketches between the first and last frame. The bottom row is the SketchBetween output.

This attribute can encourage us to consider possible inputs from the start, impacting the design of our models. In Sketchbetween, we developed a tool to color in the in-between frames for hand drawn animation using a vector-quantized variational auto-encoder (VQ-VAE) [Loftsdottir and Guzdial, 2022]. We visualize inputs from human animators and Sketchbetween's outputs in Figure 1. To mirror the human design process, our generator takes in two colored keyframes and sketched in-between frames. While many Generative AI models could take an input image prompt like this they may struggle with consistent output [Saravanan and Guzdial, 2022; Löwenström, 2024]. In comparison, SketchBetween outperformed more general image inpainting baselines. We have continued to find success when taking this tact, specifically designing generators to align their valid input sets with user expectations [Halina and Guzdial, 2021b; Cooper and Guzdial, 2023; Halina and Guzdial, 2023].

We provide an example of the consequences of not achieving this attribute in Anhinga, a puzzle design tool based on the game *Snakebird* [Guzdial *et al.*, 2021]. In Anhinga, we included an exhaustive PCG system capable of generating variations of a current puzzle that maximally altered an input puzzle's difficulty while minimally altering its structure [Sturtevant *et al.*, 2020]. Crucially, Anhinga would run this generative process any time a user made any change to a puzzle, functionally meaning that every single user change represented an input to our generator. Based on a human subject study, we found that designers using Anhinga did find the generator easy to use but significantly preferred a version of the tool without the system that responded to every user input. We consider this to be due to cognitive overload based on the rate of feedback. Of greater concern was that in a comparative ranking, designers believed a puzzle they made without the generator would be more challenging than one they made with it, but we found objective evidence for the opposite. This suggests that lacking human-centered input alignment may harm a user's ability to evaluate the outputs of a generation process.

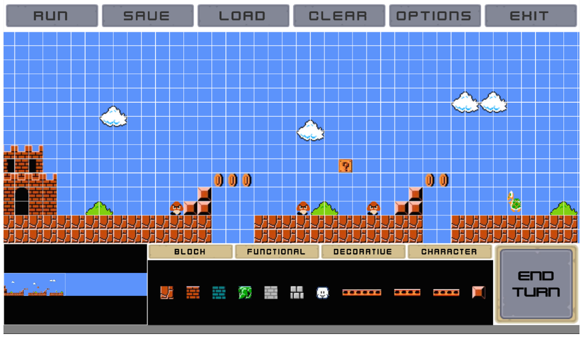

Figure 2: Screenshot of the Morai Maker level editor.

4 Adaptability

Modern generative AI and PCG tools rarely adapt to a user in a significant fashion. In generative AI tools, this is due to the large financial and time cost of retraining models, which makes such an adaptation ill-suited to real-time interactive sessions. In PCG tools, significant adaptation requires an additional development effort in terms of modelling the designer, which may not be feasible except in a post-hoc fashion [Alvarez *et al.*, 2022]. However, due to the diversity in the design practices of human designers, and more generally the differences across humans, we have found a strong positive response to models that adapt to users [Mahmoudi-Nejad *et al.*, 2021].

We developed Morai Maker as a platform to study human-AI interactions in level generation and design for the game *Super Mario Bros.*, we visualize its interface in Figure 2. We employed a turn-based interaction modality [Guzdial and Riedl, 2019], which meant that a human designer and an AI agent took turns placing down pieces of level structure, with the human ending their turn with the large “End Turn” button at the bottom right of the interface. We ran an initial comparative study between three non-adaptive machine learning agents (Markov Chain, Bayes Net, and LSTM), and found no significant difference between user preferences across the approaches [Guzdial *et al.*, 2019], despite measurable differences between the possibility spaces Y_g of each generator [Summerville, 2018]. In a follow-up study, we replaced these static, non-adaptive agents with a deep reinforcement learning agent that learned from implicit human designer feedback [Guzdial *et al.*, 2018], and found that users could identify and preferred the adaptation [Guzdial *et al.*, 2019].

