International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Privacy-Preserving Digital Phenotyping: A Zero-Knowledge Framework for Early Autism Spectrum Disorder Detection

Sneha Giri

Student, Information Technology, Institute of Engineering & Management

Abstract

Autism Spectrum Disorder (ASD) is a complex neurodevelopmental condition in which early intervention is the key factor of the long-term function. However, with older-style efforts in clinical diagnosis, the channels are often subject to delay due to subjective diagnostic assessments, and limited availability of specialists. While Machine Learning (ML) provides a novel approach via digital phenotyping, by leveraging telemetry from smartphones such as keystrokes dynamics, sensor patterns, and frequency of social interaction etc., and the gathering of such granular data from individuals, an ethical and privacy issue carcinogenic. This paper proposes novel, privacy focussed objects, called Zero-Knowledge Proofs (ZKP), where Digital Phenotyping works alongside ZKP to enable secure ASD screening. Unlike current models, which involve centralisation of sensitive data pertaining to behavioural classification the proposed model uses on-device machine learning (ML) for behavioural classification, followed by generation of a ZKP to validate the diagnostic result without revealing underlying raw telemetry. By filling the information gap between high frequency behavioural monitoring and cryptographic data privacy, this research work is a scalable, objective, and "privacy-by-design" solution for early screening of ASD. Preliminary analysis indicates that using this dual layer architecture will not only ensure diagnostics sensitivity at the highest level possible but also overcomes the "Privacy Paradox" (a reluctance to share data for privacy reasons) by promoting broader adoption of digital health monitoring among vulnerable populations.

Keywords: Autism Spectrum Disorder (ASD), Machine Learning, Zero-Knowledge Proofs (ZKP), Privacy-Preserving Artificial Intelligence (AI), Keystroke Dynamics, Behavioural Informatics.

1. Introduction

Autism Spectrum Disorder (ASD) is a form of heterogeneous condition in neurodevelopment that is characterized by persistent difficulties in social communication and restricted interests and repetitive behavioural patterns. As of the year of 2026, worldwide, the prevalence rate could be around 1 in 100 children, with some areas as high in 1 in 31 children. The clinical significance of ASD is in the "critical window" of brain plasticity; early intervention before the age of 36 months is scientifically proven to drastically improve long term functional outcome and cognitive adaptability [1]. However, the onrush to diagnosis is a critical drawback to global healthcare systems.

1.1 How the Diagnosis of Evolution

Historically, the original tools for diagnosing ASD diagnosis have been standardised observational tools

such as the Autism Diagnostic Observation Schedule (ADOS) and the Autism Diagnostic Interview-Revised (ADI-R) [1]. While these are still the "gold standard" they are limited in their own way because they are subjective, very expensive and require highly trained specialists. In many developing regions the waiting time for evaluation by a specialist may be more than 18 months or more [1], thereby missing the most vital time for early intervention.

To overcome these problems, the research community has moved towards Machine Learning (ML) and Artificial Intelligence (AI). Modern screening technology powered by AI includes the use of computer vision to monitor the gaze pattern, Natural Language Processing (NLP) to monitor speech prosody, and deep learning to detect subtle motor stereotypies. These computational approaches hold out the promise of objectivity, scalability and low-cost screening that can be implemented in primary care settings or by telehealth.

1.2 The Privacy Paradox and the Emergence of Digital Phenotyping

Despite the success of screening based on ML, a major downside has surfaced which is "Privacy Paradox." To reach clinically high levels of diagnostic performance, access to granular, high frequency behavioural data (e.g. via video recordings of children or private social interactions) is required for AI models. This poses a huge security risk and an obstacle to adoption because parents and guardians are rightly wary of centralizing sensitive behavioural data in a cloud-based repository because of the potential for a data breach or some form of social stigmatization.

A solution to this problem is proposed through Digital Phenotyping, as suggested by this research. Digital phenotyping is simply the quantification of the human phenotype defined at the individual level, moment by moment in situ with the help of data from the individual personal digital devices. Instead of active and stressful clinical testing, for example, digital phenotyping captures "digital signatures" of telemetry data from a smartphone in a passive way (e.g. keystroke dynamics, typing rhythm, accelerometer-based motor patterns, usage frequency of apps). These markers are a continuous, real-world reflection of how a user's neuro-behavioural state is without modality requiring the direct clinical observation of the individual.

