Computational Vision & Acts of Imagination

Katja Novitskova in conversation with Nora N. Khan
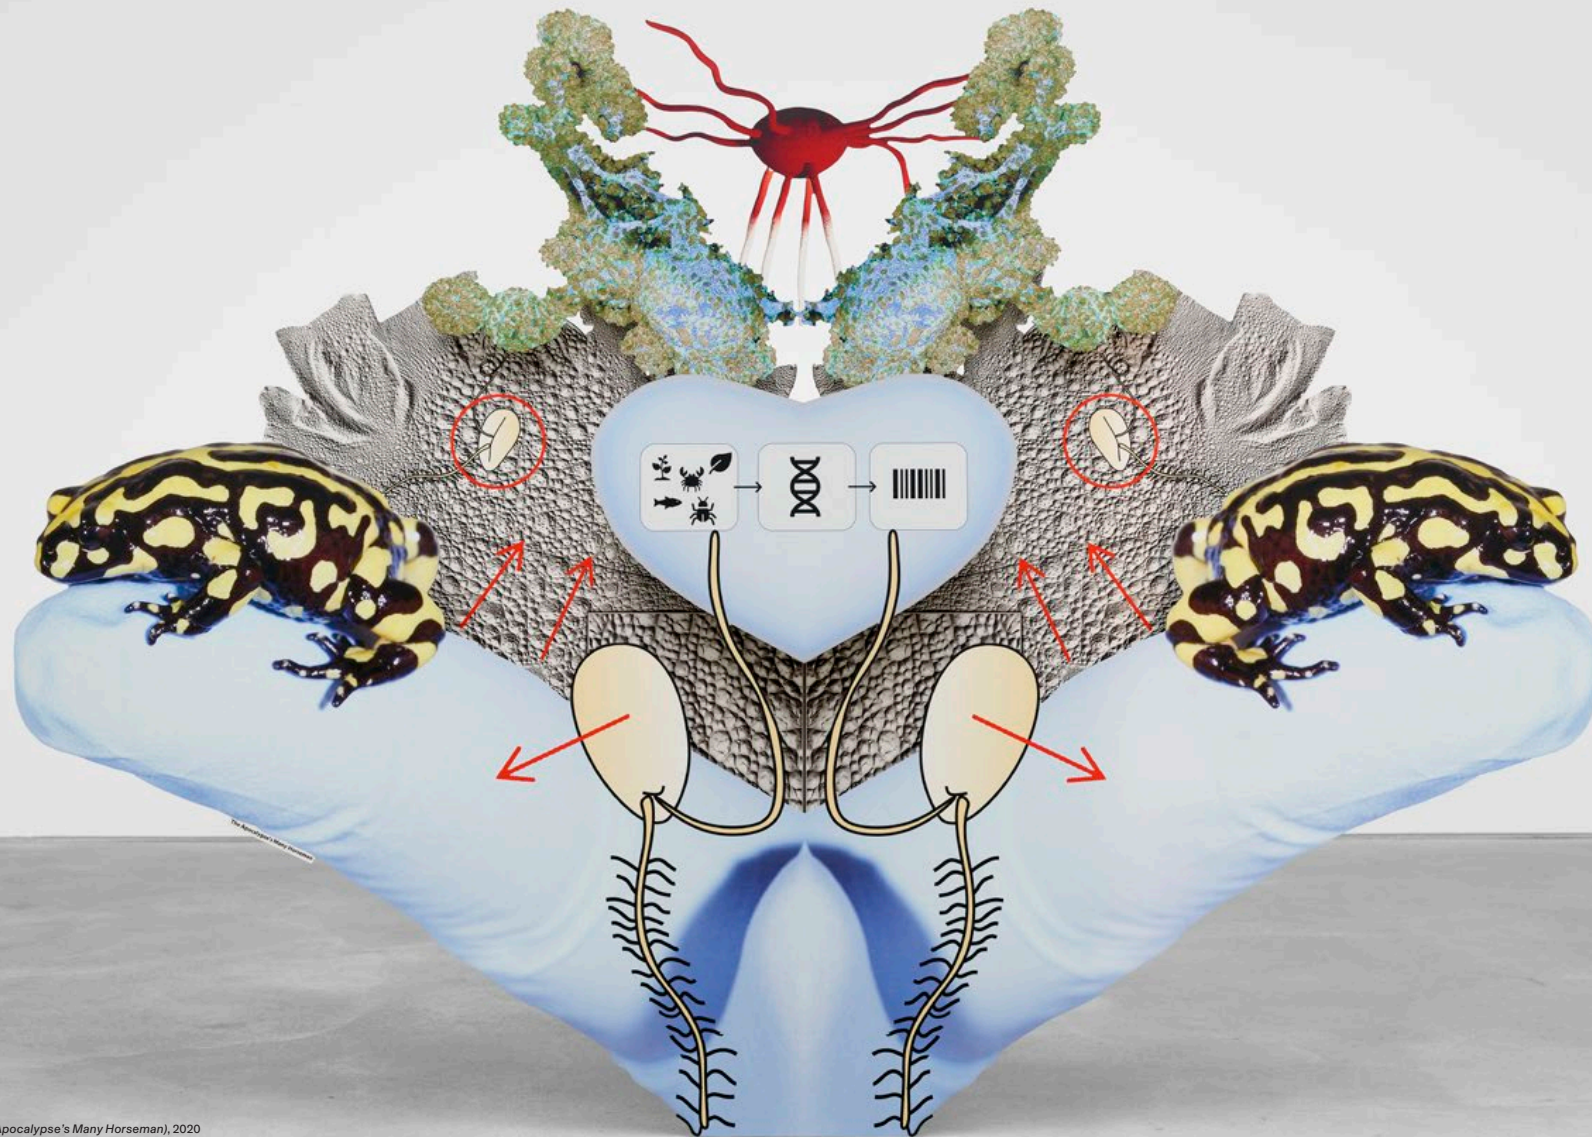
Mamaroo (Eggs), 2017
Electronic baby swing, silicon stress eggs, plastic hoses, lasers, epoxy clay, digital print, polyurethane resin, aluminium folding stand
Courtesy the artist and Kraupa-Tuskany Zeidler. Photo: def image

