

LoGAN: Generating Logos with a Generative Adversarial Neural Network Conditioned on Color

Citation for published version (APA):

Mino, A., & Spanakis, G. (2018). LoGAN: Generating Logos with a Generative Adversarial Neural Network Conditioned on Color. In *17th IEEE International Conference on Machine Learning and Applications, ICMLA 2018, Orlando, FL, USA, December 17-20, 2018* (pp. 965-970) <https://doi.org/10.1109/ICMLA.2018.00157>

Document status and date:

Published: 01/12/2018

DOI:

[10.1109/ICMLA.2018.00157](https://doi.org/10.1109/ICMLA.2018.00157)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Download date: 26 Apr. 2024

LoGAN: Generating Logos with a Generative Adversarial Neural Network Conditioned on color

Ajkel Mino

Department of Data Science and Knowledge Engineering
Maastricht University
Maastricht, Netherlands
ajkimino@gmail.com

Gerasimos Spanakis

Department of Data Science and Knowledge Engineering
Maastricht University
Maastricht, Netherlands
jerry.spanakis@maastrichtuniversity.nl

Abstract—Designing a logo is a long, complicated, and expensive process for any designer. However, recent advancements in generative algorithms provide models that could offer a possible solution. Logos are multi-modal, have very few categorical properties, and do not have a continuous latent space. Yet, conditional generative adversarial networks can be used to generate logos that could help designers in their creative process. We propose LoGAN: an improved auxiliary classifier Wasserstein generative adversarial neural network (with gradient penalty) that is able to generate logos conditioned on twelve different colors. In 768 generated instances (12 classes and 64 logos per class), when looking at the most prominent color, the conditional generation part of the model has an overall precision and recall of 0.8 and 0.7 respectively. LoGAN’s results offer a first glance at how artificial intelligence can be used to assist designers in their creative process and open promising future directions, such as including more descriptive labels which will provide a more exhaustive and easy-to-use system.

Index Terms—Conditional Generative Adversarial Neural Network, Logo generation

I. INTRODUCTION

Designing a logo is a lengthy process that requires continuous collaboration between the designers and their clients. Each drafted logo takes time and effort to be designed, which turns this creative process into a tedious and expensive endeavor.

Recent advancements in generative models suggest a possible use of artificial intelligence as a solution to this problem. Specifically, Generative Adversarial Neural Networks (GANs) [1], which learn how to mimic any distribution of data. They consist of two neural networks, a generator and a discriminator, that are contended against each other, trying to reach a Nash Equilibrium [2] in a minimax game.

Whilst GANs and other generative models have been used to generate almost anything, from MNIST digits [1] to anime faces [3], Mario levels [4] and fake celebrities [5], logo generation has yet to receive thorough exploration. One possible explanation is that logos do not contain a hierarchy of nested segments, which networks can learn and try to reproduce. Furthermore, their latent space is not continuous, meaning that not every generated logo will necessarily be aesthetically pleasing. To the best of our knowledge, Sage,

et al. [6] to be only one to have tackled this problem thus far. They propose a clustered approach for dealing with multi-modal data, specifically logos. Logos get assigned synthetic labels, defined by the cluster they belong to, and a GAN is trained, conditioned on these labels. However, synthetic labels do not provide enough flexibility. The need for more descriptive labels arises: labels that could better convey what is in the logos, and present designers with more detailed selection criteria.

This paper provides a first step to more expressive and informative labels instead of synthetic ones, which is achieved by defining logos using their most prominent color. Twelve color classes are introduced: black, blue, brown, cyan, gray, green, orange, pink, purple, red, white, and yellow. We propose LoGAN: an improved Auxiliary Classifier Wasserstein Generative Adversarial Neural Network (with Gradient Penalty) (AC-WGAN-GP) that is able to generate logos conditioned on the aforementioned twelve colors.

The rest of this paper is structured as follows: In Section 2 the related work and background will be discussed, then in Section 3 the proposed architecture be conveyed, followed by Section 4, where the dataset and labeling process will be explained. Section 5 presents the experimental results, consisting of the training details and results, and finally we conclude and discuss some possible extensions of this work in Section 6.

II. BACKGROUND & RELATED WORK

First proposed in 2014 by Goodfellow et al. [1], GANs have seen a rise in popularity during the last couple of years, after possible solutions to training instability and mode collapse were introduced [7]–[9].

