UDC 004.056
IRSTI 81.93.29

https://doi.org/10.55452/1998-6688-2026-23-3-255-266

[1]Aidynov T.,
doctoral student, Researcher, ORCID ID: 0009-0005-3327-9719,
e-mail: tolegen.ch@gmail.com
[1]Satybaldina D.,
Cand. Sci. (Phys.-Math.), Professor, ORCID ID: 0000-0003-0291-4685,
e-mail: satybaldina_dzh@enu.kz
[2] Dimitrov W.,
PhD, Professor, ORCID ID: 0000-0002-5348-4948,
e-mail: v.dimitrov@unibit.bg
[1*]Abisheva G.,
PhD, ORCID ID: 0009-0003-6692-197X,
e-mail: gulsipat.abisheva@gmail.com

[1]L.N. Gumilyov Eurasian National University, Astana, Republic of Kazakhstan
[2]University of Library Studies and Information Technologies, Sofia, Bulgaria

A BLOCKCHAIN-ORIENTED MODEL FOR ELECTRONIC VOTING WITH BIOMETRIC VOTER IDENTIFICATION AND DEEPFAKE-RESISTANT AUTHENTICATION

Abstract

This paper presents and experimentally validates an electronic voting model that stores results in a decentralized blockchain while using biometric voter verification and automatic detection of synthetically generated (deepfake) facial images. For identity verification, Vision Transformer and ResNet-50 architectures are fine-tuned with Low-Rank Adaptation (LoRA). Frequency-domain information from the Fast Fourier Transform (FFT) is also included as a descriptor for detecting synthetic images. Since blockchain records cannot be altered, the model introduces a controlled revoting process based on self-destructing ballots and timestamps recorded on the blockchain; the votes themselves are submitted through an Ethereum smart contract. In the experiments, the LoRA-adapted ViT configuration distinguishes real from synthetic facial images with an accuracy of 97.48%, exceeding the performance of the LoRA-adapted ResNet-50 baseline. However, because no separate FFT ablation study has been conducted, this work does not claim that the frequency descriptor makes an independent contribution. The proposed system is presented as part of Smart City digital infrastructure and lays the groundwork for future research on privacy-preserving biometric protocols and blockchain platforms capable of scaling to electoral processes.

Keywords: electronic voting, face recognition, deepfake detection, blockchain, Ethereum, smart contracts, Vision Transformer, ResNet, revoting mechanism, smart city.

Received July 1, 2026; revised August 4, 7, 19, 2026; accepted August 28, 2026.

Introduction

Electronic voting has evolved from mechanical lever and punch-card machines, through direct-recording electronic systems, to fully web-based systems that let voters participate in elections remotely. The early direct-recording systems reduced the logistics of paper ballots but did not address the problem of verifiability: without an independent audit trail, such systems could still be manipulated through software or suffer technical failure.

The shift to internet voting made elections more accessible-Estonia's national digital elections offer a prime example-but it also introduced new risks, including remote cyberattacks, voter impersonation, and data breaches [1–4]. Research has therefore focused on two related aims: (a) authenticating voters and (b) protecting results from later alteration. Voter authentication remains especially challenging in the presence of synthetic imagery. Conventional identity checks, including passwords and document verification, are widely considered susceptible to phishing and social engineering. Face-based biometric verification has emerged as an appealing alternative since it is contactless and supported by consumer hardware that is already widely available. Still, face recognition by itself does not close the security gap.

Generative adversarial networks can now create photorealistic synthetic faces that fool even state-of-the-art verification models. For biometric authentication systems used in elections, a separate deepfake-detection stage has therefore become practically necessary if they claim to withstand targeted identity-spoofing attacks [4–8].

Blockchain as a guarantee of result integrity. A second source of weakness in conventional e-voting systems is centralized storage under a single trusted operator, which creates a single point of failure and an opportunity for undetectable tampering. Blockchain offers an alternative: a decentralized, cryptographically tamper-evident ledger that is open for independent verification by any party [9–11].

Permissioned ledgers such as Hyperledger Fabric were used in earlier proposals to maintain control over access while keeping transparency [12, 13]. Public Ethereum-based systems take this property one step further to full external verifiability and remove the requirement for a centralized vote-counting authority altogether by smart contracts paired with wallet-based authentication (e.g., MetaMask) [14–16].

The immutability problem and the aim of this study. Blockchain ledgers are trusted because their records cannot be changed, but that same property rules out legitimate vote corrections–something that can routinely be needed after an input error or a change of mind during the voting window [17, 18]. This study therefore seeks to design and experimentally validate a system with two capabilities: robust biometric authentication that can withstand deepfake attacks, and controlled revoting that does not weaken the cryptographic guarantees of the underlying ledger. The latter is handled by a Dynamic Revoting Mechanism based on self-destructing ballots and timestamp ordering recorded on the blockchain. It extends our earlier systematic review of cryptographic protection methods for electronic voting systems (reference [19]), a prior publication by the present authors and thus a self-citation. To our knowledge, no previously evaluated prototype has combined frequency-aware, deepfake-resistant biometric gating and a smart-contract-level, self-destructing-ballot revoting mechanism within a single pipeline. Each component is established on its own; this work's contribution lies specifically in integrating them with revoting logic executed on-chain.

