

Generating Gameplay-Relevant Art Assets with Transfer Learning

Adrian Gonzalez,¹ Matthew Guzdial,² Felix Ramos¹

¹Department of Computer Science, Cinvestav IPN, Unidad Guadalajara, México

²Computing Science Department, University of Alberta

adrian.glez@cinvestav.mx, guzdial@ualberta.ca, felix.ramos@cinvestav.mx

Abstract

In game development, designing compelling visual assets that convey gameplay-relevant features requires time and experience. Recent image generation methods that create high-quality content could reduce development costs, but these approaches do not consider game mechanics. We propose a Convolutional Variational Autoencoder (CVAE) system to modify and generate new game visuals based on their gameplay relevance. We test this approach with Pokémon sprites and Pokémon type information, since types are one of the game's core mechanics and they directly impact the game's visuals. Our experimental results indicate that adopting a transfer learning approach can help to improve visual quality and stability over unseen data.

Introduction

Game development is a demanding task. Gameplay systems generally include numerous elements to make them stand out from similar titles, as well as to provide variety and balance. On the other hand, designing compelling visual assets that quickly and consistently convey those gameplay-relevant features (play-style, difficulty, weaknesses, etc.) is not trivial, especially while striving to preserve project-wide artistic cohesion. This is also an important consideration when creating variations on existing content, such as characters' alternative appearances or *skins*, enemy sub-classes (e.g., Mario's Dry Bones are visual and mechanical variations of the Koopas), or player customization systems. Most of these processes are iterative and time demanding, further increasing development costs (Rebouças Serpa and Formico Rodrigues 2019).

Automating the visual design process could help to improve asset quality and reduce development time. Recent general-purpose deep-learning models for image creation provide high-quality results; however, these approaches are limited to particular tasks with large training sets, such as face, character, or landscape generation (Isola et al. 2016; Karras, Laine, and Aila 2018; Simon 2020). Outside of

fully autonomous generation, some approaches identify latent vectors to grant users the ability to explore the possibilities of a model's learned latent space (Burgess et al. 2018; Voynov and Babenko 2020). We identify two main issues with both of these approaches: they require large amounts of data and their controllability is not influenced by gameplay-relevant aspects like mechanics.

To study how we can generate images that relay gameplay-relevant information, we decided to work with images from the Pokémon series (Nintendo 2019). Pokémon games have clearly defined gameplay elements that are present in their art style (the *type information*), which help to communicate each Pokémon's strengths and weaknesses to players (Liapis 2018). The main Pokémon titles are turn-based role-playing videogames in which players make their companions –Pokémon– battle. Understanding each type's weaknesses and resistances is crucial for victory.

In this article, we present a Variational Autoencoder (VAE) for Pokémon type swapping, which modifies the inputs' visual designs to transmit user-defined types in a controllable manner. To combat the problem of low training data, we adopted a transfer learning approach. Our system's intended behavior would change a specified Pokémon's appearance to convey a given target type (such as *fire* or *water*), according to the visual attributes commonly exhibited by Pokémon of that target type. For instance, if the yellow-colored electric-type Pokémon *Pikachu* were changed to fire-type, its colors would shift towards red.

Related Work

In this section, we briefly discuss computational approaches to automated or assisted visual design generation. We also present an overview of mainstream generative models and their applications to the production of visual game assets, and how those works relate to our proposed approach.

Procedural Content Generation

Procedural content generation or PCG refers to the creation of game content using algorithms with limited or indirect user input (Shaker, Togelius, and Nelson 2016). More closely related to our proposal's objective, there is Visual PCG (Guzdial et al. 2017), which involves the generation

of visual components for games, and PCG via machine learning (PCGML) (Summerville et al. 2018). We present some PCG-based works that create visual game elements and how a particular game’s mechanics affect their creation processes.

