Designing Trust: Creating Voting Systems that Inspire Confidence in the Electoral Process

CS4991 Capstone Report, 2025

Alexander Davis
Computer Science
The University of Virginia
School of Engineering and Applied Science
Charlottesville, Virginia USA
wfy8cn@virginia.edu

ABSTRACT

Elections are the cornerstone of democracy, yet more than ever people lack trust in this process. My proposal focuses on how we can design voting systems in a way that encourages security and promotes trust in elections. Some methods include using open source software that can be vetted by neutral third parties to ensure transparency in the voting process. This paper proposes a general purpose LLM trained to both assist election officials and inform voters. Preliminary implementation of some of these methods shows increased trust by the general population in the electoral process. Future work should focus on communication of these technologies to increase technical literacy and ensure trustworthy voting systems are actually trusted by the general public.

1. INTRODUCTION

Well-run elections are essential for a functioning democracy, and failure in the election process can have disastrous consequences for society at large. A growing number of threats to democracy have made many in both the United States and around the world distrustful of elections and institutions as a whole. These systems can only function if the general population is invested and participates in them, so it is paramount that we work to reestablish trust in our elections.

It is important to run elections that are secure. A study by Steward (2022) suggested that we should distinguish between two distinct aspects of security: trustworthiness and trust (2022). Trustworthiness refers to how secure an election system actually is, and is a measure of how much we should trust the system. This could be influenced by vulnerabilities in voting machines, the level of security protocols followed by an election locality, problems with software running on election equipment, or similar factors.

Trust refers to how people perceive the security of the election system, and may or may not be related to how secure the system actually is. While security vulnerabilities that decrease trustworthiness are problematic, it can be messier to determine how trust impacts elections. Elections require voters to put faith in the results, so people distrusting the outcomes of an election or believing it was not administered fairly, regardless of whether any actual malpractice occurred, can be just as damaging as an actual problem with the process. In some ways, combatting distrust can be more difficult than restoring trustworthiness. While a software bug or lapse in security protocol provides an easy target to fix, it is not always obvious how to change how people feel about elections.

2. RELATED WORKS

Perhaps the most well documented security concern surrounding voting machines is the use of Direct Recording Machines (DREs) as opposed to standard paper ballot-

based machines. One study compared the average wait time for voters in locations using each machine type, and found that DREs resulted in longer wait times and a poorer voting experience (Wadowski et al., 2023). Another study documented a cybersecurity hacking conference where several experts were tasked with breaking various DRE machines. Every machine tested was compromised with relative ease, with some vulnerabilities as simple as an unchangeable and publicly available default password. Paper ballots can also serve as a paper trail if any verification of the count is necessary, so DREs can further add doubt into the security of the system (Blaze et al, 2017). While no system is perfect, replacing these machines is one important step localities can take to improve security.

Americans' confidence in institutions as a whole is at an all time low, and trust in elections specifically is declining (Rainie & Perrin, 2019). While combating mistrust can be murkier than fixing bugs in software, there are some common reasons for this mistrust. Research suggests that trust can be influenced by demographic factors including race and political affiliation. Greater knowledge about both the electoral process and technology involved in elections is associated with higher levels of trust (Alvarez et al., 2008).

Political losing is also a reason, referring to the trend that voters supporting losing candidates distrust election outcomes more than other voters, regardless of any other factors (Mauk, 2022). Different groups of people may have different reasons for skepticism, so any attempt to improve trust needs to consider the intended audience. Some studies have found that including nonpartisan election observers can dramatically raise trust. Dedicating more resources to local election offices and taking steps such as increasing poll worker training can improve voter experience and repair trust at a grassroots level (Bush & Prather, 2023).

Another potential solution could be using open source software for voting machines. In New Hampshire, some local officials explored switching to such a system, arguing that it would reassure voters of election integrity. Open source software can often increase security, as anyone can review source code and contribute to patching vulnerabilities, which would allow voters to review the software and feel confident their votes were counted correctly (Mestel, 2024).

3. PROPOSAL DESIGN

The proposed design is to make an encompassing tool that can be used by both localities and voters to get information about election systems and guidance on what they can do to improve election security. The overall design involves fine tuning an LLM on election data to find a general model that can be given to localities. Each locality would then adapt the model to work with their specific data. The model would have two different versions, one meant for elections workers, which would have heightened permissions and provide more detailed and specific answers, and another meant for voters inquiring about election proceedings. Based on both my research and experience working with an elections office, the biggest challenge is not a lack of solutions to security concerns, but rather a lack of a cohesive plan and resources to implement these solutions. By providing an adaptive tool to localities, we can assist in implementing some of these solutions and in promoting trust in the general public.

The first step would be to train a large language model (LLM) to answer election related questions. I propose using LLaMA, an LLM that is open source and lightweight. Being open source would enable anyone to analyze the design of the model and align with the goal of promoting trust, and the lightweight nature would allow for use even by low resource localities. The main difficulty would be creating a reasonable training set to build the model on. Since I do not have access

to government records, I would likely be unable to build such a dataset myself, but I could create a proof of concept to show how such a model could function when built on a larger and more accurate dataset. To do this, I would need to scrape information from public election sites and train the model on this information. I would need to include data from different localities and data discussing different topics (a broader scope than just election machines or just software) to show the model can generalize for different contexts. During training, the model would use masking to hide subsets of the data and attempt to reconstruct sentences/paragraphs. Once trained, the model would iteratively predict the next word based on past context.