I first spoke with Katja Novitskova ten years ago, when I was a critic just starting out. Her research and work with strange images, soft images – the eyes of animals caught at night, cryptic texts and diagrams culled from the depth of the internet – captured and held me. I struggled to find the right words to describe her work. Over the long course of her practice and career, she has consistently investigated the impact of machine vision, machine recognition, and generative models, as each of these systems of capture and prediction have evolved, by paying attention, in large part, to their artifacts. Belugas look curiously at underwater camera drones, deer kissing in the woods: Katja turns us to focus on the eye, visage, and interiors of nonhuman beings, mediated many times over. In many of her works, we're invited to consider animals' and machines' inner worlds, to approximate their interiority, and also consider our continual translation, partial interpretations, and guesswork about the world around us. I hoped to speak to her as she prepares for a new show a decade later, on the whole arc of her practice. Our conversation launches from her core curiosity about the encounter of nonhuman and machine as extension of the human, out towards the whole world of her approximations.

Installation view: exhibition *Random Forest*, Temnikova & Kasela Gallery, Tallinn, 2023
Front: *Soft Approximation (Looking Glass Deers Kissing 03)*, 2023
Polyurethane resin, UTR-8100 resin, water-based pigments, powder pigments, carnelian, quartz, aluminium wire

Courtesy the artist and Temnikova & Kasela Gallery. Photo: Stanislav Stepaško

"The end of a civilization …", 2019
UV-resistant ink transfer on synthetic clay, epoxy clay, nail polish, aluminium frame
Courtesy the artist. Photo: Philipp Ottendörfer

NORA N. KHAN

In 2016, you were working on your book *Dawn Mission*, published by Mousse, soon to be ten years ago! We met then as I had the happy chance to write some texts for the book. I remember having a file folder of countless animal camera images from you, which I spent hours looking at, imagining the gaze of the animal back into the camera. What did the animal see? How did they see the camera? How did the machine see the animal? There was a sense of communion between these two nonhuman intelligences, a conversation that these images created a space to consider.

A few years later I would learn about Eduardo Kohn's *How Forests Think*, which opens with Kohn being advised to sleep *face up* in the foothills of Sumaco Volcano in Ecuador. His host, Juanicu, tells him that if a jaguar runs upon you and can see you gazing back at him, he'll saunter on. He'll only attack if you sleep facedown, meaning a jaguar has to see you "as a being capable of looking back, a self like himself, a *you*," to leave you alone.[1] We are known through the eyes of other humans, yes, but also animals, and machines. Kohn argues that these encounters force us into a position. To understand that seeing and representing are not human affairs would change how we live in the world.

I'd love to start here. Your works are full of animals and beings gazing back, figures gazing back, from the dark of night, the dark of the sea, the dark of space. How does this encounter – with other kinds of beings – figure in your work? What have you learned from working so consistently with this encounter over the past 10, 15 years?

KATJA NOVITSKOVA

What a wonderful passage from Eduardo Kohn; thank you for introducing it to me! There are a couple of points to my answer. I think even after all this time, there is still a sense of deep mystery to me in the moment when two gazes meet, especially human and nonhuman. It's a profound activation of consciousness, a moment when two self-self membranes touch.

On the other hand, so many of the images I work with have one thing in common: the animal encounters not a human being, but rather a machine – often an automatic camera – that is our extension. These automatic cameras coupled with machine vision and pattern processing algorithms have been developing at a rapid pace in the last few decades. The animal eye often functions differently to ours, yielding alternative experiences of the world.

Machine vision has gotten quite advanced in how and what it can detect, and so these nonhuman interactions are becoming more frequent and more complex. We can only approximate the experience of the other self, but we know it is there because it has eyes, and thus, an inner world. Any meeting point of membranes is a point of translation and interpretation. What does a beluga whale see and think when it is looking directly into the lens of an underwater camera drone? How does a self-driving car see a deer running across the road at night? I find this deeply fascinating.

