NEUROMATION

NEUROMATION

Where androids dream of electric sheep

White Paper V3.2

www.neuromation.io

Legal Disclaimer

THIS DOCUMENT DOES NOT GIVE PERSONAL, LEGAL, OR FINANCIAL ADVICE. YOU ARE STRONGLY ENCOURAGED TO SEEK YOUR OWN PROFESSIONAL LEGAL AND FINANCIAL ADVICE.

1. The Neuromation NeuroToken (NTK) White Paper (hereinafter “the White Paper” or “the Document”) is presented for informational purposes only.
2. Nothing in the Document shall be construed as an offer to sell or buy securities in any jurisdiction, a solicitation for investment, or investment advice. The Document does not regulate any sale and purchase of the NeuroToken (NTK; as referred to in the White Paper). The sale and purchase of the NeuroToken (NTK) is governed by the Terms and Conditions.
3. Several estimates, phrases, and conclusions incorporated in the White Paper constitute forward-looking statements. Such statements or information concern matters that involve uncertainties and risks, which may result in material differences from the results anticipated.
4. The White Paper may be updated or altered, with the latest version of the Document prevailing over previous versions, and we are not obliged to give you any notice of the fact or content of any changes. The latest version of the White Paper in English is available at the website https://neuromation.io/en/. Although we make every effort to ensure that all data submitted in the White Paper is accurate and up to date at the point in time that the relevant version has been disseminated, the proposed Document is not an alternative to consulting an independent third-party opinion.
5. The White Paper and the related documents may be translated into languages other than English. Should a conflict or an inconsistency arise between the English-language version and a foreign-language version, the English-language version of the Document shall govern and prevail.
6. The White Paper does not constitute an agreement that binds Neuromation. Neuromation, its directors, officers, employees, and associates do not warrant or assume any legal liability arising out of or related to the accuracy, reliability, or completeness of any material contained in the White Paper. To the fullest extent permitted by any applicable law in any jurisdiction, Neuromation shall not be liable for any indirect, special, incidental, consequential, or other losses arising out of, or in connection with, the White Paper, including, but not limited to, loss of revenue, loss of income or profits, and loss of data.

Persons who intend to purchase NeuroTokens (NTK) should seek the advice of independent experts before committing to any action set out in the White Paper.

7. You do not have the legal right to participate in the NeuroToken (NTK) digital asset public sale if you are a citizen, a resident of (tax or otherwise), or a green card holder of the United States of America (including Puerto Rico, the US Virgin Islands, and any other protectorate of the United States) or other representatives of the United States or any jurisdiction where the issue of NeuroTokens (NTK) would be illegal or subject to any requirement for registration, licensing, or lock-up. "A representative of the United States" means a naturalized person resident in any of those jurisdictions or any institution organized or registered in accordance with the laws of any of those jurisdictions. According to the laws of the United States, citizens of the United States living abroad can also be considered "US representatives" under certain conditions.

Legal Disclaimer

8. You agree that you purchase, receive, and hold the NeuroTokens (NTK) at your own risk and that the NeuroTokens (NTK) are provided on an 'as is' basis without warranties of any kind, either express or implied. It is your responsibility to determine if you are legally allowed to purchase the NeuroTokens (NTK) in your jurisdiction and whether you can then resell the NeuroTokens (NTK) to another purchaser in any given jurisdiction. You bear the sole responsibility for determining or assessing the tax implications of your participation in the crowdsale, purchasing, or receiving and holding the NeuroTokens (NTK) in all respects and in any relevant jurisdiction.

9. No regulatory authority has examined or approved of any of the information provided in this White Paper. No such action has been or will be taken under the laws, regulatory requirements, or rules of any jurisdiction.

10. The regulatory status of tokens and distributed ledger technology is unclear or unsettled in many jurisdictions. It is difficult to predict how or whether regulatory agencies may apply existing regulation with respect to such technology and its applications, including the Neuromation Platform and NeuroTokens (NTK) . It is likewise difficult to predict how or whether legislatures or regulatory agencies may implement changes to laws and regulations affecting distributed ledger technology and its applications, including the Neuromation Platform and NeuroTokens (NTK). Regulatory actions could negatively affect the Neuromation Platform and NeuroTokens (NTK) in various ways, including, for purposes of illustration only, through a determination that the purchase, sale, and delivery of NeuroTokens (NTK) constitutes unlawful activity or that NeuroTokens (NTK) are a regulated instrument that requires registration or the licensing of some or all of the parties involved in the purchase, sale, and delivery thereof. The Neuromation Platform may cease operations in a jurisdiction in the event that regulatory actions, or changes to laws or regulations, make it illegal to operate in such jurisdiction or commercially undesirable to obtain the necessary regulatory approval(s) to operate in such jurisdiction.

11. Given that **NeuroTokens** (NTK) **are** based on the Ethereum protocol, any malfunction, breakdown**,** or abandonment of the Ethereum protocol may have a material adverse effect on **NeuroTokens** (NTK). Moreover, advances in **cryptography or** technical advances**,** such as the development of quantum computing, could present risks to **NeuroTokens** (NTK) and the Neuromation Platform, including the utility of **NeuroTokens** (NTK), by rendering ineffective the cryptographic consensus mechanism that underpins the Ethereum protocol.

12. As with other decentralized cryptographic tokens based on the Ethereum protocol, **NeuroTokens** (NTK) **are** susceptible to attacks by miners in the course of validating **NeuroToken** (NTK) transactions on the Ethereum blockchain, including, but not **limited to,** double-spend attacks, majority mining power attacks, and selfish-mining attacks. Any successful attacks present a risk to the Neuromation Platform **and NeuroTokens** (NTK), including, but not limited to, accurate execution and recording of transactions involving **NeuroTokens** (NTK).

Contents

Introduction

Artificial Intelligence (AI) is seeing another spring and becoming a ubiquitous part of society. Researchers and engineers have made important steps towards solving practical problems in machine learning, particularly in such applications as vision, speech, machine translation, and decision making.

Due to technological advancements coupled with progress in research and innovation, AI is quickly becoming a part of every company's operations. This progress brings its own challenges, such as supply of sufficient talent, resources, and data. The most pressing problem is the availability of well-labeled data to support the supervised learning process. Even when unstructured data is abundant, converting it to a labeled dataset is an extremely laborious and costly process. There is a need to address the challenges of practical AI adoption wholesale -- on a platform level. Until this is done, most businesses will be struggling to adopt AI. Due to the disparity in resources and focus among companies, we will continue to see a widening gap between early adopters and legacy operators, putting the latter at an extreme strategic disadvantage.

Neuromation offers a unique solution that unites market resources, the scientific community, and commercial and private entities into an integrated marketplace - the **Neuromation Platform**. To address the most critical issue, Neuromation focuses on synthetic datasets that have been proven to yield compelling results. Using synthetic datasets in machine learning will further decrease the cost of--and will ease--widespread AI adoption.

Introduction

The AI ecosystem is founded on three pillars:

- Datasets – large datasets of structured (or unstructured) data used for deep learning algorithms to learn a specific task.
- Computing power – a standalone server or a farm of interconnected GPUs designed to provide the necessary framework for training AI models with existing datasets to achieve the objective.
- Machine learning models – a set of algorithms designed to process datasets in the Neuromation platform environment.

Each of these domains poses a unique set of challenges for both the supply and demand side of the AI ecosystem. Neuromation offers a new and effective approach of connecting supply and demand for each of these domains in a marketplace, bringing together the competencies of researchers, engineers, and designers and connecting them with commercial sector companies.