Pollite (Guzdial et al. 2017) is an artificial abstract artist based on a convolutional neural network (CNN) that learns to associate features, like shapes and colors, to emotions, from tagged real-world pictures. It can create and modify images to express feelings such as anger or joy. Both their work and our proposal involve concepts to alter visuals (emotions and type information, respectively). However, adopting an approach similar to Pollite’s would require tagging real-world scenes with gameplay elements that, in many cases, would demand manually annotating them or creating a system to do it instead, thus reducing the expected development benefits.

The evolutionary algorithm developed by (Liapis 2018) modifies Pokémon sprites’ colors based on type information. It uses the associations given between color palettes and the different Pokémon types, e.g., *fire* type is related to red tones. Although their approach and ours aim to assist artists in their tasks, their system evolves a sprite’s color palette and then assigns them a type, in contrast, we propose to define a type (or types) and then change the Pokémon’s colors and shape (and even textures) to fit the new type information. This increases our model’s expressiveness, since it is not limited to palette swaps, and allows its users to make specific requests, such as a *fire*-type *Pikachu*.

Deep Generative Models

In this subsection, we mention works based on two deep-learning architectures applied to image generation: Variational Autoencoders (VAEs) (Kingma and Welling 2013) and Generative Adversarial Networks (GANs) (Goodfellow et al. 2014). A VAE consists of two networks: first, an encoder that generates a mean and a variance of a Gaussian distribution per latent space dimension, then the inputs’ latent representations are obtained by sampling from such distributions, and second, a decoder that reconstructs those representations back to the input data space (Kingma and Welling 2013; Larsen, Sønderby, and Winther 2015; Pihlgren, Sandin, and Liwicki 2020).

Works that use VAEs for visual design in games are uncommon. Nonetheless, VAEs have been employed for image and texture synthesis (Chandra et al. 2017; Kingma and Welling 2019), and level generation (Guzdial et al. 2018; Snodgrass and Sarkar 2020). We decided to use a VAE since points sampled from the latent space near a known input tend to resemble it, which is useful to create variations of existing content. However, as stated in (Kingma and Welling 2019), generative VAEs are known to produce blurry results. Therefore, we consider exploring GANs as future work.

In (Rebouças Serpa and Formico Rodrigues 2019), the authors proposed a deep-learning asset generation system for pixel art sprites for a 2D fighting game using line art sketches. Their tool, which is built upon the *pix2pix* architecture (Isola et al. 2016), produces semi-final sprites that must be fine-tuned by a human artist, therefore reducing the

production time for each image. Unlike our proposal, they do not include gameplay-related information in their model.

Artbreeder (Simon 2020), a tool based on a deep convolutional GAN (DCGAN), allows its users to manipulate numerous parameters to adjust the creation of images such as human faces, landscapes, and characters. Another relevant DCGAN model is presented in (Horsley and Liebana 2017), which generates sprites of faces, characters, and creatures using small amounts of training data. The system developed by (Jin et al. 2017) permits its users to create anime faces using a GAN architecture, by providing parameters, such as hair and eye color and style, to control the image generation process. However, none of these models explicitly consider gameplay-specific features to control the generation process.

A notable related work is *pokemon2pokemon* (Wong 2019), which uses CycleGAN (Zhu et al. 2017) to modify the color of Pokémon images given a type and shows positive results. However, it does not modify the Pokémon’s shape, which might be due to CycleGAN’s difficulties when handling geometric changes.

Pokémon

The main Pokémon titles are turn-based role-playing games in which humans command creatures named Pokémon during one-on-one or two-on-two battles. Types are a core mechanic and there exist 18 types: *Bug*, *Dark*, *Dragon*, *Electric*, *Fairy*, *Fighting*, *Fire*, *Flying*, *Ghost*, *Grass*, *Ground*, *Ice*, *Normal*, *Poison*, *Psychic*, *Rock*, *Steel*, and *Water*. Each type possesses weaknesses and resistances to *attacks* from other types. Every Pokémon has one or two types and four attacks (each with its own type). An attack’s damage depends on the attacked Pokémon’s types weaknesses and resistances. Pokémon who use an attack that matches their type gain increased effect. This makes understanding types crucial to win. For instance, fire-type Pokémon are weak against water-type attacks, but resistant to grass-type ones. Therefore, conveying the types of a Pokémon through its design is crucial, especially in the earlier games of the series in which players were not shown the types of a Pokémon unless they owned it.

