powe(R)OC: A power simulation tool for eyewitness lineup ROC analyses

Eric Y. Mah

Department of Psychology, University of Victoria

Author Note

I have no known conflicts of interest to disclose. Correspondence concerning this article should be addressed to Eric Y. Mah, Department of Psychology, University of Victoria, Victoria, BC, V8W 2Y2.

Contact: ericmah@uvic.ca <https://orcid.org/0000-0001-5730-6498>

Abstract

Receiver operating characteristic (ROC) curve analyses have become increasingly common in eyewitness lineup experiments, yet statistical power for these analyses is not well-understood. *powe(R)OC* is a free, open-source R Shiny web app that allows users with minimal programming and statistical knowledge to conduct simulation-based power analyses for two-condition (e.g., simultaneous vs. sequential lineups) eyewitness lineup experiments (and certain other recognition memory experiments) that use ROC analysis. *powe(R)OC* uses existing data (either user-uploaded or publicly-available open data) as a basis for simulation, allows powering for analyses of partial area under the curve (pAUC) and *deviation from perfect performance* (Smith et al., 2019), allows users to specify various simulation parameters (e.g., effect sizes, sample sizes, number of lineups, pAUC truncation, test tails), view simulation results uploaded by other users, view ROC effect sizes in the literature, and download simulation results in a summary report. This report describes ROC analyses, challenges in simulating ROC data, methods implemented in *powe(R)OC* and their underlying assumptions and limitations, and possible future directions.

powe(R)OC: A power simulation tool for eyewitness lineup ROC analyses**ROC analysis in the eyewitness literature**

Receiver operating characteristic curves are used most commonly in the medical field to index the diagnostic accuracy of a test (Mickes et al., 2012). ROC analysis has only recently seen widespread adoption as a measure of lineup procedure diagnosticity/discriminability in the field of eyewitness identification. This adoption was largely spurred by criticisms of prior measures (e.g., diagnosticity ratios) and the sequential versus simultaneous lineup debate (Mickes et al., 2012; Wixted & Mickes, 2012).

In eyewitness studies, ROC curves are constructed by plotting pairs of correct and false ID rates across a range of operating points, with each point representing different levels of confidence (Wixted & Mickes, 2012). Here, “correct ID” refers to lineup identifications of the culprit in culprit-present/target-present lineups, and “false ID” refers to lineup identifications of the innocent suspect in culprit-absent/target-absent lineups. The treatment of filler identifications (picks of lineup members who are not the culprit or suspect) in ROC analysis differs, but typically filler IDs are treated as equivalent to lineup rejections. ROC curves for different lineup procedures or experimental manipulations (e.g., simultaneous & sequential lineups) can be compared via analyses of *area-under-the-curve* (AUC), or more typically, *partial area-under-the-curve* (pAUC), which truncates both curves at the lowest observed false ID rate (Wixted & Mickes, 2014). Two hypothetical ROC curves are shown below in Figure 1:

Figure 1

Example eyewitness ROC curves

In the figure above, hypothetical lineup Condition B bows further away from the dashed line (representing chance diagnostic accuracy/discriminability), resulting in greater pAUC (i.e., greater empirical discriminability) than hypothetical lineup Condition A.

Since the papers of Mickes, Wixted and others (2012), the use of ROC analysis in the field has become more common. For instance, Figure 2 below shows Google Scholar entries including the keywords “eyewitness”, “lineup”, and “ROC”, plotted as a proportion of the number of entries including the keywords “eyewitness” and “lineup”:

Figure 2

Proportion of eyewitness lineup papers mentioning ROC, over time

Note. Ribbon = bootstrapped 95% percentile CI on the proportions

The increased interest in ROC has prompted lively debate on issues of measurement in the field, with proponents of ROC arguing that analysis of partial area-under-the-curve (pAUC) provides a valid, theory-free measure of eyewitness discriminability (Wixted & Mickes, 2015, 2018), and opponents arguing that the widely adopted pAUC analyses leave out important information about discriminability (Wells, Smalarz, et al., 2015) and that other measures (ROC-based or otherwise) are needed to capture lineup outcomes of interest (Smith et al., 2019, 2020; Wells, Yang, et al., 2015). For instance, Smith et al. (2019) propose *deviation from perfect performance* (DPP) as an alternative ROC-based measure. DPP involves the calculation of diagnostic utility at each operating point on the ROC curve, and averaging utility across all points provides a measure of *average utility* that is similar to pAUC. However, Smith et al. (2019) argue that DPP more consistently indicates which of two lineup procedures is diagnostically superior, especially in cases where two procedures

result in differential false ID rates. Regardless of whether one falls into the pAUC or DPP camp (or some other camp, e.g., *full ROC curves*; Smith et al., 2020), it seems likely that ROC-based analyses of eyewitness lineup data will continue to be widely-used.

