
NeuRoN: Decentralized Artificial Intelligence, Distributing Deep Learning to the Edge of the Network

Walter De Brouwer¹ and Mason Borda¹

{walter, mason}@doc.ai

September 7, 2017

¹doc.ai

Abstract

NeuRoN is a *decentralized artificial intelligence network*. It is a distributed and permissionless network of users with data ownership necessary to train neural network models. A publicly accessible network democratizes access to artificial intelligence. Thus, simultaneously protecting user data and incentivizing the use of their data for the greater good of humanity.

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. Furthermore, NeuRoN will function as a token-based economy. Participants will be required to prove ownership of tokens and receive tokens for participating. This enables broader participation and incites greater desire to interact amongst those qualified to discover new insights within our decentralized artificial intelligence platform.

Contents

1	Introduction	4
1.1	Overview	4
1.2	Blockchain	4
1.3	Decentralized Artificial Intelligence	4
2	An Edge Learning Network	4
2.1	Post-Cloud: Towards Fog and Edge	4
2.2	Multi-agent Systems	5
2.3	Artificial Intelligence Today	5
2.4	Participants	5
2.5	Learning	5
2.6	Proposal	6
2.7	Training	6
2.8	Synchronization	7
2.9	Payout	7
2.10	Stochastic Gradient Descent	8
2.11	Downpour SGD	8
3	Focused Domain: Personalized Biology	9
4	Problems and Opportunities	10
5	Three Use Cases	11
5.1	Parents Solve Problem Together	11
5.1.1	Scenario	11
5.1.2	How it works	11
5.2	The Clinical Trial of One	12
5.2.1	Scenario	12
5.2.2	How it works	12
5.3	AI-first Health Insurance Company	13
5.3.1	Scenario	13
5.3.2	How it works	13
6	Train Personalized Models	13
7	Conclusion	16
8	Securities Disclaimer	16

List of Figures

1	Learning phases (Loop)	6
2	Learning phases (Sequential)	6
3	Proposal	7
4	Training	7
5	Payout	8
6	Downpour SGD: Single Client	9
7	Downpour SGD: Multiple Clients	9
8	Personal Portable Biology File	14
9	Onboarding module	14
10	Blood Test Decoder	15
11	Genomics Test Decoder	15
12	Medicine Decoder	16

List of Algorithms

1	Downpour SGD: Processing by worker i	9
2	Downpour SGD: Processing by master j	10

1 Introduction

1.1 Overview

Mobile phones and tablets are now the primary computing devices for many people. In many cases, these devices are rarely separated from their owners, and the combination of rich user interactions and powerful sensors means they have access to an unprecedented amount of data, much of it private in nature. Models learned on such data hold the promise of greatly improving usability by powering more intelligent applications, but the sensitive nature of the data means there are risks and responsibilities to storing it in a centralized location. NeuRoN is an Ethereum-based blockchain that tackles the paradigm of decentralizing artificial intelligence.

1.2 Blockchain

The blockchain enables parties to interact based on a set of agreed upon rules. For example, such a rule can define payment transfers. These generic rules are referred to as smart contracts. The Ethereum network [1] is a platform that enables the creation of peer-to-peer applications based on smart contracts. This allows developers to create cryptographically enforceable relationships. Thus, NeuRoN will use smart contracts to provide the network infrastructure it needs, namely incentivizing developers, end users, and research organizations to interact in our artificial intelligence environment through the use of a token.

1.3 Decentralized Artificial Intelligence

Decentralized artificial intelligence has a couple of benefits; namely, data privacy and creating a collaborative atmosphere. Data privacy is ensured by communicating machine learning models back and forth and keeping data on the end user's device. Furthermore, once models have matured, they are accessible to everyone on the network. Resulting in the removal of a centralized proprietary organization, which currently ultimately determines the destiny of future discoveries.

