

Published by Inter Arts Center, Lund University

Text © The authors, 2024

Editorial: Hedvig Jalhed and Margot Edström



This text is licensed under an Attribution-NonCommercial-NoDerivatives 4.0 International License. This license allows users to download and share the article for non-commercial purposes, so long as the article is reproduced in the whole without changes, and the original authorship is acknowledged (See full terms and conditions here: <https://creativecommons.org/licenses/by-nc-nd/4.0/>)

DOI: <https://10.37852/oblu.362.c877>

1, 2, 3 Playtime

ISBN:

978-91-531-2612-6 (print)

978-91-531-2613-3 (digital)

Publications from Inter Arts Center, no 1

ISSN: 3035-7586

Suggested citation: de Seta, G. & Øen Nordahl, M. 2024, 'Replicator Report', in M. Edström & H. Jalhed (red.), *1, 2, 3 Playtime*, Inter Arts Center, Malmö.
<https://doi.org/10.37852/oblu.362.c877>



Workshop participants gathering at Aldea (top); 3D model by Andreas Zißler (bottom). Bergen Centre for Electronic Arts (BEK). Bergen, Norway, 2023.

Replicator Report

Gabriele de Seta & Magnhild Øen Nordahl

In early November 2023, we ran a three-day event called *Replicator Workshop: A practical speculation on generative AI for physical production*. This workshop was supported by BEK, the Bergen Centre for Electronic Arts, and co-hosted by the Aldea Center of Contemporary Art, Design and Technology. The main idea behind the *Replicator Workshop* was to explore the emerging intersections between generative artificial intelligence and physical production through speculative research and interdisciplinary collaboration between researchers, artists and curators. Our initial pitch was to gather a small group of dedicated participants, research and experiment with the tools available at that moment and figure out a way to create a ‘common object’ which would function as an anchor for us to understand the process leading from generative AI modelling to material production. Whatever the results, we planned to share them a couple of weeks later, during the yearly BEK symposium *The Only Lasting Truth Is Change*. As this brief report documents, not everything went as planned, which we count as part of our experiment’s success.

Our interest in the relationship between generative AI, 3D-modelling and physical production developed organically from our previous collaborative encounters and individual practices – Gabriele being an academic researcher working on the socio-cultural life of machine learning, and Magnhild being a visual artist working on the sculptural translations of abstraction in everyday life. Throughout 2023, generative AI was rapidly moving from the domains of text and images into the less charted possibilities opened up by multimodal models, which include new pipelines such as image-to-video or text-to-3D synthesis. The alluring promise of being able to prompt a machine learning model with a simple line of text and obtain a 3D-printed object seemed to finally approximate the famous Replicator technology depicted in *Star Trek*: a

technological box capable of creating things on demand. What could we find out if we tried to actually build a Replicator with available models and tools? As often happens, we started by jotting down a set of questions:

What are the current limits of generative AI? How does the seemingly endless variation in digital outputs interface with physical production processes? How close is generative AI to the science fictional fabrication afforded by the Star Trek replicator, and what can we speculate about the near future of generative AI for physical production through currently available tools?

In the weeks running up to the workshop, we did all the reading we could – academic articles, technical documentation, press releases, tech news – and compiled lists of machine learning models and generative AI tools. We put out an open call and were stunned by the amount of applications we received. There was definitely more interest in the topic than we thought, and picking three names to join us was not an easy task. Besides some shared readings to bring us all on the same page, we asked all workshop participants to come to Bergen with a ‘prompt’ that they would hold onto for the three days, a sort of lifeline to avoid losing ourselves as we explored unstable diffusion processes and uncharted latent spaces. Funda Zeynep Ayguler, a media artist based in Germany with a background in animation, architecture and design, brought the sentence ‘machines like rats are part of nature and our lives are embedded within each other’. Berlin-based curator, artist and writer Silver Carls-son brought an image as prompt: a painting by Hans Holbein the Younger. Architect and artist Andreas Zißler brought something physical: a rusted chip of metal from a Richard Serra sculpture, gifted to him by a friend. Magnhild’s prompt of choice was a photo of one of her sculptures which she had used to fine-tune a machine learning model to generate new shapes. And Gabriele brought an image of a seven-hole doughnut, the topological structure of the human body.

