

Over the last decade, we have witnessed a revolution in machine learning and artificial intelligence: *deep learning*, the art of designing deep neural networks and the craft of training them efficiently on huge datasets, has completely changed the landscape of machine learning. Every week brings exciting new results and technologies based on deep neural networks: self-driving cars by *Tesla* and *Google*, voice assistants with speech recognition by *Apple* and *Google*, machine translation by *Google*, image recognition breakthroughs by *Facebook* and *Google* and programs beating world champions in Go by *DeepMind*, acquired a few years ago by *Google*. You get the picture. Why is it that in a world of budding AI startups (such as *Neuromation*) full of talented researchers and imaginative entrepreneurs, a world where it has never been easier to train a neural network with millions of weights (often all it takes to get a state of the art model is to pull a repository from GitHub), the best results in so many areas still come from a handful of large companies? *Google* has the right culture and a lot of the right people, but why is it ahead at *everything*?

We believe that a large part of the answer is simply its technological advantage: large companies have *more data* and *more computational power*. Yes, you can find a state of the art machine translation model on GitHub or implement it yourself relatively easily—but where will you find the huge parallel corpus of translations that *Google* has painstakingly collected over the years of operating *Google Translate*? Yes, you can download a pre-trained object detection model and fine-tune it on your dataset of images—but labeling this dataset will take a lot of expensive manual labor, and experiments with state of the art computer vision models on a desktop with 1–2 modern GPU's will take days or weeks of processing time.

Here at *Neuromation*, we plan to address both of these bottlenecks. As for the data, we are working on many exciting applications of *synthetic data* with artificially generated labeled datasets that deep learning models can train on. For example, our first large revenue stream is a large contract intended to revolutionize the world of retail: we train object detection models for items on supermarket shelves. Excellent object detection models have already been developed by the deep learning community, and they keep getting better... but to get them to work you need a labeled dataset. It would cost millions of dollars and probably years of labor to get a labeled dataset sufficiently recognizing the 170,000 different items that appear in just one regional retail catalogue. To avoid this manual labor, we develop 3D models of the items that need to be recognized. At a relatively small upfront cost in manual labor (a 3D model of a bottle of Pepsi is quite simple as 3D models go, and then you can reuse the same bottle with dozens of different labels), the outcome is an endless source of perfectly labeled data. Synthetic rendered pictures can have pixel-perfect labeling and even additional data that humans cannot provide at all, like the relative rotation angles of different objects or their depth (z-coordinate) in the picture below. We are continually improving the quality of our renderings, and they already look pretty much like studio photos. Would you be able to tell that this is all synthetic?

No wonder the neural networks also do not care too much. In our experiments, modern computer vision models train on these synthetic pictures just fine, with excellent capabilities for transfer and generalization to real photos.

Synthetic data is not the answer to everything—for example, at this point in time you cannot reliably generate parallel corpora for machine translation, although automated

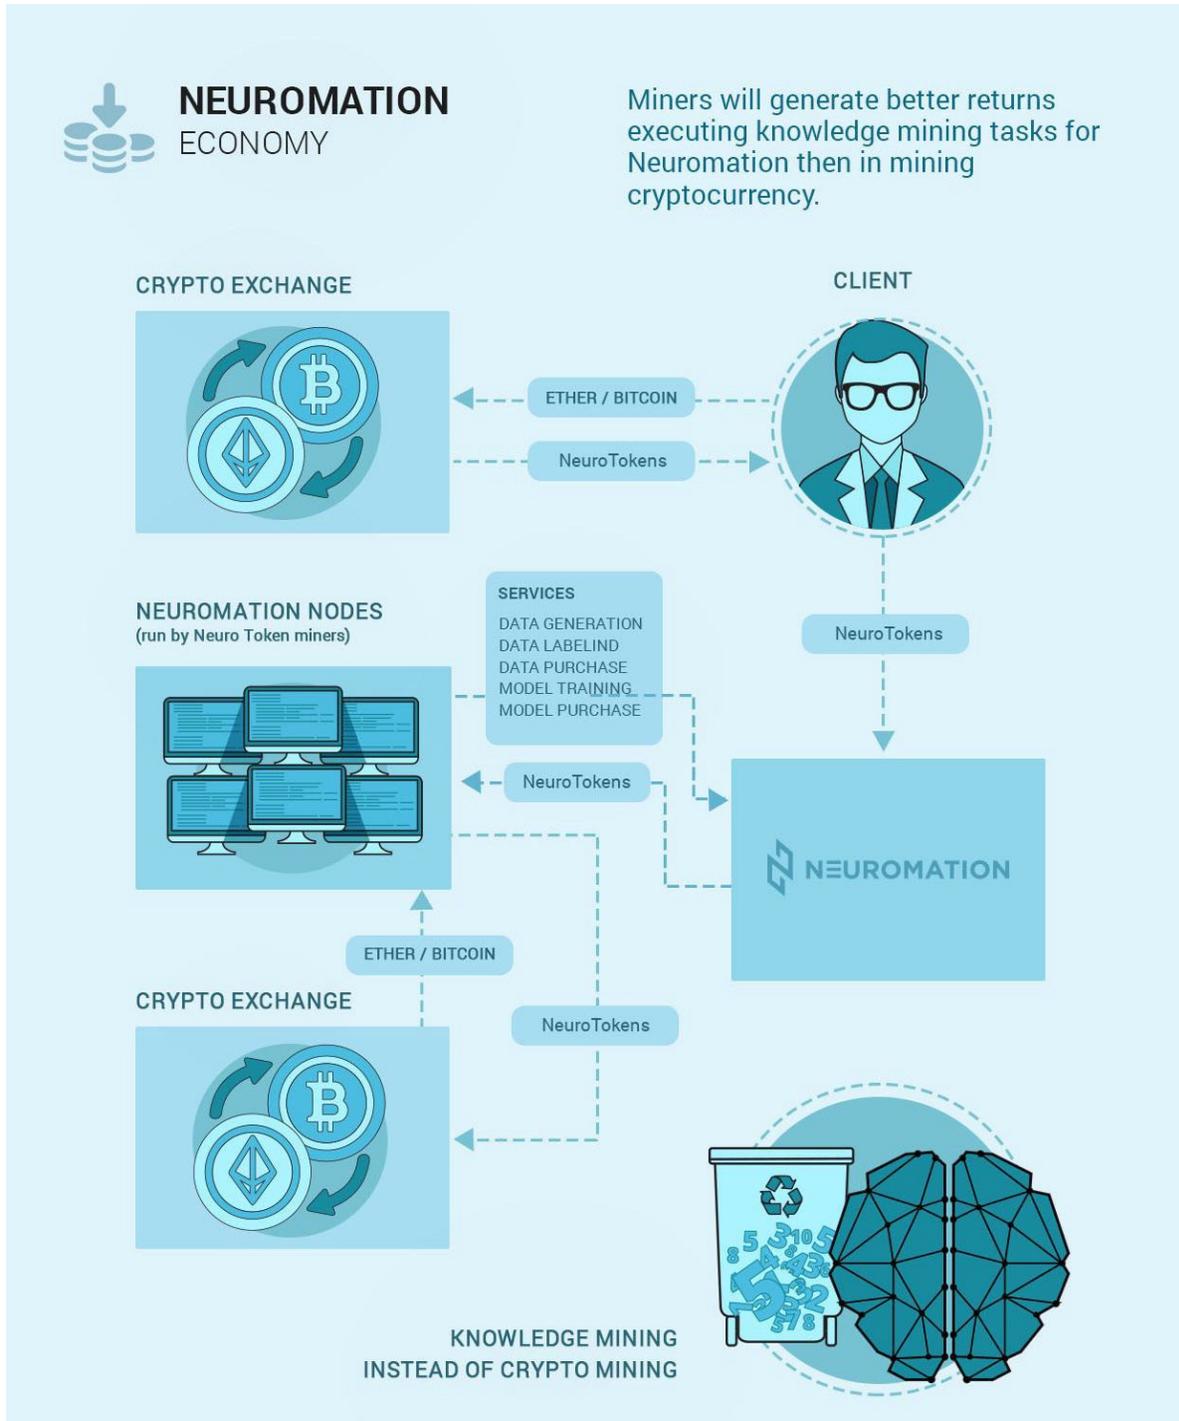
data augmentation techniques exist in natural language processing too. But synthetic data is definitely a great fit for computer vision, including secondary applications like training self-driving cars or flying drones.

