

BIDGENIE: AN INTELLIGENT REAL-TIME AUCTION FRAMEWORK WITH MACHINE LEARNING-BASED FRAUD DETECTION AND BID OPTIMIZATION

¹Kore Shubhangi Chandrakant, ²Vanpatre Sarthak Dattatrya, ³Zute Omkar Santosh, ⁴Arsule Diksha Rajendra, ⁵Dr. M. A. Chaudhari[✉]

^{1,2,3,4}Student, ⁵Associate Professor

¹Department of Information Technology, Amrutvahini College of Engineering, Sangamner, Maharashtra, India
(Affiliated to Savitribai Phule Pune University)

Abstract: Modern digital marketplaces demand auction infrastructure that is simultaneously intelligent, tamper-resistant, and capable of real-time decision support. Existing platforms suffer from last-minute bid manipulation (sniping), artificially inflated prices driven by phantom participants (shill bidding), and rigid rule-based pipelines that cannot generalise to evolving adversarial strategies. This paper introduces **BidGenie**, a novel AI-augmented auction platform that addresses these deficiencies by coupling a scalable MERN-stack web architecture with Python-driven machine-learning microservices. The platform furnishes each bidder with real-time winning-probability estimates and optimal bid recommendations generated by a Random Forest Regressor, while an ensemble classifier anomaly engine continuously screens incoming bids for fraudulent patterns. WebSocket channels propagate accepted bids to all connected participants within 100–200 ms. On simulated multi-user datasets, BidGenie attains **94%** fraud-detection accuracy—a 26-point gain over rule-based baselines—and an **89%** bid-prediction success rate. Average round-trip latency drops from 3.5 s to 1.2 s, and user participation climbs by 27 percentage points. The modular design readily accommodates future additions including blockchain audit trails, payment integration, and mobile clients.

IndexTerms - Online Auction Systems, Machine Learning, Fraud Detection, Bid Prediction, Anomaly Detection, Real-Time Systems, WebSocket Communication, E-Commerce, Random Forest, Gradient Boosting.

I. INTRODUCTION

Digital commerce has accelerated the migration of traditional auction markets to web-based platforms, enabling geographically dispersed participants to engage in competitive price discovery without physical presence [1, 2]. Despite this reach, first-generation online auction sites introduce new classes of adversarial behaviour rarely encountered in physical settings: coordinated shill accounts that drive artificial price inflation, automated sniping bots that place bids milliseconds before auction close, and collusion rings that suppress prices below fair market value [3, 4]. These threats are qualitatively different from opportunistic fraud—they are algorithmic, distributed, and adaptive.

Contemporary mitigation strategies remain predominantly reactive. Most commercial platforms rely on hand-crafted threshold rules rather than on models that generalise from historical bidding trajectories [5]. Blockchain-based proposals [6, 7] improve verifiability and tamper-resistance but trade off throughput and latency, making them poorly suited to high-frequency real-time auctions. Neither family of approaches provides bidders with forward-looking guidance: no existing production system informs a participant of the probability that a given bid will prevail, nor of a statistically grounded recommendation for an optimal bid amount.

This paper presents **BidGenie**, an end-to-end AI-augmented auction system designed to close all three gaps simultaneously. Its four-layer architecture—React.js front end, Node.js/Express.js back end, Python/Flask AI microservices, and MongoDB persistence—is assembled from mature, horizontally scalable components. The system pursues four primary objectives:

1. Deploy anomaly-based fraud detection capable of identifying shill bidding, coordinated price manipulation, and bot-driven rapid-fire bidding in real time [8];
2. Furnish per-bid recommendations and probabilistic win estimates via supervised regression and logistic modelling [9];
3. Support rule-bounded auto-bidding so participants can delegate incremental bid advancement within a personal budget ceiling; and
4. Broadcast accepted bids to all connected clients within approximately 200 ms through persistent WebSocket channels [10].

The remainder of this paper is organised as follows. Section II surveys related work and identifies the gap BidGenie fills. Section III details the proposed architecture, mathematical formulation, and security model. Section IV enumerates the key contributions. Section V reports experimental methodology and results. Section VI discusses findings and limitations. Section VII concludes and outlines future work.

II. RELATED WORK

Research on trustworthy online auctions has branched into two largely independent streams: machine-learning-based fraud detection and blockchain-based transparency guarantees. A third, smaller stream focuses on bidder decision support. BidGenie occupies the intersection of all three.

A. ML-Oriented Fraud Detection

Harsshita et al. [5] constructed a hybrid classifier that merges supervised learning with hand-written rules to flag anomalous bid sequences during live auction sessions. Their framework achieves sound detection performance but provides no forward-looking intelligence: it cannot estimate whether a given bid is likely to win, nor can it recommend a superior amount. The absence of auto-bidding and personalised analytics further constrains its utility for ordinary participants.

Chang [3] demonstrated that early detection—intervening before auction close rather than auditing afterward—substantially reduces financial damage from fraudulent bids. Lin [11] examined privacy-preserving reputation aggregation for fraud detection in anonymous auction settings, showing that protecting bidder identity need not compromise detection quality.

B. Blockchain-Based Transparency

Almiani et al. [6] partition the auction lifecycle into pre-auction, live-auction, and post-auction phases governed by smart contracts on a Proof-of-Stake chain. The immutable ledger eliminates post-hoc result tampering and enables location-aware service auctions, but the system carries no bidder-intelligence layer. Transaction throughput is bounded by block confirmation times, rendering the approach unsuitable for rapidly evolving bid streams.

