CS230

Fake news: generating nonsense newspapers in the name of art

Miranda Li*
Department of Computer Science
Stanford University
mirandal@stanford.edu

Abstract

We present an artistic project employing a PGAN pipeline to generate gibberish newspapers. This project aims to contrast the antiquity of print newspaper design aesthetics with the technical modernity of GANs, and unite these two seemingly disjoint things to create a visual result that is simultaneously meaningful and nonsensical, familiar and foreign, absurd and mundane. View the project here.

1 Introduction

Print design, particularly of newspapers, has changed so little since the days of manual printing that its visual signature has become ubiquitous.

(a) Front page of The New York Times for October 28, 2020.

(b) Front page of the Hickory Daily Record for October 28, 1920.

Figure 1

In figure 1, we juxtapose the front page of the New York Times for the day of October 28, 2020 with a newspaper from October 28, 1920, exactly a century ago. Note the similarities – the small note about the weather in the top right corner, the justified body text, the large picture front and center.

*https://mirandali707.github.io/

CS230: Deep Learning, Winter 2018, Stanford University, CA. (LateX template borrowed from NIPS 2017.)

The consistency of newspaper design is by no means surprising: newspapers primarily function as a vehicle for the information they contain. The aesthetics of a newspaper are not the focus; print editors will be familiar with the saying "content drives design" – the content of the articles is the focus, and the design only aims to make absorption of this content as seamless as possible.

We aim to defy the notion that newspaper design must always serve its content, and instead showcase print design as an art form in and of itself. By stripping the front page of a newspaper of its meaning, we can begin to see the aesthetics of its design in a new, language-agnostic light, which is no longer obfuscated by our human urge to read words when we see them.

Generative adversarial networks (GANs) are uniquely suited for this artistic endeavor because the generator (metaphorically, the forger) will learn from scans of man-made newspapers, but with no human intuition for what a newspaper represents or contains. The discriminator evaluates a given newspaper, real or generated, with a complete lack of language understanding or conceptual context. To the discriminator, a nonsense, well-crafted fake is just as good as the real thing. While any literate human viewer, burdened with the conceptual understanding that newspapers contain words which are meant to be read, has no choice but to read the words on the newspaper even subliminally (is it even possible to see text in an entirely language- and meaning-agnostic way once one is fluent in that language?), a GAN trained only on image data of newspaper scans is subject to no such compulsion. In this way, the GAN itself becomes an integral conceptual piece of this artwork – the GAN functions as a man-made way to divorce ourselves completely from our unshakeable man-made understanding of what a newspaper is.

These GAN-generated images are subsequently presented via a website containing both a gallery view and a digital zine. The true final product, however, is a physical zine, containing generated newspapers and Photoshopped images of people reading them. This zine (which, at the time of writing, has yet to be printed and distributed, but will be very soon) aims to recall the sensory memories we have of leafing through newspapers and return what was once physical (the historical newspapers which were scanned by the Library of Congress to produce our training set) but has since been void of its physicality (downloaded, digitally cropped and scaled, processed as a collection of pixels through a GAN, regurgitated by this same GAN as a new, slightly different collection of pixels) back to the physical realm once again.

These generated newspapers, stripped of their cultural significance and linguistic meaning, distorted just enough to feel unsettling but not enough to be unrecognizable, represent a warring (or maybe a uniting) of worlds: analog and technological, ubiquitous and novel, meaningful and meaningless. These generated newspapers highlight the advances of AI but also its limitations: computers, unlike humans, don't read unless they're told to.

2 Related work

Lately, GANs have been popular for the creation of digital art. Perhaps most famously, This Person Does Not Exist utilizes a GAN to produce high-resolution, highly realistic images of faces [9]; other creatives have employed GANs for font generation [4], large-scale, dynamic installations [10], and more [11] [8]. None to the author's knowledge have used newspapers in particular, though work exists in applying deep learning to analyze the layout of print documents [7].

This project expands upon this existing, rapidly growing body of GAN art, but also pays homage to the highly sensory, physical associations we have with newspapers by presenting the final artwork in printed form (a zine). Similarly, this circularity of training a GAN on images of physical things and returning its generated, digital output to the physical realm can also be seen in, for example, works by the Obvious collective such as Nike By You [17], in which a collection of Nike sneakers was manufactured based on the output of a GAN trained on Nike sneaker images, and Facets of AGI [16], in which wooden African masks were produced based on the output of a GAN trained on images of African masks.

