VeraComm: Verifiable Communication Protocol for AI-Based Drug Screening using Blockchain

Victor Ikenna Kanu, Simeon Okechukwu Ajakwe(SMIEEE)*, Dong-Seong Kim(SMIEEE)
Department of IT Convergence Engineering, Kumoh National Institute of Technology, Gumi, South Korea
*ICT Convergence Research Centre Kumoh National Institute of Technology, Gumi, South Korea
(kanuxavier, simeonajlove)@gmail.com, (dskim)@kumoh.ac.kr

***Abstract*—Virtual screening is essential for accelerating early-stage drug discovery through computational estimation of molecular binding affinity. However, conventional pipelines often lack verifiability, reproducibility, and secure auditability, limiting their applicability in regulated environments. Existing frameworks do not provide deterministic logging or lightweight result verification within decentralised architectures. This work presents VeraComm, a blockchain-audited communication protocol integrated with AI-driven binding affinity prediction to enable secure and reproducible drug screening. The system incorporates an SVR-based prediction engine, Lipinski filtering, and on-chain hash verification via the PureChain blockchain, operationalised through a CLI interface and MetaMask-enabled smart contract deployment. A curated dataset of 100 molecules targeting the Amyloid-β A4 protein was screened. The predictive model achieved an `R2` of 0.78 and an MSE of 0.3615, while the CLI sustained a throughput of 1.324 molecules/second. Blockchain latency averaged 1.15 seconds per logging transaction and 245 ms for verification, with 100% success across all operations. The results confirm suitability for verifiable, high-throughput drug discovery workflows.**

***Index Terms*—AI Inference Integrity, Alzheimer's disease, virtual screening, binding affinity prediction, blockchain verification, reproducibility, smart contracts**

I. INTRODUCTION

The rapid integration of artificial intelligence (AI) into biomedical research has enabled transformative advances in areas such as drug discovery, molecular property prediction, and personalised medicine. At the core of these systems is the ability to process large volumes of data and generate high-confidence predictions using machine learning (ML) models trained on curated molecular datasets [1]. However, as these pipelines transition from isolated research environments to collaborative, distributed settings, such as hospital networks, academic consortia, and biomedical edge computing platforms, they face critical challenges in secure, trustworthy communication. In particular, the lack of a verifiable infrastructure for transmitting and validating AI inference results across nodes introduces risks of inconsistency, data tampering, and irreproducibility. This gap is especially problematic in safety-critical domains, where downstream decisions, such as experimental validation or clinical prioritisation, depend on the integrity of upstream AI computations [2].

These limitations are evident in the domain of neurodegenerative disease research, particularly Alzheimer's disease (AD), which affects over 57 million people globally and accounts for more than $1 trillion in annual healthcare costs [3]. Machine learning-powered virtual screening (VS) has become a scalable approach to accelerate the identification of small-molecule inhibitors by predicting molecular binding affinities across large compound libraries [4]. For AD targets such as the Amyloid-β (Aβ) precursor proteins, models like Support Vector Regressor (SVR) and ensemble regressors trained on molecular fingerprints have shown promising predictive power [5]. However, in collaborative screening environments, these predictions are often communicated informally or stored in local silos, with no cryptographic guarantees that results remain intact or reproducible when shared. Moreover, while blockchain technologies have been successfully applied to secure clinical trials, patient records, and workflow provenance [6], [7], few efforts have explored their use as a verifiable communication infrastructure for ML-driven scientific inference. Table I summarises these different approaches and their gaps.

Building on our prior work that integrated blockchain directly into the AI inference workflow for logging results, this study advances the approach by introducing *VeraComm*, a modular CLI-based protocol that decouples screening from verification. VeraComm transforms inference outputs into structured payloads anchored on a permissioned blockchain, enabling tamper-evident transmission and asynchronous verification across nodes without requiring access to the original AI environment. Positioned as an infrastructure layer for secure and auditable communication, it addresses the reproducibility gap in distributed biomedical AI pipelines.