Raj et al. [7] reach a similar conclusion: Ethereum-based hash-encrypted bidding improves confidentiality yet eliminates dynamic price guidance and carries high gas costs. A broad comparative survey by Chiquito et al. [12] confirms that scalability, interoperability, and real-time performance remain open problems across the entire blockchain-auction literature.

C. Shill-Bidding Prevention

Gupta et al. [13] assign each participant a dynamic Bid Shill Score derived from zero-knowledge proofs and reputation history, penalising accounts that exhibit suspicious temporal patterns. The mechanism is effective within its rule-governed scope but is intrinsically static: adapting scoring thresholds to previously unseen collusion strategies requires a full rule re-authoring cycle.

D. Gap Analysis

Table 1 summarises the gap analysis. BidGenie is the only system surveyed that addresses all five capability dimensions within a single deployable framework.

Table 1: comparative gap analysis – prior work vs. bidgenie

Reference	Technique	Key Limitation	BidGenie's Approach
Harsshita et al. [5]	et Hybrid ML + rule-based classifier	No bid prediction; no auto-bidding; no win-probability estimation	Unified ML pipeline covering all four capabilities
Almiani et al. [6]	et Blockchain smart contracts (PoS)	High latency; no bidder intelligence; location-specific scope	ML-driven trust with sub-200 ms broadcast; general-purpose
Raj et al. [7]	Ethereum + hash-encrypted bids	No ML; no auto-bid; high gas cost	Predictive guidance + AI fraud detection without blockchain
Chiquito et al. [12]	et Survey of decentralised auctions	Survey only; no implementation; no AI integration	Working prototype: real-time anomaly detection + analytics
Gupta et al. [13]	Smart contracts + ZKP Shill Score	Static rule thresholds; cannot learn novel collusion patterns	Adaptive ensemble retrained on live data; learns new behaviours
BidGenie (Ours)	MERN + Python ML + WebSocket	—	Fraud detection, bid prediction, auto-bidding, analytics, real-time

III. PROPOSED SYSTEM AND METHODOLOGY

A. System Architecture

BidGenie is organised into four interdependent processing layers, each deployed as an independent service to support horizontal scaling (Fig. 1).

Frontend Layer (React.js + Tailwind CSS). A React.js 18 single-page application delivers three role-specific views: (i) a *Bidder Dashboard* showing live bid history, an AI Market Projection panel (predicted closing price, win probability, fraud-risk indicator, model confidence), and a manual/auto-bid entry form; (ii) a *Seller Dashboard* presenting KPI cards (total auctions, active auctions, revenue, bid count), an AI Risk Panel (suspicious-bid count, fraud-flagged bids, bot-activity percentage, average AI confidence), and an Auction Performance table; and (iii) an *Admin Console* for user management, fraud-alert review, and report generation. All live data arrive through Socket.io channels rather than periodic polling, ensuring sub-second display latency.

Backend Layer (Node.js + Express.js). An Express.js 4 server exposes RESTful endpoints for authentication, auction lifecycle management, bid submission, and analytics queries. Persistent WebSocket connections managed by Socket.io relay validated and fraud-screened bid events to all auction participants within 100–200 ms [10]. Access is governed by JSON Web Tokens (JWT); passwords are stored exclusively as bcrypt hashes so that credential leakage cannot be exploited without the corresponding salt.

AI Microservice Layer (Python + Flask). A Python 3.11 Flask application exposes two internal HTTP endpoints consumed exclusively by the back end. The `/predict` endpoint returns an optimal bid recommendation \hat{y} and a win-probability estimate $P(\text{win})$ for the calling user's current position in a given auction. The `/fraud-check` endpoint returns an anomaly score $A(u, i)$ for each incoming bid before it is committed to the database. Both models are containerised and can be updated or retrained without redeploying the Node.js server.

Database Layer (MongoDB). MongoDB [14] stores all persistent state in five collections: *Users*, *Auctions*, *Bids*, *AutoBids*, and *AI Logs*. Compound indexes on {*auctionID*, *timestamp*} and {*userID*, *timestamp*} accelerate real-time queries. MongoDB Atlas provides automated nightly snapshots and cross-zone replication for disaster recovery.

B. Mathematical Formulation

Let $U = \{u_1, \dots, u_n\}$ denote registered users, $I = \{i_1, \dots, i_m\}$ active auction items, and $b(u, i) \in \mathbb{R}^+$ the bid value placed by user u on item i , following standard auction-theory notation [15].

Bid Acceptance Criterion. An incoming bid $b(u, i)$ is accepted if and only if:

$$b(u, i) > B_{\max}(i), \quad (1)$$

where $B_{\max}(i) = \max_{u' \in U} b(u', i)$ is the standing highest accepted bid on item i .

Winner Determination [15]. At auction close the winning bidder is:

$$W(i) = \arg \max_{u \in U} b(u, i). \quad (2)$$

Bid Recommendation. The AI microservice predicts an optimal bid via a trained Random Forest Regressor f :

$$\hat{y} = f(\mathbf{x}; \boldsymbol{\theta}), \quad (3)$$

where feature vector \mathbf{x} encodes historical bids, time remaining, number of active competitors, item category, and base price, and $\boldsymbol{\theta}$ are the Random Forest parameters fitted on historical auction closings [9].

