Cortex - AI on Blockchain

The Decentralized AI Autonomous System

Ziqi Chen
zchen@cortexlabs.ai

Weiyang Wang
weiyang.wang@cortexlabs.ai

Xiao Yan
xiao.yan@cortexlabs.ai

Jia Tian
mickey@cortexlabs.ai

Abstract

In the current blockchain world, the chain of built-in Turing Complete smart contracts is widely used, attracting a large number of application developers. However, due to the high cost of over-idealized World Computer concept, smart contracts limit their capabilities at design stage and do not fully exploit Turing Complete immense computational potential. As a result, developers are limited to write short programs and access only a very small amount of resources. While the proliferation of common smart contracts depends on the performance gains of new technologies, some extremely useful routines can be introduced ahead of time and can be applied with reasonable optimization and hardware support. This article describes a new public chain, Cortex. By revising and extending the instruction set, Cortex adds AI algorithms support for smart contracts so that anyone can add AI to their smart contracts. At the same time, Cortex has proposed an incentive mechanism for collective collaboration that allows anyone to submit and optimize models in Cortex, and the contributors to the models can be rewarded. Just as what's happening in some areas, thanks to Cortex's openness and sharing features, Cortex is set to create many models that transcend human capabilities. At the same time, as a social experiment, we also look forward to the Artificial General Intelligence (AGI) being born on the Cortex.

CRYP2NOVA

Archived copy · cryp2nova.com

IMPORTANT: YOU MUST READ THE FOLLOWING DISCLAIMER IN FULL BEFORE CONTINUING

The sale ("**Token Sale**") of the Cortex Coin ("**Cortex Coins**"), the exchange medium for participants on the Cortex platform as detailed in this whitepaper (the "**Whitepaper**") is only intended for, made to or directed at, only certain persons. Moreover, this Whitepaper is not a prospectus or offer document of any sort and is not intended to constitute an offer of securities of any form, units in a business trust, units in a collective investment scheme or any other form of investment, or a solicitation for any form of investment in any jurisdiction. No regulatory authority has examined or approved of any of the information set out in this Whitepaper. This Whitepaper has not been registered with any regulatory authority in any jurisdiction.

By accessing and/or accepting possession of any information in this Whitepaper or such part thereof (as the case may be), you represent and warrant to Cortex Labs Pte. Ltd. (Singapore Company Registration No.: 201803307C) (the "**Project Company**") and NeoCortex Global Limited (BVI Company No.: 1967270) (the "**Token Issuer**") that:

(a) you are not located in the People's Republic of China and you are not a citizen or resident (tax or otherwise) of, or domiciled in, the People's Republic of China;

(b) you are not located in the United States of America and you are not a citizen, resident (tax or otherwise) or green card holder of, or domiciled in, the United States of America unless you are a U.S. Qualified Person (as defined herein);

(c) you are not located in a jurisdiction where the Token Sale is prohibited, restricted or unauthorized in any form or manner whether in full or in part under its laws, regulatory requirements or rules;

(d) you agree to be bound by the limitations and restrictions described herein; and

(e) you acknowledge that this Whitepaper has been prepared for delivery to you so as to assist you in making a decision as to whether to purchase the Cortex Coins.

CRYP2NOVA

Archived copy · cryp2nova.com

Nomenclature

AI DApps	Decentralized Artificial Intelligence Applications
ASGD	Asynchronous Stochastic Gradient Descent
BFT	Byzantine Fault Tolerant
CNN	Convolutional Neural Network
CVM	Cortex Virtual Machine
ERC 20	The Ethereum token standard - Ethereum Request for Comment
EVM	Ethereum Virtual Machine
FHE	Fully Homomorphic Encryption
GAN	Generative Adversarial Network
HE	Homomorphic Encryption
Kaggle	A centralized platform for machine learning researchers to submit and compete models
PoS	Proof of Stake
PoW	Proof of Work
RL	Reinforcement Learning
RNN	Recurrent Neural Network
VAE	Variational Autoencoder
zk-SNARKs	Zero-Knowledge Succinct Non-Interactive Argument of Knowledge
zk-STARKs	Zero-Knowledge Succinct Transparent ARguments of Knowledge
ZSL	Zero-Knowledge Security Layer

CRYP2NOVA

Archived copy · cryp2nova.com

Contents

1 Introduction **12**

1.1 Background 12

1.2 AI on Blockchain - Feasibility Analysis 13

1.3 AI DApps 14

2 System Architecture **15**

2.1 Expanding the capabilities of smart contracts and blockchain 15

2.1.1 Cortex Intelligent Inference Framework 15

2.1.2 Model Submission Framework 16

2.1.3 Smart AI contract 16

2.2 Model and Data Storage 17

2.3 Cortex Consensus Inference Criteria 17

2.4 Model Selection and Valuation 18

2.5 Consensus MechanismPoW Mining 18

2.6 Fraud Prevention and Model Screening 18

3 Software Solutions **18**

3.1 CVM: EVM + Inference 18

3.2 Cortex Core Instruction Set and Framework Standards 18

3.3 Cortex Model Representation Tool 20

3.4 Storage Layer 20

3.5 Model Indexing 20

3.6 Model Cache 20

3.7 Full Node Experiment 21

4 Hardware Solutions **22**

4.1 CUDA and RoCM 22

4.2 FPGA 23

4.3 Full Node Hardware Configuration Requirements - Multi-GPU and Legendary USB Mining 23

4.4 Hardware Modifications for Existing GPU Mining Farm 23

4.5 Mining and Computing on Mobile Devices and IoT 24

5 Applications and Future Work **24**

5.1 Application Cases 24

5.1.1 Information Service 24

5.1.2 Financial Services 25

5.1.3 AI Assistant 25

5.1.4 Simulation Environment 25

5.2 Future Work 25

Archived copy · cryp2nova.com

5.2.1 Data Privacy . . . 25
5.2.2 Block Size and TPS Improvement . . . 25
5.2.3 Adapter of AI Chips on Mobile Devices . . . 25

6 Roadmap **26**

7 Token Model **26**
7.1 Cortex Coin (CTXC) . . . 26
7.1.1 Rewards for Model Provider . . . 26
7.1.2 Cost for Model Provider . . . 26
7.1.3 Model Complexity and Endorphin Spend . . . 26
7.2 Token Distribution . . . 26
7.3 Token Issuance Curve . . . 27

8 Core Team Members **28**

9 Advisors **29**
9.1 Technical advisors . . . 29
9.2 Academic advisors . . . 29
9.3 Business advisors . . . 29
9.4 Organization partnerships . . . 30

10 Leading Investors **30**

A Cortex Mathematical Representations of AI on blockchain **30**

B Summary of Basic Types for Deep Learning **31**
B.1 Supervised Learning . . . 31
B.2 Unsupervised Learning . . . 31
B.3 Other Types of Learning . . . 32
B.3.1 Semi-supervised Learning . . . 32
B.3.2 Active Learning . . . 32
B.3.3 Reinforcement Learning . . . 32
B.3.4 Transfer Learning . . . 33

C Distributed Cloud Computing Practice for Deep Learning **34**
C.1 Model Parallel . . . 34
C.2 Data Parallel . . . 34
C.3 Others . . . 36

D Cortex Lab Existing Achievements **36**

Archived copy · cryp2nova.com

IMPORTANT NOTICE

This Whitepaper in current form is being circulated for general information and to invite investor feedback only on the Cortex platform as presently conceived, and is subject to review and revision by the directors of the Token Issuer and/or the Project Company, the advisers, and/or legal advisers of the Token Issuer and/or the Project Company. Please do not replicate or distribute any part of this Whitepaper without this note in accompaniment. No part of this Whitepaper is intended to create legal relations between a recipient of this Whitepaper or to be legally binding or enforceable by such recipient against the Token Issuer and/or the Project Company. An updated version of this Whitepaper may be published on a date to be determined and announced by the Token Issuer and/or the Project Company in due course.

PLEASE READ THIS SECTION AND THE FOLLOWING SECTIONS ENTITLED DISCLAIMER OF LIABILITY, NO REPRESENTATIONS AND WARRANTIES, REPRESENTATIONS AND WARRANTIES BY YOU, CAUTIONARY NOTE ON FORWARD-LOOKING STATEMENTS, THIRD PARTY INFORMATION AND NO CONSENT OF OTHER PERSONS, TERMS USED, NO ADVICE, NO FURTHER INFORMATION OR UPDATE, RESTRICTIONS ON DISTRIBUTION AND DISSEMINATION, NO OFFER OF INVESTMENT OR REGISTRATION AND RISKS AND UNCERTAINTIES CAREFULLY.

IF YOU ARE IN ANY DOUBT AS TO THE ACTION YOU SHOULD TAKE, YOU SHOULD CONSULT YOUR LEGAL, FINANCIAL, TAX OR OTHER PROFESSIONAL ADVISOR(S).

The Cortex Coins are not intended to constitute securities of any form, units in a business trust, units in a collective investment scheme or any other form of investment in any jurisdiction. This Whitepaper does not constitute a prospectus or offer document of any sort and is not intended to constitute an offer of securities of any form, units in a business trust, units in a collective investment scheme or any other form of investment, or a solicitation for any form of investment in any jurisdiction.

This Whitepaper does not constitute or form part of any opinion or any advice to acquire, sell, or any solicitation of any offer by the Project Company or the Token Issuer to acquire any Cortex Coins nor shall it or any part of it nor the fact of its presentation form the basis of, or be relied upon in connection with, any contract or investment decision.

The proceeds from the sale of the Cortex Coins will be deployed to support ongoing development and growth of the Cortex platform, business development, partnerships and community support activities.

No person is bound to enter into any contract or binding legal commitment in relation to the acquisition of Cortex Coins and no cryptocurrency or other form of payment is to be accepted on the basis of this Whitepaper.

Any agreement as between the Token Issuer and you as a participant in the Token Sale, and in relation to any purchase of Cortex Coins, is to be governed by only a separate document setting out the terms and conditions (the "**Token Sale Terms**") of such agreement. In the event of any inconsistencies between the Token Sale Terms and this Whitepaper, the former shall prevail.

PLEASE NOTE THAT THE TOKEN ISSUER WILL NOT OFFER OR SELL TO YOU, AND YOU ARE NOT ELIGIBLE TO PURCHASE ANY CORTEX COINS IN THE TOKEN SALE IF: (A) YOU ARE LOCATED IN THE PEOPLE'S REPUBLIC OF CHINA OR IF YOU ARE A CITIZEN OR RESIDENT (TAX OR OTHERWISE) OF, OR DOMICILED IN, THE PEOPLE'S REPUBLIC OF CHINA; (B) YOU ARE LOCATED IN THE UNITED STATES OF AMERICA OR IF YOU ARE A CITIZEN, RESIDENT (TAX OR OTHERWISE) OR GREEN CARD HOLDER OF, OR DOMICILED IN, THE UNITED STATES OF AMERICA, UNLESS YOU ARE A U.S. QUALIFIED PERSON; OR (C) SUCH TOKEN SALE IS PROHIBITED, RESTRICTED OR UNAUTHORIZED IN ANY FORM OR MANNER WHETHER IN FULL OR IN PART UNDER THE LAWS, REGULATORY REQUIREMENTS OR RULES IN THE JURISDICTION IN WHICH YOU ARE LOCATED, AT THE TIME OF YOUR INTENDED PURCHASE OR PURCHASE OF THE CORTEX COINS IN THE TOKEN SALE.

No regulatory authority has examined or approved of any of the information set out in this Whitepaper. No such action has been or will be taken under the laws, regulatory requirements or rules of any

Archived copy · cryp2nova.com

jurisdiction. The publication, distribution or dissemination of this Whitepaper does not imply that the applicable laws, regulatory requirements or rules have been complied with.

There are risks and uncertainties associated with the Project Company and the Token Issuer and their business and operations, the Cortex Coins, the Token Sale, and the Cortex platform. Please refer to the section entitled Risks and Disclosures set out at the end of this Whitepaper.

This Whitepaper, any part thereof and any copy thereof must not be taken or transmitted to any country where distribution or dissemination of this Whitepaper is prohibited or restricted.

No part of this Whitepaper is to be reproduced, distributed or disseminated without including this section and the following sections entitled "Disclaimer of Liability", "No Representations and Warranties", "Representations and Warranties By You", "Cautionary Note On Forward-Looking Statements", "Third Party Information and No Consent of Other Persons", "Terms Used", "No Advice", "No Further Information or Update", "Restrictions On Distribution and Dissemination", "No Offer of Investment Or Registration" and "Risks and Uncertainties".

DISCLAIMER OF LIABILITY

To the maximum extent permitted by the applicable laws, regulations and rules, the Project Company and/or the Token Issuer shall not be liable for any indirect, special, incidental, consequential or other losses of any kind, in tort, contract or otherwise (including but not limited to loss of revenue, income or profits, and loss of use or data), arising out of or in connection with any acceptance of or reliance on this Whitepaper or any part thereof by you.

NO REPRESENTATIONS AND WARRANTIES

The Project Company and the Token Issuer do not make or purport to make, and hereby disclaim, any representation, warranty or undertaking in any form whatsoever to any entity or person, including any representation, warranty or undertaking in relation to the truth, accuracy and completeness of any of the information set out in this Whitepaper.