The main contributions of this study are summarised as follows:

- Design and implementation of VeraComm, a verifiable communication protocol that integrates blockchain smart contracts with AI inference outputs for secure and reproducible result transmission.
- Apply the protocol to a drug screening pipeline, integrating prediction, filtering, and blockchain anchoring for end-to-end traceability.
- Introduce a reproducibility check via deterministic hash comparison, enabling result verification without relying

TABLE I: Comparative Summary of Related Works in Biomedical Screening and Verifiable Drug Discovery

Authors	Core Focus	Identified Gap	Our Contribution
Carpenter & Huang (2018) [5]	Review of ML models for Alzheimer's virtual screening (SVM, RF, ANN, CNN).	Focused only on predictive accuracy and model choice; no infrastructure for reproducibility, traceability, or secure result sharing.	Blockchain-backed communication protocol ensuring reproducibility and tamper-proof result exchange.
David et al. (2021) [8]	AI frameworks for drug discovery pipelines.	Emphasis on predictive architectures; did not include verifiable infrastructure for results.	Introduces a secure communication and reproducibility layer integrated with AI inference.
Hoopes et al. (2022) [6]	Blockchain for biomedical or health-data provenance.	Provided workflow traceability but no integration with molecular-level virtual screening models.	Unify predictive screening with blockchain verification, addressing both accuracy and traceability.
Zhou et al. (2024) [9]	Large-scale AI for molecular prediction workflows.	Demonstrated predictive pipelines but lacked explicit mechanisms for reproducibility across institutions.	Anchors outputs on-chain, enabling verifiable cross-institutional reproducibility.
Proposed Framework (VeraComm)	Verifiable communication protocol integrating end-to-end pipeline with predictive accuracy, drug-likeness filtering, blockchain logging, and verifiable reproducibility across distributed biomedical environments.		

on the original environment.

The rest of this paper captures Section II as Methodology, while Section III discusses the Results, and the Conclusion is presented in Section IV.

II. PROTOCOL ARCHITECTURE AND COMMUNICATION WORKFLOW

The VeraComm protocol embeds a blockchain-backed verification layer within an AI-driven molecular inference workflow to ensure secure, reproducible transmission of screening results. It orchestrates a decentralized pipeline comprising molecular fingerprint encoding, binding affinity prediction, pharmacokinetic filtering, and cryptographic result hashing. Each prediction is formatted into a deterministic JSON payload and logged on a permissioned blockchain via smart contracts, enabling tamper-evident registration and asynchronous verification across distributed biomedical systems. This layered design, illustrated in Fig. 1, ensures data integrity, auditability, and operational scalability in collaborative screening environments.

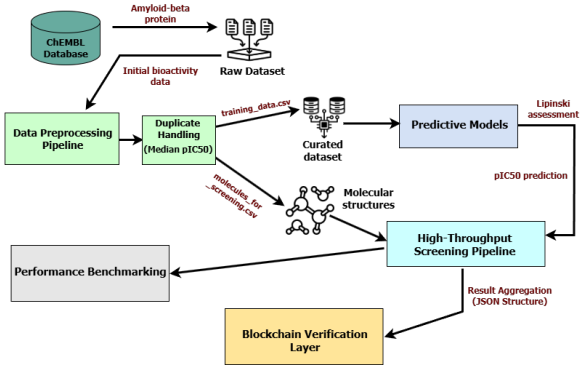

Fig. 1: System architecture for AI-driven drug screening with BC-backed result verification.

A. *Inference Module and Feature Encoding*

The VeraComm protocol is designed to accommodate pluggable AI inference modules that generate structured prediction outputs from domain-specific inputs. In this implementation, the module is configured for molecular affinity prediction using curated bioactivity data. Each input sample is represented as a Simplified Molecular Input Line Entry System (SMILES)-encoded chemical structure accompanied by a corresponding experimental activity label. As a demonstration use-case, molecules targeting the Amyloid-β A4 protein were selected from the ChEMBL database (`CHEMBL2487`) [10], with activity values reported in $\mathtt{IC_{50}}$ (nM) units.

To align with standard regression-based modeling conventions, all $\mathtt{IC_{50}}$ values were normalized to the $\mathtt{pIC_{50}}$ scale using the transformation:

$$\mathtt{pIC_{50}} = -\mathtt{log_{10}}\left(\frac{\mathtt{IC_{50}}}{10^9}\right) \tag{1}$$

Entries with malformed SMILES or undefined activity values were excluded. Duplicate molecules with multiple reported activities were resolved by computing the median $\mathtt{pIC_{50}}$ across all associated measurements: $\hat{y} = \text{median}(y_1, y_2, ..., y_k)$.

