


NEURON

Decentralized Artificial Intelligence
Distributing Deep Learning
to the Edge of the Network



In our world, where
intelligence is everything,
do we really want to end
up with a couple of
corporations owning the
world's **most intelligent**
machines?

Abstract

The Neuron Network is a decentralized learning platform for machines. It is a distributed and permissionless network of users (or peers) with data ownership necessary to teach (or train) neural network based artificial intelligence models. Humans are the users and the machines are the students. An open, publicly accessible network democratizes access to artificial intelligence. It incentivizes users to bring data to the network for use as the curriculum used to teach machines.

An initial implementation of the network will focus on decentralizing the stochastic gradient descent algorithm to enable the network to operate in a permissionless and decentralized fashion. The network is a token-based economy. Participants are required to prove ownership of tokens and receive tokens for participating. Such a network enables a broader participation and incites a greater desire to interact amongst those qualified to do so.

Overview

The advent of Machine Learning as a Service (“MLaaS”) platforms has made artificial intelligence algorithms accessible to the common developer. After a series of acquisitions by large corporations, artificial intelligence companies and associated intellectual property have quickly commoditized under the umbrella of a few incumbents.

Today, in lieu of aggregating the data and human capital necessary to build an applied artificial intelligence-based company, it is possible to build on top of existing APIs.

This includes facial and image recognition, text analysis, NLP, sentiment analysis, language translation and general prediction. It has become easier than ever to incorporate machine learning and deep learning algorithms into mobile and web applications.

In order to give way to the next wave of innovation in artificial intelligence, we propose a decentralized artificial intelligence Neuron Network of decentralized participants to generate an open, crowdsourced learning network.



Blockchain

The Ethereum network is a platform that enables the creation of peer-to-peer applications based on smart contracts. This allows developers to create cryptographically enforceable relationships. An Ethereum-based smart contract platform will provide the network infrastructure it needs. This includes incentivizing a network of developers, end users, voters and research organizations to participate. The incentive will be in the form of a token that is necessary for participating in the network.

Artificial Intelligence Decentralized

The blockchain enables parties to interact based on a set of agreed upon business rules. These business rules can define payment transfers or generic rulesets called smart contracts. A decentralized network of peers host these rule sets in order to verify proposed transactions on a smart contract.

Smart contracts can be used to define the manner in which various parties interact in a trustless, permissionless manner. A trustless, permissionless network can lay the foundation for a platform. Data aggregation and deep learning models can be developed that would otherwise be too costly for centralized institutions to undertake.

Imagine the Machine University of the Future...

Imagine the university of the future where students are machines, teachers are humans.

Eventually the teachers will also be users of the machine they have trained.

EDUCATION TODAY

Carbon-based students and teachers

Student applies to university

Student gets admitted

Student goes to centralized university campus

Student gets curriculum, buys syllabi

Student studies courses

Students are coached by professors

Student summarizes, structures

Student takes exam

Student gets result

Student graduates or retakes course (learning stops)

Student leaves to join the workforce, learning is a discrete function (now and then student returns for refresher and then goes back)

Students are coached by professors

EDUCATION TOMORROW

Silicon students, carbon teachers

Teacher (user) sends public key to blockchain

Teacher (user) gets private key to AI

NEURON is decentralized, it is virtual - both everywhere and nowhere

Machine is exposed to user (teacher) datasets

Machine classifies datasets

Users (teachers) and students are coached by smart contracts

Machine goes through an epoch (go over the complete training data once)

Machine builds a prediction model

Machine gets a score (F1)

Machine repeats and keeps getting better (learning does not stop)

DAIs never leave, learning has become a continuous function, and happens decentralized in real-time.