Power for ROC analysis

Statistical power (the probability of correctly rejecting a false H_0) is an important consideration when planning an experiment (J. Cohen, 1992), and ROC analyses are no exception. Despite this, a review of the eyewitness literature (and correspondence with several prominent researchers in the field) reveals a lack of principled methods for ROC experiment sample size planning. For example, according to Colloff and Wixted, “methods for calculating a priori power analysis for eyewitness identification experiments are not well specified” (2020, p. 126).

As a result, researchers conducting eyewitness ROC experiments have relied on sample size heuristics that have not been extensively validated (e.g., $n = 300\text{--}500$ per condition, inputting previously-observed pAUCs and test statistics into standard power calculators; Colloff & Wixted, 2020). For DPP, there is a lack of even heuristic guidance when it comes to sample size planning (understandable given the novelty of the measure). Arguably, power and sample size planning should be especially important for ROC analyses, which typically involve thousands of participants and substantial costs. Without a good understanding of power for these analyses, researchers run the risk of wasting money and participant time, both for under- and over-powered experiments.

In the medical literature (where ROC analysis is more common), there do exist some analytic methods for estimating power for ROC experiments (Hillis et al., 2011; Hillis & Scharz, 2018; Kim et al., 2014; Negida et al., 2019). However, these methods either require deep understanding of underlying ROC parameters or make assumptions that may not be

appropriate eyewitness ROC experiments (e.g., partial AUC truncation, only single or paired-ROC curve designs).

Regardless of the validity of the above analytic methods and their applicability to eyewitness experiments, there are several reasons to prefer simulation-based power estimation methods to analytic ones. Briefly, simulation-based methods involve repeated random sampling from data with pre-specified characteristics (e.g., N , M , SD), followed by the application of the statistical test(s) of interests to each random sample, after which one can examine the long-run behaviour of the statistical test (e.g., proportion of samples where the test was significant).

One major advantage of simulation-based methods is that they can provide more accurate results when data are “messy” and the assumptions of analytic methods do not hold (as is common with real-world data; Morris et al., 2019). Simulation-based methods are also often more intuitive than their analytic counterparts. (e.g., for specifying effect sizes and analysis parameters, simulated tests are the same as that used in actual analyses; Lakens & Caldwell, 2021). Finally, simulation-based methods allow tighter control over and specification of data parameters (where analytic methods may only allow for specification of a few parameters).

There have been recent strides in simulation-based analyses in the eyewitness literature, with two noteworthy examples. The first is `sdtlu`, an R (R Core Team, 2021) package for signal detection theory (SDT) analysis of lineup data (A. L. Cohen et al., 2021). `sdtlu` allows users to fit various SDT computational models to eyewitness data and calculate and compare model-derived AUC. Crucially, `sdtlu` allows users to simulate data from fitted SDT models under different conditions (e.g., sample sizes), allowing for the possibility of power analysis. Mickes et al. (2021) have developed *pyWitness*, a Python package that features eyewitness lineup data wrangling/processing, fitting of SDT-based

computational models, and data simulation. Like `sdtlu` (A.L. Cohen et al., 2021), *pyWitness* allows users to estimate SDT model parameters from data and simulate new data using those model parameters, making power simulation possible.

Any power simulation method relies on certain assumptions (i.e., the data-generating processes). In the cases of `sdtlu` and *pyWitness*, the veracity of power simulation results relies at least in part on the validity of the computational models used to generate the data. When it comes to measuring eyewitness discriminability, researchers have pointed out that theory-based discriminability estimates can often differ from theory-free empirical discriminability estimates, and that empirical discriminability is to be preferred (Wixted & Mickes, 2018). Others have argued that non-SDT & non-memorial (e.g., social) variables can and do influence eyewitness performance, and that the influence of these variables may not be captured via SDT model-based analyses (Kovera & Evelo, 2021).

This is not to say that the above tools are inappropriate for ROC power analyses. If the underlying assumptions hold, then `sdtlu` and *pyWitness* will likely provide accurate estimates. However, regardless of the validity of the results of these programs, they require a moderate level of programming knowledge, have UIs that could be more user friendly, and do not provide a simple means for modular power analysis (e.g., simultaneously varying effect sizes and *N*s).

My primary goal with *powe(R)OC* was to create a user-friendly, intuitive and accessible R Shiny application that allows users to simulate a priori power or conduct post-hoc sensitivity simulations for typical eyewitness lineup experiments (and some recognition memory experiment designs) involving ROC analyses, while maximizing the amount of control users have over simulation parameters and minimizing the number of assumptions. In the next section, I will describe two main challenges in simulating ROC data for power, the solutions I implement in the *powe(R)OC* app, and the rationale for the chosen solutions.

Challenges of simulation-based ROC analyses (and solutions)

Challenge 1: The shape of the data

The first challenge in simulating ROC data is determining *how* to build the data from which to sample. For pAUC and DPP analyses (the focus of *powe(R)OC*), one needs to specify the number of culprit-present correct IDs and culprit-absent false IDs at each level of confidence. Ideally, simulated correct and false IDs distributed across confidence levels should resemble the data we actually collect—otherwise the simulations will not be useful. So, what do these data look like in the real world?