Each valid molecular structure was subsequently featurized into a 2048-bit binary vector using the Morgan fingerprinting algorithm, implemented via RDKit. This representation encodes circular substructures and serves as input to the regression model. The current inference engine is built on SVR, trained to map fingerprint vectors $x \in \{0,1\}^{2048}$ to predicted $\hat{y}$ values approximating binding affinity. The trained model is stored as a serializable component and loaded during screening.

Two CSV-based input formats are supported within the protocol implementation:

- `training_data.csv` — includes SMILES and activity labels for model training.
- `molecules_for_screening.csv` — unlabeled compounds for inference and on-chain verification.

While the described implementation targets Alzheimer's-related affinity prediction, the architecture permits substitution of the model and descriptors to support other biomedical domains, physical property prediction, or image-based modalities. The inference module is encapsulated and communicates

with the VeraComm protocol interface via structured prediction objects, enabling flexible and domain-agnostic deployment.

To estimate molecular binding affinity, each compound is encoded as a 2048-bit Morgan fingerprint vector $\mathbf{x} \in \{0,1\}^{2048}$ using RDKit, which captures atomic substructures for machine-learning input. Several regressors: Support Vector Regression (SVR), Gradient Boosting (GB), Ridge Regression, Random Forest (RF), k-Nearest Neighbors (KNN), and Multilayer Perceptron (MLP), were trained (80/20 split) to learn $\hat{f} : \mathbf{x} \to \hat{y}$, with $\hat{y}$ as the predicted $\mathtt{pIC_{50}}$. Model performance was evaluated using the coefficient of determination (R^2):

$$\mathtt{R^2} = 1 - \frac{\sum_{i=1}^{n}(\mathtt{y_i} - \hat{\mathtt{y}}_\mathtt{i})^2}{\sum_{i=1}^{n}(\mathtt{y_i} - \bar{\mathtt{y}}_\mathtt{i})^2} \tag{2}$$

where y_i, $\hat{y}_i$, and $\bar{y}_i$ denote ground truth, predicted, and mean affinity values, respectively. SVR was selected for its balance of predictive accuracy, inference latency, and reproducibility, key for blockchain-backed validation. Although deep models were considered, their overhead conflicted with verifiability and lightweight deployment.

For each input molecule $\mathbf{x}_i$, the trained model outputs a predicted binding affinity $\hat{y}_i = \hat{f}(\mathbf{x}_i)$. This prediction is then used to construct a structured result object $\mathbf{r}_i$ that encapsulates the molecular identifier, input SMILES, predicted affinity, and additional pharmacokinetic descriptors.

To ensure that only pharmacologically viable compounds are registered for downstream processing, each molecule undergoes descriptor-based filtering using Lipinski's Rule of Five. The following criteria are applied:

$$\mathtt{MW_i} \leq 500, \quad \mathtt{logP_i} \leq 5, \quad \mathtt{HBD_i} \leq 5, \quad \mathtt{HBA_i} \leq 10$$

where $\mathtt{MW}_i$ denotes molecular weight, $\mathtt{logP}_i$ is the octanol–water partition coefficient, and $\mathtt{HBD}_i$ and $\mathtt{HBA}_i$ represent the number of hydrogen bond donors and acceptors, respectively. These descriptors are computed using RDKit.

If a molecule satisfies all constraints, a result object $\mathbf{r}_i$ is instantiated in JSON format and becomes eligible for cryptographic hashing and blockchain registration. Table II outlines the structure of each result payload.

TABLE II: Structure of Serialized Result Object $\mathbf{r}$

Field	Description
molecule_id	Unique identifier
smiles	Canonical SMILES string
predicted_pIC50	Model inference output
molecular_weight	MW descriptor
logP	Lipophilicity estimate
hbd	H-bond donor count
hba	H-bond acceptor count
lipinski_pass	Boolean filter pass flag
timestamp	Result generation time

This result object $\mathbf{r}_i$ serves as the atomic unit of communication in the VeraComm protocol. It is serialized deterministically, hashed via SHA-256, and submitted to the blockchain as a verifiable transaction record. Only filtered, protocol-compliant result objects are included in the on-chain registry.