REPRESENTATIONS AND WARRANTIES BY YOU

By accessing and/or accepting possession of any information in this Whitepaper or such part thereof (as the case may be), you represent and warrant to the Project Company and the Token Issuer as follows:

(a) you agree and acknowledge that the Cortex Coins do not constitute securities of any form, units in a business trust, units in a collective investment scheme or any other form of investment in any jurisdiction;

(b) you are not:

  (i) located in the People's Republic of China or a citizen or resident (tax or otherwise) of, or domiciled in, the People's Republic of China;

  (ii) located in the United States of America or a citizen, resident (tax or otherwise) or green card holder of, or domiciled in, the United States of America, unless you are a U.S. Qualified Person (as defined herein); OR

  (iii) located in a jurisdiction where the Token Sale is prohibited, restricted or unauthorized in any form or manner whether in full or in part under the laws, regulatory requirements or rules in such jurisdiction;

(c) if you are located in the United States of America or a citizen, resident (tax or otherwise) or green card holder of, or domiciled in, the United States of America, you are a U.S. Qualified Person, being a person (being either an individual or legal entity or a person, including without limitation a governmental authority) who is an Accredited Investor as defined in Rule 501(a) of Regulation D under the U.S. Securities Act of 1933, as may be modified, amended or supplemented from time to time;

Archived copy · cryp2nova.com

(d) you agree and acknowledge that this Whitepaper does not constitute a prospectus or offer document of any sort and is not intended to constitute an offer of securities of any form, units in a business trust, units in a collective investment scheme or any other form of investment in any jurisdiction, or a solicitation for any form of investment, and you are not bound to enter into any contract or binding legal commitment and no cryptocurrency or other form of payment is to be accepted on the basis of this Whitepaper;

(e) you acknowledge and understand that no Cortex Coin should be construed, interpreted, classified or treated as enabling, or according any opportunity to, tokenholders to participate in or receive profits, income, or other payments or returns arising from or in connection with the Cortex Coins or the proceeds of the Token Sale, or to receive sums paid out of such profits, income, or other payments or returns;

(f) you agree and acknowledge that no regulatory authority has examined or approved of the information set out in this Whitepaper, no action has been or will be taken under the laws, regulatory requirements or rules of any jurisdiction and the publication, distribution or dissemination of this Whitepaper to you does not imply that the applicable laws, regulatory requirements or rules have been complied with;

(g) you agree and acknowledge that this Whitepaper, the undertaking and/or the completion of the Token Sale, or future trading of Cortex Coins on any cryptocurrency exchange, shall not be construed, interpreted or deemed by you as an indication of the merits of the Project Company, the Token Issuer, the Cortex Coins, the Token Sale, and the Cortex platform;

(h) the distribution or dissemination of this Whitepaper, any part thereof or any copy thereof, or acceptance of the same by you, is not prohibited or restricted by the applicable laws, regulations or rules in your jurisdiction, and where any restrictions in relation to possession are applicable, you have observed and complied with all such restrictions at your own expense and without liability to the Project Company and/or the Token Issuer;

(i) you agree and acknowledge that in the case where you wish to acquire any Cortex Coins, Cortex Coins are not to be construed, interpreted, classified or treated as:

  (i) any kind of currency other than cryptocurrency;
  (ii) debentures, stocks or shares issued by any person or entity;
  (iii) rights, options or derivatives in respect of such debentures, stocks or shares;
  (iv) rights under a contract for differences or under any other contract the purpose or pretended purpose of which is to secure a profit or avoid a loss;
  (v) units in a collective investment scheme;
  (vi) units in a business trust;
  (vii) derivatives of units in a business trust; or
  (viii) any form of investment;

(j) you are legally permitted to participate in the Token Sale and all actions contemplated or associated with such participation, including the holding and use of Cortex Coins;

(k) the amounts that you use to acquire the Cortex Coins were not and are not directly or indirectly derived from any activities that contravene the laws and regulations of any jurisdiction, including anti-money laundering laws and regulations;

(l) if you are a natural person, you are of sufficient age and capacity under the applicable laws of the jurisdiction in which you reside and the jurisdiction of which you are a citizen to participate in the Token Sale;

(m) you are not obtaining or using Cortex Coins for any illegal purpose;

(n) you have a basic degree of understanding of the operation, functionality, usage, storage, transmission mechanisms and other material characteristics of cryptocurrencies, blockchain-based software systems, cryptocurrency wallets or other related token storage mechanisms, blockchain technology, and smart contract technology;

(o) you are fully aware and understand that in the case where you wish to purchase any Cortex Coins, there are risks associated with the Project Company and the Token Issuer and their business and operations, Cortex Coins, the Token Sale, and the Cortex platform;

Archived copy · cryp2nova.com

(p) you bear the sole responsibility to determine what tax implications a purchase of Cortex Coins may have for you and agree not to hold the Project Company, the Token Issuer and/or any other person involved in the Token Sale liable for any tax liability associated with or arising therefrom;

(q) you agree and acknowledge that the Project Company and the Token Issuer are not liable for any direct, indirect, special, incidental, consequential or other losses of any kind, in tort, contract or otherwise (including but not limited to loss of revenue, income or profits, and loss of use or data), arising out of or in connection with any acceptance of or reliance on this Whitepaper or any part thereof by you;

(r) you waive the right to participate in a class action lawsuit or a class wide arbitration against the Project Company, the Token Issuer and/or any person involved in the Token Sale and/or with the creation and distribution of Cortex Coins; and

(s) all of the above representations and warranties are true, complete, accurate and non-misleading from the time of your access to and/or acceptance of possession this Whitepaper or such part thereof (as the case may be).

CAUTIONARY NOTE ON FORWARD-LOOKING STATEMENTS

All statements contained in this Whitepaper, statements made in press releases or in any place accessible by the public and oral statements that may be made by the Project Company or its directors, executive officers or employees acting on behalf of the Project Company and/or the Token Issuer (as the case may be), that are not statements of historical fact, constitute "forward-looking statements". Some of these statements can be identified by forward-looking terms such as "aim", "target", "anticipate", "believe", "could", "estimate", "expect", "if", "intend", "may", "plan", "possible", "probable", "project", "should", "would", "will" or other similar terms. However, these terms are not the exclusive means of identifying forward-looking statements. All statements regarding the Project Company and/or the Token Issuer's business strategies, plans and prospects and the future prospects of the industry which the Project Company and/or the Token Issuer is in are forward-looking statements. These forward-looking statements, including but not limited to statements as to the Project Company and/or the Token Issuer's prospects, future plans, other expected industry trends and other matters discussed in this Whitepaper regarding the Project Company and/or the Token Issuer are matters that are not historic facts, but only predictions.

These forward-looking statements involve known and unknown risks, uncertainties and other factors that may cause the actual future results, performance or achievements of the Project Company and/or the Token Issuer to be materially different from any future results, performance or achievements expected, expressed or implied by such forward-looking statements. These factors include, amongst others:

(a) changes in political, social, economic and stock or cryptocurrency market conditions, and the regulatory environment in the countries in which the Project Company and/or the Token Issuer conduct their business and operations;

(b) the risk that the Project Company may be unable to execute or implement its business strategies and future plans;

(c) changes in interest rates and exchange rates of fiat currencies and cryptocurrencies;

(d) changes in the anticipated growth strategies and expected internal growth of the Project Company and the Cortex platform;

(e) changes in the availability and fees payable to the Project Company in connection with its businesses and operations or in the Cortex platform;

(f) changes in the availability and salaries of employees who are required by the Project Company and/or the Token Issuer to operate their business and operations;

(g) changes in preferences of users of the Cortex platform;

(h) changes in competitive conditions under which the Project Company operates, and the ability of the Project Company to compete under such conditions;

(i) changes in the future capital needs of the Project Company and the availability of financing and capital to fund such needs;

Archived copy · cryp2nova.com

(j) war or acts of international or domestic terrorism;

(k) occurrences of catastrophic events, natural disasters and acts of God that affect the businesses and/or operations of the Project Company and/or the Token Issuer;

(l) other factors beyond the control of the Project Company and/or the Token Issuer; and

(m) any risk and uncertainties associated with the Project Company and the Token Issuer and their business and operations, the Cortex Coins, the Token Sale, and the Cortex platform.

All forward-looking statements made by or attributable to the Project Company, the Token Issuer and/or persons acting on behalf of the Project Company and/or the Token Issuer are expressly qualified in their entirety by such factors. Given that risks and uncertainties that may cause the actual future results, performance or achievements of the Project Company and/or the Token Issuer to be materially different from that expected, expressed or implied by the forward-looking statements in this Whitepaper, undue reliance must not be placed on these statements. These forward-looking statements are applicable only as of the date of this Whitepaper.

None of the Project Company, the Token Issuer or any other person represents, warrants, and/or undertakes that the actual future results, performance or achievements of the Project Company and/or the Token Issuer will be as discussed in those forward-looking statements. The actual results, performance or achievements of the Project Company and/or the Token Issuer may differ materially from those anticipated in these forward-looking statements.

Nothing contained in this Whitepaper is or may be relied upon as a promise, representation or undertaking as to the future performance or policies of the Project Company and/or the Token Issuer.

Further, the Project Company and the Token Issuer disclaim any responsibility to update any of those forward-looking statements or publicly announce any revisions to those forward-looking statements to reflect future developments, events or circumstances, even if new information becomes available or other events occur in the future.

THIRD PARTY INFORMATION AND NO CONSENT OF OTHER PERSONS

This Whitepaper includes information obtained from various third party sources ("Third Party Information"). None of the publishers of Third Party Information has consented to the inclusion of Third Party Information in this Whitepaper and is therefore not liable for Third Party Information. While reasonable action has been taken to ensure that Third Party Information has been included in their proper form and context, neither the Project Company and the Token Issuer nor their respective directors, executive officers, and employees acting on its behalf, has independently verified the accuracy, reliability, completeness of the contents, or ascertained any applicable underlying assumption, of the relevant Third Party Information. Consequently, neither the Project Company and the Token Issuer nor their respective directors, executive officers and employees acting on their behalf makes any representation or warranty as to the accuracy, reliability or completeness of such information and shall not be obliged to provide any updates on the same.

TERMS USED

To facilitate a better understanding of the Cortex Coins being the subject of the sale conducted by the Token Issuer, and the business and operations of the Project Company and the Token Issuer, certain technical terms and abbreviations, as well as, in certain instances, their descriptions, have been used in this Whitepaper. These descriptions and assigned meanings should not be treated as being definitive of their meanings and may not correspond to standard industry meanings or usage.

Words importing the singular shall, where applicable, include the plural and vice versa and words importing the masculine gender shall, where applicable, include the feminine and neuter genders and vice versa. References to persons shall include corporations.

Archived copy · cryp2nova.com

NO ADVICE

No information in this Whitepaper should be considered to be business, legal, financial or tax advice regarding the Project Company, the Token Issuer, the Cortex Coins, the Token Sale, or the Cortex platform. You should consult your own legal, financial, tax or other professional adviser regarding the Project Company and the Token Issuer and their business and operations, the Cortex Coins, the Token Sale, and the Cortex platform. You should be aware that you may be required to bear the financial risk of any purchase of Cortex Coins for an indefinite period of time.

NO FURTHER INFORMATION OR UPDATE

No person has been or is authorised to give any information or representation not contained in this Whitepaper in connection with the Project Company and the Token Issuer and their business and operations, the Cortex Coins, the Token Sale, or the Cortex blockchain, if given, such information or representation must not be relied upon as having been authorised by or on behalf of the Project Company or the Token Issuer. The Token Sale shall not, under any circumstances, constitute a continuing representation or create any suggestion or implication that there has been no change, or development reasonably likely to involve a material change in the affairs, conditions and prospects of the Project Company and/or the Token Issuer or in any statement of fact or information contained in this Whitepaper since the date hereof.

RESTRICTIONS ON DISTRIBUTION AND DISSEMINATION

The distribution or dissemination of this Whitepaper or any part thereof may be prohibited or restricted by the laws, regulatory requirements, and rules of any jurisdiction. In the case where any restriction applies, you are to inform yourself about, and to observe, any restrictions which are applicable to your possession of this Whitepaper or such part thereof (as the case may be) at your own expense and without liability to the Project Company and the Token Issuer.

Persons to whom a copy of this Whitepaper has been distributed or disseminated, provided access to or who otherwise have the Whitepaper in their possession shall not circulate it to any other persons, reproduce or otherwise distribute this Whitepaper or any information contained herein for any purpose whatsoever nor permit or cause the same to occur.

NO OFFER OF INVESTMENT OR REGISTRATION

This Whitepaper does not constitute a prospectus or offer document of any sort and is not intended to constitute an offer of securities of any form, units in a business trust, units in a collective investment scheme or any other form of investment, or a solicitation for any form of investment in any jurisdiction. No person is bound to enter into any contract or binding legal commitment and no cryptocurrency or other form of payment is to be accepted on the basis of this Whitepaper.

No regulatory authority has examined or approved of any of the information set out in this Whitepaper. No such action has been or will be taken under the laws, regulatory requirements or rules of any jurisdiction. The publication, distribution or dissemination of this Whitepaper does not imply that the applicable laws, regulatory requirements or rules have been complied with.

RISKS AND UNCERTAINTIES

Prospective purchasers of Cortex Coins should carefully consider and evaluate all risks and uncertainties associated with the Project Company and the Token Issuer and their business and operations, the Cortex Coins, the Token Sale, and the Cortex platform, all information set out in this Whitepaper and the Token Sale Terms prior to any purchase of the Cortex Coins. If any of such risks and uncertainties develops into actual events, the business, financial condition, results of operations and prospects of the Project Company and/or the Token Issuer could be materially and adversely affected. In such cases, you may lose all or part of the value of the Cortex Coins. Please refer to the risk factors set out in pages 39 to 44 of this Whitepaper.

Archived copy · cryp2nova.com

1 Introduction

Know Thyself. *– Ancient Greek Aphorism on Temple of Apollo at Delphi*

Two possibilities exist: either we are alone in the Universe or we are not. Both are equally terrifying. *– Arthur Charles Clarke*

After hundreds of millions of years of evolution, the wisdom of mankind shines at its peak, illuminating the future in the fog, and the loneliness it accompanies is even stronger. Fear of the unknown, its own confusion and the pursuit of reality, make mankind as a whole, feel the unprecedented loneliness. Therefore, human study life, devote all efforts to create new automatic machines, hoping to surpass the speed of evolution, accelerating into the future.

On January 3, 2009, Bitcoin [1] launched a Genesis Block as a self-sustaining P2P system that ingeniously drove participants to maintain their operations and provided limited but highly disruptive financial functionalities. On June 30, 2015, Ethereum [2] went live, adding Turing-Complete smart contracts to the blockchain, allowing for consensus on the execution of short programs. Compared to Bitcoin, Ethereum can perform more complex computations and provide a richer response, yet these contracts are not self-evolving codes. Instead, they are a collection of purely rule-based and recursive programs. With reference to Conway's life game [3], the virtual currency network based on P2P technology can be defined as life on the Internet, maintaining its existence by providing financial functions. As long as there is one single full node alive, the state of the network can be sustained, and can respond to the interaction from the outside world. However, since the human longing for an intelligent network has not manifested yet, these primitive networks have remained relatively simple.

Based on this, Cortex has added a consensus on AI to the network. This facilitates all nodes working together to reach a consensus on the implementation of a smart contract that requires AI to empower the system with intelligent responses. As a stand-alone public chain compatible with EVM smart contracts, Cortex runs both existing contracts and inferred contracts with AI and will survive as a smarter web presence after the Genesis Block is released. In Cortex, due to open source and natural competition mechanisms, the best model will survive to enhance the intelligence level of the network. From a machine learning researcher's point of view, the Cortex platform brings together open models of a variety of basic smart applications with the state of the art quality, which will greatly accelerate their research and drive AI adoptions much faster. Once deployed, the chain also allows the computation of the model to get the whole network notarized automatically. It is still unknown whether aliens exists or not, but human beings are no longer alone when they are accompanied by the AI.

1.1 Background

Existing blockchain contracts can only perform simple smart contract computations, which cannot satisfy the application of real-world AI.

The blockchain addresses the transmission and accounting of decentralized value networks. Bitcoin uses a SHA256 Proof of Work (PoW) for consensus on the computational contribution, with each block divided into three parts, namely:

1. The hash value of the last block serves as the block header of the current block;
2. Pending transactions (t_1, t_2,$\cdots$,t_n) within the time window T will be hashed into block coinbase; and
3. Including miner's address, which is normally the address of mining pool, the X, as the input of hash functions, will be dispatched by pool server to each miner, who will complete certain computations. The goal is to find $H(X, \text{nonce}) <$ Target Difficulty, where nonce is an appended randomized guessing number. The computation result will be verified by the whole network nodes so as to get the reward out of the block to the exact miner's address, and then the whole network enters the computation of the next block, thus forming a chain eventually. In addition, there are some other information, such as version number, Merkle tree, timestamp, etc.

Archived copy · cryp2nova.com

The whole mining process can be summarized as:

$$\mathrm{SHA256}(\mathrm{SHA256}(\mathrm{version} + \mathrm{prev_hash} + \mathrm{merkle_root} + \mathrm{ntime} + \mathrm{nbits} + \mathrm{nonce})) < \mathrm{TARGET}$$

Ethereum makes use of uncle blocks [4] to improve network concurrency. In particular, both the Ethereum network and the Rootstock [5] network are designed for smart contracts on the chain. The consensus and tamper-proof of make blockchain automatically ensures the enforcement of the contract, contract execution, and funds allocation, thus eliminating the trust and dependence on people or other third parties.