A. Generative Adversarial Networks

GANs (architecture depicted in Fig.1(a)) consist of two different neural networks, a generator and a discriminator, that are trained simultaneously in a competitive manner.

The generator is fed a noise vector ($G(z)$) from a probability distribution (p_z), and outputs a generated data-point (fake image). The discriminator takes its input either from the generator (the fake image) ($D(G(z))$) or from the training set (the real image) ($D(x)$) and is trained to distinguish

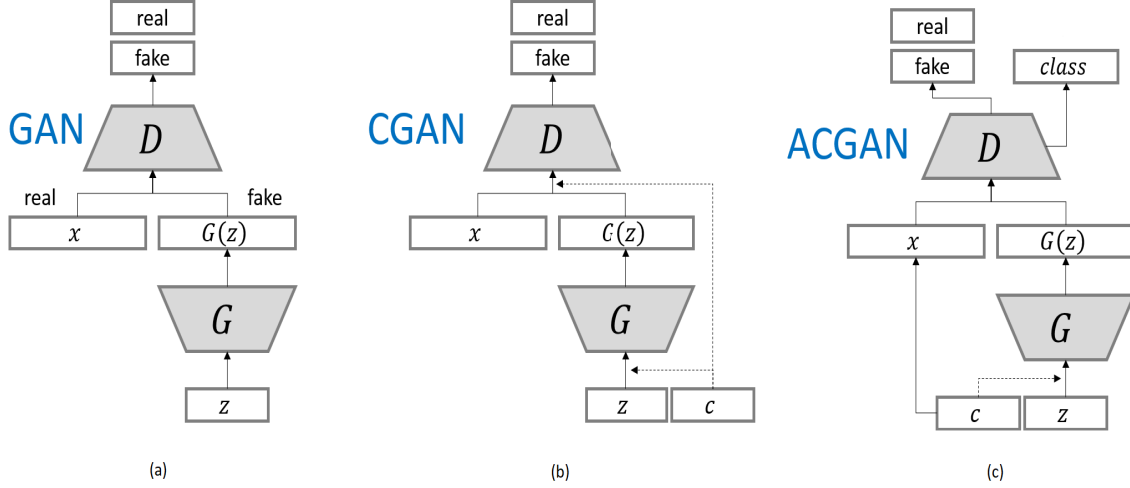


Fig. 1. GAN, conditional GAN (CGAN) and auxiliary classifier GAN (ACGAN) architectures, where x denotes the real image, c the class label, z the noise vector, G the Generator, and D the Discriminator.

between the two. The discriminator and the generator play a two-player minimax game with value function (1), where the discriminator tries to maximize V , while the generator tries to minimize it.

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{data}} [\log D(x)] + \mathbb{E}_{z \sim p_z} [\log(1 - D(G(z)))] \quad (1)$$

1) *Objective functions*: While GANs are able to generate high-quality images, they are notoriously difficult to train. They suffer from problems like training stability, non-convergence and mode collapse. Multiple improvements have been suggested to fix these problems [7]–[9]; including using deep convolutional layers for the networks [10] and modified objective functions i.e. to least-squares [11] or to Wasserstein distance between the distributions [12]–[14].

B. Conditionality

Conditional generation with GANs entails using labeled data to generate images based on a certain class. The two types that will be discussed in the subsections below are Conditional GANs and Auxiliary Classifier GANs.

1) *Conditional GANs*: In a conditional GAN (CGAN) [15] (architecture depicted in Fig.1(b)) the discriminator and the generator are conditioned on c , which could be a class label or some data from another modality. The input and c are combined in a joint hidden representation and fed as an additional input layer in both networks.

2) *Auxiliary Classifier GANs*: Contrary to CGANs, in Auxiliary Classifier GANs (AC-GAN) [16] (architecture depicted in Fig.1(c)) the latent space z is conditioned on the class label. The discriminator is forced to identify fake and real images, as well as the class of the image, irrespective of whether it is fake or real.

C. GAN Applications

These different GAN architectures have been used for numerous purposes, including (but not limited to): higher-quality image generation [5], [17], image blending [18], image super-resolution [19] and object detection [20].