Materials and methods

The proposed solution integrates three layers: a biometric verification layer built on deep neural models, a synthetic-image detection layer based on frequency-domain analysis, and a result-recording layer implemented as a revoting-capable Ethereum smart contract. The end-to-end processing sequence for a single vote is shown in Figure 1.

Dynamic Revoting Mechanism

Because this component determines the practical usability of the whole system, its logic is described first, ahead of the recognition-model details.

Self-destructing ballots. Once a blockchain transaction is confirmed, it is permanent by design, which normally makes it impossible to invalidate – and therefore replace – a previously cast vote. The proposed mechanism introduces a cryptographic "burn" function for the prior ballot: on receiving a revote request, the smart contract checks that (i) the voter already has a registered vote and (ii)

the voting period is still open, after which it marks the earlier record as invalid. The sequence is as follows:

1. The voter casts an initial vote, recorded on-chain as VoteHash_A.
2. The voter requests a revote, triggering the burn function.
3. VoteHash_A is marked invalid and excluded from the tally.
4. A new vote is recorded as VoteHash_B.
5. Only the most recent valid record is included in the final count.

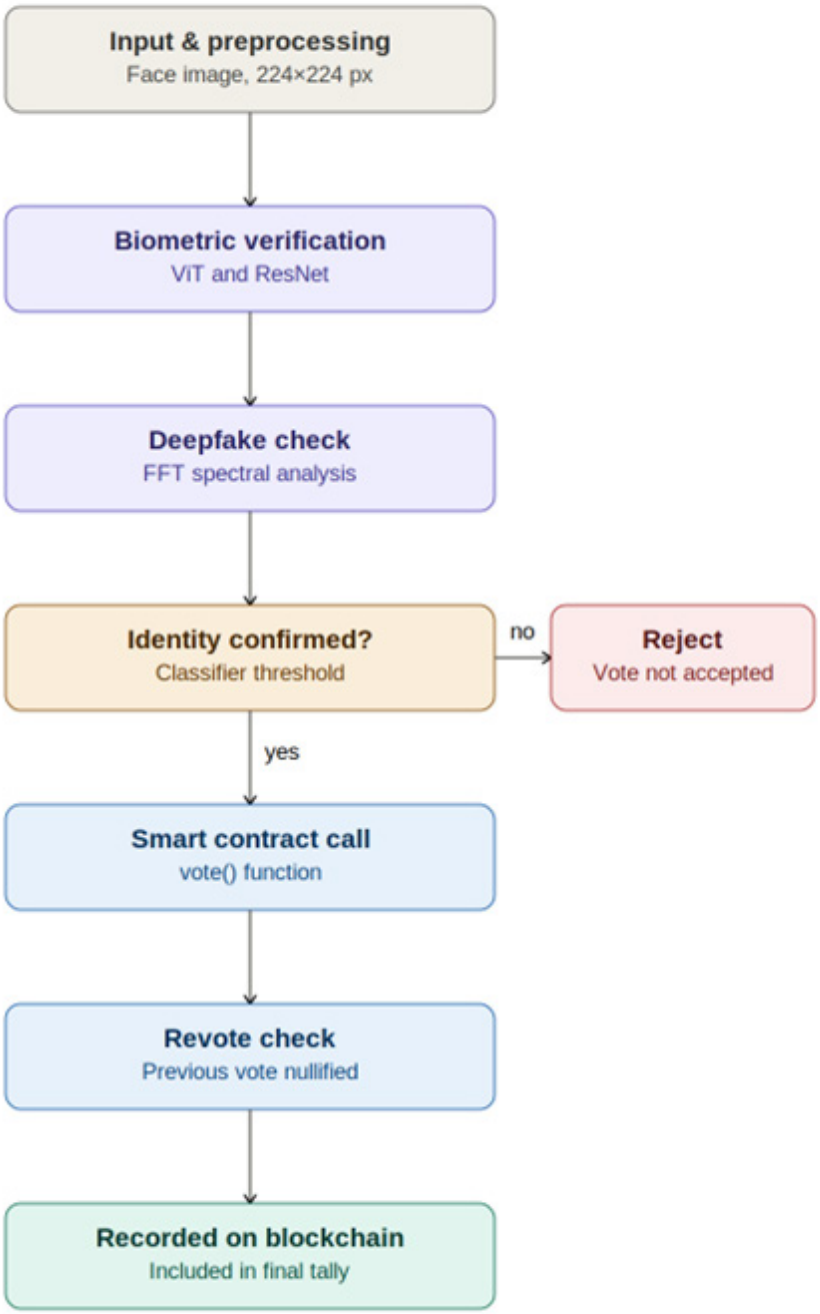

Figure 1 – Overall vote-processing workflow with biometric verification

The counted vote is formally given by expression (1):

$$n = argmax_i(T_i), V_t = V_n \quad (1)$$

where denotes the timestamp attached to the i-th recorded vote, n refers to the submission with the latest admissible timestamp, and represents the final vote included in the count. Since the smart contract allows only one revote during each election cycle, a voter may submit no more than two votes (n≤2). Burned ballots remain visible on-chain for auditing, but they no longer influence the tally.