The substantial increase in computing power has led to the rise of AI in recent years. Machine learning problems in the field of AI can be generalized to the following form: For a problem Q and the input data set D for this problem, a metric P is given, from which the model M is obtained, so that the evaluation of the model M in this problem is improved. In this form, all machine learning problems can be attributed to the following elements: input, where D stands for; output, where Y stands for; metric, where P stands for. The learning algorithms need to solve the problem:

Maximize or Minimize: $\mathrm{P}(Y, \mathrm{M}(D))$. Appendix A and B give more detailed descriptions and math statements on the specific issues.

In the process of optimizing the objective function, various numerical methods are often used to iteratively gradient descent to find the global optimal solution. Large-scale distributed learning often adopts ASGD (asynchronous stochastic gradient descent) to optimize the results. Sometimes, for particular problems, the training can only obtain a sub-optimal solution at a certain distance from the global optimum according to a distribution.

Due to the innate central tendency in machine learning training, a centralized large cluster has an unparalleled advantage over a loosely coupled decentralized cluster. The Project Company is committed to building a cluster off the chain for training and thus optimizing the on-chain model.

1.2 AI on Blockchain - Feasibility Analysis

There are currently a number of AIs on blockchain projects, which are conceptualized in a basic fashion and present concrete practical solutions. The directions are summarized as follows:

1. Attempting to combine reinforcement learning with distributed mining.
2. Establishing a platform for data exchange and machine learning tasks publishing.
3. Providing all chains machine learning API call.
4. In the process of training, homomorphic encryption is used to keep the users' data privacy confidential.

The current research interest of blockchain lies in the guarantee of privacy - namely fungibility. The homomorphic encryption scheme can protect both user data and the model from being stolen by others in the cloud computing. The so-called homomorphic encryption refers to an encryption method that makes the result of an operation on plaintext equivalent to the result of an operation on a ciphertext (equivalent to the encryption operator $\circ$ unchanged):

$$E(x) \circ E(y) = E(x \circ y)$$

For FHE (Full Homomorphic Encryption), it means invariance is met for any operator $\circ$. Although there are theoretical papers that prove the feasibility of completely homomorphic encryptions, due to the amounts of computation being too large, there is no practical industrial application. In addition, there is an encryption method called SWHE (Somewhat Homomorphic Encryption) [6] that supports only certain functions such as polynomials, algebraic multiplications, additions, and so on.

For machine learning problems, there are currently three situations:

1. Data encryption, model parameters are not encrypted. In this case, the user is more concerned about data confidentiality, such as a patient's CT data.
2. Data is not encrypted, model parameters are encrypted. In this case, the user simply wants to train or test his own model and does not care about the data.

Archived copy · cryp2nova.com

3. Both Data and model parameters are encrypted.

For the third case, if the encryption algorithm or the private key is different, the data cannot be trained. For example, if user A provides a model and user B provides the data, the training process cannot be carried out.

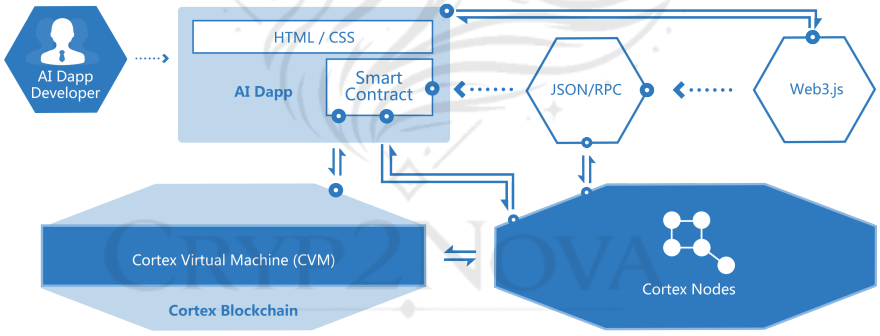
For data or model encryption, the exact value of the loss function will be unknown because the last residual function is computed as being encrypted. At this moment, the model cannot be assessed, cross-validation cannot be done to avoid over-fitting, and hyper-parameter adjustments (such as critical learning rates) cannot be made. All hyper-parameter adjustments must be confirmed by the publisher after decryption.

For the first two cases, FHE is not considered at the moment due to the high computational complexity. Only the SWHE method with less computational complexity is usable. After testing, the computational complexity of the state of the art is still very high. Compared with the traditional method, if we use the homomorphic encryption to train and infer the model, the computational complexity increases by 2-3 magnitudes, which is unacceptable.

For the decentralized parallel training, due to the current difficulties in terms of network bandwidth, synchronization parameter difficulty and training progress consensus, Cortex has introduced the world's top notch experts in distributed machine learning to solve the issue.

Cortex's main mission is to provide the best in class machine learning models on the blockchain, and users can infer using smart contracts on the Cortex blockchain. The Cortex goal also includes implementing a machine learning platform that allows users to post tasks on the platform, submit models, make inferences by calling intelligent contracts, and create their own AI DApps (Artificial Intelligence Decentralized Applications).

Figure 1: AI DApp



1.3 AI DApps

Cortex brings AI to smart contracts, making the following applications possible:

Information Services: Personalized Recommendation System, Search Engine, News Writing Services.

Finance: Credit, Intelligent Investment Advisory.

AI Assistant: Automatic Q&A, Industry Knowledge Map, Speech Synthesis, Face Attribute Prediction.

Simulation Environment: Auto-driving, Go and other Reinforcement Learning Applications.

Archived copy · cryp2nova.com

2 System Architecture

2.1 Expanding the capabilities of smart contracts and blockchain

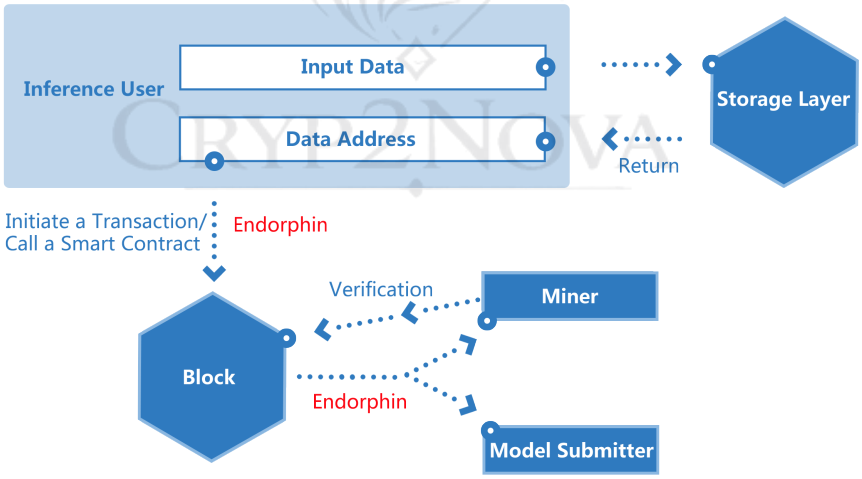
2.1.1 Cortex Intelligent Inference Framework

Model providers will not be limited to those under the Project Company. Machine learning researchers around the world can upload well-trained corresponding data models to the storage layer. Other users who need this AI models can make inferences by using models and paying their providers. At each inference, a full node synchronizes the model and the data from the storage tier to the local site. Making an inference using Cortex's unique virtual machine, CVM (Cortex Virtual Machine), synchronizes the results to the whole network, and returns the result.

Given a model, the output, being the results obtained from data plugged into the model, is an intelligent inference process. Every time a user initiates a transaction, performs a smart contract, or performs an intelligent inference, the user needs to pay a certain amount of Endorphin. Endorphin is the pricing unit that is required to be paid for dealing transactions or executing AI smart contracts on the Cortex chain. The amount of Endorphin required for each payment depends on the difficulty of the model operation and the model ranking. The price of Endorphin is market-driven, which is a dynamic conversion rate between Endorphin and Cortex Coin, reflecting Cortex's cost of executing smart contracts. The more the number of people who initiate transactions or execute AI smart contracts at the same time, the higher the price of Endorphin. It depends on the gaming between miners and people who initiate transactions or execute AI smart contracts. The equivalent Cortex Coin for the amount of Endorphin required is split into two parts, one for the miners' cost of packaging the transactions in a block and the other for the model provider who attracts smart contracts to request inference services. Cortex will set a cap on this ratio for the percentage paid to the model provider.

Cortex will add an Infer instruction to the smart contract, making it possible to support the use of Cortex blockchain models in smart contracts.

Figure 2: Inference Process



The following pseudocode shows how to use Infer in a smart contract. The data is inferred on the model once a user calls these smart contracts:

Archived copy · cryp2nova.com

Inference Pseudocode

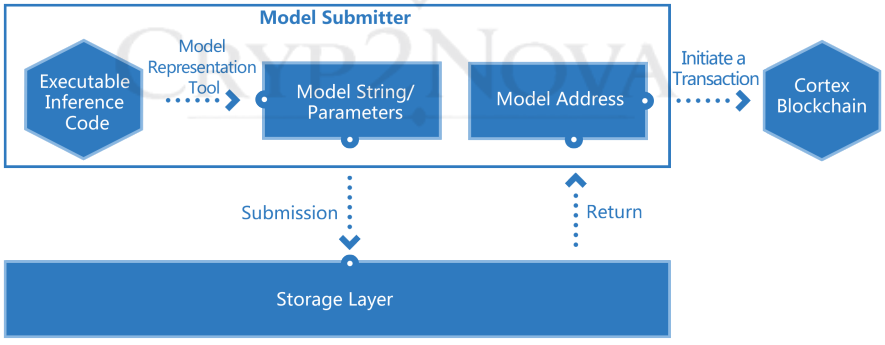
```
contract MyAIContract{
  InferType res;
  ...
  function myAIFunction(){
    ...
    res=infer(model_hash, data_hash);
    ...
  }
  ...
}
contract InferType{
  ...
}
```

2.1.2 Model Submission Framework

In the previous section, we analyzed the pitfalls and feasibility of training on the blockchain. Cortex also provides a submission interface for training off-chain, including the instruction-interpreting virtual machine for models. This will set up a bridge between computing power providers and algorithm providers for trading and collaboration.

The user parses the model into a model string via Cortex's CVM, parses the parameters up to the storage layer, and publishes a generic interface to let the smart contract programmers call. The model provider needs to pay a certain storage fee to ensure that the model can be persistently saved in the storage layer. Part of the fee charged for inference by calling this model in a smart contract is delivered to the model provider. The provider can also withdraw and update accordingly if needed. In the case of withdrawals, in order to ensure that the smart contract calling this model works properly, the Project Company will host it according to the usage of the model and keep invoking the model for a fee equal to the storage and maintenance costs. Cortex also provides an interface to upload the model to the storage layer and obtain the model hash. The provider then initiates a transaction, executing a smart contract to write the model hash into the storage, so that all users are apprised of the model's input and output status.

Figure 3: Model Submission Process



2.1.3 Smart AI contract

Cortex allows users to write machine learning programs on the blockchain, and submit some interactions that depends on other contracts. For example, the interactions between pets on the electronic pet Cryptokitties running on Ethereum can be dynamic, intelligent, and evolutionary. Through the uploaded reinforcement learning model, given the smart AI contracts, users can easily achieve a variety of similar AI applications.

Archived copy · cryp2nova.com

At the same time, Cortex provides AI interface calls for other blockchains. For example, on the blockchains of Bitcoin Cash and Ethereum, Cortex provides the wallet address analysis results based on smart AI contracts; the models being used for addresses analysis will not only be helpful for regulation technologies but will also provide general public with a risk-based assessment of the transaction recipient.

2.2 Model and Data Storage

Rather than actually storing models and data, the Cortex blockchain stores the hash values of the model and the data. The key-value storage system is off-chain. New models and new data are available to the Cortex blockchain after there are enough copies spreading over the whole network.

2.3 Cortex Consensus Inference Criteria

When a user initiates a transaction to a contract, the full node needs to execute the code of the smart contract. The difference between Cortex and ordinary smart contracts is that intelligent contracts may involve inference instructions, and subsequently require all nodes to agree on the result of this inferred result. The full node implementation process is:

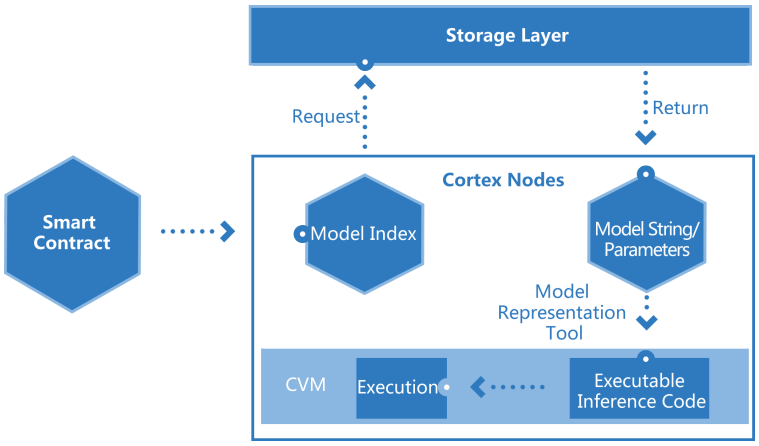
1. The full node locates the model at the storage layer by querying the model index and downloads the model string and the corresponding data parameter of the model.
2. The model string is translated into executable code using the Cortex model representation tool.
3. Through the virtual machine CVM provided by Cortex, the implementation of executable code, the results of all node broadcast consensus.

The Cortex model representation tool can be divided into two parts:

1. Using the model representation tool, model providers convert the model code, which can be written in machine learning framework familiar to them such as MXNet or TensorFlow, into a model string such that the string could be submitted to the storage layer.
2. After all the full nodes download the model string, the string needs to be converted through Cortex model representation tool into executable code to perform the inference on the CVM.

The role of the CVM is that every inference execution on all full nodes is deterministic, chapter 3.2 and chapter 3.3 describes the implementation details of the Cortex Model Presentation Tool and CVM.

Figure 4: Inference Process on Full Nodes



Archived copy · cryp2nova.com

2.4 Model Selection and Valuation

Cortex blockchain will not restrict the model, the user can rely on the number of inference calls as a relatively objective model evaluation criteria. When Cortex users have high-precision demands on the model inference regardless of the computational cost, Cortex supports retaining the original model parameters using floating-point numbers. Thus, an official or third-party agency can both rank the model by defining its own mechanism (recall, accuracy, computation speed, benchmarking dataset, etc.) and display it on third-party websites or applications.

2.5 Consensus MechanismPoW Mining

For a long time, one-CPU-one-vote is the ideal goal in digital cryptocurrency community. However, this vision has not been realized. The reason is that ASIC's special design significantly improves the computation speed. Communities and academics have explored a number of memory hard algorithms to make graphics and CPU mining more user-friendly, which aims to make spending large amount of capital on buying specialized mining equipment less advantageous. Recent community practice shows that Dagger-Hashimoto [7] by Ethereum and Equihash [8] by Zcash are algorithmic practices for more successful GPU priorities.

Cortex will further give priority to one-machine-one-vote, using Cuckoo Cycle PoW to further narrow the gap between CPU and GPU speedup ratio. At the same time Cortex blockchain will fully explore the performance of smart phone GPU, making the difference between mobile phones and desktop GPUs more in line with the ratio of the general hardware platform evaluation tools (such as GFXBench): for example, the best consumer level GPU is 10-15 times more powerful than the best mobile phone GPU. Taking into account the lower power consumption of mobile computing, it is more feasible for a larger number of mobile phones charging at nighttime to mine Cortex blockchain.

