At the Intersection of Conceptual Art and Deep Learning: The End of Signature

Divya Shanmugam*
MIT CSAIL
Cambridge, MA 02139
divyas@mit.edu

Kathleen M Lewis*
MIT CSAIL
Cambridge, MA 02139
kmlewis@mit.edu

Jose Javier Gonzalez Ortiz*
MIT CSAIL
Cambridge, MA 02139
josejg@mit.edu

Agnieszka Kurant
Tanya Bonakdar Gallery
Manhattan, NY 10011
agnieszka.kurant@gmail.com

John Guttag
MIT CSAIL
Cambridge, MA 02139
guttag@mit.edu

Figure 1: Two newly constructed buildings display “collective signatures" as large-scale LED and neon steel sculptures. Each sculpture is a synthetic signature generated from a large collection of real signatures. Agnieszka Kurant, *The End of Signature*, 2021 MIT Collection. Commissioned with MIT Percent-for-Art funds, and produced in collaboration with a team of computer scientists. Photos: Charles Mayer Photography, courtesy of the artist and the MIT List Visual Arts Center.

Abstract

MIT wanted to commission a large scale artwork that would serve to “illuminate a new campus gateway, inaugurate a space of exchange between MIT and Cambridge, and inspire our students, faculty, visitors, and the surrounding community to engage with art in new ways and to have art be part of their daily lives.”[2] Among other things, the art was to reflect the fact that scientific discovery is often the result of many individual contributions, both acknowledged and unacknowledged. This paper details a collaboration between a widely-exhibited artist, computer scientists, and the broader community to produce a set of collective signatures. After collecting signatures from two communities—the university, and the surrounding city—the computer scientists developed generative models and a human-in-the-loop feedback process to work with the artist to create an original signature-like structure representative of each community. These signatures are now large-scale steel, LED, and neon light sculptures that appear to sign two new buildings in Cambridge, MA.

*Denotes equal contribution

[2]Paul Ha, Director MIT List Visual Arts Center

Workshop on Broadening Research Collaborations in ML (NeurIPS 2022).

1 Introduction

In the *End of Signature* (EOS) series of work, conceptual artist Agnieszka Kurant explores *collective intelligence*. The work discussed here is intended to reflect the fact that scientific discoveries are often built on the labor of multiple generations of scientists, teams of scientists, or several scientists working simultaneously on the same subject in various parts of the globe. EOS evokes this concept with two different collective signatures realized as large-scale animated light sculptures that appear to sign and re-sign two new buildings in the Kendall Square area of Cambridge (Figure 1). One signature represents the broad population of Cambridge, and the other the MIT community. By signing an MIT building with the collective signature of many scientists, academics, and students from MIT (including signatures from earlier generations of MIT scientists), the artwork honors not only well-known scientists but also anonymous interns, uncredited students, post-docs, and members of the community in which they worked. Furthermore, signing a second building with the collective signatures of members of the community in which MIT sits, honors the importance of the influence of the community on the university. These signatures are also intended to highlight the growing influence of technology on art and society; as handwritten signatures are replaced by digital signatures, the meaning of these sculptures will evolve.

Kurant believes that the meaning of authorship in artistic creation is evolving to become a "complex, hybrid authorship," –like other forms of collective intelligence found in science, nature, and online communities. This belief led her to produce EOS in close collaboration with a team of computer scientists.

In keeping with her goal of honoring the science produced at MIT and her continuing interest in the intersection of human and artificial intelligence, the artist decided to explore the use of generative models to produce artwork. She worked with closely with a faculty member and three doctoral students at MIT CSAIL. Working within the artist's conceptual ideas and building off her previous work [1], the computer scientists designed a neural network to learn how to combine signatures without copying any single signature. This turned out to be more challenging in practice than in concept. The experience highlights the importance of *shared language* and *conceptual flexibility* in collaborations between the machine learning community and artists.