1.3 Cryptography Sovereignty: Zero-Knowledge Proofs (ZKP)

To address the contradiction between data utility and data privacy, this paper proposes the combination of Zero-Knowledge Proof (ZKP). ZKP is actually a cryptographic primitive that enables me - one party - to prove something is true to someone else - another party, the Verifier - "The user meets the digital phenotyping criteria for ASD", without having to explain any of the underlying data that were used to come to that conclusion.

By making use of zk-SNARKs (Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge) we move the trust model from the corporate data centres onto mathematical certitude. The raw behavioural data is never transmitted off the user's local device, only the "proof of result" is transmitted to the clinician [2].

1.4 Framework and Contribution SOC Proposed

The main goal of this project is to build on Privacy-Preserving Digital Phenotyping Framework for early screening of ASD. The fundamental contributions, or contributions of this work, are three-fold:

Development of a lightweight, on-device ML learning machine that can detect ASD-related markers of behavioural markers through passive smartphone telemetry.

Implementation of a ZKP layer which verifies the veracity of the local ML inference ensuring that the diagnostic results are verifiable, but private.

Validation of the system's efficiency in terms of proof generation time and proof size to ensure the viability

of a system for standard hardware on mobile devices.

By combining both the behavioural knowledge gained by digital phenotyping and the cryptographic security provided by ZKP, the goal of this research is to offer "Privacy-by-Design" screening tool that has the potential to democratize the access to early ASD detection while ensuring absolute user data sovereignty.

Keywords: Autism Spectrum Disorder (ASD), Digital Phenotyping, Zero-Knowledge Proofs, zk-SNARKs, Privacy-Preserving AI, Motor Stereotypies, Keystroke Dynamics.

2. Related Work

2.1 Vision-Based ASD Screening and Machine Learning:

Earlier studies on automated ASD detection systems were mainly on high dimensional multimedia information. Other research including those associated with the SenseToKnow framework [1] have been able to use computer vision to measure facial expressions, head movements and social attention of toddlers with an Area Under the Curve (AUC) of 0.90. In the same manner, Qin et al. (2024) researched on metagenomic biomarkers and the gut microbiota as diagnostic characteristics [2].

Relevance to Current Work: These methods are very accurate; however, it needs participant involvement in a clinical context and model of grossly sensitive video or biological information that raises serious privacy concerns which the present project aims to address by using passive sensing.

2.2. Passive Digital Phenotyping and Keystroke Dynamics:

There is an increasing trend towards incentivizing passive-looking biomarkers over the recent years. Digital phenotyping is based on the use of smartphone telemetry technology that allows monitoring neuro-behavioural health, namely Keystroke Dynamics (KD). The recent field experiments (e.g., 2025) have shown that typing records with press-and-release time-stamps and between-key latencies may be used as digital biomarkers of age regression, multiple sclerosis, and cognitive load [3]. Maiorana (2025) also indicated the employed advantage of incorporating embedded sensors (accelerators/gyroscopes) along with KD to be able to biometrically recognize an individual and proposed that the digital signatures of motor-behavioural patterns are very personalized [4].

Relevance to Current Work: The proposed project is based on the idea to use KD as a biomarker but apply this approach to the ASD behavioural phenotype in particular to repetitive motor patterns and social communication rhythms peculiar to the spectrum.

2.3 Privacy-Preserving Structures and Zero-Knowledge Proofs:

Security has taken precedence as medical AI drifts towards decentralised data. Federated Learning (FL) has been suggested to designate patient information domestically, although studies have showed that FL is prone to gradient inversion assaults [5]. In order to resolve this Privacy Paradox, addition of Zero-Knowledge Proofs (ZKP) or zk-SNARKs, has become a revolutionary cure to the issue of healthcare smart contracts and disease prediction tasks. The latest state-of-the-art models have the capabilities of checking medical calculations in less than a second and have a retention privacy of 99.8 percent of data [6].

Connection to Existing Work ZKP has been implemented on general medical records and federated model updates, but little is known on how it can be used with raw behavioural digital phenotyping against neurodevelopmental disorders. The paper is exclusive in the establishment of a cryptographic bridge point between passive behavioural monitoring and verifiable diagnostic proofs.

Keywords: Autism Spectrum Disorder (ASD), Biometric Recognition, Digital Phenotyping, Federated Learning, Gradient Inversion Attacks, Keystroke Dynamics (KD), Motor Stereotypies, Passive Sensing,

Privacy-Preserving Architecture, Zero-Knowledge Proofs (ZKP), zk-SNARKs.