System Overview

Our objective is a system that allows its users to create variations on existing Pokémon. The users select a Pokémon design and one or two types, then the image is modified to make it show distinctive features of the given types.

To achieve this, we employ a convolutional VAE. Our process to train our VAE is as follows. First, we collect a set of Pokémon images and their type information for training. Second, given the lack of data, we use the Anime Face Dataset (AFD for short) (Churchill 2019) as a source dataset for a transfer learning approach. Given our final goal of controllability through Pokémon type information, we assign type labels to the AFD’s images based on how similar they are to each type’s Pokémon designs. Then, we train the VAE on the now-labeled dataset. Finally, we transfer the learned weights from the anime samples and fine-tune them via training on the Pokémon images.

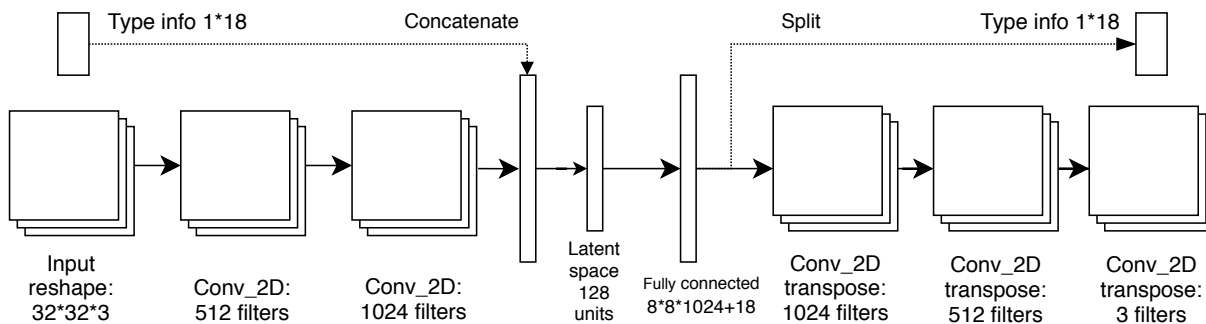


Figure 1: Architecture of the proposed convolutional variational autoencoder (CVAE). Type information is represented as a one-hot-encoded vector of size 18. The type information and the last convolution’s result are passed to the latent space’s fully connected layers.

Dataset Collection

The Pokémon images and type information were retrieved from (Churchill 2017) and (Subbiah 2018), respectively, and updated with resources from (Serebii.net 2020). Some elements, such as the *Pikachu* variations, were omitted to avoid over-representing features in the set. Our final set contained 974 Pokémon and is available online¹. Additionally, we used the Pokémon *regional variants*, which are variations of Pokémon but with different types and slightly distinct designs (to convey their modified types), to build a special test dataset. This set provides us with useful comparison data for our model. Since a Pokémon can have one or two types, and there exist 18 Pokémon types, we handled type information as one-hot-encoded vectors of size 18, and used 0.5 in two positions for Pokémon with two types. We experimented with a one-hot encoding for types, but found it less effective for our needs.

All images were resized to 32*32 pixels using bicubic filtering, the same size as in (Krizhevsky 2009), and converted to the Hue, Saturation, Value (HSV) format, like in (Liapis 2018; Lim, Liapis, and Harrell 2016). Since dark-colored Pokémon were showing poor results, we opted to use four different background colors for each sample: black, white, and two random noise backgrounds (for training samples only). As in (Rebouças Serpa and Formico Rodrigues 2019), we only used horizontal flips for data augmentation; thus, we had eight images per Pokémon, for a total of 7204 instances in our dataset.