2 The Machine Learning Pipeline

The final collective signatures were the result of a multi-stage collaboration between the artist, members of the community, and the computer scientists. First, the artist worked with the university to collect a dataset of roughly 1000 signatures, using pen and paper (Figure 3). The computer scientists pre-processed the manually collected signatures to format them for input to a neural network (Sec. 2.1). The artist and computer scientists then developed a human-in-the-loop method to combine creative vision of the artist with ML models capable of generating evocative images. Generative models have been shown to produce photorealistic images in specific domains [2], and they were a natural choice to include in our approach. We experimented with several generative models and decided on a WGAN [3]. We built an interactive pipeline that would allow us to incorporate the artist's judgement as the WGAN's critic evolved. A key component of this pipeline was a visualization tool that could be used to visualize the impact of various changes to the loss function for the GAN's discriminator. Our human-in-the-loop process is outlined below and expanded upon in future sections:

1. Preprocess the collected signatures and create a community-specific dataset.
2. Train for a variable number of epochs, producing various samples at each one. We were concerned by the possibility of over-fitting and specifically, the possibility that the generative model simply memorized examples in the training set. To ensure that no generated example too closely resembled any single individual's signature, we used a divergence metric to test for similarity between a generated image and the training set.
3. The artist used the visualization tool to note which samples she liked and which she did not.
4. The computer scientists would use this feedback to adjust the architecture and loss function. Steps 2-4 would repeat until the WGAN was consistently producing satisfactory signatures.
5. The computer scientists post-processed the chosen images so that they were suitable for manufacturing.

6. The artist was then presented with two final sets of signatures, one for each community. She then chose one from each set to be manufactured.

2.1 Preprocessing

Volunteers signed within a box on a paper form; the forms were then scanned into the computer. This method posed some challenges including faint writing, different stroke widths and colors of signatures (depending on whether a pen, pencil, or marker was used), and signatures that extended beyond the box on the form. We pre-processed images with thin strokes by applying a dilation filter, to normalize the stroke width of the signatures across the dataset. Additionally, we applied hysteresis thresholding and median filtering to standardize stroke intensity. Training the WGAN without these pre-processing steps led to excessively blurry images.

2.2 Creating Community-Weighted Dataset

Datasets used to train WGANs typically have thousands of examples [2]. We had access to far fewer examples, which posed a barrier to realizing the project's vision for signatures specialized to each community. In order to remain faithful to this goal, the computer scientists, in collaboration with the artist, chose to train both models on signatures from both communities, and up-weight examples from one community. Though not part of the original concept, this collaborative decision led to a new vision in which each collective signature is both informed by the whole, and specialized to individuals from a particular community.

2.3 The Generative Model

We explored multiple model classes including a variational autoencoder (VAE) [4] and a recurrent neural network (Draw-RNN) [5]. We began with a VAE [4], but while signature-specific structures appeared in samples from the latent space, the VAE failed to capture the finer structures. Seeking an approach validated on learning latent-spaces for handwritten letters, we experimented with Draw-RNN [5], which applies a sequential variational autoencoder framework to mimic the reconstruction of letters. Ultimately, we decided on the WGAN because the samples produced exhibited the fine lines and flourishes characteristic of signatures. Examples from the alternate architectures are included in Figure 4.

Generative Adversarial Networks (GANs) consist of two models, termed the generator and the critic, which are optimized jointly in a zero-sum adversarial game [6]. The generator produces synthetic images, while the critic attempts to classify images as synthetic or real. Each model is trained in competition with the other: the generator attempts to produce images that the critic is unable to distinguish as synthetic, and the critic attempts to differentiate between generated images and real images. Wasserstein GANs (WGAN) improve upon this by introducing a different loss function with smoother gradients, which often leads to more stable training [3]. We implement the generator as a 7-layer neural network, where each layer is a 2D convolution paired with a Leaky-ReLU activation function, followed by an upsampling factor of 2. Similarly, we implement the critic as a 7-layer neural network, where each layer is a 2D convolution. We follow [7] and train the critic using a loss function with two terms. The first refers to the difference in loss between the synthetic and real images, and the second is a gradient penalty that has been shown to stabilize training. The code will be made public.

