Unsupervised Doodling and Painting with Improved SPIRAL

John Mellor   Eunbyung Park   Yaroslav Ganin   Igor Babuschkin   Tejas Kulkarni

Dan Rosenbaum   Andy Ballard   Theophane Weber   Oriol Vinyals   S. M. Ali Eslami

Abstract

We investigate using reinforcement learning agents as generative models of images (Ganin et al., 2018). A *generative agent* controls a simulated painting environment, and is trained with rewards provided by a discriminator network simultaneously trained to assess the realism of the agent's samples. Compared to prior work, we make a number of improvements to the architectures of the agents and discriminators that lead to intriguing and at times surprising results. We find that when sufficiently constrained, generative agents can learn to produce images with a degree of visual abstraction, despite having only ever seen real photographs and no trajectories. And given enough time with the painting environment, they can produce images with considerable realism. These results show that, under the right circumstances, some aspects of human drawing can emerge from simulated embodiment, without the need for external supervision, imitation or social cues. See `https://learning-to-paint.github.io` for further analysis and videos.

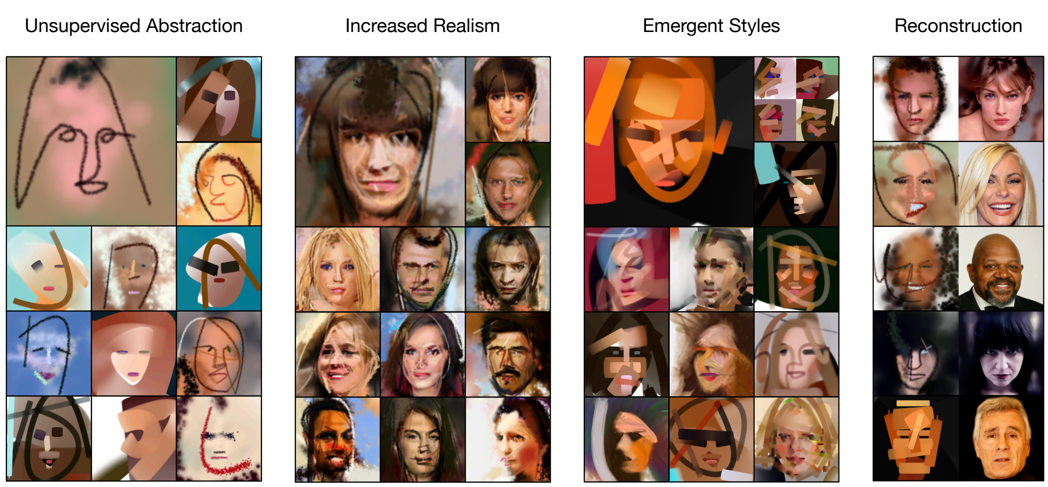

Figure 1: Despite not seeing examples of human drawings, agents can produce images with a degree of visual abstraction, in a diverse array of styles, and they can scale to approach realistic results.

Humans have been representing and reconstructing their visual sensations, using physical tools to draw and sculpt depictions of those sensations, for at least 60,000 years (Hoffmann et al., 2018), well before the development of symbolic writing systems. And abstraction from raw sensations and use of tools are thought to be essential components of certain human cognitive abilities (Clark and Chalmers, 1998; Lake et al., 2015, 2017). We set out to create generative models that similarly use physical

Figure 2: Agents with varying hyper-parameters were trained to generate images (a, b, c) in 17 steps and (d) in 1000 steps. (e) 19-step agent from Ganin et al. (2018) for comparison.

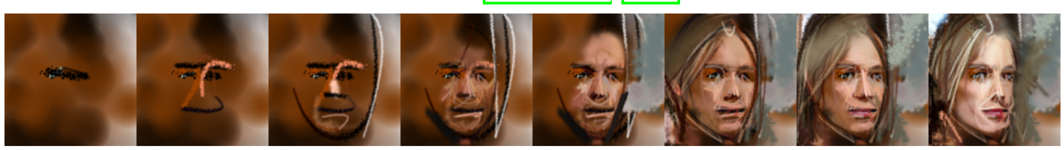

Figure 3: Agents make good use of hundreds of timesteps, controlling the brush with precision. See `https://learning-to-paint.github.io` for videos and further analysis.

grounding. In this work, we equip artificial agents with digital brushes, pens and spray cans. We train these agents with reinforcement learning to interact with digital painting environments (Renold, 2004; Li, 2017). The agents act by placing strokes on a simulated canvas and changing the brush size, pressure and colour as they do so. Building on Ganin et al. (2018), we consider a setting where the agents' rewards are specified by jointly-trained adversarial discriminator networks. In contrast to Ganin et al. (2018) we use *1.* the non-saturating GAN objective, *2.* spectral normalisation in the discriminator, *3.* automatic reward redistribution, *4.* a new form of conditioning, *5.* a spline-based version of the action space and *6.* a shared population discriminator (see appendix B for full details).

