Game Design using Creative AI

Anurag Sarkar
Northeastern University
sarkar.an@husky.neu.edu

Abstract

Creative AI refers to the application of AI techniques, primarily machine learning, for performing creative tasks such as generating art and music. In this paper, we present exploratory work that demonstrates and argues for leveraging such creative AI techniques for game design. We train variational autoencoders (VAEs) on levels from the games *Super Mario Bros.* and *Kid Icarus* and show how the affordances of the learned models can help inform co-creative game and level design, similar to applications of creative AI and machine learning in visual art and music.

1 Introduction and Background

Machine learning approaches for creative tasks like image and music generation and style transfer have become increasingly popular. Often referred to using the umbrella term *Creative AI* [18], such methods exist at the confluence of computational creativity and artificial intelligence, and have found success in informing several co-creative processes and applications. However, most such uses have been restricted to the domains of visual art and music, and have not been widely adopted for game design. An emerging subfield in games research is Procedural Content Generation via Machine Learning or PCGML [31] which refers to the automated production of content for games using generative models trained on existing game data. Such techniques have found success in controllable level generation for games such as *Super Mario Bros.* [8, 23, 24, 28] and *The Legend of Zelda* [27, 29] but apart from a few examples [11–13], such approaches have focused on automated level generation rather than on how the affordances of such systems could assist designers in co-creative settings, akin to creative AI techniques for visual art and music. We thus propose combining existing creative AI approaches with recent advances in PCGML to inform co-creative game and level design systems of the future. As simple exemplars of what is possible, we demonstrate exploratory results of repurposing simple techniques applied in the visual arts, namely, reverse image search, shortest path between images and image interpolation, for the game design domain. We present results using VAEs trained on level data from the games, *Super Mario Bros.* and *Kid Icarus*, and conclude by discussing future directions for such work.

2 Approach

We gathered level data from the Video Game Level Corpus [30], widely used in games research, and used 1 level each from *Super Mario Bros.* (SMB) and *Kid Icarus* (KI). Both are platformer games developed by Nintendo with SMB levels progressing horizontally and KI levels progressing vertically. Since our goal is to demonstrate feasibility, we used only 1 level per game for fast training. While GANs [6] have traditionally been the de-facto generative model for creative AI, VAEs [15] may be better suited for co-creative level design as they allow the designer to supply known inputs as opposed to GANs which only accept latent vectors as inputs to generation. To create samples for training, we slid a 16x16 window horizontally across the SMB level and vertically across the KI level, obtaining 187 SMB samples and 191 KI samples. Thus, for each game, a VAE was trained to generate 16x16 level segments of that game. Both VAEs consisted of an encoder with 2 strided convolutional layers using batchnorm and LeakyReLU activation, a 32-dimensional bottleneck layer and a decoder with 1

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

non-strided followed by 2 strided convolutional layers with batchnorm and ReLU activation. Both models were trained for 5000 epochs and implemented using PyTorch [17]. To demonstrate our approach, we present results of performing level design equivalents of three simple co-creative image generation techniques adapted from Gene Kogan's ML4A course [16], namely reverse level search, shortest paths between levels and level interpolation, based off of reverse image search, shortest path between images and image interpolation respectively.

2.1 Reverse Level Search

This can enable designers to search for levels based on an input level and some defined objective. We used 4 metrics - cosine distance, non-linearity, density and difficulty - to search for vectors in the latent space of the VAEs. *Non-linearity* measures how closely the topology of a segment follows a straight line. *Density* measures the proportion of solid tiles in a level. Finally, *Difficulty* measures the number of enemies and hazards in a segment. We present results of searching for segments with the closest and furthest values for these metrics compared to a given input segment as recommended in [26]. Search was done within the training segments as well as within segments generated using 10000 random latent vectors. Results are shown in Supplementary Figures 1 and 2.

2.2 Shortest Paths between Levels

Latent spaces allow constructing graphs where nodes represent levels and edges represent different relations between levels based on latent vector operations. Designers could use such graphs to make specific queries for desired levels and segments. As an example operation, we present finding shortest paths between given level segments. For each game, we constructed a graph where each node represented a segment. We then added an edge between a node and each of its k nearest neighbors based on cosine distance between their corresponding latent vectors and then computed the shortest paths between any two given nodes. Shortest paths for training and generated segments are shown in Supplementary Figures 3 and 4 respectively, using k=5 for the former and k=10 for the latter.