To find out, I downloaded open data from OSF for five recent, large-*N* eyewitness lineup papers comprising 8 separate experiments and 22 unique experimental conditions (Akan et al., 2021; Colloff et al., 2018, 2020; Colloff, Seale-Carlisle, et al., 2021; Colloff, Wilson, et al., 2021; Morgan et al., 2019)¹. Correct/false ID proportion data from 20/22 experimental conditions are displayed below in Figure 3 (two conditions from Akan et al., 2021 that used 3-point confidence scales excluded to restrict y-axis range):

Figure 3

Suspect ID proportions at each level of confidence in open datasets

¹ These open datasets are used in numerous examples here and in the application proper. References for which open data were used are denoted in the References list via a ¹.

Note. Proportions are out of all suspect identifications rather than all lineup responses. Colors denote different experiments, different lines indicate different experimental conditions, and the solid black line and gray ribbon are LOESS-smoothed best-fit lines and 95% CIs on the fits.

There are general trends in the distributions of ID proportions across confidence levels—a curvilinear relationship for culprit-absent suspect IDs and a positive trend for culprit-present suspect IDs. Despite these trends (which generally align with our current understanding of confidence-accuracy relationships; Brewer et al., 2021), there is a decent amount of condition-to-condition variability (e.g., in the mid-confidence range for culprit-absent IDs and the high-confidence range for culprit-present IDs). Importantly, it is likely

that this variability is due not only to pure sampling variability across experiments but substantive factors inherent to the experimental conditions/manipulations.

In simulation-based analyses, it is common to create synthetic datasets by simulating parameters of interest from specified distributions. For example, one might simulate data for a two-group t test by randomly sampling from two Gaussian distributions with specified M s and SD s. In the case of ROC analysis, one would need a distribution to sample vectors of correct/false ID probabilities across all confidence levels (with the constraint that the vector must sum to 1). Although there is a distribution that satisfies these requirements—the Dirichlet distribution—the specification of distribution parameters is not intuitive (e.g., it requires specifying a separate α parameter for each confidence level).

An alternative to the distribution-based method, and the one I have adopted in *powe(R)OC*, is to use pre-existing data as the basis for simulation. That is, users either upload their own data (e.g., pilot data) or make use of one of the aforementioned open datasets, which are included for convenience in the app. For instance, a user might conduct a smaller-scale pilot study and upload the data, or make use of an open dataset with a similar sample/manipulation to the planned experiment. The use of pre-existing data for ROC data simulation is not novel—this is the method also adopted in *sdtlu* (A.L. Cohen et al., 2021) and *pyWitness* (Mickes et al., 2021), perhaps in part due to the difficulties in creating ecologically-valid ROC data from scratch.

There are several advantages to using pre-existing data. For one, it does not require extensive knowledge of underlying distributions and parameters (and arguably, the specification of a Dirichlet distribution based on prior data already essentially requires a full pre-existing dataset anyway). In fact, one need not specify *any* data distributions or parameters with this method. Pre-existing data also allows one to simulate not only *suspect ID* vs. *not*, but the full complement of lineup identification decisions (i.e., *suspect ID* vs. *filler*

ID vs. rejection). Simultaneous simulation of Dirichlet distributions for all three decision types involves additional assumptions and parameter value decisions. Although the distinction between filler IDs and rejections is not relevant for pAUC or DPP analyses, it is for other ROC-based measures (e.g., full ROC curves; Smith et al., 2020; certain computational models; Wixted et al., 2018).

Additionally, the use of pre-existing data means that one need not make any assumptions about the data-generating processes or any commitment to theoretical mechanisms (e.g., whether the data were generated according to Gaussian signal detection theory processes, within-participant correlations between performance on multiple lineups²). Instead, the key assumption is that the pre-existing dataset that one uses resembles the data that one plans to collect. If this assumption does not hold, power estimates may be inaccurate, although the natural variation of Monte Carlo sampling provides some safeguards against estimates that are hyper-specific to a given pre-existing dataset. This key assumption is already made whenever one uses an effect size or other sample parameters from a previous study for power analysis. Additionally, the a priori simplicity of power analytic methods that do *not* use pre-existing data masks the fact that these analyses often make assumptions about the structure of the rest of the data (e.g., the nature of distributions at the response level). By using pre-existing data, these assumptions are made more explicit. Finally, the use of pre-existing data encourages pilot testing and (hopefully) wider sharing of eyewitness ROC data.

Challenge 2: Effect size specification

The next challenge in simulating power for ROC analysis is effect size specification. If one uses pre-existing data with two conditions, this is not an issue, as the effect of interest

² Standard ROC analyses do not consider within-participant correlations between performance on multiple lineups.

exists in the data. However, this means that users cannot power for effects that do not already exist in another dataset.

A more ideal solution would be to use pre-existing data (avoiding the challenges of ROC data specification) with the ability to apply custom effect sizes to the condition(s) of interest. Like many other effect size metrics (e.g., ANOVA's F), the test statistic for pAUC analysis, D , is not the most intuitive (Lakens & Caldwell, 2021). Specifically:

$$D = \frac{AUC1 - AUC2}{s}$$

...where s is the standard deviation of the bootstrapped differences between the two AUCs (Robin et al., 2011).