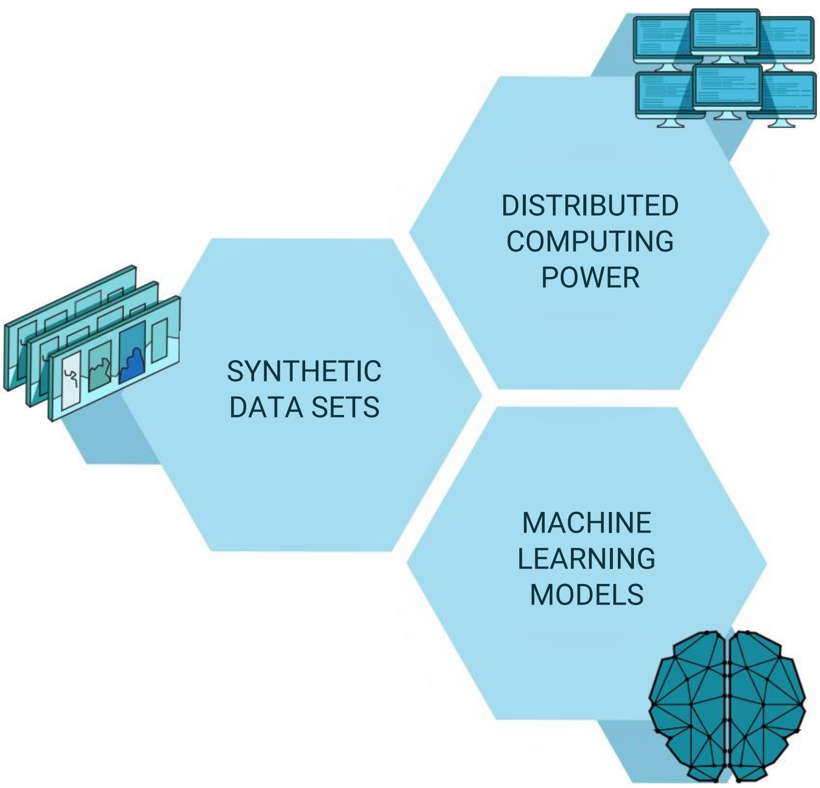

Overview

We are building a critical component of the budding AI ecosystem and plan to capitalize on the first-mover advantage: The Neuromation platform provides an exchange and ecosystem where participants can either contribute to or purchase the components of an AI model.

The Neuromation platform will use distributed computing along with blockchain proof-of-work tokens to revolutionize AI model development. It will combine all the components necessary to build deep learning solutions with synthetic data in one place. Platform service providers, commercial or private, will provide specific resources for the execution and development of synthetic datasets, distributed computing services, and machine learning models, addressing the “three pillars” of AI in the previous slide. Each executed service (dataset generation or model computation) or data piece sold (dataset or data generator) is associated with a reward, numerated in our currency, the NeuroToken (NTK). The Neuromation platform will run based on an auction model that allows customers to negotiate prices directly with service providers.

Imagine a place where you can go and easily address all requests to acquire AI capability. A vendor will create the data generator for you, then a group of Neuromation nodes will use the generator to quickly create a massive virtual dataset. You can then select a set of deep learning architectures to train on that data. Then another group of Neuromation nodes will do the training in record time!