Users (teachers) and students are coached by smart contracts

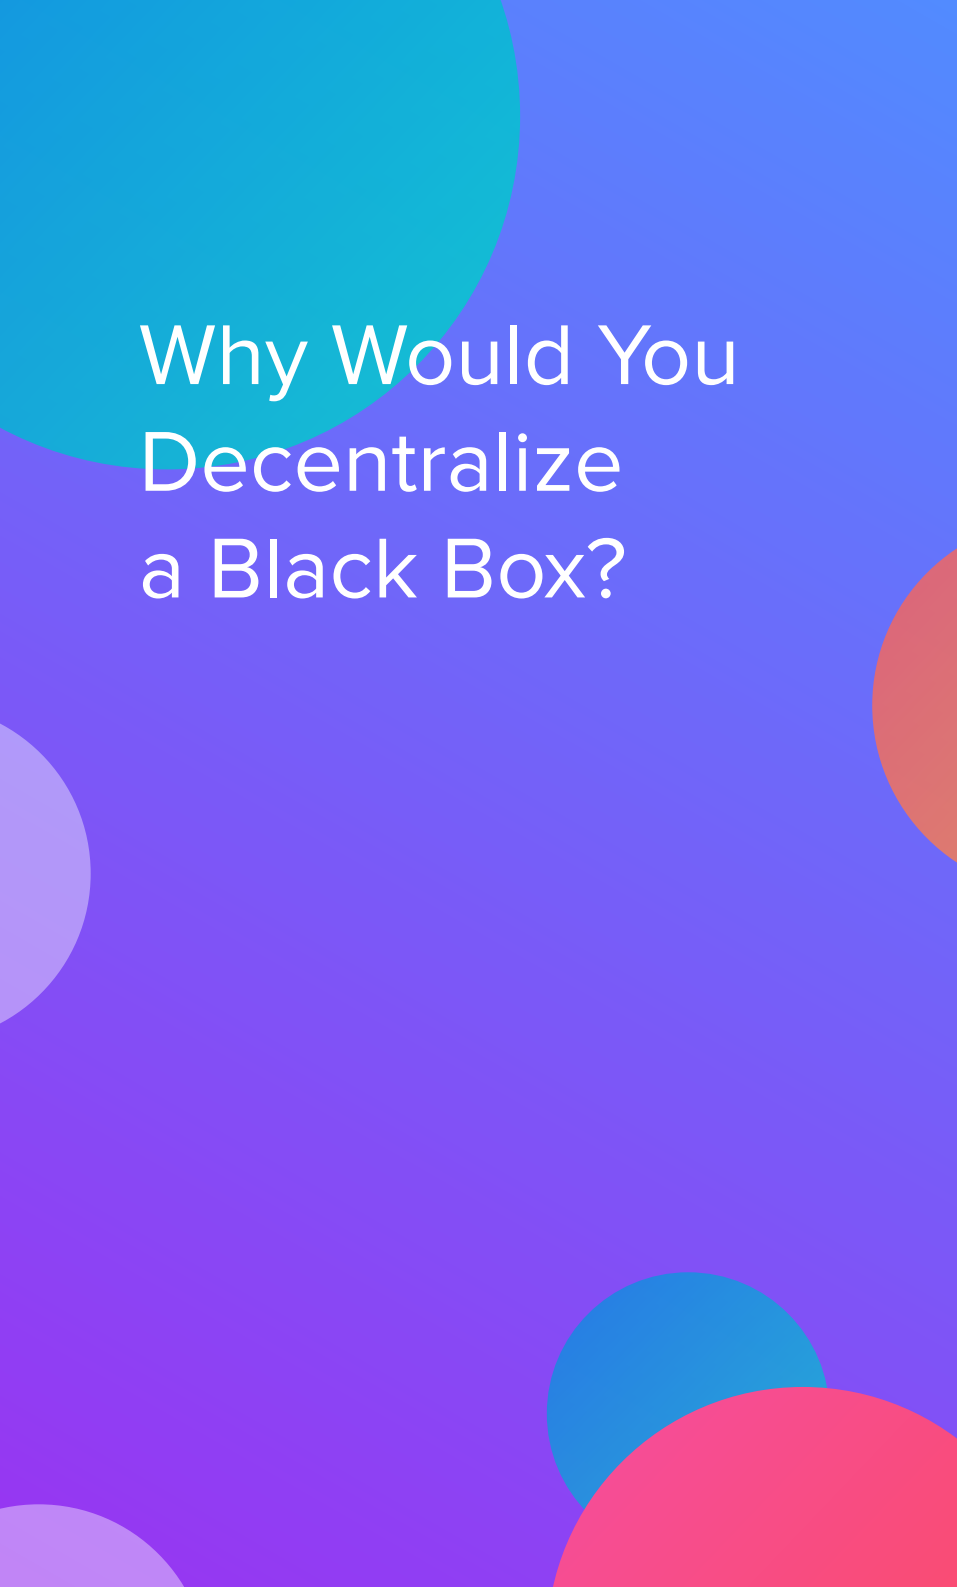
What is a Decentralized Black Box?

What sets them apart from other models of machine learning:

1. They are biologically (neurologically) inspired (connectionism).
2. They have numerous layers (deep) with reduced abstraction.
3. They may behave in surprising (unfiltered) and permissionless behavior, which is why sometimes they are called 'black boxes'.
4. Their performance improves per layer and per epoch; the learning never stops.
5. Exposed to new data, they can learn unsupervised.
6. They use randomness in the initial values (stochastic).
7. They improve because of backpropagation through time (BPTT).
8. The AI component is moved as far as possible to the edge of the network by performing prediction and computation at or near the source of the data (on laptops, smartphones, tablets and sensors).
9. The population of the data is build, via conversational AI with the machine.
10. The provenance of the data is timestamped cryptographically on a trustless, decentralized blockchain.
11. LSTMs and CNNs can be used in one architecture. With enough training (on the image side) the LSTM can work together with a CNN to clinically explain in natural language a diagnosis, at least on a primary level to a patient through a text interface. An example underneath by the Bengio group "Show, Attend and Tell: Neural Image Caption Generation with Visual Attention".