Given that the results obtained with the Pokémon images were not sharp nor detailed enough in initial tests, we adopted a transfer learning approach using a dataset that shared some visual traits with our target domain. Pokémon designs resemble some Japanese manga and anime styles; hence, we decided to work with the Anime Face Dataset (AFD) (Churchill 2019), which contains about 63,000 illustrations of anime-style character faces. We augmented these by flipping horizontally as well.

Transfer Learning Process

The AFD does not possess type information, which is crucial for our intended system. To provide the AFD with the types required for the transfer learning process, and to ensure that the distribution of the types in both datasets was equal, we did the following: first, we obtained the mean HSV value for each of the 18 Pokémon types, considering only non-background pixels in each image. Second, for every element in the anime set, we calculated the mean HSV value of its pixels and computed its mean squared distance with respect to each of the types’ average HSV values. Third, we used these distances as preferences (the lowest one being the most preferred), and then employed the Gale-Shapley algorithm (Gale and Shapley 1962) to assign the types. In the current implementation, each image was given only one type.

We decided to use a VAE architecture because they can reproduce given inputs with slight modifications. This behavior is beneficial to our goal since we want the modified designs to be recognizable as variations of the original one; thus, some of the source’s characteristics must be preserved, and the changes made should be enough to convey the new type information.

The proposed convolutional VAE (CVAE) model is shown in Figure 1. It is similar to the CVAE shown in (Tensorflow 2020) but adapted for HSV format, and the type information is handled like the *level design pattern labels* used in (Guzdial et al. 2018). Our model receives the images in HSV format plus the vector of type information. The model’s encoder consists of two convolution operations (512 and 1024 filters respectively) with 2x2 filters and 2x2 strides (instead of *max-pooling* (Horsley and Liebana 2017)). The filter sizes are small because the resulting images lacked detailed features (cloudiness). The second convolution’s output and the given type information are fed to the two fully connected layers with 128 units each for the latent space (mean and standard deviation). The decoder takes a vector of size 128 as input, which is passed to a fully connected layer of size 8*8*1024+18; after that, we split the last 18 values to reconstruct the type information. The remainder is passed through two deconvolutions (Zeiler et al. 2010) with 1024 and 512 filters, respectively. Finally, it passes to another deconvolu-

¹<https://github.com/EtreSerBe/PokeAE>

tion with only three filters for the output image’s HSV values. We used the Adam optimizer, and our loss function was the reduced mean of the sum of the cross-entropy and the Kullback-Leibler divergence, as proposed in (Kingma and Welling 2013). All layers’ activation functions were leaky relu, except for the latent space, which was linear, and the output’s activation, which used relu.

The model was trained first on the AFD with the added type information. The initial training stage consisted of 10 epochs with a learning rate of 0.0001 and a batch size of 128. In later stages, we fine-tuned the model by decreasing the learning rate to 0.00001 and training for 50 more epochs. Then, we fine-tuned the model on the Pokémon dataset, with a learning rate of 0.0001. We trained with a batch size of 256 for three rounds of 100 epochs each, with learning rates of 0.00005, 0.00002, and 0.00001, respectively. The low learning rates were used so the convolutional filters learned from the anime samples did not change abruptly, as they would lose the benefits of transfer learning.

We randomly split the Pokémon dataset into 6616 training instances and 588 test instances, with 827 and 147 different Pokémon, respectively. Our type-labeled AFD contains 125,130 training images and only 2,000 for testing. To be used as a baseline for comparison, we trained another instance of our model using only the Pokémon data. It was trained for 50 epochs with a learning rate of 0.0001 and a batch size of 128. Later, we fine-tuned it by training for ten rounds of 50 epochs each, with a batch size of 256.