The white eye glow: that happens due to retinas reflecting the light from the flash, a visual reflection of the camera that is capturing an animal. It does not occur naturally. Speaking in philosophical terms, a photograph of an animal with glowing eyes is a synthetic assemblage that emerged with the technologies of vision. It is always a representation of the technology we use and of us as the viewer besides the animal form itself.

[1] From the introduction of *How Forests Think* by Eduardo Kohn, which is accessed at: https://www.anth.ucsb.edu/sites/secure.lsit.ucsb.edu.anth.d7/files/sitefiles/Kohn%20-%20How%20Forests%20Think%20-%20Introduction.pdf

Approximation (The Apocalypse's Many Horseman), 2020
Digital print on aluminium, cutout display
Courtesy the artist and Kraupa-Tuskany Zeidler. Photo: def image

Approximation III (beluga), 2013
Digital print on aluminium, cutout display
Courtesy the artist and private collection. Photo: def image

NORA N. KHAN

Your interest in the logic of machine vision extends, importantly, to our increased understanding of the natural world *through* the logic of systems. That is, we recognize animals and other beings through the images captured by thermal cameras, skeleton algorithms, movement capture, recognition models, and so on. There's high aesthetic and conceptual drama in this computational drive to capture and know fully, as the gap between the embodied and felt experience is huge. We understand this intuitively. But the computational version of the deer increasingly overrides our experiential knowledge of the deer. The computational version *becomes* our knowledge. You've admirably tracked and traced this push throughout your practice and are particularly foregrounding the intensity and all-encompassing nature of this collapse in these current works through layering and pastiche. Generated deer lick at twisted skeletal outlines. Could you please talk about your process of layering and how you understand the collapse of the computational with the experiential?

KATJA NOVITSKOVA

One of the earliest "epiphanies" I had in my practice is the understanding that, in some cases, there are more images of certain animal species than actual individuals in the wild. I'm never engaging with actual animals for my work, only with their digital representations. The drive to visually capture wildlife – creating digital, flat approximations of them in the form of images – is correlated to the physical decline of the actual embodied populations and their habitats. The icon – the model – is replacing the object, speaking in semiotic terms. And of course, the models are objects themselves: they are digital matrices that sit on servers and consume energy. These computational versions of nature become the default form, and they also keep "evolving" within the logic of computational systems.
    With the arrival of LLMs, it became possible to further explore this tension, or collapse, as you say, between the two worlds. I am "breeding" photographs of animals with other images to generate something that can reveal the differences between the model and the actual in the most intriguing ways. I then proceed to translate the generated images into physical artworks, closing the loop. Of course, there is a vast difference between a living, breathing being and a synthetic simulation of it in 3D-printed resin and aluminium wires, but the artwork is the materialization of the whole system of relations described above.

NORA N. KHAN

I love "approximations" as your term for your best-known works – I remember seeing them at Greene Naftali years ago, how they violently cut through our visual space. You've spoken on *loser images* and *soft approximations* more recently in an interview with Joanna Zylinska. Can you speak about your current approximations, and approximating as a form of forecasting? How do these approximations you are currently creating predict - and what are they predicting?

KATJA NOVITSKOVA

"Approximation" is a term that I believe is critical to understanding models and images as models. Function approximation is a critical concept in machine learning, enabling algorithms to generalize from limited experience to a broader set of states and actions. Intuitively, approximation is something that approaches the real thing but is itself incomplete. This incompleteness is necessary because the model is always a reduced echo of the thing it attempts to capture. The reduction in itself is not a
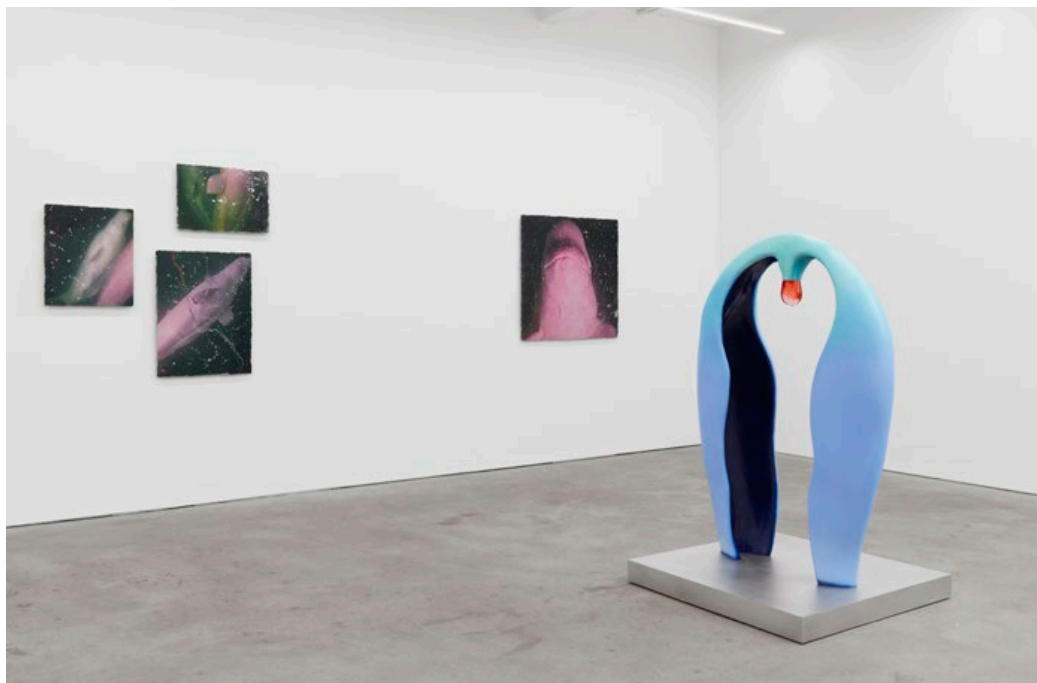
Installation view: exhibition *Soft Approximations*,
Kraupa-Tuskany Zeidler, Berlin, 2022
Front: *Approximation (Looking Glass Penguins)*, 2022
Polymer 3D print, steel structure, polyurethane,
PU foam PU resin, acrylic paint, raw aluminium and MDF
Courtesy the artist and Collection Köser. Photo: def image

