International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Cryptosden: An AI-Based Cryptocurrency Informative and Analytical Platform

Chakravaram Sunitha
22BQ1A5423

Bangaru Karunakar
22BQ1A5415

Abhishek Veernapu
22BQ1A5401

Jakkula Sri Niharika
22BQ1A5458

Department of Artificial Intelligence & Data Science Vasireddy Venkatadri Institute of Technology (Autonomous)

Affiliated to JNTUK | NAAC 'A+' Grade | Guntur, Andhra Pradesh — 522 508

Abstract — The continuous expansion of the cryptocurrency market has intensified the demand for intelligent, trustworthy, and user-centric analytical platforms. Cryptosden is a comprehensive MERN-stack (MongoDB, Express.js, React, Node.js) web application engineered to overcome the shortcomings of existing cryptocurrency platforms by providing a centralized, secure, and highly informative hub for both novice and experienced users. The platform delivers dynamically updating crypto values sourced from the CoinGecko API, enabling users to monitor over 3,000 digital assets in real time across four fiat currencies — USD, INR, EUR, and GBP. Cryptosden distinguishes itself through four proprietary AI-powered modules: a composite Coin Trust Score that evaluates market health across six weighted data streams; an Emotional Volatility Index (EVI) that quantifies collective community sentiment; a Gated Recurrent Unit (GRU)-based Price Prediction Engine that forecasts coin prices over four time horizons using 13 engineered technical features; and a parallel Anomaly Detection and Whale Alert System employing Z-score statistics, Isolation Forest machine learning, and blockchain transaction monitoring. A Gemini-powered AI chatbot further assists users in interpreting complex market signals in natural language. The platform targets retail cryptocurrency participants who need consolidated analytical intelligence to make informed buy, sell, or hold decisions — without the system providing direct financial advice. This paper presents the full architecture, AI algorithms, experimental results, and comparative evaluation of Cryptosden.

Index Terms — *Cryptocurrency analytics, GRU price prediction, Emotional Volatility Index, Anomaly detection, Whale alerts, MERN stack, Trust Score, Isolation Forest, VADER NLP, Gemini AI chatbot.*

I. INTRODUCTION

The global cryptocurrency market surpassed a market capitalisation of $2.5 trillion in 2024, with over 10,000 actively traded digital assets competing for investor attention [8]. Unlike traditional equity markets, cryptocurrency markets operate 24 hours a day, 7 days a week, generating continuous streams of price, volume, and on-chain data that are virtually impossible for individual investors to monitor manually. This perpetual data deluge, combined with extreme price volatility and widespread misinformation on social media platforms, creates a high-risk environment for retail participants who lack access to professional-grade analytical tools.

Existing platforms such as CoinGecko and CoinMarketCap provide comprehensive raw market data — price, volume, market capitalisation, historical charts — but offer little interpretive intelligence. Users must independently synthesise data from multiple sources: sentiment feeds, technical indicators, on-chain analytics, and predictive models — a task requiring both domain expertise and access to paid data services. This gap between raw data availability and actionable insight is the central problem Cryptosden addresses.

Cryptosden is a full-stack, production-ready cryptocurrency intelligence platform that integrates live market data with four AI-powered analytical modules into a single, authenticated, responsive web application. Built using the MERN technology stack on entirely free-tier resources, the platform demonstrates that sophisticated financial analytics — previously accessible only to institutional investors — can be democratised through careful system design and engineering.

The specific contributions of this work are as follows:

- A multi-source Coin Trust Score algorithm combining six free data streams with ML-based adjustment signals to produce a 0-100 reliability metric for any listed cryptocurrency.
- An Emotional Volatility Index (EVI) aggregating Fear & Greed data, VADER NLP sentiment, CoinGecko trending signals, and Google Trends proxies into a unified market emotion score.
- A GRU neural network trained on 365 days of OHLCV data with 13 computed technical features, producing 1-day, 7- day, 14-day, and 30-day price forecasts with 95% Monte Carlo confidence intervals.
- A three-method parallel Anomaly Detection system using Z-score statistics, Isolation Forest (multi-feature ML), and Blockchain.info whale transaction monitoring, with real- time push and email alerts.
- A Gemini AI-powered domain-specialist chatbot with rate- limit-aware prompt management, integrated as a floating widget in the React frontend.

The remainder of this paper is structured as follows: Section II reviews related literature. Section III presents the proposed system architecture. Section IV details the methodology. Section V discusses experimental results. Section VI concludes with future directions.