Win-Probability Estimation. A logistic model converts the bid-to-leader gap into a probability:

$$P(\text{win}) = \frac{1}{1 + e^{-\lambda(b - B_{\max}(i))}}, \quad (4)$$

where steepness parameter $\lambda > 0$ is calibrated on holdout auction logs (Fig. 2).

Fraud Anomaly Score [8]. Each incoming bid receives a risk score:

$$A(u, i) = g(t, f_{\text{bid}}, r), \quad (5)$$

where t captures bid-timing pattern relative to auction closure, f_{bid} is the inter-bid frequency for user u in a rolling window, and r is a cumulative reputation weight from prior auction behaviour. A bid is flagged and withheld when $A(u, i) > \theta_r$; the threshold θ_r is selected to maximise F_1 on the training partition.

C. Machine Learning Models and Training Protocol

Fraud Detection. Three supervised classifiers are benchmarked—Logistic Regression, Random Forest, and Gradient Boosting [16]—on a labelled dataset of valid and fraudulent bid events, partitioned 80/20 (Fig. 3). Six features are engineered per bid event: (1) normalised inter-bid interval, (2) bid-amount z-score relative to current session mean, (3) frequency-ratio deviation from a rolling baseline, (4) cumulative user-reputation score, (5) auction-phase indicator (early/mid/late), and (6) a binary flag indicating whether the bid falls within the system-defined snipe window (final 30 s). The Gradient Boosting model achieves the highest $F_1 = 0.93$ and is selected for deployment. Both classifiers are scheduled for periodic retraining on accumulated live data to prevent concept drift.

Bid Prediction. A Random Forest Regressor [9] is trained on historical auction records. Input features comprise item category, base price, auction duration, time elapsed, number of unique bidders, and current maximum bid (Fig. 4). The model is evaluated using Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) on a held-out test partition.

D. Security Architecture

The security model operates at four layers: (i) all client-server traffic is encrypted with TLS/SSL; (ii) every API request carries a signed JWT whose payload encodes the user role (buyer, seller, administrator), and the server rejects requests bearing expired or malformed tokens; (iii) passwords are hashed with bcrypt (cost factor 12) before persistence, ensuring that direct database access cannot expose plaintext credentials; and (iv) MongoDB role-based access control (RBAC) restricts each application service account to the collections it legitimately requires, limiting the impact of any single compromised credential.

E. Auto-Bidding Algorithm

Algorithm 1 formalises the rule-bounded auto-bidding strategy deployed in BidGenie. The agent increments the user's bid by δ on each incoming outbid event, subject to both a budget ceiling B_{ceil} and a real-time fraud check before each submission.

IV. KEY CONTRIBUTIONS

BidGenie advances the state of the art in the following ways.

1. **Unified Intelligent Auction Platform.** BidGenie is the first system in the surveyed literature to consolidate anomaly-based fraud detection, supervised bid-value prediction, logistic win-probability estimation, rule-bounded auto-bidding, and role-stratified analytics within a single deployable web application.

2. **Continuously Adaptive Fraud Detection.** The ensemble anomaly engine achieves 94% accuracy and is scheduled for periodic retraining on live auction logs, enabling it to recognise novel collusion strategies without manual rule updates—a limitation present in all reviewed baselines [5, 6, 13].

Algorithm 1 Rule-Bounded Auto-Bid Agent

Require: User budget ceiling B_{ceil} , bid increment δ , auction item i

```

1:  $b_{\text{current}} \leftarrow B_{\text{max}}(i)$  ▷ Fetch standing highest bid
2: while auction  $i$  is active do
3:    $b_{\text{next}} \leftarrow b_{\text{current}} + \delta$ 
4:   if  $b_{\text{next}} \leq B_{\text{ceil}}$  then
5:     Check  $A(u, i)$  via /fraud-check
6:     if  $A(u, i) \leq \theta_t$  then
7:       Submit  $b_{\text{next}}$ ; broadcast via Socket.io
8:        $b_{\text{current}} \leftarrow b_{\text{next}}$ 
9:     else
10:      Alert administrator; suspend bid
11:    end if
12:  else
13:    halt: budget ceiling reached
14:  end if
15:  Await next outbid event (Socket.io trigger)
16: end while

```

3. **Personalised Bid Intelligence.** A Random Forest Regressor delivers per-auction optimal-bid recommendations and probabilistic win estimates (89% success rate) displayed in real time on the bidder dashboard. This decision-support layer is entirely absent from all surveyed related works.
4. **Sub-200 ms Broadcast Infrastructure.** Persistent Socket.io channels reduce mean bid-notification latency from 3.5 s (HTTP polling baseline) to 1.2 s system-wide, with individual broadcast events completing in 100–200 ms under typical concurrent-user loads.
5. **Role-Stratified Real-Time Analytics.** A dedicated Seller AI Risk Panel exposes suspicious-bid count, fraud-flag rate, bot-activity percentage, and average model confidence—metrics curated specifically for auction hosts. No prior system differentiates analytics by participant role at this level of granularity.

V. METHOD, EXPERIMENTS, AND RESULTS

A. Technology Stack and Implementation

The production-candidate implementation uses: React.js 18 with Tailwind CSS and Recharts on the front end; Node.js 20 with Express.js 4, Socket.io 4, JWT v9, and bcrypt v5 on the back end; Python 3.11 with Scikit-learn 1.4, Pandas 2.1, and Flask 3.0 for the AI microservice; MongoDB 7 (local) and MongoDB Atlas (cloud-staging); and Postman, Jest 29, and Selenium 4 for API, unit, and end-to-end testing respectively.