A Decentralized Black Box needs:

1. Large amounts of data.
2. Powerful Neural Network algorithms (CNNs and or LSTMs).
3. Deep learning architectures (Tensorflow, PyTorch etc.).
4. Powerful machines (GPUs).
5. A large community of seeders, developers, curators and data donors.
6. For conversational AI, extensive socializing with human users, is how it learns to fill its state machines with dialog roll-outs.



Why Would You Decentralize a Black Box?

John Mattison (CIO of Kaiser Permanente) famously said “The great question of the 21st century is going to be ‘Whose black box do you trust?’” Black boxes are unstoppable because their results are indisputably better than any other algorithm, which will make them ideal to place our trust in those systems whose methods for making decisions we do not understand and can reflect the bias of their creators.

Our answer to the Mattison question is: “I trust MY black box. I trust the black box that I have trained, the one into which I have input data and the one I can audit.” To achieve this goal, we need to democratize access to black boxes, which remains diabolically difficult for the uninitiated.

NEURON will provide a suite of the tools to everyone and show them step by step how to educate and train their own AI.

Problems and Opportunities

1. THE INTEGRATIONAL BURDEN

Our solution will guide the participant in finding and collecting its own medical data. Most people do not have access to their medical information, they do not know where to start and if they do, their knowledge of computer science is limited. We will make use of several technologies that doc.ai has built already. We will also foster and support an open-source developer community to help innovate tools on our tool platform stacks, to facilitate the integration and collection of data and to present algorithms to interpret the data: a decentralized Kaggle for personalized biology.

2. ORGANIC DATA

To make great predictions, we need to be able to audit the provenance of the data and perform data forensics and KYD (Know Your Data) processes. Too many algorithmic systems and neural nets exhibit undesired behavior because of noise or “self-reported” data.

We will use the blockchain as triple-entry accounting so we can track and authenticate the data sources.

3. UNBIASED DATA

“If you’re not a white male, artificial intelligence’s use in healthcare could be dangerous. Data coming from randomized control trials are often riddled with bias. The highly selective nature of trials systemically disfavor women, the elderly, and those with additional medical conditions to the ones being studied; pregnant women are mostly entirely ignored.”

4. PRIVACY CONCERNS

People may be hesitant to share their medical data over networks where strangers can be lurking. By decentralizing the data over all the users, the data is cryptographically timestamped and becomes immutable.

Furthermore we will address HIPAA requirements by maintaining the data on the edge device (and not in the cloud or on a centralized server).

5. VERTICAL INTELLIGENCE

While we remain far from creating a doctor in a machine, in medicine there are generalists (GPs) and Specialists. Generalists are analogous to GAI (General AI), unreachable at this point of technological development. But specialists are like vertical AIs and closer to realization. The ABMS (American Board of Medical Specialties) lists more than 150 medical specialties and subspecialties.

USE CASE 1

Parents Solve a Problem Together

In San Francisco, there is an informal organization of parents called RADIUS who have one thing in common: they have children with epilepsy, the seizures ranging from 4 to 78 per day.

Most have exhausted all current medical possibilities for decreasing seizures. Some keep diaries with values on their children, others have given up, and some have turned to alternative medicine.

All have data but this data is noisy and not uniform. Nobody has a model how these seizures work or what triggers them. Their only hope left is that they get used to it. One member tries NEURON.

How it Works

1. One member of RADIUS tries out NEURON and sets up the doc.ai application for her child.
2. The doc.ai application enables the mother to upload information about her child to create structured datasets.
3. By running the application, the mother aggregates a number of tokens.
4. At the next meeting of RADIUS she has made a slide show explaining what she did and how others can all join in.
5. Seeking to learn more about their child's seizures, RADIUS proposes a problem statement and associated bounty to the network.
6. The proposal is structured that anyone seeking to learn about the same problem may participate by funding the project.
7. RADIUS propose a bounty of \$500 per participant expecting to hit a target of \$30,000 and 60 participants with similar datasets.
8. RADIUS broadcasts the proposal on the network and 900 parents (a mesh group) join in. That makes the total bounty \$450,000!
9. Because the bounty is so big they decide to split it into 4 milestones, from great to greater to awesome to miraculous. Find the triggers!
10. Now the bounty is communicated to the doc.ai data scientist network with the goal to find a model that looks for an avoidable trigger of the seizures.
11. Once initial models are ready for broadcasting to the Neuron network, the network receives them and enhances the models using their own data by means of a distributed variant of the stochastic gradient algorithm.
12. Several data scientists answer the challenge and a leaderboard is produced with F1 scores.
13. The highest F1 score is selected, the prediction model is sent to the parents and the prize is awarded.
14. The parents now have the results of their own clinical trial and can consult with their own physicians about the implications.
15. They can even approach pharma to develop a new drug based on the new target, and offer their trial in return for equity of the drug they will buy.