There has additionally been some work in using creative coding frameworks like Processing to design zines [18].

3 Dataset and Features

We used images from the Newspaper Navigator dataset from the Library of Congress [1][2]. The dataset comprises scans of 16,358,041 historic newspaper pages.

Our final dataset consists of 1081 greyscale images, each with dimensions 1024x1024. We chose to only use the front pages of newspapers for which their publishers published their first issues after January 1, 1900 – headlines didn't really become prominent on the front pages of newspapers until the turn of the century, and we felt the headline comprises an integral part of the character and visual signature of a modern newspaper. The Newspaper Navigator dataset had 1123 unique publishers which began publishing after January 1, 1900; we used one front page from each of these publications and were able to successfully retrieve a total of 1081 images.

We subsequently processed these images to be greyscale, cropped off the bottoms of the images (we chose this instead of a center crop so that we could preserve the top of the newspaper, which is arguably the most visually defining part), and resized them down to 1024x1024.

4 Methods

We ran PGAN [12] using the Facebook Research implementation in Pytorch [13]. PGAN, or ProGAN, refers to an architecture by which a GAN is trained to produce successively larger output until the final output resolution is achieved. Our model begins by training for 48,000 iterations to produce outputs of 4x4 resolution (scale 0), then trains for 96,000 iterations to produce outputs of 8x8 resolution (scale 1), and so on, doubling at each scale, until it finishes training at scale 8 for 200,000 iterations and outputs the final 1024x1024 images.

We subsequently generated 100 images from the trained PGAN and hand-picked some of the more intriguing output to incorporate into a zine. For 8 of these images, we found high-resolution, free usage photos online of people reading newspapers and used Adobe Photoshop to splice the generated newspapers into the photos. Zine design was done in Adobe InDesign, and printing was done by Omega Printing in Palo Alto.

5 Experiments

First, we ran PGAN on an initial dataset of 200 images, with the default configuration, which outputs an RGB image. After training to scale 3, we observed the following input:

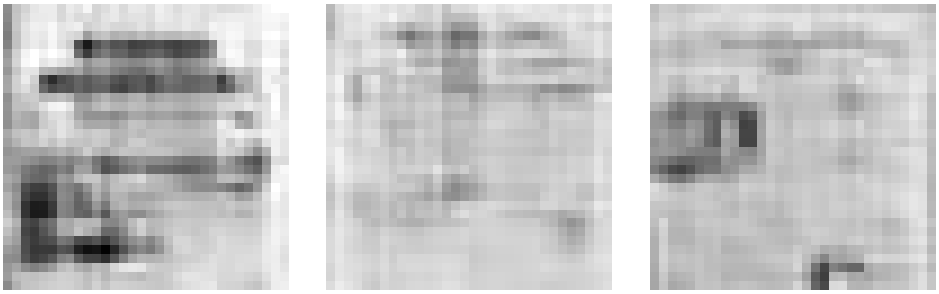

Figure 2: RGB PGAN output on initial dataset of 200, 32x32

This output seemed to be going in the right direction (the generated pages were all black and white, with vaguely rectangular blobs), but we decided to stop training on this dataset at scale 3 (output image dimensions 32x32) because the rest of our dataset (for a total of 1081 images) had been successfully processed.

Our second PGAN model was run on the full dataset, with a smaller number of iterations per scale: we used [1000, 2000, 2000, 2000, 2000, 2000, 2000, 2000, 4800], versus the default max iterations [48000, 96000, 96000, 96000, 96000, 96000, 96000, 96000, 200000]. The output of this model at scale 5 (128x128) is shown below.

Figure 3: RGB PGAN output on full dataset with reduced iterations, 128x128

Evidently, the decreased number of iterations caused PGAN to move onto further scales before previous ones had converged, as we can see by the significant greenish tint and diagonal squiggly patterns. For this reason, we stopped training this model.

After fixing a bug with the Facebook Research PGAN implementation, we configured PGAN to produce greyscale output and ran it with the default number of iterations per scale. This model was our final version, and produced the output we used for our final product.

Due to memory constraints we reduced batch sizes in the final PGAN implementation from the default 16 to 8 for scale 6, 4 for scale 7, and 2 for scale 8.