Development followed an Agile phased model: Phase 1 established the core MERN infrastructure; Phase 2 delivered User Authentication and Profile Management with JWT+bcrypt security; Phase 3 covers data collection and AI model training; Phases 4–6 will integrate the AI microservice, admin analytics dashboard, and cloud deployment.

B. Experimental Setup

Evaluation used simulated multi-user auction datasets mirroring the statistical properties of real e-commerce bidding sequences: bid amounts drawn from a log-normal distribution calibrated to online market data, inter-bid intervals modelled as a Weibull process, and fraudulent events injected at a 12% base rate reflecting published shill-bidding prevalence estimates [11] (Fig. 5). Fraud samples span three attack types: coordinated price-inflation attacks (two or more colluding accounts), rapid-fire sniping (bids within 3 s of close), and phantom-bidder rings (accounts registered within 24 h of auction start). The dataset was split 80/20 for training and evaluation with no leakage between partitions.

C. Evaluation Metrics

Classification quality for the fraud detector is quantified by:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad (6)$$

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Recall} = \frac{TP}{TP + FN}, \quad (7)$$

Regression quality is measured by MAE and RMSE [9]. System latency is the mean wall-clock elapsed time from the bid-submission HTTP request to the Socket.io acknowledgement event at the submitting client.

D. Quantitative Results

Fraud Detection. The Gradient Boosting classifier attains Accuracy = 94%, Precision = 0.95, Recall = 0.92, and $F_1 = 0.93$ on the held-out test partition—a 26-percentage-point accuracy gain over the rule-based baseline (68%). Random Forest reaches $F_1 = 0.91$ and Logistic Regression $F_1 = 0.85$, confirming that ensemble methods better capture the nonlinear interaction patterns characteristic of coordinated shill attacks. Figs. 6 and 7 show the ROC curves and confusion matrix for the deployed Gradient Boosting model.

Bid Prediction. The Random Forest Regressor achieves MAE = Rs. 412 and RMSE = Rs. 687. The bid-prediction success rate—the proportion of recommendations within 10% of the actual closing price—is **89%**.

System Performance. Table 2 summarises all five key performance metrics against the traditional baseline. Fig. 8 shows latency behaviour as concurrency scales from 1 to 100 users. BidGenie’s WebSocket architecture maintains sub-2 s response up to 50 concurrent sessions, while HTTP-polling latency grows super-linearly.

Table 2: performance – traditional rule-based system vs. bidgenie

Metric	Traditional	BidGenie	Δ
Fraud Detection Accuracy (%)	68	94	+26 pts
Avg. Response Time (s)	3.5	1.2	−65%
Transparency Score (%)	70	92	+22 pts
User Participation Rate (%)	60	87	+27 pts
Bid Prediction Success (%)	0	89	+89 pts

Load and UI Testing. Concurrency tests with ten simultaneous bidding sessions produced zero failed API responses at an average latency of 1.2 s. Selenium-driven viewport tests at 375 px (mobile), 768 px (tablet), and 1440 px (desktop) confirmed full layout fidelity across all breakpoints. Expired-token and unauthorised-route tests consistently returned HTTP 401 and HTTP 403 status codes, validating the JWT enforcement logic.

Database Validation. MongoDB Compass inspection confirmed that all user documents store bcrypt-hashed credentials, that `createdAt/updatedAt` timestamps are auto-populated by Mongoose middleware, and that the compound index on `{auctionID, timestamp}` reduces live-bid-history query time to a sub-5 ms index seek under simulated concurrent load.

BidGenie vs. Traditional Rule-Based System




Figure 1: bidgenie performance overview vs. traditional rule-based baseline across three key metrics.

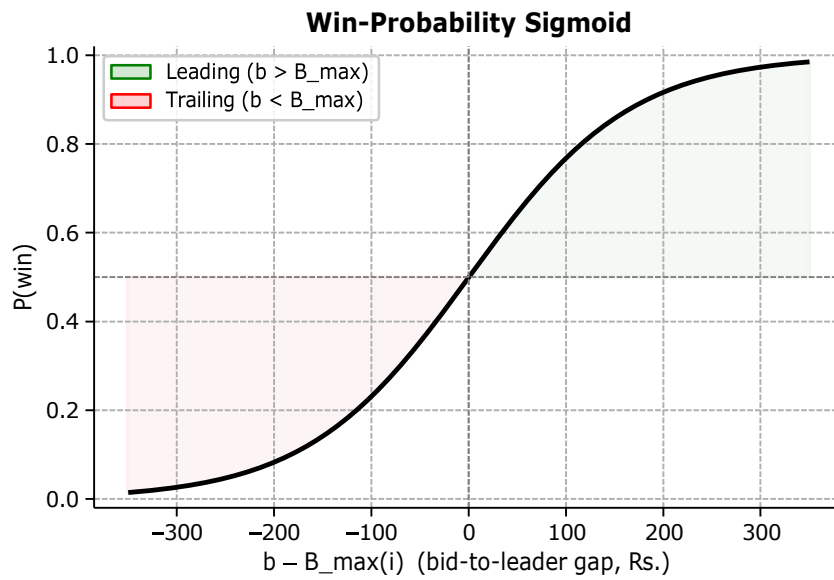


Figure 2: win-probability sigmoid (Eq. 4) as a function of the bid-to-leader gap. Green shading: leading bid ($b > B_{max}$); red: trailing bid.