We show highlights of our experiments on Celeba-HQ (Karras et al., 2017) in Figure 1 and Figure 2. Amongst all the framework's settings, the one we observed to have the most profound impact on the agents' behaviour was the number of times they were allowed to interact with the canvas in each episode to generate an image. Agents that were constrained with short episodes learned qualitatively different policies than those that could afford numerous interactions, producing images with a degree of visual abstraction, not unlike how humans do when similarly constrained (Selim, 2018; Fan et al., 2019). SPIRAL agents (Ganin et al., 2018) only learn to perform meaningful actions in the last 20 to 50 steps or so of the episode, regardless of the episode length. However, our improved SPIRAL++ agents make full use of episodes of up to 1000 steps, leading in turn to qualitatively different generation policies (Figure 3). See the appendix for extensive experiments and ablations.

To the best of our knowledge, this is the first time that an unsupervised artificial system has discovered visual abstractions of this kind, abstractions that resemble those made by children and novice illustrators. And whilst techniques such as GANs (Goodfellow et al., 2014) and style transfer (Gatys et al., 2015) are increasingly being used for creative investigation in the artistic community, generative agents have the ability to generate novel aesthetic styles with no human input except for a choice of brush (see *e.g.* Figures 25–30 in the appendix), potentially enabling new modes of artistic output. The extent to which any of these systems can be considered *creative* is open for discussion (Hertzmann, 2018), however we hope this framework can provide new points for consideration in that debate.

Acknowledgements We would like to thank Aishwarya Agrawal, Nicolas Heess, Simon Kohl and Adam Kosiorek for insightful discussions and comments on this manuscript.

References

Agrawal, A., M. Malinowski, F. Hill, S. M. A. Eslami, O. Vinyals, and T. Kulkarni
2018. Generating diverse programs with instruction conditioned reinforced adversarial learning. *arXiv preprint arXiv:1812.00898*.

Brock, A., J. Donahue, and K. Simonyan
2019. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*.

Clark, A. and D. Chalmers
1998. The extended mind. *analysis*, 58(1):7–19.

Cohn, G.
2018. Ai art at christie's sells for $432,500. https://www.nytimes.com/2018/10/25/arts/design/ai-art-sold-christies.html.

El Helou, M., S. Mandt, A. Krause, and P. Beardsley
2019. Mobile robotic painting of texture. Technical report.

El-Nouby, A., S. Sharma, H. Schulz, D. Hjelm, L. E. Asri, S. E. Kahou, Y. Bengio, and G. W. Taylor
2018. Keep drawing it: Iterative language-based image generation and editing. *arXiv preprint arXiv:1811.09845*.

Ellis, K., D. Ritchie, A. Solar-Lezama, and J. Tenenbaum
2018. Learning to infer graphics programs from hand-drawn images. In *Advances in Neural Information Processing Systems*, Pp. 6059–6068.

Eslami, S. A., D. J. Rezende, F. Besse, F. Viola, A. S. Morcos, M. Garnelo, A. Ruderman, A. A. Rusu, I. Danihelka, K. Gregor, et al.
2018. Neural scene representation and rendering. *Science*, 360(6394):1204–1210.

Fan, J., R. Hawkins, M. Wu, and N. Goodman
2019. Pragmatic inference and visual abstraction enable contextual flexibility during visual communication. *arXiv preprint arXiv:1903.04448*.

Fernando, R. et al.
2004. *GPU gems: programming techniques, tips, and tricks for real-time graphics*, volume 590. Addison-Wesley Reading.

Frans, K. and C.-Y. Cheng
2018. Unsupervised image to sequence translation with canvas-drawer networks. *arXiv preprint arXiv:1809.08340*.

Ganin, Y., T. Kulkarni, I. Babuschkin, S. M. A. Eslami, and O. Vinyals
2018. Synthesizing programs for images using reinforced adversarial learning. In *Proceedings of the 35th International Conference on Machine Learning*.

Gatys, L. A., A. S. Ecker, and M. Bethge
2015. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576*.

Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio
2014. Generative adversarial nets. In *Advances in neural information processing systems*, Pp. 2672–2680.

Gregor, K., I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra
2015. Draw: A recurrent neural network for image generation. In *ICML*.

Guez, A., T. Weber, I. Antonoglou, K. Simonyan, O. Vinyals, D. Wierstra, R. Munos, and D. Silver
2018. Learning to search with MCTSnets. In *Proceedings of the International Conference on Machine Learning (ICML)*.

Gulrajani, I., F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville
2017. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems 30*.

Ha, D. and D. Eck
2017. A neural representation of sketch drawings. *arXiv preprint arXiv:1704.03477*.