Evaluation

The evaluation is focused on measuring our system’s outputs’ visual quality (since the Pokémon must be detailed enough to be recognized as the one given as input), and controllability based on the type information set by the user. We performed three evaluation tasks. To determine the generated images’ quality, we compared them to the input images provided, over both test and train sets. We used two comparison metrics: Mean Squared Error (MSE) in the RGB images and Structural Similarity Index (SSIM) (Zhou Wang et al. 2004), in YUV format, with filter size=11, filter sigma=1.5, k1=0.01, and k2=0.03.

On the other hand, the procedure to evaluate the controllability or *type-swap task* consisted of setting the type for every sample to a single target type and passing them to the system, instead of their original types. This process is shown in Figure 2. Note that only one type was used since the anime dataset samples were only assigned one type each. We tested this with four types: *fire*, *grass*, *water*, and *fairy*. We used the first three since most of the Pokémon of those types are red, green, and blue respectively. We used the fairy type because it was the second most preferred type during the anime faces type assignment before applying the Gale-Shapley algorithm.

Additionally, we performed a third evaluation that involves the previous two and a special *regional variants* test dataset. This dataset was composed of in-game variations of existing Pokémon where they possess different types and designs (to convey their modified types). The *original-to-regional* task consisted of comparing the visual similarity

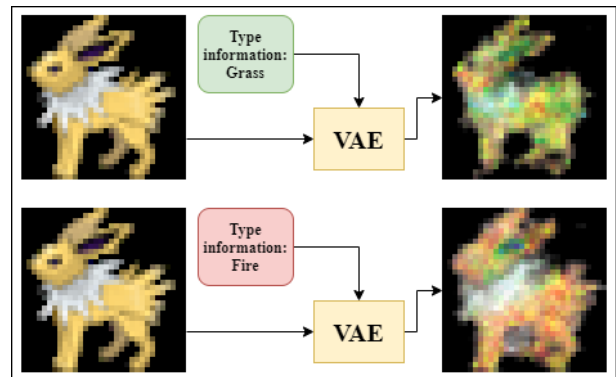


Figure 2: Type swap task example. The electric-type Pokémon on the left is modified to show *grass* (above) and *fire* (below) types. The one above shows greener tones that are common in grass-type Pokémon, while the one below has bright red colors like many fire-type Pokémon.




Figure 3: Reconstruction accuracy comparison. Results over training data are of similar quality, but, as the SSIM comparison indicates, the images are structurally more accurate when using the transfer approach.

between the regional variants’ images and our system’s output after type-swapping their non-regional versions to the variants’ types.

Results

The Pokémon reconstruction visual quality scores are shown in Tables 1 and 2. For the MSE results in Table 1 lower is better, and the transfer learning model consistently outperforms the non-transfer one. For the SSIM results in Table 2 higher is better, which indicates at least a 1% increase in visual similarity to the inputs when using the transfer learning approach, even though the two datasets are vastly different. Note that for all of this article’s figures all images shown, including the Pokémon inputs, were resized from 32x32 pixels to 128x128 using the nearest neighbor method. An example of both systems’ outputs is presented in Figure 3. Both models present blurry outcomes, a known drawback of autoencoders, which will be improved in future work.

Initial type-swap task outcomes were barely distinct from

Model version	Test	Train	Test and train
Transfer learning	0.03216	0.01409	0.01692
Non-transfer	0.03349	0.01438	0.01733

Table 1: Image comparison between actual Pokémon and reconstructed Pokémon images using Mean Squared Error.

Model version	Test	Train	Test and train
Transfer learning	0.3940	0.6496	0.6109
Non-transfer	0.38361	0.6268	0.5895

Table 2: Image comparison values using the Structural Similarity Index (SSIM).

the inputs. To make them more evident, we increased the type vector’s magnitude from 1 to 20. Representative results of the type-swap task are shown in Figures 4, 5, and 6, but the ones for fairy-type were omitted because no consistent changes were noticeable. The lack of changes may have been caused by the large number of AFD’s images that preferred the fairy type but were reassigned to another one by the Gale-Shapley algorithm (13717, which is 10.88% versus the final 4.05% or 5109). This reassignment likely caused several images’ fairy-like features to be ignored.