We found further support for the importance of adaptability in human-AI interactive generation with KiaiTime. KiaiTime was a level design tool we developed for the game *Taiko no Tatsujin* [Halina and Guzdial, 2021a]. We ran a study with KiaiTime comparing the naive, implicit adaptation strategy we found success with in Morai Maker to Threshold Designer Adaptation (TA) [Halina and Guzdial, 2022]. In TDA, we learn a domain-dependent adaptation hyperparameter by approximating interactive design sessions from pre-existing outputs, treating these outputs as if they were from KiaiTime $Y'_{KiaiTime}$. We found that designers enjoyed working with the more advanced adaptatibility approach and that this led to a greater variety of final outputs.

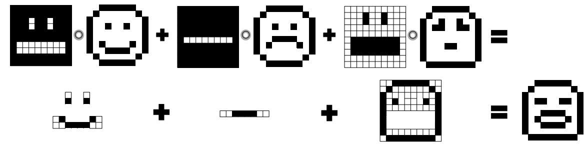

Figure 3: Visualization of a Conceptual Expansion combining a series of bitmoji faces to produce a new one. The alpha filters to the left of each face determine what parts of each given feature to combine into the final output.

5 Novelty

Out-of-distribution generation exists as a term for approaches that alter a generator such that its generative distribution shifts to encompass an area outside of its initial generative distribution [Lotfollahi *et al.*, 2020]. However, this approach typically relies on training data from a secondary distribution and a transfer learning methodology. We do consider this an example of the novelty attribute in human-AI interactive generation, but also look beyond it to cases where we may have no data from the out-of-distribution region that encompasses our desired outputs [Sarkar *et al.*, 2023]. This is related to the distinction between P-creativity and H-creativity [Boden, 2009], we consider cases where examples of the kinds of output we wish to generate may not exist in history. While this may sound impossible, it is possible to approximate such a generator through secondary features of the desired outputs.

Combinational creativity, also called conceptual combination, is a common human cognitive practice for approximating novel knowledge by combining existing knowledge [Wisniewski and Gentner, 1991; Boden, 2010]. Existing approaches have attempted to replicate this process in computers, most commonly through an approach referred to as conceptual blending [Fauconnier and Turner, 2003]. Conceptual blending generally functions by combining two existing pieces of content to produce a novel piece of content, essentially acting as a generator g that takes two inputs and returns an output that represents their combination [Goguen and Harrell, 2004]. We instead propose taking two arbitrary generators g_i and g_j and combining them, which can be understood as an example of PCG via Knowledge Transformation (PCG-KT) [Sarkar *et al.*, 2023].

We initially proposed directly combining learned video game level generation models via combining learned Bayes Nets for *Super Mario Bros.* level generation via conceptual blending [Guzdial and Riedl, 2016]. We found that the “blended” Bayes Nets better-approximated blended game levels from human expert. We followed this work by proposing a novel combinational creativity algorithm we called conceptual expansion, which was better suited to dealing with matrix-based, machine-learned features.

$$CE(F, \mathrm{A}) = \alpha_1 * f_1 + \alpha_2 * f_2 ... \alpha_n * f_n \quad (2)$$

We define a conceptual expansion as a linear combination of potentially repeated features F and α filters A. We can then optimize this combination of existing features and α filters to approximate novel features, typically through a search-based process. We visualize a conceptual expansion in a toy example in Figure 3. In an initial experiment replicating Bayes Net *Super Mario Bros.* level design task, we found that conceptual expansion outperformed conceptual blending and other PCG-KT approaches in terms of the variety of possible outputs, the size of the Y_g set [Guzdial and Riedl, 2018c].