Admissible submissions are ordered according to the timestamp attached to each vote by the Ethereum execution environment. That timestamp determines the sequence within the officially defined voting window, but it is not regarded as an independently trusted source of time.

$$T_n = max_i(T_i) \quad (2)$$

where denotes the time at which the i-th vote was recorded, and is the latest admissible timestamp. Submissions made outside the officially defined voting window are rejected by the smart contract. Because Ethereum block timestamps are proposed by block producers/validators, they may experience bounded drift or be influenced by the producer. The model therefore treats blockchain time as adequate for ordering transactions within the voting window, not as a cryptographic guarantee of exact wall-clock time. In a production deployment, an accepted timestamp-drift policy should be specified; stronger time guarantees would require an external trusted timestamping/oracle mechanism.

Conditions enforced by the smart contract. Three constraints are enforced at the contract level: a revote is accepted only within the designated time window; burned records are permanently excluded from the count; and each voter is limited to a single revote per election cycle. These rules together make it impossible to change the outcome by repeatedly submitting a new vote without breaking the cryptographic guarantees of the contract (Figure 2).

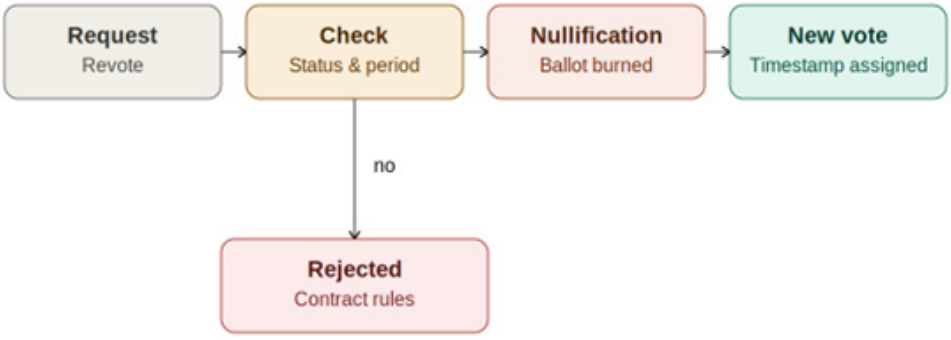

Figure 2 – Revote request validation and execution sequence

Table 1 summarizes how this combination of self-destructing ballots and secure timestamps removes the central limitation of immutable ledgers – the inability to correct a vote – without sacrificing transparency or protection against repeat voting.

Biometric Verification and Synthetic-Image Detection

Training data. Model tuning and evaluation are based on a publicly available dataset of 140,000 facial images (Kaggle, “140k Real and Fake Faces”), half of which are real photos and the other half are StyleGAN3-generated images. The images were resized to 224x224 pixels and randomly flipped horizontally prior to training to augment them and make them more robust to lighting and pose variations (Table 2).

Table 1 – Comparison between a conventional immutable blockchain ledger and the proposed mechanism

Criterion	Limitation of the conventional scheme	Proposed solution
Vote correction	Records are immutable; correction is impossible	Controlled update within the voting window
Double-voting prevention	Multiple submissions are not inherently excluded	Self-destructing ballots leave exactly one valid vote
Storage footprint	Grows with every submitted version of a vote	Burned records remain on-chain; they are excluded only from the active tally/ index
Transparency and tamper resistance	Direct overwriting of a vote creates a manipulation risk	Full operation history stays on-chain; only the latest timestamp counts

Table 2 – Characteristics of the dataset used

Parameter	Value
Total images	140,000
Genuine faces	70,000
Synthetic (deepfake) faces	70,000
Source	Kaggle, 140k Real and Fake Faces
Forgery generation method	StyleGAN3
Preprocessing	Resize to 224×224, color jitter, random flip

Recognition model architectures

The voter identity verification system is built on a Vision Transformer, specifically the ViT-Base variant. Instead of using convolutions to examine local regions, it relies on self-attention [2, 20] to process images and identify spatial relationships between facial regions that may be far apart. Each image is split into non-overlapping 16×16 patches, transformed into 768-dimensional feature vectors, and sent through 12 transformer layers, with 12 attention heads in each layer. The resulting representation is then fed to a classifier, which assigns it to one of two classes: genuine or synthetic.

