Proceedings of the Sixteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE-20)

Artificial Intelligence as an Art Director

Adrian Gonzalez
Department of Computer Science
Center for Research and Advanced Studies of the National Polytechnic Institute
(Cinvestav IPN) Unidad Guadalajara, México
Zapopan, Jalisco 45019
adrian.glez@cinvestav.mx

Abstract

My research focuses on the design of an artificial intelligence-based system that can perform the tasks of a videogame art director. I began with the analysis of game developers' design processes, workflows, and methodologies. I observed they involve fast-paced proposal-evaluation-correction cycles. This, along with my interest in generative methods, led me to consider using Variational Autoencoders (VAEs), which also include similar cycles, and do not require a high number of samples to produce satisfactory results. Currently, I am using VAEs on experiments over Pokémon images, with positive results.

Introduction

In the context of entertainment, creating unique, compelling, and high-quality assets is expensive, time-consuming, and requires ever-increasing knowledge and skills from diverse areas of expertise. Despite that, audience's expectations in terms of asset quality keep growing. To satisfy such demands, big companies often hire large teams of experts; in contrast, smaller developers tend to sacrifice some of the desired asset properties mentioned above: a risky move in a highly competitive market.

Recent progress in content generation methods and techniques has enabled other alternatives to satisfy those demands (Procedural Content Generation, Machine Learning, Deep Learning, Reinforcement Learning, etc.) (Shaker, Togelius, and Nelson 2016; Khalifa et al. 2020; Summerville et al. 2017; Gravina et al. 2019; Kingma and Welling 2013; Karras, Laine, and Aila 2018). These techniques allow for the analysis and creation of content (visuals, audio, levels, and even games) (Rebouças Serpa and Formico Rodrigues 2019; Torrado et al. 2019; Guzdial and Riedl 2018; Hoover et al. 2015; Cook, Colton, and Gow 2017) of considerable quality quickly. Titles such as The Division 2 (Ubisoft 2019) and the Borderlands series (Gearbox-Software 2020) employ some of these approaches. However, their inclusion in the games industry has not been widespread. Also, some types of content, like levels, have had more impact than others, such as visuals, in which I am interested. I propose that a co-creative system specifically designed to fit into a game's early visual development stages can encourage studios to embrace these methods.

Related Work

Artificial intelligence (ML, PCG, etc.) methods can be applied to assist during the development process (Liapis, Yannakakis, and Togelius 2013; Guzdial et al. 2017; Valenzuela, Matamala, and Germanidis 2018), improve workflow, and reduce the technical knowledge required for a task. Commonly used functions include evaluating or suggesting changes to enhance the work's quality (Gravina et al. 2019; Davis et al. 2018). Common drawbacks include lack of control (or ignoring the artist's desired style or intention), confusing parameters, and annoying artifacts, all of which can limit their utility (Rebouças Serpa and Formico Rodrigues 2019). A notable and relevant visual content generator is Artbreeder (Simon 2020). It is an online, GAN-based, image-generation tool that allows users to control the generation process via several pre-established parameters; however, some of them, such as *Ninja* or *Suit*, are often unclear or unpredictable.

Current Work

This work's main research topic is: Can a program perform the tasks of a videogame art director? There exist separate approaches that have accomplished some of these tasks, which could be adapted and merged to create an AI art director. My main objective is to design a system capable of evaluating and generating visual content. The system will provide feedback and inspiration to the artist, according to user-established guidelines or frames (Liapis et al. 2019) that represent a desired art direction or style. Its feedback will possibly be in the form of editable visual suggestions or natural language explanations and directions. My proposed sub-objectives are as follows:

1. Analyze the human creative process to evaluate and propose approaches and models that can function similarly. *Achieved*.
2. Determine the development stages in which my proposal could be most beneficial (considering their com-

mon drawbacks). For this, I studied artists' and other creatives' workflows and processes. *Achieved*.

3. Establish clearly the activities an art director carries out, to break them down into steps that can be more easily represented in a computational model. *In progress*.
4. Propose the system's architecture and workflow for each art direction task, as well as their interactions. *Incomplete*.
5. Implement programs that perform the steps of each art direction task. *Incomplete*.
6. Design and perform experiments and case studies for each task. *Incomplete*.

By studying cognitive models of the design process (Mekern, Hommel, and Sjoerds 2019; Augello et al. 2016; Gonçalves, Cardoso, and Badke-Schaub 2014) I noticed the presence of a proposal-evaluation-correction cycle. On the other hand, various methodologies adopted by game developers, especially during early stages (sketches, concept art, etc.), use short and numerous iterations for proposal and correction, even among different game facets (Liapis et al. 2019). I argue that, during those stages, a co-creative tool can provide the artists with useful feedback or inspiration, regardless of the drawbacks mentioned above.