2 An Edge Learning Network

Distributed deep learning is the most intriguing future-proof challenge today because the world of AI and Blockchain have not yet intersected meaningfully. When they do, intelligence and the product of intelligence, prediction, itself will be attached to every known entity with open access. To start this adventurous journey we have chosen a field that is very close to the human experience: our health, and our unique biological data that makes us all the same but entirely different. Neuron will provide a suite of distributive, AI tools accessible to everyone with step by step instructions how to educate and train their own AI, to eventually understand the dynamic processes and concepts that underlie prevention and lead to health as a continuous function and not a discrete one. NeuRoN will integrate economic incentives for everyone in the network to be rewarded in NeuRoN tokens for contributing to the complex brain that will eventually bind all AIs into one.

2.1 Post-Cloud: Towards Fog and Edge

The Cloud has become global data warehouse. But two new technologies are already waiting to provide a better alternative: Fog and Edge [2]. Fog computing extends cloud computing and services to the edge of the network, bringing the advantages and power of the cloud closer to where data is created and computed. Sometimes Fog and Edge are used as synonyms but they are in fact different when we talk about AI, compliance and security. Edge reduces the points of failure, as each device independently operates and determines which data to store locally and which data to send to the cloud for further analysis. With AI more and more real-time processing will be required, which means more intelligence is needed in or near to local devices. This local processing is known as the 'intelligent edge'.

2.2 Multi-agent Systems

Multi-agents systems are used to solve problems that are difficult to solve by individual agent. Multiple-agent communication technologies can be used for management and organization of computing fog and act as a global, distributed operating system, which combines decentralized P2P general-purpose computing tasks distribution, multiple-agents communication protocol and smart-contract based rewards, powered by a blockchain.

The Blockchain of EVM is an example of swarm intelligence. May miners work as reactive agents, who without any control of agent-supervisor carry out the work on maintenance of the network, moving only in accordance with their own motivation [3].

2.3 Artificial Intelligence Today

Current work in machine learning has shown that larger models can dramatically improve overall performance [4]. With the advent of deep learning, the field is rapidly expanding. However, large neural network models face infrastructure limitations. These limitations can be overcome in several ways. One approach is to develop more robust Graphical Processing Units (GPUs) and Tensor Processing Units (TPUs) that can handle the computational load. Another approach is to distribute the load to clients speaking to a central server for batch training updates. Thus it is possible to crowdsource the computational power necessary to train models and incentivize participation via a token-based network.

2.4 Participants

The parties involved in the NeuRoN learning phase are defined as the following:

1. Research organizations - Parties interested in using NeuRoN to train their models.
2. Users - Participants interested in connecting their devices and data to the network.
3. Neural Network Models - Models trained with the user's data.

These three parties interact in order to train the students or neural network models that form the basis of the artificial intelligence. The phases which they collectively undergo constitute the learning. The phases involved in the *learning* are proposal, training, synchronization and payout.

2.5 Learning

Learning consists of a model being broadcast to the clients to learn from the data. A learning cycle has several phases:

1. Proposal
2. Training
3. Synchronization
4. Payout

The proposal phase consists of servers listening and monitoring to determine whether a target dataset exists for training (“teaching”) a neural network model. The dataset \mathcal{D} in this network consists of n clients each with their own data point to contribute. After the dataset and associated clients have been targeted, the training phase begins. The training phase consists of some initial weight w_0 being broadcast to the clients. Client c_i then uses this initial weight to calculate the associated gradient weight w_i . This yields a set W of n gradients with which to update the initial weight w_0 .

When the iteration is completed, the gradients are sent back to the server for a synchronization phase. The server receives and compiles each gradient w_i from all associated clients into a set W of n gradients with which to perform the stochastic updates on w_0 as described in the “Stochastic Gradient Descent” section. This updated weight is broadcasted to the same network of clients for a synchronization of the new model. Following a synchronization and initial epoch on the dataset, the clients are paid out. NeuRoN interactions are depicted as a loop in Figure 1, and sequentially in Figure 2.