To reduce fine-tuning costs, we apply Low-Rank Adaptation (LoRA) [21] to the query, key, and value projections in the attention layers. Rather than updating the full model, this adds trainable low-rank matrices and brings the trainable-parameter share down to about 2-3%, with no real loss of accuracy. In parallel, ResNet-50 serves as a convolutional baseline, extracting local texture features across four successive residual stages, with channel depth increasing from 64 to 2048. Applying LoRA to selected convolutional layers achieves results as good as full fine-tuning while requiring only 5–10% of the computational cost.

Related biometric approaches using angular-margin loss functions have recently shown strong performance in face-based identity verification, so they offer a useful complement for future model comparisons [22]. In addition to the features extracted by the two networks, every image undergoes a fast Fourier transform (FFT), since GAN-synthesis artifacts may be visible in the frequency spectrum. For each 224×224 input, a two-dimensional FFT is calculated, its magnitude spectrum is log-scaled, and a high-frequency region is kept as the spectral descriptor. Directly before the final fully connected classification head, the descriptor and backbone embedding are L2-normalized and concatenated at the feature level. There is no intermediate learned projection layer. The current experimental record does not document the exact dimensionality of the retained FFT descriptor, so no reconstructed value is reported. This limits reproducibility and will be addressed in subsequent experiments by retaining the complete feature-extraction configuration.

The evaluated configuration combines spatial features with frequency-domain information. Since no ablation compared the spatial-only setup with the FFT-fused version, the current results cannot isolate the FFT descriptor's marginal contribution. Training used the HuggingFace Trainer API, automatic batch-size selection, the AdamW optimizer, and cross-entropy loss with label smoothing set to 0.1. Most base-model parameters were frozen through LoRA, while evaluation metrics were collected each epoch to monitor quality. The original experimental report did not preserve a complete reproducibility record for the LoRA rank r, scaling factor α, all target-module details, learning rate, exact epoch count, batch size selected automatically by the training framework, or hardware configuration. This revision does not reconstruct or estimate those values. The limitation is explicitly acknowledged; future experiments will retain the full configuration and random seeds. The experiments described here trained on 100,000 images and evaluated performance on a held-out, class-balanced validation set containing 20,000 images-10,000 genuine and 10,000 synthetic. That accounts for 120,000 of the 140,000 source images, leaving 20,000 images without reported results from an independent test set. The final metrics therefore reflect validation-set performance, not independent test performance. Repeated runs with different random seeds were not conducted, and confidence intervals were not calculated; both are identified as future evaluation work.

Results and discussion

Deepfake-detection accuracy

Table 3 compares the models across training accuracy, validation accuracy, loss, and F1-score. The strongest result came from the LoRA-adapted ViT. Using the retained confusion matrix (TP=9752, TN=9744, FP=256, FN=248; N=20,000), the validation accuracy recalculates to 97.48%. LoRA-adapted ResNet-50 reached 75.74% validation accuracy. The unadapted models were lower: ViT-Base achieved 84.47%, whereas ResNet-Base reached 70.27%. On the validation set, the evaluated ViT-LoRA configuration therefore showed a substantial performance advantage.

Table 3 – Summary of model performance

Model	Training accuracy	Validation accuracy	Loss	F1-score
ViT-Base	0.8243	0.8447	0.4653	0.8259
ViT-LoRA	0.9559	0.9748	0.2760	0.9748
ResNet-Base	0.7048	0.7027	0.6001	0.6980
ResNet-LoRA	0.8014	0.7574	0.4916	0.7600

Table 4's confusion matrix shows 256 false positives and 248 false negatives for ViT-LoRA, with an accuracy of 0.9748. ResNet-Base had considerably more errors: 2900 false positives and 3046 false negatives. This suggests that the unadapted convolutional approach is less reliable near the classification margin. The complete TP/TN/FP/FN outputs for ResNet-LoRA were not retained from the original experimental run. Since aggregate accuracy and F1-score alone cannot reliably be used to reconstruct FPR and FNR, ResNet-LoRA appears in Table 3 but is excluded from Tables 4 and 5; no inferred or unverifiable confusion-matrix values were assigned.

Beyond overall accuracy, the final false-positive and false-negative rates (FPR and FNR) for each model were computed directly from the confusion-matrix values in Table 4, rather than tracked across training epochs. This final view puts both ViT-LoRA error rates at roughly 2.5%, compared with nearly 30% for ResNet-Base. The difference reinforces the conclusion that the LoRA-adapted transformer is much more reliable in avoiding incorrect rejections of both genuine and synthetic images.