Neuromation Platform

The Neuromation platform will be the destination for data and services for AI

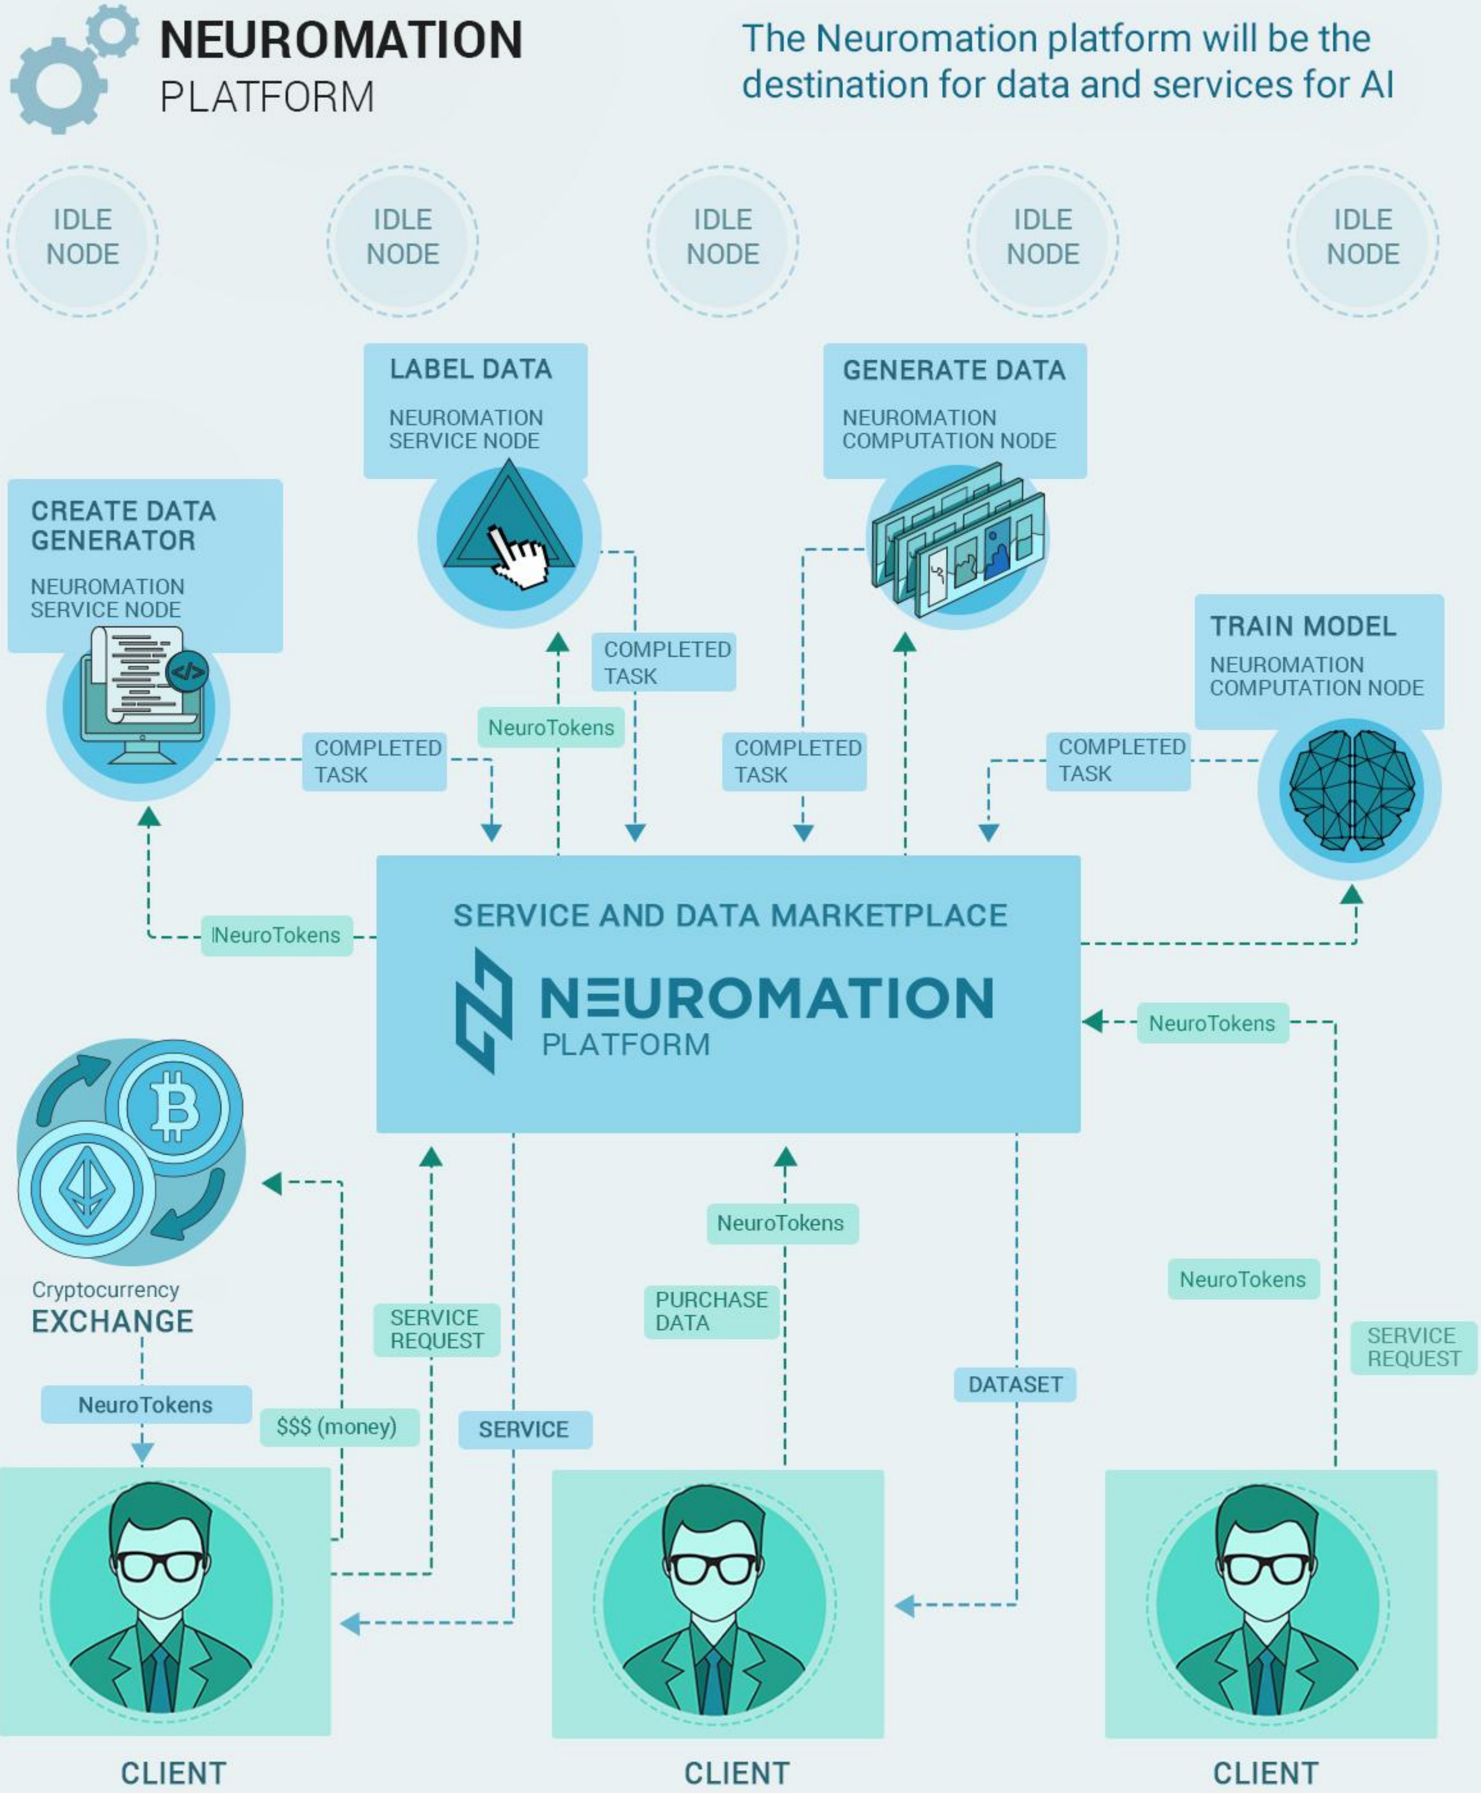

Platform Elements Overview

The Neuromation platform will combine technological and business elements to efficiently train large-scale deep learning systems on synthetic data.

Platform user experience is divided into the following categories:

Synthetic Dataset Module:

PROCESSES:

- Create a data generator
- Order dataset from the data generator
- Request data labeling

LIBRARIES:

- A repository of datasets for deep learning
- A repository of data generators
- Datasets (marketplace)

Machine Learning Module:

PROCESSES:

- Define deep learning model
- Import the model (e.g., clone TensorFlow code)
- Order training on selected dataset
- Request a custom model from the marketplace

LIBRARIES:

- A repository of deep learning models (marketplace)

User Module:

PROCESSES:

- Purchase NeuroTokens
- Register to use or provide processing power
- Register to use and provide services
- Download and install Neuromation node software

LIBRARIES:

- User data
- User models

Platform Elements Overview

The platform's backend technology is divided into the following components:

Neuromation Node:

- Node analysis (Processing price floor — minimum amount in NeuroTokens accepted per computing unit. Node efficiency limit based on bandwidth/processing power/storage)
- Distributed data generation package
- Distributed learning package
- Node synchronization middleware

Market Module:

- Enables efficient matching of buying and selling orders for datasets, models, and labeling services
- Enables liquidity in the system so that prices across the assets are conducive to scaling the system

AI Marketplace - Exchange Overview

All value exchange operations on the Neuromation Platform will use NeuroTokens (NTK).

Available services:

Buy Data	Label Data	Generate Data	Buy Model	Train Model

Platform Price Setting:

The price for each service will be determined by the aggregate setting of the Neuromation nodes (price per unit of computation). Each node will have a minimum token price-floor setting. The minimal price floor can also be adjusted dynamically via an algorithm that will maximize the total tokens earned for the node. The Neuromation platform will determine the resources required for each requested task and select the most efficient node pool (minimizing the price for the customer). Between nodes “sniffing” out the market and customers hunting for the most efficient price, the platform will find equilibrium in supply and demand.

Token Purchase for Clients:

To complete transactions on the Neuromation platform, a client will need to buy NeuroTokens (NTK). To simplify the purchase mechanics, Neuromation will provide a client portal that will make NeuroToken (NTK) purchases a one-click process.

Neuromation Token Flow

NEUROMATION ECONOMY

Miners will generate better returns executing knowledge mining tasks for Neuromation than in mining cryptocurrency