USE CASE 2

The Clinical Trial of One

Marcus O' Connor is retired and is taking care of his partner who suffers from "metabolic syndrome": a cluster of conditions like increased blood pressure, high blood sugar, excess body fat around the waist, and abnormal cholesterol or triglyceride levels, that occur together, increasing one's risk of heart disease, stroke and diabetes.

Marcus is tired of going through the medical system without real progress and decides to do something. He is inspired by a new healthcare model called precision medicine, which integrates genetic information, microbiome data, and information on patients' environment and lifestyle to provide custom-tailored solutions.

He starts an experiment, which he calls "The Clinical Trial of One", on his wife. Marcus has no medical background.

How it Works

1. Marcus wants to begin by constructing an actionable, iterative medical (quantified biological) profile for his wife and do his own clinical trial.
2. He contacts NEURON and the AI welcomes him, listens to his story, and starts documenting and learning. He uses his laptop together with his wife to have long conversations in a text-based question and answer format.
3. The doc.ai application requests:
 - her facebook login, a picture of herself and her medicine cabinet, her to use some cheap wellness devices available on Amazon, her to fill in an online questionnaire
 - her latest blood results from LabCorp or Quest, a cheap genomic and microbiomic test (spit in a tube send in the mail)
4. By doing this Marcus now has a robust real-time medical record and biological profile on his wife, and the AI begins to give targeted education and helps them develop an initial plan of action.
5. Marcus has also been awarded tokens, which he can now use to buy a prediction study about her health.
6. He proposes a problem statement and associated bounty to the network.
7. Find me how I can improve my wife's situation and a plan for improving the quality of her life.
8. The proposal is structured that anyone seeking to learn about the same problem (and put in their own datasets) may participate by funding the project.
9. He proposes a bounty of \$100 per participant expecting to hit a target of \$5,000 and 50 participants with similar datasets.
10. He starts to recruit friends and acquaintances for the clinical study of one. The empowerment is exhilarating.
11. He broadcasts the proposal on the network and 100 interested parties join in. That makes the total bounty \$10,000.
12. Now the bounty is communicated to the doc.ai data scientist network.
13. Once an initial model is ready for broadcasting to the Neuron network, the network receives it and enhances the model using its own data by means of a distributed variant of the stochastic gradient algorithm.
14. Several data scientists answer the challenge and a leaderboard is produced with F1 scores.
15. The highest F1 score is selected, the prediction model sent to the participants and the prize is awarded.
16. The members now have the results of their own clinical trial and go and seek medical opinions.
17. The group is contacted by a patient organization and together they prepare a therapeutic plan for all those interested in metabolic syndrome.
18. Marcus and his wife still enjoy daily the advice and education of their own medical AI, which they have now put into speak mode instead of text mode and which coaches them through life. The information is automatically updated in real time based on changes in the data sets (blood results, weight, resting heart rate, sleep quality, etc.).
19. Now and then the AI connects with other AIs in the group which has now grown to be 23,000 members and exchanges new tips and techniques and latest publications on metabolic syndrome.

USE CASE 3

AI-First Health Insurance Company

A health insurance company takes up the idea to become an AI-first Health Insurance by offering its members something unique: their own medical black box, a neural net that knows everything about their health and can feed them coaching, care and prediction, even behavioral change.

They can talk to this “deep mind” 24/7 and it will get to know them better and better. It will also flawlessly connect with their doctors, their pharmacist and other carbon-based units in the system.

The insurer does not want access to this information, and cannot, because the data resides on the edge device (smartphone) of the customers.

Their theory is that this service will significantly decrease the healthcare costs of both the customers and the insurance company.