I decided to study this problem from the perspective of co-creative tools (Guzdial, Liao, and Riedl 2018). Currently, I am using Variational Autoencoders (VAEs) (Kingma and Welling 2013) and Generative Adversarial Networks (GANs) (Goodfellow et al. 2014; Karras, Laine, and Aila 2018). I selected them since their training includes proposal-evaluation-correction steps much like the human's design process, and, once trained, they produce content quickly. However, both VAEs and GANs require a large number of training samples, which can be difficult or expensive to procure, especially for small game creators. Since VAEs require less samples, I began working with them to evaluate their *controllability*, which is essential for co-creativity.

In the games industry, the differences between artists and producers are blurry (Riedl and Zook 2013), specially for art directors, who commonly perform management tasks. Nonetheless, I chose to focus only on the director's design activities to prioritize the generation of visual content. The art direction tasks identified so far are: shape language, silhouettes, color theory and composition, lighting, research and inspiration (*moodboard* that can be updated), and style. Each of these tasks could be carried out and evaluated by smaller, more specialized generators; however, as stated in (Liapis et al. 2019), the integration of their results must be coherent according to the user-provided general description or parameters. One of this work's main challenges will be defining such coherence evaluation methods. Other crucial consideration for my proposal is that art direction is about "providing directions, not answers" (Rogic 2017); thus, I will favor generating meaningful sketches over finalized assets.

I came up with an experiment in which the generative system will provide suggestions to the user. The system is trained over a set of well-known characters that have clear

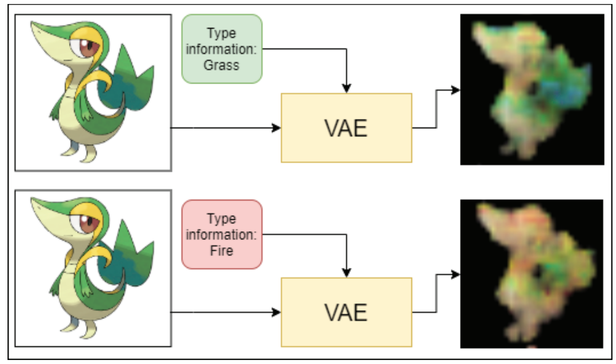

Figure 1: Pokémon type-swap example. Given the Pokémon's image and type information on the left, my system generates the image on the right. The first row keeps its original type (*grass*). The second row changes the type to *fire*, causing the decoder to alter the output.

design intentions (to convey certain gameplay-relevant elements through visual cues). Then, when the gameplay-relevant element is changed, the system will provide the artist with potential adjustments for the characters to reflect such change. Specifically, I took Pokémon (Nintendo 2019) images that contain visual cues to convey their type(s), such as *fire* or *grass*, and created a program that outputs an image of the same Pokémon, but altered to express a given target type. I trained a VAE with the images and added the type information before the embedding layer as a modifier. So far, there are positive results, but the extent of the changes and the level of detail can be improved, as shown in Figure 1. I am currently using transfer learning to try to enhance the results' quality. This experiment's goal is similar to the one from (Liapis 2018), but they used decision trees for classification, and evolutionary methods for palette generation.

Regarding the low number of game character-focused datasets, I am developing an avatar creation application, very similar to the avatar customization interface *Heroes of Elibca* presented in (Lim, Liapis, and Harrell 2016). It records each generated character and its attributes (strength, wisdom, etc.) along with some demographic data about its creator. My objective with this is to publicly provide a collection of human-designed characters for research. This application lacks some features like choosing color variations but compensates for it by reducing the time and information required of the user.

Future Work

This work's goal is creating a system that will help developers to create art quickly, without extensive theoretical or practical knowledge. I acknowledge that, due to time constraints, I will only be able to tackle one or two art direction tasks, therefore, I consider that modeling the *shape language* and *silhouettes* art direction tasks should pose a challenging but feasible next step in my research. Finally, for each task developed, I would like to perform user studies on with artists of different experience levels, to gain insight on how my proposal impacts their workflow.

References

Augello, A.; Infantino, I.; Lieto, A.; Pilato, G.; Rizzo, R.; and Vella, F. 2016. Artwork creation by a cognitive architecture integrating computational creativity and dual process approaches. *CoRR* abs/1601.00669.