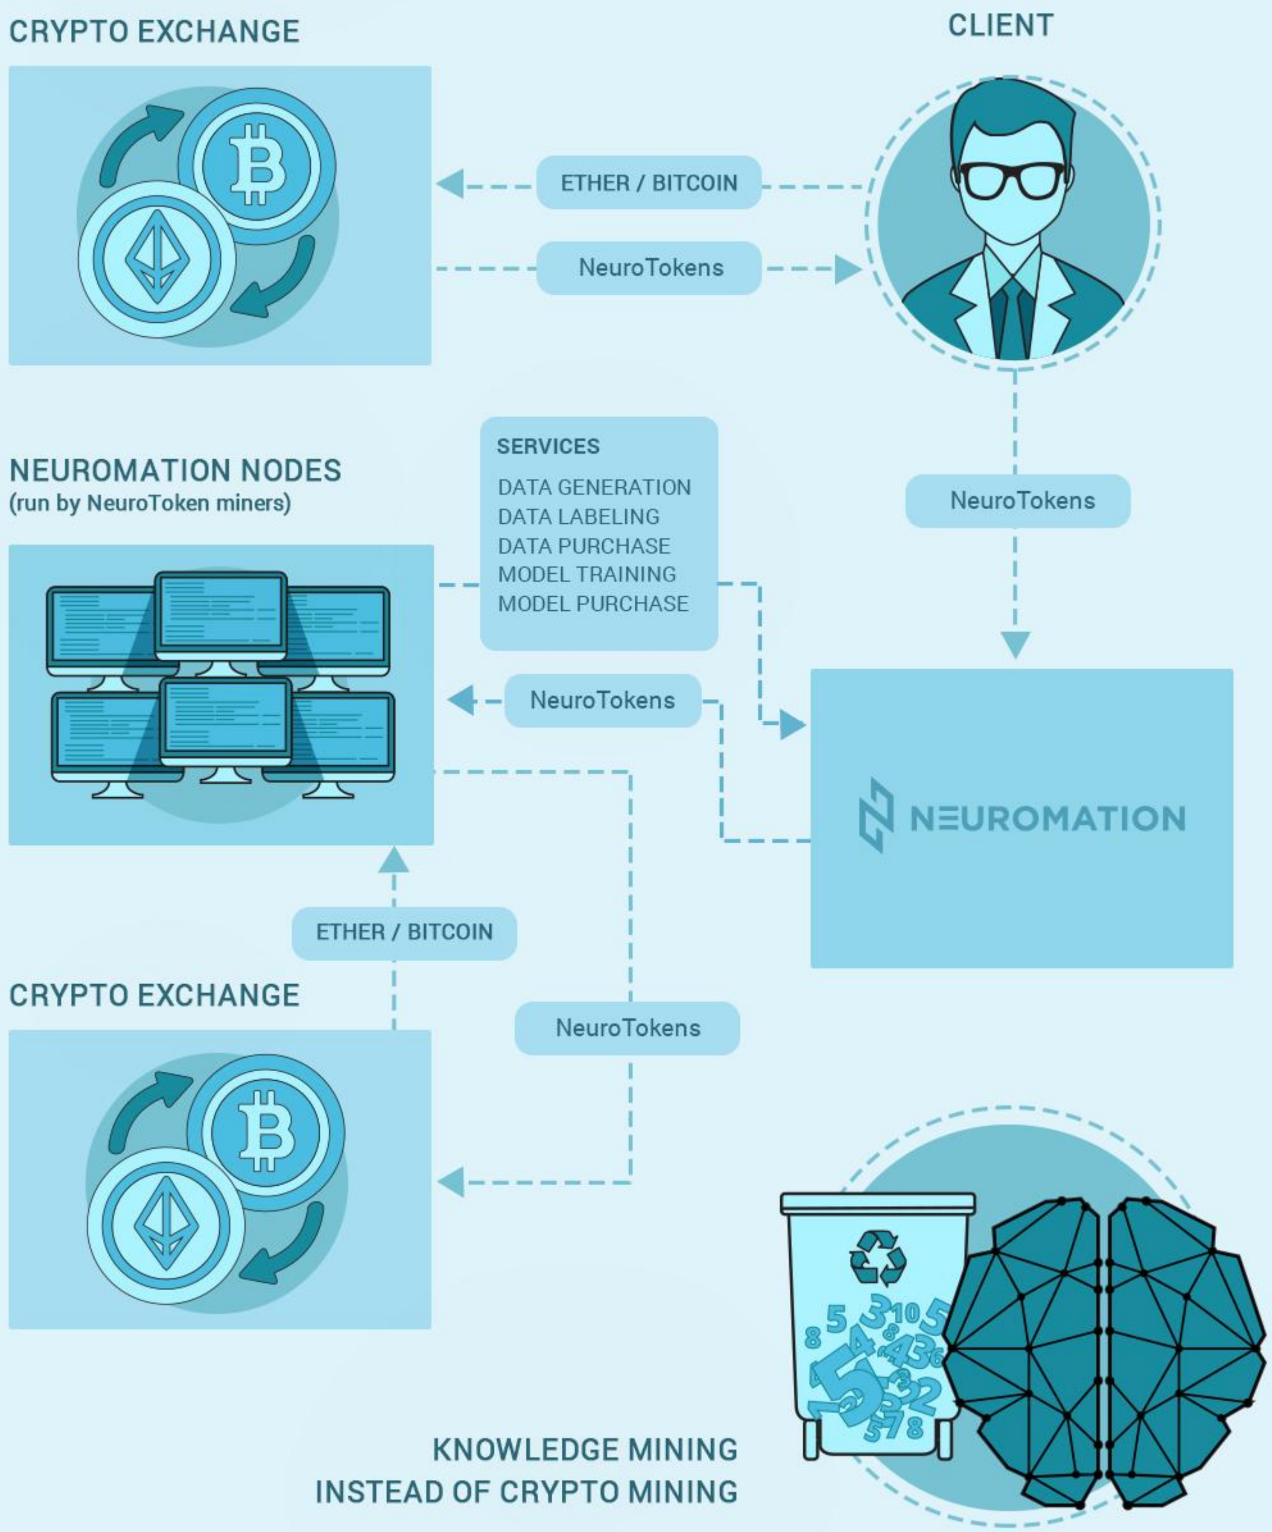

Neural Networks and Deep Learning

Deep learning employs extremely large artificial neural networks; training such networks often requires vast datasets with highly accurate labeling. Collecting large datasets of images, text, or sound is easy, but describing and annotating data to make it usable has traditionally been challenging and costly. Crowdsourcing was applied to the problem of dataset collection and labeling a few years ago, employing large numbers of humans to provide much-needed annotations and improve label accuracy. It proved slow and expensive and introduced human bias. We propose a solution where accuracy is guaranteed by construction: large datasets with perfectly accurate labels are synthesized artificially.The benefits of synthetic data are manifold. It is fast to synthesize and render, perfectly accurate, tailored for the task at hand, and can be modified to improve the model and training itself.

Industry is already using synthetic data to some degree, but its utility is limited. There is currently no general tool set available to facilitate the use of synthetic data on a large scale. **The Neuromation platform will change this.**

Neural Networks and Deep Learning

Ten years ago, machine learning experienced a revolution. Although neural networks are one of the oldest paradigms in artificial intelligence, researchers and practitioners had been unable to train large-scale deep neural networks reliably and efficiently until 2009. Due to new insights from Geoffrey Hinton, Yoshua Bengio, and Yann LeCun, neural architectures quickly became state of the art in image processing, speech recognition, and natural language processing. This led to breakthrough results in machine learning and redefined artificial neural networks as **deep learning**. Today, deep learning is the *lingua franca* of machine learning, spanning domains from biometrics and self-driving cars to playing Go and synthesizing speech.

The basic building block of a neural network is the perceptron, a simple mathematical model of a neuron first proposed by Frank Rosenblatt in 1957. A perceptron computes a weighted sum of its inputs (dendrites in biological neurons), whose weights represent model parameters to be learned, followed by a nonlinear activation function. The neural network itself is a large composition of perceptrons, whose parameters can be jointly optimized using gradient descent computed by backpropagation, a technique first proposed in the 1960s and brought to full generality in the 1970s.

Neural Networks and Deep Learning

The compositional structure of neural networks allows for parallelization in both the training and inference of such models. Parallel processing of the networks' layers advanced GPU-based computational paradigms, leading to breakthroughs in high-performance computing. Currently, virtually all industrial applications of neural networks are trained and run on GPUs, and the next jump in efficiency will involve specialized hardware, such as Google's TPU and other embedded hardware platforms.

Even the simplest neural architectures (e.g., a two-layer feedforward network), are known to be universal approximators — they have the theoretical ability to approximate any given mathematical function. In turn, complex neural networks must have sufficient expressive power (capacity).

When trying to efficiently train deep neural networks with large capacity, one quickly arrives at a problem: "Where can we find a lot of well-labeled data?" and "How noisy are the labels?"

This problem is solved by the Neuromation platform.

Synthetic Data

One of the most important problems of modern deep learning is where to get *data*. Common neural architectures are usually intended for *supervised learning* (training to solve a specific classification or a regression problem on a labeled dataset). There are techniques for *unsupervised learning* within the deep learning framework (e.g., with stacks of restricted Boltzmann machines and generative adversarial networks) and approaches to convert unsupervised feature extraction into a supervised problem (via autoencoders). However, fine tuning with supervised datasets is still required. A lot of effort in modern deep learning research goes into *transfer learning* (reusing trained models for other tasks on similar data), *domain adaptation* (reusing trained models for similar tasks on datasets with different properties), and *one-shot learning* (using a handful of training samples to learn a new class once the network has trained on other classes), but these efforts are still far from mainstream.
As for the data, although it is often easy to get a virtually unlimited dataset of unlabeled samples (e.g., random photos off the Internet or human speech records), it is very laborious and expensive to label such data manually. For the last few years, crowdsourcing has been virtually the only way to create and label datasets; it employed large numbers of people to correct mistakes and improve accuracy. This method proved slow and expensive and introduced human bias.

An average retail set has 150,000+ items. A deep NN needs thousands of labeled photo examples for each item. Doing this by hand would take years.

Moreover, some data (e.g., per pixel depth values in images) is virtually impossible to label by humans (we are not good at absolute depth perception) and sophisticated methods and tools are required to obtain the ground truth (accurate labels).


ACCURACY


SPEED


AUTOMATION

Synthetic Data

We propose a solution where perfect label quality is guaranteed by construction: large datasets with perfectly accurate labels are synthesized artificially. The benefits of synthetic data are manifold:

- once the environment is ready, it is fast and cheap to create as much data as needed;
- the data is perfectly accurate and tailor-made for the task at hand, with labeling that might be impossible to obtain by hand (for example, an illumination map of the environment or depth values);
- it can be modified to improve the model and training in a constant feedback loop between the model and the synthetic environment.

In our solution, synthetic data is used for training, and a small validation set is comprised of real, manually labeled data. In essence, this makes using synthetic data for a *transfer learning* problem possible: we need to reuse models trained on one kind of dataset (synthetic) and apply them on another kind of dataset (actual images). Our approach, however, has several important advantages that greatly simplify transfer learning in this case. First, the synthetic training dataset is not given as part of the problem but rather generated by us - we can and do try to make it match real data.

> **Data is the new oil. Getting new hand-labeled datasets is like finding new oil deposits. Synthetic data is like cheap and ubiquitous synthetic oil.**

Second, and even more importantly, the above-mentioned feedback loop between the model and the dataset makes it much simpler to perform transfer learning: we are able to tune not only the model but also the training set, a luxury seldom available to machine learning practitioners.