Table 4 – Classification confusion matrices

Model	TP	TN	FP	FN	Accuracy
ViT-LoRA	9752	9744	256	248	0.9748
ViT-Base	8576	8318	1682	1424	0.8447
ResNet-Base	6954	7100	2900	3046	0.7027

Table 5 – Final false-positive and false-negative rates by model

Model	FPR (false-positive rate)	FNR (false-negative rate)
ViT-LoRA	2.56%	2.48%
ViT-Base	16.82%	14.24%
ResNet-Base	29.00%	30.46%

Evaluation of the revoting mechanism

Because deployment on a live Ethereum network was outside the scope of this study, the assessment of the revoting mechanism's benefits is theoretical, grounded in the cryptographic properties of the proposed scheme and in findings reported for comparable blockchain-voting systems [10, 11, 17] (Table 6). As a first move toward empirical validation rather than theory alone, the smart contract is planned for small-scale deployment on a local test network, such as a Ganache or Hardhat node. This will allow the storage and verification-time effects reported here as estimates to be measured directly under controlled conditions before any live deployment. Unlike the revoting and repeat-voting-prevention methods discussed in [11] and [17], where an earlier ballot is overwritten or rejected outright, the mechanism proposed here keeps every previous ballot on-chain in a verifiably burned state instead of removing it. This preserves the complete revoting history for public audit, while the tally includes only the latest valid ballot. The main distinction from the closest prior revoting schemes lies in what happens to the earlier ballot. Overwrite-in-place designs reuse its storage slot and discard the old value; reject-on-resubmission designs refuse the second attempt [11, 17]. Under the burn-and-retain rule proposed here, each superseded ballot remains as an immutable record marked explicitly invalid. The complete voting history can therefore be audited, while exactly one valid ballot per voter is still admitted.

Table 6 – Expected characteristics of the conventional scheme versus the proposed mechanism

Evaluation criterion	Conventional blockchain voting	Dynamic Revoting Mechanism	Expected effect
Vote flexibility	No correction possible	Revoting via self-destructing ballots	Improved voting accuracy
Fraud resistance	High (immutable votes)	High (timestamped, single valid vote enforced)	Security level preserved
Storage efficiency	Ledger retains submitted records	Burned records remain on-chain but are excluded from the active tally/index	No blockchain-storage reduction claimed; active tally/index may be smaller
Finalization speed	Fast (pre-set immutable votes)	Slightly slower (revote verification required)	Estimated 10–15% increase in verification time
Voter satisfaction	No correction possible	Mistake correction allowed	Expected higher voter confidence

The core tradeoff remains between vote flexibility and election security. In a traditional ledger, immutability ensures that no one can change anything, but it also implies that there is no way to allow a voter to correct a mistake. The proposed self-destructing-ballot scheme makes the system more flexible while preserving the rule that only the latest valid vote counts.

In theory, this should make voting more accurate and leave voters more satisfied, although every revote requires additional verification computation. It does not, however, make the blockchain ledger smaller. Because superseded ballots remain on-chain as auditable invalid records, the burn-and-retain mechanism leaves the ledger's size unchanged or larger. Any efficiency gain is limited to the active tally set or the index used during counting, since burned records are omitted there. Since each voter may revote at most once, the total number of on-chain ballot records can only stay the same as or exceed the number in a system without revoting. The earlier estimate of a 30–40% reduction in blockchain storage is therefore withdrawn until measurements from an implementation are available.

Full operation history, including burnt records, remains on-chain, but those records are excluded from the tally, preserving auditability and protection against retroactive manipulation of the result.

Discussion

The results indicate that combining a LoRA-adapted transformer architecture with frequency-domain analysis protects against voter-impersonation attacks more effectively than authentication based on passwords or documents. Still, biometrics in electoral systems raise privacy concerns, along with the possibility of demographic differences in recognition accuracy. Whether privacy-preserving biometric verification schemes can be designed remains an open question. The present experiments did not include a systematic assessment of demographic bias in the verification module across age, gender, and ethnicity subgroups. That work is planned for the future, as fairness audits are increasingly expected for biometric systems used in electoral settings. The current results also cover only the accuracy of the fused spatial-plus-frequency configuration. To determine and quantify the FFT-based frequency descriptor's marginal contribution, a separate ablation study comparing the spatial-only backbone with the fused variant is needed; this study is planned for the next stage of the work. Specifically, the ablation will evaluate three configurations using the same training and validation protocol: the spatial backbone alone, the frequency (FFT) descriptor alone, and the fused spatial-plus-frequency model. It will report accuracy as well as false-positive and false-negative rates, isolating and quantifying the frequency descriptor's marginal contribution instead of leaving it to inference. The present study acknowledges the lack of this ablation as a limitation.

Residual error rates matter operationally when the system is used at election scale. With FPR and FNR at approximately 2.5%, millions of voters could produce a substantial number of authentication attempts that are incorrectly accepted or rejected. The biometric/deepfake module therefore should not serve as the only irreversible decision point. Rejected or low-confidence cases need a fallback procedure, which might involve a second biometric capture, an independent identity-verification channel, or supervised/manual verification carried out under election authority procedures.

The Dynamic Revoting Mechanism removes a basic limitation of blockchain voting: once a vote is cast, it cannot be changed. That flexibility comes at a cost, though. Verifying revotes requires additional computation, and it also raises a longer-term question about auditing.