Approximation (Looking Glass Beluga), 2022
PLA 3D print, epoxy clay, acrylic paint, raw aluminium and mdf pedestal
Courtesy the artist and private collection. Photo: def image

problem, as long as the approximation can still act similar enough to have predictive potential.

Climate change models approximate climate based on laws of physics, chemistry, and large datasets of sensory information from around the world. The same models can therefore forecast the development of climate systems by "fast-forwarding" the model into the future. Besides these fundamentals, I'm interested in creative, poetic, uncanny, speculative, and contradictory aspects of using something as an approximation for something else, or approximating algorithms. A flat, printed cutout sculpture of a beautiful animal is constructed through a chain of cultural choices and technological conditions: access to the animal, the photographer's choice of composition, camera lens and resolution, digital post-production, aluminium dibond display. The resulting visual pattern becomes a technology that can be used as an attention device (for example, as content online, an advertisement meant for TV screens, or as an artwork).

My intuitive extrapolation from more than ten years ago was that an image of two penguins bonding (or any animal for that matter) is a closer approximation of future technologies or AI to come than any image of contemporary tech at the time. With *Soft Approximations*, I made the next step, employing a more hybrid and layered approach of generating new images that mutate and reconfigure the initial signals in new, unexpected ways.

NORA N. KHAN
I'm curious to hear about the evolution of your aesthetic across your use of different systems and models. You work consistently through this process of scanning, new image generation, re-feeding of the algorithm – in parallel to genetic processes, as you describe. And the tool, or system, you choose seems most interesting for the artifacts it creates, whether Midjourney, DALL·E, or newer diffusion models.

KATJA NOVITSKOVA
My method has remained largely the same since my days as a graphic design student, with the exception of added materiality, largely thanks to a residency I did at Rijksakademie in the Netherlands. I situate myself as a person behind the computer, with access to the Internet, and a desire to understand the ecological reality unfolding in the world today through this mediated window. I use images I find online, openly available texts and accounts, consumer grade software, and recently, various LLMs, in order to, as you say, layer these pre-existing signals into a new understanding of the world.

This is not too different perhaps from a conspiracy theorist who is obsessively drawing circles and arrows around blurry objects in low-res images (a style I also often reference). With the development of new tools and the availability of large datasets, I shifted my focus from attention economies towards a more fundamental relationship between our environment, our ongoing expansive attempts to map, capture and mutate every living thing, and how this impulse is reflected through images, words, scientific graphs, corporate logos, and so on. There are a lot of self-referential elements that travel and transform from work to work, and I see my practice as a one connected body, with many chambers.

NORA N. KHAN
Could you speak about how you think about the use of different generations of AI image software and systems in this process you deploy, layering their outputs, chaining them together? I'm thinking of the beluga whale, among your earlier "approximations," which is twisted over

Approximation (Biobanks), 2022
Digital print on aluminium, cutout display
Courtesy the artist and Marta Herford Museum. Photo: def image

Installation view: exhibition *Soft Approximations*, Kraupa-Tuskany Zeidler, Berlin, 2022
Courtesy the artist and Kraupa-Tuskany Zeidler. Photo: def image

the use of different generative models. We still manage to recognize elements of the original, practicing our own kind of pattern recognition. So: I recognize the eye, the surface texture of the beluga, even in the semi-grotesque current form. How do you consider our training *by* these models? What do you expect your audience might see, if these systems help us to track and predict patterns even as they predict us, our movements?

KATJA NOVITSKOVA

For these recent works, I fed the documentation photos of two of my older series, *Approximation* and *Pattern of Activation*, to an image-generating AI algorithm (Looking Glass, a reverse-engineered unofficial DALL·E knock-off) as input. The resulting images mutated and reconfigured the original artworks in ways that a human artist perhaps would not be able to do – definitely not me. I've been trying to bridge the domain of synthetic biology and approximations that occur within art, towards synthetic artmaking, using biological patterns and shapes as input. Looking Glass is a messy and fuzzy tool with a low-resolution degree of detail, which makes it interesting to use. It helps to cook up a soup of meanings, patterns, and references and it exposes the preferences of the algorithm towards certain patterns, a machine-human visual apophenia. The resulting sculptures can be seen as echoes and ghosts of the original animals whose photos were taken, the sculptures I made using those photographs, the documentation images of the sculptures, the dataset and my input. They are perhaps like data zombies that crawl and re-emerge after being turbulently shuffled between physical and digital domains, between breathing bodies, creating powerful emotional responses and machinic ambivalence. I'm not quite sure why I have the urge to make these, but I think it's about the exploration of the possibility space that lies within this new reality where code, image, and biology can be so deeply intertwined.