Synthetic Data

In a current application of this approach, we use state-of-the art object detection and image segmentation models to recognize objects on the shelves of a grocery store/supermarket. The objects include food items, such as bottles of soft drinks (Pepsi/Mirinda), cartons of juice, and packs of potato chips. This specific application is particularly interesting for training models on synthetic data. The objects themselves are relatively easy to render: a Pepsi bottle makes for a much simpler 3D model than a human face does.

Sufficient hand-labeled data is very hard to obtain due to the following reasons:

- large-scale food companies, such as Pepsico, may have thousands of different items, and these catalogues change frequently;
- the objective is to be able to recognize these objects in photos of supermarket shelves with hundreds of items on each; hand-labeling even one such photo is a laborious process, and millions of them are needed;
- high-quality learning also requires a massive null set — labeled objects that are similar to ones being learned but not identical to them; typically, 10x the number of objects in the training set is required.

Synthetic Data

We have trained several different models for *object recognition* (SSD, YoLo, and Fast(er) R-CNN) and *object segmentation* (SegNet and FCN for segmentation), and are currently working on an implementation of Mask R-CNN for *instance segmentation*. All of these models contain a convolutional network that does the bulk of feature extraction work from the image and differ in the extra layers on top of these features that capture the actual semantics and produce class labels, bounding box candidates, segmentation, and other outputs.

We have already obtained excellent object recognition results on both synthetic test sets (which would be expected) and test sets with real photographs (which is encouraging and serves as proof of concept).

In further work, we plan to improve synthetic data-generation pipelines with machine learning algorithms, leading to a true active learning framework: Synthetic data is used to improve a model that is verified on real data. At the same time, the learned model improves the synthetic generation pipeline to further benefit learning.

Distributed Computing and the Neuromation Economy

The deficiency in computing power required to scale synthetic dataset generation is the key impediment for widespread adoption. The Neuromation platform has an economically viable proposition to cryptocurrency miners that would attract their computing resources and help resolve the problem.

Background:

Crypto miners are facing pressure from the ever-decreasing efficiency of their proof-of-work computations required to mine Ether or other coins. Growing competition, consolidation of miner pools, and increasing demand for computing power and resources will soon make it difficult for many smaller miners to keep their efforts economically feasible.

Neuromation economic incentive for miners:

We want miners to have options. In addition to their existing mining software, they can load up Neuromation Computation Node. When a Neuromation task is available, each node can bid to participate. If a node wins the bid, it will switch computing power from mining Ether or another coin to a task on the Neuromation blockchain platform. The node will generate synthetic data or train a deep learning model. As a reward, the miner will receive our NeuroTokens (NTK). Once the task is complete, the Neuromation node will exit and the miner will proceed with mining crypto. Initially, our NeuroTokens (NTK) will extend Ether. We plan to migrate to our own blockchain later on once the Neuromation platform economy matures.

Our preliminary estimates indicate that miners will earn up to 3 to 5 times more by mining NeuroTokens (NTK) versus cryptocurrency. Miners will not be engaged in Neuromation tasks 100% of the time, so mining NeuroTokens (NTK) will be an efficiency boost to their existing set up. You can see the difference in effective yield below for similarly configured rigs: one running a cryptocurrency mining algorithm and one running deep learning/data rendering tasks.

Type of Computing	Yield/Time
Ether Mining	$ 7 - 8 USD/**Day**
Amazon Deep Learning	$ 3 - 4 USD/**Hour**

Machine Learning Models

- There are many publicly available models: **VGG**, **GoogLeNet**, and **ResNet** for computer vision, **Wavenet** for audio, and so on.
- They come in a variety of frameworks and libraries: **TensorFlow**, **Caffe**, **Torch**, **pyTorch**, **Theano**, **Keras** etc.; converters exist between the frameworks.
- The Neuromation Platform will support all major deep learning and automated differentiation frameworks and libraries and provide support for major off-the-shelf pretrained model formats.
- Training results will be made available for the user to download.

- Moreover, models can be compressed to speed up response times or allow using them on mobile devices. Common **model compression** tools include:
  - after a (large) model has finished training, many of its layers often prove to be superfluous and can be removed, compressing the network by a factor of 2-3; this especially applies to very deep computer vision networks based on residual connections; this technique is known as **model distillation**;
  - some networks, called **mobile nets**, are designed from the ground up to be small (e.g., SqueezeNet in computer vision).

- Requests for custom models and architectures. If existing models and tools are not sufficient for the given applications, customers can request a custom model design for a specific problem. This request will be posted to the **Neuromation Exchange** and will be filled by either Neuromation experts or other certified suppliers.

Neuromation Labs

Our strategy is not to develop our platform in isolation but to work with partners in select industries to bring our vision to life organically. We are developing “Neuromation Labs” that will develop synthetic data and train deep learning models on live applications. Each “lab” will be a study of a specific problem in partnership with a category leader. As our platform is fleshed out, we will move parts of generation and training there, allowing us to organically test parts of our vision in real-life scenarios.

Industry collaboration is key to our success.

The Labs will seed the Neuromation platform market with initial data generators and datasets. We will also encourage our Labs partners to provide further services through the platform, thus building the initial market. We have already launched two Labs and have more in progress. Working off real-world examples allows us to generalize similar approaches across industries and incorporate these approaches into our platform, making it more powerful with each Lab.

Neuromation Model Example — classifying merchandise on supermarket shelves

Electricity transformed almost everything 100 years ago. Today, I actually think AI is the new electricity [slightly paraphrased] - Andrew Ng

Retail Automation Labs

Neuromation has partnered with some of the industry's leading brands to solve the problem of finding consumer goods on store shelves. Using synthetic data of every product, we are able to create large, perfectly labeled datasets for hundreds of thousands of SKUs in the retail industry. Deep learning vision algorithms trained on these massive datasets are able to efficiently analyze and label shelf availability, percentage of the shelf, layout accuracy, and other metrics.

Retail Automation Lab (live)

Our Retail Technology Lab is the first to go live in the field with select partners in European, Middle Eastern, North American, and South American markets. Ideally, every shelf will feed a constant stream of analytic data to our retail partners. Using our synthetic data approach, we are able to retrain analytic models quickly and accurately, making AI for retail agile and useful for a multitude of applications.

Synthetic Data Example — procedurally generated 3D store shelves:

Exponential growth is projected to continue for the store inventory and logistics robotics market.

Retail Automation lab (live)

What is especially fascinating is that in only about a month of working on the inventory recognition problem, we have achieved 95%+ accuracy, a result others have spent years of effort and millions of dollars on. Notably, the model performs well on real-life data without seeing anything but synthetic datasets during training.
This breakthrough proves the viability and efficiency of our approach.
With the Neuromation platform, the gateway to easy AI training at scale will finally be opened.

(Breakthrough) Model taught on synthetic data recognizes real data.

The potential demand for image recognition alone in the retail industry is enormous.

In fact, by 2020, 85% of customer interactions in retail will be managed by artificial intelligence, according to Gartner.

It amounts to more than 40 billion images per year, according to ECR research of 72 of the biggest retailers and suppliers. Going forward, we plan to create datasets that mimic human interaction with the shelf. We will be able to track customer flow and intent, creating a full simulator of the retail store with a multitude of possible applications.

Enterprise Automation Labs

A typical difficulty with industrial environments is the highly dynamic nature of performed tasks. Models have to be agile enough to work with dynamic, real-time data with high-precision output. This processing is done on the fly (sometimes literally — some tasks are performed by flying drones). Real-time datasets cannot be built fast enough in a dynamic environment to process and output them in a timely manner. The answer to this problem is the Neuromation platform, which can procedurally create virtual environments where the models can be trained. Once the models are trained in these virtual sandboxes, they can be quickly calibrated on real data. As a result, agile hypothesis testing becomes possible.

Neuromation plans to open an Enterprise Automation Lab where the synthetic data approach will help implement solutions in a variety of industries, including manufacturing, supply chain, financial services, and agriculture.

Pharma/Medical/Biotech

There are many interesting targets for the Neuromation approach in pharma/medicine/biotech. In general, any problem where construction is easier than recognition (output is easy to calculate from input, but the reverse is hard, like a hash function in computer science) will yield well to our synthetic data approach. We estimate that a substantial amount of hypothesis building in drug discovery and biological system dynamics can be handled with deep learning models. To kick off this lab, we have partnered with a leading provider of infant monitoring solutions.