Up until the writing of this paper, logo generation has previously only been investigated by Sage et al. [6], who accomplish three main things:

- Define the Large Logo Dataset (LLD) [21]
- Synthetically label the logos by clustering them using the ResNet Classifier network
- Build a GAN conditioned on these synthetic labels

However, as these labels are defined from computer-generated clusters, they do not necessarily provide intuitive classes for a human designer. More descriptive labels, that could provide designers with more detailed selection criteria and better convey what is in the logos are needed.

This paper offers a solution to the previous problem by:

- 1) Using twelve colors to define the logo classes
- 2) Defining LoGAN: An AC-WGAN-GP that can conditionally generate logos

III. PROPOSED ARCHITECTURE

The proposed architecture for LoGAN is an Auxiliary Classifier Wasserstein Generative Adversarial Neural Network with gradient penalty (AC-WGAN-GP), depicted on Fig.2. LoGAN is based on the ACGAN architecture, with the main difference being that it consists of three neural networks, namely the discriminator, the generator and an additional classification network¹. The latter is responsible for assisting the discriminator in classifying the logos, as the original classification loss

¹Code for this paper is available via <https://github.com/ajki/LoGAN>

from AC-GAN was dominated by the Wasserstein distance used by the WGAN-GP.

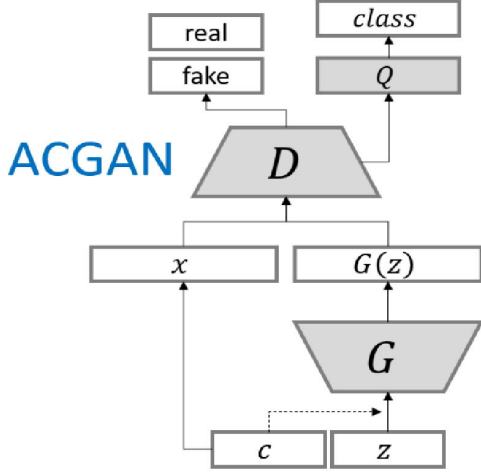


Fig. 2. LoGAN architecture, where x denotes the real image, c the class label, z the noise vector, G the Generator, D the Discriminator and Q the Classifier

In an original AC-GAN [16] the discriminator is forced to optimize the classification loss together with the real-fake loss. This is shown in equations (2) & (3), which describe the loss from defining the source of the image (training set or generator) and the loss from defining the class of the image respectively. The adversarial aspect of AC-GAN comes as the discriminator tries to maximize $\mathcal{L}_{Class}^{ACGAN} + \mathcal{L}_{Source}^{ACGAN}$, whilst the generator tries to maximize $\mathcal{L}_{Class}^{ACGAN} - \mathcal{L}_{Source}^{ACGAN}$ [16].

$$\mathcal{L}_{Source}^{ACGAN} = \mathbb{E}[\log P(S = real|X_{real})] + \mathbb{E}[\log P(S = fake|X_{fake})] \quad (2)$$

$$\mathcal{L}_{Class}^{ACGAN} = \mathbb{E}[\log P(C = c|X_{real})] + \mathbb{E}[\log P(C = c|X_{fake})] \quad (3)$$

Since WGAN-GP is more stable while training, the proposed architecture makes use of the WGAN-GP loss, instead of the AC-GAN loss. The loss function for the discriminator and generator of LoGAN are the same as a WGAN-GP [14], stated in equations (4) & (5).

$$\mathcal{L}_D^{WGANGP} = -\mathbb{E}_{x \sim p_d}[D(x)] + \mathbb{E}_{\hat{x} \sim p_g}[D(\hat{x})] + \lambda \mathbb{E}_{\hat{x} \sim p_g}[(\|\nabla D(\alpha x + (1 - \alpha)\hat{x})\|_2 - 1)^2] \quad (4)$$

$$\mathcal{L}_G^{WGANGP} = -\mathbb{E}_{\hat{x} \sim p_g}[D(\hat{x})] \quad (5)$$

The loss of the additional classifier network, is defined as:

$$\mathcal{L}_Q^{LoGAN} = \mathbb{E}_{x,y}[\log Q(y|x)] \quad (6)$$

To avoid instability during training and mode collapse certain measures were taken. The generator and classifier were

trained once for every 5 iterations of discriminator training, as suggested by Gulrajan et al. [14], z was sampled from a Gaussian Distribution [22], and batch normalization was used [23].