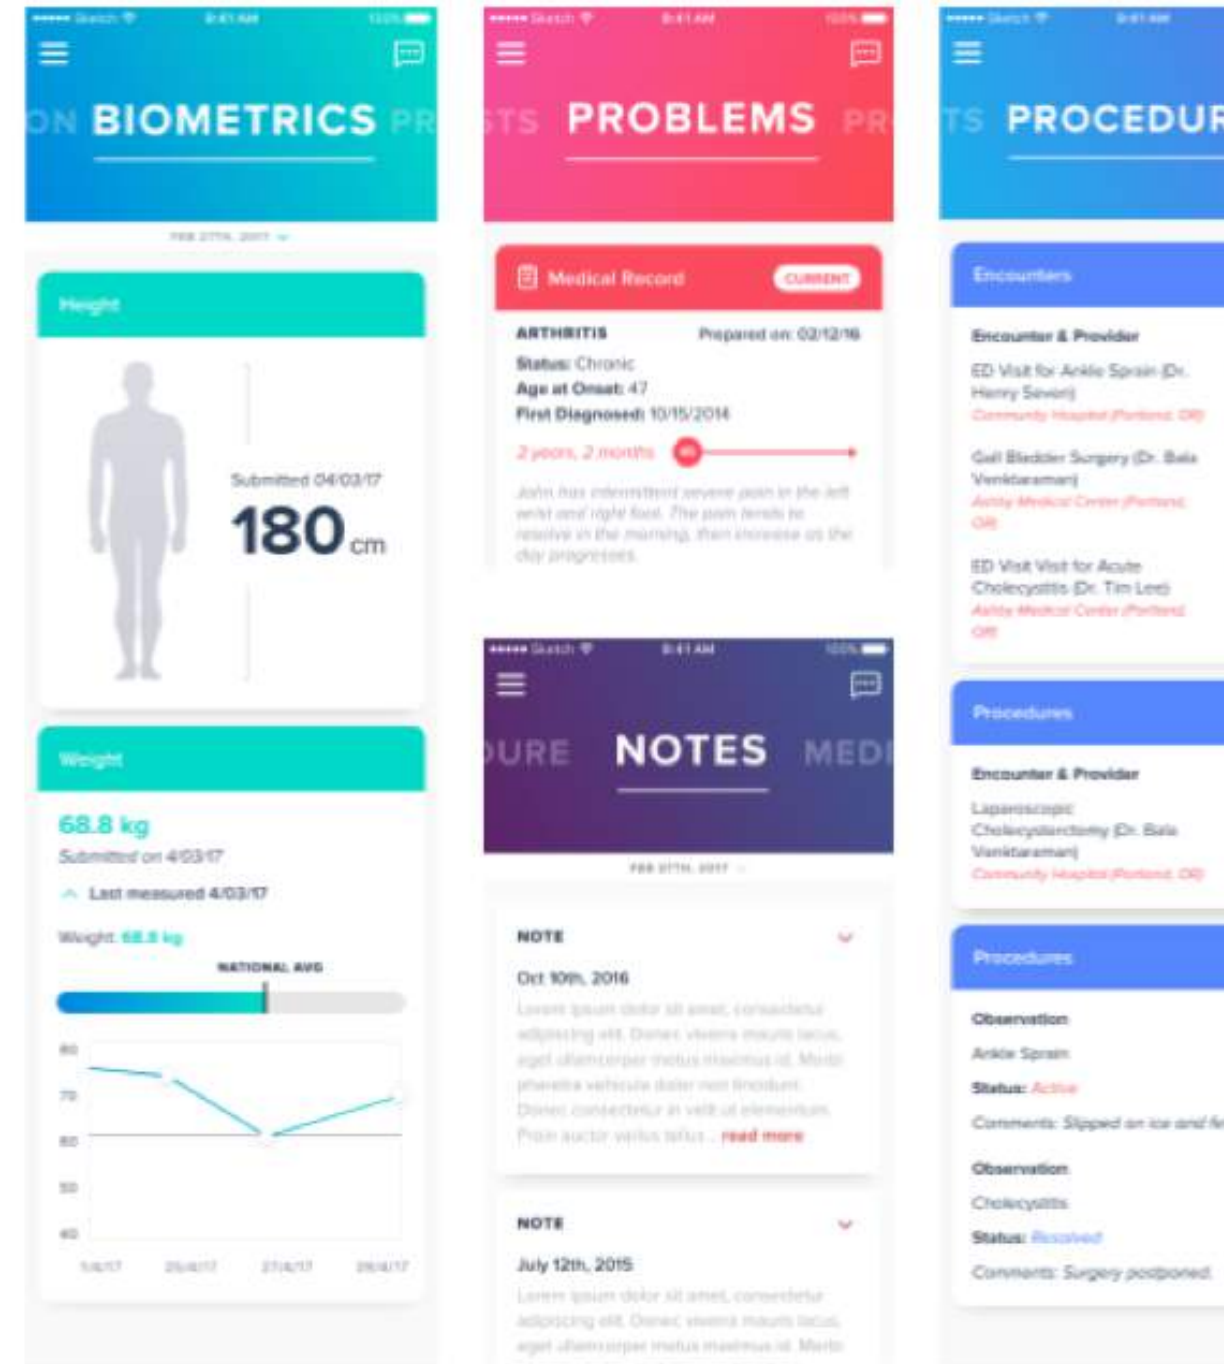
How it Works

1. The Insurer employs a consultancy company to help the first 100,000 customers to get AI-ready, envisioning the network will then scale by itself because of network effects.
2. The insurance company sets up a waiting list for the first 100,000 and sends out a mailing campaign to its customers: be one of the first in the world to get your own medical AI. They have two weeks to sign in on the Neuron Network. This is a free service and they will have their own robo-doctor 24/7. In the future they will even be able to monetize their health data on a monthly basis by leasing it to pharma companies.
3. The consultancy company starts in batches of 10,000 users. Everyone gets a number in the queue. For the account of the insurer they buy tokens to use the doc.ai application to onboard the members. They set up interviews with them on Skype to walk them through the process and distribute instruction videos on YouTube. They have an overall monitoring screen with progress bars per user.
4. The customers aggregate tokens by onboarding, which they can now use to buy prediction studies about their health or just cash them in on an exchange for dollars.
5. After a while the first 20,000 are onboarded, and the insurer now launches a competition on the network with an associated bounty: cluster features of my members and connect predictions to it.
6. Several data scientists answer the challenge and a leaderboard is produced with F1 scores.
7. The highest F1 score is selected and the prediction model is sent to the insurer, who sends it to its members. All members in the network belong to a cluster or a mesh where they can compare with others and receive predictions and plans to optimize their health.
8. The members can now also connect their robo-doctors horizontally in their clusters and start social groups to help each other.
9. The members now have the results of their own clinical trial and go and seek medical opinions.
10. The insurer monitors the progress of the project and calculates the results.

How to Train Your Own Black Box