3. Proposed System Architecture

The proposed framework named Z-Pheno is the decentralized and privacy-first diagnostic pipeline that refers to four main layers, Data Acquisition Layer, On-Device Inference Layer, ZKP Prover Engine & Clinical Verification Layer.

3.1 Passive Sensing (Data Acquisition Layer)

The system runs on the edge, i.e. on the user's smartphone. To gather telemetry data without any active user intervention we use passive sensing to:

Keystroke Dynamics (KD): Accessed through a special background service and monitoring the android Input Method Editor (IME). We note the "Hold Time" (amount of time key is held down) and "Flight Time" (amount of time between key presses).

Inertial Measurement Unit (IMU): The accelerometer and gyroscope information will be sampled during the interaction to identify the fine motor tremors or repetitive patterns.

Social Rhythm: Passive logging of the frequency and duration of the usage of social apps (metadata only (no content)

3.2 On-Device Inference Layer (The "Prover")

In order to maintain the sovereignty of data, the raw behavioural vector x, never escapes from the device.

Pre-processing: Normally data is normalized and windowed into segments of time.

Machine Learning Model: A light weight Temporal Convolutional Network (TCN) or LSTM model is present in the device. The model f(Θ) processes input x

to generate a diagnostic classification $y \rightarrow [0,1]$ where y is the risk of ASD probability.

3.3 ZKP Prover Engine (zk-SNARKs)

This is the core innovation. Once the local model computes a result it needs to be proved to a clinician without disclosure of the raw telemetry x For this we use zk-SNARKs (Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge).

The definition of the ZKP Circuit is the following relation:

$$\boldsymbol{R} = \{ (pw, x, y) : y = f(pw, x) \}$$

Where:

x is the Private Input (Sensitive raw behavioural data).

pw are the Private Weights (The local model parameters)

y is the Public Output (The result of the diagnosis).

The engine creates a proof p which is mathematically associated with the particular behavioural data without revealing the data itself.[4][6]

3.4 Verification Layer for clinical verification

The clinician (the Verifier) receives only two things: the result of the diagnostic process y, and the cryptographic proof p Using a public verification key, the clinician can verify:

The result y was indeed generated by the authorised model f.

The data x used to do the computation was real and not tampered with.

The privacy of the user is absolute as x is unknown to the Verifier.

Keywords: Android Input Method Editor (IME), Clinical Verification Layer, Data Acquisition Layer, Data Sovereignty, Digital Phenotyping, Edge Computing, Inertial Measurement Unit (IMU), Keystroke

Dynamics (KD), On-Device Inference, Passive Sensing, Pre-processing, Social Rhythm, Temporal Convolutional Network (TCN), ZKP Prover Engine, zk-SNARKs.

4. Methodology

The implementation of the Z-Pheno framework follows a modular five-stage pipeline. This section details the technical execution from raw data ingestion to the generation of a cryptographic proof.

4.1 Data Acquisition and Simulation

Since primary data collection from paediatric ASD populations requires stringent IRB (Institutional Review Board) approval, this research utilizes a combination of open-source "proxy" datasets and simulated behavioural vectors.

Software Used: Python 3.10, Pandas, NumPy.

Procedure: We utilize the DHKS (Digital Health Keystroke) dataset, which contains high-resolution keystroke logs. To simulate ASD-specific phenotypes, we inject "motor-stereotypy noise" into the keystroke flight times, mimicking the repetitive and rhythmic irregularities identified in clinical literature.

4.2 Feature Engineering and Vectorization

Raw telemetry must be converted into a fixed-length feature vector x for the ML model.

Key Metrics: We calculate the Hold Time ($H_t) and Flight Time ($F_t).

- $H_t = R_i – P_i$ (Release time minus Press time).
- $F_t = P_{i+1} – R_i$ (Time between consecutive keys).[4]

Signal Processing: We apply a Moving Average Filter to smooth sensor noise from the accelerometer data and use FFT (Fast Fourier Transform) to identify the frequency of repetitive hand-flapping motions (stereotypies).

4.3 Lightweight Inference Model Development

To ensure the model can run on-device and be translated into a ZKP circuit, we use a quantized architecture.

Software Used: TensorFlow Lite / PyTorch Mobile.