Many of the WGAN design choices were made in response to the artist's feedback. Specifically, the computer scientists experimented with the size of the latent space and kernel size in response to the artist's feedback. We found that smaller latent space sizes produced higher quality images. In contrast to traditional choices for latent space dimensions (256-512), we found that a latent space size of 5 produced the least blurry and qualitatively useful images. This could be explained by the relative lack of complexity in black and white signatures compared to images of, for example, faces. The kernel size governs the receptive field of parameters in the discriminator. Smaller kernel sizes translate to fewer parameters and can lead the critic to focus on more granular features. In response to the artist's preference for thinner generated strokes, the computer scientists settled on a small kernel size.

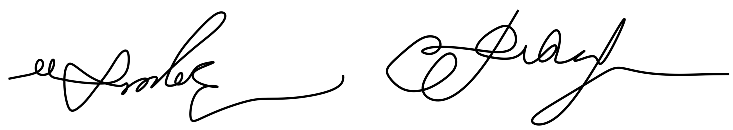

(a) MIT signature (b) Cambridge signature

Figure 2: Final signatures produced by our process.

2.4 Post-processing

While the WGAN produced a static image, the actual sculpture needed to be dynamic. There were also manufacturing constraints that had to be respected. The post-processing step mapped WGAN samples to manufacturable mock-ups of a continuous signature. Outputs of the WGAN were post-processed by fitting a b-spline curve to hand-annotated anchor points within the signature. The computer scientists chose the anchor points in close collaboration with the artist, and found the use of interactive tools helpful in facilitating the collaboration.

3 The Final Product

The artist chose two final signatures representing the university community and the city community. The signatures are simultaneously visible from a large public plaza adjacent to the Kendall Square T stop in Cambridge, MA.

The city signature was installed in 2021 on a canopy covering the entrance of a graduate dormitory and faces downwards (Figure 1). It is made of 18mm coated cobalt blue neon tubing, magnetic neon transformers and a custom programmed controller. The dimensions are 193 x 580 inches.

The university signature (Figure 3a) was installed in the spring of 2022 on a building housing commercial laboratory spaces. It is suspended from an overhang, and is made of LEDs, acrylic lens, steel, paint, and a custom programmed controller. It appears to sign itself in either black or white roughly once a minute. The dimensions are 234 x 697 x 16 inches.

4 Reflections

The experience was both enjoyable and educational, but it was not without bumps.

Early on, we discovered that there was a considerable *terminology gap* separating the artist and the computer scientists. The computer scientists struggled to map the artist's comments on the aesthetics of the images generated by the WGAN to something that could be tangibly represented as part of the model optimization. The various models produced many results. Batches of results were sent to the artist with a request to note which were most appropriate for this project. The artist expressed a strong preference for the ones that had more diverse shape, with lines going down and other lines going up. The preference was then incorporated in the loss function. Near the end of the project the artist was presented with sets of signatures and exercised her discretion as an artist to choose signatures that she felt would have the desired visual effect when enlarged by more than 85 times and mounted as planned on the two buildings. Transparent communication around the capabilities and limitations of machine learning models was crucial to ensuring a productive collaboration.

Another issue was reconciling technical constraints with the original conceptual ideas. For example, it took some time for the artist to conclude that training on a combined data set with community-specific weighting was consistent with her concept for the twin pieces. Discussions during the artistic process helped the computer scientist and the artist find technically feasible solutions within the conceptual ideas. Through this "artist/scientists-in-the-loop" development process, the computer scientists and the artist used both technical and conceptual constraints to develop the final signatures. The process was deeply iterative, and benefited from bespoke tools to facilitate this iteration (i.e., the interactive visual tool, and functionality to determine anchor points). Future work at the intersection of machine learning and creativity would benefit from collaborative tools that prioritize this type of iteration, especially for particularly time-consuming steps in the machine learning pipeline (including data pre-processing and model training).

Since the signatures were installed, it has been interesting to observe the reaction of those viewing them. As of this writing, no explanation of the art is in place. People do seem to understand that they are signature-like, and often speculate on what the signatures are. Talking to people about them is similar to talking to those viewing cumulus clouds or rock formations–viewers impose their own perspective on the signatures–often seeing aspects of their own name in one or the other. The collaboration on *End of Signature* illustrates how machine learning can serve as a provocative tool for artists to provide perspective on technology's societal impact and offers practical insight into how to support such collaborations.