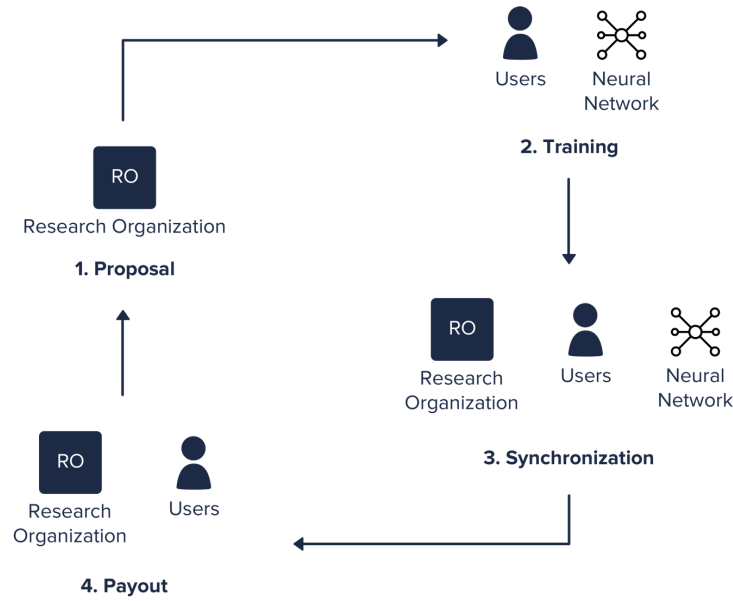


Figure 1: Learning phases (Loop)

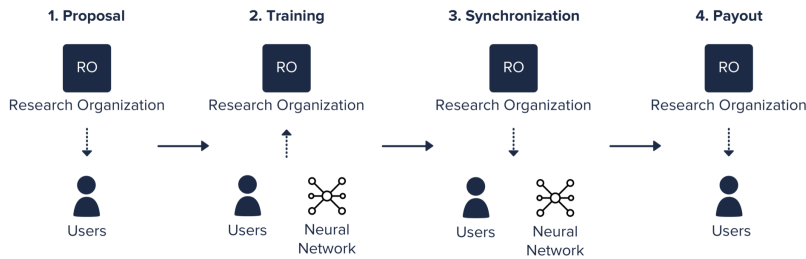


Figure 2: Learning phases (Sequential)

2.6 Proposal

The proposal phase, shown in Figure 3, consists of a broadcast of the curriculum or the dataset available for training. When the server determines there are enough peers with a particular curriculum, the server will initiate an proposal. This requires the server to select the users or peers that they wish to initiate a learning with, and to undergo an initial teaching. In order for the stochastic gradients to be properly calculated, the dataset which is iterated over must be of a fixed sample size. As a result, the users must be present for the duration of the learning and the curriculum or dataset must be fixed. The initial teaching phase consists of an initial weight being calculated.

This initial weight, along with the model, is passed to the users for the teaching phase. For the proposal period, the server must determine how it would like to compensate its users. Users can be paid per epoch or for completing a set of epochs on a particular dataset.

2.7 Training

The initial training involves the transfer of parameters necessary to undertake the training on the first client. The initial training is the first in a sequence of identical experiences where the model is trained by users asynchronously. The first training consists of the transmission of the initial stochastic gradient descent weight and model to be passed to the device of the users.

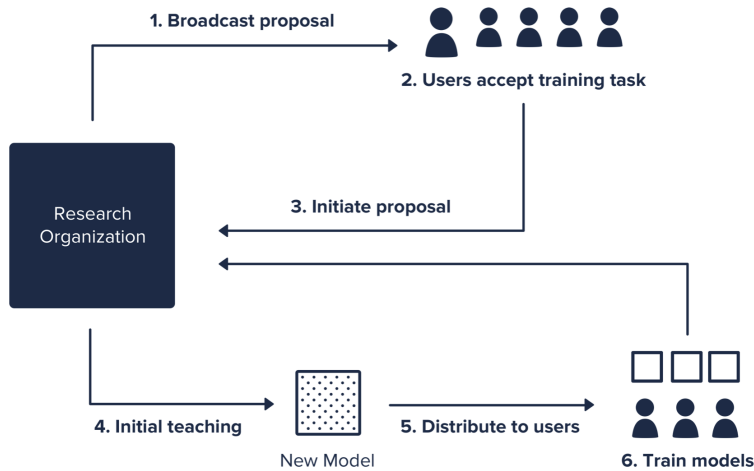


Figure 3: Proposal

The user will teach the neural network based on the locally stored dataset. The neural network will return only its calculated delta in the weight between the knowledge it initially entered with and the knowledge it has attained for the day (or iteration on the dataset), shown in Figure 4.

