

---

## **awesome-zkml**

A place where you can find content, codebases, scientific papers, projects and applications related to ZKML. We always appreciate contributions and suggestions. If you have any changes you'd like to suggest, please feel free to open an issue or submit a pull request.

### **Learn ZK**

Here you can find a list of very popular resources to get started with zero-knowledge cryptography (courtesy of Ingopedia made by Ingonyama):

- Ingopedia
- zkProof Standards - Resource
- ZK Mesh - resource
- Curated list of ZKP implementations
- Awesome - Matter labs - ZK proofs
- Awesome - Mikerah - Privacy on Blockchains
- Resource: Awesome\_Plonk
- ZK research 0x
- ZK canon
- Proofs, Args and ZK - Justin Thaler

### **Learn ML**

Here you can find a list of very popular resources to get started with machine learning.

- awesome-machine-learning
  - books
  - courses
  - content
  - events
  - meetups

---

## Content

### ZKML community calls

We are organizing bi-weekly ZKML community calls. For details on when and where the next call is happening, check the pinned messages on our Telegram channel. Recorded versions of the calls are available on our YouTube channel and notes summarizing what was discussed can be found [here](#).

- ZKML community call #0

### Articles and podcasts

- Zero Knowledge Machine Learning - Remco Bloemen
- Zero-Knowledge Proofs and Their Applications to Machine Learning (video)
- ZK Machine Learning
- ZK for ML
- Zero Knowledge Podcast: Episode 246: Adversarial Machine Learning Research with Florian Tramèr
- Zero-Knowledge Machine Learning by Jason Morton (video)
- Modulus Labs (Twitter)
  - Chapter 1: How to Put Your AI On-Chain
  - Chapter 2: Why Put Your AI On-Chain?
  - Chapter 3: The World's First On-Chain AI Trading Bot
  - Chapter 4: Blockchains that Self-Improve
  - Chapter 5: The Cost of Intelligence
- Trustless Verification of Machine Learning (Daniel Kang, Tatsunori Hashimoto, Ion Stoica, Yi Sun)
- ZK Podcast - episode 265: Where ZK and ML intersect with Yi Sun and Daniel Kang
- Linear A Research
- An introduction to zero-knowledge machine learning - Worldcoin
- Zero Gravity (The Weight is Over) - ZKHack Lisbon
- Zero-Knowledge Decision Tree Prediction (ZK-DTP) - ZKHack Lisbon
- Open-sourcing zkml: Trustless Machine Learning for All - Daniel Kang
- Checks and balances: Machine learning and zero-knowledge proofs - Elena Burger @ a16z
- ZKML: Bridging AI/ML and Web3 with Zero-Knowledge Proofs - Cathie So
- Do language models possess knowledge (soundness)? - Tarun Chitra
- Balancing the Power of AI/ML: The Role of ZK and Blockchain - SevenX Ventures
- The Ultimate Guide to the ZKML ecosystem (Twitter thread) - SevenX ventures
- Verified Execution of GPT, Bert, CLIP, and more - Daniel Kang

- 
- zkML: Evolving the Intelligence of Smart Contracts Through Zero-Knowledge Cryptography - 1kx
  - Dcbuilder - Zero-Knowledge Machine Learning and its use cases (Jul 2023)
  - TensorPlonk: A “GPU” for ZKML, Delivering 1,000x Speedups - Daniel Kang
  - ZK10: ZKML with EZKL: Where we are and the future - Jason Morton (ZKonduit)
  - ZK10: ZKML Endgame: Specialized ZK Proving with GKR - Ryan Cao (Modulus)
  - ezkl blog
  - Modulus blog
  - Giza blog
  - The promise and challenges of crypto + AI applications - Vitalik Buterin

## Codebases

- zk-mnist - @hopanml @sunfishstanford @henripal (2022)
- zk-ml/demo - @liaopeiyuan (2021)
- circomlib-ml - @socathie (2022)
  - Gitcoin Grant Proposal
- proto-neural-zkp - @recmo @dcbuid3r (2022)
  - zkml experiments with plonky2 at Worldcoin
- RockyBot - @ModulusLabs (2022)
  - RockyBot is the first ever fully on-chain AI trading bot
- ezkl by Jason Morton (2022+)
  - ezkl is a library and command-line tool for doing inference for deep learning models and other computational graphs in a zk-snark.
- keras2circom (@socathie) (2023)
  - keras2circom is a python tool that transpiles a tf.keras model into a circom circuit.
- Zator - Verified inference of a 512-layer neural network using recursive SNARKs.
- Otti (2022)
  - Compiler and zkSNARK for optimization problems including LP, SDP, SGD. Includes ZK proof of the full training of a perceptron on real-world datasets.
- Linear A - tachikoma (2022+)
- Linear A - uchikoma (2022+)
- zk-dtp - Zero Knowledge Decision Tree Predict is designed to address this pressing issue by offering privacy-preserving predictions using decision tree models, built on top of RISC Zero’s zkVM.

