
Botto: A Decentralized Autonomous Artist

Mario Klingemann
BottoDAO
quasimondo@gmail.com

Simon Hudson
BottoDAO
hudsonsims@gmail.com

Zivvy Epstein
MIT Media Lab
Cambridge, MA 02139
zive@mit.edu

Abstract

Botto project¹ is an experiment in creating a decentralized autonomous artist that generates art based on community feedback. Every week, Botto creates and sells artwork via a series of models and crowdsourced evaluations of over 5,000 people who also decide how to manage the artist and its sales. We present a formal description of how Botto works and its implications for creative machine learning.

1 Introduction

For hundreds of years, people have imagined autonomous entities that are able to be creative or co-creative in some form. Examples of machines or contraptions that could be seen as some precursor of the idea of AI can be found throughout literature [Kantosalo et al., 2021]. Today, text-to-image models like Stable Diffusion [Rombach et al., 2022] and Midjourney [Holz, 2022] can produce high-fidelity imagery in a matter of seconds. Popular discourse has framed these models as artistic agents themselves, capable of creating their own artwork [Epstein et al., 2020]. Conversations that take up this framing are linked to concerns that these technologies will replace human jobs [Frank et al., 2019], and the framing been shown to affect perceptions of credit and responsibility for people involved in the creation [Epstein et al., 2020].

Yet, virtually all cases of these generative tools have human artists providing the ideas and filtering out the interesting from the noise and the banal. Botto is an experiment in creating an agent that can autonomously create art works of cultural impact and financial reward without direct human intervention. To date, Botto [Botto, 2021] has sold 42 artworks through weekly auctions, ranging from 7ETH to 100ETH, totalling over 770 ETH in sales (or \$2.47M USD if proceeds were converted at time of sale). To put in perspective, it is the 17th best selling artist of all time on SuperRare, a premier digital auction site for one-of-a-kind NFTs. The project makes use of NFTs, DAOs, and DeFi to sell the work and pay for the labor that provides feedback on Botto’s images, which helps guide its aesthetic development. This labor is influential only in ways that do not violate Botto’s agency and render it the tool of a single human artist, such as by allowing individuals alter the images or prompts. As an experimental community, economy and creator, the Botto project must navigate issues around anthropomorphized AI, machine creativity, governance, ownership and credit.

2 Approach

The Art Engine - The system for generating image “fragments” has several components. The first is a custom prompt generator to mine textual prompts. The prompt generator creates a combination of random words and full sentences to seed image generation. The second stage is a text-to-image model that generates 4000 weekly images from the prompts; currently Botto uses both VQGAN + CLIP [Radford et al., 2021] and Stable Diffusion. The final ingredient is a computational taste model

¹The author Mario Klingemann originally conceived of the concept and designed Botto’s art engine. The author Simon Hudson is one of the operators of the BottoDAO and a contributor to the project’s ongoing design.

for filtering out 350 results to present to the community for feedback via voting. Both the prompt and taste models are trained by the feedback and fit their outputs to a predicted distribution of voting, along with some out-of-distribution outputs to help prevent getting stuck in a niche.

The Botto project is strictly against any “cheating” or human guidance other than voting. That means the prompts are random, there are no seed images of existing real-world images used, and the selection of fragments to present to the community are entirely controlled by the system itself. The only human direction Botto got at the outset was from adding a small amount of pre-curated prompts to the entirely random ones generated by the algorithm to get the engine started. These inputs were not iteratively developed so as to avoid unnecessarily narrowing the latent space available in the generative model it uses.

Community Governance - Holders of the cryptocurrency \$BOTTO are able to vote on the 350 images, picking one of two options presented to them in successive order. Users can stake the strength of their vote based on their \$BOTTO holdings to determine a weekly community winner and influence Botto’s training.

\$BOTTO holders lobby on Discord and Twitter for individual pieces throughout the week, and especially during the final 24 hours to pick a final winner. Broader discussions also take place in the community on how to manage Botto as an artist. These include budgeting decisions, artist and gallery collaborations, and, most critically, how to evolve the protocol of Botto’s art engine and economy. As Botto’s governance is decentralized, it is up to the community to work together to decide how to help Botto evolve over time with new generative models and collaborations without violating its agency.²

A Sustaining Economy - Each week, the selected winning image is minted as an NFT on the Ethereum blockchain and auctioned for the cryptocurrency ETH on premium online auction house SuperRare. The auction start is announced on Twitter from @bottoproject (26k followers to date).

At the outset, proceeds of the sale in ETH were used to buy \$BOTTO on the open market and remove them permanently from circulation so as to increase the value of the remaining tokens. This mechanism functioned to reward the voters for their contribution, thus incentivizing them to vote on visuals that would sell for high values. Since launch, community governance has iterated on this economic policy to weight rewards based on different types of contributions, voting behavior, as well as saving some of the proceeds to pay for servers and ad hoc tasks approved by the protocol members.

3 Discussion

Along with Gene Kogan’s Abraham project [Kogan, 2016, 2019] (conceptualized as an “autonomous artificial artist”), Botto is part of a new class of systems that challenge the notion of authorship. To sustain their own production and pay for the resources they consume, Abraham and Botto use blockchain economies to create, own, and sell their art. The idea of an autonomous artistic agent independent of a human creator is consistent with the notion of the machine condition [Colton et al., 2020], whereby machines can creatively express their existence.