To move beyond video game levels, we next attempted to generate whole novel video games. We machine learned symbolic rule-based models of three video games from gameplay video, which we might now consider examples of world models [Guzdial *et al.*, 2017]. We applied conceptual expansion over these world models, combining them to produce novel world models for unseen games. In an initial leave-one-out experiment, we found that conceptual expansions could better approximate held-out, human-authored games than relevant game generation baselines [Guzdial and Riedl, 2018a]. We then applied this approach to generate fully novel games, and found in a human subject study that participants ranked the generated games as more surprising than games authored by human expert game developers and equally as creative [Guzdial and Riedl, 2021].

We found deep neural networks (DNNs) to be an excellent fit for conceptual expansion due to their learned matrices of features. We initially found that conceptual expansion on DNNs could allow for low-data transfer learning for image classification and generation tasks [Guzdial and Riedl, 2018b]. We since have applied this approach to recombine features from MusicVAE to create a model for generating Iranian Folk Music, which outperformed transfer learning baselines [Doosti and Guzdial, 2023].

We do not anticipate that conceptual expansion is the only approach for achieving novelty. However, we have demonstrated across a number of models and domains that it is a robust approach for approximating unseen or nonexistent knowledge in generation processes. We therefore identify it as a valuable comparison point for future approaches at truly novel content generation.

6 Future Work

In this paper, we have surveyed recent work and findings related to three key attributes of human-AI interactive generation. However, there are still a large number of open problems related to each attribute. In addition, there are likely many more underexplored or even unidentified attributes of similar importance to these three.

For human-centered input alignment, the research thus far has been piecemeal, considering domain-dependent input modalities and evaluating their impact independently. There is the possibility of deriving a theory to guide this practice, but we anticipate full automation may be more useful. As with Threshold Designer Adaptation, we imagine extracting beneficial input modalities and other relevant design features from existing content or records of creative processes [Halina and Guzdial, 2022] We identify this as a similar impulse to automating science, with many of the same methods likely relevant [Waltz and Buchanan, 2009].

For adaptability, we have demonstrated consistent results supporting user preference for adaptive generators [Guzdial *et al.*, 2019; Mahmoudi-Nejad *et al.*, 2021; Halina and Guzdial, 2022]. However, we have evidence that adaptation can fail to outperform random behavior [Yu *et al.*, 2021]. As such, it's clear that further research is required to differentiate between these cases. Outside of this more general need, there are specific questions in terms of when to adapt, how to adapt, and by how much. We anticipate significant overlap with other areas of human-AI interaction, such as teachable robots [Thomaz and Breazeal, 2008].

For novelty, we have identified that computational models of combinational creativity or conceptual combination can successfully lead to out-of-distribution generation without any training data [Guzdial and Riedl, 2016; Guzdial and Riedl, 2021]. We can more broadly consider these tasks to be related to zero-shot transfer learning. As such, we anticipate these approaches may be beneficial for more general zero-shot transfer learning tasks, and we have already found evidence for this in the financial domain [Mahajan and Guzdial, 2022]. We hope to more fully explore potential applications, particularly in low data domains like medicine, but we also anticipate more fundamental research into novelty in generation. In particular, we hope to determine how human users can best guide and benefit from novelty.

We have defined three beneficial attributes in terms of our human-AI interactive generation lens, informed by prior work. However, we do not believe this is an exhaustive set. In particular, each attribute focused on only one aspect of our lens, but we anticipate benefits from studying the interplay and relationships between these aspects.

7 Conclusions

In this paper we introduced the lens of human-AI interactive generation to identify connections between the Generative AI and PCG paradigms. Based on this lens we defined three attributes: human-centered input alignment, adaptability, and novelty. We surveyed recent work and results related to these attributes, and outlined opportunities for future work. Our hope is that this lens will contribute to an ongoing dialogue between Generative AI and PCG researchers, in order to share future contributions across both fields.