IV. DATA

The dataset used to train LoGAN is the LLD-icons dataset [21], which consists of 486'377 32×32 icons.

1) *Labeling*: To extract the most prominent color from the image a k-Means algorithm with $k = 3$ was used to define the RGB values of the centroids. The algorithm makes use of the MiniBatchKMeans implementation from sci-kit learn package, and the package webcolors was used to turn the RGB values into color words. The preliminary colors consisted of X11 color names², which were grouped into 12 main classes: black, blue, brown, cyan, gray, green, orange, pink, purple, red, white, and yellow. The class distribution can be observed in Fig.3.

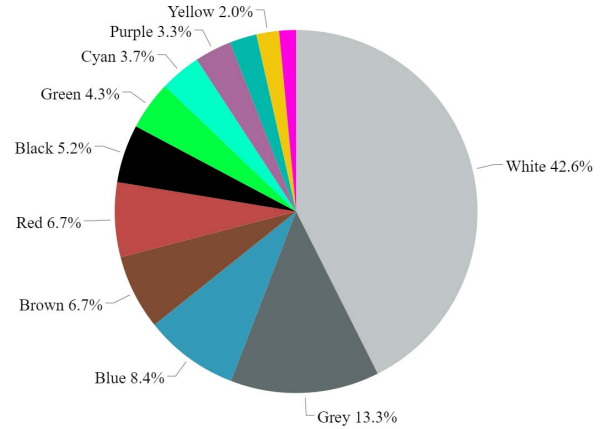


Fig. 3. Dataset distribution by class.

V. EXPERIMENTAL RESULTS

In this section, the training process, quality evaluation and the results obtained from the model³ will be presented.

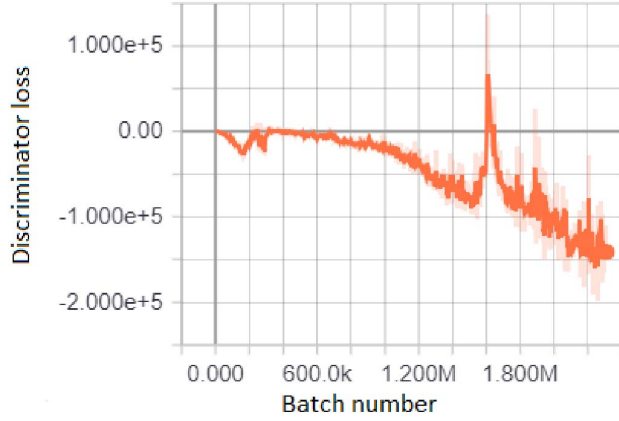
A. Training

Fig.4 shows the loss graphs per batch for the discriminator 4a, generator 4b and classifier 4c. It can be observed that both the discriminator and the generator have not converged as the loss graphs have a downward trend. This does, however, not imply improper training as neither WGAN nor WGAN-GP are guaranteed to reach convergence [8].

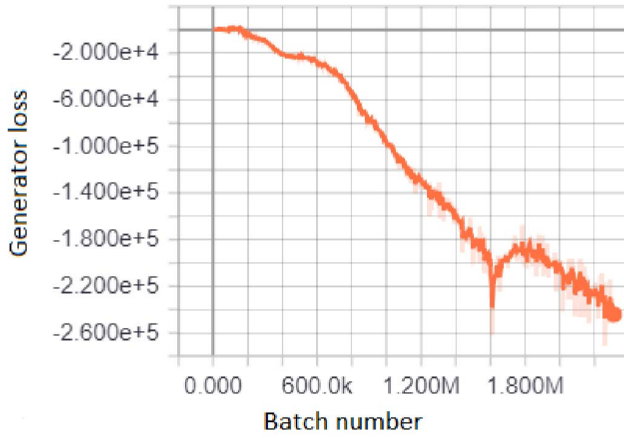
The classifier on the other hand has converged with a loss value close to 1. Further investigation into the classification loss shows that the classification loss for fake images has

²A list of the X11 color names, and the class grouping of RGB values can be found on Wikipedia at : https://en.wikipedia.org/wiki/Web_colors#X11_color_names