It is particularly noteworthy that there is no direct link between consensus algorithms for block encryption and the computation of intelligent inference on the chain. PoW guarantees more fairness on the hardware of the miners involved in mining while the intelligent inference contract automatically provides PVI (Public Verifiable Inference).

2.6 Fraud Prevention and Model Screening

Since the model is completely open source, it can be copied or plagiarized. Under normal circumstances, if it is a very good model, there will often be a lot of usage after going live, and copying these models would not have great advantages. However, in some special cases, where there is obvious plagiarism or full reproduction, the Project Company will intervene and arbitrate, through the blockchain oracle for publicity.

3 Software Solutions

3.1 CVM: EVM + Inference

As mentioned earlier, Cortex has its own virtual machine called the Cortex Virtual Machine (CVM). The CVM instruction set is fully EVM-compatible, and in addition, CVM provides support for inference instructions. Cortex will add a new INFER instruction at 0xc0. The input to instruction is the inferred code, the output is the inferred result. The contents of the virtual machine instructions used by CVM are described in Table 1.

3.2 Cortex Core Instruction Set and Framework Standards

Typical AI applications - image problems, speech / semantic / textual problems, and reinforcement learning problems require the following tensors without exceptions. The cost of Cortex's tensor operations is a potential anchor for Endorphin billing, which parses the core instruction set for machine learning and deep learning. In different computing frameworks, this term is often referred as the network layer or operator.

- Tensor Computational Operations:
  - Tensor numerical four operations: input - tensors, numerical and four operators

Archived copy · cryp2nova.com

Table 1: CVM Instruction Set

Prefix	Instructions
0s	Stop and Arithmetic Operations
10s	Comparison & Bitwise Logic Operations
20s	SHA3
30s	Environmental Information
40s	Block Information
50s	Stack, Memory, Storage and Flow Operations
60s & 70s	Push Operations
80s	Duplication Operations
90s	Exchange Operations
a0s	Logging Operations
c0s	Cortex Operations
0xc0	INFER
f0s	System operations

  - Four bitwise operation between tensors: input - two tensors and four operators
  - Tensor bitwise function operations: input - tensor and power functions, trigonometric functions, power and logarithmic functions, inequality functions, random number generating functions, rounding functions and so on.
  - Dimensional reduction of tensor: input - tensors and meet the binding law and exchange law.
  - Broadcast operation between tensors: input - tensors, padding for lower dimensional tensors for bitwise operations
  - Multiplication between tensors: including operations such as tensor and matrix, matrix multiplication / matrix multiplication with vectors, and matrix multiplication. For example, the NCHW / NHWC tensor storage mode.

- Tensor Reconstruction Operations:
  - Dimensionality exchange, expansion and compression
  - Sort by dimensionality
  - Value padding
  - Join by channel
  - Splicing or cropping by image

- Neural Network Specific Operations:
  - Full connection
  - Neural network activation function mainly depends on the tensor bitwise function operations.
  - 1D / 2D / 3D convolution (including convolutional kernels with different scales, holes, grouping, etc.)
  - 1D / 2D / 3D deconvolution operation and linear interpolation operation through upsampling/unpooling
  - Common auxiliary operations (such as the first / second order statistics BatchNorm)
  - Image-based aided computing (such as the deformation of convolutional network parameter module)
  - Specific tasks aided computing (such as ROIPooling, ROIAlign module)

The mainstream AI computing frameworks have been covered by the Cortex core instruction set. Subject to the implementation of BLAS on different platforms, Cortex supports the conversion of a Cortex model with floating-point (Float32, Float16) parameters to a fixed-point number (INT8, INT6) parameter model Wu et al. [9]Han et al. [10], such that the consensus across frameworks is achieved.

Archived copy · cryp2nova.com

3.3 Cortex Model Representation Tool

The Cortex MRT (Model Representation Tool) creates an open, flexible standard that enables deep learning frameworks and tools to be interoperable. It enables users to migrate deep learning models from one framework to another, making them easier to put into production. As the blockchain is in an open ecosystem, it makes AI more accessible and valuable: developers will choose the right framework for their tasks, framework authors will focus on innovation and enhancements, and hardware vendors will simplify optimization. For example, developers can use frameworks like PyTorch to train complex computer vision models and infer using CNTK, Apache MXNet, or TensorFlow.

The Cortex MRT was designed for

- Representation: Mapping strings to mainstream neural network models, the finest granular instructions supported by the probabilistic model
- Organization: Mapping the instruction set to the main neural network framework code
- Transfer: Providing isomorphic detection tools that allow the same models to migrate to each other in different machine learning / neural network frameworks

3.4 Storage Layer

Cortex can use any key-value storage system, such as IPFS and libtorrent, to save the model. The abstraction layer of Cortex's data storage does not depend on any specific distributed storage solution. Distributed hash tables or IPFS can be used to solve storage problems. For different devices, Cortex adopts different strategies:

- The full node always stores the blockchain data model
- The mobile phone node takes a Bitcoin SPV mode, with only a small full-size model

Cortex is only responsible for consensus inference, and does not store any training sets. To help contract authors filter the model and avoid over-fitting data model cheats, contract authors can submit test sets to the Project Company, which publishes the model results.

A call from the contract level will be queued in the memory pool, waiting for the block, and will be packaged into the block as a confirmed transaction. The data is broadcast to the full node during caching, including the mining pool. For models and data that exceed Cortex's current storage limits, such as medical hologram data, which could be dozens of GBs, communities would have to wait for Cortex updating protocols for storage limits and additional support in future. Cortex is able to cover the vast majority of AI applications such as pictures, voices, texts, short videos, etc.

For Cortex's full nodes, it requires more storage space than existing Bitcoin and Ethereum to keep cached data test sets and data models. Taking into account Moore's Law, storage prices will continue to decline, and thus will not constitute too much of an obstacle.

For each data model, annotation information is created within the Metadata for retrieval of on-chain calls. The format of Metadata is expressed in Table 2

3.5 Model Indexing

Cortex stores all the models. At each node, for each transaction that needs to be verified, it is necessary to quickly retrieve the corresponding model for inference in the smart contract if it involves consensus inference. In the memory of full node, Cortex will index locally stored models and retrieve them based on addresses within smart contracts.

3.6 Model Cache

Cortex's full node storage capacity is limited, so it cannot save all the models of the entire network. Cache is introduced to solve this problem. A full node maintains a model cache locally. Replacement strategies vary by implementation, with the most commonly used, LRU (Least Recently Used), FIFO (First In First Out) and others typically available, as well as any other solution to increase the hit ratio.

Archived copy · cryp2nova.com

Table 2: Metadata

	Keyword	Example
Model:		
	MD5	8ac7b335978cf2fe8102c5c43e380ca1
	Field	Speech, Image, etc.
	Method	NasNet Large
	Is Training	False if model is deployed else True
	Loss Function	Softmax, Hinge loss, KL divergence
	Rank	1
	Model Size	$10^4 - 10^{10}$ Byte
Input:		
	Dataset	ImageNet22k + Place365
	Dimension	Dim of Input, e.g. 2 or 3 for Image
	Size	e.g. $1024 \times 1920 \times 3$
	Dtype	Float32
Output:		
	Range	[0,1],[0,255]
	Predict	Top 5 Predictions
	Dimension	Dim of output
	Size	e.g. 1

3.7 Full Node Experiment

This chapter describes the results of some experiments on a single machine for the throughput of inference instructions executed by all nodes. Test platform configuration is:

- CPU: E5-2683 v3
- GPU: 8x1080Ti
- RAM: 64 GB
- Disk: SSD 960 EVO 250 GB

The test code used in the experiment is based on python 2.7 and MXNet, and mainly contains the following models:

- CaffeNet
- Network in Network
- SqueezeNet
- VGG16
- VGG19
- Inception v3 / BatchNorm
- ResNet-152
- ResNet101-64x4d

All models can be found in the MXNet documentation [1].

Experiments were performed on the CPU and GPU to test the inference speed of these models on the platform. These tests do not consider the speed of reading the model, with all models loaded into host or GPU memory in advance.

The test results are shown in Table 3, the Batch Size in parentheses (the amount of data samples imported in one batch), and all GPU test results are based on the single card.

[1] `http://mxnet.incubator.apache.org/model_zoo/index.html`

Archived copy · cryp2nova.com

Table 3: Inference Performance

Model	Size	CPU Infer (1)	GPU Infer (1)	GPU Infer (64)
CaffeNet	232 MB	196 ms	2.23 ms	39.98 ms
Network in Network	28.97 MB	115 ms	2.12 ms	42.90 ms
SqueezeNet v1.1	4.72 MB	130 ms	2.16 ms	46.18 ms
VGG16	527.79 MB	657 ms	5.84 ms	177.95 ms
VGG19	548.05 MB	681 ms	6.70 ms	205.26 ms
Inception v3/BN	43.19 MB	1084 ms	8.53 ms	80.61 ms
ResNet-152	230.18 MB	4050 ms	23.93 ms	253.08 ms
ResNet101-64x4d	283.86 MB	2650 ms	14.19 ms	182.73 ms

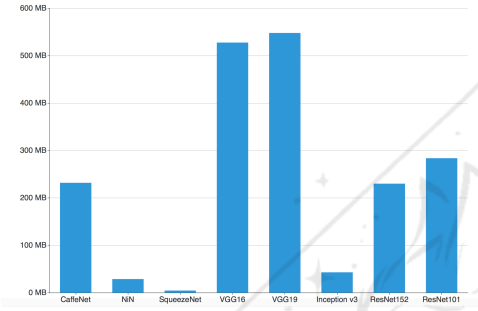


Figure 5: Model Size

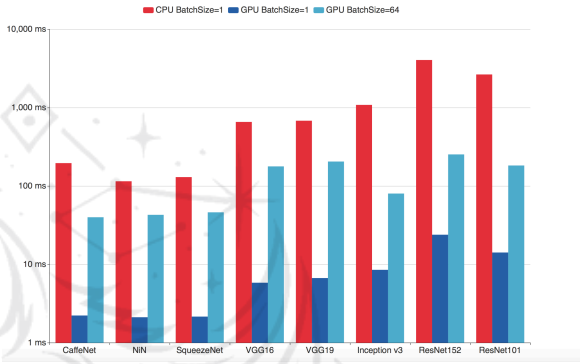


Figure 6: Different models inference time contrast (Logarithmic coordinates)

The above is the result of a single machine test. In order to simulate the real situation, the experimental platform is continuously running inference on a dataset stream containing about 100K images. Each inference is performed on a randomly selected model and the Batch Size is 1, the picture is distributed to 8 GPU cards with load balancing. For two situations:

1. All the models have been read and stored in GPU memory, and the average speed of inference on a single picture is 3.16 ms.
2. Each time the cached data was re-read (including the model and the input data) instead being loaded into GPU memory ahead of time, the average speed of inference per single picture was 113.3 ms.

Conclusions All nodes support load balancing after the model has been pre-fetched to GPU memory, and parallel inference between GPU cards for the same model results in approximately 300 inferences of a single inference per second. If, in extreme cases, no GPU memory is read ahead and only cached, which means each time the model is re-read and the data re-loaded, a single inference can be made 9 times per second. The above experiments are all unoptimized computations. One of Cortex's goals is to continuously optimize and improve inference performance.

4 Hardware Solutions

4.1 CUDA and RoCM

Cortex's hardware solutions deeply involve the use of NVidia's CUDA driver and CUDNN library as the GPU development framework. At the same time, the AMD OpenMI software project uses the RoCM driver and HIP/HCC AI library R&D program, Cortex plans to support the computations by the end of 2018 once they are available.

Archived copy · cryp2nova.com

Figure 7: Bitcoin USB Miner and Neural Network Computing Stick



4.2 FPGA

FPGA products are characterized by low fixed-point arithmetic (INT8 or even INT6 arithmetic) and lower latency, but higher computational power consumption and poorer flexibility. Deep learning tasks already have a good deployment in the area of autonomous driving and cloud services by using FPGA. Cortex plans to provide inference support for Xilinx and Altera products.

4.3 Full Node Hardware Configuration Requirements - Multi-GPU and Legendary USB Mining

Make mining great again!
Unlike traditional Bitcoin and Ethereum nodes, Cortex has a higher hardware requirement for full node. This requires a relatively larger amount of hard disk storage and a multi-GPU desktop host for the best possible speed of confirmation, but this is not a must. In the field of Bitcoin mining, the USB miner used to be a popular plug and play small ASIC mining device. Before the large-scale mining farms emerged, this decentralized mining mode was extremely popular. Cortex full node in the absence of GPU can be configured to have the similar neurocomputers with special AI chips, and computing stick, which have matured in the market. Unlike Bitcoin USB Mining, the computing stick is the complementary hardware to verify the full node, not the equipment needed in the specific process of mining.

4.4 Hardware Modifications for Existing GPU Mining Farm

For an existing GPU mining farm, and in particular for high-end GPU configurations, the Cortex Project Company provides retrofitting consulting services and total technical solutions that enable the farm to have the same level of AI computing center comparing the world's leading AI companies. The hardware cost-effectiveness will far exceed the existing commercial GPU cloud; it is extremely exciting to leverage those millions of GPU to join in AI computing competition globally! The multi-center mining farms have the opportunity to sell computing power to algorithmic providers and researchers, or to generate data models the same quality as a world-class AI company in a collaborative manner.

Some concrete retrofitting strategies are:

- Motherboard and CPU custom strategies to meet the multiple lanes of PCI-E to expend data transmission bandwidth for deep learning
- 10 Gigabyte switches and networking card solutions
- Storage hardware and network bandwidth solutions
- Customized software automatically switch computing jobs among training, mining the Cortex main chain, and the other competitive GPU currencies.
- Customized mobile applications monitoring mining revenue, manual switching and other remote management over GPU device.

Archived copy · cryp2nova.com

Figure 8: Big Basin from Facebook AI Research utilizes Wedge 100-32X Network Switch for the Distributed Training Farm in the Permissioned Chain



4.5 Mining and Computing on Mobile Devices and IoT

Balancing the ratio of revenue among heterogeneous computing chips, such as CPU, GPU, FPGA and ASIC, to make mining more decentralized, has been the main difficulties in the blockchain PoW design. In particular, the community hopes to help those relatively weak computing power devices, like smartphones and IoT devices, to join in the mining process. At the same time, as mobile devices on the market have already appeared in support of AI computing chip or computing library, computing framework on mobile AI chips can also participate in AI inference. Compared to a full node holding larger data models, the mobile terminal needs to be customized to screen off large sized data models. Cortex main chain will release Android and iOS client applications to implement:

- The idle GPU computing device on the SoC, ARM architecture CPU / GPU to participate in mining. For example, a TV box's GPU performance is actually very good, but it is basically idling 90% of the time.
- Smartphones can participate in mining while charging in office during the day or while its owner is sleeping at night, as long as the algorithm provides the GPU on mobile devices to have a fair revenue competitiveness.
- Smartphones or other devices equipped with an AI chip will automatically switch between the main blockchain and AI inference.

The inference computational power on mobile devices may be limited by the software technology of the chip supplier. Different software vendors are encapsulating the computational protocol in their methods. The Project Company will be responsible for the preparation of the abstraction layer interface and filter out smart contracts to remain lightweight for mobile devices.