The primary challenge here lies in the mismatch between the effect size metric and the structure of the data. That is, how does one translate a D statistic into differential proportions of correct/false IDs at each confidence level? Even with the AUCs and s , there does not seem to be a clear way to translate this into differences at the base data level. And even if there was a way, D depends on where the ROC curves are truncated. For example, an absolute AUC difference of .001 means very different things if the ROCs terminate at a very low false ID rate (substantial difference) versus a very high false ID rate (trivial difference).

The solution I have adopted in *powe(R)OC* was to use a different effect size metric. This metric relates directly to pAUC values and works as follows: An effect size of 0 means the two pAUCs are equivalent, an effect size of .6 means that one condition has 60% more pAUC than the other, and an effect size of -.4 means that one condition has 40% less than the other. This effect size can be easily calculated from previous studies, and is more intuitive than D .

How to actually implement these effect sizes in simulation, at the data level? For any two identical ROC curves, applying a multiplier X to the number of correct IDs at each level of confidence also results in a pAUC that is X times as large as the original. Thus, this effect

size readily translates to changes at the base data level required for ROC data simulation³. A graphical example of practical effect size application (using data from Colloff, Wilson et al., 2021) is shown below in Figure 4.

Figure 4

Example effect sizes (using Colloff, Wilson, et al., 2021, Exp. 1, “Medium-similarity fillers” condition as the base)

This effect size measure allows users to reap the advantages of using pre-existing data while allowing flexibility in the kinds of differences they are interested in.

³ There are no modifiers applied to false IDs, primarily because the focus of the app is pAUC, which typically truncates/equates false ID rates. To apply custom effect sizes to data with differing false ID rates, currently users of *powe(R)OC* must upload a data file containing conditions with differing false ID rates, after which custom effect size(s) can be applied to the 2nd condition in the dataset.

Of course, this effect size metric relies on an assumption about empirical data—namely, that when two ROCs differ, the ratio of condition correct IDs will be relatively constant across confidence levels. To test the validity of this assumption, I examined the ratios in the open datasets. Figure 5 below plots the correct ID ratios for 13 condition comparisons.

Figure 5

Condition correct ID ratios by confidence in the literature

Note. Ribbons represent bootstrapped 50% percentile CIs on the proportions. 95% CIs were not used because ratios were often based on very few data points (e.g., < 10 IDs at a given level of confidence), resulting in inflated distribution tails. Comparisons are ordered by N (highest-to-lowest, top-to-bottom).

The figure above demonstrates variability in the correct ID ratio across confidence levels, although there do not seem to be any clear patterns. And for the highest- N studies (where ratios are more likely to be stable/accurate), the ratios are relatively consistent (albeit small) outside of the low-confidence range. Thus, when simulating data it may not be unreasonable to apply a constant effect size to all confidence levels, sacrificing some level of verisimilitude for parsimony and interpretability.

That said, there are clearly cases in which the effect size is *not* constant, and cases where there are theoretical reasons to predict varying effects at different confidence levels (e.g., one condition increases hits relative to the other, but only at higher confidence levels). The sample of datasets I've looked at here is also small, restricting any conclusions I might draw about the constancy (or not) of effects across confidence levels. As such, the current instantiation of *powe(R)OC* allows the specification of separate effect sizes at each confidence level in the base data. In this case, the effect size does not map 1:1 onto pAUC, and should instead be interpreted as applying to the # of hits at each confidence level. This trade-off in effect size interpretability allows additional flexibility in specifying ROC curves (e.g., if one has a good idea of how effects might differ across confidence levels).

powe(R)OC

powe(R)OC incorporates the methods and assumptions described above in an open-source, freely accessible application programmed in *R* (R Core Team, 2021) using the

shiny package (Chang et al., 2021). The app (including source code) can be accessed at <https://github.com/E-Y-M/poweROC>⁴.

Summary of current features

The current version of *powe(R)OC* allows users to:

- Upload their own dataset or select an open dataset to use as the basis for simulation, and visualize ROC curves from the selected dataset.
- Set a number of simulation parameters, including: specification of between- or within-subjects analysis, effect sizes (either constant or varying across confidence levels, with visualization of ROC curves), sample sizes to test, number of TA and TP lineups per subject, number of random samples & pROC/DPP tests to simulate for each combination of effect size and N , the number of bootstraps resamples to conduct for the pROC/DPP test, how partial AUCs are truncated, whether the tests are one-tailed or two-tailed, and the Type I error rate:

Figure 6

powe(R)OC simulation parameters page

⁴ The online version of the app linked on the main GitHub page is currently hosted on the free, open-source version of Shiny server which supports a limited number of concurrent users/sessions. Thus, it is recommended that users download and run a local copy (see the GitHub README file for instructions on how to do this).

- The app includes details about each parameter and where appropriate, recommended settings based on app testing & validation.