To address how the model can be modified to work for a specific locality, I plan to use retrieval augmented generation (RAG). RAG is a technique used to augment an LLM to retrieve information from a trusted database rather than solely the training data. Due to the sensitive nature of election information, responses need to have a high degree of accuracy, and using a regulated database ensures answers can accurately reflect the election system. There may be certain contexts where a very specific answer is required, and this can force the model to give such an answer. It can also provide certain regional context, such as what vendors a locality buys their equipment from. For a question on certain hardware, the RAG model could find specific information such as model and permissions (an election official would be given more information than a standard voter), and give this as additional context to help the LLM produce a response. Implementing RAG would allow the LLM to determine how to generally respond to a question, while the RAG model would augment it with context specific to a locality.

4. ANTICIPATED RESULTS

I would use various tools such as lexical substitution, which tests a model's domain knowledge by asking for synonyms to substitute in a sentence, to evaluate how well the model performs. One goal here would be to understand whether or not an LLM can be adapted to the domain of elections, or if a more supervised approach would be necessary to create a useful model. A result could be that the model produced accurate results on 98% of the training set, but this could vary depending on what metric is used. I would also want to perform user testing to determine how effective it is perceived to be. This would need to involve voters and elections officials and survey how accurate and effective the model was. This survey, examining the model's accuracy and effectiveness, would involve voters and elections officials.

If successful, this project could be implemented by localities in the United States. This would be a relatively cost-effective solution because it can be applied anywhere, rather than just a subset of locations. Election workers could use the tool to help assess the state of their election security system, and create a plan to improve based on their specific needs and resources. One benchmark to measure this aspect could be analyzing change in compliance with LESS. This would also streamline the process of repairing trust by creating a nationally generalizable tool that anyone could use to learn about election security.

5. CONCLUSION

This proposal is inspired in part by a summer internship in which I worked with the Arlington County elections office to improve security standards in preparation for the 2024 election. We primarily focused on increasing compliance with the Local Election Security Standards (LESS), which provides guidelines for election localities in Virginia to follow in order to be certified as prepared to run elections. The most significant takeaway from both my own experience and from discussing with interns working at other localities was how important funding and resources could be

to determining how well the locality met these standards. While LESS provides an idealistic guideline, many localities face difficult tradeoffs in choosing which standards to fix with the limited resources they have.

While I left my locality more prepared for the election than they had been, I felt limited in the scope of what I could accomplish. Localities often lack the staff to properly address security concerns, and many localities deal with the same kinds of problems. After completing the internship, I wanted to find a way to address all of these issues and help localities efficiently divide the resources they have and find some kind of plan that suits them. I also wanted to find a way to not only help localities address their security concerns, but ensure these improvements are conveyed to the average voter in a way that improves trust. This project could help to address all of these problems by creating an accessible tool that can be used by anyone to build trust.

6. FUTURE WORK

In the future, the tool should be improved to ensure it can effectively communicate with users. Further research should focus on improving the data set used to train the model. Increasing the scope of the dataset or the accuracy would allow for better generalization and decrease errors. Depending on the results of this study once implemented, the model design may need to be altered. Future work could also explore different base models or how different architectures influence results. Finally, much of future work would be less technical and involve continued efforts to bridge the gap between technical and non-technical actors in the election process.

REFERENCES

Wadowski, G., Otte, L., Bernardo, N., & Macht, G. (2023). A comparative study of electronic voting and paper ballot systems in modern elections. https://www.example.comhttps://esra-conference.org/files/election-science-conference/files/a_comparative_study_of_electronic_voting_and_paper_ballot_systems_in_modern_elections_wadowski_-_uri_club_tennis.pdf

Blaze, M., Braun, J., Hursti, H., Hall, J., MacAlpine, M., & Moss, J. (2017, September). Defcon 25 voting Machine Hacking Village. DEFCON. https://defcon.org/images/defcon-25/DEF%20CON%2025%20voting%20village%20report.pdf

Mauk, M. (2022). Electoral integrity matters: How the electoral process conditions the relationship between political losing and political trust. *Qual Quant* 56, 1709–1728. https://doi.org/10.1007/s11135-020-01050-1

Alvarez, R., Hall, T., & Llewellyn, M. (2008). Are Americans confident their ballots are counted? *The Journal of Politics*, 70(3), 754–766. doi:10.1017/S0022381608080730

Bush, S., & Prather, L. (2023, January 13). How to restore trust in U.S. election results. Greater Good. https://greatergood.berkeley.edu/article/item/how_to_restore_trust_in_us_election_results

Stewart, C., III. (2022). Trust in elections. Daedalus, 151(4), 234–253. https://doi.org/10.1162/daed_a_01953

Mestel, S. (2024). How open source voting machines could boost trust in US elections. MIT Technology Review. https://www.technologyreview.com/2024/03/07/1089524/open-source-voting-machines-us-elections/

Rainie, L., & Perrin, A. (2019, July 22). Key findings about Americans' declining trust in government and each other. Pew Research Center. https://www.pewresearch.org/short-reads/2019/07/22/key-findings-about-americans-declining-trust-in-government-and-each-other/