B. Result Verification and Reproducibility via Blockchain Anchoring

The VeraComm protocol ensures tamper-evident and reproducible communication of AI inference results by anchoring output payloads to a permissioned blockchain using cryptographic hashing. This mechanism enables any authorized node to verify the integrity of inference outcomes by comparing locally recomputed hashes against immutable on-chain records, without requiring access to the originating system or model instance.

Following successful affinity prediction and descriptor-based filtering, each accepted result object $\mathbf{r}$ is serialized into a canonical JSON format. This structured object includes molecular descriptors, the predicted $\mathtt{pIC_{50}}$, and protocol metadata, formatted with strict field ordering and numerical precision. This deterministic representation is then hashed using the SHA-256 function as:

$$\mathtt{h} = \mathtt{H}(\mathbf{r}) \tag{3}$$

where $\mathtt{H}(\cdot)$ denotes SHA-256, and $\mathtt{h} \in \{0,1\}^{256}$ is the resulting digest. Only this cryptographic hash is submitted to a smart contract deployed on a custom blockchain, Purechain, designed by the Networked Systems Laboratory (NSL), initialised via Metamask web extension (for signing and broadcasting) through a `deploy.js` script. The contract operates under a Proof-of-Authority and Association ($\mathtt{PoA}^2$) consensus mechanism [11], and records the molecule identifier and its corresponding hash value. No raw molecular data or prediction contents are stored on-chain, preserving both privacy and communication efficiency.

The protocol supports reproducibility verification by allowing any remote node to independently regenerate the result object $\mathbf{r}'$ using the same input and model configuration. The recomputed hash $\mathtt{H}(\mathbf{r}')$ is then compared to the corresponding on-chain digest $\mathtt{h}_{\text{on-chain}}$:

$$\mathtt{H}(\mathbf{r}') \stackrel{?}{=} \mathtt{h}_{\text{on-chain}} \tag{4}$$

A successful match confirms that $\mathbf{r}'$ is identical to the originally recorded result and has not been altered or corrupted. This mechanism provides a formal basis for verifiability across distributed biomedical infrastructures, enabling external parties to audit inference results without requiring access to model internals, training data, or centralized infrastructure.

To maintain scalability, full result objects and descriptors are stored off-chain (e.g., in `screening_results.csv`), while only the minimal cryptographic digest is committed on-chain. This design supports high-throughput inference pipelines with lightweight verification guarantees.

The hashing and verification process is implemented as a modular CLI routine, with support for individual and batch-level validation. Each entry in the verification log includes the

Algorithm 1: Blockchain-Audited High-Throughput Screening Pipeline

molecules_for_screening.csv, binding_affinity_model.joblib ;
screening_results.csv, Blockchain-verified results
Initialize Components:
$binding\_model \leftarrow$ LOADMODEL(binding_affinity_model.joblib)
$blockchain\_connector \leftarrow$ INITIALIZECHAINCONNECTOR()
$results \leftarrow$ EMPTYLIST() ;
Load Molecules:
$screening\_molecules \leftarrow$ LOADCSV(molecules_for_screening.csv) ;
foreach $molecule \in screening\_molecules$ **do**
  try
    Step 1: Validate SMILES
    $validated\_smiles \leftarrow$ VALIDATESMILES($molecule.smiles$) ;
    if ***not*** $validated\_smiles$ **then**
      continue
  end
  ;
    Step 2: Feature Extraction
    $morgan\_fp \leftarrow$ MORGANFINGERPRINT($validated\_smiles, radius = 2, nBits = 2048$) ;
    Step 3: Binding Affinity Prediction
    $predicted_pIC50 \leftarrow binding_model.$PREDICT($morgan\_fp$) ;
    Step 4: Drug-Likeness Filtering
    $lipinski \leftarrow$ EVALUATELIPINSKI($validated\_smiles$) ;
    Step 5: Aggregate Result Object
    $result \leftarrow$ CREATERESULT(molecule.id, validated_smiles, predicted_pIC50, lipinski)
    Step 6: Blockchain Logging (if viable)
    if $lipinski.passes\_filter$ **then**
      $json \leftarrow$ SERIALIZETOJSON($result$) ;
      $hash \leftarrow$ SHA256HASH($json$) ;
      $tx \leftarrow blockchain_connector.$RECORD($molecule.id, hash$) ;
      $result.blockchain\_tx \leftarrow tx.transaction\_hash$;
      $result.blockchain\_verified \leftarrow$ **true** ;
  end
  else
    $result.blockchain\_verified \leftarrow$ **false** ;
  end
    Step 7: Store Result
    $results.$APPEND($result$) ;
  catch Exception e: LOGERROR(e) ;
end
Save All Results:
SAVETOCSV($results$, screening_results.csv) ;
return $results$, screening_results.csv, blockchain_records