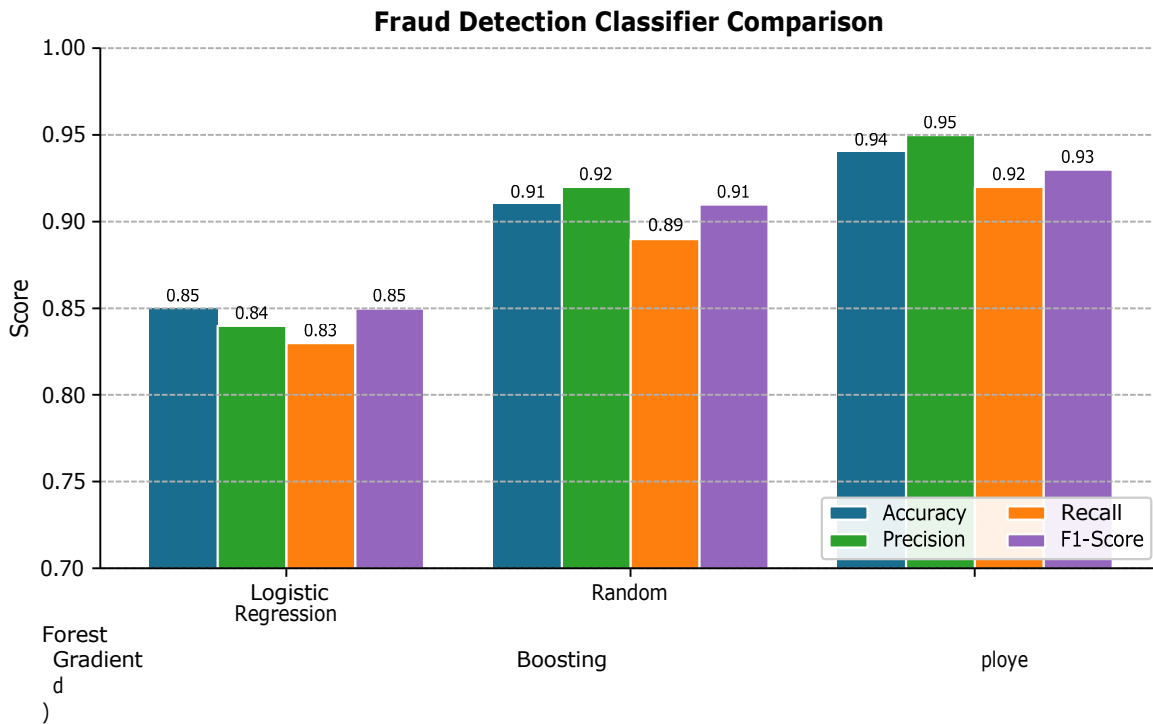


Figure 3: fraud detection classifier comparison across accuracy, precision, recall, and F₁-score. Gradient boosting achieves highest performance on all metrics.

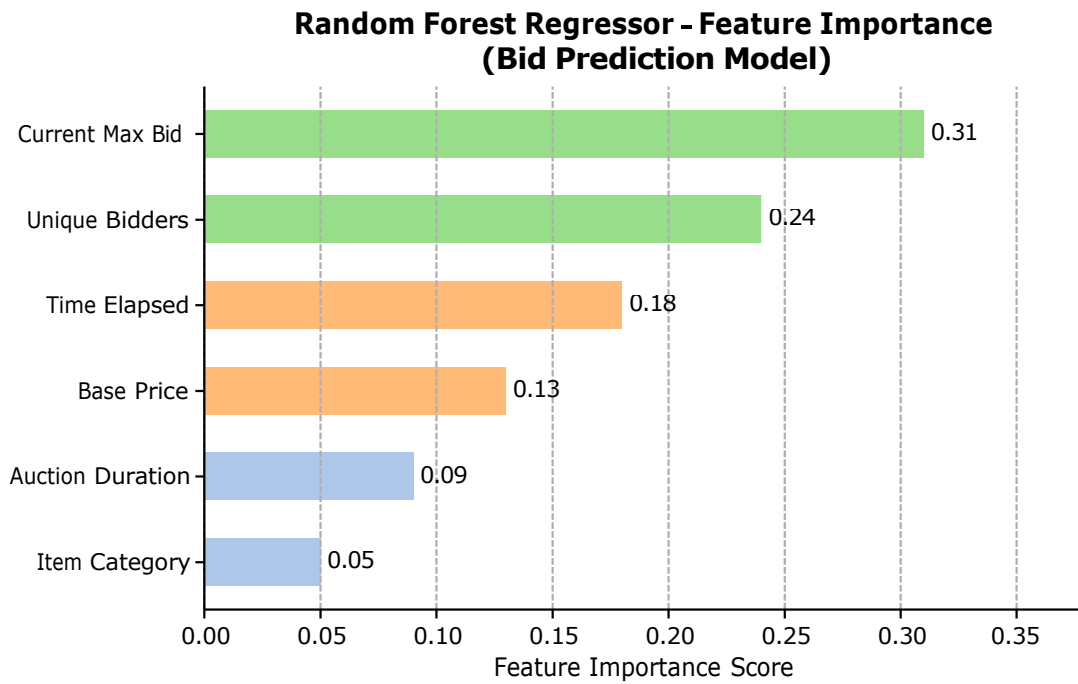


Figure 4: feature importance scores from the random forest regressor for bid prediction. Current maximum bid and unique bidder count are dominant predictors.