We have several products in beta that have been mentioned in the use cases. These will guide and teach users how to train their decentralized AI (basically, how to train the trainer). These products will show users how to construct datasets of their health, and where and how to access these datasets.

Personal Portable Biology File. The instruction set needed to construct a personalized biology profile which is global, portable, user-friendly and quantifiable. This file has multi-signature capabilities.





Onboarding Module with Computer Vision

Who likes to fill in data on a screen the size of your hand or even on a desktop? Nobody. Our tools populate themselves with frictionless help from their users.

Selfie2BMI module uses state-of-the-art Deep Neural Networks and optimization techniques to predict a variety of anatomic features including height, weight, BMI, age and gender from face. Besides these vital anatomies, it also monitors 23 facial attributes like skin, receding hairlines, wrinkles, teeth and other attributes.

Blood Test Decoder

Deep conversational agent designed to enhance the post-blood test experience, enabling the user to discuss and answer any question on 400+ blood biomarkers.

It is trained on hundreds of thousands of medical documents and common FAQs to answer complex questions about blood results. The agent can personalize the conversation based on the user's age, gender and pre-conditions to provide relevant answers and educate using interactive content.





Genomics Test Decoder

Deep conversational agent designed to enhance the genetic counseling experience by providing answers from simple educational questions to complex personalized questions. It has a memory and remembers every visit and recommendation it has given.

When it does not know, it goes to look for the answer in a massive dataset. It can also pass you on to its carbon-based colleague.

Medicine Decoder

Deep conversational agent trained on medication dosage, side effects and other guidelines to answer personalized questions.

It will be connected to the pharmacogenomic recommendation engine if genomic results are present.