The neural network will visit every user's device until it has learned or calculated a stochastic gradient descent weight from each. This completion will signal a full epoch on the dataset.




Figure 4: Training

2.8 Synchronization

A synchronization occurs after the neural network has completed an epoch. This consists of the server aggregating all of the weights or learnings and calculating a new model or a new initial weight to provide to the network. This initial weight is again broadcast to the users if the curriculum and users are still available for a second epoch.

2.9 Payout

During the initial proposal period, users are provided with the amount in tokens that the server is willing to pay. When the number of required epochs by the server are completed, users are compensated for teaching the neural networks.

Compensation is designed to operate as a free market in order to reflect the supply and demand of the data available to the network. When a proposal is received, the users have the opportunity to accept or reject the offer. The intent is to encourage the research organizations to pay out a fair and attractive amount to users in order to encourage participation in the network.

A key attribute for any open network is a mechanism for encouraging good behavior or alternatively discouraging bad behavior on the network. Bitcoin [5] requires a fee per attempted transaction in




Figure 5: Payout

order to limit bad behavior. If a malicious party attempts to perform a DDoS or Sybil attack on the network, it will cost a fee per transaction. If the fee is not attractive to the network of miners, the transaction will simply not be accepted in a proposed block. If the attacker proposes transactions that contain an attractive fee, the fee acts as a disincentive to propose the transaction. If the attacker proposes a high volume of transactions in a short period of time, the demand on the mining network will increase. This will increase the cost of participating in the network [6].

NeuRoN will require a similar mechanism as the network will not be able to objectively assess the quality of data or the intent of every participant in the network. In order to reduce the potential for DDoS or Sybil attack on the network, a buy-in will be required to participate with the network in order to enforce a monetary barrier to entry to DDoS or Sybil attack. These tokens will be returned at the end of the training phase or when the research organization determines to conclude the learning.

2.10 Stochastic Gradient Descent

Stochastic Gradient Descent (SGD) is a machine learning algorithm which iterates through a data set several times (epochs) to fit the best weights to a model. In order to learn, the model must compute a loss function that maps a datapoint’s prediction and its ground-truth value to a quantitative value representing the “cost” or “correctness,” depending on the sign of the loss function. The loss function itself varies based on the domain of the task. Then, the gradient of the loss is back-propagated to the model weights, marking an update. Unlike traditional gradient descent, SGD makes updates after each datapoint. This allows faster convergence, computational efficiency and potential to be decentralized over a network.

2.11 Downpour SGD

Applying a variant of SGD, referred to as Downpour SGD [7], allows for further decentralization. Specifically, the update framework is no longer constrained to perform updates after each datapoint. Downpour SGD can asynchronously compute gradients for each datapoint and perform a global update after each epoch.

We define the algorithm for Downpour SGD to be $w' = w - \eta \nabla w$. Where the parameters are defined as follows:

1. w – The parameters of the model
2. η – Learning rate or the size of the steps we take to reach a local minimum
3. ∇w – Gradient of the objective function with respect to the parameters
4. w' – The updated newly calculated parameters of the the model

Applying Downpour SGD to a single client-server model, the client has the ability to train on an individual iteration of the data. The calculated gradients (weight updates) are then sent back to the server. The server then performs a global model weight update, and propagates the new model back to the client. The single client-server network design is depicted in Figure 6 and a multi-agent topology is shown in Figure 7.

Due to the nature of the Downpour SGD, the weights that are passed back to the server need not be in order. Such a network can enable an edge learning network where artificial intelligence models compensate the clients for data and NeuRoN participation. As the models are trained, the user is compensated via tokens for the learning performed on a per epoch basis. The worker and master algorithms are sketched out in Algorithm 1 and Algorithm 2, respectively.