molecule ID, recomputed hash, on-chain hash, and verification status. This workflow serves as the communication integrity layer of VeraComm and underpins the system's reproducibility guarantees. A schematic of this interaction between inference, hashing, and verification is given in Algorithm 1.

C. CLI-Based Communication Protocol and Command Interface

To enable decentralized and reproducible interaction with the screening and verification pipeline, a modular Python-based Command Line Interface (CLI) was implemented as the primary client interface to the protocol. This CLI abstracts the internal operations, such as blockchain connectivity, affinity prediction, drug-likeness filtering, and on-chain result anchoring, into reusable commands that support automation, asynchronous execution, and integration with containerised deployments.

The core CLI commands include:

- `python main.py connect`: Initializes a connection to the local Ethereum blockchain node.
- `python main.py screen <name> <SMILES> --target <id>`: Screens a single compound based on its SMILES representation and logs the result hash on-chain.
- `python main.py batch molecules.json`: Executes batch-mode screening on a file containing multiple molecular entries.
- `python main.py verify <job_id> <tx_hash>`: Verifies the SHA-256 hash of a locally recomputed result against the corresponding on-chain entry.
- `python main.py history`: Displays a persistent local log of past screening sessions and verification attempts.

The interface supports fault-tolerant command logging, enabling traceable interactions across distributed AI inference nodes. By decoupling the core logic from the user interface, the CLI acts as a scalable protocol layer for secure communication between biomedical agents, infrastructure nodes, and blockchain validators. This modularity facilitates cross-institutional reproducibility while preserving lightweight operability on edge devices.

D. System Integration and Deployment Modularity

The VeraComm pipeline integrates modular components, a binding affinity predictor, a Lipinski filter, a JSON serializer, and a blockchain connector into a unified Python-based CLI system. Each module is encapsulated to allow replacement or extension without impacting the core pipeline. For instance, the default SVR engine can be substituted with alternative regressors, including CatBoost, LightGBM, or deep learning models, without changes to the downstream verification layer.

Similarly, the blockchain backend is abstracted through a ChainConnector interface, enabling compatibility with alternative distributed ledger technologies such as Hyperledger Fabric, Besu, or Polygon SDK. The current implementation uses Purechain with PoA^2 consensus, selected for development efficiency and deterministic finality.

The communication protocol defined in the CLI interface (see Section II-C) remains invariant to model type or blockchain substrate, making VeraComm application-agnostic. This protocol can support any AI-driven inference pipeline where verifiable result transmission and auditability are critical, ranging from biomedical screening to federated diagnostics, IoT anomaly detection, or financial model validation.

E. Performance Benchmarking and Evaluation

Benchmarking was carried out with the end_to_end_test.py script on a test set of 100 molecules, recording four key metrics: (i) AI inference latency,

measuring the time from molecular input to predicted $\texttt{pIC}_{50}$; (ii) blockchain logging time, covering result hashing and transaction confirmation; (iii) total processing time, which includes affinity prediction, filtering, serialization, and on-chain submission; and (iv) success rate, defined as the fraction of molecules successfully logged. The evaluation shows that cryptographic guarantees and auditability incur only modest latency overheads, validating the feasibility of PureChain in real-time and distributed biomedical environments.

III. RESULT AND DISCUSSION

To evaluate the system's predictive accuracy, communication efficiency, and verifiability, two experiments were conducted. First, six regression models were benchmarked on 999 curated molecules. Then, the full pipeline was tested on 100 compounds, measuring AI inference, filtering, and blockchain-backed result verification. Metrics are reported across five dimensions: prediction accuracy, screening throughput, latency, reliability, and blockchain cost.