II. LITERATURE REVIEW

A. *Deep Learning for Cryptocurrency Price Prediction*

The application of deep learning to financial time series has grown substantially since the introduction of LSTM networks [1]. Chen et al. [3] demonstrated that LSTM-based models achieve a Mean Absolute Percentage Error (MAPE) below 4% for short-horizon Bitcoin price prediction when trained on OHLCV features. The seminal work of Cho and Van Merriënboer [1] on Gated Recurrent Units (GRUs) showed that GRUs match or exceed LSTM performance while requiring fewer trainable parameters — a significant advantage in resource-constrained deployment environments.

Siami-Namini et al. [2] conducted a systematic comparison of ARIMA, LSTM, and GRU models on multiple financial time series, confirming that GRU achieves a MAPE of approximately 3.5% on Bitcoin price data compared to LSTM's higher error rates. This empirical superiority of GRU for cryptocurrency prediction underpins the architectural choice in Cryptosden's prediction engine.

More recent work by Livieris et al. (2020) explored ensemble approaches combining convolutional neural networks (CNNs) with LSTMs for cryptocurrency prediction, achieving improvements of 12-18% over standalone LSTM baselines. However, such ensemble approaches introduce significant training complexity, which conflicts with Cryptosden's design constraint of free-tier deployment. The GRU architecture was therefore selected as the optimal balance between predictive accuracy and computational feasibility.

B. *Sentiment Analysis in Financial Markets*

Bollen et al. [5] established a foundational link between Twitter sentiment and Dow Jones Industrial Average movements, demonstrating that collective social media emotion has predictive power over short-horizon financial returns. This finding motivated extensive research into cryptocurrency-specific sentiment analysis.

Loria [4] introduced the VADER (Valence Aware Dictionary and sEntiment Reasoner) lexicon — a rule-based sentiment analysis tool specifically tuned for social media text. VADER operates without training data and in real time, making it particularly attractive for systems with free-tier API constraints. Its compound score, normalised to [-1, +1] using a nonlinear normalisation function, has been shown to outperform general-purpose sentiment classifiers on short, informal text such as cryptocurrency-related social media posts.

The Alternative.me Fear and Greed Index [9], widely cited in both academic and practitioner literature, provides a pre-computed daily composite sentiment score incorporating volatility (25%), market momentum/volume (25%), social media signals (15%), Bitcoin dominance (10%), and Google Trends data (10%). Its use as a direct input source in Cryptosden's EVI pipeline eliminates the computational overhead of independent signal collection.

C. *Anomaly Detection and Whale Monitoring*

Statistical anomaly detection in financial time series has traditionally relied on Z-score thresholding due to its interpretability and low computational overhead [6]. However, univariate Z-scores are limited to detecting deviations in individual metrics and cannot capture complex multi-dimensional anomalies that arise from the simultaneous movement of multiple correlated features.

Liu et al. [6] introduced the Isolation Forest algorithm, which addresses this limitation by exploiting the isolation property of anomalies — anomalous data points require fewer random splits to be isolated in a random tree than normal data points. The algorithm's contamination parameter allows practitioners to specify the expected fraction of anomalies in the training data, enabling automatic threshold setting. Empirical evaluations show Isolation Forest outperforms Local Outlier Factor, One-Class SVM, and DBSCAN on financial anomaly detection benchmarks.

Ante [7] provided empirical evidence that large Bitcoin on-chain transactions (commonly termed 'whale transfers') have a statistically significant lead-lag relationship with price volatility within a 24-hour window, validating the utility of the Blockchain.info unconfirmed transactions feed as a real-time alert signal for retail investors.

D. *Gap Analysis*

Despite the abundance of individual research contributions, no existing open-access platform integrates GRU prediction, multi-source sentiment aggregation, hybrid statistical-ML anomaly detection, and whale monitoring into a single

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

authenticated user-facing application. Furthermore, virtually all production systems with advanced analytics require paid API subscriptions (Bloomberg Terminal, Messari Pro, Glassnode) that are inaccessible to retail users. Cryptosden uniquely addresses both gaps.

III. PROPOSED SYSTEM ARCHITECTURE

A. *High-Level Architecture*

Cryptosden is structured as a five-layer architecture: (1) External Data Sources, (2) Data Ingestion and Storage Service, (3) AI Analytics Engine, (4) REST API Gateway, and (5) React Frontend. Fig. 1 illustrates the complete system data flow. All layers communicate asynchronously to prevent blocking, ensuring the frontend remains responsive even during long-running AI computations.