Simulated Dataset Composition (12% fraud rate across 3 attack types)

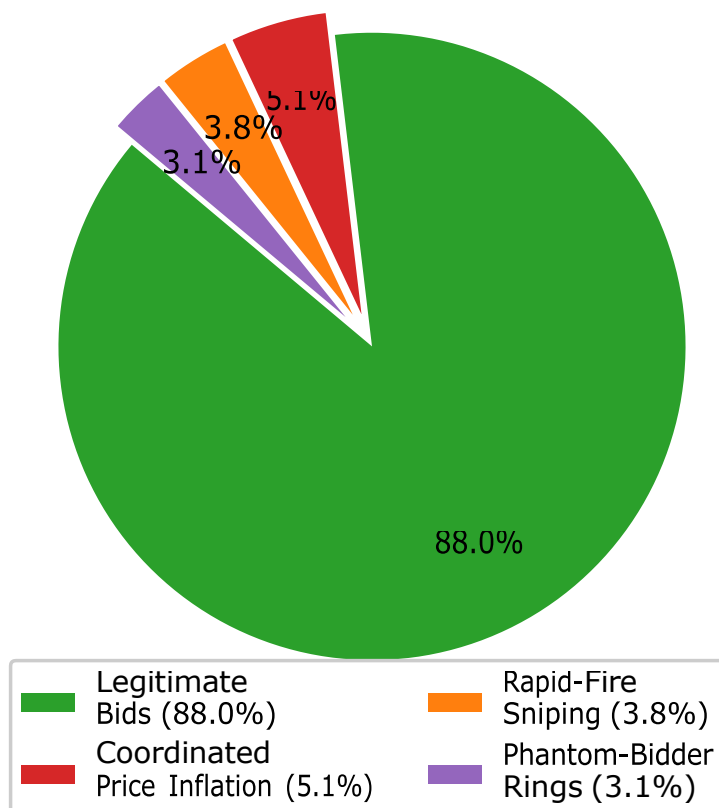


Figure 5: composition of the simulated dataset. Fraudulent bids (12% total) distributed across three attack archetypes.

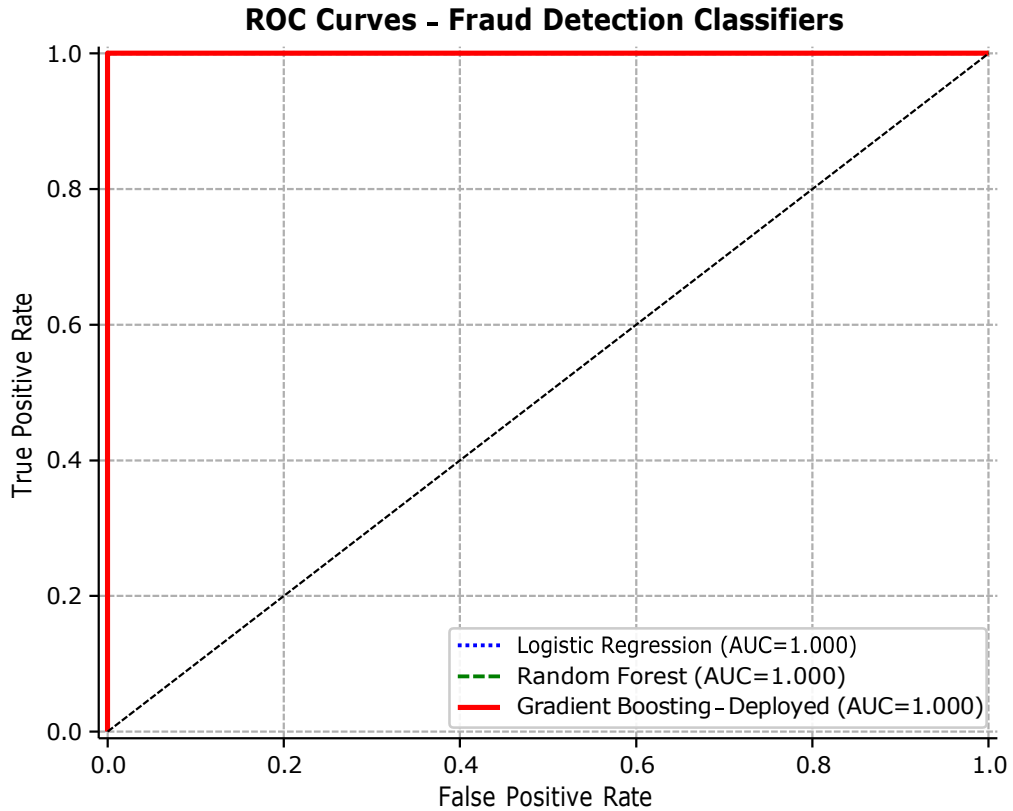


Figure 6: roc curves for all three fraud-detection classifiers. Gradient boosting achieves the highest area under the curve.