6 Results/Discussion

Our PGAN model was able to learn many defining features of newspapers – all the generated images contained some semblance of straight, horizontal lines of "text" (the resulting text isn't legible by any means, but is clearly recognizable as text; there are variable-length "words" with spaces between them, and "paragraphs" end with lines of differing lengths), most learned 6 or 7 columns of text, and most learned evenly spaced margins around the top, left, and right of the image (the bottom of the newspapers were cropped so that the inputs could be square). We can look to figure 4 to see a good example of several of these newspaper features – note the neat separation of rectangular "images" from the body text, the subheads in larger, bolder font than the body text, what appear to be brief captions beneath the subheads (center-justified text with the same weight as body text, spaced slightly farther apart), the large, prominent headline / publication name at the top of the page, and even the smaller details at the top of the page – the thin horizontal bar where the publication date would go, the small boxes in the top margin on the left and right where the weather, a slogan, or a brief listing of upcoming events might live (see the top left and right corners of the New York Times or the Hickory Daily Record in figure 1).

Figure 4

As a brief note, we considered but decided against computing an FID score for these results due to time constraints and concerns that the resulting score would be misleading because of our small training set (1081 images; the official Tensorflow implementation [19] recommends at least a 10k sample size).

Despite overall satisfying results, some of the generated images seemed to evidence a slight failure to converge in certain cases, and we observed instances of mode collapse. Some outputs had large black splotches, frequently columns were curvy and distorted,

Figure 5

and occasionally the number of columns was inconsistent, even Escher-esque (note that in figure 5, the top of the page appears to have 7 columns, but at the bottom we count only 6). Others had too few columns, and appeared more like manifestos or brochures than newspapers (figure 6). A close examination of the dataset is needed in order to discern whether this is a failure of the model or similarly evident in the training dataset.

Upon examining 100 images generated from the trained model, we found 2 instances of mode collapse, shown in paired images in figure 7. Even in a small sample of 100 generated images, the model is producing images which are highly similar in the output space. Though this mode collapse isn't debilitating, given that we have an adequate quantity of otherwise diverse outputs, further analysis and experimentation will be necessary to remedy this slight mode collapse in the future.

Figure 7: Mode collapse in generated output

Despite these errors, a fascinating result of this project was the successful illusion of legibility — when presented with these generated newspapers, a small sample of literate human beings were immediately inclined to attempt to read them. Reactions included: "What language is that?", "This almost looks like Urdu, if you squint hard enough", and "It looks like a newspaper someone spilled water on", despite entirely nonsensical "text" and aforementioned failures of convergence (wiggly columns, spooky black streaks, etc.).

Though we initially intended to run PGAN as a baseline and employ a more state-of-the-art StyleGAN2 [14] [15] architecture as our final model, we decided after observing human reactions to the generated images to stick with the PGAN output, because despite its failures, it successfully learned enough of a newspaper's visual features to trick the human eye, and in this way accomplished the task we set out to achieve with flying colors. Just as much this project aims to highlight the power of GANs to create visually meaningful images, it also aims to highlight the limitations of GANs and AI at large in producing meaning without explicit directions towards what is and isn't meaningful. The outputs are, upon close inspection, nonsensical, but at a brief glance are immediately recognizable as newspapers. This delicate balance between coherence and incoherence is what we had sought to begin with, and in this way the PGAN succeeded fully not despite but instead *because of* its failures.

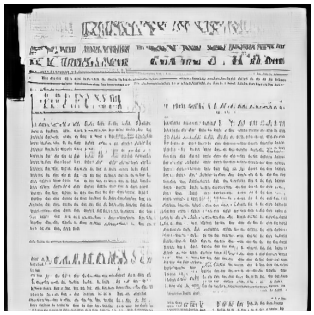

Figure 6

7 Conclusion/Future Work

If we wanted to improve the realistic quality of our output, we could have employed a StyleGAN2 or other, newer, architecture instead, and given more time we would have. A StyleGAN2 model would allow for disentanglement of the latent space, which might shed interesting insight into what the features the GAN is learning. We would then be able to compare those learned features with print designers' long-held intuitions and traditions. We could also attempt to remedy the divergent loss (see loss figures in the Appendix) of our existing PGAN model via "unrolling", in which the discriminator is only updated after k updates of the generator.