Our Team



Founders



Key Advisors



Artificial Intelligence - Computer Science - Applied Cryptography



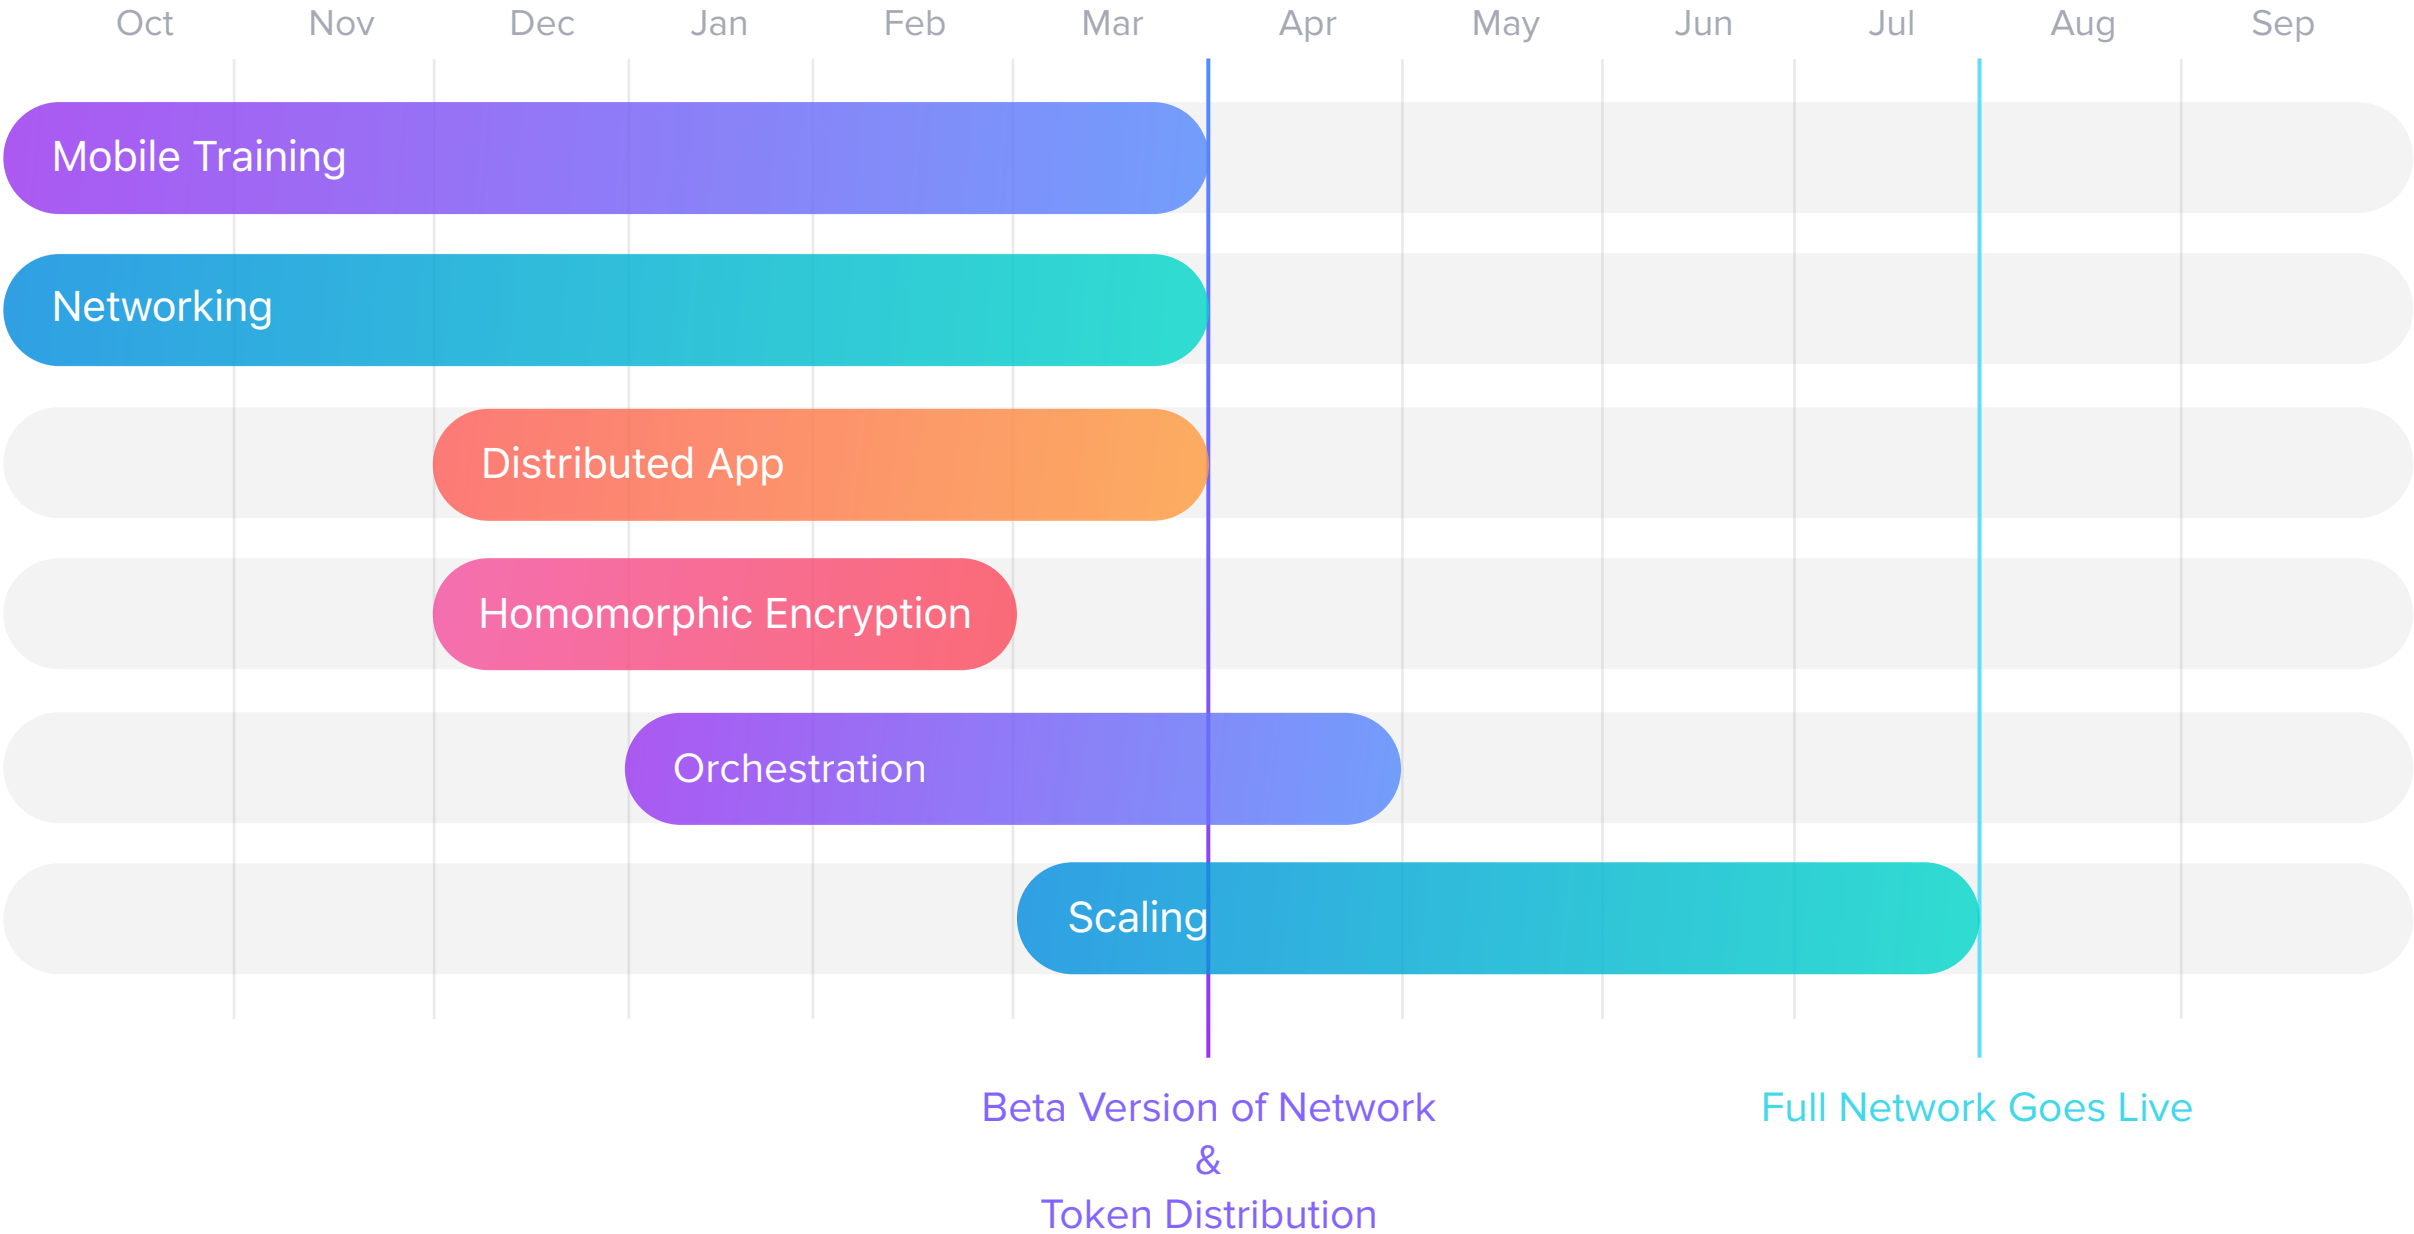
Benign Leadership



Belief System

1. It is time to decentralize machine intelligence from large corporations to the hands of consumers.
2. People are ready for a new healthcare paradigm, where they will own their data on intelligent edge devices.
3. Healthcare can be personalized by training an individual's data through deep layered intelligence.
4. 24/7 healthcare by machine intelligence will eliminate medical errors and improve health literacy.
5. Machine intelligence should be evenly distributed and open source.
6. Decentralization protects us from cybersecurity threats in the future.
7. Distributed deep learning is revolutionary and an indispensable step towards a non-dystopian future.


Research Roadmap Oct. 2017 - Sep. 2018



Use of Funds

		TODAY	\$5M	\$10M	\$15M	\$20M	\$25M	\$30M+
TEAM	AI/ML	4	5	6	7	8	9	10
	NLP/CL	1	2	2	2	3	3	3
	Blockchain	3	3	4	4	4	5	5
	Comp. Science	3	4	4	5	5	6	6
	Design	1	2	2	2	3	3	3
	Life Science	2	3	3	4	4	4	4
	Finance	1	1	1	2	2	2	2
	Community	0	1	1	3	4	5	6
OUR VISION	Community	Hackathon, Meetups	Dev Conference Palo Alto, US	Dev Conference Shanghai, China	Dev Conference Paris, France	Grants to students Distributive AI systems	Repository/p2p Decentralized AI articles	Endowed professorship
	Edge Networking	Centralized Products	Decentralized architecture	Distributive architecture	Data on edge	Classification on edge	Computing on edge	Training on edge
	Education	Forums	Video Tutorials	Dev program	MOOC	Partner with universities	Train the (human) trainer	Train the (machine) trainer
	Product	Apps	Dapps	IDE	Fully automated AI	Frictionless AI with predictions	Model available	Predictions in your hand
	Platform Users	200	10,000	50,000	100,000	200,000	300,000	400,000

Any financial or operational projections expressed are for illustrative purposes only. Actual results may differ materially. See "Forward-Looking Statements" and "Risk Factors" in our private placement memorandum. While based on assumptions that management believes are reasonable, you are cautioned not to place undue reliance on any forward-looking statements.



Contact Us

info@doc.ai