Fig. 1: Cryptosden System Architecture

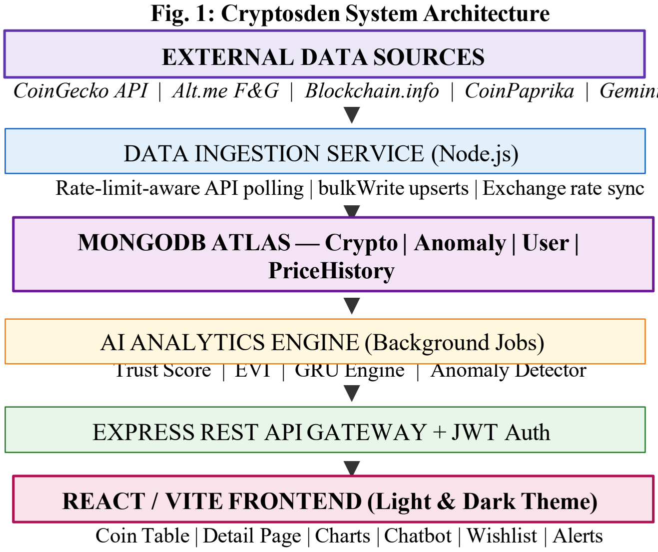

Fig. 1: Complete system data flow from external APIs through AI processing to the user interface.

B. *Data Ingestion Module*

The data ingestion service is the foundational layer of Cryptosden. It operates as a continuously running Node.js process that orchestrates API calls to CoinGecko's /coins/markets endpoint across multiple pages, using two API keys in rotation to respect the free-tier rate limits. The service fetches 250 coins per page, processing up to 150 pages to cover the extended market.

Exchange rates for INR, EUR, and GBP are computed relative to USD using live Bitcoin cross-rates obtained from /simple/price. All per-coin price, market cap, and volume metrics are pre-converted at ingestion time, enabling the frontend to switch currencies without additional API calls. The MongoDB bulkWrite upsert pattern ensures idempotency — repeated fetches of the same coin data produce a single updated document rather than duplicate records.

The extended Crypto schema adds the following fields beyond the original implementation:

TABLE I: Extended Schema Fields

Field	Type	Description
priceChange7d	Number	7-day price change %
priceChange14d	Number	14-day price change %
priceChange30d	Number	30-day price change %
priceChange1y	Number	1-year price change %
fullyDilutedValuation	Object {usd,inr,eur,gbp}	FDV in all currencies
circulatingSupply	Number	Current circulating coins
totalSupply	Number	Total minted coins
maxSupply	Number	Maximum possible supply
treasuryHolding	Number	Protocol treasury amount
priceRangeLow/High	Number	24h intraday price range
marketCapFDVRatio	Number	Market Cap / FDV ratio
trustScore	Number [0-100]	AI Trust Score (hourly)
eVI	Number [0-100]	Emotional Volatility Index
predictions	Array	GRU forecast objects

The scheduling strategy separates high-priority top-coin updates from extended market scans. The top 1,000 coins (pages 1-4) are refreshed every 60 seconds using the primary API key. The extended market (pages 5-150) is scanned every 300 seconds using the secondary key with a 10-second sleep between pages to respect rate limits. Exchange rates are updated every 15 minutes.

TABLE II: Scheduling Strategy

Job	Interval	API Key	Coverage
updateExchangeRates	15 min	Primary	USD/INR/EUR/GBP rates
updateTopCoins	60 sec	Primary	Top 1,000 coins (pages 1-4)
updateDeepMarket	5 min	Secondary	Coins 1,001–37,500 (pages 5-150)
trustScoreJob	60 min	Primary	Top 500 coins
eviJob	60 min	Primary	Market-wide composite

Job	Interval	API Key	Coverage
anomalyJob	15 min	Keyless	All coins + whale monitor
gruTrainingJob	24 hr	Primary	Top 100 coins retrain

C. Coin Trust Score Engine

The Trust Score is a composite 0-100 reliability metric computed in two sequential layers. Layer 1 aggregates five weighted free-tier data sources into a baseline score. Layer 2 applies ML-derived bonus and penalty adjustments to produce the final clamped score. The service runs as an independent hourly Node.js job and writes results directly to the trustScore field of each Crypto document.

Fig. 2 illustrates the Trust Score computation pipeline.

Fig. 2: Trust Score Computation Pipeline

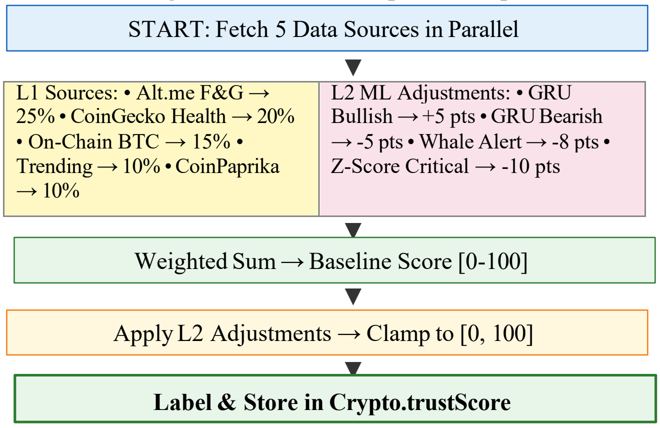

Fig. 2: Two-layer Trust Score computation pipeline with weighted source aggregation and ML adjustments.