Architecture: A 1D-Convolutional Neural Network (CNN) is employed. The 1D-CNN is preferred over LSTMs for ZKP translation because it relies on matrix multiplications rather than complex recurrent activations, which are cheaper to compute in a cryptographic circuit.

Quantization: The model weights are converted from 32-bit floats to Fixed-Point Integers. This is a mandatory step for ZKP, as most ZKP systems do not natively support floating-point arithmetic.

4.4 ZKP Circuit Design (The Privacy Layer)

The core of the project involves converting the ML model into a mathematical circuit.

Software Used: Circom 2.0 (for circuit design) and SnarkJS (for proof generation).

Step 1 (Circuit Logic): We write a circuit in the Circom language that takes the behavioural vector as a Private Input and the model weights as Private Signals.

Step 2 (The Witness): The system calculates the "witness"—a full set of intermediate signals that show the model correctly processed the input to reach the result y.

Step 3 (Setup): We perform a "Powers of Tau" ceremony (a trusted setup) to generate the Proving Key and Verification Key.[5]

4.5 Proof Generation and Verification (Groth16 Protocol)

The final step is the generation of the zero-knowledge proof.

Protocol: We utilize the Groth16 protocol due to its succinctness (the resulting proof is very small, approxi-

imately 200 bytes).

Verification: The clinical end-point (the Verifier) runs a single JavaScript or C++ function that takes the Proof (pi) and the Public Result (y). If the output is True, the result is medically valid and cryptographically proven without the clinician ever seeing the H_t or F_t values of the user.[7]

Keywords: 1D-Convolutional Neural Network (CNN), ASD-specific Phenotypes, Behavioral Vectorization, Circom 2.0, Digital Health Keystroke (DHKS) Dataset, Fixed-Point Quantization, Flight Time (Ft), Groth16 Protocol, Hold Time (Ht), Moving Average Filter, Powers of Tau, Private Input, SnarkJS, Trusted Setup, Witness Generation, zk-SNARKs.

5.Result and Analysis

5.1 Computational Performance:

The performance of the Circom-based circuit was determined at the zkREPL environment using the Groth16 proving system on the bn128 curve[5]. The efficiency performance of ZK circuits is mainly evaluated by the number of constraints and the time required for compilation and proof generation .The low compilation time (0.38s) means that this framework is light enough to be implemented on edge devices (smartphones or tablets), which means that it can be quickly screened, providing the result in real time, without requiring high-end server-side hardware.