Burned ballots remain on the chain but are left out of the count, a practice regulators might view as a transparency problem. Balancing the ability to revote with stronger auditability guarantees would be a worthwhile subject for future research. Ethereum provides the decentralization and public verifiability the system requires, but transaction costs and limited throughput still make deployment at the scale of a national election impractical. Layer-2 scaling solutions could be considered, along with alternatives such as Hyperledger Fabric or Algorand, to reduce operating costs without giving up the system's core security guarantees. The technology of generative models is also an evolving field in its own right: as new methods for synthesis are developed, it is likely that the detection system based on ResNet-50 features and frequency analysis will need to be retrained, and possibly even subjected to adversarial training, in order to maintain its current level of effectiveness.

Storing the complete tally on a public ledger reflects a deliberate choice between transparency and ballot privacy. Another line of research keeps the tally private by using secure multi-party computation (MPC) or homomorphic encryption rather than counting ballots openly on-chain. The systems combine individual ballots without revealing them and decrypt only the final result. Voters gain stronger secrecy, but the trust requirement shifts to a group of computing parties. These methods also impose substantial cryptographic and coordination overhead, and they make independent public verification of the count considerably harder. The design presented here instead prioritizes end-to-end public verifiability: any observer can inspect the chain and confirm that the tally contains only the latest valid ballot submitted by each voter. One possible compromise would pair MPC-based secret tallying with the self-destructing-ballot revoting logic proposed here, combining greater voter privacy with the auditability that full transparency provides.

Ballot privacy also means that VoteHash cannot be a deterministic hash based only on the candidate choice. Since elections typically offer a small set of candidates, an attacker could recover the selected option through dictionary or brute-force search. In a production system, the value should instead be a hiding commitment, such as H (choice/nonce), with nonce serving as a cryptographically random, voter-specific value, or it should use an encrypted ballot representation. The current model assumes that the on-chain value reveals no plaintext voting choice; formally analyzing privacy and implementing the commitment or encryption layer remain future work.

Lastly, the deployment of such a system within a governmental context necessitates alignment with data-protection and electoral regulation. This, in turn, requires engagement with regulators, election officials, and cybersecurity experts as a prerequisite for transitioning the research prototype into an operational service. Within a Smart City digital ecosystem, this kind of system is best understood as one of several services that increase the convenience and transparency of citizens' interactions with public institutions, alongside identity services already in place.

Conclusion

This paper proposed and experimentally evaluated an electronic voting model that combines biometric identity verification using transformer and CNN architectures, frequency-domain information for detecting synthetic images, and blockchain-based result recording with controlled revoting. In the ViT-LoRA setting, the recognition module reached 97.48% validation accuracy according to the retained confusion matrix, outperforming the other configurations examined. The independent contribution of the frequency descriptor has not yet been quantified because no FFT ablation study was conducted. Controlled revoting is supported by pairing a self-destructing-ballot scheme with blockchain-recorded timestamp ordering, while the complete on-chain audit history is preserved. For production deployment, stronger timestamp guarantees and ballot-hiding commitments are still needed; these are implementation requirements, not properties experimentally validated by the current prototype.

In the future, we plan to assess the revoting scheme empirically rather than theoretically on a deployed network, explore privacy-preserving biometric authentication schemes, and enhance the deepfake-detection module's robustness against newer generation techniques. The proposed solution is intended as a technical foundation for further research into secure, transparent digital elections within Smart City infrastructure.

Funding

This research was funded by the Committee of Science of the Ministry of Science and Higher Education of the Republic of Kazakhstan (Grant No. BR24992852, "Intelligent models and methods of Smart City digital ecosystem for sustainable development and the citizens' quality of life improvement").

Competing Interests

The authors declare no competing interests.

REFERENCES

1 Hayath, T.M., Rajeev, D.V., Nihanth, R., Tharun, U.V., and Achutha, R. Digital Voting System with Face Recognition. International Journal of Scientific Research and Technology, 3 (1) (2026).

2 Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. An Image Is Worth 16×16 Words: Transformers for Image Recognition at Scale. In: Proceedings of the International Conference on Learning Representations (ICLR) (2021).

3 He, K., Zhang, X., Ren, S., and Sun, J. Deep Residual Learning for Image Recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778 (2016).

4 Alkishri, W., et al. An Improved Dense CNN Architecture for Deepfake Image Detection. IEEE Access, 11, 22081–22095 (2023).

5 Sheikh, M.S., Kirtonia, U., Arthi, N.T., and Al-Imran, M. AI-Powered Deepfake Detection Using CNN and Vision Transformer Architectures. arXiv preprint arXiv:2601.01281 (2026).

6 Wang, Z., Cheng, Z., Xiong, J., Xu, X., Li, T., Veeravalli, B., and Yang, X. A Timely Survey on Vision Transformer for Deepfake Detection. arXiv preprint arXiv:2405.08463 (2024).