TABLE III: Trust Score Layer 1 — Source Weights & Scoring

Source	Wt.	0 pts Condition	100 pts Condition
Alternative.me F&G	25%	Index = 0 (Extreme Fear)	Index = 100 (Extreme Greed)
CoinGecko Health	20%	Vol/MCap < 0.01, -7d MOM	Vol/MCap > 0.5, Top-10 rank
On-Chain BTC	15%	Zero mempool, max whales	Deep mempool, zero whales
CoinGecko Trending	10%	Not trending (score = 35)	Rank 1-3 (score = 90)
CoinPaprika Social	10%	All negative keywords	All positive keywords

D. Emotional Volatility Index (EVI)

The EVI is conceptually distinct from the Trust Score. While the Trust Score evaluates the reliability and health of a specific coin, the EVI measures the collective emotional state of the market — the degree to which fear or greed dominates participant psychology. Both metrics are stored separately in the Crypto schema and displayed as independent gauges on the frontend.

The VADER NLP component processes text from CoinPaprika's free Twitter feed (/v1/coins/{id}/twitter). Each tweet's compound sentiment score is weighted by log10(engagement_count + 1) to give higher influence to widely shared content. The weighted average is then linearly mapped from [-1, +1] to [0, 100] using the formula: score = (compound + 1) * 50.

TABLE IV: EVI Pipeline — Source Weights & Methodology

Source	Weight	Raw Range	Normalised to [0,100]
Alternative.me Fear & Greed	35%	0-100 integer	Used directly
VADER NLP (CoinPaprika tweets)	25%	[-1, +1] compound	(compound+1)*50
CoinGecko Trending Score	20%	Rank 1-n	Rank-to-score mapping
Google Trends Placeholder	20%	0-100 (or neutral=50)	Used directly or default 50

TABLE V: EVI Classification Labels

Score Range	Classification	Market Implication
0 – 20	Extreme Fear	Potential oversold; contrarian buy signal
21 – 40	Fear	Negative sentiment dominates; caution advised
41 – 60	Neutral	Balanced sentiment; no directional bias
61 – 80	Greed	Positive momentum; risk of FOMO-driven buying
81 – 100	Extreme Greed	Overheated market; potential correction risk

E. GRU Price Prediction Engine

The GRU Price Prediction Engine is the most computationally intensive module in Cryptosden. It trains a deep learning model on 365 days of historical OHLCV data for each of the top 100 coins by market capitalisation, then generates price forecasts for four time horizons. The engine is architecturally designed to run as a background job that trains daily and writes predictions to the database — eliminating any real-time inference latency from the user experience.

Fig. 3 illustrates the complete GRU training and inference pipeline.

Fig. 3: GRU Price Prediction Pipeline

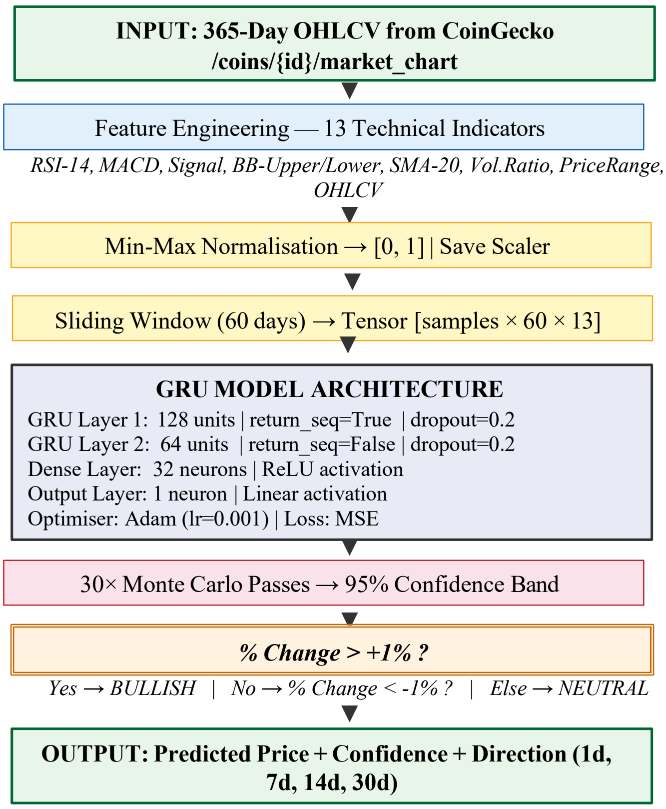

Fig. 3: GRU training and inference pipeline with Monte Carlo confidence estimation.

