Bayesian Election Audits in One Page

Audrey Malagon (amalagon@vwu.edu) and Ronald L. Rivest (rivest@mit.edu)
2018-07-13

Philip Stark (2007) invented and promulgates risk-limiting audits (RLAs), which sample cast paper ballots to upper bound the (worst-case) risk that the audit fails to detect and correct an incorrect reported outcome. Bayesian audits, due to Rivest and Shen (2012), are similar, but upper bound the probability (average risk) that similar collections of paper ballots give outcomes different than the reported one.

Before the audit, election officials tabulate the collection C of n cast paper ballots to obtain a **reported outcome R ,** which may be wrong (due to error or fraud), but which should be correct and thus satisfy

$$R = \text{Outcome}(C) .$$

A **statistical election audit** (RLA or Bayesian) aims to give confidence in R's correctness starting with a **hand examination and interpretation** of a random **sample** S of some size s of ballots drawn without replacement from C:

$$S \leftarrow \text{Sample}(C, s)$$

A **Bayesian statistical election audit** uses S to estimate the probability *p* that collections C' similar to C have outcomes different than R . We call *p* the **expected loss** of the audit. The audit stops and accepts R as correct if *p* is less than a given **loss limit** (e.g., 1%). Otherwise the audit repeats using a larger sample, likely escalating to examine all ballots if R is incorrect. A smaller loss limit may require a larger sample but gives more confidence in the correctness of the election outcome.

A Bayesian audit estimates *p* by probabilistically "reversing the sampling" ("restoring") to obtain hundreds of ballot collections C' similar to C, and estimating *p* as the fraction for which

$$R \neq \text{Outcome}(C') .$$

Restoring starts with S , then successively adds n-s votes back, in a random manner, to obtain C':

$$C' \leftarrow \text{Restore}(S, n)$$

Which votes does Restore add back? To ensure that C' is similar to C , Restore adds copies of votes randomly selected from the growing sample. To enable restoration of votes even for candidates not in the sample, Restore adds to S one vote for every candidate (even those with votes) when it starts, and removes one vote for every candidate when it ends. These extra votes determine the **Bayesian prior.**

Since Restore picks votes to copy at random, it may return a somewhat different result C' each time. But each such C' should be similar to the original C, and this similarity improves with the size *s* of the initial sample S . Variations in C' reflect the uncertainty the auditor has about C , and thus about Outcome(C) , knowing only S .

A laptop can restore S to C' quickly, as restoration does not sample or examine by hand any paper ballots. We call constructing C' in this way *drawing from the posterior distribution* defined by S and the prior. Methods based on Dirichlet-multinomial distributions give even greater efficiency.

The ability of a Bayesian audit to probabilistically reverse the sampling process is a powerful tool for auditing. Since a Bayesian audit uses only vote copying, it is independent of the tabulation method and works for complex voting methods like IRV. It extends to handle stratified audits, ballot-level comparison audits (stratifying by reported vote), multi-jurisdiction audits, and audits where jurisdictions have different types of equipment. It may serve as a useful alternative to an RLA. You can find more details here.