NORA N. KHAN

How do you think of the interiority of these animals captured by machine vision – their (arguable across 22+ theories of) animal consciousness, presence, and being in the world, as it is flattened or reframed by machine vision in this computational regime? There is a continual tension between the flat approximations, say, of the penguins lowering their heads and bowing and what I imagine is happening between them. As the penguins' bowing becomes abstracted further in your recent sculptures, something else is communicated by the generative-source image – a red nose, a blue body, the curve of their necks. A different kind of emotion. Do you think this emotion we might feel seeing the generative image-sourced, mutated sculpture as different in order from our emotion seeing the original image, or similar? And in what way? What exactly are we identifying?

KATJA NOVITSKOVA

You are asking about the penguins, but it is perhaps easier to answer using the deer as a reference. With the series of kissing deer, I've been taking these images to an extreme. There was a moment in 2018-10-07 at 11:15pm when these two animals kissed each other in the woods in Veluwe nature park, which was captured by an automatic camera. I'm not sure if the animals were aware of the camera, but they were facing it. I see this moment of emotional and physical closeness as deeply alive: that's what life is to me, these moments of quiet joy and love. And their moment was "captured" by our technological eye.

I can recognize this bonding between two animals from a low-resolution black and white jpeg. This image ended up in a giant dataset
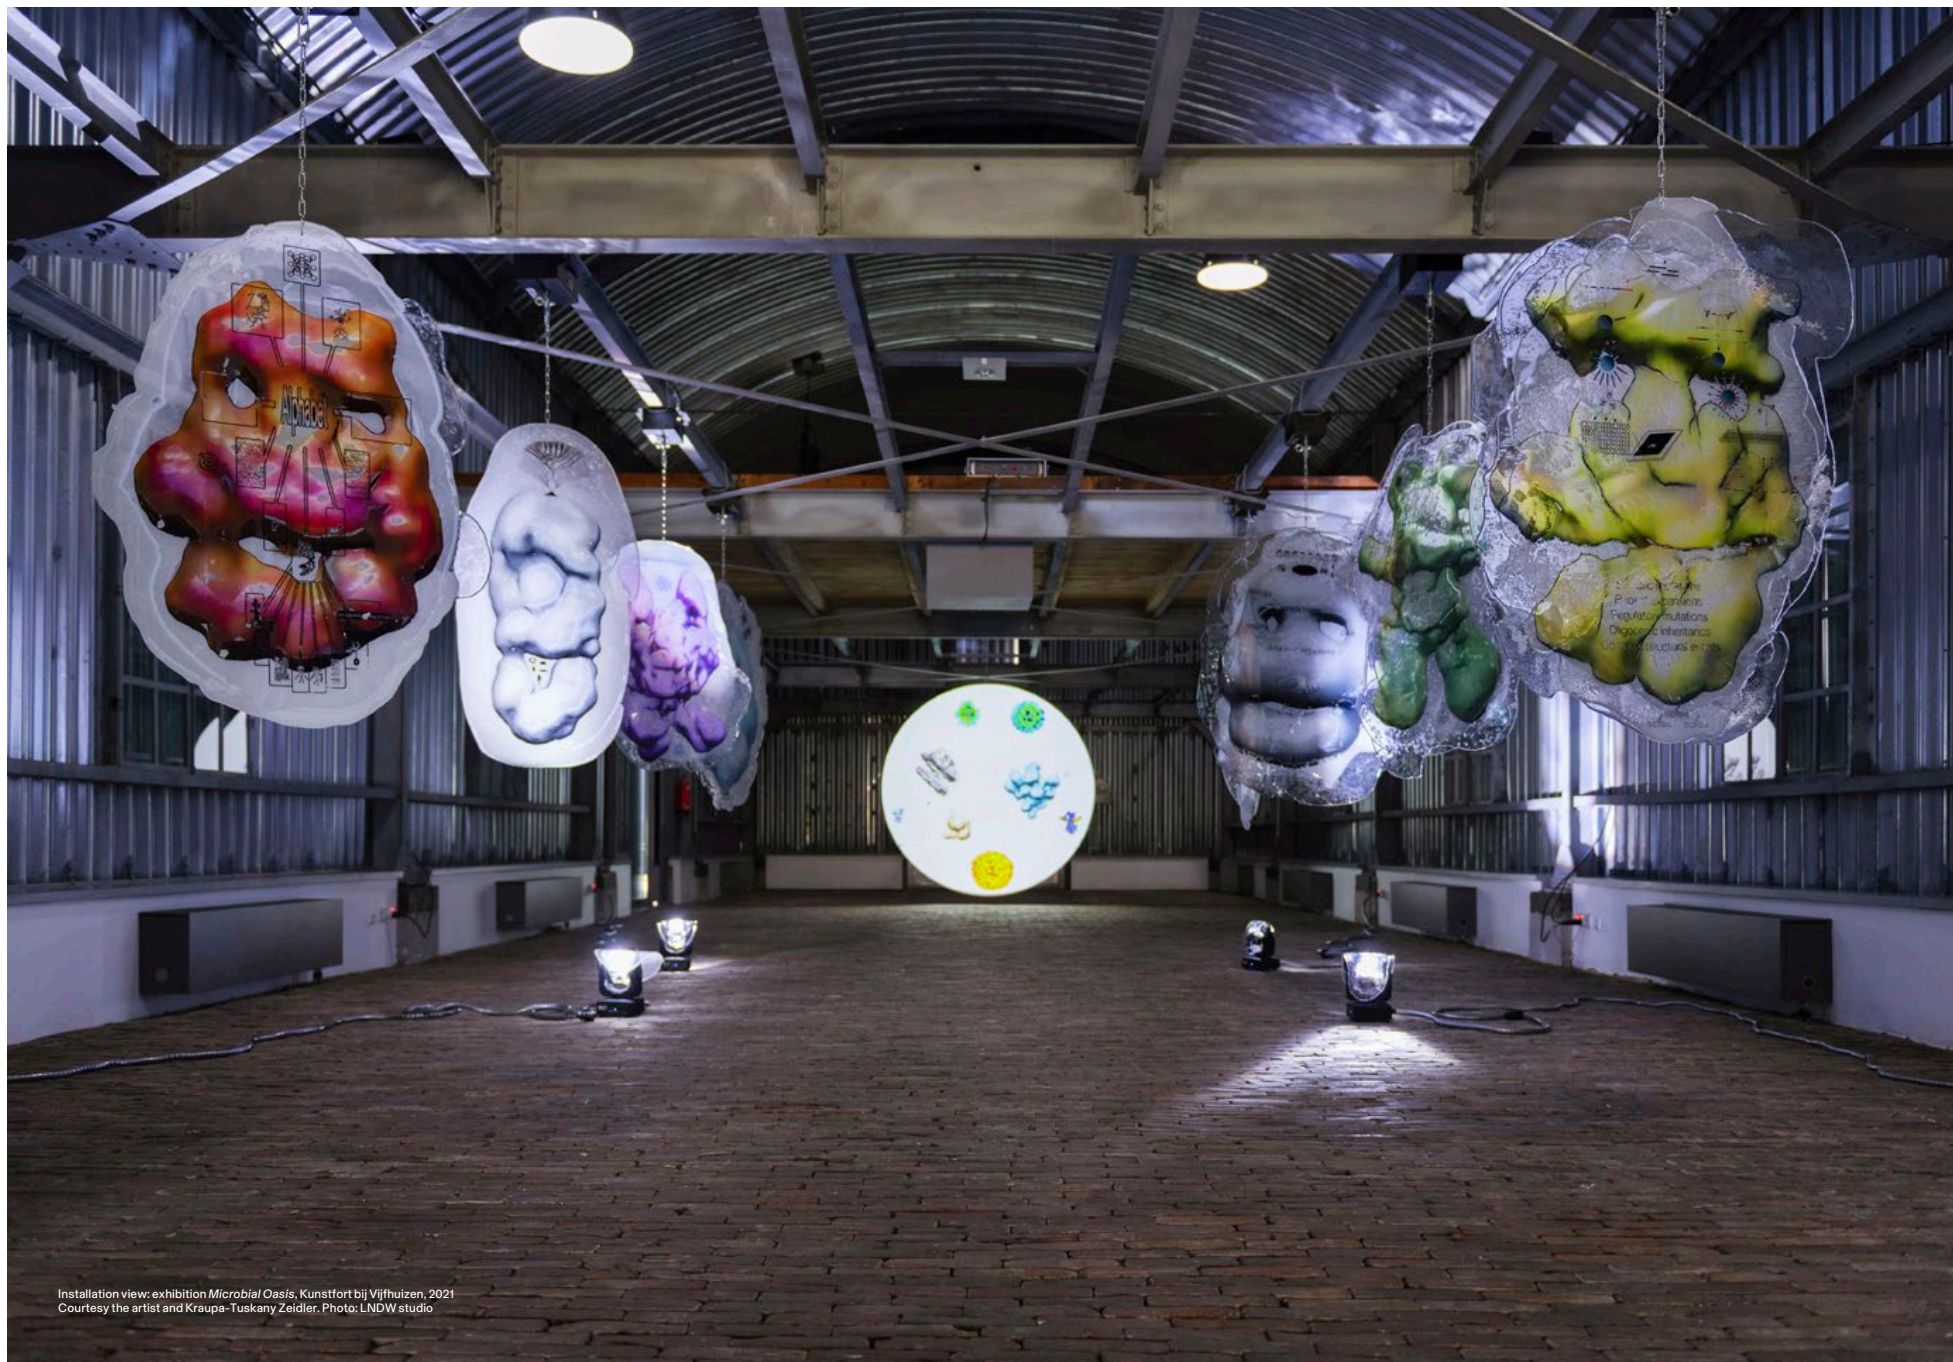
Installation view: exhibition *Microbial Oasis*, Kunstfort bij Vijfhuizen, 2021
Courtesy the artist and Kraupa-Tuskany Zeidler. Photo: LNDW studio
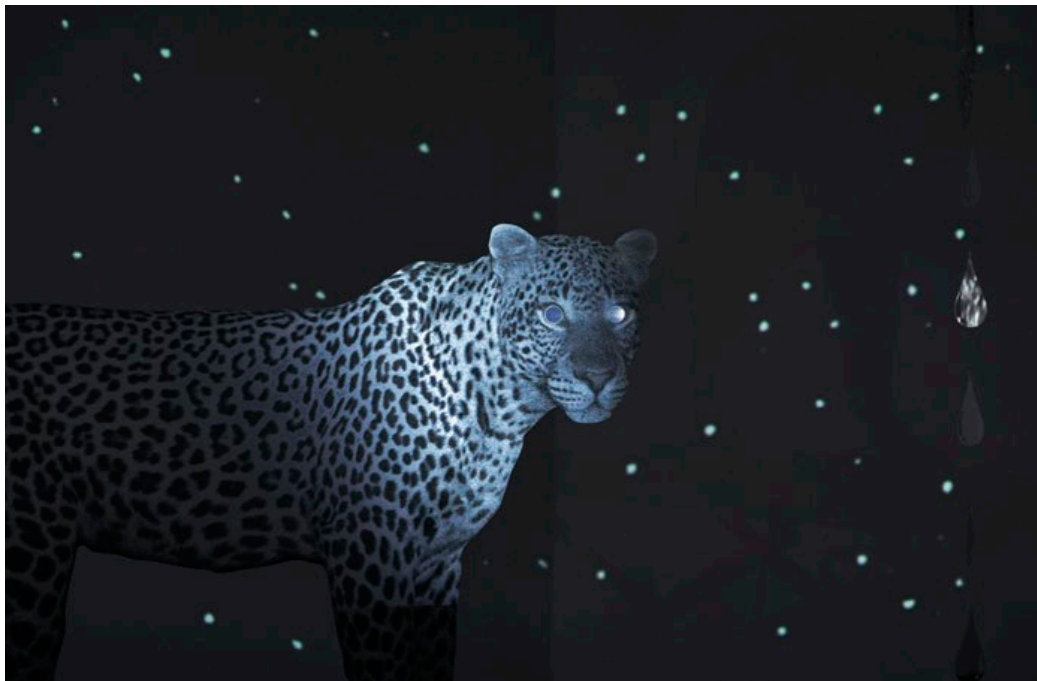
Detail: *Pattern of Activation (Eyes of the World)*, 2017
Courtesy the artist and KAT_A Collection. Photo: Tõnu Tunnel