After a morning of introductions and a rough idea of where we would like to get by the evening, we started playing around with our individual prompts. Most of the first day was spent around a roundtable with laptops and screens, files and links sent across platforms, many failed attempts and occasional exclamations when this or that model produced an interesting output or glitch. In this phase, the key mode of collaborative work was bricolage: approaching generative AI tools through trial-and-error, downloading and running machine learning models on the fly, skimming documentation,

jumping from free demos on corporate websites to GitHub repositories, from step-by-step walkthroughs to Colab notebooks. This process was hardly systematic, and we struggled to track every experiment and observation – most of it is to figure out what we want to achieve, and which tools could bring us there. For example, Magnhild fed her artwork photo into multiple image-to-text models to obtain a variety of descriptions which she then used to prompt commercial text-to-3D services like Alpha3D or 3DFY. Among dozens of AI-generated objects, a specific shape resulting from the image-to-text generated prompt ‘creative wall clock makes your wall look more stylish’ struck us, as it seemed to repeat with a certain consistency: a flat, organic oval with mysterious holes, a weird topographical surface and a jagged rim. ‘*This is what is actually interesting,*’ Magnhild concluded – not the shape itself nor its



3D printing the swirlvoid. Bergen Centre for Electronic Arts (BEK). Bergen, Norway, 2023.
Photo: Gabriele de Seta & Magnhild Øen Nordahl

aesthetic qualities, but the relationship between initial prompt and final output, the recurrence of a morphological entity that we could not yet explain.

Another approach that emerged as productive was comparison: at first, Andreas stuck to a textual prompt describing the object he brought, generating multiple 3D models with different commercial tools. When he felt like he was starting to get repetitive results, he moved to text-to-image and image-to-3D models to verify if their outputs offered any useful resemblance; then, he changed the prompt, at first just by replacing a few words, then generalising or abstracting it completely ('sculpture'), ending up on combinations of made-up words (like 'gupe grouty virtish thirb'). After Andreas placed each output side by side in a grid, we slowly started to notice recurring patterns and features, as well as formal similarities: 3D objects generated by machine learning models often occupy a cubic volume, never straying too far from this invisible frame; some features, like the jagged edges or a blobby protrusion, seem to be linked in some way to the textual prompt; the generative tools can't really produce truly flat surfaces; and so on. Producing outputs with a fixed input until they become repetitive and then tweaking the input became a sequence of lateral moves similar to the parable of the blind men and the elephant¹ – only, in our case the elephant was a hyper-dimensional manifold. This approach also led us to gather around each other's screens and formulate questions: How do these systems generate 3D objects? Are these models trained on similar datasets? Do they rely on similar natural language processing frameworks to interpret text?

By the second day, we had consolidated some provisional intuitions. First, each generative 3D tool left a very recognisable imprint on its outputs – some commercial ones were clearly fine-tuned to create standardised objects like sofas or shoes, and could not process other object categories; others were less predictable and seemed to veer towards unrecognisable things unless prompted with very clear instructions. Second, the capacity of generative AI systems to produce 3D models was rather limited and, in Andreas' words, *uncreative*: the outputs were often boxed in, repetitive, approximative, scaleless and lacking the captivating weirdness of 2D image generation. As we moved from digital objects to physical production, something else started becoming evident: the 3D models we had generated were hardly ready to be manufactu-

¹ This parable, documented as early as 500 BCE India and recurring in Buddhist and Hindu religious texts, revolves around a group of blind men encountering an elephant for the first time, and trying to figure out its shape by only touching a part of it – obviously, the result is that they all gather a different impression of the elephant.

red. Some of them, like the ones that Funda created with the OpenAI Shap-E model, were very interesting in terms of shapes, but they were not stable structurally, and would break up and crumble during the process of 3D-printing. While we started printing some of the simpler objects on Aldea's Prusa machines, Funda spent hours retouching her 3D models of organic rat-shaped machines to make them printable, adding joints between floating parts, supports under heavy jutting structures, and reinforcing brittle connections. In some cases, even this was not enough: after hours spent printing them, a couple of Funda's models still fell apart while we cleaned them up from the 3D-printing support structures. 'All these objects have horrible geometries, they are very inefficient,' she concluded.

