DOI:10.1145/2977335

Election Auditing and Verifiability

Overall, the Inside Risks Viewpoint "The Risks of Self-Auditing Systems" by Rebecca T. Mercuri and Peter G. Neumann (June 2016) was excellent, and we applaud its call for auditing systems by independent entities to ensure correctness and trustworthiness. However, with respect to voting, it said, "Some research has been devoted to end-to-end cryptographic verification that would allow voters to demonstrate their choices were correctly recorded and accurately counted. However, this concept (as with Internet voting) enables possibilities of vote buying and selling." This statement is incorrect.

While Internet voting (like any remote-voting method) is indeed vulnerable to vote buying and selling, end-to-end verifiable voting is not. Poll-site-based end-to-end verifiable voting systems use cryptographic methods to ensure voters can verify their own votes are correctly recorded and tallied while (paradoxically) not enabling them to demonstrate how they voted to anyone else.

Mercuri and Neumann also said, "[end-to-end verifiability] raises serious questions of the correctness of the cryptographic algorithms and their implementation." This sentence is potentially misleading, as it suggests confidence in the correctness of the election outcome requires confidence in the correctness of the implementation of the cryptographic algorithms. But end-to-verifiable voting systems are designed to be "fail safe"; if the cryptographic algorithms in the voting system are implemented incorrectly, the audit will indeed fail. Poor crypto implementations in the voting system will not allow an audit to approve an incorrect election outcome.

Finally, we note that end-to-end verifiable election methods are a special case of "verifiable computation," whereby a program can produce not only a correct result but also a "proof" that it is the correct result for the given inputs. Of course, the inputs need to be agreed upon before such a proof makes sense. Such methods may thus be useful not only for election audits but elsewhere.

Joseph Kiniry, Portland, OR, and
Ronald L. Rivest, Cambridge, MA

Authors Respond:

We cannot fully elucidate here the flaws in each of the many proposed cryptographically verifiable voting subsystems. Their complexity and that of the surrounding systems environments undemocratically shifts the confirmation of correct implementation to a scant few intellectually elite citizens, if even accomplishable within an election cycle. However, all of these methods have vulnerabilities similar to the Volkswagen emission system; that is, stealth code can be triggered situationally, appearing correct externally while internally shifting vote tallies in favor of certain candidates over others. We have previously discussed the incompleteness of cryptographic solutions embedded in untrustworthy infrastructures, potentially enabling ballot contents to be manipulated or detected via vote-selling tags (such as write-in candidates or other triggers). The mathematics of close elections also requires that a very high percentage of ballots (over 95%) be independently checked against the digital record, which is not likely to occur, leaving the results unverified.

Rebecca T. Mercuri, Hamilton, NJ, and
Peter G. Neumann, Menlo Park, CA

Unintended Consequences of Trusting AIs

Toby Walsh's Viewpoint "Turing's Red Flag" (July 2016) raised very good points about the safety of increasingly human-like AI and proposed some common-sense law to anticipate potential risks. It is wise to discuss such protections before the technology itself is perfected. Too often the law trails the technology, as with the Digital Millennium Copyright Act in response—perhaps a decade late—to illegal file sharing.

Walsh primarily addressed the potential threat of autonomous systems being mistaken for humans, but what about the reverse? Humans could gain an unfair or even a dangerous advantage by impersonating an AI. For instance, in a world where autonomous vehicles are allowed smaller following distances and prompt extra caution from nearby human drivers, a human could install an "I am autonomous" identity device in order to tailgate and weave through traffic with impunity, having won unearned trust from other drivers and vehicles.

A similar situation could arise with the advent of bots that act as intermediaries between humans and online services, including, say, banks. As bots become more trusted, a human-in-the-middle attack could compromise everyone's private data.

At perhaps the outer reaches of techno-legal tension, we could even imagine the advent of identity theft where the individual is an AI, lovingly brought to life by a Google or an Amazon, and the thief to be punished is a human impersonator. Is this the route through which AIs might someday become legal persons? In a world where the U.S. Supreme Court has already extended constitutional free speech rights to corporations, this scenario seems quite plausible.

Mark Grossman, Palo Alto, CA

Author Responds:

Grossman makes a valid point. Just as we do not wants bots to be intentionally or unintentionally mistaken for human—as I suggested in my Viewpoint—we also do not want the reverse. The autonomous-only lane on the highway should not have humans in it pretending to be, say, the equivalent of more-capable autonomous drivers.

Toby Walsh, Berlin, Germany

More to Asimov's First Law

In his Viewpoint (July 2016), Toby Walsh argued for some sort of preliminary indication in cases in which a human is interacting with a robot. I suggest he check Isaac Asimov's classic science fiction novels *Caves of Steel* (1953) and *The Naked Sun* (1957) for an earlier treatment of the topic. In the latter work especially,