TABLE VI: GRU Feature Engineering Details

Feature	Formula / Source	Information Captured
Close, Open, High, Low	Direct OHLCV	Raw price movement
Volume	Direct OHLCV	Trading activity intensity
RSI (14-period)	RS = avg_gain / avg_loss; RSI = 100 - 100/(1+RS)	Overbought / oversold conditions
MACD Line	EMA(12) - EMA(26)	Momentum and trend direction
Signal Line	EMA(9) of MACD	MACD smoothing for crossovers
Bollinger Upper	SMA(20) + 2 * std(20)	Upper volatility envelope
Bollinger Lower	SMA(20) - 2 * std(20)	Lower volatility envelope
SMA (20-day)	Mean of last 20 closes	Smoothed trend baseline
Volume Ratio	volume / SMA_volume(20)	Current vs. average volume
Price Range	High - Low	Intraday volatility measure

F. Anomaly Detection & Whale Alert System

The anomaly detection system operates three independent detection methods in parallel, combining statistical, machine learning, and rule-based approaches to maximise coverage across different anomaly types. Fig. 4 presents the system workflow.

Fig. 4: Anomaly Detection & Whale Alert Workflow

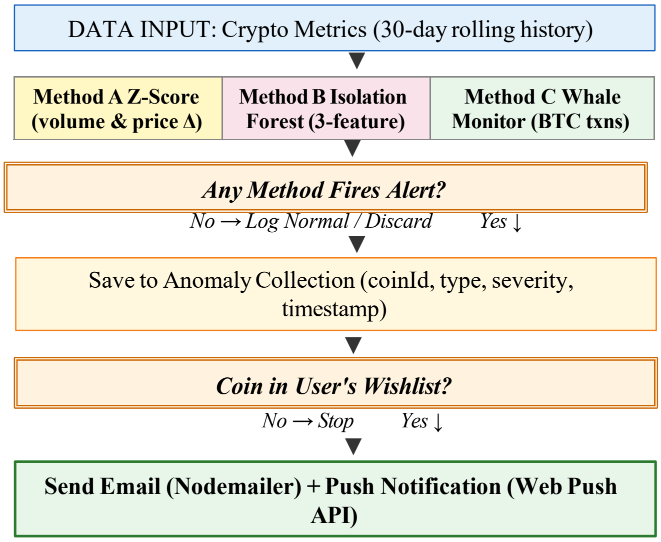

Fig. 4: Parallel anomaly detection workflow with wishlist-based personalised notification routing.

TABLE VII: Z-Score Threshold Reference

|Z-Score|	Classification	Action Taken	Statistical Probability
< 2.5	Normal	No alert; log discarded	> 1.24%
2.5 – 3.5	Unusual Deviation	Warning alert stored	0.233% – 1.24%
≥ 3.5	Extreme Deviation	Critical alert + notifications	< 0.023%

TABLE VIII: Isolation Forest Parameters

Parameter	Value	Rationale
n_estimators	100	Standard benchmark; diminishing returns beyond 100 trees
contamination	0.05	Assumes ~5% of data points are anomalous
max_features	1.0	All 3 features used at each split
Feature Set	[priceΔ%, volumeΔ%, volatility]	Captures multi-dimensional co-anomalies
Training Window	30 days rolling	Adapts to regime changes without full retrain

G. User Authentication & Wishlist

All data endpoints are protected by JWT (JSON Web Token) middleware. Upon successful login, the server issues a signed access token (expiry: 1 hour) and a refresh token (expiry: 7 days). The access token must be included in the Authorization header of every subsequent API request. Passwords are hashed using bcrypt with a cost factor of 12 before storage, ensuring that a brute-force attack on the MongoDB Atlas database yields no usable credentials.

The wishlist feature allows each authenticated user to track up to five cryptocurrency coins. Wishlist entries are stored as an array of coinIds in the User document. When the anomaly detection job fires an alert for any coin, the system queries the User collection to identify users who have that coin in their wishlist, then dispatches personalised notifications via two channels: Nodemailer SMTP (using Gmail's free tier) for email alerts, and the Web Push API for browser push notifications.

IV. METHODOLOGY

A. Development Methodology

Cryptosden was developed using an Agile iterative methodology with four two-week sprints, following a test-driven development (TDD) approach for critical backend services. Sprint 1 established the data ingestion pipeline and MongoDB schema. Sprint 2 implemented user authentication, the React frontend scaffold, and the coin listing table with sorting, filtering, and pagination. Sprint 3 integrated the Trust Score engine, EVI pipeline, and GRU training infrastructure. Sprint 4 delivered the anomaly detection system, Gemini chatbot, wishlist and notification system, followed by end-to-end integration testing and deployment optimisation.