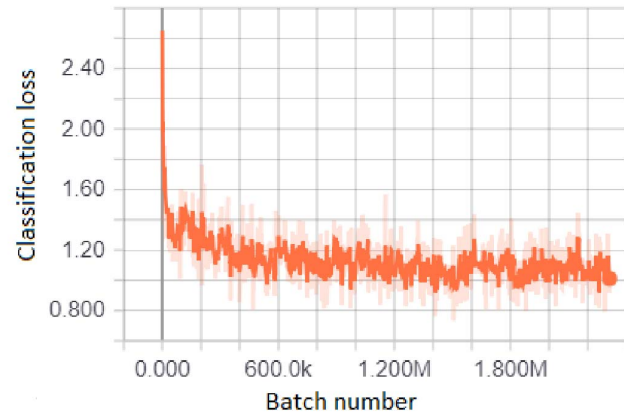
³The model was trained on a Windows 10 machine, with a Tesla K80 GPU, for 400 epochs, which lasted around three full days.



(a)



(b)



(c)

Fig. 4. Discriminator, Generator and Classifier losses, where the X-axis represents the batch number, and the Y-axis denotes the loss value for that certain batch.

converged to zero (loss depicted on Fig.5). This means that the generated images get classified correctly.

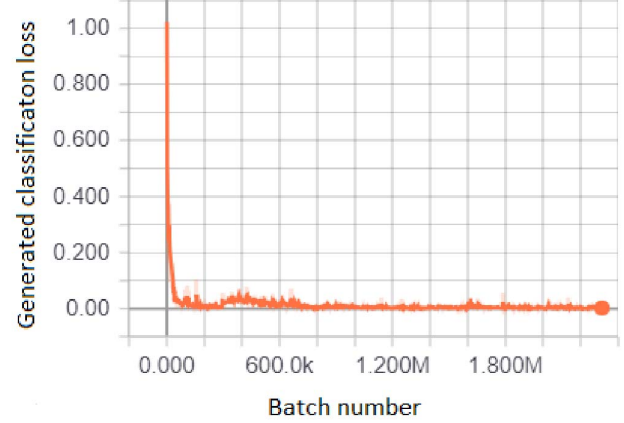


Fig. 5. Classification loss for generated images.

B. Quality evaluation

The network is expected to generate logo-resembling images that have clear color definition. In each epoch, 64 logos will be generated per class. The top three most prominent colors in the logos will be extracted, and the generated pairs and triplets will be analyzed. For each class (c), the precision (7), recall (8) and F1-score (9) will be measured for the most prominent color in the logo.

$$Precision = \frac{correctly_generated_as(c)}{total_generated_as(c)} \quad (7)$$

$$Recall = \frac{correctly_generated_as(c)}{total_actual(c)} \quad (8)$$

$$F1 = 2 * \frac{Precision * Recall}{Precision + Recall} \quad (9)$$

C. Results

The results for the class conditioned generation after 400 epochs are shown in Fig.6. As expected, a slight blurriness can be noticed on the generated logos. This is because the training images are only 32×32 pixels. Despite the blurriness, the generated logos resemble feasible ones. The generated images are mainly dominated by round and square shapes. Even irregular shapes have been generated, for example the heart and the x in the array of white logos. A look-alike of the Google Chrome logo is also present in the midst of the cyan class.

The precision, recall and F1-score of the class conditioned generation after 400 epochs is shown on Table I. The precision scores are relatively high with the exception of white and gray. This is because the most predominantly white logos are small logos with a lot of white or transparent space around. Consequently many small logos generated in other colors are going to be classified as white. Similarly, because of its neutrality the gray class also appears a lot in other classes.

The recall values are lower than the precision overall, with red having the highest recall at 0.92 and pink having the lowest




Fig. 6. Results from the generation of 64 logos per class after 400 epochs of training. Classes from left to right top to bottom: green, purple, white, brown, blue, cyan, yellow, gray, red, pink, orange, black.

TABLE I
PRECISION, RECALL AND F1-SCORE OF THE MOST PROMINENT COLOR IN EACH OF THE LOGOS IN FIG.6.

Class	Precision	Recall	F1
black	0,95	0,86	0,90
blue	0,73	0,69	0,71
brown	0,63	0,47	0,55
cyan	0,98	0,66	0,79
gray	0,57	0,50	0,53
green	1	0,80	0,89
orange	0,96	0,80	0,87
pink	0,95	0,30	0,45
purple	0,65	0,41	0,50
red	0,84	0,92	0,88
white	0,24	0,83	0,38
yellow	0,96	0,78	0,86
Average	0,79	0,67	0,69

at 0.3. Most of the pink labeled logos generated belong to class white.