In Figure 4, the non-transfer model generates more evident visual changes. However, its effects can be uncontrollable and make it difficult to recognize from the original, which might be caused by the amount of variety of samples. Both models’ results showed the expected red colors that most fire Pokémon possess, but the transfer learning one also presented some unexpected vivid green tones, likely due to how the types were assigned to the AFD’s images.

In Figure 5, both models presented noise on the white background during type swap experiments, especially with type vectors of large magnitude. The transfer learning model’s results present the expected green coloring. On the other hand, the baseline model generated very noisy images. Several AFD’s samples have white background, which might explain why the transfer approach had less isolated pixels in this experiment.

Finally, in Figure 6, the non-transfer model’s results present considerable changes, but also noticeable artifacts. The transfer model’s outputs over testing data were slightly more blue, but did not show the expected blue tones when using red or yellow Pokémon. This problem could be caused by the small number of AFD’s images labeled as water-type before the Gale-Shapley algorithm was used (about 2000).

Model version	MSE	SSIM
Transfer learning	0.1000	0.1704
Non-transfer	0.0952	0.1633

Table 3: *Original-to-regional* visual quality task results. The transfer learning version presents slightly less pixel-wise errors but higher structural similarity with the regional variants.

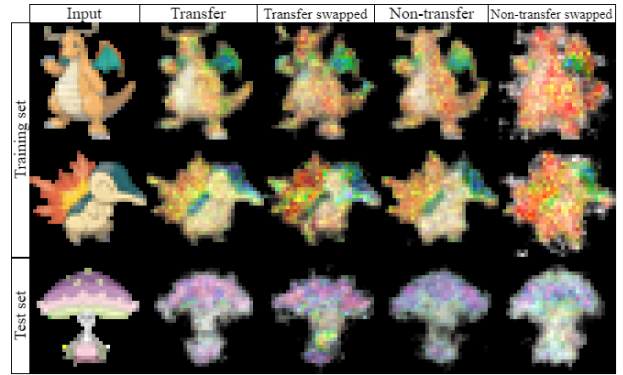


Figure 4: *Fire* type swap task results. The non-transfer model generates more evident visual changes but its effects are uncontrollable. Their original types were: *dragon-flying*, *fire*, and *grass-fairy*.

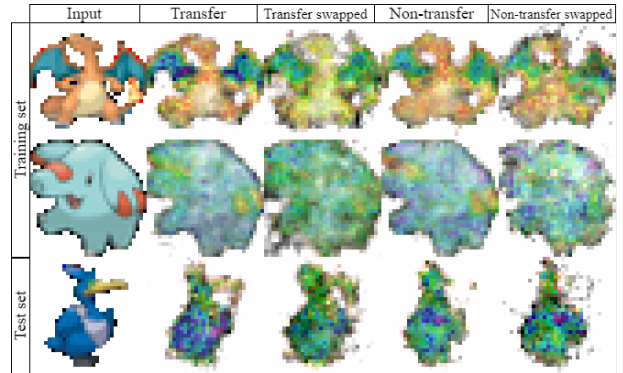


Figure 5: *Grass* type swap task results using white background samples. Both models had troubles with that background during type swap tasks. Their original types were: *fire-flying*, *ground*, and *water-flying*.

That is less than 2% of the samples, compared to the 11.5% water Pokémon, therefore, that difference was compensated with anime faces that might not have been suitable as water-type. While the non-transfer version of the model shows more evident changes in this task, it also presents considerably more noise, both inside and outside of the Pokémon, which is not desirable in the final content. Therefore, with our current results, we argue that pursuing a transfer learning approach is more suitable if the generated content is aimed towards the general public, especially when the training data is scarce. However, we acknowledge that both models’ results over testing data have plenty of room for improvement.