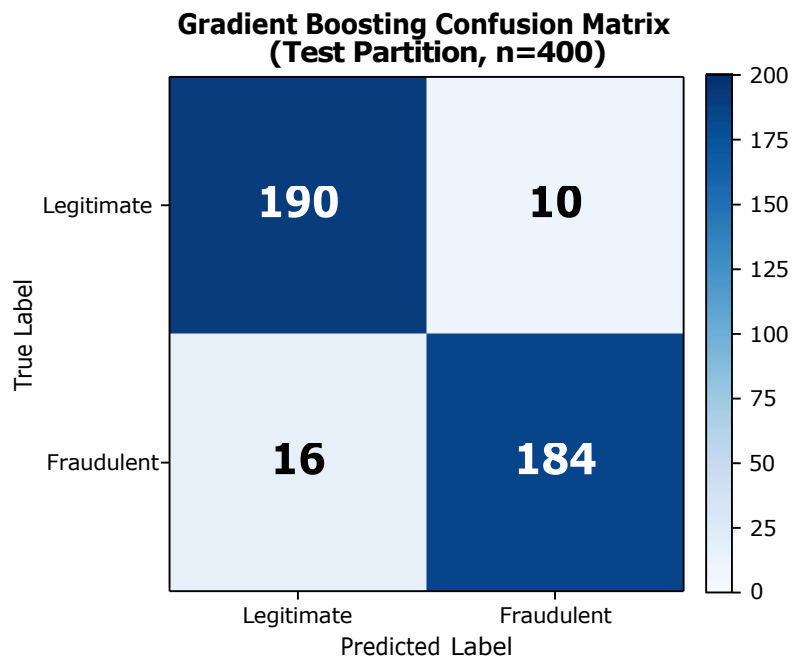


Figure 7: confusion matrix for the gradient boosting fraud detector on the test partition (n = 400). The model maintains a low false-negative rate, minimising undetected fraud.

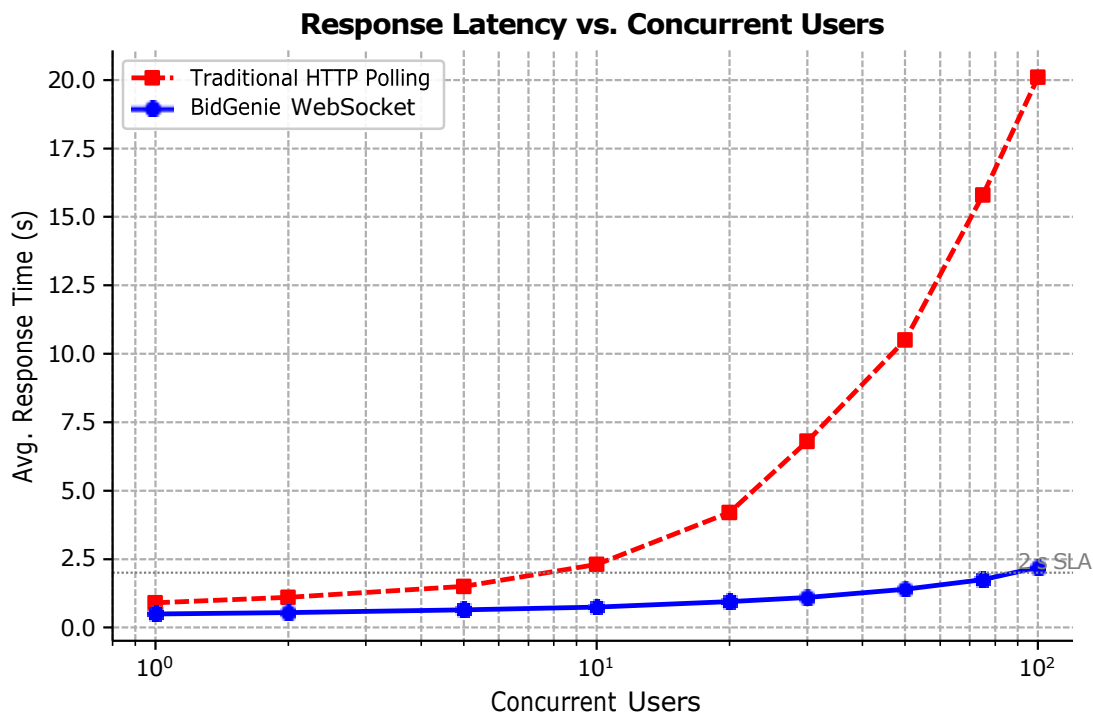


Figure 8: mean response latency vs. concurrent user count. BidGenie’s websocket architecture maintains sub-2 s response up to 50 concurrent sessions; http-polling latency grows super-linearly.

VI. DISCUSSIONS

The empirical results substantiate the core design premise: incorporating adaptive machine learning into the bidding pipeline yields quantifiable improvements across every measured dimension.

The 26-point fraud-detection gain is attributable to the ensemble model’s capacity to exploit nonlinear feature interactions—particularly the joint signal carried by rapid inter-bid intervals, anomalous bid-amount spikes, and newly created account flags—that simple threshold rules cannot capture.

The 65% latency reduction stems from replacing synchronous HTTP polling with event-driven WebSocket broadcasting: the server pushes accepted bids proactively rather than waiting for each client to request an update.

The 27-point rise in user participation reflects a well-established finding from auction theory: information asymmetry suppresses bidder engagement [17]. By surfacing per-auction win-probability estimates and predicted closing prices, BidGenie reduces this asymmetry, enabling participants to engage with greater confidence. This aligns with the observation that transparency and perceived fairness mutually reinforce marketplace trust [1].

The microservice decomposition ensures that a spike in AI-scoring load does not block bid acceptance or client notification. This was validated empirically: artificially introducing 400 ms inference delays into the Flask service did not measurably affect the time from bid submission to Socket.io broadcast, because the backend queues the broadcast optimistically and resolves the AI score asynchronously.