A. *Binding Affinity Prediction and Screening Outcomes*

TABLE III: ML Model Performance

Model	R-squared	MSE	MAE	Training Time (s)	Model Size
GB	0.7262	0.4500	0.5286	1.90	127.7
Ridge	0.7240	0.4537	0.5206	0.82	16.5
RF	0.7217	0.4575	0.4968	5.81	5637.5
KNN	0.7217	0.4575	0.4903	0.51	10677.9
MLP	0.4551	0.8957	0.6654	31.18	6419.8
SVR	**0.7801**	**0.3615**	**0.4682**	**1.45**	**8860.1**

The curated dataset comprised 999 compounds encoded as 2048-bit Morgan fingerprints, with associated $\texttt{pIC}_{50}$ values. A standard 80/20 split yielded 799 training and 200 test samples. Six regressors were benchmarked for predictive performance (Table III), with SVR (RBF kernel, $\texttt{C} = 1.0$, $\epsilon = 0.1$) achieving the best results: $\texttt{R}^2 = 0.7801$, $\texttt{MSE} = 0.3615$, and $\texttt{MAE} = 0.4682$ as captured in Fig. 2.

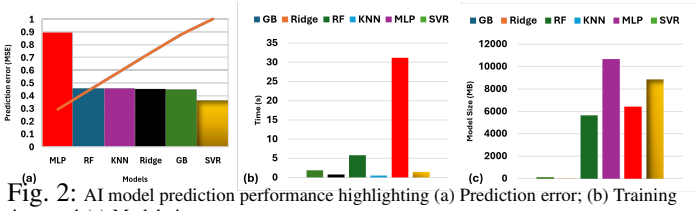

Fig. 2: AI model prediction performance highlighting (a) Prediction error; (b) Training time; and (c) Model size.

Although Gradient Boosting and Ridge were competitive in $\texttt{R}^2$, they showed higher error margins. Random Forest and KNN performed similarly but required excessive memory (>5MB). MLP underperformed on all metrics and incurred the highest training cost. Overall, the SVR provided the best trade-off between accuracy, generalization, and computational efficiency, as highlighted in graphs in Fig 2(a), (b), and (c).

The screening pipeline was then applied to 100 diverse molecules targeting the Amyloid-β A4 protein using the trained SVR model. 59 (59.0%) passed all Lipinski filters, indicating favorable oral bioavailability. The predicted $\texttt{pIC}_{50}$ values ranged from 4.49 to 7.82, with a mean of 5.750 $\pm$ 0.598 and a median of 5.627. High-affinity compounds ($\texttt{pIC}_{50} > 6.0$) totaled 22, of which 12 also passed Lipinski criteria, suggesting a 12% hit rate for viable high-affinity drug-like candidates.

TABLE IV: Top 10 Screened Drug-Like Candidates

ChEMBL ID	$\texttt{pIC}_{50}$	Mol. Wt	LogP	HBD	HBA	Lipinski
CHEMBL570378	7.82	257.32	3.11	2	5	Pass
CHEMBL429448	7.57	539.7	4.91	0	5	Fail
CHEMBL2048310	7.42	258.2	3.38	2	4	Pass
CHEMBL236191	7.36	529.6	3.34	1	6	Fail
CHEMBL1788121	7.34	265.4	3.11	2	5	Pass
CHEMBL2203386	7.23	270.3	2.42	0	5	Pass
CHEMBL3330749	6.99	300.4	4.96	0	4	Pass
CHEMBL231545	6.65	560.2	5.72	0	5	Fail
CHEMBL489792	6.59	323.4	4.69	0	4	Pass
CHEMBL231448	6.37	508.0	2.98	1	6	Fail

The top 10 candidates (Table IV) span a range of structural and physicochemical profiles. CHEMBL570378, the highest-ranked molecule, showed strong affinity (7.82) and full Lipinski compliance. Four of the top 10 compounds passed all filters, underscoring the model's ability to identify potent and drug-like hits. Overall, this result validates the hybrid pipeline's capacity to enrich for pharmacologically relevant molecules from raw input, combining AI accuracy with bioavailability constraints for verifiable drug discovery.