The F1-score marks black on top with 0.90, and white on the bottom with 0.38. Based on Equation (9), and the fact that white has the lowest precision, its F1-score is also expected to be low.

Fig.7 shows the distribution of the top three generated colors in each class. This distribution is calculated by extracting the top three colors from each logo, and distinguishing the most prominent ones in each category. As expected from the results in Table I, white and gray are present in the top three for many of the classes. Some interesting combinations include the orange class, where the color brown appears, and the yellow class, where the color blue appears. There are three classes which get generated using only shades from their own class (blue, brown, purple). However, if we consider black and white as shades of gray, that number rises to five.

VI. CONCLUSION & FUTURE WORK

In this paper, it was shown how designers can be assisted in their creative process for logo design by modern artificial

intelligence techniques, namely Generative Adversarial Networks. The proposed model can successfully create logos if given a certain keyword, which in our case consisted of the most prominent color in the logo. This class of keywords can be considered descriptive as it provides a property of the logo that is easy for humans to distinguish.

The proposed architecture consists of an Auxilliary Classification Wasserstein GAN with gradient penalty (AC-WGAN-GP) and generates logos conditioned on twelve colors. This helps designers in their creative process and while brainstorming, making logo design cheaper and less time-consuming. As the generated logos have very low resolution, they can serve as a very rough first draft of a final logo, or as a means of inspiration for the designer. Regarding the results, the classifier converged, and the generated logos meet the requirements of the class they belong to. This is backed up by precision, recall, and comparison with other logos.

Despite the promising results, a limitation of the approach is the blurriness of the generated logos. At the same time, color is not a stand-alone keyword for defining a logo. Higher resolution training images and other labels such as the shape of the logo or the focus of the company would provide valuable input, thus improving the results.

Possible extensions to this work include conditioning on more labels, such as the shape of the logo. However, as logos do not always have clear geometrical shapes, possible issues could include logos with text only or logos with irregular shapes. These issues could be overcome by splitting the dataset into two main groups, logos with an obvious geometrical shape, e.g. quadrilateral, circular or triangular; and logos with a non-regular shape, e.g. text, letters, hearts, etc. Extracting shapes from the first group can be easily done using packages like OpenCV. The non-regular group of shapes may be a bit more challenging to label. A possible path is to use a tool such as `tesseract-ocr` perform optical character recognition to extract the text in the logo, and use it as the label. As for the irregularly shaped logos, those could be put in an extra class