For the original-to-regional task, the visual quality comparison results are shown in Table 3. The scores are considerably lower than for the reconstruction task, in part because the regional variants not only change colors but also size, shape and pose. We also present some positive visual results in Figure 7. We note that changing the fourth Pokémon (*Slowbro*) from *water-psychic* to *poison-psychic* introduced the same purple splotches seen in the real regional variant.

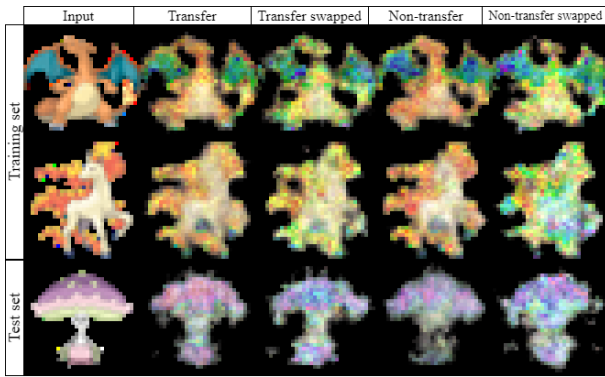


Figure 6: *Water* type swap task results. Their original types were: *fire-flying*, *fire*, and *grass-fairy*.




Figure 7: *Original-to-regional* task representative results. The bottom row elements are existing type-swapped versions of the Pokémon in the first row. The middle row ones are our system’s proposed type-swapped designs. The target types are above each column. The first column shows darker tones common in dark-type Pokémon, and the third one seems more white (similar to most ice-types). Their original types are, from left to right, *normal*, *electric*, *fire*, *psychic-water*, and *ground*.

We illustrate how our VAE interpolates between two different Pokémon in Figure 8.

Limitations and Future Work

We identify three crucial aspects to improve our proposed system. First, we consider that using a different method to assign the type information to the Anime Face Dataset could lead to improvements for the type-based controllability of the transfer model. This is because one HSV tuple for each type cannot hold information about spatial or structural features (for instance, most Flying-type Pokémon have wings). Moreover, the Saturation and Value channels in the type were confined to small ranges [0.22, 0.37] and [0.51, 0.69] (in contrast, the Hue values ranged between [0.19, 0.51]), which resulted in undesired type overlaps. Additionally, several images from the AFD have white background, which causes the type assignment process to favor types with



Figure 8: Interpolation between two different Pokémon’s latent representations. The *fire*-type Pokémon on the left changes towards the *electric*-type one on the right. This is useful as we could explore which dimensions affect features related to each type, to improve the type swap results.

brighter palettes such as *fairy* and *fire*.

Second, the current network architecture used is simple. We argue that exploring other alternatives, such as using image patches, discriminator modules, or moving to different architectures (such as GANs), could lead to better image quality and controllability results. Third, we require evaluating how human artists respond to this kind of system, which will provide us with crucial feedback about future work directions.

On a more distant horizon, we also consider the exciting possibility that this kind of approach could be applied in games. For instance, it could be used to automatically generate visual indicators for Pokémon that are affected with special status effects, such as burned, frozen, or poisoned, which are currently only shown as a written text. Visual indicators like these could help to enhance a player’s understanding of the game’s mechanics.

Conclusions

We propose a Convolutional Variational Autoencoder (CVAE) system to modify Pokémon sprites according to a target Pokémon type. Our experimental results indicate that adopting a transfer learning approach, using a type-labeled version of the Anime Face Dataset, can help to improve visual quality and stability over unseen data, despite the considerable differences between both domains. While the presented models’ outcomes might be usable during very early stages of the design process, their quality and controllability are not yet suitable for game development beyond that point. However, we expect that this problem will diminish in future versions of the system.

Acknowledgements

This work was supported by CONACYT through the doctoral scholarship number CVU-777880. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Alberta Machine Intelligence Institute (Amii).