The idea of synthetic data blends seamlessly into the second main objective of *Neuromation*: providing computational resources for everyone. Synthetic data makes it easier on the manual labor and human resources costs, but requires quite a lot of computational resources to generate. And naturally, training modern state of the art deep neural networks with millions of weights has never been easy. Where will the computational power come from?

Our basic idea is simple but powerful. Many AI startups train their models on cloud-based providers like *Amazon Web Services*; on AWS, an instance with 4–6 GPUs suitable for modern deep learning will cost you about \$5–7 per hour. At the same time, there are huge numbers of GPUs in cryptocurrency mining farms. By mining, e.g., Ether tokens (ETH), the same “farm” with 4–6 GPUs earns for its owner about \$5–7 *per day*. See the opportunity here? We plan to bridge this gap.

The *Neuromation* platform will become a global unified marketplace for all sorts of computational needs for AI, especially deep learning. We will provide a unified platform of smart contracts for computational power, be it for training deep neural networks, synthetic data generation, or any other needs you may have. A miner will be able to sell the computation time on a GPU farm, getting more money than he could ever get from mining cryptocurrencies... along with the fuzzy feeling of getting his machines to actually do something useful rather than searching for pointless hash collisions. And an AI practitioner will be able to buy this computational power much cheaper than any cloud-based provider can afford to sell it right now.

At present, the easiest way to build this global network of trustworthy smart contracts is to introduce a blockchain-based token of exchange. Hence, we have minted the *Neurotoken*, an Ethereum-based coin for our future platform. We have already reached preliminary agreements with large mining pools, so computational power will be plentiful, and we will be able to provide it for everyone at a much lower cost than centralized cloud-based services — revolutionizing the market of AI model training.



Synthetic data also fuels this idea. First, synthetic data requires significant computational resources to produce (the rendering you saw above does need a couple of GPU seconds to generate), so it is a new and potentially important source of contracts for the *Neuromation* platform. Secondly, training on synthetic data simplifies this kind of outsourcing for training. If you train on synthetic data, you don't have to upload huge datasets to some anonymous mining pool—you simply upload the code for the model together with a data generator, and the data itself is created remotely. This alleviates the

large costs of network transfers that cloud-based providers like *Amazon* spend millions to reduce.

Do you want to be an early adopter and reap the benefits of getting *Neurotokens* at the very low initial offering prices? Check out our pre-sale, which starts on October 15th, and be sure to participate in the upcoming *Neuromation ICO* in November.

Sergey Nikolenko, Chief Research Officer at *Neuromation*

Sergey Nikolenko is a researcher in the field of machine learning (deep learning, Bayesian methods, natural language processing and more) and analysis of algorithms (network algorithms, competitive analysis), Sergey has authored more than 120 research papers, several books, courses “Machine learning”, “Deep learning”, and others. Extensive experience with industrial projects (Neuromation, SolidOpinion, Surfingbird, Deloitte Analytics Institute).