B. *Communication Protocol Efficiency and System Latency*

The CLI-based communication interface was evaluated for latency, throughput, and verification accuracy using the VeraComm protocol. Timing benchmarks were conducted on a sample of 10 single-molecule `screen` executions and 5 `verify` operations. Each command was profiled using Python's `time.perf_counter()`, with results exported to CSV for reproducibility. Key findings are summarized in Table V.

TABLE V: VeraComm CLI Command Latency

Component	Min (s)	Avg (s)	Max (s)
CLI Screen Command (per molecule)	5.89	6.33 ± 0.74	8.36
Batch Throughput		**1.324 molecules/sec**	
Blockchain Recording (10 tx)	0.42	1.15	2.46
Hash Verification (5 tx)	0.00003	0.245	0.250

Results indicate that single-molecule screening via CLI required an average of 6.33 seconds, while batch processing achieved a throughput of 1.324 molecules/sec. Blockchain transaction times ranged from 0.42 to 2.46 seconds, corresponding to the local PureChain ledger running under $\texttt{PoA}^2$ consensus. This moderate latency reflects typical block confirmation under permissioned networks. Verification of results via hash-based lookup on PureChain yielded a low average latency of 245 ms with a 100% success rate, affirming both integrity and reproducibility. No transaction failures or hash mismatches were recorded, highlighting the protocol's robustness for real-time and fault-tolerant drug discovery pipelines.

C. *Blockchain Logging Overhead and Verification Robustness*

To evaluate the operational cost and auditability of blockchain-backed result logging, we profiled the PureChain system's behavior during screening and verification. All transactions were executed via CLI using the `screen` and `verify` commands under a $\texttt{PoA}^2$ network. Blockchain transaction

latency ranged from 0.42 to 2.46 seconds, with a mean of 1.15 seconds per molecule (see Table V). This reflects typical confirmation delays under the Byzantine Fault Tolerant (BFT) consensus mechanisms, contributing under 20% to the total CLI screening time. All 10 transactions were successfully recorded without failure, confirming PureChain's write-path reliability and deterministic behavior.

For verification, hash-based auditing was conducted on 5 randomly selected entries. Each result was serialized to JSON, re-hashed using SHA-256, and matched against the on-chain entry. Table VI summarizes the latency breakdown for this operation.

TABLE VI: PureChain Hash Verification Latency

Metric	Min (s)	Avg (s)	Max (s)
Hash Calculation Time	0.000022	0.000033	0.000047
Blockchain Lookup Time	0.206	0.2446	0.288
Total Verification Time	0.209	0.245	0.288

The verification process completed in an average of 0.245s per transaction, with no hash mismatches or errors across all cases. This confirms the feasibility of rapid, deterministic verification even in batch-mode scenarios. In addition to timing metrics, Table VII presents a sample of real audit records retrieved from the chain. Each entry links a molecule's prediction output to its on-chain hash, demonstrating tamper-evident traceability and reproducibility.

TABLE VII: Sample Audit Trail for Molecules Recorded on PureChain

Molecule ID	pIC_{50}	SMILES	Blockchain Tx Hash (truncated)
mol_1	7.99	CN(C)CCOC(=O)c1ccc2[nH]c(=O)ccc2c1	0xb49b8e…bce48c
mol_2	9.05	CN(C)CCOC(=O)c1cccc2nccc(=O)c12	0x03f25a…eada15
mol_3	6.62	Cc1ccc(cc1)C(=O)NC(C)C	0x5273b7…9bbddc

These results validate PureChain's ability to deliver blockchain-backed verifiability with minimal overhead, full audit traceability, and deterministic performance, all essential traits for scalable AI-driven drug discovery pipelines with cryptographic result guarantees.

While current benchmarking uses Amyloid-beta A4 to demonstrate proof-of-concept, the VeraComm protocol is agnostic to the molecular target. The pipeline can accommodate other targets by re-training or swapping the inference model, and descriptor filtering remains valid across diverse scaffolds. Future work will extend the evaluation to multi-target settings and further explore scalability in high-throughput virtual screening scenarios.