References

[1] *Agnieszka Kurant The End of Signature*, 2015. `https://www.guggenheim.org/artwork/33834`.

[2] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.

[3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv: 1701.07875*, 2017.

[4] Diederik P Kingma and Max Welling. An introduction to variational autoencoders. *arXiv preprint arXiv:1906.02691*, 2019.

[5] Karol Gregor, Ivo Danihelka, Alex Graves, Danilo Jimenez Rezende, and Daan Wierstra. Draw: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015.

[6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

[7] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. *arXiv preprint arXiv:1704.00028*, 2017.

A Supplementary Material

Checklist

1. Include example signature form in the appendix
2. Include discussion of the consent users gave

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section discuss privacy concerns
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]

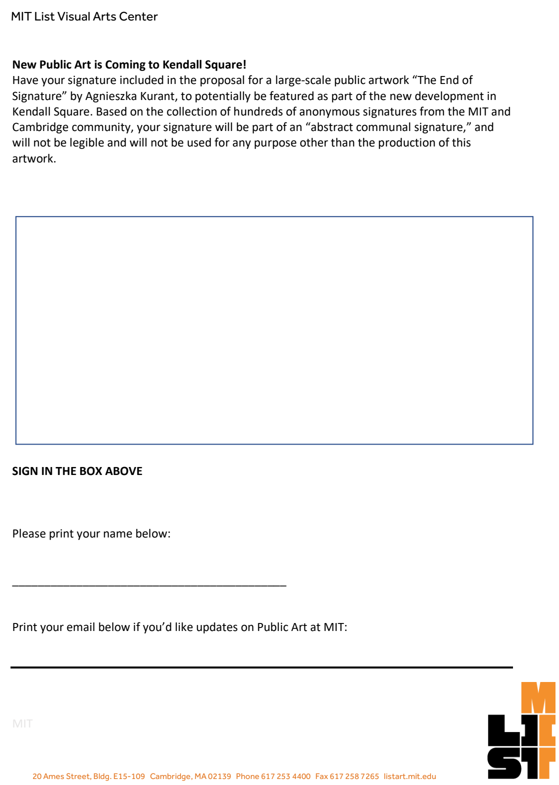
MIT List Visual Arts Center

New Public Art is Coming to Kendall Square!
Have your signature included in the proposal for a large-scale public artwork "The End of Signature" by Agnieszka Kurant, to potentially be featured as part of the new development in Kendall Square. Based on the collection of hundreds of anonymous signatures from the MIT and Cambridge community, your signature will be part of an "abstract communal signature," and will not be legible and will not be used for any purpose other than the production of this artwork.

SIGN IN THE BOX ABOVE

Please print your name below:

____

Print your email below if you'd like updates on Public Art at MIT:

____

MIT

20 Ames Street, Bldg. E15-109 Cambridge, MA 02139 Phone 617 253 4400 Fax 617 258 7265 listart.mit.edu

Figure 3: The paper form used to collect signatures from the MIT and Cambridge communities.

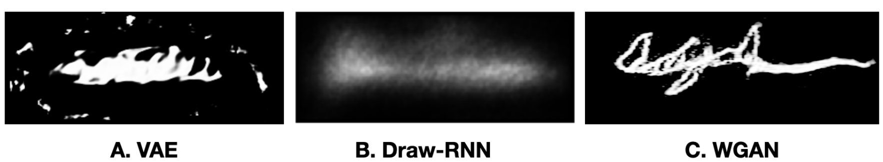

Figure 4: Comparison of signatures generated by three different model architectures.

(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]

(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A] Compute resources and GPU types are not critical to our central claims on broadening research collaborations with artists.

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...

    (a) If your work uses existing assets, did you cite the creators? [N/A]

    (b) Did you mention the license of the assets? [N/A]

    (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

    (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] See Section

    (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] See Section

5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [Yes] See Section
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [Yes] See Section
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]