Figure 6: Downpour SGD: Single Client

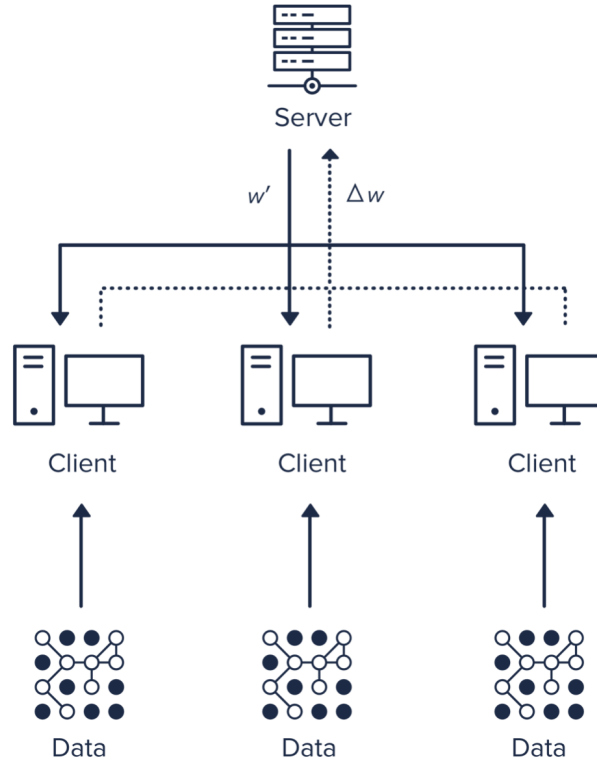


Figure 7: Downpour SGD: Multiple Clients

Algorithm 1 Downpour SGD: Processing by worker i

Initialize: $w_i = \tilde{w}$

repeat

 Receive \tilde{w} from the master: $w_i \leftarrow \tilde{w}$

 Compute gradient ∇w_i

 Send ∇w_i to the master

until forever

3 Focused Domain: Personalized Biology

The medical field is experiencing a large influx of statistical modeling and machine learning. Already we are discovering an increased number of insights. This will likely continue and will grant many benefits to millions of people. Furthermore, the exponential scaling of data will result in a positive

Algorithm 2 Downpour SGD: Processing by master j

Input: learning rate η
Initialize: \tilde{w} is initialized randomly, $v_i = 0$
repeat
 Receive ∇w_i asynchronously from some worker
 Update: $\tilde{w} = \tilde{w} - \eta \nabla w_i$
until forever

feedback loop of greater insights. As such, we choose personalized biology as the initial use case for following reasons:

1. High impact domain.
2. High data growth (expected to double every 73 days by 2020 [8]).
3. High priority privacy.

NeuRoN allows research organizations and individuals to keep their data private with the option to monetize by participating in studies that will help the advancement of medicine. A research organization may propose a medical-oriented artificial intelligence problem to the individuals on the network with an associated bounty for participation. If the proposal is accepted by the network, the artificial intelligence models are first generated with the research organization. Once the models are ready for broadcasting, the network receives them and enhances the model using its own data by means of a distributed variant of the stochastic gradient algorithm. The doc.ai application will sit atop the NeuRoN network, enabling individuals to cooperatively participate and learn more about their personal biology.

4 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 ¹. 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 [9] 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, elderly, and those with additional medical conditions to the ones being studied; pregnant women are mostly entirely ignored.” [10]
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

¹See page in this white paper page 16, “How to Train your Own Black Box”

and closer to realization. The ABMS (American Board of Medical Specialties) lists more than 150 medical specialties and sub-specialties ².

5 Three Use Cases

5.1 Parents Solve Problem Together

5.1.1 Scenario

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.

5.1.2 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 being 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.

² ABMS guide to medical specialties: http://www.abms.org/media/114634/guide-to-medicalspecialties_04_2016.pdf

5.2 The Clinical Trial of One

5.2.1 Scenario

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.

5.2.2 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.
 - (a) her facebook login
 - (b) a picture of herself and her medicine cabinet
 - (c) her to use some cheap wellness devices available on Amazon
 - (d) her to fill in an online questionnaire
 - (e) her latest blood results from LabCorp or Quest
 - (f) 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.

5.3 AI-first Health Insurance Company

5.3.1 Scenario

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.

5.3.2 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.

6 Train Personalized Models

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. See Figure 8.