- 
- zkp-gravity/0g - ZeroGravity - Zero Gravity is a system for proving an inference run (i.e. a classification) for a pre-trained, public Weightless NN and a private input. (2023)
  - ddkang/zkml - zkml is a framework for constructing proofs of ML model execution in ZK-SNARKs.
  - ZKaggle and ZKaggleV2 - @socathie (2023)
    - POC of a decentralized bounty platform for hosting, verifying, and paying out bounties, similar to Kaggle, but with the added benefit of privacy preservation

## Papers

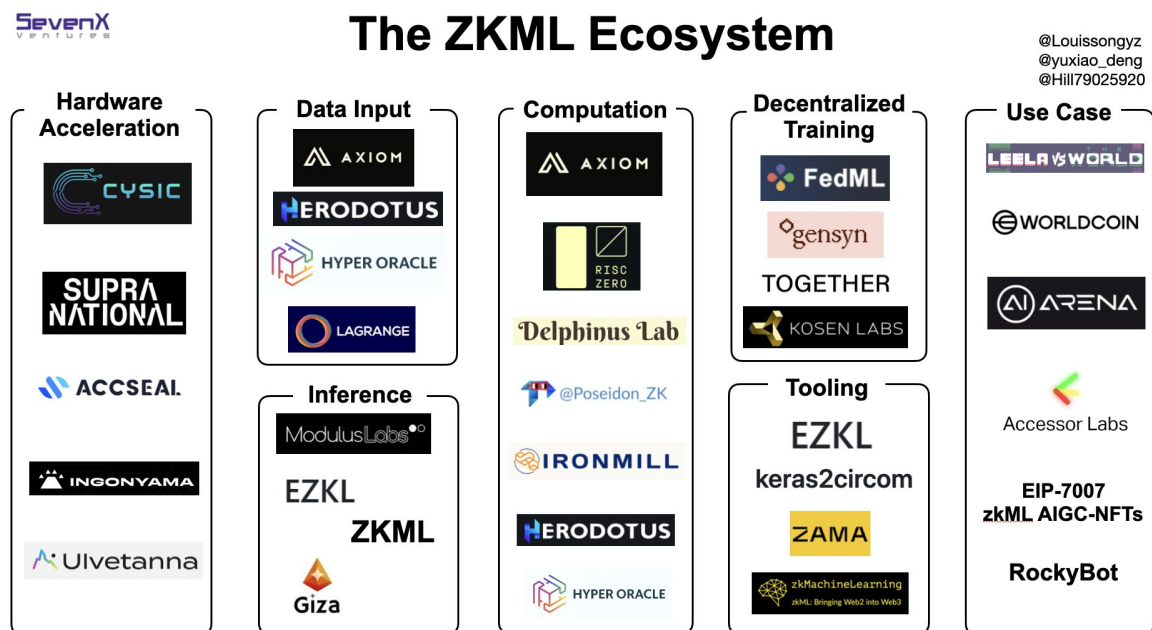
- Justin Thaler (2013). “Time-Optimal Interactive Proofs for Circuit Evaluation”
- Pengtao Xie, Misha Bilenko, Tom Finley, Ran Gilad-Bachrach, Kristin Lauter, Michael Naehrig (2014). “Crypto-Nets: Neural Networks over Encrypted Data”
- Nathan Dowlin, Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, John Wernsing (2016). “CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy”
- Zahra Ghodsi, Tianyu Gu, Siddharth Garg (2017). “SafetyNets: Verifiable Execution of Deep Neural Networks on an Untrusted Cloud”
- Payman Mohassel and Yupeng Zhang (2017). “SecureML: A System for Scalable Privacy-Preserving Machine Learning”
- Jian Liu, Mika Juuti, Yao Lu, and N. Asokan (2017). “Oblivious Neural Network Predictions via MiniONN transformations”
- Seunghwa Lee, Hankyung Ko, Jihye Kim, and Hyunok Oh (2020). “vCNN: Verifiable Convolutional Neural Network based on zk-SNARKs”
- Ramy E. Ali, Jinhyun So, A. Salman Avestimehr (2020). “On Polynomial Approximations for Privacy-Preserving and Verifiable ReLU Networks”
- Boyuan Feng, Lianke Qin, Zhenfei Zhang, Yufei Ding, and Shumo Chu (2021). “ZEN: An Optimizing Compiler for Verifiable, Zero-Knowledge Neural Network Inferences”
- Tianyi Liu, Xiang Xie, and Yupeng Zhang (2021). “zkCNN: Zero Knowledge Proofs for Convolutional Neural Network Predictions and Accuracy”
- Chenkai Weng, Kang Yang, Xiang Xie, Jonathan Katz, and Xiao Wang (2021). “Mystique: Efficient Conversions for Zero-Knowledge Proofs with Applications to Machine Learning” (slides)
- Jiasi Weng, Jian Weng, Member, IEEE, Gui Tang, Anjia Yang, Ming Li, Jia-Nan Liu (2022). pvCNN: Privacy-Preserving and Verifiable Convolutional Neural Network Testing
- Sebastian Angel, Andrew J. Blumberg, Eleftherios Ioannidis, Jess woods (2022). Efficient Representation of Numerical Optimization Problems for SNARKs
- Daniel Kang, Tatsunori Hashimoto, Ion Stoica, Yi Sun (2022). Scaling up Trustless DNN Inference