5 Applications and Future Work

5.1 Application Cases

Cortex blockchain could host AI DApp as follows:

5.1.1 Information Service

- Personalized recommendation system: based on user profile and show / click log in the transactions, to recommend news that matches interest.
- Image search engine: based on the image data, retrieve similar images.
- News writing: based on the text corpus, to generate a similar style of text.
- Automatically summarize: based on the text corpus, to generate a summary.

Archived copy · cryp2nova.com

5.1.2 Financial Services

- Credit score: Calculate the credit rating based on the user's online data.
- Intelligent investment advisory: Automatic trading decisions based on financial data.

5.1.3 AI Assistant

- Automatic Q&A: chat robot, based on the user dialogue to generate answers.
- Industry knowledge graph: expert system, can be used in medical, consulting and other industries
- Speech Synthesis: Generate another voice of similar style based on the user's speech data
- Face attribute prediction: based on user-uploaded face data on age, gender, emotion and other attributes to make a judgment

5.1.4 Simulation Environment

The application of reinforcement learning can predict the output of the model through the environment to realize intelligent decision-making and verification for automatic driving, the game Go, games and etc.

5.2 Future Work

The goal of Cortex is to be a leader in the field of machine learning on blockchain. For new technologies that may be invented in the future, the Project Company will continue to innovate and integrate. This chapter describes some potential future work directions.

5.2.1 Data Privacy

Cortex implements the blockchain of inferences consensus, however, typical scenarios for machine learning applications such as paramedical assistance, bioinformatic recognition, speech recognition, etc., require strict privacy protection mechanisms to protect data privacy and model intellectual property. Related technical directions are: Differential Privacy [11], Zero-Knowledge Proof [12], Homomorphic Encryption [13] and etc. The Project Company closely pays attention to the progress and integrates when new technologies become practically available.

5.2.2 Block Size and TPS Improvement

Due to the limitations of the consensus mechanism of PoW, Cortex will also encounter the problem of block size and TPS issue. Currently, possible solutions are PoS, DAG, cross-chain communication and so on. Essentially, scaling directly into the blockchain will be a trade-off due to the limitations of the distributed system CAP theorem [14], among system consistency, availability, and persistence. The Project Company will also study scaling issues and improve network performance without sacrificing core assumptions.

5.2.3 Adapter of AI Chips on Mobile Devices

Since the latest mobile devices are often equipped with dedicated AI chips, Cortex's inference framework can also call the AI chip computing interface of mobile devices, adapt the instruction set of each chip, and develop the lightweight intelligent inference in the virtual machine.

Archived copy · cryp2nova.com

6 Roadmap

2018 Q1	Issue ERC20 token
2018 Q1	Listed on multiple mainstream exchanges
2018 Q3	Testnet for mining, namely Bernard
2019 Q1	Testnet for smart AI contracts, namely Dolores
2019 Q2	Mainchain, namely Arnold genesis block. Burn ERC20 Coin to start mining the Cortex Coin, covertible on a 1:1 ratio

7 Token Model

7.1 Cortex Coin (CTXC)

7.1.1 Rewards for Model Provider

In contrast to the traditional blockchain, in which the reward for each packaging block is paid directly to the miners, to motivate developers to submit richer and better models, Endorphins invoking the contract, are not only allocated to miners who help package the block on full node, but also used to pay the model providers. The proportion of fees charged will be determined by the market gaming price, similar to the Ethereum Gas mechanism. AI smart contracts having higher max Endorphins price have higher priority to be executed. A single execution of a contract will cause a fee of Endorphins Limit * max Endorphins Price.

7.1.2 Cost for Model Provider

To prevent overwhelming model committers from submitting and storing attacks - such as arbitrarily submitting almost unusable models or submitting the same model frequently to consume storage resources - each model submitter must pay for storage. More calls to use the model, more revenue the providers can earn, thereby encouraging the model provider to submit better models.

7.1.3 Model Complexity and Endorphin Spend

Endorphin is a measurement of the amount of computing resources spent on a hardware level within a virtual machine when bringing a data model into a contract during inference. Generally speaking, the cost of Endorphin is proportional to the size of the model. Cortex also sets an upper bound of 8GB on the parameter size of the model, corresponding to up to about 2 billion Float32 parameters.

7.2 Token Distribution

In respect to the few natural constants of the physical universe, we choose the speed of light propagating in vacuum as total amount of token, 299792458. 60,000,000 CTXC (about 20.01%, all percentages are two-decimal approximation, the same below) will be allocated to early investors.

Cortex Coin Total	299,792,458	100%
Cortex Coin Miners (Mining Reward)	150,000,000	50.03%
Investors (Genesis Block)	60,000,000	20.01%
Project Foundation (Genesis Block)	74,792,458	24.95%
Cortex Lab	45,000,000	15.01%
Project Marketing Lock	27,000,000	9.01%
Challenge Bounty	2,792,458	0.93%
Advisors/Academia/Community (Genesis Block)	15,000,000	5.00%

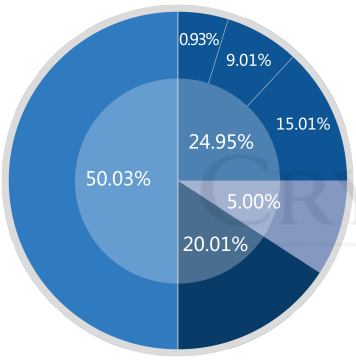
Table 4: Token distribution details 1

Table 4 and Table 5 describe the details of the distribution of Cortex Coins.

Archived copy · cryp2nova.com

50.03%		Cortex Coin Miners (Mining Reward)	Miners receive rewards for contributing the computational power, maintaining the Cortex blockchain, and running the full node of AI DApps.
20.01%		Investors (Genesis Block)	Investors' financial support is used for Cortex main blockchain development, business development, partnerships, and community support.
24.95%		Project Foundation (Genesis Block)	
	15.01%	Cortex Lab	Responsible for the development of Cortex protocol, data model maintenance, software upgrades, hardware solutions.
	9.01%	Project Marking Lock	Quarterly accelerated unlock 1: 2: 3: 4 to airdrop / marketing, reward data scientist / Kaggle communities and new lab members for development.
	0.93%	Challenge Bounty	Used to reward community developers for interesting projects like BTC, XMR address fungibility analysis.
5.00%		Advisors/Academics/Community (Genesis Block)	AI for business and academia, scholar community and open source framework developer community.

Table 5: Token distribution details 2



50.03% Cortex Coin Miners (Mining Reward)

20.01% Investors (Genesis Block)

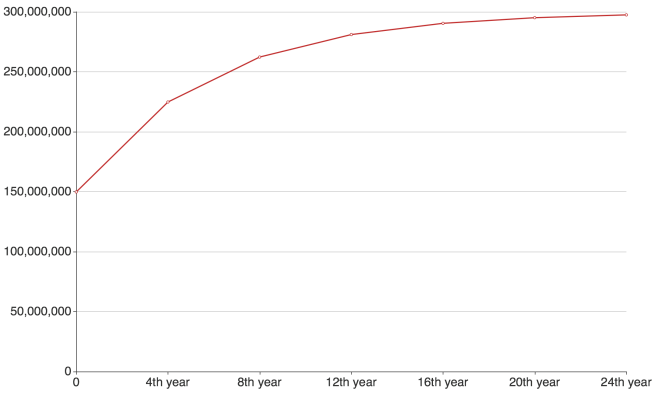
24.95% Project Foundation (Genesis Block)

5.00% Advisors/Academics/Community (Genesis Block)

7.3 Token Issuance Curve

Cortex Coins will be classically capped 299,792,458 coins. 150,000,000 Cortex Coins will be issued by mining.

Archived copy · cryp2nova.com



The first 4 years	75,000,000
The second 4 years	37,500,000
The third 4 years	18,750,000
The fourth 4 years	9,375,000
The fifth 4 years	4,687,500

And so on, halves every four years classically as Bitcoin.

8 Core Team Members

Ziqi Chen *co-founder CEO* B.E. in Civil Engineering from Tsinghua University. M.S in Civil Engineering from Carnegie Mellon University and M.S. in Computer Science from University of California, Santa Cruz (USCS). In USCS, the birthplace of AdaBoosting and Online Learning, Ziqi was advised by Prof. David P. Helmbold in theoretical machine learning and various applications, including the Go algorithm. Before that, he served as Principle Research Scientist in SFTC company, and was responsible for the mesh generation algorithms in finite element method, which was used in aerospace and weapons research. With rich e-commerce business experience and years of practice in blockchain, Ziqi started Waterhole.io. He has solid understanding and first hand experience in mining, computing and digital wallets and a deep understanding of mining devices, consensus algorithms and public blockchain ecosystem. In Bitcoin, Ethereum, Zcash and other communities, his company operates a considerable amount of hash power and the wallet serves many cryptocurrency users.

Weiyang Wang *co-founder CTO* National Mathematics Contest winner to be a recommended student enrolling in Tsinghua University. B.E. in Aerospace Engineering and minor in Economics from Tsinghua University. M.S. in Statistics from University of Chicago. Co-advised by Prof. Peter McCullagh, Logistic Regression inventor winning 1990 COPSS Presidents' Award, and Yali Amit, the founder of Random Forest. Weiyang's research interests lie in statistical machine learning and game theory in University of Chicago. By implementing algorithms of CNN / RNN algorithms from many leading papers in deep learning field, the DeepInsight team, where he was in, had developed projects of OCR / Transfer Learning / Face Recognition / Distributed Computing, being selected into the Awesome MXNet Project. He had won the silver medal in the Kaggle Contest. He is also responsible for translation services of many AI reports of O'Reilly. As a certified FRM and CFA Level II candidate, he had won championship in the 2017 Dorahack Hacking Marathon , leading a team creating ABS (Asset-Backed Securities) system prototype , based on the IBM Hyperledger. Weiyang is also mastering number theory, encryption and other related fields. He has worked for Jingdong Finance and Wanda Research Institute.

Xiao Yan *Chief deep learning engineer* National Olympiad in Informatics(NOI) winner to be a recommended student enrolling in Tsinghua University. B.E. in Electronic Engineering from Tsinghua University. Currently 5th year Ph.D. candidate in Computer Science in Tsinghua University. During Ph.D. program, he has published many top papers, including ACM Siggraph. He is mastering in OI coding competition, research and development in CUDA GPU, FPGA hardware programming, cryptography, and etc. Xiao's research area includes finite element analysis, tensor analysis, turbulence, deep learning and reinforcement learning. He has also worked in several applications, including machine learning, blockchain, quantitative trading. He managed over 10 million US dollars in PE fund and gained substantial return on investment. He is interested in quantum computers and financial macro analysis and has deep understanding in mining software, distributed computing and consensus algorithms. He has served as Founder / CTO or technical director, in a number of AI companies and private equity funds.

Yang Yang *Blockchain chief-engineer* B.E. in Computer Science from Dalian University of Technology, and M.S. in Computer Science from Tsinghua University. Yang has deep understanding in

Archived copy · cryp2nova.com

distributed databases, Hyperledger / EEA. He worked in British Telecom and IBM. He developed multiple social software such as Dial Wizard/Lan Xin and so on. He has over 10 years of Java and C ++ backend experience, including develop VisiDB, and is able to implement a database in C / C ++ by himself. Yang is an expert in blockchain source code and encryption technologies. He has development experience in a variety of cryptographic digital currencies. Yang has first hand experience in backend design and development of high persistent TCP communications in commercial softwares, similar scale with WhatsApp and WeChat. Yang also has expertise in all kinds of smart contracts and hard fork coins.

9 Advisors

9.1 Technical advisors

Jia Tian *Chief-Scientist* National Physics Contest and Biology Contest winner to be a recommended student enrolling in Tsinghua University. B.E. and M.S. in Computer Science from Tsinghua University. Jia was a distributed system expert with years of academic and industry experience. He worked at Baidu and Alibaba, then the architect of so.com, a search engine with over 100 million daily active users, and build another recommender system. As a co-founder of multiple high tech startup companies, Jia has first-hand experience in several fields, including search engine, recommender systems, AI, FinTech and etc. His first company, Wolong Cloud, was acquired by Alibaba. Later, he joined Beijing Machine Learning Information Technology Co., as CTO, he led projects in recsys, chatbot, medical imaging and etc. After that, he joined Pony.ai, an autonomous vehicle startup company, which is invested by Sequoia Capital, IDG Capital Partners in angel round. He previously served as Chief Scientist for BitFund, and blockchain advisor for several ICO project. He is also one of the earliest investors in Bitcoin Zcash and Bitfinex. Jia's research interests includes quantum computing, nuclear fusion and computational neuroscience.

Matt Branton *Engineering Advisor* Matt Branton is a distributed computing expert, with over ten years of experience in payment and trading system design and architecture. He was a founding partner at World Financial Desk, a high frequency trading firm specialized in fixed income and derivative products, and responsible for trillions in volume. An early Bitcoin adopter, he founded Coinlock for encrypted micro-payment content delivery and has spent the last five years building smart contract and distributed ledger systems at Ember Financial. He has multiple blockchain patents pending, including his work on synthetic mining, and novel derivatives products.

9.2 Academic advisors

Whitfield Diffie *Professor* Cryptographer and one of the pioneers of public-key cryptography. Diffie received an honorary doctorate from the Swiss Federal Institute of Technology in 1992. He is also a fellow of the Marconi Foundation and visiting fellow of the Isaac Newton Institute. In July 2008, he was also awarded a Degree of Doctor of Science (Honoris Causa) by Royal Holloway, University of London. He was also awarded the IEEE Donald G. Fink Prize Paper Award in 1981, The Franklin Institute's Louis E. Levy Medal in 1997 a Golden Jubilee Award for Technological Innovation from the IEEE Information Theory Society in 1998, and the IEEE Richard W. Hamming Medal in 2010. Diffie was elected a Foreign Member of the Royal Society (ForMemRS) in 2017. Together with Martin Hellman, Diffie won the 2015 Turing Award, widely considered the most prestigious award in the field of computer science.

9.3 Business advisors

Vincent Zhou *Founder of FBG Capital*

Yahui Zhou *CEO and Chairman of Beijing Kunlun Tech Co. (300418.SZ)*

Heting Shen *Retired CEO and Chairman of China Metallurgical Group Corporation 601618 (SHA) and 01618.HK (MCC)*

Archived copy · cryp2nova.com

Guy Corem *CEO of DAGlabs*

9.4 Organization partnerships

Tsinghua University Laboratory of Brain and Intelligence

Shanghai Jiaotong University School of Mathematical Sciences

University of California, Berkeley Blockchain Labs

Stanford University

10 Leading Investors

Bitmain A world-class business provider of cryptographic currency mining hardware and blockchain software solutions. The company is also investing heavily in the research and development of AI chips.

FBG Capital FBG Capital, founded in 2015, with primary focus on investment and trading digital asset including but not limited to Bitcoin. FBG Capital is one of the most active hedge funds in Asia. FBG Capital also incubates promising blockchain projects and companies through investments, strategic advisory and technical support. FBG Capital is formed with professionals from both financial industry and blockchain technology space.