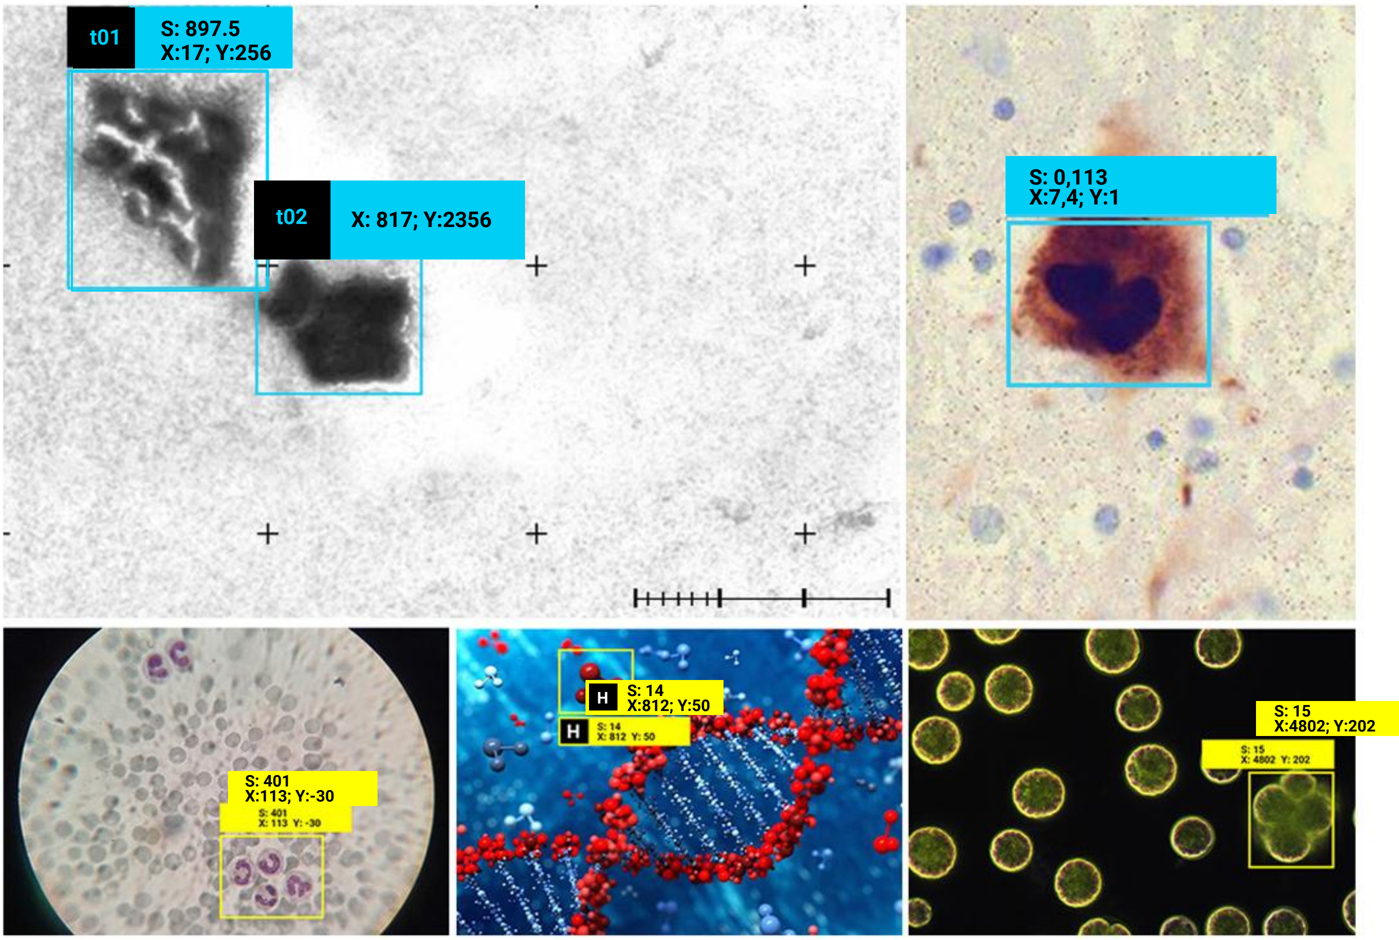

We will jointly develop a smart camera that will send alerts when the baby has switched positions, changed breathing patterns, woken up, etc.

Other Neuromation Labs

Neuromation plans to open other Labs in partnership with industry. Our strategy is to have a partner in every domain.

Team

Max Prasolov
CEO

As a long-time entrepreneur, senior manager, and producer for the past 16 years, Maxim has produced more than 50 animated commercials, 3D movies, commercial and industrial applications, and computer games. He has worked with international retail and industrial brands, such as Unilever, Yukos, TPE, Metro Cash & Carry, Severstal Group, Ferrexpo, commercial banks, and investment and insurance companies. Maxim was a member of the IPO team of Ferrexpo, the largest iron ore producer in Ukraine, which was listed on the London Stock Exchange. Since 2014, he has been investing in drone development, AI, and augmented reality multimedia startups.

Fedor Savchenko
CTO

Fedor has more than 20 years of experience in leading complex software development projects with an emphasis on computer graphics, 3D engines, CGI production, and VR environments. He holds degrees in both mathematics and graphic design. This unique perspective allows him to build software that is both functional and beautiful. He is an innovator who is always on the cutting edge of the latest trends.

Sergey Nikolenko
Chief Research Officer

As a researcher at the Steklov Mathematical Institute at St. Petersburg in the field of machine learning (deep learning, Bayesian methods, natural language processing, and more) and the analysis of algorithms (network algorithms and competitive analysis), Sergey has authored more than 120 research papers, several books, courses on machine learning and "deep learning," and other publications. He has extensive experience with various industrial projects, including Neuromation, SolidOpinion, Surfingbird, and the Deloitte Analytics Institute.

Constantine Goltsev
Investor / Chairman

A serial entrepreneur and online advertising industry veteran, Constantine has more than 20 years of experience in software and product development. He is the former CEO and founder of the pioneering video advertising network AdoTube, which grew to 200 employees, selling in 23 markets with 13 offices worldwide. AdoTube was subsequently sold to Exponential Interactive. Constantine is also a founder and the president of SolidOpinion — a company that has innovated engagement in online communities around content and commenting.

Team

Denis Popov
VP of Engineering

Denis has over 15 years of software engineering experience and has led large technical projects from inception to completion. He has served as the CTO, CSA, VP of engineering, and adviser to dozens of successful projects and startups. Denis has expertise in the management of technical processes and in high-volume concurrency, big data storage, and real-time data analysis. He was a technical lead at the facial recognition company Viewdle, which was acquired by Google.

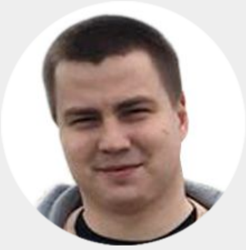

Kyryl Truskovskyi
Lead researcher

Kyryl has over four years of experience in the field of machine learning. For the bulk of this time, he has helped build startups from inception. He has also developed expertise in selecting and implementing full architecture and large-scale solutions. Kyryl has a strong mathematical background and a passion for research that has prompted him to pursue an academic path. He is currently a PhD student in computer science.

Esther Katz
VP Communication

Esther is a PR professional with 20 years of activity in major markets with a focus on biotech and fintech. She is the former head of Jacob Peres Office, a leading PR agency in Israel. Esther is an early cryptocurrency adopter and blockchain enthusiast. She lives and works in Israel.

Evan Katz
CRO

Evan is a leader in the AI Driven Media Industry. He has worked at AnyClip (one of Israels fastest growing companies) and Taptica (International Media Company listed on UK Stock Exchange). He has also founded and owned Multiverse (AI Driven Media Company) and Tribal Marketing (Import & distribution of Sunglasses worldwide). His education includes Business Management & International Trade and part of a Mechanical Engineering Degree at Technion Israel. He is an exceptional Cook, who enjoys riding Motor Cycles, playing Rugby, Expanding Blockchain community in Israel. He is socially very aware and volunteered to help mentor young entrepreneurs and did extensive fund Raising for young entrepreneurs in South Africa.He is an ambitious, determined and goal orientated asset who brings enthusiasm, charisma and charm to each and every project making him immediately approachable and successful.