Screenshot from a Zooniverse citizen science platform dataset, 2019
Courtesy the artist

that was published on a citizen science platform, where I indexed it as part of the volunteer program that helps scientific and environmental projects to process large datasets and train algorithms. In contrast to the penguins image, which was originally a professional photo taken for *National Geographic*, this one would have stayed completely obscure and perhaps only seen by ten people in total. I took a screenshot of the image and kept returning to it in my creative process. I was wondering how the algorithm that would be trained on this dataset would perceive these bonding animals. I used the screenshot to create several art-works that increasingly get more and more transformed by an algorithmic grinder. This is also a key question to me: once we encounter this de-formed amalgamation of signals that vaguely refer to the original moment, what is left of this initial life form and their world? Do we still re-cognize life in it, if there is any life left? Will the algorithm learn some-thing new about the world from me feeding it these non-trivial data points? Where is a limit to the apophenia within our affective response to an object or artwork?

NORA N. KHAN

You spoke on *zombification* and *mutation*; zombification allows us to see different patterns, overlaps, and relationships. Mutation and twisting images to be uncanny: this has been written on in the last years in critical discourse as a kind of defamiliarization, a momentary disjunct in our expectation of what we'll see. When we experience a twist or inversion of our expectation, as we often do with synthetic images, we're, it's argued, taken out of our known world of images and directed *elsewhere*. I'm curious about the most recent protein fold works, the way they take a combination of aesthetics – memes, rave culture, protein folds not found in nature, and face-like apophenia-happy images. These stylistic influences are im-mediately recognizable to a contemporary eye. I take them as both lightly sinister and extremely funny, even as the material is deadly serious. I'd love to hear about the role of humor in your work, which I don't think you get enough credit for!

KATJA NOVITSKOVA

Hehe, thanks for noticing! I don't want to be seen as "irony-pilled," but I appreciate the subtle absurdity in the uncanny that surrounds us. I have a dry sense of humor, and often my go-to instinct while making new work is, "I cannot believe I'm doing this, but I'm doing this."
There is a risk involved in making a new body of work: a risk of being completely misunderstood, a risk of being repetitive, derivative. When the work tingles my sense of the uncanny, I feel like I'm on the right path, at least when it comes to my conscience. There is usually a bit of humor either in my backend ideas or in the presentation. This humor gives me confidence. I also tend not to overproduce things, so there is not too much slickness, and the works are honest about their assembled materiality.