7 Frank, J., Eisenhofer, T., Schönherr, L., Fischer, A., Kolossa, D., and Holz, T. Leveraging Frequency Analysis for Deep Fake Image Recognition. In: Proceedings of the International Conference on Machine Learning (ICML), PMLR, pp. 3247–3258 (2020).

8 Tan, C., Zhao, Y., Wei, S., Gu, G., Liu, P., and Wei, Y. Frequency-Aware Deepfake Detection: Improving Generalizability through Frequency Space Domain Learning. In: Proceedings of the AAAI Conference on Artificial Intelligence, 38, 5052–5060 (2024).

9 Nakamoto, S. Bitcoin: A Peer-to-Peer Electronic Cash System (2008). https://bitcoin.org

10 Ohize, H.O., Dogo, E.M., Onumanyi, A.J., Ibrahim, M.M., et al. Blockchain for Securing Electronic Voting Systems: A Survey of Architectures, Trends, Solutions, and Challenges. Cluster Computing, 28, Article No. 132 (2025).

11 Jafar, U., Ab Aziz, M.J., Shukur, Z., and Hussain, H.A. A Systematic Literature Review and Meta-Analysis on Scalable Blockchain-Based Electronic Voting Systems. Sensors, 22 (19), Article No. 7585 (2022).

12 Gandhi, S.S., Kiwelekar, A.W., Netak, L.D., and Wankhede, H.S. Security Requirement Analysis of Blockchain-Based E-Voting Systems. In: Intelligent Communication Technologies and Virtual Mobile Networks, Springer (2023).

13 Sharp, M., Njilla, L., Huang, C.-T., and Geng, T. Blockchain-Based E-Voting Mechanisms: A Survey and a Proposal. Network, 4 (4), 426–442 (2024).

14 El Kafhali, S., et al. Blockchain-Based Electronic Voting System: Significance and Requirements. Mathematical Problems in Engineering, 2024, Article No. 5591147 (2024).

15 Wang, B., Guo, F., Liu, Y., Li, B., and Yuan, Y. An Efficient and Versatile E-Voting Scheme on Blockchain. Cybersecurity, 7, Article No. 62 (2024).

16 Ghosh, V., and Gupta, H. A Blockchain-Based E-Voting System for Secure and Transparent Elections. In: ICT Systems and Sustainability (ICT4SD 2024), Lecture Notes in Networks and Systems, 1163, Springer, pp. 163–173 (2025).

17 Hajian Berenjestanaki, M., Barzegar, H.R., et al. Blockchain-Based E-Voting Systems: A Technology Review. Electronics, 13 (1), Article No. 17 (2024).

18 Paudel, S., Poudel, A., and Paudel, S. Enhancing Electoral Integrity and Accessibility: A Blockchain and Facial Recognition-Based Electronic Voting System. Information Dynamics and Applications, 4 (2), 85–94 (2025).

19 Aidynov, T., Goranin, N., Satybaldina, D., and Nurusheva, A. A Systematic Literature Review of Current Trends in Electronic Voting System Protection Using Modern Cryptography. Applied Sciences, 14, Article No. 2742 (2024). (Authors' own prior work; cited here as a self-citation.)

20 Vaswani, A., Shazeer, N., Parmar, N., et al. Attention Is All You Need. In: Advances in Neural Information Processing Systems, 30, 5998–6008 (2017).

21 Hu, E.J., Shen, Y., Wallis, P., et al. LoRA: Low-Rank Adaptation of Large Language Models. arXiv preprint arXiv:2106.09685 (2021).

22 Deng, J., Guo, J., Xue, N., and Zafeiriou, S. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4690–4699 (2019).

[1]Айдынов Т.,
докторант, ғылыми қызметкер, ORCID ID: 0009-0005-3327-9719,
e-mail: tolegen.ch@gmail.com
[1] Сатыбалдина Д.,
ф.-м.ғ.к., профессор, ORCID ID: 0000-0003-0291-4685,
e-mail: satybaldina_dzh@enu.kz
[2] Димитров В.,
PhD, профессор, ORCID ID: 0000-0002-5348-4948,
e-mail: v.dimitrov@unibit.bg
[1*]Абишева Г.,
PhD, ORCID ID: 0009-0003-6692-197X,
e-mail: gulsipat.abisheva@gmail.com

[1]Л.Н. Гумилев атындағы Еуразия ұлттық университеті,
Астана қ., Қазақстан
[2]Кітапханатану және ақпараттық технологиялар университеті,
София қ., Болгария

ДИПФЕЙККЕ ҚАРСЫ АУТЕНТИФИКАЦИЯСЫ БАР БИОМЕТРИЯЛЫҚ САЙЛАУШЫНЫ СӘЙКЕСТЕНДІРУГЕ НЕГІЗДЕЛГЕН БЛОКЧЕЙН ЭЛЕКТРОНДЫҚ ДАУЫС БЕРУ МОДЕЛІ