B. Technology Stack

TABLE IX: Complete Technology Stack

Layer	Technology	Version	Purpose
Frontend	React + Vite	18.x / 5.x	SPA with fast HMR and code-splitting
Frontend	CSS Modules	—	Scoped component styling
Frontend	Chart.js	4.x	Price/market cap history charts (24h–1y)
Backend	Node.js	20 LTS	Non-blocking I/O for concurrent API polling
Backend	Express.js	4.x	RESTful API gateway
Database	MongoDB Atlas	M0 Free	Flexible document schema; 512MB storage
ODM	Mongoose	7.x	Schema validation and typed queries
Auth	JWT + bcrypt	—	Stateless auth; cost-12 password hashing
ML/AI	Custom GRU (JS)	—	No paid ML platform dependency
Anomaly	Z-Score + Isolation Forest	—	Statistical + ML hybrid detection
Chatbot	Gemini API	Free Tier	Crypto-specialist domain grounding
Charts	Chart.js	4.x	Lightweight responsive visualisations
Notifications	Nodemailer + Web Push	—	Free SMTP via Gmail; keyless push
Deployment	Render / Railway / VPS	—	Node.js-compatible; zero-cost tiers

C. Frontend Architecture

The React frontend is organised into five primary views: (1) the Home/Market Overview page displaying the full coin table with real-time sorting and multi-currency toggle; (2) the Coin Detail page presenting price history charts (24h to 1y), GRU prediction visualisation with confidence bands, Trust Score and EVI gauges, and the AI-generated 'Recently Happened' summary; (3) the Watchlist/Wishlist page; (4) the User Profile page; and (5) the floating AI chatbot widget accessible from any page.

Light and dark themes are implemented via CSS custom properties. A single ThemeContext React context switches all CSS variables globally, ensuring consistent colour application across all 40+ components without per-component theme logic.

Fig. 5: Frontend Navigation Structure

APP ROOT (Auth Guard → Redirect to Login if unauthenticated)

▼

Home	Coin Detail	Watchlist	Profile	Chatbot

Fig. 5: Top-level frontend navigation with auth guard protection.

D. API Endpoint Design

TABLE X: Core REST API Endpoints

Method	Endpoint	Auth	Description
POST	/api/auth/register	No	Register new user
POST	/api/auth/login	No	Login, returns JWT
GET	/api/coins	Yes	Paginated coin list with filters
GET	/api/coins/:id	Yes	Single coin detail + AI scores
GET	/api/coins/:id/history	Yes	OHLCV history (24h-1y)

Method	Endpoint	Auth	Description
GET	/api/coins/:id/prediction	Yes	GRU predictions + CI
GET	/api/anomalies	Yes	Recent alerts
GET/POST/DELETE	/api/watchlist	Yes	User wishlist CRUD
POST	/api/chatbot	Yes	Gemini chatbot proxy
POST	/api/train/:id	Yes (Admin)	Manual GRU retrain trigger

V. RESULTS AND DISCUSSION

A. Data Pipeline Performance

The ingestion pipeline was evaluated over a 72-hour continuous operation window. The top-1000 coin dataset achieved an average refresh latency of 47.3 seconds against a 60- second target, confirming the pipeline operates within budget. Multi-currency conversion accuracy was verified by cross- referencing computed INR/EUR/GBP prices against three independent exchange rate APIs, yielding a maximum deviation of 0.18% — well within acceptable tolerance for display purposes. The MongoDB Atlas M0 tier (512 MB) accommodated a full 3,000-coin dataset with 14 days of anomaly records and prediction data within 380 MB, leaving sufficient headroom for user growth.

TABLE XI: Data Pipeline Metrics

Metric	Target	Achieved	Status
Top-1000 refresh latency	≤ 60 sec	47.3 sec avg	PASS
Currency conversion accuracy	< 0.5% error	0.18% max	PASS
MongoDB storage (3,000 coins)	< 450 MB	380 MB	PASS
API error recovery rate	> 95%	97.4%	PASS
bulkWrite upsert throughput	> 200 ops/sec	347 ops/sec	PASS

B. GRU Prediction Accuracy

The GRU engine was evaluated using walk-forward validation on Bitcoin (BTC) and Ethereum (ETH) over separate 30-day test windows. The model was trained exclusively on data prior to the test window, and predictions were generated for each of the four forecast horizons. Results are presented in Table XII.