The outputs of the project can be found at https://mirandali707.github.io/nonsense_newspapers/index.html. The initial batch of 100 generated images (no handpicking) can be viewed in the gallery, and the 8 handpicked images can be found alongside their corresponding painstakingly Photoshopped edits in the digital zine. We are in the process of working with a printing press to turn the digital zine into a small batch of physical ones, printed on newsprint just as any bona fide newspaper should be.

8 Acknowledgements

We would like to sincerely thank Sharon Zhou for her valuable guidance throughout the entire process, as well as the entire CS230 teaching team. We'd also like to thank Benjamin Charles Germain Lee for his advice and guidance on data preprocessing. We are also grateful for Jesselyn Hernandez at Omega Printing for aiding in the production of the printed zine, and for the various friends who were subjected to texts of unsolicited fake news(papers).

References

[1] LibraryOfCongress. "LibraryOfCongress/Newspaper-Navigator." GitHub, 12 Sept. 2020, github.com/LibraryOfCongress/newspaper-navigator. Accessed 29 Sept. 2020.

[2] "Newspaper Navigator | Experiments | Work | Library of Congress." The Library of Congress, 2015, news-navigator.labs.loc.gov/. Accessed 29 Sept. 2020.

[3] Karras, Tero, Samuli Laine, and Timo Aila. "A style-based generator architecture for generative adversarial networks." Proceedings of the IEEE conference on computer vision and pattern recognition. 2019.

[4] Art. "AIfont: AI-Generated Typeface." Process — Studio for Art and Design, 10 Oct. 2018, process.studio/works/aifont-ai-generated-typeface/. Accessed 30 Sept. 2020.

[5] "Alternative Ways of Writing English." Omniglot.Com, 2020, www.omniglot.com/conscripts/english.htm#english. Accessed 30 Sept. 2020.

[6] Ghiasi, Golnaz, et al. "Exploring the structure of a real-time, arbitrary neural artistic stylization network." arXiv preprint arXiv:1705.06830 (2017).

[7] A. Antonacopoulos, C. Clausner, C. Papadopoulos and S. Pletschacher, "Historical Document Layout Analysis Competition," 2011 International Conference on Document Analysis and Recognition, Beijing, 2011, pp. 1516-1520, doi: 10.1109/ICDAR.2011.301.

[8] E. Wallner, Mlart.Co, 2020, mlart.co/. Accessed 29 Oct. 2020.

[9] "This Person Does Not Exist." Thispersondoesnotexist.Com, 2019, thispersondoesnotexist.com/. Accessed 29 Oct. 2020.

[10] Refikanadol.Com, 2019, refikanadol.com/. Accessed 29 Oct. 2020.

[11] Golan Levin and Collaborators. "Flong - Interactive Art by Golan Levin and Collaborators." Flong.Com, 2020, www.flong.com/. Accessed 29 Oct. 2020.

[12] Karras, Tero, et al. "Progressive Growing of GANs for Improved Quality, Stability, and Variation." ArXiv.Org, 2017, arxiv.org/abs/1710.10196. Accessed 29 Oct. 2020.

[13] facebookresearch. "Facebookresearch/Pytorch_GAN_zoo." GitHub, 14 May 2020, github.com/facebookresearch/pytorch_GAN_zoo. Accessed 29 Oct. 2020.

[14] Karras, Tero, et al. "Analyzing and Improving the Image Quality of StyleGAN." ArXiv.Org, 2019, arxiv.org/abs/1912.04958. Accessed 29 Oct. 2020.

[15] NVlabs. "NVlabs/Stylegan2." GitHub, 8 Oct. 2020, github.com/NVlabs/stylegan2. Accessed 29 Oct. 2020.

[16] “Facets of AGI – Obvious.” Obvious-Art.com, 2019, obvious-art.com/african-mask/. Accessed 14 Nov. 2020.

[17] “Obvious x Nike – Obvious.” Obvious-Art.com, 2019, obvious-art.com/nikebyyou/. Accessed 14 Nov. 2020.

[18] Grajewski, Malte. “Behance.” Behance.net, 2019, www.behance.net/gallery/82382419/thesis-project-robozine-a-generative-magazine. Accessed 14 Nov. 2020.

[19] bioinf-jku. “Bioinf-Jku/TTUR.” GitHub, 29 Sept. 2020, github.com/bioinf-jku/TTUR. Accessed 16 Nov. 2020.

9 Appendix