- View a growing database of effect sizes (in the app metric) in the literature to inform choices for simulation⁵. These effect sizes are taken from the included open datasets and a selection of other papers. Maintenance of this database is an ongoing project that carries value outside of its usefulness for power simulations (e.g., meta-analyses, literature reviews).
- Run the actual simulations, using the following method:

Figure 7

Illustration of simulation method

⁵ Papers from which effect sizes were taken are denoted in the reference list via a ².

For i effect sizes, j sample sizes, and k simulations:

Note. ROC curves are constructed via the method described in Gronlund et al., 2014⁶. pAUC is analyzed using the bootstrap method of `roc.test` in pROC (Robin et al., 2011). DPP is analyzed using the bootstrap method (and analysis code) in Smith et al. (2019).

- View simulation results and download a summary report (see Supplementary Material⁷ A for a full example of a report using Colloff, Wilson, et al., 2021, Exp. 1, “Medium-similarity fillers” condition as the base), including the hypothetical ROC

⁶ See also <https://mickeslab.com/handy/roc-tutorial/> for an excellent tutorial & walkthrough.

⁷ Supplementary Material can be downloaded at [https://github.com/E-Y-M/poweROC/blob/main/powe\(R\)OC%20manuscript%20Supplementary%20Material.docx](https://github.com/E-Y-M/poweROC/blob/main/powe(R)OC%20manuscript%20Supplementary%20Material.docx)

curves tested, simulation parameters for reproducible analyses, exact power estimates with empirical simulation pAUC and DPP condition values (and bootstrapped 95% percentile CIs), and power curves to visualize computed power as a function of effect and sample sizes.

- Anonymously upload power simulation results to a database linked to the app (via Google Sheets integration). These results do not include the user-uploaded raw data or ROCs but do list simulation parameters and pAUC/DPP values, so these results can be used to inform simulation parameter choices.
- View previous power simulation results uploaded by other users, with the ability to import simulation parameters (e.g., N_s , effect sizes) into the main simulation tab.
- View app testing & validation results, including tests of the ability of simulations to recover reported pAUC/DPP values, behaviour of the simulations under different parameter settings, an example of a full power simulation, and general use recommendations (see Supplementary Material B for a copy of these results).

Finally, although the primary focus of *powe(R)OC* is eyewitness lineup studies, there is nothing in principle preventing the use of this app for old/new recognition memory experiments where the primary outcome measures are ROC curves. In this case, “target-present” and “target-absent” lineups refer to the number of “old” and “new” trials per participant, with the only possible outcomes being suspect IDs and rejections.

Limitations & caveats

As mentioned previously, *powe(R)OC* operates on the core assumption that the data used as the basis for simulation will resemble the data one plans to collect. To the extent that this assumption does not hold, power estimates may not be accurate. This is not necessarily an assumption unique to *powe(R)OC*—any analytic or simulation-based power analysis will rely on the assumption that the collected data conform to assumptions made in the

calculations/simulations. This is a fact that is not often explicitly acknowledged but important to consider when conducting any power analysis.

That said, there are ways to both mitigate deviations from assumptions and determine the potential effects of deviations after the fact. For instance, *powe(R)OC*'s random sampling ensures that samples and ROC curves maintain important general characteristics while still allowing for variation around characteristics of the base data. The figure below illustrates simulation ROC curve variation across random samples for different sample sizes (again using data from Colloff, Wilson et al., 2021) is shown below in Figure 8:

Figure 8

Randomly-sampled ROC curves for a single condition (using Colloff, Wilson, et al., 2021, Exp. 1, "High-similarity fillers" condition as the base)

Note. Dark black line represents the actual ROC curve from the data, colored lines represent randomly-sampled ROC curves across 200 simulations, for condition *ns* of 250, 500, and 750.

These random samples replicate the experimental process of sampling from a population (i.e., the base data) with differing levels of precision (i.e., varying N s).

Aside from the natural variability inherent to random sampling, users can also run multiple simulations using different plausible candidate datasets as the base to examine how power estimates might vary (e.g., using data from two experiments examining unfair lineups). As far as *systematic* differences between simulated and collected data go, the effects of some differences are predictable. For example, a lower false ID rate in collected vs. simulated data will decrease power and a higher false ID rate will increase power, all else equal.

A more practical limitation of *powe(R)OC* is analysis time. Because of the multiplicative nature of the simulation loops (each N must be examined at each effect size), the large number of random samples to be generated, and the fact that for each sample the pAUC and DPP values must themselves be bootstrapped (the main simulation bottleneck), simulations can often take several hours to complete. *powe(R)OC* also requires a bit more work on the user side than other power analysis tools (e.g., G*Power; Faul et al., 2007), both in terms of background knowledge and parameter setting. However, the additional time and effort costs are arguably justified given the high cost and large samples necessary for most ROC experiments. The more demanding nature of *powe(R)OC* has the potential to translate into substantial savings in participant costs and data collection time. Additionally, users with minimal experience with R can download a copy of *powe(R)OC* to run analyses locally, reducing the risk of internet disruption during simulations.