More investors see official website at : http://www.cortexlabs.ai

A Cortex Mathematical Representations of AI on blockchain

The traditional AI definitions and Cortex instructions can be described mathematically

i: Cortex instruction

I: $\{i\}$, Cortex instruction set

s: Cortex instruction parameters, $s \in R^{N^d}$

S: Cortex status space of instruction parameters

r: $\in R$, Cortex string of model representations

R: Cortex effective space of strings for model representations

m: $(i_1, i_2, ...i_k) \in M$, sorted Cortex instruction set, Cortex string model representation full descriptions of r

M: Space of sorted effective Cortex instruction set

$\mathfrak{S}$: Space of Cortex model parameters M^S

$\mathfrak{F}$: Supported AI computing framework mapping set

$\mathfrak{f} \in \mathfrak{F}$: AI computing framework mapping functions $\mathfrak{f} : f(m) \to c$

d: Data samples

D: $\{d\}$, Dataset

$P : \{p : D \to (D_{\text{train}}, D_{\text{test}}, D_{\text{val}})\}$, A partition of dataset D

$\mathfrak{M}$: The measurement of model training and model ranking service $\mathfrak{M} : f(\mathfrak{I})(d) \to R$

Archived copy · cryp2nova.com

$A : \{a\}$, Set of model execution strategies A complete Machine Learning / AI problem $q \in Q$, computable on a physical device is a four-tuple:

$Q : (\mathfrak{F}(M), P(D), \mathfrak{M}, A)$, Complete definitions of problem set

It has a specific framework to select a model, with the form of probability model or neural network model $\mathfrak{f}(M)$

It has a partition $p(D)$ that contains specific data samples for training, inferencing, or checking for generalization

It has a measurement waiting to be optimized, or need sorting $\mathfrak{m}$

It has a machine learning strategy a, to determine the current process need only inference, consensus or training

B Summary of Basic Types for Deep Learning

Referred machine learning frameworks in chapter 1.1 express different machine learning problems, here are a few classic machine learning modes:

B.1 Supervised Learning

Supervised Learning is defined under the settings which use training dataset as data input and the label of these data samples as the target output are given simultaneous. Machine learning algorithms process, "train" these data samples according the input and output, in order to make prediction for those data samples in the test dataset which are not well labeled, or in the validation dataset as a double-check for the correctness of the machine learning model. As in hand-written digits image recognition problem, the form of data samples in dataset D could be 64×64 bitmap image with each pixel has 1 or 3 channels of 8 bit information. The labels, $Y\prime$ stands for the according class for that data point, in this case, digits classification is a 10-way classification problem, digit 0 has a one-hot label vector $Y\prime$ as $[1, 0, 0, 0, 0, 0, 0, 0, 0, 0]^T$, and so is $1 - 9$. The form of P could be $||Y\prime - M(D)||$, and model M could be selected from neural networks or simply a fully-connected bayesian classifier.

B.2 Unsupervised Learning

Unsupervised Learning is defined under the settings which the label of a dataset is totally absent. By utilizing the unsupervised algorithms, we hope to find the hidden structure from the dataset itself instead of making accurate predictions. Sometimes in the form of dimension reduction(by PCA, AutoEncoding, etc.) with respect to minimizing the reconstruction loss, or in the Clustering problems, $Y\prime$ is no longer provided, as in a conventional K-means, $P(Y\prime, M(D))$ is degraded to

$$\arg\max_{S} \sum_{i}^{k} \frac{1}{2|S_i|} \sum_{x,y \in S_i} ||x - y||^2$$

so we can make abstract classes $A, B, C, ...$ to fit the data samples which are in clusters in spaces of particular dimensions.

Some typical unsupervised learning algorithms are

- **Generative Adversarial Networks**
  Generative Adversarial Networks, also known as GAN is one mainstream family of unsupervised neural networks, which is widely utilized in image and sequence generation. It is composed by a discriminator deep neural network model$\mathbb{D}$ as well as a generator deep neural network$\mathbb{G}$. $\mathbb{G}$ network take a data sample from a fake picture $\boldsymbol{z}$ under certain sampling algorithms as input (usually from a high-dimensional random gaussian distribution), after some level of unpooling operations (sometimes known as decoding network), a feature $\mathbb{G}(\boldsymbol{z})$ output is given to $\mathbb{D}$ as input. The role of $\mathbb{D}$ is to discriminate between data samples of $\mathbb{G}(\boldsymbol{z})$ and real data samples d. In short, $\mathbb{G}$ maximize the probability to deceive $\mathbb{D}$, and $\mathbb{D}$ simultaneously discriminate the true data and the fake data that $\mathbb{G}$ generated halfway. The

Archived copy · cryp2nova.com

loss function is in the form like:

$$\min_{\mathbb{G}} \max_{\mathbb{D}} \mathbb{E}_{\boldsymbol{x}\in\chi}[\log \mathbb{D}(\boldsymbol{x})] + \mathbb{E}_{\boldsymbol{z}\in Z}[\log(1 - \mathbb{D}(\mathbb{G}(\boldsymbol{z})))]$$

- **Variational AutoEncoder**
  VAE Network is one of mainstream generative networks in practical. An encoder part of VAE project the data samples into simple multi-dimensional space, (e.g. $\mathcal{N}(0, 1)$) and the decoder part samples from this distribution directly to generate new data. The metric is written as following:

  $$\mathcal{L}(\phi, \theta, \boldsymbol{x}) = D_{KL}(\boldsymbol{q}_\phi(z|x)||p\theta(\boldsymbol{z})) - \mathbb{E}_{\boldsymbol{q}_\phi(\boldsymbol{z}|\boldsymbol{x})}(\log p_\theta(\boldsymbol{x}|\boldsymbol{z}))$$

  ϕ,θ are training hyperparameters, $\boldsymbol{x}$ is conform to the distribution of input data, $\boldsymbol{z}$ is the feature vector after the encoding process. Encoder optimizes the first part of the right hand side of the equation, while decoder optimized the second part.
- **Manifold Learning**
  The core of manifold learning is to find a decent manifold projection to find out the topological structure from data. By utilizing these kind of algorithm like t-SNE, the distance or similarity between data sample pairs are well-kept, but after these projections, data samples in different classes shows a linearly divisible pattern.

B.3 Other Types of Learning

B.3.1 Semi-supervised Learning

With respect to Semi-supervised learning [15], partial labels are given in the training dataset, and only some basic hypothesis are made to those data samples which are not labeled.

B.3.2 Active Learning

In active learning, there are multiple stages of label querying, training and prediction, so we can mine information from those hard samples (but not bad cases, if human interaction is also present in the data pipeline). The hard cases (wrongly predicted data, or those samples falling in the margin region when implementing SVM algorithm) allows us to learn faster from data. A feasible metric is

$$P = P\prime(Y\prime_{D_1}, M_1(D_1)) - P\prime(Y\prime_{D_1+D_2}, M(D_1 + D_2))$$

.

B.3.3 Reinforcement Learning

The main target of Reinforce Learning is maximizing the expected reward for an agent within a particular environment E along the time span T, by take reasonable action A according to current state S, or past state series.

[16] Definitions:

- State S: State of the environment, i.e. current chess setting during a chess game.
- Action A: An operation drawn from an action set, can be taken with respect to the current state and hereby change the current state to another, i.e. move a Castle during a chess game.
- Transfer Function $T : S \times A \times S\prime \to [0, 1]$ stands for the state transfer probability of state S is transferred to another state $S\prime$ by taking an action A, noted as $T(S, A, S\prime)$
- Reward Function $R : S \to R$, action reward:$S \times A \times S \to R$
  Action pattern of agents sometimes are defined in a closed form known as Markovian Decision Process, here MDP is a quadruple (S, A, T, R), with maximizing the total reward as its optimizer's main target.
  During the process, a decay coefficient is given as part of the problem setting, such that

  $$\text{Total Value } V = \text{Current Reward } R + \text{Decay Coefficient } \gamma \times \text{Value in the last time step } v$$

- A policy, could be a definite mapping function $S \to A$ or a random choice according to distributionπ s.t. $\pi(s, a),\ s.t.\ \sum \pi(s, a) = 1$

Archived copy · cryp2nova.com

With respect to reinforcement learning problems, if MDP is known, the whole process is dynamic programming; else in the case where MDP is unknown, we should implement deep learning.

Two common methods are known as Q Learning and Policy Gradient.

The most import thinking in Q Learning is from dynamic programming. The process of computing $Q : S \times A \to R$ through Bellman equation

$$v(s) = \mathbb{E}[R_{t+1} + \gamma v(S_{t+1})|S_t = s]$$

Algorithm could continuously upheave the reward function until convergence. In order to avoid state space overwhelming, we introduce Deep Q-Learning (DQN) to fit Q function.

Policy Gradient simulates possible actions according to input states. To maximize the expectation of reward sum during multiple runs of simulation, we tend to take gradient updates to achieve an optimized state if the reward gets higher. The whole model can be seen as a End-To-End $S \to A$ policy optimizer.

Q-Learning finds a best action under a certain state based on its expectation, but Policy Gradient directly predicts the action. Actor-Critic combines the policy based and value based approach, which actor performs like ordinary PG, and critic makes value predictions, hence Asynchronous Advantages Actor-Critic approach(A3C) is a smoother optimizer which makes an accurate assessment for every current state along with the environment. Another approach which combines the two methods is known as DDPG, which we would not give the details here.

B.3.4 Transfer Learning

Transfer Learning is a quick developing branch of machine learning which transfers the past models to the new domain without excessive computing power, even the underlying distribution of embedding vector in the domain differs from the source domain. This ensures the re-usability of existing models and domain knowledge, avoiding repeat training on older datasets.

- Inductive Transfer Learning

  Among Kagglers, or competitors in entry-level competitions, a classic process which utilizes an existing model to accelerate training new models on the dataset competition provided is called "Fine-Tuning", and the existing model is called "Pre-trained Model". If we already have well performing models trained on a sufficiently large and clean dataset, we could speed up model training significantly, this is not limited to basic classifiers, but also those with specific structures like object detectors and semantic segmentation models, to acquire a new model with a higher performance than ones trained on this small dataset, even in the case that source domains differs from target domains. A state-of-the-art model, trained carefully on high quality and large dataset in HPC center, can setup a new baseline for the researchers in their field.
- Transductive Transfer Learning

  Transductive transfer learning makes the assumption that the dataset of source domain and target domain may vary, while data samples from source domain are well labeled, the ones from target domain is not. This kind of transfer learning can dominantly boost the performance compared to simply training model on source domain, or simple fine-tuning.

  A typical proof of it is J-MMD Algorithm. By introducing unsupervised MMD loss together with the classic cross-entropy loss utilized by softmax classifier, we can get models performing equally well on dataset composed from images of different quality and resolution. Also, neural network structures like detectors and semantic segmentation algorithms get boosted too.
- Unsupervised Transfer Learning

  Generalized transfer learning can also be used in pure unsupervised learning when labels of the whole dataset are absent, or unnecessary. Introducing similarity between samples or distribution into loss function can perform particular function in some special problems.

  In recent several years, an AI application called Neural Style Transfer appeared. This application can apply styles in S (as Style Picture) like Vincent van Gogh's Starry Night or Edvard Munch's The Scream to any user uploaded content C(= Content Picture) , to form a new image M (Mixed Picture), such that M has similarity in different scales with S and

Archived copy · cryp2nova.com

C. Although there's no way to find a mature business model in this technique, Ostagram from Russia, Prisma App from America and Philm App from P.R.China are very successful cases on attracting users to try out this technique, and soon Meitu and many camera apps will join this battle. Inspired by this algorithm, AI aesthetics are taking attention from artist and AI engineers. By researching images of human faces as well as the nature environment, they can modify and cartoonize any kind of image and videos by any kind of AI filters.

C Distributed Cloud Computing Practice for Deep Learning

Although many deep learning and machine learning frameworks already implemented mainstream algorithm utilizing single device in single machine, practically they support multiple devices in single machine / multiple devices in multiple machines parallel computing , avoiding the memory constraints happened in the cases using single device. Most of them transfer parameters between devices through PCI-E v3.0 bus, or using 10G to 56G Ethernet. This chapter is NOT for personal devices, but mat be feasible for data centers with sufficiently high specifications. We warmly welcome these HPC centers to join us in the alliance chain instead of public chains. They can sell their computing power to those in need of AI model training services.

Three main policies on distributed deep learning and machine learning are: Model parallel, data parallel, and other approaches.

C.1 Model Parallel

Model parallel refers to a single node - single GPU, KNL, or host memory etc. It does not contain the whole learning model, but rather, is distributed on different devices. Different pieces first compute on their own, before finally reducing to one result. Since the connection between submodels is in a great number, distributing models to different nodes often incurs huge communication overhead, which makes model parallel unadaptive to many computing scenarios. Based on different type of dividing models, model parallel are in two typical classes.

1. The first class model parallel divides different layers in neural network, or different stage in computing pipeline into subsets. In other words, different nodes are responsible for different layers/stages. RNN is adaptive to this kind of model parallel, in the case that single device doesn't fit the whole model, this approach is easy and effective to implement. It is unusual in training Convolutional Neural Networks, since the model is relatively small compared to the data and features in Graphic memory or host memory. However, in Google's One Weird Trick Krizhevsky [17], removing the fully-connected layer from CNN backbone remarkably reduced communication overhead during training. This engineering trick is proven to be effective in large scale CNN deployment like Baidu Image Search Service.
2. The second class model parallel divides each layer into several subsets, with reduction operations taking place after each node processes its own subset. This is very unusual in that it may cause serious network traffic problems, so that we seldom observe this kind of model parallel in practical RNNs and CNNs implementations. An exception is DSSTNE, open-sourced by Amazon dss [18], a distributed sparse DNN implementation used in recommender systems. Granted for sparsity, no huge overhead in communication is observed, while the problem setting is still a case outside RNNs and CNNs and not general enough.

In all, model parallel hardly provide a stable training acceleration approach according to its communication issues. Some deep learning framework provide support to this feature like TensorFlow Apache SINGA, but it's unusual in practice. For large scale distributed training of machine learning models, data parallel shows a more stable behavior.

C.2 Data Parallel

There are two main starting point of data parallel : System architecture design, and optimizing algorithm design. Respect to system architecture, mainstream distributed machine learning framework utilizes MPI or Parameter Server. The Parameter server is adaptable to many techniques to reduce bandwidth cost incurred in communication, it's the de facto baseline in the research of system architecture design. Eric Xing gives an introductory review related to this field Xing et al. [19], classifies

Archived copy · cryp2nova.com

data parallel paradigm into BSP (Bulk Synchronous Parallel), ASP (Asynchronous Parallel) and SSP (Stale Synchronous Parallel). BSP is the naive solution in the early stage of parallel computing developed, the main feature of BSP is it simultaneously updates parameter on each node in the cluster which consumes much Ethernet bandwidth. The margin acceleration rate of distributed algorithms is lower as the number of nodes grow; ASP , an idea introduced after HOGWILD! , avoids bandwidth consumption that happens when parameters are transferred simultaneously. This extreme bandwidth-saving policy may cause model gaps between nodes and a slow convergence speed, hence slowing down the whole training process. DistBelief, the first generation deep learning framework from Google chose this paradigm. To alleviate this, Model AverageZinkevich et al. [20], which is also widely used in BSP, takes the average of each model and propagates the parameters to each node, which makes convergence smoother. In general, the acceleration curve is concave ; SSP is a tradeoff between bandwidth saving and model convergence given limited asynchronous condition, is one option of best practices. We should notice that the convergence proof of SSP is indeed a breakthrough in the field of distributed machine learning, but Eric Xing's work focused on those conventional models, not deep neural networks. Experiments are ongoing in frameworks like MXNet and PaddlePaddle etc, and additional proof about the convergence of deep models is given by Eric Xing laterKeuper [21]. But considering all the conditions, BSP is still the mainstream in distributed deep neural network training, and MPI kept its glory in most of the computing frameworks.