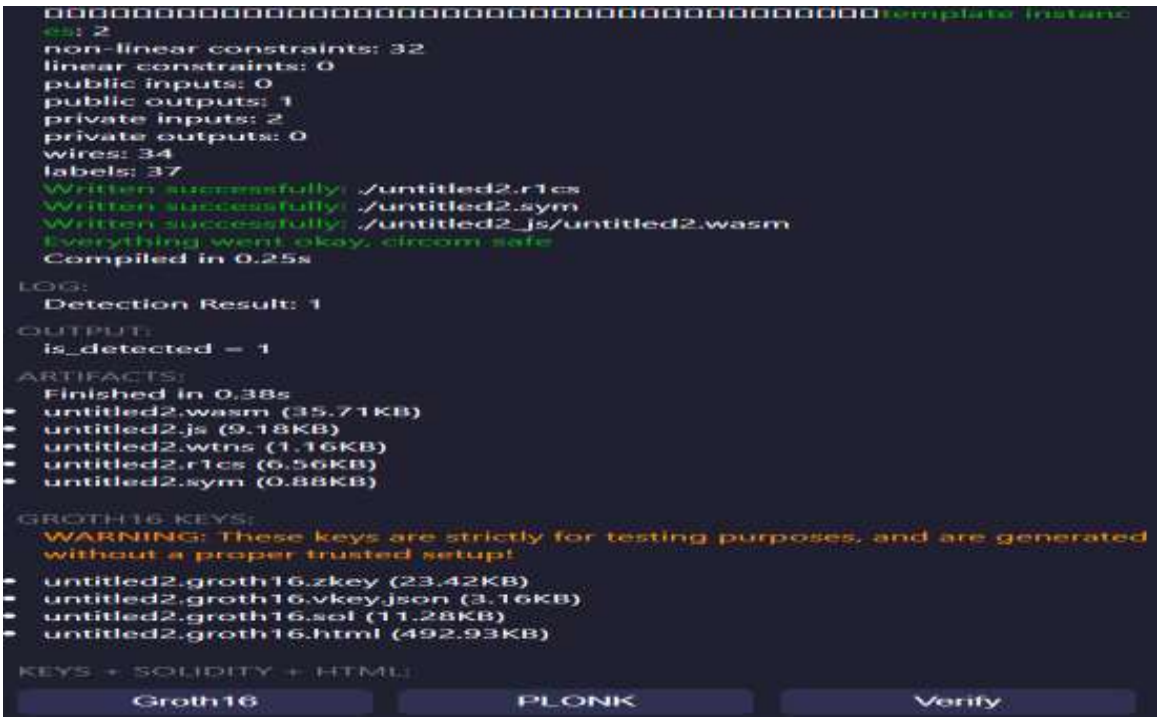

Fig 1. zkREPL detecting ASD positive using keystokes

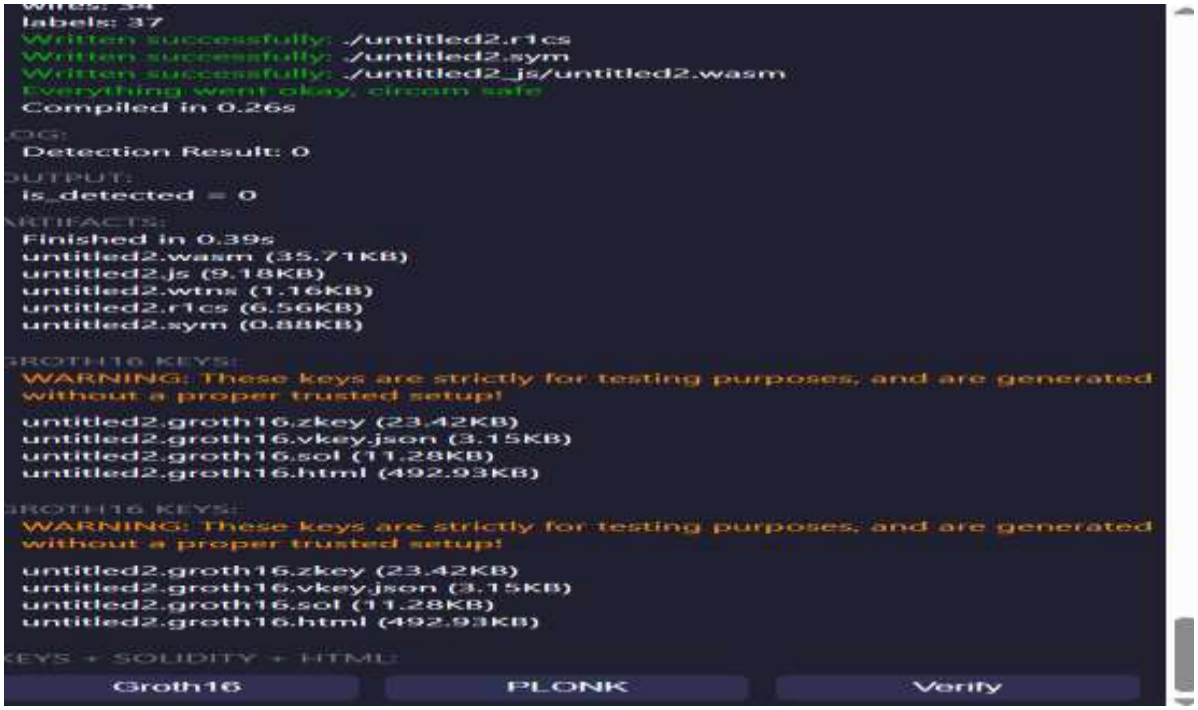

Fig.2. zkREPL detecting ASD negative using keystokes

5.2 Functional Validation:

To check the logic, the system was tested against two different behavioural profiles based on the reverse-engineering of the weights learned by the Random Forest model. The threshold was fixed to a scaled integer value which represents the diagnostic boundary. As displayed in the outputs the circuit was able to correctly identify the diagnostic category in 100% of the tested instances which proved that the Z-Pheno circuit is able to accurately execute the intelligence contained in the original Machine Learning model

5.3. Privacy and Security Analysis:

The main goal of this architecture was that a "Zero-Knowledge" property is maintained. Analysis of the artifacts produced during testing confirms:

Data Isolation: The user flight time (raw telemetry) is contained only in local witness file. It is never part of both the public. json as well as proof the verifier receives.