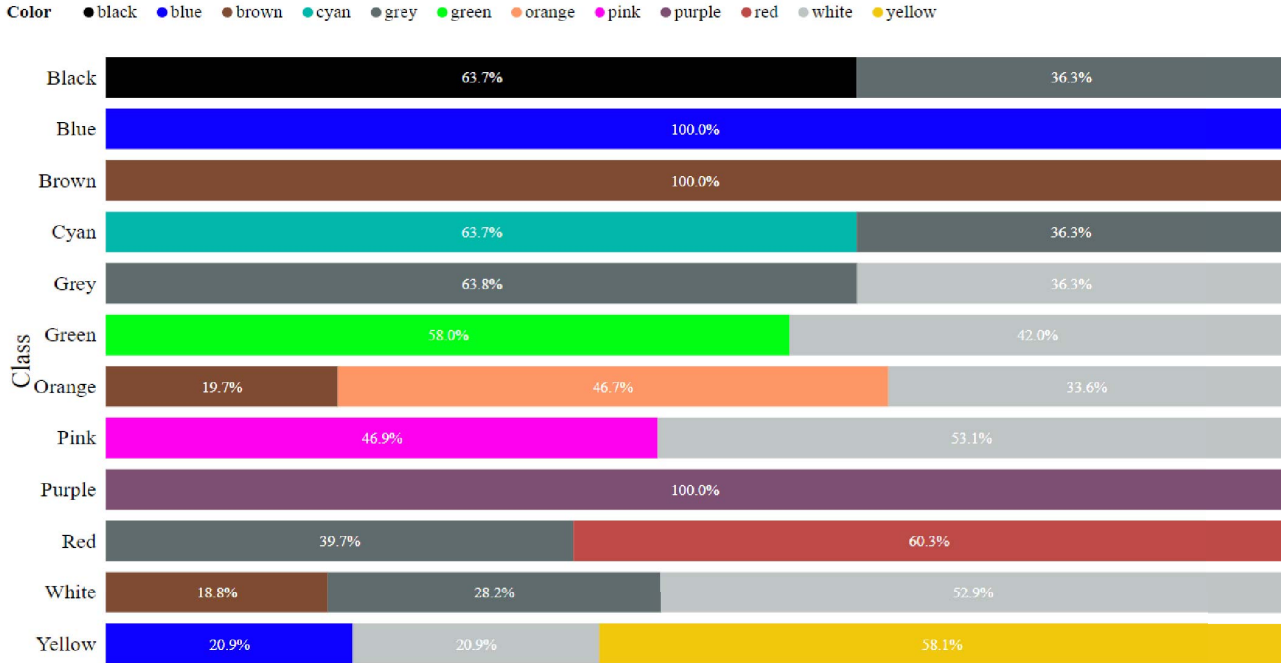


Fig. 7. Distribution of the top three generated colors per class.

with a label such as 'others'.

Finally, another possible set of labels for the logos could be gathered for the LLD-logos dataset (which also provides higher quality images) by gathering the most used words to describe the company the logo belongs to. Combining these labels with word embedding models could potentially introduce a semantic meaning to the logos, further boosting the interpretability of the current approach.

REFERENCES

- [1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [2] J. F. Nash *et al.*, "Equilibrium points in n-person games," *Proceedings of the national academy of sciences*, vol. 36, no. 1, pp. 48–49, 1950.
- [3] Y. Jin, J. Zhang, M. Li, Y. Tian, H. Zhu, and Z. Fang, "Towards the automatic anime characters creation with generative adversarial networks," unpublished.
- [4] V. Volz, J. Schrum, J. Liu, S. M. Lucas, A. M. Smith, and S. Risi, "Evolving mario levels in the latent space of a deep convolutional generative adversarial network," in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2018)*. New York, NY, USA: ACM, July 2018. [Online]. Available: <http://doi.acm.org/10.1145/3205455.3205517>
- [5] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of gans for improved quality, stability, and variation," unpublished.
- [6] A. Sage, E. Agustsson, R. Timofte, and L. Van Gool, "Logo synthesis and manipulation with clustered generative adversarial networks," unpublished.
- [7] N. Kodali, J. Abernethy, J. Hays, and Z. Kira, "On convergence and stability of gans," unpublished.
- [8] L. Mescheder, A. Geiger, and S. Nowozin, "Which training methods for gans do actually converge? arxiv preprint," unpublished.
- [9] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training gans," in *Advances in Neural Information Processing Systems*, 2016, pp. 2234–2242.
- [10] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," unpublished.
- [11] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley, "Least squares generative adversarial networks," in *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE, 2017, pp. 2813–2821.
- [12] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein gan," unpublished.
- [13] D. Berthelot, T. Schumm, and L. Metz, "Bgan: Boundary equilibrium generative adversarial networks," unpublished.
- [14] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of wasserstein gans," in *Advances in Neural Information Processing Systems*, 2017, pp. 5769–5779.
- [15] M. Mirza and S. Osindero, "Conditional generative adversarial nets," unpublished.
- [16] A. Odena, C. Olah, and J. Shlens, "Conditional image synthesis with auxiliary classifier gans," unpublished.
- [17] H. Zhang, T. Xu, H. Li, S. Zhang, X. Huang, X. Wang, and D. Metaxas, "Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks," in *IEEE Int. Conf. Comput. Vision (ICCV)*, 2017, pp. 5907–5915.
- [18] H. Wu, S. Zheng, J. Zhang, and K. Huang, "Gp-gan: Towards realistic high-resolution image blending," unpublished.
- [19] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network," unpublished.
- [20] J. Li, X. Liang, Y. Wei, T. Xu, J. Feng, and S. Yan, "Perceptual generative adversarial networks for small object detection," in *IEEE CVPR*, 2017.
- [21] A. Sage, E. Agustsson, R. Timofte, and L. Van Gool, "Lld - large logo dataset - version 0.1," <https://data.vision.ee.ethz.ch/cvl/lld>, 2017.
- [22] T. White, "Sampling generative networks: Notes on a few effective techniques," unpublished.
- [23] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," unpublished.