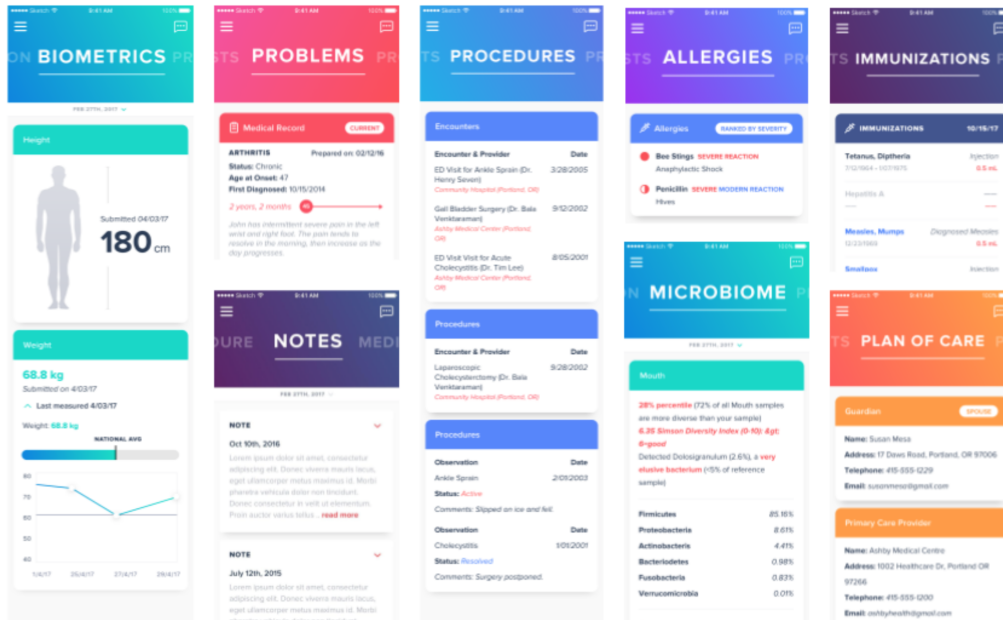


Figure 8: Personal Portable Biology File

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. Resulting in the output displayed in 9.




Figure 9: Onboarding module

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. Interactions are depicted in Figure 10.

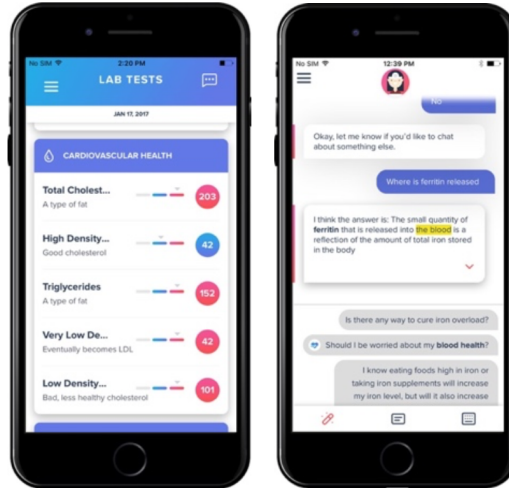


Figure 10: Blood Test Decoder

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 it in a massive dataset. It can also pass you on to its carbon-based colleague. This is realized in Figure 11.

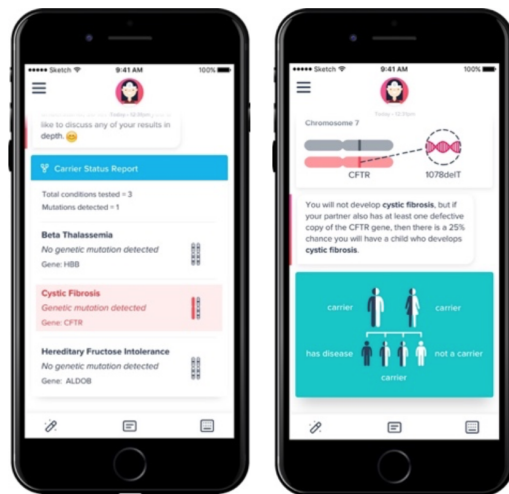


Figure 11: Genomics Test Decoder

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. A sample interface is shown in Figure 12.