Hertzmann, A.
2003. A survey of stroke-based rendering. *IEEE Computer Graphics and Applications*, (4):70–81.

Hertzmann, A.
2018. Can computers create art? In *Arts*, volume 7, P. 18. Multidisciplinary Digital Publishing Institute.

Hoffmann, D. L., C. Standish, M. García-Diez, P. B. Pettitt, J. Milton, J. Zilhão, J. J. Alcolea-González, P. Cantalejo-Duarte, H. Collado, R. De Balbín, et al.
2018. U-th dating of carbonate crusts reveals Neandertal origin of Iberian cave art. *Science*, 359(6378):912–915.

Huang, Z., W. Heng, and S. Zhou
2019. Learning to paint with model-based deep reinforcement learning. *arXiv preprint arXiv:1903.04411*.

Jaderberg, M., V. Dalibard, S. Osindero, W. M. Czarnecki, J. Donahue, A. Razavi, O. Vinyals, T. Green, I. Dunning, K. Simonyan, et al.
2017. Population based training of neural networks. *arXiv preprint arXiv:1711.09846*.

Kar, A., A. Prakash, M.-Y. Liu, E. Cameracci, J. Yuan, M. Rusiniak, D. Acuna, A. Torralba, and S. Fidler
2019. Meta-sim: Learning to generate synthetic datasets. *arXiv preprint arXiv:1904.11621*.

Karras, T., T. Aila, S. Laine, and J. Lehtinen
2017. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.

Karras, T., S. Laine, and T. Aila
2018. A style-based generator architecture for generative adversarial networks. *arXiv preprint arXiv:1812.04948*.

Kingma, D. P. and J. Ba
2014. Adam: A method for stochastic optimization.

Kingma, D. P. and M. Welling
2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.

Kotani, A. and S. Tellex
2019. Teaching Robots to Draw. In *IEEE International Conference on Robotics and Automation (ICRA)*.

Lake, B. M., R. Salakhutdinov, and J. B. Tenenbaum
2015. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338.

Lake, B. M., T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman
2017. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40.

Lázaro-Gredilla, M., D. Lin, J. S. Guntupalli, and D. George
2018. Beyond imitation: Zero-shot task transfer on robots by learning concepts as cognitive programs. *arXiv preprint arXiv:1812.02788*.

Li, D.
2017. Fluid paint. http://david.li/paint/.

Li, M., Z. Lin, R. Mech, E. Yumer, and D. Ramanan
2019. Photo-sketching: Inferring contour drawings from images. In *WACV 2019: IEEE Winter Conf. on Applications of Computer Vision*.

Liu, Y., Z. Wu, D. Ritchie, W. T. Freeman, J. B. Tenenbaum, and J. Wu
2018. Learning to describe scenes with programs.

Louppe, G. and K. Cranmer
2017. Adversarial variational optimization of non-differentiable simulators. *arXiv preprint arXiv:1707.07113*.

Miyato, T., T. Kataoka, M. Koyama, and Y. Yoshida
2018. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*.

Nakano, R.
2019. Neural painters: A learned differentiable constraint for generating brushstroke paintings. *arXiv preprint arXiv:1904.08410*.

Ng, A. Y., D. Harada, and S. J. Russell
1999. Policy invariance under reward transformations: Theory and application to reward shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning (ICML 1999)*.

Oord, A. v. d., N. Kalchbrenner, and K. Kavukcuoglu
2016. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*.

Pathak, D., P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros
2016. Context encoders: Feature learning by inpainting. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Renold, M.
2004. Mypaint. http://mypaint.org/.

Ruiz, N., S. Schulter, and M. Chandraker
2018. Learning to simulate. *arXiv preprint arXiv:1810.02513*.

Russakovsky, O., J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei
2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252.

Schmidhuber, J.
1997. Low-complexity art. *Leonardo*, 30(2):97–103.

Selim, M.
2018. Spiderman 10 min 1 min 10 sec speed challenge. https://www.youtube.com/watch?v=x9wn633vl_c.

Sutton, R.
2019. The bitter lesson. http://www.incompleteideas.net/IncIdeas/BitterLesson.html.

Welling, M.
2019. Do we still need models or just more data and compute? https://staff.fnwi.uva.nl/m.welling/wp-content/uploads/Model-versus-Data-AI-1.pdf.

Williams, R. J.
1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. In *Reinforcement Learning*. Springer.

Zeng, K., M. Zhao, C. Xiong, and S. C. Zhu
. From image parsing to painterly rendering.

Zheng, N., Y. Jiang, and D. Huang
2018. Strokenet: A neural painting environment.

Zhou, T., C. Fang, Z. Wang, J. Yang, B. Kim, Z. Chen, J. Brandt, and D. Terzopoulos
2018. Learning to sketch with deep q networks and demonstrated strokes. *arXiv preprint arXiv:1810.05977*.