IV. CONCLUSION

This study introduces a verifiable and reproducible drug screening pipeline that combines AI-driven binding affinity prediction with blockchain-backed result validation. By integrating SVR-based regression models, Lipinski filtering, and PureChain-enabled hash verification into a CLI-controlled architecture, the system ensures deterministic processing, traceable audit logs, and end-to-end transparency. Experimental evaluation demonstrated strong predictive accuracy (R^2 = 0.78), consistent screening throughput (1.324 molecules/sec), and low-latency blockchain interaction (1.15 s logging, 245 ms verification), all with 100% success rates across 100 compounds. These findings validate the feasibility of embedding secure, tamper-evident mechanisms into virtual screening pipelines without compromising computational efficiency. The framework presents a scalable solution for trustworthy molecular discovery in high-stakes biomedical applications, setting the stage for future extensions involving smart contract automation, multi-target screening, and deployment in decentralized research ecosystems.

ACKNOWLEDGMENT

This work was partly supported by Innovative Human Resource Development for Local Intellectualization program through the IITP grant funded by the Korea government (MSIT) (IITP-2025-RS-2020-II201612, 33%) and by Priority Research Centers Program through the NRF funded by the MEST (2018R1A6A1A03024003, 33%) and by the MSIT, Korea, under the ITRC support program (IITP-2025-RS-2024-00438430, 34%).

REFERENCES

[1] S. O. Ajakwe, I. U. Ajakwe, T. Jun, D.-S. Kim, and J.-M. Lee, "Ciswqms: Connected intelligence smart water quality monitoring scheme," *Internet of Things*, vol. 23, p. 100800, 2023.

[2] S. O. Ajakwe, K. L. Olabisi, and D.-S. Kim, "Multihop intruder node detection scheme (minds) for secured drones' fanet communication," *IET Intelligent Transport Systems*, vol. 19, no. 1, p. e70080, 2025.

[3] K. L. Lanctôt, J. H. Hahn-Pedersen, C. Eichinger, C. Freeman, A. Clark, L. Tarazona, and J. Cummings, "Burden of illness in people with alzheimer's disease: a systematic review of epidemiology, comorbidities and mortality," *The journal of prevention of Alzheimer's disease*, vol. 11, no. 1, pp. 97–107, 2024.

[4] B.-l. Sun, Y. Chen, D.-y. Fan, C. Zhu, F. Zeng, and Y.-j. Wang, "Critical thinking on amyloid-beta-targeted therapy: challenges and perspectives," *Science China Life Sciences*, vol. 64, pp. 926–937, 2021.

[5] K. A. Carpenter and X. Huang, "Machine learning-based virtual screening and its applications to alzheimer's drug discovery: a review," *Current pharmaceutical design*, vol. 24, no. 28, pp. 3347–3358, 2018.

[6] R. Hoopes, H. Hardy, M. Long, and G. G. Dagher, "Sciledger: A blockchain-based scientific workflow provenance and data sharing platform," in *2022 IEEE 8th International Conference on Collaboration and Internet Computing (CIC)*. IEEE, 2022, pp. 125–134.

[7] V. I. Kanu, S. O. Ajakwe, and D. Kim, "Blockchain-verified deterministic framework for protein structure and binding site analysis in drug discovery distributed system," in *Proceedings of the Korean Institute of Communications and Information Sciences (KICS) Summer 2025 Conference*. Landing Shinhwa World, Jeju Island, South Korea: KICS, Jun. 2025, conference paper.

[8] D. E. Graff, E. I. Shakhnovich, and C. W. Coley, "Accelerating high-throughput virtual screening through molecular pool-based active learning," *Chemical science*, vol. 12, no. 22, pp. 7866–7881, 2021.

[9] G. Zhou, D.-V. Rusnac, H. Park, D. Canzani, H. M. Nguyen, L. Stewart, M. F. Bush, P. T. Nguyen, H. Wulff, V. Yarov-Yarovoy *et al.*, "An artificial intelligence accelerated virtual screening platform for drug discovery," *Nature Communications*, vol. 15, no. 1, p. 7761, 2024.

[10] ChEMBL, "Target: Amyloid-beta A4 protein (CHEMBL2487) — ebi.ac.uk," https://www.ebi.ac.uk/chembl/explore/target/CHEMBL2487, [Accessed 14-07-2025].

[11] I. U. Ajakwe, V. I. Kanu, S. O. Ajakwe, and D.-S. Kim, "eBCTC: Energy-Efficient Hybrid Blockchain Architecture for Smart and Secured K-ETS," *Cleaner Engineering and Technology*, 2025.