---

with Zero-Knowledge Proofs

- Haodi Wang, Thang Hoang (2022). ezDPS: An Efficient and Zero-Knowledge Machine Learning Inference Pipeline
- Modulus Labs - The Cost of Intelligence: Proving Machine Learning Inference with Zero-Knowledge

## Projects interested in ZKML



- Axiom - Axiom provides smart contracts trustless access to all on-chain data and arbitrary expressive compute over it. Like GPUs do for CPUs, Axiom augments blockchain consensus with zero-knowledge proofs
- 0xPARC - The 0xPARC Foundation promotes application-level innovation on Ethereum and other decentralized platforms
  - zkMnist
- Worldcoin - A Privacy-Preserving Proof-of-Personhood Protocol
  - proto-neural-zkp
- Gizatech - Fully on-chain artificial intelligence on Starknet
- Modulus Labs - Bringing powerful ML models on-chain
- Risc Zero - The General Purpose Zero-Knowledge VM

- 
- Supranational - A product and service company developing hardware-accelerated cryptography for verifiable and confidential computing
  - Ingonyama (Hardware) - Zero Knowledge ASICs (ZPU)
  - Zama.ai (FHE ML / FHE-ZK ML) - FHE tooling for machine learning, blockchain and more. ZK-FHE is an interesting research area. FHE.org is a very interesting community with a lot of potential for collaboration.
  - zkMachineLearning - ZKML tooling for Circom
  - Aleo - Platform for building fully private and programmable Web applications.
  - PSE team @ Ethereum Foundation - Some ZKML research initiatives here
  - Ion Protocol - Lending protocol for staked & restaked assets. They partnered with Modulus to build a risk engine that analyzes validator credit risk. Read more here

## Use cases

Decision tree for a use case that would use ZKML -> Intersection of {needs privacy, computational integrity} and {heuristic optimization problem solved by ml}.

- Computational integrity
  - Modulus Labs
    - \* On-chain verifiable ML trading bot - RockyBot
    - \* Blockchains that self-improve vision (examples):
      - Enhancing the Lyra finance options protocol AMM with intelligent features
      - Creating a transparent AI-based reputation system for Astraly
      - Working on the technical breakthroughs needed for contract-level compliance tools using ML for Aztec Protocol (a zk-rollup with privacy features)
  - ML as a Service (MLaaS) transparency ([link](#))
  - Worldcoin
    - \* Verifying that a user has created a valid and unique WorldID locally by running the IrisCode model on self-hosted biometric data and is calling `_addMember(uint256 groupId, uint256 identityCommitment)` function on the WorldID Semaphore identity group with a valid identityCommitment. -> Makes protocol more permissionless
    - \* Making the Orb trustless, provide proof that fraud filters are applied
    - \* Enable IrisCode upgradeability
  - ZK anomaly/fraud detection
    - \* Creates the ability for creating a ZK proof of exploitability/fraud. Anomaly detection models could be trained on smart contract data and agreed upon by DAOs as interesting metrics to be able to automate security procedures such as preventively pausing

---

contracts in a more proactive way. There are startups already looking at using ML models for security purposes in a smart contract context, so ZK anomaly detection proofs feel like the natural next step.

- Generic SNARK for ML inference: ability to easily prove and verify that an output is the product of a given model and input pair.
- Privacy
  - Decentralized Kaggle: proof that model has greater than x% accuracy on some test data without revealing weights
  - Privacy-preserving inference: medical diagnostics on private patient data get fed into the model and the sensitive inference (i.e. cancer test result) gets sent to the patient. (vCNN paper, page 2/16)