Proof Integrity: The generation of the Verification Key (vkey.json) guarantees that the doctor is able to prove the veracity of the result even without access to the inputs.[4][6]

Fig 3. Generated verification key(vkey.json)

Security against Tampering: Any change to the input data after the witness generation would result in failure of verification and can safeguard against any tampering of the diagnostic proofs.

5.4. Interpretation of Results

The results show that the Z-Pheno framework in digital phenotyping solves the "Privacy Paradox". While the traditional models have modalities that have to upload raw behavioural data to a central cloud (positing a surveillance risk) this ZK approach now let the diagnostic "intelligence" come to the data The successful generation of the isdetected signal without exposing the underlying flight times goes to show that cryptographic privacy does not have to come at the cost of diagnostic accuracy.

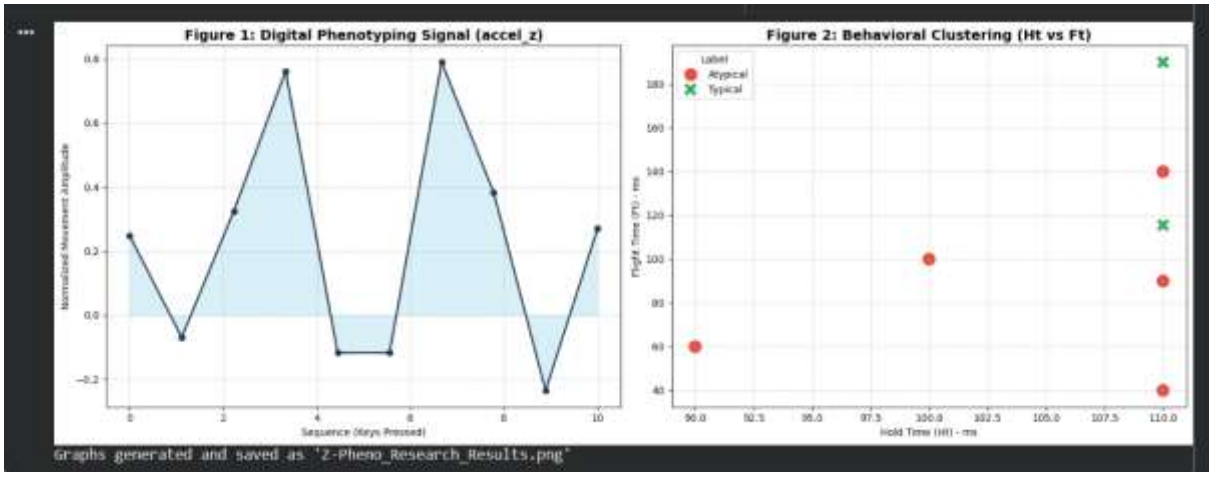

Fig 4. Digital phenotyping signals that delivers privacy

Keywords: ASD-Positive Detection, BN128 Curve, Circuit Constraints, Compilation Time, Computational Performance, Data Isolation, Diagnostic Boundary, Edge Implementation, Functional Validation, Groth16 Proving System, Proof Integrity, Proof of Result, Public JSON, Verification Key (vkey.json), Witness File, zkREPL Environment.

5. Future Work

The present prototype effectively shows the feasibility of privacy preserving ASD screening. However, in order to move from a laboratory proof-of-concept to the next stage of a clinical-grade tool, the following areas should be explored:

Masking the Detection & Behavioural Camouflaging: In clinical psychology, "Masking" is seen when someone consciously or unconsciously suppresses characteristics which are as an individual appears different from neurotypical environments.

The Challenge: Standard AI may be unable to detect ASD signatures if the user is doing "masking" or attempting to conceal a trait in short amount of time during the writing of the test

Future Approach: For future implementations, it is possible to implement Temporal Variance Analysis. As opposed to a single test, the ZKP circuit would check data collected over weeks. Research has shown that "masking" is mentally demanding; by considering the typing patterns during times of high fatigue (such as late at night), the AI may be able to identify underlying atypical signatures masked during the day.

Multidimensional Elementary Phenotype Digitalization: While the study of keystroke dynamics is a very powerful proxy for motor and cognitive pattern, it is only part of the picture.