SSP created a new research direction in the field of distributed learning. The followers introduced multiple ideas combined with various optimization methods. Taifeng Wang from MSRA Zheng et al. [22] mentioned that ASP is fundamentally the same as SSP which caused stale state of node, and worsened the convergence of algorithms. A Taylor expansion in additional to naive SGD compromises the convergence accuracy and speed. This idea is straightforward and effective before having a rigorous proof, and accepted and implemented in frameworks like MXNet. These kind of works inspired many researchers who seeks to make progress with optimization methods. By now, most of the deep learning algorithms take advantage of SGD with a fast convergence speed slightly slower than GD, with a sublinear time complexity $O(1/\epsilon^2)$ under the circumstances when the loss function is strongly convex. Although GD as a option provides $O(n/\epsilon)$ time complexity convergence speed, its natively memory inefficient which caused OOM (out of memory) of GPU. SGD as the mainstream approach, have many descendants like Adagrad, Adam, Momentum etc, which doesn't remarkably improve the convergence speed, until the publication of SVRG Johnson and Zhang [23]. It maintains a global grad and asynchronously modify it (may have a lag of $1\tilde{5}$ epochs). The local grad and current grad works together on updates of parameters. Wujun Li's team from Nanjing University successfully applied SVRG on distributed machine learning systems Zhao et al. [24] and acquired a linear speedup. More attempts about applications of SVRG is still in progress. In theory, a series of work is trying hard to find a "best" tradeoff between architecture and optimizer to achieve the combination of "asynchronous" and "linearly speeding up".Huo and Huang [25] is a typical example. We may see a widespread acceptance of SVRG during deep neural networks training.

In practice, BSP keeps its dominant position in deep learning paradigm, while more works focus on bandwidth saving by two main approaches, either from structure designing or optimizer design. Lecun introduced EASGD Zhang et al. [26], which can be applied to BSP or ASP, which computes Elastic Difference of the model parameters between Worker and Server in Parameter Server before communication happen. In advance, Microsoft CNTK introduced 1-bit SGD Seide et al. [27], by quantization process, parameters in Float32, Float16, Int8 are clipped to 1 bit, so the bandwidth saving is more effective. Additional pretraining is necessary to ensure convergence. According to the experiment from CNTK, 1-bit SGD is effective in bandwidth saving, but beyond the incurred loss of accuracy, multi-machine parallel acceleration is harmed. The default setting CNTK recommended is Block Momentum Chen and Huo [28]. Starting from Model AverageZinkevich et al. [20], it refines the model update stage using filtering a series of updates collected from training history, so that in each step of parameter update it shows a smoother pattern of training than ordinary. Block Momentum itself is not a work focusing on bandwidth compression, but rather a technique making better convergence curve to accelerate the whole training process. This technique ensures CNTK to be one of the most efficient framework which implements the multi-machine version of training deep neural networks.

A breakthrough in structure , Poseidon Xie et al. [29] Zhang et al. [30] from Eric Xing's team from Carnegie Mellon University, introduce P2P communication above Parameter Servers' Master-Worker Structure. In this structure, each worker can dynamically decide to choose either PS or any

Archived copy · cryp2nova.com

other worker to communicate with, according to bandwidth usage surveillance and prediction. The core reason that renders the P2P structure feasible, is the introduction of the idea of SFB (Sufficient Factor Broadcasting) introduced by Poseidon. SFB factorized the parameter matrix to transfer into the product of two low rank matrices, called SF(Sufficient Factor). The throughput of SF sometimes is remarkably lower than the original matrix. The structure of P2P communication with SF between workers alleviate bandwidth blocking problems occurring in parameter server, and this work is tested in real cases: Inception Networks can be accelerated linearly in an ordinary Ethernet network.

Engineering teams from industries mostly choose optimization for BSP structure. Work from Baidu. Inc Gibiansky [31] and Uber's TensorFlow version based on Baidu. Inc's work, Horovod hor [32], changed the conventional All-Reduce operation in BSP to Ring All-Reduce, such that a much larger scale of cluster can be accelerated on model training. In addition, IBM's Power AI Cho et al. [33] introduced Multi Ring All-Reduce and claim they achieved near perfect performance on 256 devices, however, IBM did not reveal much details of their work.

C.3 Others

From the above, the key to improve parallel training performance lies in the balance between bandwidth saving and effective convergence. Therefore, from another perspective, part of the model compression technology itself can also be classified as parallel deep learning, except for approximating the parameters by accuracy loss.

In addition to Dettmers [34], an ICLR 2018 double-blind review of authors [35] deserves our attention. By using methods of momentum correction, local gradient cropping, momentum factor masking, and warm up train, the exchanged parallel training gradient parameters were compressed hundreds times. This could be used for large-scale distributed training on cheap 10 Gigabit Ethernet.

D Cortex Lab Existing Achievements

The Cortex lab's world-leading models include:
OCR(ReCaptcha and Algebra Equations 99%+, Simplified Chinese 98%+)

Models that are expected to be developed and exposed to the Cortex eco-community include:

Instance-Aware Segmentation and Mask-RCNN
Keypoint Detection
Image Captioning, AIC competetion top 20%
Go algorithms

CRYP2NOVA

References

[1] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. https://bitcoin.org/bitcoin.pdf, 2009.

[2] Ethereum white paper. https://github.com/ethereum/wiki/wiki/White-Paper.

[3] Conway's game of life. https://en.wikipedia.org/wiki/Conway%27s_Game_of_Life.

[4] Yonatan Sompolinsky and Aviv Zohar. Secure high-rate transaction processing in bitcoin. In *Financial Cryptography and Data Security - 19th International Conference, FC 2015, San Juan, Puerto Rico, January 26-30, 2015, Revised Selected Papers*, pages 507–527, 2015. doi: 10.1007/978-3-662-47854-7_32. URL https://doi.org/10.1007/978-3-662-47854-7_32.

[5] Rsk:bitcoin powered smart contracts. http://www.the-blockchain.com/docs/Rootstock-WhitePaper-Overview.pdf.

[6] Craig Gentry, Shai Halevi, and Nigel P. Smart. Homomorphic evaluation of the aes circuit. Cryptology ePrint Archive, Report 2012/099, 2012. https://eprint.iacr.org/2012/099.

Archived copy · cryp2nova.com

[7] Dagger-hashimoto. https://github.com/ethereum/wiki/blob/master/Dagger-Hashimoto.md.

[8] Alex Biryukov and Dmitry Khovratovich. Equihash: Asymmetric proof-of-work based on the generalized birthday problem. Cryptology ePrint Archive, Report 2015/946, 2015. https://eprint.iacr.org/2015/946.

[9] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. *CoRR*, abs/1512.06473, 2015. URL http://arxiv.org/abs/1512.06473.

[10] Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. *CoRR*, abs/1510.00149, 2015. URL http://arxiv.org/abs/1510.00149.

[11] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. *Calibrating Noise to Sensitivity in Private Data Analysis*, pages 265–284. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006. ISBN 978-3-540-32732-5. doi: 10.1007/11681878_14. URL https://doi.org/10.1007/11681878_14.

[12] S Goldwasser, S Micali, and C Rackoff. The knowledge complexity of interactive proof-systems. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, STOC '85, pages 291–304, New York, NY, USA, 1985. ACM. ISBN 0-89791-151-2. doi: 10.1145/22145.22178. URL http://doi.acm.org/10.1145/22145.22178.

[13] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 169–178, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-506-2. doi: 10.1145/1536414.1536440. URL http://doi.acm.org/10.1145/1536414.1536440.

[14] Seth Gilbert and Nancy Lynch. Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59, June 2002. ISSN 0163-5700. doi: 10.1145/564585.564601. URL http://doi.acm.org/10.1145/564585.564601.

[15] Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. *Semi-Supervised Learning*. MIT Press, 2006.

[16] Martijn van Otterlo and Marco Wiering. *Reinforcement Learning and Markov Decision Processes*, pages 3–42. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. ISBN 978-3-642-27645-3. doi: 10.1007/978-3-642-27645-3_1. URL https://doi.org/10.1007/978-3-642-27645-3_1.

[17] Alex Krizhevsky. One weird trick for parallelizing convolutional neural networks. *CoRR*, abs/1404.5997, 2014. URL http://arxiv.org/abs/1404.5997.

[18] Amazon dsstne: Deep scalable sparse tensor network engine. 2016.

[19] Eric P. Xing, Qirong Ho, Pengtao Xie, and Dai Wei. Strategies and principles of distributed machine learning on big data. *Engineering*, 2(2):179 – 195, 2016. ISSN 2095-8099. doi: https://doi.org/10.1016/J.ENG.2016.02.008. URL http://www.sciencedirect.com/science/article/pii/S2095809916309468.

[20] Martin Zinkevich, Markus Weimer, Lihong Li, and Alex J. Smola. Parallelized stochastic gradient descent. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 2595–2603. Curran Associates, Inc., 2010. URL http://papers.nips.cc/paper/4006-parallelized-stochastic-gradient-descent.pdf.

[21] Janis Keuper. Distributed training of deep neuronal networks: Theoretical and practical limits of parallel scalability. *CoRR*, abs/1609.06870, 2016. URL http://arxiv.org/abs/1609.06870.

Archived copy · cryp2nova.com

[22] Shuxin Zheng, Qi Meng, Taifeng Wang, Wei Chen, Nenghai Yu, Zhiming Ma, and Tie-Yan Liu. Asynchronous stochastic gradient descent with delay compensation for distributed deep learning. *CoRR*, abs/1609.08326, 2016. URL http://arxiv.org/abs/1609.08326.

[23] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *International Conference on Neural Information Processing Systems*, pages 315–323, 2013.

[24] Shen Yi Zhao, Ru Xiang, Ying Hao Shi, Peng Gao, and Wu Jun Li. Scope: Scalable composite optimization for learning on spark. 2016.

[25] Zhouyuan Huo and Heng Huang. Asynchronous stochastic gradient descent with variance reduction for non-convex optimization. 2016.

[26] Sixin Zhang, Anna Choromanska, and Yann LeCun. Deep learning with elastic averaging SGD. *CoRR*, abs/1412.6651, 2014. URL http://arxiv.org/abs/1412.6651.

[27] Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 1-bit stochastic gradient descent and application to data-parallel distributed training of speech dnns. In *Interspeech 2014*, September 2014.

[28] Kai Chen and Qiang Huo. Scalable training of deep learning machines by incremental block training with intra-block parallel optimization and blockwise model-update filtering. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 5880–5884, 2016.

[29] Pengtao Xie, Jin Kyu Kim, and Eric P. Xing. Large scale distributed multiclass logistic regression. *CoRR*, abs/1409.5705, 2014. URL http://arxiv.org/abs/1409.5705.

[30] Hao Zhang, Zeyu Zheng, Shizhen Xu, Wei Dai, Qirong Ho, Xiaodan Liang, Zhiting Hu, Jinliang Wei, Pengtao Xie, and Eric P. Xing. Poseidon: An efficient communication architecture for distributed deep learning on GPU clusters. *CoRR*, abs/1706.03292, 2017. URL http://arxiv.org/abs/1706.03292.

[31] Andrew Gibiansky. Bringing hpc techniques to deep learning.

[32] Horovod. https://github.com/uber/horovod, 2017.

[33] Minsik Cho, Ulrich Finkler, Sameer Kumar, David S. Kung, Vaibhav Saxena, and Dheeraj Sreedhar. Powerai DDL. *CoRR*, abs/1708.02188, 2017. URL http://arxiv.org/abs/1708.02188.

[34] Tim Dettmers. 8-bit approximations for parallelism in deep learning. *CoRR*, abs/1511.04561, 2015. URL http://arxiv.org/abs/1511.04561.

[35] Anonymous authors. Deep gradient compression-reducing the communication bandwidth for distributed training. *ICLR*, 2018.

Archived copy · cryp2nova.com

Risk Factors

RISKS RELATING TO PARTICIPATION IN THE TOKEN SALE

Investments in startups such as the Token Issuer and the Project Company involve a high degree of risk.

Financial and operating risks confronting startups are significant and Token Issuer and the Project Company are not immune to these. Startups often experience unexpected problems in the areas of product development, marketing, financing, and general management, among others, which frequently cannot be solved.

The Token Issuer and/or the Project Company may be forced to cease operations.

It is possible that, due to any number of reasons, including, but not limited to, an unfavorable fluctuation in the value of cryptographic and fiat currencies, the inability by the Token Issuer and/or the Project Company to establish the Cortex platform or the Cortex Coins' utility, the failure of commercial relationships, or intellectual property ownership challenges, the Token Issuer and/or the Project Company may no longer be viable to operate and the Token Issuer and/or the Project Company may dissolve or take actions that result in a dissolution of the Token Issuer and/or the Project Company.

The tax treatment of the Token Sale Terms, the purchase rights contained therein and the Token Sale is uncertain and there may be adverse tax consequences for purchasers upon certain future events.

The tax characterization of the Token Sale Terms and the Cortex Coins is uncertain, and each purchaser must seek its own tax advice in connection with an investment in the Cortex Coins. An investment pursuant to the Token Sale Terms and the purchase of Cortex Coins pursuant thereto may result in adverse tax consequences to the purchaser, including withholding taxes, income taxes and tax reporting requirements. Each purchaser should consult with and must rely upon the advice of its own professional tax advisors with respect tax treatment of an investment in the Cortex Coins pursuant to the Token Sale Terms.

There is no prior market for Cortex Coins and the Token Sale may not result in an active or liquid market for the Tokens

Prior to the Token Sale, there has been no public market for the Cortex Coins. In the event that the Cortex Coins are traded on a cryptocurrency exchange, there is no assurance that an active or liquid trading market for the Cortex Coins will develop or if developed, be sustained after the Cortex Coins have been made available for trading on such cryptocurrency exchange. There is also no assurance that the market price of the Cortex Coins will not decline below the consideration at which the purchaser acquired the Cortex Coins at. Such purchase consideration may not be indicative of the market price of the Cortex Coins after they have been made available for trading on a cryptocurrency exchange. A Cortex Coin is not a currency issued by any central bank or national, supra-national or quasi-national organisation, nor is it backed by any hard assets or other credit. The Token Issuer is not responsible for nor does it pursue the circulation and trading of Cortex Coins on the market. Trading of Cortex Coins merely depends on the consensus on its value between the relevant market participants, and no one is obliged to purchase any Cortex Coin from any holder of the Cortex Coin, nor does anyone guarantee the liquidity or market price of Cortex Coins to any extent at any time. Accordingly, the Token Issuer cannot ensure that there will be any demand or market for Cortex Coins, or that the purchase consideration is indicative of the market price of Cortex Coins after they have been made available for trading on a cryptocurrency exchange.

Future sales of the Cortex Coins could materially and adversely affect the market price of Cortex Coins

Any future sale of the Cortex Coins (which were not available for sale in the Token Sale) would increase the supply of Cortex Coins in the market and this may result in a downward price pressure on the Cortex Coin. The sale or distribution of a significant number of Cortex Coins outside of the

Archived copy · cryp2nova.com

Token Sale, or the perception that such further sales or issuance may occur, could adversely affect the trading price of the Cortex Coins.