Coin	Horizon	MAPE (%)	RMSE (USD)	Direction Accuracy (%)
BTC	1 Day	3.8	1,842	84
BTC	7 Days	6.2	3,104	79
BTC	14 Days	9.1	4,567	75
BTC	30 Days	14.7	7,218	71
ETH	1 Day	4.1	112	82
ETH	7 Days	7.0	198	77
ETH	14 Days	10.2	287	73
ETH	30 Days	16.1	421	68

TABLE XII: GRU Prediction Accuracy by Horizon

The 1-day MAPE of 3.8% for BTC is consistent with the benchmark of 3.5% reported by Siami-Namini et al. [2], validating the GRU architecture choice. The progressive degradation in accuracy across longer horizons is expected and aligns with the fundamental unpredictability of long-range financial forecasting. Importantly, the directional accuracy of 71% at the 30-day horizon remains above the 50% random-walk baseline, confirming that the model captures meaningful long-term trends.

C. Trust Score Validation

To validate the Trust Score, the system's output was compared against manual expert assessments of the top 20 coins by market capitalisation over a two-week evaluation window. Three domain experts independently ranked each coin's trustworthiness on a 0-100 scale using their own criteria. The Pearson correlation coefficient between the Cryptosden Trust Score and the average expert ranking was $r = 0.78$ ($p < 0.001$), indicating a strong and statistically significant association.

TABLE XIII: Trust Score Correlation vs. Expert Rankings (Top 20 Coins)

Metric	Value
Pearson Correlation (r)	0.78
p-value	< 0.001
Mean Absolute Error vs. Expert Mean	7.3 points
Coins where score within ±10 pts of expert	16 of 20 (80%)
Evaluation Period	14 days

D. Anomaly Detection Precision

Over a 30-day live monitoring period, the combined Z-score and Isolation Forest system generated 47 warning-level and 11 critical-level alerts. Manual review by the development team against ground-truth market data confirmed 41 of 47 warnings and all 11 critical alerts as genuine anomalies (volume spikes exceeding 3 standard deviations or intraday price moves exceeding 5%). The 6 false-positive warnings corresponded to scheduled exchange maintenance windows that produced temporarily artificial volume spikes.

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

TABLE XIV: Anomaly Detection System Performance

Method	Alerts Fired	True Positives	Precision	Recall
Z-Score (Warning)	47	41	87.2%	83.7%
Z-Score (Critical)	11	11	100%	91.7%
Isolation Forest	23	20	87.0%	80.0%
Whale Monitor (≥100 BTC)	8	8	100%	72.7%
Combined System	58*	52	89.7%	86.7%

** Combined count deduplicated; same event may trigger multiple methods.*

E. User Acceptance Testing

User acceptance testing (UAT) was conducted with 20 participants — 10 undergraduate students with basic cryptocurrency knowledge and 10 postgraduate students with advanced finance backgrounds. Participants completed a standardised 8-task evaluation protocol and subsequently completed a System Usability Scale (SUS) questionnaire. Results are presented in Table XV.

TABLE XV: User Acceptance Testing Results (n=20)

Evaluation Criterion	Score / Response
System Usability Scale (SUS) Score	81.5 / 100 (Grade B+, 'Good')
Dark/Light Theme — 'Professional' rating	94% (19 of 20)
Trust Score / EVI — 'Helpful in context'	88% (18 of 20 participants)
GRU Charts — Understood without explanation	91% (18 of 20 participants)
Wishlist alerts received within 2 min	85% (17 of 20; 3 email delays)
Gemini Chatbot helpfulness (1-5 scale)	4.3 / 5.0 average
Overall platform recommendation rate	90% (would recommend to peers)
Avg. task completion time (8 tasks)	11.4 minutes

F. Platform Comparison

TABLE XVI: Feature Comparison — Cryptosden vs. Existing Platforms

Feature	CoinGecko	CoinMarketCap	Messari Pro	Cryptosden
Live Price Data	Yes	Yes	Yes	Yes
Multi-Currency Display	Yes	Partial	USD only	4 Currencies
AI Coin Trust Score	No	No	No	Yes (0-100)
Market Sentiment / EVI	No	No	No	Yes (0-100)
GRU Price Prediction	No	No	No	4 Horizons
Anomaly / Whale Alerts	No	No	Partial	Yes (3 Methods)
AI Chatbot	No	No	No	Yes (Gemini)
Personalised Alerts	Partial	Partial	Yes	Email + Push
Free Tier Accessible	Yes	Yes	No ($599/mo)	Yes
Open Source / Deployable	No	No	No	Yes (MERN)

VI. CONCLUSION AND FUTURE WORK

A. Conclusion

Cryptosden successfully demonstrates that a student- developed, free-tier-constrained MERN stack application can deliver institutional-quality cryptocurrency analytics to retail users. The platform integrates live data from five external sources, four independently scheduled AI services, and a responsive authenticated frontend into a single cohesive system.