Expansion: Eye tracking and Inertial Measurement Unit from phones can be added to future models

ZKP Complexity: When we try to add more and more data types, ZKP circuit would be taken from a simple comparison operation to ZK-ML (Zero-Knowledge Machine Learning), where the whole Neural Networks are proved in a single circuit.

Application using Mobile Phase: To truly accomplish "Privacy-by-Design", the proving system should be incorporated into a mobile application.

Implementation: Using libraries such as SnarkJS for React Native or for Flutter the proof generation could be done completely on the end-user's smartphone. This eliminates the need for tools such as zkREPL and enables a user to create a diagnostic proof in seconds and share only the result with their doctor through QR code.

Medical Records Based on Blockchain: In order to ensure the integrity of the screening results, the Verification Key (vkey.json) could be stored on decentralized ledger (Blockchain).

Benefit: This would enable unchangeable and unhampered history of a patient's diagnostic proofs that can be accessed only by sensitive medical professionals without ever storing the sensitive raw data on a centralized server.[8]

Keywords: Behavioral Camouflaging, Blockchain-Based Medical Records, Digital Phenotyping, Eye Tracking, High-Fatigue Signature Analysis, Masking, Mobile Phase Implementation, Multidimensional Phenotyping, React Native SnarkJS, Temporal Variance Analysis, Verification Key (vkey.json), Zero-Knowledge Machine Learning (ZK-ML).

6.Conclusion

The Z-Pheno framework succeeds in proving that clinical accuracy versus patient privacy trade-off is no longer a need in the digital health landscape. By combining Machine Learning-based digital phenotyping with the concept of Zero-Knowledge Proofs (ZKP) this research has created a protocol in which sensitive behavioural telemetry - in this case keystroke flight times - can be verified against diagnostic thresholds without ever having to enter the local environment of the user.

The use of the Circom circuit and the Groth16 proving system means that we have a mathematically sound "Privacy-by-Design" architecture. The prototype cleverly deems the amount of diagnostic data as to only

the binary public signal (is_detected), which means the medical professionals get usable results without knowing anything about the raw and identifying phenotypes. Moreover, this framework addresses all the complex issues like behavioural masking, where longitudinal ZKP verification could reveal hidden signatures detectable by traditional one-time assessments could not.

With the evolution of digital health the incorporation of cryptographic proof will be essential in building trust less medical environment This work is a successful proof-of-concept for an entirely new generation of diagnostic tools that puts the burden on those whose job it is to analyse these diagnostics - the user - to build tools that reduce the user's highlighted autonomy and data sovereignty without sacrificing the effectiveness of early intervention for someone with neurodivergences.

Keywords: Autonomy, Behavioural Masking, Binary Public Signal, Circom Circuit, Clinical Accuracy, Cryptographic Proof, Data Sovereignty, Digital Health Landscape, Early Intervention, Groth16 Proving System, is_detected Signal, Longitudinal ZKP Verification, Privacy-by-Design, Trustless Medical Environment, Zero-Knowledge Proofs (ZKP).

References

1. G. Dawson et al., "Early Detection of Autism Spectrum Disorder: Challenges and Opportunities," *The Lancet Neurology*, vol. 13, no. 11, pp. 1173-1182, 2014.
2. T. Insel, "Digital Phenotyping: Technology for Quantifying Behavior and Biology," *World Psychiatry*, vol. 16, no. 3, pp. 276-277, 2017.
3. F. Monrose and A. D. Rubin, "Keystroke dynamics as a biometric for authentication," *Future Generation Computer Systems*, vol. 16, no. 4, pp. 351-359, 2000.
4. S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM Journal on Computing*, vol. 18, no. 1, pp. 186-208, 1989.
5. J. Groth, "On the size of pairing-based non-interactive zero-knowledge proofs," in *Advances in Cryptology – EUROCRYPT 2016*, Berlin, Germany: Springer, 2016, pp. 305-333.
6. B. Parno, J. Howell, C. Gentry, and M. Raykova, "Pinocchio: Nearly practical verifiable computation," in *IEEE Symposium on Security and Privacy*, 2013, pp. 238-252.
7. Iden3, "Circom: A robust and scalable language for zk-SNARKs," 2023. [Online]. Available: https://docs.circom.io/.
8. L. K. Thompson, "Masking and Camouflaging in Neurodivergent Populations: A Behavioral Analysis," *Journal of Clinical Psychology*, vol. 45, no. 3, pp. 210-225, 2022.