Acknowledgments

The work included in this article was funded through an NSERC Discovery Grant and Canada CIFAR AI Chair at the Alberta Machine Intelligence Institute (Amii). I am deeply grateful to all of my collaborators, students and otherwise, who made this work possible.

References

[Alvarez *et al.*, 2022] Alberto Alvarez, Jose Font, and Julian Togelius. Toward designer modeling through design style clustering. *IEEE Transactions on Games*, 14(4):676–686, 2022.

[Boden, 2009] Margaret Boden. Creativity: How does it work. *The idea of creativity*, 28:237–50, 2009.

[Boden, 2010] Margaret A Boden. *Creativity and art: Three roads to surprise*. Oxford University Press, 2010.

[Cooley, 2000] Mike Cooley. Human-centered design. *Information design*, pages 59–81, 2000.

[Cooper and Guzdial, 2023] Seth Cooper and Matthew Guzdial. path2level: Constraint-based level generation from paths. In *2023 IEEE Conference on Games (CoG)*, pages 1–4. IEEE, 2023.

[DeVries and Taylor, 2018] Terrance DeVries and Graham W Taylor. Learning confidence for out-of-distribution detection in neural networks. *arXiv preprint arXiv:1802.04865*, 2018.

[Doosti and Guzdial, 2023] Anahita Doosti and Matthew Guzdial. Transfer learning for underrepresented music generation. *arXiv preprint arXiv:2306.00281*, 2023.

[Fauconnier and Turner, 2003] Gilles Fauconnier and Mark Turner. Conceptual blending, form and meaning. *Recherches en communication*, 19:57–86, 2003.

[Gallotta *et al.*, 2024] Roberto Gallotta, Graham Todd, Marvin Zammit, Sam Earle, Antonios Liapis, Julian Togelius, and Georgios N Yannakakis. Large language models and games: A survey and roadmap. *arXiv preprint arXiv:2402.18659*, 2024.

[Goguen and Harrell, 2004] Joseph A Goguen and D. Fox Harrell. Style as a choice of blending principles. In *AAAI Technical Report*, volume 7, pages 49–56, 2004.

[Guzdial and Riedl, 2016] Matthew Guzdial and Mark Riedl. Learning to blend computer game levels. *arXiv preprint arXiv:1603.02738*, 2016.

[Guzdial and Riedl, 2018a] Matthew Guzdial and Mark Riedl. Automated game design via conceptual expansion. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 14, pages 31–37, 2018.

[Guzdial and Riedl, 2018b] Matthew Guzdial and Mark O Riedl. Combinets: Creativity via recombination of neural networks. *arXiv preprint arXiv:1802.03605*, 2018.

[Guzdial and Riedl, 2018c] Matthew J Guzdial and Mark O Riedl. Combinatorial creativity for procedural content generation via machine learning. In *Workshops at the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[Guzdial and Riedl, 2019] Matthew Guzdial and Mark Riedl. An interaction framework for studying co-creative ai. *arXiv preprint arXiv:1903.09709*, 2019.

[Guzdial and Riedl, 2021] Matthew Guzdial and Mark O Riedl. Conceptual game expansion. *IEEE Transactions on Games*, 14(1):93–106, 2021.

[Guzdial *et al.*, 2017] Matthew Guzdial, Boyang Li, and Mark O Riedl. Game engine learning from video. In *IJCAI*, pages 3707–3713, 2017.

[Guzdial *et al.*, 2018] Matthew Guzdial, Nicholas Liao, and Mark Riedl. Co-creative level design via machine learning. *arXiv preprint arXiv:1809.09420*, 2018.

[Guzdial *et al.*, 2019] Matthew Guzdial, Nicholas Liao, Jonathan Chen, Shao-Yu Chen, Shukan Shah, Vishwa Shah, Joshua Reno, Gillian Smith, and Mark O Riedl. Friend, collaborator, student, manager: How design of an ai-driven game level editor affects creators. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–13, 2019.