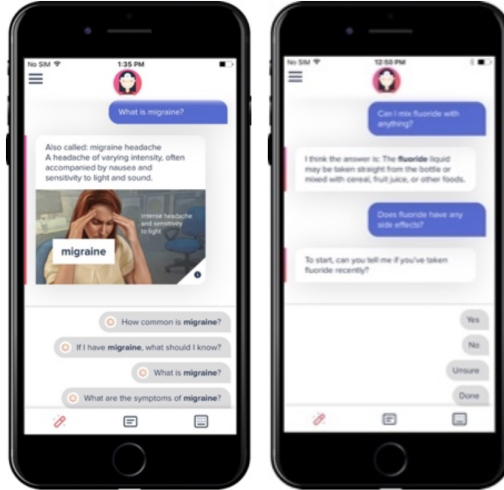


Figure 12: Medicine Decoder

7 Conclusion

We propose a Ethereum-based, open decentralized artificial intelligence platform where all can participate in collectively training neural network models. NeuRoN will play an essential role in the democratization and broad adoption of artificial intelligence. This freedom will allow artificial intelligence to succeed in making the world a better place, when provided with an unrestricted amount of data, development, and peer-to-peer supervision.

8 Securities Disclaimer

This document is for informational purposes only and does not constitute an offer or solicitation to sell securities. Any such offer or solicitation will be made only by means that are in compliance with applicable securities and other laws, including Rule 506(c) of Regulation D. No information or opinions presented herein are intended to form the basis for any purchase decision, and no specific recommendations are intended. Accordingly, this document does not constitute investment advice or counsel or solicitation for investment in any security. This document does not constitute or form part of, and should not be construed as, any offer for sale or subscription of, or any invitation to offer to buy or subscribe for, any securities, nor should it or any part of it form the basis of, or be relied on in any connection with, any contract or commitment whatsoever. DOC.AI expressly disclaims any and all responsibility for any direct or consequential loss or damage of any kind whatsoever arising directly or indirectly from: (i) reliance on any information contained in this document, (ii) any error, omission or inaccuracy in any such information, and (iii) any action resulting therefrom.

References

- [1] Ethereum, “ethereum/wiki,” 2014. [Online]. Available: <https://github.com/ethereum/wiki/wiki/White-Paper>
- [2] N. Mohan and J. Kangasharju, “Edge-fog cloud: A distributed cloud for internet of things computations,” *CoRR*, vol. abs/1702.06335, 2017. [Online]. Available: <http://arxiv.org/abs/1702.06335>
- [3] S. Ponomarev and A. E. Voronkov, “Multi-agent systems and decentralized artificial superintelligence,” *CoRR*, vol. abs/1702.08529, 2017. [Online]. Available: <http://arxiv.org/abs/1702.08529>
- [4] D. W. E. Noga, “8 ways machine learning is improving companies’ work processes,” Jun 2017. [Online]. Available: <https://hbr.org/2017/05/8-ways-machine-learning-is-improving-companies-work-processes>
- [5] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” May 2009. [Online]. Available: <http://www.bitcoin.org/bitcoin.pdf>
- [6] “Sybil attack,” Jul 2017. [Online]. Available: https://en.wikipedia.org/wiki/Sybil_attack
- [7] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, M. aurelio Ranzato, A. Senior, P. Tucker, K. Yang, Q. V. Le, and A. Y. Ng, “Large scale distributed deep networks,” in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2012, pp. 1223–1231. [Online]. Available: <http://papers.nips.cc/paper/4687-large-scale-distributed-deep-networks.pdf>
- [8] “Ibm’s new watson health unit changes the game,” 2015. [Online]. Available: <https://www-03.ibm.com/press/us/en/pressrelease/46580.wss>
- [9] “Triple entry accounting,” 2005. [Online]. Available: http://iang.org/papers/triple_entry.html
- [10] Commentary, “If you’re not a white male, artificial intelligence’s use in health-care could be dangerous,” Jul 2017. [Online]. Available: <https://qz.com/1023448/if-youre-not-a-white-male-artificial-intelligences-use-in-healthcare-could-be-dangerous/>