Advisors

Andrew Rabinovich
Adviser

Andrew is a world-leading scientist in deep learning and computer vision research. He has been studying machine learning with an emphasis on computer vision for over 15 years. He is the author of numerous patents and peer-reviewed publications. He also founded a biotechnology startup that was subsequently acquired. Andrew received a PhD in computer science from the University of California, San Diego in 2008. He has held several leading R&D positions at Google and Magic Leap, where he is currently the director of deep learning.

David Orban
Adviser

David is the founder and managing partner of Network Society Ventures, a seed-stage global investment firm focused on innovative ventures at the intersection of exponential technologies and decentralized networks. He is an entrepreneur, visionary, guru, author, blogger, keynote speaker, and thought leader of the global technology landscape. His entrepreneurial accomplishments span several companies that were founded and grown over more than twenty years. David is also the founder and a trustee of Network Society Research, a London-based global think tank.

Yuri Kundin
Adviser

Yuri Kundin is the director at the KPMG US San Francisco office with a specialization in risk and compliance advisory services. He has over 15 years of experience in advising on regulatory compliance matters to address requirements of various US agencies (e.g., the SEC, OCC, and FDIC) and the Monetary Authority of Singapore, Hong Kong Monetary Authority, and Russian Central Bank. Lately, Yuri has been specializing in developing frameworks for risk, compliance, and attestation in blockchain, cryptocurrency, and token sale ecosystems.

Dave Carlson
Adviser

Dave Carlson is the CEO of the Giga-Watt, one of the largest turnkey mining services and custom packages service providers in North America. With a longstanding software engineering background and successful entrepreneurial track record, Dave is one of the pioneers of the mining community. Prior to becoming the CEO of Giga-Watt, he served for 4 years as the CEO and founder of MegaBigPower, one of the largest single-operator mining facilities in the world.

NEUROMATION

Revenue Sources - Platform

The Neuromation platform intends to be the premier destination for AI services for the world's businesses. We project around $71 million USD gross in transactions on the platform. Depending on the type of service received through the platform, Neuromation will take a 5% to 15% commission from each transaction.

We anticipate that platform usage will grow 3x to 5x a year from 2018 to 2022. We project that the Neuromation platform will generate over $100 million in yearly revenue from commissions in three years.

ESTIMATED PLATFORM ACTIVITY BREAKDOWN (*disclaimer)

Job on the platform	Price ETH	Price NeuroToken (NTK)	Price USD	Time
Create data generator	100.00	100 000.00	20 000	4 weeks
Generate data (10k units)	0.17	174.22	100	1 day
Model training	0.01	87.11	50	0.5 days
Data labeling (1k set)	0.88	871.09	500	1 day
Data set sales	10.00	10.00	2 000	NA
Model sales	100.00	100.00	20 000	NA

EXPECTED 2018 PLATFORM ACTIVITY TRANSACTIONS VOLUME (*disclaimer)

Job on the platform	2018 Expected jobs	Vol. ETH	Vol. NeuroToken (NTK)	Vol. USD
Create data generator	1000	100 000.00	100 000 000.00	20 000 000
Generate (10k units)	10 000	1 742.16	1 742 160.28	1 000 000
Model training	100 000	1 451.80	1 451 800.80	5 000 000
Data labeling (1k set)	10 000	8 710.80	8 710 801.39	5 000 000
Data set sales	10 000	100 000.00	100 000 000.00	20 000 000
Model sales	1000	100 000.00	100 000 000.00	20 000 000
Total	**132 000**	**311 904.76**	**311 904 760.00**	**71 000 000**

* These numbers are estimates and may vary substantially from actual performance.

Revenue Sources — Neuromation Labs

Through Neuromation Labs, we partner with industry leaders to spearhead AI solutions with them. Neuromation generates revenue by selling API (SaaS) or data to the industry partner.

Examples:

- **Retail Lab**

In our Retail Lab, we have partnered with OSA HP (http://osahp.com/en/) and ECR (http://ecr-all.org/?lang=en) to supply a 170,000-item synthetic retail dataset for Eastern Europe. The contract involves building deep learning models that can recognize the objects on the shelves and be an integral part of OSA analytics. The contract potential is over €4.25 million over an 18-month time frame. Here is an example of the model deployment on a mobile device: https://youtu.be/QO1KGwQKoSU

- **Medical Devices Lab**

In our Medical Devices Lab, we have partnered with MonBaby (https://monbaby.com/), a leading infant tracking technology provider, to create a smart camera to augment their existing offering. The camera will track the motion of the child and give caretakers valuable analytics. Neuromation will receive a revenue share from each device sold. The revenue potential is over €2 million over the next few years.

- **Other Labs**

Several other Labs are slated to open in early 2018. These industry partnerships should give Neuromation additional revenue opportunities.

Our overall expectation is that we will make around 30% of our total revenue through Neuromation Labs partnerships.

Neuromation Road Map

NEUROMATION

Q4 2016	Team building
Q1 2017	Seed round completed. Work started
Q2 2017	Launch of R&D in Retail Lab to test synthetic data hypothesis
Q3 2017	Proof of hypothesis on synthetic data completed. Retail synthetic data generator proven. Model development complete (95% of accuracy reached). Marketing and first sales to retail clients. Token sale preparation.
Q4 2017	Neuromation token sales begin. Neuromation platform v1 launched with 1,000-GPU capacity deployed. Medical Devices Lab started.
Q1 2018	Active sales and marketing of Neuromation platform. NeuroTokens (NTK) are tradable on exchanges. Industrial Automation Lab is started.
Q2 2018	Neuromation platform v2 launched with 10,000-GPU capacity deployed.
Q3 2018	NeuroToken (NTK) will migrate to its own blockchain. Neuromation platform v3 launched with 100,000-GPU capacity deployed.

Beyond 2018, we expect Neuromation to become a definitive destination for businesses to develop their AI capability. Our NeuroTokens (NTK) are the primary exchange mechanism enabling synthetic data generation, distributed model training, data labeling, and other AI services.

We envision being a global resource pool for synthetic data by providing an ever-growing library of datasets for every conceivable use case. As we tie service providers into the platform, Neuromation will be an enabling ecosystem for all AI.

Token Sale

NEUROMATION

Timing:

RECESS

Oct 25th 2017
Pre-sale
(Whitelist registration)

Jan 1st 2018
Pre-sale ends

Jan 7th 2018
Sale starts

Feb 15th 2018
Sale ends

Goals:

2,000 ETH
Retail + Med Devices Labs.
Basic Neuromation Platform.

5,000 ETH
Industrial Innovation Labs.
1,000-GPU cluster enabled
on the platform.

10,000 ETH
Large Nauromation Platform
(full distributed data
Generation and computation).
10,000-GPUs pool tied
To the platform.

60,000 ETH
Neuromation Fund,
100,000-GPU pool.
World wide adoption.

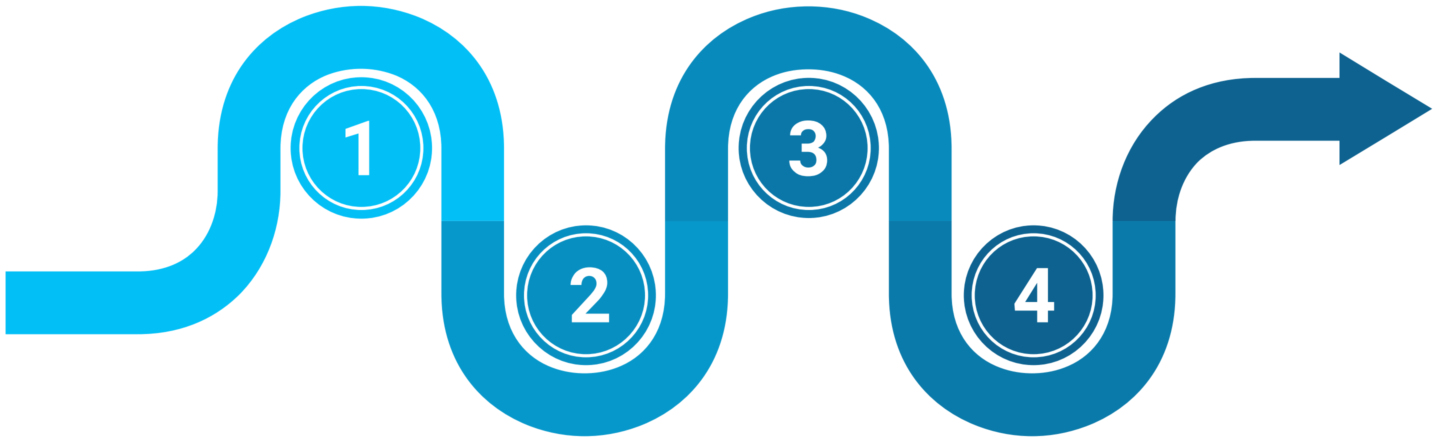

Token Sale

Jurisdiction: Estonia. Neuromation is compliant with Estonian crowdfunding laws.