Cook, M.; Colton, S.; and Gow, J. 2017. The angelina videogame design system—part i. *IEEE Transactions on Computational Intelligence and AI in Games* 9(2):192–203.

Davis, N.; Maher, M. L.; Grace, K.; and Eltayeby, O. 2018. Surprise walks: Encouraging users towards novel concepts with sequential suggestions. In *ICCC*.

Gearbox-Software. 2020. Borderlands series. https://borderlands.com.

Gonçalves, M.; Cardoso, C.; and Badke-Schaub, P. 2014. What inspires designers? preferences on inspirational approaches during idea generation. *Design Studies* 35(1):29 – 53.

Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'14, 2672–2680. Cambridge, MA, USA: MIT Press.

Gravina, D.; Khalifa, A.; Liapis, A.; Togelius, J.; and Yannakakis, G. N. 2019. Procedural content generation through quality diversity. *CoRR* abs/1907.04053.

Guzdial, M., and Riedl, M. 2018. Automated game design via conceptual expansion. *CoRR* abs/1809.02232.

Guzdial, M.; Long, D.; Cassion, C.; and Das, A. 2017. Visual procedural content generation with an artificial abstract artist. In *Proceedings of ICCC Computational Creativity and Games Workshop*.

Guzdial, M.; Liao, N.; and Riedl, M. 2018. Co-creative level design via machine learning. *CoRR* abs/1809.09420.

Hoover, A. K.; Cachia, W.; Liapis, A.; and Yannakakis, G. N. 2015. Audioinspace: Exploring the creative fusion of generative audio, visuals and gameplay. In *EvoMUSART*.

Karras, T.; Laine, S.; and Aila, T. 2018. A style-based generator architecture for generative adversarial networks. *CoRR* abs/1812.04948.

Khalifa, A.; Bontrager, P.; Earle, S.; and Togelius, J. 2020. Pcgrl: Procedural content generation via reinforcement learning.

Kingma, D. P., and Welling, M. 2013. Auto-encoding variational bayes.

Liapis, A.; Yannakakis, G. N.; Nelson, M. J.; Preuss, M.; and Bidarra, R. 2019. Orchestrating game generation. *IEEE Transactions on Games* 11:48–68.

Liapis, A.; Yannakakis, G. N.; and Togelius, J. 2013. Sentient sketchbook: Computer-aided game level authoring. In *FDG*.

Liapis, A. 2018. Recomposing the pokémon color palette. In *EvoApplications*.

Lim, C.-U.; Liapis, A.; and Harrell, F. D. 2016. Discovering social and aesthetic categories of avatars: A bottom-up artificial intelligence approach using image clustering. In *DiGRA/FDG #3916 - Proceedings of the First International Joint Conference of DiGRA and FDG*. Dundee, Scotland: Digital Games Research Association and Society for the Advancement of the Science of Digital Games.

Mekern, V.; Hommel, B.; and Sjoerds, Z. 2019. Computational models of creativity: a review of single-process and multi-process recent approaches to demystify creative cognition. *Current Opinion in Behavioral Sciences* 27:47–54.

Nintendo. 2019. Pokemon series. https://www.pokemon.com.

Rebouças Serpa, Y., and Formico Rodrigues, M. A. 2019. Towards machine-learning assisted asset generation for games: A study on pixel art sprite sheets. In *2019 18th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames)*, 182–191.

Riedl, M. O., and Zook, A. 2013. Ai for game production. In *2013 IEEE Conference on Computational Intelligence in Games (CIG)*, 1–8.

Rogic, D. 2017. Art direction bootcamp: Pantalones! things i wished i knew prior to becoming an art director. https://youtu.be/L4-KoapvUcs.

Shaker, N.; Togelius, J.; and Nelson, M. J. 2016. *Procedural Content Generation in Games*. Springer Publishing Company, Incorporated, 1st edition.

Simon, J. 2020. Artbreeder. https://artbreeder.com.

Summerville, A.; Snodgrass, S.; Guzdial, M.; Holmgård, C.; Hoover, A. K.; Isaksen, A.; Nealen, A.; and Togelius, J. 2017. Procedural content generation via machine learning (pcgml).

Torrado, R. R.; Khalifa, A.; Green, M. C.; Justesen, N.; Risi, S.; and Togelius, J. 2019. Bootstrapping conditional gans for video game level generation.

Ubisoft. 2019. Tom clancy's the division 2. https://www.ubisoft.com/en-us/game/the-division/the-division-2.

Valenzuela, C.; Matamala, A.; and Germanidis, A. 2018. Runway: Adding artificial intelligence capabilities to design and creative platforms.