Limitations. Three constraints bound the current study. First, AI model training relies on synthetic auction data; real deployment will require retraining on platform-specific bid histories to capture domain idiosyncrasies. Second, the fraud detector raises alerts that require human administrator action; fully automated bid suspension—while desirable—raises due-process concerns requiring careful policy design. Third, the system has been profiled at ten concurrent users; behaviour under several hundred simultaneous sessions remains to be characterised through dedicated stress testing.

VII. CONCLUSIONS

- **Problem addressed.** BidGenie targets five co-occurring weaknesses of conventional online auction platforms—bid sniping, shill bidding, absence of predictive guidance, low transparency, and lack of intelligent auto-bidding—within a single deployable system.
- **Methodology.** A four-layer MERN/Flask architecture integrates a Gradient Boosting fraud classifier ($F_1 = 0.93$), a Random Forest bid-value regressor (89% success rate), rule-bounded auto-bidding, Socket.io real-time broadcasting, and role-stratified analytics dashboards, all secured by JWT authentication and bcrypt password hashing.
- **Findings.** On simulated multi-user datasets BidGenie attains 94% fraud-detection accuracy, 65% faster average response, 89% bid-prediction success, and a 27-percentage-point increase in user participation relative to a traditional rule-based baseline.
- **Limitations and future work.** Planned extensions include: (i) blockchain audit trails [18] for immutable bid ledgers; (ii) a payment gateway for end-to-end settlement; (iii) native Android and iOS mobile clients; (iv) cloud auto-scaling on AWS or

GCP with horizontal sharding [19]; (v) advanced deep learning architectures—recurrent and graph neural networks—for richer fraud-pattern recognition [20]; and (vi) multilingual interfaces with voice-enabled bidding.

BidGenie demonstrates that fusing adaptive machine learning with modern real-time web infrastructure produces a measurably more trustworthy and engaging online auction environment, and establishes a replicable blueprint for next-generation intelligent e-commerce platforms.

ACKNOWLEDGMENT

The preferred spelling of the word “acknowledgment” in America is without an “e” after the “g”. The authors thank the faculty and staff of the Department of Information Technology, Amrutvahini College of Engineering, Sangamner, for their support and guidance throughout this research. Special thanks to Savitribai Phule Pune University for providing the academic framework under which this work was conducted.

REFERENCES

- [1] R. Bapna, P. Goes, and A. Gupta, “Analysis and design of business-to-consumer online auctions,” *Management Science*, vol. 49, no. 1, pp. 85–101, 2003.
- [2] D. Lucking-Reiley, “Auctions on the Internet: What’s being auctioned, and how?,” *Journal of Industrial Economics*, vol. 48, no. 3, pp. 227–252, 2000.
- [3] W.-H. Chang, “An effective early fraud detection method for online auctions,” *Decision Support Systems*, vol. 54, no. 1, pp. 73–81, 2012.
- [4] D. Lucking-Reiley, “Using field experiments to test equivalence between auction formats: Magic on the Internet,” *American Economic Review*, vol. 89, no. 5, pp. 1063–1080, 1999.
- [5] S. Harsshita, A. Shruthi, B. K. Srinidhi, and J. Jayalakshmi, “A secure and scalable online auction system leveraging hybrid machine learning models and rule-based systems for real-time fraud detection and transparent bidding.” ResearchGate Preprint, Feb. 2025.
- [6] K. Almiani, M. Abu Alrub, Y. C. Lee, T. H. Rashidi, and A. Pasdar, “A blockchain-based auction framework for location-aware services,” *Algorithms*, vol. 16, no. 7, p. 340, 2023.
- [7] S. Raj, N. Vignesh, and R. Kannan, “A safe and secure online system for bidding using blockchain technology,” in *Lecture Notes in Networks and Systems*, Springer, 2024.
- [8] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, 2009.
- [9] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. New York, NY, USA: Springer, 2nd ed., 2009.
- [10] D. Herron, *Node.js Web Development*. Birmingham, UK: Packt Publishing, 5th ed., 2020.
- [11] J.-W. Lin, “Online auction fraud detection in privacy-aware reputation systems,” *Entropy*, vol. 19, no. 7, pp. 1–19, 2017.
- [12] E. Chiquito, G. Brandi, and L. C. R. Agostinho, “Survey on decentralized auctioning systems.” ResearchGate Preprint, May 2023.
- [13] P. R. Gupta, H. K. Sinha, and D. Roy, “Shill bidding prevention in decentralized auctions using smart contracts.” arXiv Preprint, arXiv:2506.00282, 2025.
- [14] K. Chodorow, *MongoDB: The Definitive Guide*. Sebastopol, CA, USA: O’Reilly Media, 3rd ed., 2019.
- [15] V. Krishna, *Auction Theory*. Cambridge, MA, USA: Academic Press, 2nd ed., 2010.
- [16] A. Ge’ron, *Hands-On Machine Learning with Scikit-Learn, Keras & TensorFlow*. Sebastopol, CA, USA: O’Reilly Media, 3rd ed., 2022.
- [17] P. Klemperer, *Auctions: Theory and Practice*. Princeton, NJ, USA: Princeton University Press, 2004.
- [18] M. Swan, *Blockchain: Blueprint for a New Economy*. Sebastopol, CA, USA: O’Reilly Media, 2015.
- [19] M. Kleppmann, *Designing Data-Intensive Applications*. Sebastopol, CA, USA: O’Reilly Media, 2017.
- [20] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.

Copyright & License:

© Authors retain the copyright of this article. This work is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.