As we moved to the third day of the workshop, we found ourselves straddling a widening gap between the many models we generated and the tools we had available to manufacture them. In order to approximate a Replicator, a lot of bridging work was necessary. This included: format conversions, model cleaning and resizing, object orientation and slicing, material testing, and even more mundanely, someone watching over the 3D printers in case they got clogged or misaligned. All this work, which is often hidden from the promotional narratives of generative AI, happened through established tools such as 3D-modelling software and 3D printer utilities. And while these procedures, which took a large part of our final workshop hours, are fundamental for bringing digital objects into the physical world, they are also key decision points, as every step requires making choices that lie outside of the capacities of current generative AI: How big is this object? How is it actually oriented in space? Which kind of support structure would work best to print it? What kind of filling would make it less brittle? How can material be saved while maintaining structural integrity? Magnhild's attempt at conjoining two of our printed objects by designing a third, flexible shape, exemplifies this complex, layered process of human-driven analysis and decision-making:

In order to connect the two objects prompted with 'Time' on Alpha3D and MasterpieceX, I used the new AI-assisted automated modelling 'connect two faces' function in the Fusion360 software. But in order for it to work, I had to simplify the meshes of the time-objects, since they were too complex for the software to deal with. The one that looks like a clock, for example, looks like a low-poly object but when you zoom into it, you realise that it actually has a very complex wireframe, which is suboptimal.

Before wrapping up and heading out for a celebratory dinner, we gathered all the objects we managed to print together for a group photo, and then decided to try one last experiment: we took individual photos of each object, and shared them with ChatGPT, asking it to come up with an appropriate name for it. All of the names ChatGPT came up with – the Spiralock, Ser-raknife, Squiggleblock, Frostone, Swirlovoid, Froggler, Holecube – could be the beginning of a story of its own, demonstrating how the recursive feeding back of outputs into inputs, from one generative model to the other, through different modes of physical and digital mediation, emerged as a productive research method. Did we manage to build a Replicator? If we take the canonical Star Trek device as a benchmark, the answer is clearly no. But by forcing ourselves to move from open-ended experimentation with generative AI for 3D modelling to the material constraints of physical production, we succeeded in getting a sense of the state of the field in late 2023, accumulating several hypotheses, insights and conclusions about how these systems work, where they fail, and what kind of bridging work their coordination requires.

Two weeks after the *Replicator Workshop*, we presented a draft of this report at the annual BEK symposium *The Only Lasting Truth is Change*. While we showed photos of our process and shared our experiences, the audience was handed our 3D-printed objects to touch, inspect and pass around. Our conclusion, on that day, was that the dream of machines capable of manufacturing almost anything a human desires remained steadfastly in the realm of science fiction; and yet, through our exploration of currently available technologies, we could at least offer some suggestions for how to approach generative 3D modelling for productive speculation about its near future. Our work on the Replicator is ongoing: in March 2024, Gabriele led a condensed version of the workshop for the *1, 2, 3 Playtime* project at Inter Arts Center, Lund University, encouraging participants to devise experimental approaches to generative AI. In October 2024, a selection of the objects we manufactured, our *Replicator Probes*, was exhibited at the *More Than Meets AI* exhibition in Bergen, complete with identification labels detailing their originating prompt, the model that created them, and their ChatGPT-generated name.



Replicator Workshop: A practical speculation on generative AI for physical production. 3D printed results. Bergen Centre for Electronic Arts (BEK). Bergen, Norway, 2023.