Future directions

Version 1.0 of *powe(R)OC* includes all the features necessary for basic power simulation for pAUC/DPP analysis of eyewitness lineup data. However, there are several additional features that could be implemented in future versions. One is the inclusion of other eyewitness lineup measures. For instance, Smith et al. (2019, 2020) have criticized partial

AUC as a) treating filler picks and rejections as equivalent when they provide different diagnostic information and b) confounding discriminability with response bias by artificially truncating one of the ROC curves. In addition to DPP, these authors have more recently proposed full ROC curves (Smith et al., 2020) as an alternative that avoids these criticisms. Although *powe(R)OC* does not currently include simulation of full ROC curves (primarily due to additional challenges in simulating full ROC effect sizes), full ROC simulation is not in principle out of reach.

The relative merits of these measures and the ROC measurement debate in general are beyond the scope of this report, but there would certainly be value in incorporating these and other measures into *powe(R)OC* (if only to allow for simulation to better understand the statistical properties of these new measures). Eyewitness researchers are often interested in both theoretical outcomes (e.g., underlying discriminability) and applied outcomes (e.g., posterior probability of guilt), and use of multiple measures occasions power considerations for each measure of interest. Additional, non-ROC measures could be included, assuming there is a means of inferentially testing them (e.g., *CAC curves*⁸; Juslin et al., 1996; Mickes, 2015, diagnosticity ratios, d' and *multi- d'* ; Lee & Penrod, 2019, Bayesian measures such as *information-gain curves*, Wells, Yang, et al., 2015, applied outcome measures such as *expected costs*; Yang et al., 2019). Aside from the obvious benefit of being able to simulate power for multiple measures, the inclusion of multiple measures in *powe(R)OC* would offer the means to examine the relationships between the various measures (i.e., when they provide similar vs. different conclusions).

powe(R)OC's compendium of literature effect sizes provides useful benchmarks for users, but is far from comprehensive. As mentioned previously, the maintenance of this

⁸ CAC curves may already be amenable to simulation via existing programs that include t tests or ANOVA (e.g., Superpower; Lakens & Caldwell, 2021), despite the lack of an extant area-under-the-curve measure.

compendium is an ongoing effort. A complete picture of effect sizes in the literature would provide better guidelines for researchers and insights that are valuable beyond their use in power analysis. The same is true of the current collection of open datasets and user-uploaded simulation results. The more available datasets there are, the more options users have for power simulation. And as the number of uploaded simulation results increases, so too does our ability to examine relationships between power, pAUC/DPP, sample sizes, etc. For instance, it may be possible to extract reasonably accurate non-simulation-based heuristics from the simulation results. Ultimately, I hope that *power(R)OC* will serve as a useful tool and code resource, not only for informed sample size planning of eyewitness experiments, but also for further statistical exploration of ROC-based measures and general simulation of eyewitness lineup data.

Acknowledgments

I would like to acknowledge Adam Krawitz, Farouk Nathoo, and Steve Lindsay for their helpful feedback and suggestions in developing this app, and Ryan Fitzgerald for his invaluable feedback on both the app and manuscript draft.

Declarations

Funding

No funding was received to assist with the preparation of this manuscript.

Conflicts of interest

I have no known conflicts of interest to declare.

Ethics approval

Not applicable

Consent to participant

Not applicable

Consent for publication

Not applicable

Open Practices Statement

All *powe(R)OC* application source code and related materials (e.g., dataset testing, app validation results, and reports) are available at <https://github.com/E-Y-M/poweROC>. The program itself is open source under the GNU General Public License (Version 3). Also included at the link are the various open datasets used for application testing and demonstrations (i.e., those prepended with ¹ in the following references list).

References

- ^{1,2} Akan, M., Robinson, M. M., Mickes, L., Wixted, J. T., & Benjamin, A. S. (2021). The effect of lineup size on eyewitness identification. *Journal of Experimental Psychology: Applied*, 27(2), 369–392. <https://doi.org/10.1037/xap0000340>
- Mickes, L., Seale-Carlisle, T., Chen, X., & Boogert, S. (2021). *pyWitness*. <https://lmickes.github.io/pyWitness/index.html>
- Brewer, N., Lucas, C., Sauer, J. D., & Palmer, M. (2021). Measuring the relationship between eyewitness identification confidence and accuracy. In A. Smith, M. Toggia, and J. M. Lampinen (Eds.), *Methods, Measures, and Theories in Eyewitness Identification Tasks*. Taylor and Francis Group.
- ²Carlson, C. A., & Carlson, M. A. (2014). An evaluation of lineup presentation, weapon presence, and a distinctive feature using ROC analysis. *Journal of Applied Research in Memory and Cognition*, 3(2), 45–53. <https://doi.org/10.1016/j.jarmac.2014.03.004>
- Chang, W., Cheng, J., Allaire, J.J., Sievert, C., Schloerke, B., Xie, Y., Allen, J., McPherson, J., Dipert, A., & Borges, B. (2021). shiny: Web Application Framework for R. R package version 1.6.0.
- Cohen, A. L., Starns, J. J., & Rotello, C. M. (2021). sdtlu: An R package for the signal detection analysis of eyewitness lineup data. *Behavior Research Methods*, 53(1), 278–300. <https://doi.org/10.3758/s13428-020-01402-7>
- Cohen, J. (1992). Statistical Power Analysis. *Current Directions in Psychological Science*, 1(3), 98–101. <https://doi.org/10.1111/1467-8721.ep10768783>
- ^{1,2}Colloff, M. F., Flowe, H. D., Smith, H. M. J., Seale-Carlisle, T. M., Meissner, C. A., Rockey, J. C., Pande, B., Kujur, P., Parveen, N., Chandel, P., Singh, M. M., Pradhan, S., & Parganiha, A. (2020). *Active exploration of faces in police lineups increases discrimination accuracy* [Preprint]. PsyArXiv. <https://doi.org/10.31234/osf.io/tvga4>