[Guzdial *et al.*, 2021] Matthew Guzdial, Nathan Sturtevant, and Carolyn Yang. The impact of visualizing design gradients for human designers. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 17, pages 18–25, 2021.

[Guzdial *et al.*, 2022a] Matthew Guzdial, Sam Snodgrass, and Adam J Summerville. Mixed-initiative pcgml. In *Procedural Content Generation via Machine Learning: An Overview*, pages 181–199. Springer, 2022.

[Guzdial *et al.*, 2022b] Matthew Guzdial, Sam Snodgrass, and Adam J Summerville. *Procedural content generation via machine learning: An overview*. Springer, 2022.

[Halina and Guzdial, 2021a] Emily Halina and Matthew Guzdial. A demonstration of kiaitime: A mixed-initiative pcgml rhythm game editor. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 17, pages 240–242, 2021.

[Halina and Guzdial, 2021b] Emily Halina and Matthew Guzdial. Taikonation: Patterning-focused chart generation for rhythm action games. In *Proceedings of the 16th International Conference on the Foundations of Digital Games*, pages 1–10, 2021.

[Halina and Guzdial, 2022] Emily Halina and Matthew Guzdial. Threshold designer adaptation: Improved adaptation for designers in co-creative systems. In Lud De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 4943–4949. International Joint Conferences on Artificial Intelligence Organization, 7 2022. AI and Arts.

[Halina and Guzdial, 2023] Emily Halina and Matthew Guzdial. Tree-based reconstructive partitioning: a novel low-data level generation approach. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 19, pages 244–254, 2023.

[Ji *et al.*, 2023] Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Jiayi Zhou, Zhaowei Zhang, et al. Ai alignment: A comprehensive survey. *arXiv preprint arXiv:2310.19852*, 2023.

[Khalifa, 2020] Ahmed Khalifa. *General Level Generation*. PhD thesis, New York University Tandon School of Engineering, 2020.

[Liu *et al.*, 2021] Jialin Liu, Sam Snodgrass, Ahmed Khalifa, Sebastian Risi, Georgios N Yannakakis, and Julian Togelius. Deep learning for procedural content generation. *Neural Computing and Applications*, 33(1):19–37, 2021.

[Liu, 2023] Vivian Liu. Beyond text-to-image: Multimodal prompts to explore generative ai. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–6, 2023.

[Loftsdottir and Guzdial, 2022] Dagmar Loftsdottir and Matthew Guzdial. Sketchbetween: Video-to-video synthesis for sprite animation via sketches. In *Proceedings of the 17th International Conference on the Foundations of Digital Games*, pages 1–7, 2022.

[Lopes and Bidarra, 2011] Ricardo Lopes and Rafael Bidarra. Adaptivity challenges in games and simulations: a survey. *IEEE Transactions on Computational Intelligence and AI in Games*, 3(2):85–99, 2011.

[Lotfollahi *et al.*, 2020] Mohammad Lotfollahi, Mohsen Naghipourfar, Fabian J Theis, and F Alexander Wolf. Conditional out-of-distribution generation for unpaired data using transfer vae. *Bioinformatics*, 36(Supplement_2):i610–i617, 2020.

[Löwenström, 2024] Paul Löwenström. Guiding generation of 2d pixel art characters using text-image similarity models: A comparative study of generating 2d pixel art characters using pixeldraw and diffusion model guided by text-image similarity models, 2024.

[Mahajan and Guzdial, 2022] Anmol Mahajan and Matthew Guzdial. Modeling individual humans via a secondary task transfer learning method. *Federated and Transfer Learning*, pages 259–281, 2022.

[Mahmoudi-Nejad *et al.*, 2021] Athar Mahmoudi-Nejad, Matthew Guzdial, and Pierre Boulanger. Arachnophobia exposure therapy using experience-driven procedural content generation via reinforcement learning (edpcgrl). In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 17, 2021.