NORA N. KHAN

I'd like to ask about your interest in the deep learning algorithm Alpha-Fold2, and de novo protein structure prediction, in which, as you explain, a protein's tertiary structure is predicted from its amino acid primary sequence. As AlphaFold2 gets better at "approximating the folding of a protein based on its sequence," a seemingly 1:1 relation-ship between the patterns and structures that are algorithmically predictable, deep in our DNA and proteins, emerges. You've noted how AlphaFold2 can manifest, recognize, and "see" novel protein folded structures into existence – a "kind of computational vision and an act of
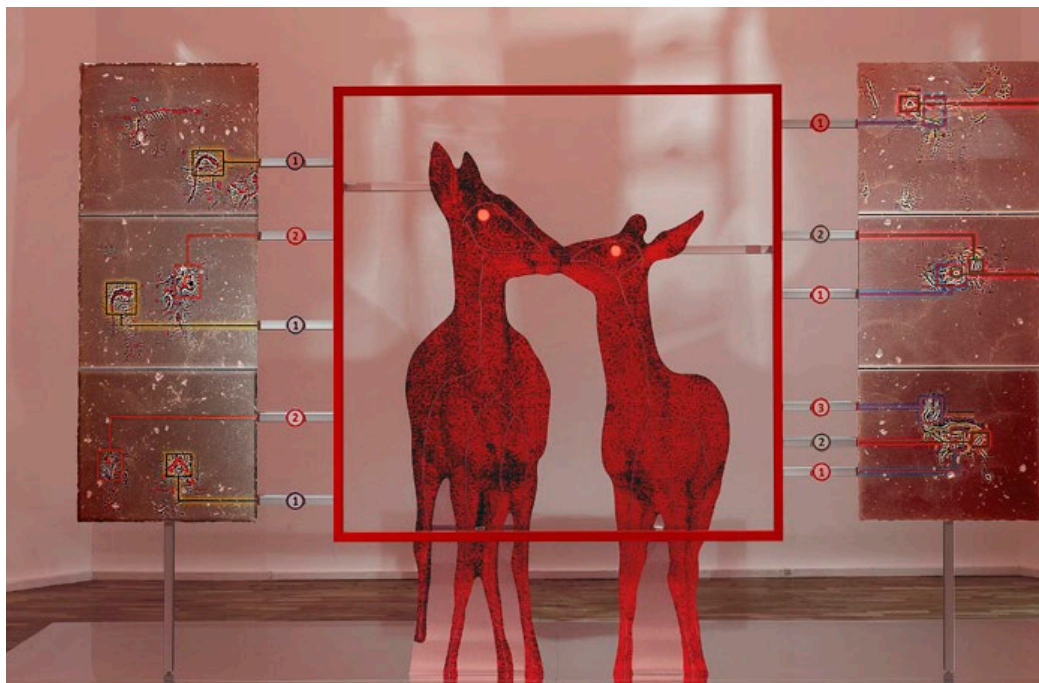
Pattern of Activation (Deers Kissing, Vision Mode), 2023
UV print on aluminium dibond, powder coated aliminum frame, UV ink, epoxy
clay, nail polish, aluminium frames, bolts and screws, plexiglass, mirror panels
Courtesy the artist and Stedelijk Museum, Amsterdam. Photo: Philipp Ottendörfer

Soft Approximation (Looking Glass Deers Kissing 02), 2023
Polyurethane resin, UTR-8100 resin, UTR 3000 resin, water-based
pigments, powder pigments, carnelian, aluminium wire
Courtesy the artist and private collection. Photo: Philipp Ottendörfer

imagination in one." Can you talk about the specific kind of imagination that computational vision produces for us, for you? Where do you see your work moving next in exploring this imagination?

KATJA NOVITSKOVA
Currently, I'm finishing a new series of sculptures that will be ready by the time this interview is published: a small population of computationally imagined and synthetically materialized creatures that form a novel branch of life.

Perhaps this is the most complicated question to answer because I see it as very multidimensional. On the surface level, I find algorithmic hallucinations, artifacts, and glitches to be the source of what is interesting and terrifying about computational vision and algorithmic prediction. In my *Microbial Oasis* project, not only is the algorithm able to generate visual protein models, but it can also hallucinate things, like a bunch of weird faces, as a byproduct. And there is *always* a byproduct. The line can be blurred, and this becomes very serious anytime computational vision has to deal with human beings. These mistakes can often be and often are fatal or life-destroying. On a conceptual level, I like that machine vision and machine intelligence form a new telescope for humanity. These are tools that will reveal things about our world that we had no prior knowledge of. They will stretch our understanding of reality. This is what projects like AlphaFold attempt to do. The fun part for me is to figure out what they will hallucinate, and how some of these hallucinations will no doubt materialize. The question of how similar hallucination and imagination are is perhaps something you can help me answer. :)

The interview was conducted in April 2025.