2.3 Interpolation

Finally, latent spaces enable the use of interpolation to generate new levels not in the training data. For each game, we selected two training segments, obtained their corresponding latent vectors by feeding them through the VAE's encoder and then did linear interpolation between these vectors to obtain segments in between the input segments. Results are shown in Supplementary Figure 5.

3 Discussion and Future Work

We demonstrated a number of affordances enabled by leveraging creative AI for game design, allowing designers to 1) search for new segments relative to input segments using various metrics, 2) find segments that are similar to each other and 3) discover segments that do not exist in the games used for training. These examples barely scratch the surface in terms of the possibilities enabled by creative AI in the level and game design domain. Models like pix2pix [14] and CycleGAN [32] that have been widely used for creative AI in visual art could be repurposed for game design to enable more complex applications such as game blending [7, 9], game style transfer [25] and feature extraction for levels. To this end, VAEs [22] and LSTMs [21] have been used to blend levels from different games, albeit by ignoring playability. While [10] uses machine learning to create new playable games by recombining existing ones, generating playable blended games may still benefit from and require learning disentangled representations of game data that separate rules and levels, similar to separate representations of style and content in visual art [2, 5]. In addition to blending, this could thus enable feature and style transfer between games. Though popular in visual art, style transfer in games is complicated due to their hierarchical structure (i.e. mechanics, aesthetics, dynamics), similar to that of music (i.e. timbre, rhythm, melody, performance etc.) and thus performing game style transfer may involve adapting the ideas for performing music style transfer put forth in [1]. For example, approaches like hierarchical VAEs for music generation as in [19] may be useful. In general, future work in creative AI for game design could take inspiration from applications of creative AI for music such as PianoGenie [3], GANSynth [4] and MusicVAE [20]. Finally, future advances in game blending and style transfer could enable *game design arithmetic* in the latent space, allowing us to extract vector representations of games and generate novel games that are a result of operations such as *Mario - Zelda + Metroid* for example. Ultimately, our results demonstrate that learned latent level design spaces hold promise for new creative applications for game and level design.

References

[1] Shuqi Dai, Zheng Zhang, and Gus G. Xia. Music Style Transfer: A Position Paper. In *Proceedings of the International Workshop on Musical Metacreation (MUME), Salamanca, Spain*, 2018.

[2] Chris Donahue and Julian McAuley. Disentangled Representations of Style and Content for Visual Art with Generative Adversarial Networks. In *NIPS Workshop on Machine Learning for Creativity and Design*, 2017.

[3] Chris Donahue, Ian Simon, and Sander Dieleman. Piano Genie. In *ACM Conference on Intelligent User Interfaces*, 2019.

[4] Jesse Engel, Kumar Krishna Agarwal, Shuo Chen, Ishaan Gulrajani, Chris Donahue, and Adam Roberts. GANSynth: Adversarial Neural Audio Synthesis. In *International Conference on Learning Representations*, 2019.

[5] Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Image Style Transfer using Convolutional Neural Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

[6] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, 2014.

[7] Jeremy Gow and Joseph Corneli. Towards Generating Novel Games using Conceptual Blending. In *Proceedings of the Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*, 2015. URL http://www.aaai.org/ocs/index.php/AIIDE/AIIDE15/paper/download/11619/11388.

[8] Matthew Guzdial and Mark Riedl. Game Level Generation from Gameplay Videos. In *Proceedings of the Twelfth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2016. URL http://www.aaai.org/ocs/index.php/AIIDE/AIIDE16/paper/download/14008/13593.

[9] Matthew Guzdial and Mark Riedl. Learning to Blend Computer Game Levels. In *Proceedings of the Seventh International Conference on Computational Creativity*, 2016. URL https://arxiv.org/pdf/1603.02738.pdf.

[10] Matthew Guzdial and Mark Riedl. Automated Game Design via Conceptual Expansion. In *Proceedings of the Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2018.

[11] Matthew Guzdial, Jonathan Chen, Shao-Yu Chen, and Mark Riedl. A General Level Design Editor for Co-Creative Level Design. In *Proceedings of the AIIDE Workshop on Experimental AI in Games*, 2017.

[12] Matthew Guzdial, Joshua Reno, Jonathan Chen, Gillian Smith, and Mark Riedl. Explainable PCGML via Game Design Patterns. In *Proceedings of the AIIDE Workshop on Experimental AI in Games*, 2018.