Leveraging the cryptocurrency \$BOTTO, Botto itself is a novel approach to tackling decentralized ownership. By tying the voting process to a stake in the currency and the value of the currency to the success of the project, Botto has found a way for voters to meaningfully gain a share of ownership. However the legal frameworks undergirding this process remains exciting future work [Eshraghian, 2020, Gordon et al., 2022].

The Botto project is an example of how AI art applications can both produce outputs of artistic value while also helping build understanding of AI’s broader context and effects for both researchers and the broader public. Botto’s performance has attracted an audience’s imagination while at the same time calling attention to the underlying mechanical functions at work and their dependence upon human labor. We believe this project can shed light on the themes of social influence, the data labor economy, and collective intelligence. Further, Botto and its corresponding DAO provide a case study highly relevant to the ongoing discussion about digital governance and organization [Frey et al., 2019, Schneider et al., 2021, Seering, 2020, Zhang et al., 2020, Jahani et al., 2018]. Through decentralized

²Developments are under way for the project to be able to iterate on new experiments within the conceptual framework of a decentralized autonomous artist. These include: adding the rapidly emerging generative models coming online and allowing them to compete within a single system, social rewards for contributors, tying financial rewards to the social rewards, and collaborations with other human artists.

governance, the Botto community actively functions as manager of the artist and protocol, which includes developing the economy surrounding Botto to keep it attractive to protocol stewards (with e.g. social awards and differentiated financial rewards) that ensure long-term sustainability.

Appendix: Ethical Considerations

The emergence and popularity of distributed socio-technical systems for artistic creation has raised thorny questions of ownership and authorship. Attempting to frame a machine as having agency can obscure the large human role in the system, and its success can legitimize marketing efforts the recycle this kind of agentic framing. Recent work has shown that these frames can actually undermine perceptions of credit to the *human actors* involved in the system [Epstein et al., 2020]. Further, the project may serve to stoke the anxious concern of AI art eliminating artist jobs. Finally, Ethereum until only very recently has had a large environmental impact with each mint, making an automated minting an unsustainable model of creativity to promote.

References

- Anna Kantosalo, Michael Falk, and Anna Jordanous. Embodiment in 18th century depictions of human-machine co-creativity. *Frontiers in Robotics and AI*, 2021.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, June 2022.
- David Holz. Midjourney documentation. <https://midjourney.gitbook.io/docs/>, 2022.
- Ziv Epstein, Sydney Levine, David G Rand, and Iyad Rahwan. Who gets credit for ai-generated art? *Isience*, 23(9):101515, 2020.
- Morgan R. Frank, David Autor, James E. Bessen, Erik Brynjolfsson, Manuel Cebrian, David J. Deming, Maryann Feldman, Matthew Groh, José Lobo, Esteban Moro, Dashun Wang, Hyejin Youn, and Iyad Rahwan. Toward understanding the impact of artificial intelligence on labor. In *Proceedings of the National Academy of Sciences*, volume 116, pages 6531—6539, 2019.
- Botto. Botto documentation. <https://docs.botto.com/>, 2021.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- Gene Kogan. Artist in the cloud: Towards the summit of ai, art, and autonomy. <https://medium.com/@genekogan/artist-in-thecloud-8384824a75c7>, 2016.
- Gene Kogan. Artist in the cloud: Towards an autonomous artist. *Machine Learning for Creativity and Design Workshop*, 2019.
- Simon Colton, Alison Pease, Christian Guckelsberger, Jon McCormack, and Maria Teresa Llano. On the machine condition and its expression. *Machine Learning for Creativity and Design Workshop*, page 342–349, 2020.
- Jason K Eshraghian. Human ownership of artificial creativity. *Nature Machine Intelligence*, 2(3): 157–160, 2020.
- Skylar Gordon, Robert Mahari, Manaswi Mishra, and Ziv Epstein. Co-creation and ownership for ai radio. *ICCC 2022*, 2022.
- Seth Frey, PM Krafft, and Brian C Keegan. "this place does what it was built for" designing digital institutions for participatory change. In *Proceedings of the ACM on Human-Computer Interaction*, volume 3, page 1–31. CSCW, 2019.

Nathan Schneider, Primavera De Filippi, Seth Frey, Joshua Z Tan, and Amy X Zhang. Modular politics: Toward a governance layer for online communities. In *Proceedings of the ACM on Human-Computer Interaction*, volume 5, page 1–26. CSCW, 2021.

Joseph Seering. Reconsidering community self-moderation: the role of research in supporting community-based models for online content moderation. In *Proceedings of the ACM on Human-Computer Interaction*, volume 4, page 107. CSCW, 2020.

Amy X Zhang, Grant Hugh, , and Michael S Bernstein. Policykit: building governance in online communities. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, page 365–378, 2020.

Eaman Jahani, Peter M Krafft, Yoshihiko Suhara, Esteban Moro, , and Alex Sandy Pentland. Scam-coins, s*** posters, and the search for the next bitcoin: Collective sensemaking in cryptocurrency discussions. In *Proceedings of the ACM on Human-Computer Interaction*, volume 2, page 1–28. CSCW, 2018.