References

- Burgess, C. P.; Higgins, I.; Pal, A.; Matthey, L.; Watters, N.; Desjardins, G.; and Lerchner, A. 2018. Understanding disentangling in -vae.
- Chandra, R.; Grover, S.; Lee, K.; Meshry, M.; and Taha, A. 2017. Texture synthesis with recurrent variational auto-encoder.

- Churchill, S. 2017. Pokemon images dataset. <https://www.kaggle.com/kvpratama/pokemon-images-dataset>.
- Churchill, S. 2019. Anime face dataset. <https://www.kaggle.com/splcher/animefacedataset>.
- Gale, D., and Shapley, L. S. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69(1):9–15.
- 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.
- Guzdial, M.; Long, D.; Cassion, C.; and Das, A. 2017. Visual procedural content generation with an artificial abstract artist. In *Proceedings of ICCG Computational Creativity and Games Workshop*.
- Guzdial, M.; Reno, J.; Chen, J.; Smith, G.; and Riedl, M. 2018. Explainable PCGML via game design patterns. *CoRR* abs/1809.09419.
- Horsley, L., and Liebana, D. P. 2017. Building an automatic sprite generator with deep convolutional generative adversarial networks. *2017 IEEE Conference on Computational Intelligence and Games (CIG)* 134–141.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2016. Image-to-image translation with conditional adversarial networks.
- Jin, Y.; Zhang, J.; Li, M.; Tian, Y.; Zhu, H.; and Fang, Z. 2017. Towards the automatic anime characters creation with generative adversarial networks.
- Karras, T.; Laine, S.; and Aila, T. 2018. A style-based generator architecture for generative adversarial networks. *CoRR* abs/1812.04948.
- Kingma, D. P., and Welling, M. 2013. Auto-encoding variational bayes.
- Kingma, D. P., and Welling, M. 2019. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning* 12(4):307–392.
- Krizhevsky, A. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.
- Larsen, A. B. L.; Sønderby, S. K.; and Winther, O. 2015. Autoencoding beyond pixels using a learned similarity metric. *CoRR* abs/1512.09300.
- 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.
- Nintendo. 2019. Pokemon series. <https://www.pokemon.com/>.
- Pihlgren, G. G.; Sandin, F.; and Liwicki, M. 2020. Improving image autoencoder embeddings with perceptual loss.
- 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.
- Serebii.net. 2020. Serebii.net. <https://serebii.net/>.
- 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/>.
- Snodgrass, S., and Sarkar, A. 2020. Multi-domain level generation and blending with sketches via example-driven bsp and variational autoencoders. *International Conference on the Foundations of Digital Games*.
- Subbiah, V. 2018. Pokemon image dataset. <https://www.kaggle.com/vishalsubbiah/pokemon-images-and-types?select=pokemon.csv>.
- Summerville, A.; Snodgrass, S.; Guzdial, M.; Holmgård, C.; Hoover, A. K.; Isaksen, A.; Nealen, A.; and Togelius, J. 2018. Procedural content generation via machine learning (pcgml). *IEEE Transactions on Games* 10:257–270.
- Tensorflow. 2020. Convolutional variational autoencoder. <https://www.tensorflow.org/tutorials/generative/cvae>.
- Voynov, A., and Babenko, A. 2020. Unsupervised discovery of interpretable directions in the gan latent space. *ArXiv* abs/2002.03754.
- Wong, R. N. 2019. pokemon2pokemon: Using neural networks to generate pokemon as different elemental types. <https://www.rileywong.com/blog/2019/5/22/pokemon2pokemon-using-cycleGAN-to-generate-pokemon-as-different-elemental-types>.
- Zeiler, M. D.; Krishnan, D.; Taylor, G. W.; and Fergus, R. 2010. Deconvolutional networks. In *In CVPR*.
- Zhou Wang; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13(4):600–612.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. *2017 IEEE International Conference on Computer Vision (ICCV)* 2242–2251.