- ^{1,2}Colloff, M. F., Seale-Carlisle, T. M., Karoğlu, N., Rockey, J. C., Smith, H. M. J., Smith, L., Maltby, J., Yaremenko, S., & Flowe, H. D. (2021). Perpetrator pose reinstatement during a lineup test increases discrimination accuracy. *Scientific Reports*, *11*(1), 13830. <https://doi.org/10.1038/s41598-021-92509-0>
- ²Colloff, M. F., Wade, K. A., & Strange, D. (2016). Unfair Lineups Make Witnesses More Likely to Confuse Innocent and Guilty Suspects. *Psychological Science*, *27*(9), 1227–1239. <https://doi.org/10.1177/0956797616655789>
- ^{1,2}Colloff, M. F., Wade, K. A., Strange, D., & Wixted, J. T. (2018). Filler-Siphoning Theory Does Not Predict the Effect of Lineup Fairness on the Ability to Discriminate Innocent From Guilty Suspects: Reply to Smith, Wells, Smalarz, and Lampinen (2018). *Psychological Science*, *29*(9), 1552–1557. <https://doi.org/10.1177/0956797618786459>
- ^{1,2}Colloff, M. F., Wilson, B. M., Seale-Carlisle, T. M., & Wixted, J. T. (2021). Optimizing the selection of fillers in police lineups. *Proceedings of the National Academy of Sciences*, *118*(8), e2017292118. <https://doi.org/10.1073/pnas.2017292118>
- Colloff, M. F., & Wixted, J. T. (2020). Why are lineups better than showups? A test of the filler siphoning and enhanced discriminability accounts. *Journal of Experimental Psychology: Applied*, *26*(1), 124–143. <https://doi.org/10.1037/xap0000218>
- ²Dobolyi, D. G., & Dodson, C. S. (2013). Eyewitness confidence in simultaneous and sequential lineups: A criterion shift account for sequential mistaken identification overconfidence. *Journal of Experimental Psychology: Applied*, *19*(4), 345–357. <https://doi.org/10.1037/a0034596>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*(2), 175–191. <https://doi.org/10.3758/BF03193146>

- ²Gronlund, S. D., Carlson, C. A., Neuschatz, J. S., Goodsell, C. A., Wetmore, S. A., Wooten, A., & Graham, M. (2012). Showups versus lineups: An evaluation using ROC analysis. *Journal of Applied Research in Memory and Cognition*, 1(4), 221–228.
<https://doi.org/10.1016/j.jarmac.2012.09.003>
- Gronlund, S. D., Wixted, J. T., & Mickes, L. (2014). Evaluating Eyewitness Identification Procedures Using Receiver Operating Characteristic Analysis. *Current Directions in Psychological Science*, 23(1), 3–10. <https://doi.org/10.1177/0963721413498891>
- Hillis, S. L., Obuchowski, N. A., & Berbaum, K. S. (2011). Power Estimation for Multireader ROC Methods. *Academic Radiology*, 18(2), 129–142.
<https://doi.org/10.1016/j.acra.2010.09.007>
- Hillis, S. L., & Scharzt, K. M. (2018). Multireader sample size program for diagnostic studies: Demonstration and methodology. *Journal of Medical Imaging*, 5(04), 1.
<https://doi.org/10.1117/1.JMI.5.4.045503>
- Kim, E., Zhang, Z., Wang, Y., & Zeng, D. (2014). Power calculation for comparing diagnostic accuracies in a multi-reader, multi-test design: Analysis of Correlated ROC Areas. *Biometrics*, 70(4), 1033–1041. <https://doi.org/10.1111/biom.12240>
- Kovera, M. B., & Evelo, A. J. (2021). Eyewitness Identification in Its Social Context. *Journal of Applied Research in Memory and Cognition*, S2211368121000292.
<https://doi.org/10.1016/j.jarmac.2021.04.003>
- Lakens, D., & Caldwell, A. R. (2021). Simulation-Based Power Analysis for Factorial Analysis of Variance Designs. *Advances in Methods and Practices in Psychological Science*, 4(1), 251524592095150. <https://doi.org/10.1177/2515245920951503>
- Lee, J., & Penrod, S. D. (2019). New signal detection theory-based framework for eyewitness performance in lineups. *Law and Human Behavior*, 43(5), 436–454.
<https://doi.org/10.1037/lhb0000343>