[Parisi *et al.*, 2019] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural networks*, 113:54–71, 2019.

[Rabii and Cook, 2023] Younès Rabii and Michael Cook. Why oatmeal is cheap: Kolmogorov complexity and procedural generation. In *Proceedings of the 18th International Conference on the Foundations of Digital Games*, pages 1–7, 2023.

[Saravanan and Guzdial, 2022] Akash Saravanan and Matthew Guzdial. Pixel vq-vaes for improved pixel art representation. In *Experimental AI in Games (EXAG), AIIDE 2022 Workshop*, 2022.

[Sarkar *et al.*, 2023] Anurag Sarkar, Matthew Guzdial, Sam Snodgrass, Adam Summerville, Tiago Machado, and Gillian Smith. Procedural content generation via knowledge transformation (pcg-kt). *IEEE Transactions on Games*, 2023.

[Shaker *et al.*, 2016a] Noor Shaker, Julian Togelius, and Mark J Nelson. Procedural content generation in games. 2016.

[Shaker *et al.*, 2016b] Noor Shaker, Julian Togelius, Mark J Nelson, Antonios Liapis, Gillian Smith, and Noor Shaker. Mixed-initiative content creation. *Procedural content generation in games*, pages 195–214, 2016.

[Smith, 2017] Gillian Smith. Procedural content generation: An overview. *Level Design*, pages 159–183, 2017.

[Sturtevant *et al.*, 2020] Nathan Sturtevant, Nicolas Decroocq, Aaron Tripodi, and Matthew Guzdial. The unexpected consequence of incremental design changes. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 16, pages 130–136, 2020.

[Summerville, 2018] Adam J Summerville. *Learning from Games for Generative Purposes*. University of California, Santa Cruz, 2018.

[Thomaz and Breazeal, 2008] Andrea L Thomaz and Cynthia Breazeal. eachable robots: Understanding human teaching behavior to build more effective robot learners. In *Artificial Intelligence*, volume 172, pages 716–737, 2008.

[Togelius *et al.*, 2011] Julian Togelius, Georgios N Yannakakis, Kenneth O Stanley, and Cameron Browne. Search-based procedural content generation: A taxonomy and survey. *IEEE Transactions on Computational Intelligence and AI in Games*, 3(3):172–186, 2011.

[Trajkova *et al.*, 2024] Milka Trajkova, Duri Long, Manoj Deshpande, Andrea Knowlton, and Brian Magerko. Exploring collaborative movement improvisation towards the design of luminai—a co-creative ai dance partner. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–22, 2024.

[Vimpari *et al.*, 2023] Veera Vimpari, Annakaisa Kultima, Perttu Hämäläinen, and Christian Guckelsberger. "an adapt-or-die type of situation": Perception, adoption, and use of text-to-image-generation ai by game industry professionals. *Proceedings of the ACM on Human-Computer Interaction*, 7(CHI PLAY):131–164, 2023.

[Waltz and Buchanan, 2009] David Waltz and Bruce G Buchanan. Automated hypothesis generation based on mining scientific literature. *Science*, 324(5923):43–44, 2009.

[Wisniewski and Gentner, 1991] Edward J Wisniewski and Dedre Gentner. On the combinatorial semantics of noun pairs: Minor and major adjustments to meaning. In *Advances in psychology*, volume 77, pages 241–284. Elsevier, 1991.

[Yu *et al.*, 2021] Kristen Yu, Matthew Guzdial, Nathan R Sturtevant, Morgan Cselinacz, Chris Corfe, Izzy Hubert Lyall, and Chris Smith. Adventures of ai directors early in the development of nightingale. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 18, pages 70–77, 2021.

[Zhang *et al.*, 2021] Xingxuan Zhang, Peng Cui, Renzhe Xu, Linjun Zhou, Yue He, and Zheyan Shen. Deep stable learning for out-of-distribution generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5372–5382, 2021.