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

[14] Phillip Isola, Jun-Yan. Zhu, Tinghui Zhou, and Alexei Efros. Image-to-Image Translation with Conditional Adversarial Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

[15] Diederik P. Kingma and Max Welling. Auto-encoding Variational Bayes. In *The 2nd International Conference on Learning Representations (ICLR)*, 2013.

[16] Gene Kogan. Machine Learning for Artists, 2019. URL http://ml4a.github.io/.

[17] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic Differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.

[18] Roelof Pieters and Samim Winiger. Creative AI: On the Democratisation and Escalation of Creativity. 2016. URL https://medium.com/@creativeai/creativeai-9d4b2346faf3.

[19] Adam Roberts, Jesse Engel, and Douglas Eck. Hierarchical Variational Autoencoders for Music. In *NIPS Workshop on Machine Learning for Creativity and Design*, 2017.

[20] Adam Roberts, Jesse Engel, Colin Raffel, Curtis Hawthorne, and Douglas Eck. A Hierarchical Latent Vector Model for Learning Long Term Structure in Music. In *International Conference on Machine Learning*, 2018.

[21] Anurag Sarkar and Seth Cooper. Blending Levels from Different Games using LSTMs. In *Proceedings of the AIIDE Workshop on Experimental AI in Games*, 2018.

[22] Anurag Sarkar, Zhihan Yang, and Seth Cooper. Controllable Level Blending between Games using Variational Autoencoders. In *Proceedings of the AIIDE Workshop on Experimental AI in Games*, 2019.

[23] Sam Snodgrass and Santiago Ontañón. Experiments in Map Generation using Markov Chains. In *Proceedings of the 9th International Conference on the Foundations of Digital Games*, 2014.

[24] Sam Snodgrass and Santiago Ontañón. Controllable Procedural Content Generation via Constrained Multi-Dimensional Markov Chain Sampling. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, 2016. URL http://www.ijcai.org/Proceedings/16/Papers/116.pdf.

[25] Sam Snodgrass and Santiago Ontañón. An Approach to Domain Transfer in Procedural Content Generation of Two-Dimensional Videogame Levels. In *Proceedings of the Twelfth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE-17)*, 2017. URL https://www.aaai.org/ocs/index.php/AIIDE/AIIDE16/paper/view/13985/13598.

[26] Adam Summerville. Expanding Expressive Range: Evaluation Methodologies for Procedural Content Generation. In *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2018.

[27] Adam Summerville and Michael Mateas. Sampling Hyrule: Sampling Probabilistic Machine Learning for Level Generation. In *Proceedings of the Eleventh Artificial Intelligence and Interactive Digital Conference*, 2015. URL http://www.aaai.org/ocs/index.php/AIIDE/AIIDE15/paper/download/11570/11395.

[28] Adam Summerville and Michael Mateas. Super Mario as a String: Platformer Level Generation via LSTMs. In *Proceedings of the 1st International Joint Conference on DiGRA and FDG*, 2016. URL https://arxiv.org/pdf/1603.00930.pdf.

[29] Adam Summerville, Morteza Behrooz, Michael Mateas, and Arnav Jhala. The Learning of Zelda: Data-Driven Learning of Level Topology. In *Proceedings of the 10th International Conference on the Foundations of Digital Games*, 2015. URL https://pdfs.semanticscholar.org/5c31/cc6d908b261dbc794d55a3bdae9a216522b4.pdf.

[30] Adam Summerville, Sam Snodgrass, Michael Mateas, and Santiago Ontañón. The VGLC: The Video Game Level Corpus. *Proceedings of the 7th Workshop on Procedural Content Generation*, 2016.

[31] Adam Summerville, Sam Snodgrass, Matthew Guzdial, Christoffer Holmgård, Amy Hoover, Aaron Isaksen, Andy Nealen, and Julian Togelius. Procedural Content Generation via Machine Learning (PCGML). In *Foundations of Digital Games Conference*, 2017. URL https://arxiv.org/pdf/1702.00539.pdf.

[32] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei Efros. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.