- Mickes, L. (2015). Receiver operating characteristic analysis and confidence–accuracy characteristic analysis in investigations of system variables and estimator variables that affect eyewitness memory. *Journal of Applied Research in Memory and Cognition*, 4(2), 93–102. <https://doi.org/10.1016/j.jarmac.2015.01.003>
- ²Mickes, L., Flowe, H. D., & Wixted, J. T. (2012). Receiver operating characteristic analysis of eyewitness memory: Comparing the diagnostic accuracy of simultaneous versus sequential lineups. *Journal of Experimental Psychology: Applied*, 18(4), 361–376. <https://doi.org/10.1037/a0030609>
- ^{1,2}Morgan, D. P., Tamminen, J., Seale-Carlisle, T. M., & Mickes, L. (2019). The impact of sleep on eyewitness identifications. *Royal Society Open Science*, 6(12), 170501. <https://doi.org/10.1098/rsos.170501>
- Morris, T. P., White, I. R., & Crowther, M. J. (2019). Using simulation studies to evaluate statistical methods. *Statistics in Medicine*, 38(11), 2074–2102. <https://doi.org/10.1002/sim.8086>
- Negida, A., Fahim, N. K., & Negida, Y. (2019). Sample Size Calculation Guide - Part 4: How to Calculate the Sample Size for a Diagnostic Test Accuracy Study based on Sensitivity, Specificity, and the Area Under the ROC Curve. *Advanced Journal of Emergency Medicine*, 3(3), e33. <https://doi.org/10.22114/ajem.v0i0.158>
- R Core Team (2021). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C., & Müller, M. (2011). pROC: An open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics*, 12(1), 77. <https://doi.org/10.1186/1471-2105-12-77>
- ²Seale-Carlisle, T. M., & Mickes, L. (2016). US line-ups outperform UK line-ups. *Royal*

Society Open Science, 3(9), 160300. <https://doi.org/10.1098/rsos.160300>

Smith, A. M., Lampinen, J. M., Wells, G. L., Smalarz, L., & Mackovichova, S. (2019).

Deviation from Perfect Performance Measures the Diagnostic Utility of Eyewitness Lineups but Partial Area Under the ROC Curve Does Not. *Journal of Applied Research in Memory and Cognition*, 8(1), 50–59.

<https://doi.org/10.1016/j.jarmac.2018.09.003>

Smith, A. M., Yang, Y., & Wells, G. L. (2020). Distinguishing Between Investigator

Discriminability and Eyewitness Discriminability: A Method for Creating Full Receiver Operating Characteristic Curves of Lineup Identification Performance.

Perspectives on Psychological Science, 15(3), 589–607.

<https://doi.org/10.1177/1745691620902426>

Wells, G. L., Smalarz, L., & Smith, A. M. (2015). ROC analysis of lineups does not measure

underlying discriminability and has limited value. *Journal of Applied Research in*

Memory and Cognition, 4(4), 313–317. <https://doi.org/10.1016/j.jarmac.2015.08.008>

Wells, G. L., Yang, Y., & Smalarz, L. (2015). Eyewitness identification: Bayesian

information gain, base-rate effect equivalency curves, and reasonable suspicion. *Law and Human Behavior*, 39(2), 99–122. <https://doi.org/10.1037/lhb0000125>

²Wetmore, S. A., Neuschatz, J. S., Gronlund, S. D., Wooten, A., Goodsell, C. A., & Carlson,

C. A. (2015). Effect of retention interval on showup and lineup performance. *Journal of Applied Research in Memory and Cognition*, 4(1), 8–14.

<https://doi.org/10.1016/j.jarmac.2014.07.003>

Wixted, J. T., & Mickes, L. (2012). The Field of Eyewitness Memory Should Abandon

Probative Value and Embrace Receiver Operating Characteristic Analysis.

Perspectives on Psychological Science, 7(3), 275–278.

<https://doi.org/10.1177/1745691612442906>

- Wixted, J. T., & Mickes, L. (2014). A signal-detection-based diagnostic-feature-detection model of eyewitness identification. *Psychological Review*, 121(2), 262–276.
<https://doi.org/10.1037/a0035940>
- Wixted, J. T., & Mickes, L. (2015). ROC analysis measures objective discriminability for any eyewitness identification procedure. *Journal of Applied Research in Memory and Cognition*, 4(4), 329–334. <https://doi.org/10.1016/j.jarmac.2015.08.007>
- Wixted, J. T., & Mickes, L. (2018). Theoretical vs. empirical discriminability: The application of ROC methods to eyewitness identification. *Cognitive Research: Principles and Implications*, 3(1), 9. <https://doi.org/10.1186/s41235-018-0093-8>
- Yang, Y., Smalarz, L., Moody, S. A., Cabell, J. J., & Copp, C. J. (2019). An expected cost model of eyewitness identification. *Law and Human Behavior*, 43(3), 205–219.
<https://doi.org/10.1037/lhb0000331>