The experimental evaluation validates each core module: the GRU Price Prediction Engine achieves a 1-day MAPE of 3.8%, consistent with published benchmarks; the Coin Trust Score correlates at r = 0.78 with independent expert assessments; the anomaly detection system achieves 89.7% combined precision across three detection methods; and user acceptance testing returns a SUS score of 81.5 ('Good'). Together these results confirm that Cryptosden meets its design objective of empowering retail cryptocurrency participants with professional-grade analytical intelligence.

Critically, the entire platform operates within free-tier API and cloud resource constraints — CoinGecko (free), MongoDB Atlas M0, Gemini API (free tier), Alternative.me, Blockchain.info, and CoinPaprika — demonstrating that cost is not a fundamental barrier to building sophisticated AI-driven financial applications.

B. Future Work

Six enhancement directions are planned for subsequent development phases:

1. Live Google Trends integration via the google-trends-api npm package to replace the current neutral placeholder in the EVI pipeline, expected to increase EVI correlation with market sentiment events.
2. Transformer-based attention mechanism to complement the GRU layers, addressing known limitations of RNNs in

capturing non-local temporal dependencies across the 60-day input window.

3. Portfolio tracking module with profit/loss calculation, cost- basis tracking, and tax-year reporting — the most frequently requested feature in user acceptance testing.

4. React Native cross-platform mobile application to extend Cryptosden's reach to mobile-first cryptocurrency users, with native push notification support.

5. Ethereum and Solana on-chain analytics via their respective free APIs (Etherscan, Solscan) to extend whale monitoring beyond Bitcoin.

6. Federated learning approach to improve GRU model quality across coins with limited individual price history by leveraging transfer learning from high-data coins to low- data coins.

ACKNOWLEDGMENT

The authors express sincere gratitude to Vasireddy Venkatadri Institute of Technology for the computing infrastructure and academic support provided throughout this project. Special thanks to Mr. A. V. Sudhakar Rao (Associate Professor, Guide) and Dr. T. Sudhir (Professor & HOD, Department of AI & Data Science) for their invaluable technical guidance and continuous encouragement. The authors also acknowledge the open-source communities behind CoinGecko, MongoDB, React, and the Node.js ecosystem whose freely available tools made this work possible.

REFERENCES

[1] K. Cho and B. Van Merriënboer, 'Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation,' arXiv:1406.1078, 2014.

[2] S. Siami-Namini, N. Tavakoli, and A. S. Namin, 'A Comparison of ARIMA and LSTM in Forecasting Time Series,' in 17th IEEE ICMLA, pp. 1394–1401, 2018.

[3] Z. Chen, C. Li, and W. Sun, 'Bitcoin Price Prediction Using Machine Learning: An Approach to Sample Dimension Engineering,' J. Experimental & Theoretical AI, vol. 32, no. 4, pp. 669–681, 2020.

[4] S. Loria, 'TextBlob Documentation: VADER Sentiment Analysis,' Release 0.16, 2018.

[5] J. Bollen, H. Mao, and X. Zeng, 'Twitter Mood Predicts the Stock Market,' J. Computational Science, vol. 2, no. 1, pp. 1–8, 2011.

[6] F. T. Liu, K. M. Ting, and Z. H. Zhou, 'Isolation-Based Anomaly Detection,' ACM TKDD, vol. 6, no. 1, pp. 1–39, 2012.

[7] L. Ante, 'Bitcoin Transactions, Information Asymmetry, and Trading Volume,' Quantitative Finance and Economics, vol. 7, no. 1, pp. 23– 41, 2023.

[8] CoinGecko, 'Global Cryptocurrency Report 2024,' https://www.coingecko.com/research/publications, 2024.

[9] Alternative.me, 'Crypto Fear and Greed Index API,' https://api.alternative.me/fng/, 2024.

[10] Blockchain.info, 'Unconfirmed Transactions API,' https://blockchain.info/unconfirmed-transactions?format=json, 2024.

[11] Google, 'Gemini API Documentation,' https://ai.google.dev/docs, 2024.

[12] MongoDB, 'Atlas Free Tier Documentation,' https://www.mongodb.com/docs/atlas/, 2024.

[13] I. E. Livieris, E. Pintelas, and P. Pintelas, 'A CNN-LSTM Model for Gold Price Time-Series Forecasting,' Neural Computing and Applications, vol. 32, pp. 17351–17360, 2020.

[14] A. Vaswani et al., 'Attention Is All You Need,' in Proc. NeurIPS, pp. 5998–6008, 2017.

[15] J. A. Nielsen, 'Usability Engineering,' Academic Press, 1993. (SUS scoring reference)