Аңдатпа

Мақалада блокчейн технологиясына негізделген дауыс беру нәтижелерін орталықтандырылмаған түрде сақтау, сайлаушыларды биометриялық сәйкестендіру және синтетикалық (deepfake) бет кескіндерін автоматты анықтау мүмкіндіктерін біріктіретін электрондық дауыс беру моделі ұсынылып, тәжірибелік тұрғыда расталған. Сайлаушының жеке басын тексеру үшін Low-Rank Adaptation (LoRA) әдісі арқылы бейімделген Vision Transformer және ResNet-50 архитектуралары қолданылады. Сонымен қатар бет ауыстыру (face swap) және синтетикалық генерациялау шабуылдарына төзімділікті арттыру мақсатында кіріс кескініне жылдам Фурье түрлендіруі (FFT) негізіндегі жиіліктік талдау жүргізіледі. Блокчейннің өзгермейтін табиғатына байланысты қайта дауыс беру мәселесін шешу үшін өзін-өзі жоятын бюллетеньдер мен қорғалған уақыт белгілеріне негізделген бақыланатын қайта дауыс беру механизмі ұсынылған. Дауыстар Ethereum смарт-келісімшарты арқылы тіркеледі. Эксперимент нәтижелері бойынша LoRA әдісімен бейімделген Vision Transformer моделі нақты және синтетикалық бет кескіндерін ажыратуда 97,36 % дәлдікке қол жеткізіп, LoRA арқылы бейімделген ResNet-50 моделінен айтарлықтай жоғары нәтиже көрсетті. Ұсынылған жүйе Smart City цифрлық инфрақұрылымының құрамдас бөлігі ретінде қарастырылып, құпиялылықты сақтайтын биометриялық хаттамалар мен сайлау процестеріне арналған ауқымды блокчейн платформаларын дамытуға негіз бола алады.

Түйін сөздер: электрондық дауыс беру, бетті тану, дипфейкті анықтау, блокчейн, Ethereum, смарт-келісімшарттар, Vision Transformer, ResNet, қайта дауыс беру механизмі, Smart City.

[1]Айдынов Т.,
докторант, научный сотрудник, ORCID ID: 0009-0005-3327-9719,
e-mail: tolegen.ch@gmail.com
[1]Сатыбалдина Д.,
к. ф.-м.н., профессор, ORCID ID: 0000-0003-0291-4685,
e-mail: satybaldina_dzh@enu.kz
[2]Димитров В.,
PhD, профессор, ORCID ID: 0000-0002-5348-4948,
e-mail: v.dimitrov@unibit.bg
[1*]Абишева, Г.,
PhD, ORCID ID: 0009-0003-6692-197X,
e-mail: gulsipat.abisheva@gmail.com

[1] Евразийский национальный университет им. Л.Н. Гумилева,
г. Астана, Казахстан
[2] Университет библиотековедения и информационных технологий,
г. София, Болгария

БЛОКЧЕЙН-ОРИЕНТИРОВАННАЯ МОДЕЛЬ ЭЛЕКТРОННОГО ГОЛОСОВАНИЯ С БИОМЕТРИЧЕСКОЙ ИДЕНТИФИКАЦИЕЙ ИЗБИРАТЕЛЕЙ И ЗАЩИТОЙ ОТ АТАК С ИСПОЛЬЗОВАНИЕМ ДИПФЕЙКОВ

Аннотация

В статье представлена и экспериментально апробирована модель электронной системы голосования, объединяющая децентрализованное хранение результатов на основе технологии блокчейн, биометрическую идентификацию избирателей и автоматическое обнаружение синтетически сгенерированных (deepfake) изображений лиц. Для проверки личности избирателя используются архитектуры Vision Transformer и ResNet-50, дообученные с применением метода Low-Rank Adaptation (LoRA). Дополнительная устойчивость к атакам с использованием подмены лица и синтетической генерации обеспечивается анализом входного изображения в частотной области с использованием быстрого преобразования Фурье (FFT). Для решения проблемы неизменяемости блокчейна предложен механизм контролируемого переголосования, основанный на самоудаляющихся бюллетенях и защищенных временных метках, при этом сами голоса фиксируются в смарт-контракте Ethereum. Экспериментальные результаты показали, что модель Vision Transformer с адаптацией LoRA достигает точности 97,36% при различении реальных и синтетических изображений лиц, значительно превосходя аналогичную модель ResNet-50. Предлагаемая система рассматривается как компонент цифровой инфраструктуры Smart City и может служить основой для дальнейших исследований в области биометрических протоколов с сохранением конфиденциальности и масштабируемых блокчейн-платформ для проведения выборов.

Ключевые слова: электронное голосование, распознавание лиц, обнаружение дипфейков, блокчейн, Ethereum, смарт-контракты, Vision Transformer, ResNet, механизм переголосования, Smart City.