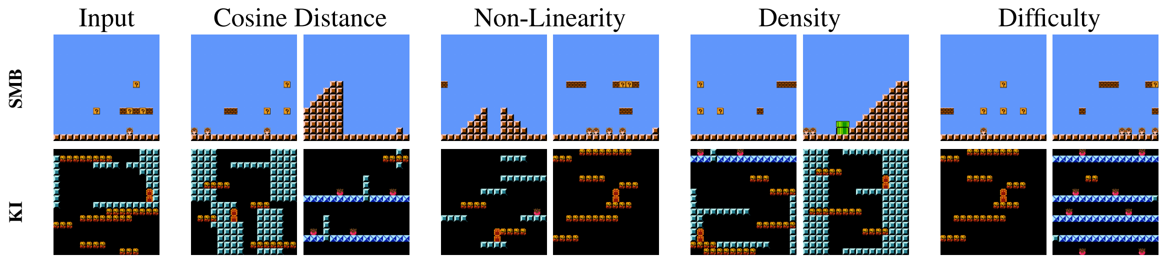

Figure 1: Reverse level search on training segments using the given input segment. Pairs consist of the closest match on the left and farthest match on the right based on the corresponding metric.

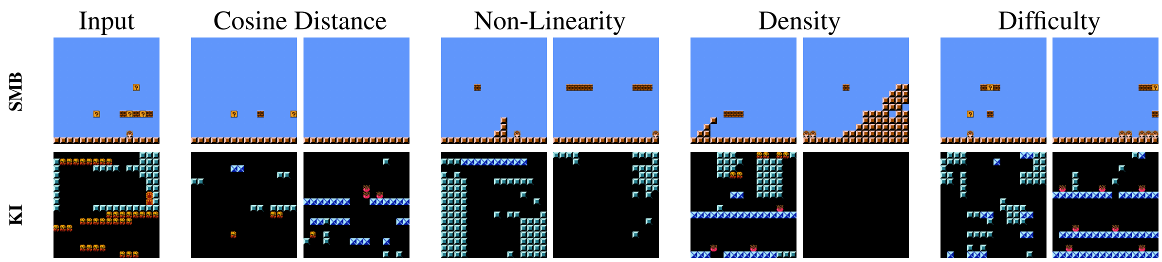

Figure 2: Reverse level search on segments generated randomly. Pairs are as described for Figure 1.

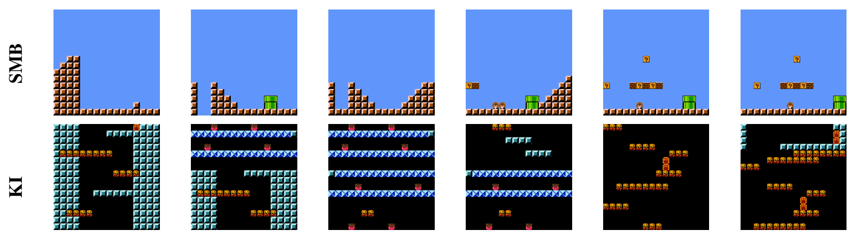

Figure 3: Shortest path between given training segments. The segments in the middle form the shortest path connecting the first and last segment.

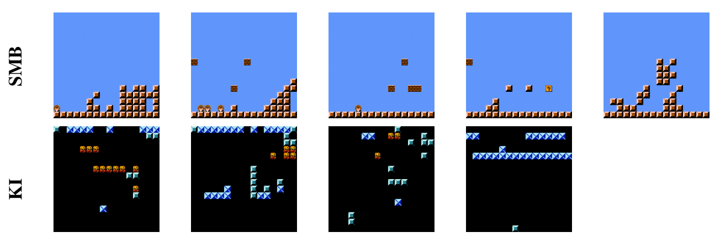

Figure 4: Shortest path between given randomly generated segments on either end.

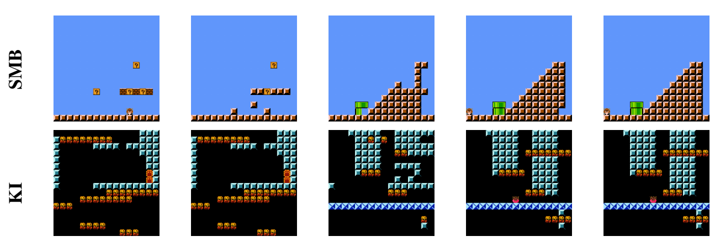

Figure 5: Linear interpolation between the first and last segments generate the segments in the middle that do not exist in the actual game.