Negative publicity may materially and adversely affect the price of the Cortex Coins

Negative publicity involving (a) the Token Issuer and/or the Project Company; (b) the Cortex platform; (c) the Cortex Coins; or (d) any of the key personnel of the Token Issuer and/or the Project Company, may materially and adversely affect the market perception or market price of the Cortex Coins, whether or not such publicity is justified.

There is no assurance of any success of Cortex platform

The value of, and demand for, the Cortex Coins hinges heavily on the performance of the Cortex platform. There is no assurance that the Cortex platform will gain traction after its launch and achieve any commercial success.

The Cortex platform has not been fully developed, finalised and integrated and is subject to further changes, updates and adjustments prior to its launch. Such changes may result in unexpected and unforeseen effects on its projected appeal to users, and hence impact its success.

While the Token Issuer has made every effort to provide a realistic estimate, there is also no assurance that the cryptocurrencies raised in the Token Sale will be sufficient for the development and integration of the Cortex platform. For the foregoing or any other reason, the development and integration of the Cortex platform may not be completed and there is no assurance that it will be launched at all. As such, distributed Cortex Coins may hold little worth or value, and this would impact its trading price.

The trading price of the Cortex Coins may fluctuate following the Token Sale

The prices of cryptographic tokens in general tend to be relatively volatile, and can fluctuate significantly over short periods of time. The demand for, and corresponding the market price of, the Cortex Coins may fluctuate significantly and rapidly in response to, among others, the following factors, some of which are beyond the control of the Token Issuer and/or the Project Company:

(a) new technical innovations;

(b) analysts' speculations, recommendations, perceptions or estimates of the Cortex Coin's market price or the Token Issuer's and/or the Project Company's financial and business performance;

(c) changes in market valuations and token prices of entities with operations similar to that of the Token Issuer and/or the Project Company that may be made available for sale and purchase on the same cryptocurrency exchanges as the Cortex Coins;

(d) announcements by the Token Issuer and/or the Project Company of significant events, for example partnerships, sponsorships, new product developments;

(d) fluctuations in market prices and trading volume of cryptocurrencies on cryptocurrency exchanges;

(e) additions or departures of key personnel of the Token Issuer and/or the Project Company;

(f) success or failure of the management of the Token Issuer and/or the Project Company in implementing business and growth strategies; and

(g) changes in conditions affecting the blockchain or financial technology industry, the general economic conditions or market sentiments, or other events or factors.

RISKS RELATING TO THE WALLET

The loss or compromise of information relating to your Wallet (as defined below) may affect your access and possession of the Cortex Coins

For purposes of receipt of your Cortex Coins, you are to establish and maintain access to a cryptocurrency wallet (Wallet). Your access to the Cortex Coins in the Wallet depends on, among other things, the safeguards to the information to such Wallet, including but not limited to the user account information, address, private key and password. In the event that any of the foregoing is lost

Archived copy · cryp2nova.com

or compromised, your access to the Wallet may be curtailed and thereby adversely affecting your access and possession to the Cortex Coins, including such Cortex Coins being unrecoverable and permanently lost.

The Wallet or Wallet service provider may not be technically compatible with the Cortex Coins

The Wallet or Wallet service provider may not be technically compatible with the Cortex Coins which may result in the delivery of Cortex Coins being unsuccessful or affect your access to such Cortex Coins.

RISKS RELATING TO THE TOKEN ISSUER AND THE PROJECT COMPANY

The Cortex platform is intended to be developed, operated and maintained by the Token issuer and/or the Project Company. Any events or circumstances which adversely affect the Token Issuer and/or the Project Company may have a corresponding adverse effect on the Token Issuer and/or the Project Company if such events or circumstances affect the Token Issuer's and/or the Project Company's ability to maintain the Cortex platform. This would correspondingly have an impact the trading price of the Cortex Coins.

The Token Issuer and/or the Project Company may be materially and adversely affected if they fail to effectively manage its operations as their business develops and evolves, which would have a direct impact on their ability to maintain the Cortex platform and consequently the trading price of the Cortex Coins.

The financial technology and cryptocurrency industries and the markets in which the Token Issuer and the Project Company compete have grown rapidly and continue to grow rapidly, and continue to evolve in response to new technological advances, changing business models and other factors. As a result of this constantly changing environment, the Token Issuer and/or the Project Company may face operational difficulties in adjusting to the changes, and the sustainability of the Token Issuer and the Project Company will depend on their ability to manage their respective operations, adapt to technological advances and market trends and ensure that they hire qualified and competent employees, and provide proper training for their personnel.

As their respective business evolves, the Token Issuer and the Project Company must also expand and adapt its operational infrastructure. The Token Issuer's and the Project Company's respective businesses rely on blockchain-based software systems, cryptocurrency wallets or other related token storage mechanisms, blockchain technology and smart contract technology, and to manage technical support infrastructure for the Cortex platform effectively, the Token Issuer and the Project Company will need to continue to upgrade and improve their data systems and other operational systems, procedures and controls. These upgrades and improvements will require a dedication of resources, are likely to be complex and increasingly rely on hosted computer services from third parties that the Token Issuer and/or the Project Company do not control. If the Token Issuer and/or the Project Company are unable to adapt their respective systems and organisation in a timely, efficient and cost-effective manner to accommodate changing circumstances, its business, financial condition and results of operations may be adversely affected. If the third parties whom the Token Issuer and/or the Project Company rely on are subject to a security breach or otherwise suffer disruptions that impact the respective services the Token Issuer and/or the Project Company utilise, the integrity and availability of their respective internal information could be compromised, which may consequently cause the loss of confidential or proprietary information, and economic loss. The loss of financial, labour or other resources, and any other adverse effect on the Token Issuer's and/or the Project Company's respective business, financial condition and operations, would have a direct adverse effect on the Token Issuer's and the Project Company's ability to maintain the Cortex platform. As the Cortex platform is the main product to which the Cortex Coins relate, this may adversely impact the trading price of the Cortex Coins.

Archived copy · cryp2nova.com

The Token Issuer and/or the Project Company may experience system failures, unplanned interruptions in its network or services, hardware or software defects, security breaches or other causes that could adversely affect the Token Issuer's and/or the Project Company's infrastructure network, and/or the Cortex platform

The Token Issuer and the Project Company are unable to anticipate when there would be occurrences of hacks, cyber-attacks, mining attacks (including but not limited to double-spend attacks, majority mining power attacks and selfish-mining attacks), distributed denials of service or errors, vulnerabilities or defects in the Cortex platform, the Cortex Coins, the Wallet or any technology (including but not limited to smart contract technology) on which the Token Issuer and/or the Project Company, the Cortex platform, the Cortex Coins and the Wallet relies or on the Ethereum blockchain or any other blockchain. Such events may include, for example, flaws in programming or source code leading to exploitation or abuse thereof. The Token Issuer and/or the Project Company may not be able to detect such hacks, mining attacks (including but not limited to double-spend attacks, majority mining power attacks and selfish-mining attacks), cyber-attacks, distributed denials of service errors vulnerabilities or defects in a timely manner, and may not have sufficient resources to efficiently cope with multiple service incidents happening simultaneously or in rapid succession.

The Token Issuer's and/or the Project Company's respective network or services, which would include the Cortex platform, could be disrupted by numerous events, including natural disasters, equipment breakdown, network connectivity downtime, power losses, or even intentional disruptions of their respective services, such as disruptions caused by software viruses or attacks by unauthorised users, some of which are beyond the Token Issuer's and/or the Project Company's control. Although the Token Issuer and the Project Company will be taking steps against malicious attacks on their respective appliances or infrastructure, which are critical for the maintenance of the Cortex platform and their respective other services, there can be no assurance that cyber-attacks, such as distributed denials of service, will not be attempted in the future, and that any of the Token Issuer's and the Project Company's intended enhanced security measures will be effective. The Token Issuer and the Project Company may also be prone to attacks on their respective infrastructure intended to steal information about their respective technology, financial data or user information or take other actions that would be damaging to the Token Issuer, the Project Company and users of the Cortex platform. Any significant breach of the Token Issuer's and/or the Project Company's intended security measures or other disruptions resulting in a compromise of the usability, stability and security of the Token Issuer's and/or the Project Company's network or services (including the Cortex platform) may adversely affect the trading price of the Cortex Coins.

The Token Issuer and the Project Company are dependent in part on the location and data centre facilities of third parties

The Token Issuer's and the Project Company's infrastructure network will be in part established through servers that which they respectively own and house at the location facilities of third parties, and servers that they respectively rent at data centre facilities of third parties. If the Token Issuer and/or the Project Company are unable to renew their respective data facility lease on commercially reasonable terms or at all, the Token Issuer and/or the Project Company may be required to transfer their respective servers to a new data centre facility, and may incur significant costs and possible service interruption in connection with the relocation. These facilities are also vulnerable to damage or interruption from, among others, natural disasters, arson, terrorist attacks, power losses, and telecommunication failures. Additionally, the third party providers of such facilities may suffer a breach of security as a result of third party action, employee error, malfeasance or otherwise and a third party may obtain unauthorised access to the data in such servers. As techniques used to obtain unauthorised access to, or to sabotage systems change frequently and generally are not recognised until launched against a target, the Token Issuer, the Project Company and the providers of such facilities may be unable to anticipate these techniques or to implement adequate preventive measures. Any such security breaches or damages which occur which impact upon the Token Issuer's and/or the Project Company's infrastructure network and/or the Cortex platform may adversely impact the price of the Cortex Coins.

Archived copy · cryp2nova.com

General global market and economic conditions may have an adverse impact on the Token Issuer's and/or the Project Company's operating performance, results of operations and cash flows

The Token Issuer and/or the Project Company could be affected by general global economic and market conditions. Challenging economic conditions worldwide have from time to time, contributed, and may continue to contribute, to slowdowns in the information technology industry at large. Weakness in the economy could have a negative effect on the Token Issuer's and/or the Project Company's respective business, operations and financial condition, including decreases in revenue and operating cash flows. Additionally, in a down-cycle economic environment, the Token Issuer and/or the Project Company may experience the negative effects of increased competitive pricing pressure and a slowdown in commerce and usage of the Cortex platform. Suppliers on which the Token Issuer and/or the Project Company rely for servers, bandwidth, location and other services could also be negatively impacted by economic conditions that, in turn, could have a negative impact on the Token Issuer's and/or the Project Company's respective operations or expenses. There can be no assurance, therefore, that current economic conditions or worsening economic conditions or a prolonged or recurring recession will not have a significant adverse impact on the Token Issuer's and/or the Project Company's respective business, financial condition and results of operations and hence the Cortex platform, which would correspondingly impact the trading price of the Cortex Coins.

The Token Issuer, the Project Company and/or the Cortex Coins may be affected by newly implemented regulations

Cryptocurrency trading is generally unregulated worldwide, but numerous regulatory authorities across jurisdictions have been outspoken about considering the implementation of regulatory regimes which govern cryptocurrency or cryptocurrency markets. The Token Issuer, the Project Company and/or the Cortex Coins may be affected by newly implemented regulations relating to cryptocurrencies or cryptocurrency markets, including having to take measures to comply with such regulations, or having to deal with queries, notices, requests or enforcement actions by regulatory authorities, which may come at a substantial cost and may also require substantial modifications to the Cortex Coins and/or the Cortex platform. This may impact the appeal of the Cortex Coins and/or the Cortex platform for users and result in decreased usage of the Cortex Coins and/or the Cortex platform. Further, should the costs (financial or otherwise) of complying with such newly implemented regulations exceed a certain threshold, maintaining the Cortex Coins and/or the Cortex platform may no longer be commercially viable and the Token Issuer and/or the Project Company may opt to discontinue the Cortex Coins and/or the Cortex platform.

Further, it is difficult to predict how or whether governments or regulatory authorities may implement any changes to laws and regulations affecting distributed ledger technology and its applications, including the Cortex Coins and the Cortex platform. The Token Issuer and/or the Project Company may also have to cease their respective operations in a jurisdiction that makes it illegal to operate in such jurisdiction, or make it commercially unviable or undesirable to obtain the necessary regulatory approval(s) to operate in such jurisdiction. In scenarios such as the foregoing, the trading price of Cortex Coins will be adversely affected or Cortex Coins may cease to be traded.

The regulatory regime governing the blockchain technologies, cryptocurrencies, tokens and token offerings such as Token Sale, the Cortex platform and the Cortex Coins is uncertain, and regulations or policies may materially adversely affect the development of the Cortex platform and the utility of the Cortex Coins

Regulation of tokens (including the Cortex Coins) and token offerings such as the Token Sale, cryptocurrencies, blockchain technologies, and cryptocurrency exchanges currently is undeveloped and likely to rapidly evolve, varies significantly among international, federal, state and local jurisdictions and is subject to significant uncertainty. Various legislative and executive bodies in Singapore and other countries may in the future, adopt laws, regulations, guidance, or other actions, which may severely impact the development and growth of the Cortex platform and the adoption and utility of the Cortex Coins. Failure by the Token Issuer, the Project Company or users of the Cortex platform to comply with any laws, rules and regulations, some of which may not exist yet or are subject to interpretation and may be subject to change, could result in a variety of adverse consequences, including civil penalties and fines. Blockchain networks also face an uncertain regulatory landscape

Archived copy · cryp2nova.com

in many foreign jurisdictions such as the European Union, China, South Korea and Russia. Various foreign jurisdictions may, in the near future, adopt laws, regulations or directives that affect the Cortex platform. Such laws, regulations or directives may directly and negatively impact the Token Issuer's and/or the Project Company's respective business. The effect of any future regulatory change is impossible to predict, but such change could be substantial and materially adverse to the development and growth of the Cortex platform and the adoption and utility of the Tokens.

New or changing laws and regulations or interpretations of existing laws and regulations may materially and adversely impact the value of the currency in which the Cortex Coins may be sold, the value of the distributions that may be made by the Token Issuer and/or the Project Company, the liquidity of the Cortex Coins, the ability to access marketplaces or exchanges on which to trade the Cortex Coins, and the structure, rights and transferability of Cortex Coins.

Cortex Coin holders will have no control on the Token Issuer or the Project Company

The holders of Cortex Coins are not and will not be entitled, to vote or receive dividends or be deemed the holder of capital stock of the Token Issuer or the Project Company for any purpose, nor will anything be construed to confer on the purchasers any of the rights of a stockholder of the Token Issuer or the Project Company or any right to vote for the election of directors or upon any matter submitted to stockholders at any meeting thereof, or to give or withhold consent to any corporate action or to receive notice of meetings, or to receive subscription rights or otherwise. Purchasers may lack information for monitoring their investment The purchasers of Cortex Coins may not be able to obtain all information it would want regarding the Token Issuer, the Project Company, the Cortex Coins, or the Cortex platform, on a timely basis or at all. It is possible that purchasers may not be aware on a timely basis of material adverse changes that have occurred. While the Token Issuer has made efforts to use open-source development for Cortex Coins, this information may be highly technical by nature. As a result of these difficulties, as well as other uncertainties, Purchasers may not have accurate or accessible information about the Cortex platform.

There may be unanticipated risks arising from the Cortex Coins

Cryptographic tokens such as the Cortex Coins are a relatively new and dynamic technology. In addition to the risks included in this section, there are other risks associated with the purchase, holding and use of the Cortex Coins, including those that the Token Issuer and the Project Company cannot anticipate. Such risks may further materialise as unanticipated variations or combinations of the risks discussed in this

CRYP2NOVA

Archived copy · cryp2nova.com