Token Details:

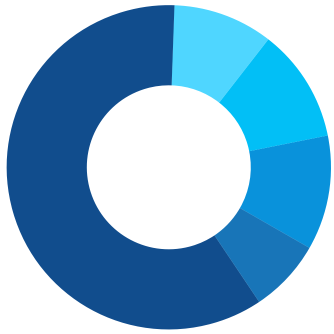

- NeuroTokens (NTK) minted: 100,000,000 (will not mint more)
- Total tokens issued in token sale: 60,000,000. Neuromation will burn unsold tokens.
- Token burn: 30% in 2018, 20% in 2019, and 10% in 2020.
  (Neuromation will burn 50% of the tokens over the next 3 years and decrease token supply)
- Price per NeuroToken (NTK) = 0.001 ETH

Token Distribution:

- 60% in the token sale
- 12% to the team
- 12% to partners
- 10% to liquidity reserve
- 6% to researchers

Token Sale Proceeds Distribution:

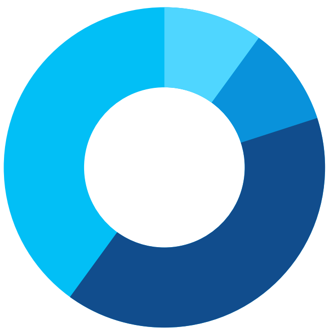

- At least 40% for platform development
- Up to 40% for liquidity reserve (prepay for server capacity)
- 10% for PR and marketing of the Neuromation service
- 10% to partners/advisers/early backers

Token Sale Bonus Program:

TIMING BONUS:

- Presale (register here to find out bonus percentage)
- First week + 15% in NeuroTokens (NTK)
- Second week + 10% in NeuroTokens (NTK)
- Third week + 5% in Neurotokens (NTK)

VOLUME BONUS (NTK):

- 30 Ethers - 1%
- 50 Ethers - 2.5%
- 100 Ethers - 5%
- 300 Ethers - 8%
- 500 Ethers - 10%
- 1000 Ethers - 15%
- 2000 Ethers - 20%
- 3000 Ethers - 30%+

Token Economics

NEUROMATION

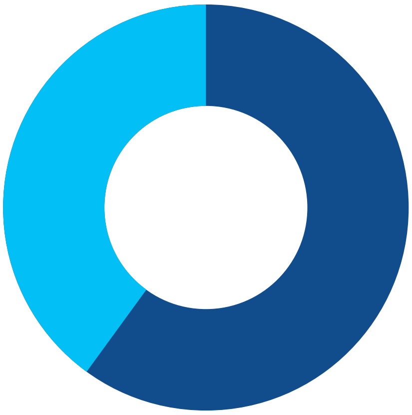

	Tokens Minted	Available in Token Sale
	100 000 000	60 000 000 (*)
	100 000 ETH	60 000 ETH

* All tokens in available pool unsold during Token Sale will be burned.

Nominal Token Price: 1 = 0,001

ESTIMATED PLATFORM ACTIVITY BREAKDOWN (*disclaimer)

Job on the Platform	Price ETH	Price NeuroToken (NTK)	Price USD	Time
Create data generator	100.00	100 000.00	20 000	4 weeks
Generate data (10k units)	0.17	174.22	100	1 day
Model training	0.01	87.11	50	0.5 days
Data labeling (1k set)	0.88	871.09	500	1 day
Data set sales	10.00	10.00	2 000	NA
Model sales	100.00	100.00	20 000	NA

PLATFORM 2018 EXPECTED ACTIVITY TRANSACTIONS VOLUME (*disclaimer)

Job on the Platform	2018 Expected jobs	Vol. ETH	Vol. NeuroToken (NTK)	Vol. USD
Create data generator	1000	100 000.00	100 000 000.00	20 000 000
Generate data (10k units)	10 000	1 742.16	1 742 160.28	1 000 000
Model training	100 000	1 451.80	1 451 800.80	5 000 000
Data labeling (1k set)	10 000	8 710.80	8 710 801.39	5 000 000
Data set sales	10 000	100 000.00	100 000 000.00	20 000 000
Model sales	1000	100 000.00	100 000 000.00	20 000 000
Total	**132 000**	**311 904.76**	**311 904 760.00**	**71 000 000**

* These are estimates only. Actual performance may differ substantially.

Token Economics

(*disclaimer)

	2018 NeuroToken (NTK) demand	2019 NeuroToken (NTK) demand (2018 x 5)	2020 NeuroToken (NTK) demand (2019 x 3)
	319 163 763.07	1 595 818 815.33	4 787 456 445.99
% of circulated tokens burned	30%	20%	10%
Total available	70 000 000.00	56 000 000.00	50 400 000.00
* Assumes 60,000,000 tokens sold in the Token Sale	**Demand for tokens is 4x the supply**	**Demand for tokens is 30x the supply**	**Demand for tokens is 90x the supply**

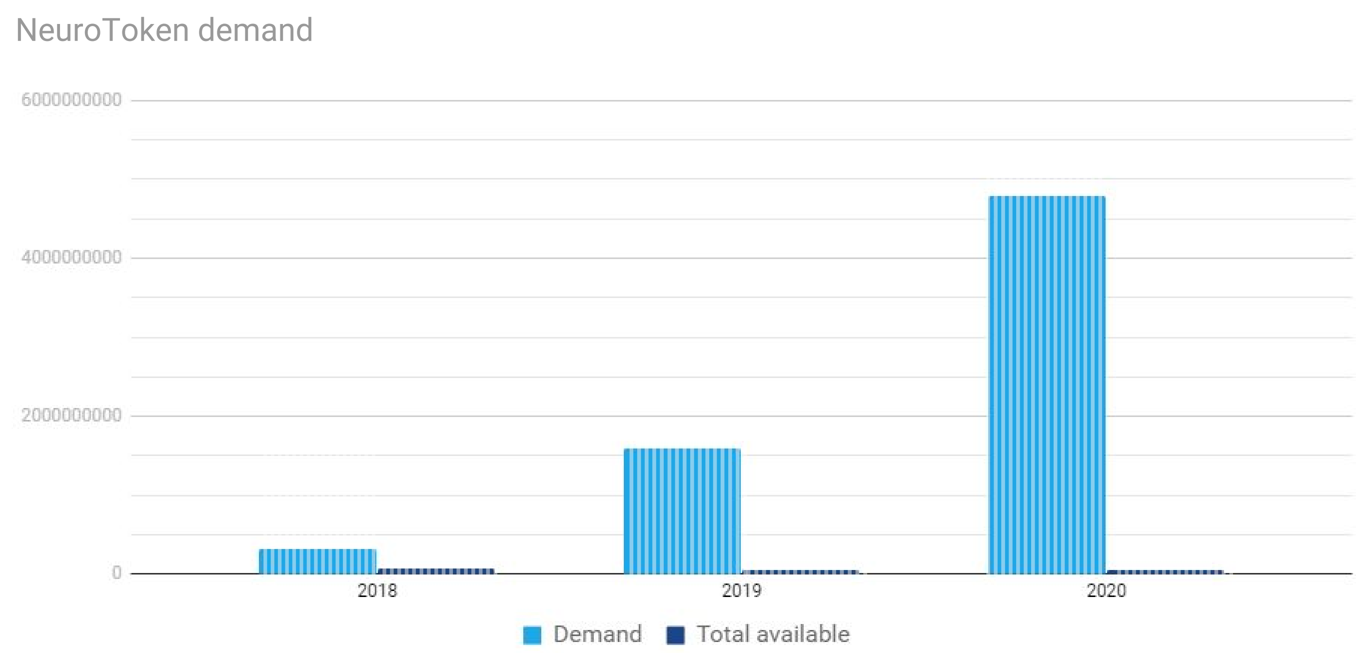

With our token burning policy, **NeuroToken** (NTK) demand on the Neuromation platform is expected to be **90 times the available token supply** by the end of 2020.

* These numbers are estimates only and may vary substantially from actual performance.