

METHODOLOGY

Open Access

Statistical stopping criteria for automated screening in systematic reviews



Max W Callaghan^{1,2,3*}  and Finn Müller-Hansen^{1,3}

Abstract

Active learning for systematic review screening promises to reduce the human effort required to identify relevant documents for a systematic review. Machines and humans work together, with humans providing training data, and the machine optimising the documents that the humans screen. This enables the identification of all relevant documents after viewing only a fraction of the total documents. However, current approaches lack robust stopping criteria, so that reviewers do not know when they have seen all or a certain proportion of relevant documents. This means that such systems are hard to implement in live reviews. This paper introduces a workflow with flexible statistical stopping criteria, which offer real work reductions on the basis of rejecting a hypothesis of having missed a given recall target with a given level of confidence. The stopping criteria are shown on test datasets to achieve a reliable level of recall, while still providing work reductions of on average 17%. Other methods proposed previously are shown to provide inconsistent recall and work reductions across datasets.

Keywords: Systematic review, Machine learning, Active learning, Stopping criteria

Background

Evidence synthesis technology is a rapidly emerging field that promises to change the practice of evidence synthesis work [1]. Interventions have been proposed at various points in order to reduce the human effort required to produce systematic reviews and other forms of evidence synthesis. A major strand of the literature works on screening: the identification of relevant documents in a set of documents whose relevance is uncertain [2]. This is a time-consuming and repetitive task, and in a research environment with constrained resources and increasing amounts of literature, this may limit the scope of the evidence synthesis projects undertaken. Several papers have developed active learning (AL) approaches [3–7] to reduce the time required to screen documents. This paper sets out how current approaches are unreliable in practice,

and outlines and evaluates modifications that would make AL systems ready for live reviews.

Active learning is an iterative process where documents screened by humans are used to train a machine learning model to predict the relevance of unseen papers [8]. The algorithm chooses which studies will next be screened by humans, often those which are likely to be relevant or about which the model is uncertain, in order to generate more labels to feed back to the machine. By prioritising those studies most likely to be relevant, a human reviewer most often identifies all relevant studies—or a given proportion of relevant studies (described by recall: the number of relevant studies identified divided by the total number of relevant studies)—before having seen all the documents in the corpus. The proportion of documents not yet seen by the human when they reach the given recall threshold is referred to as the work saved. This represents the proportion of documents that they do not have to screen, which they would have had to without machine learning.

Machine learning applications are often evaluated using sets of documents from already completed systematic

*Correspondence: callaghan@mcc-berlin.net

¹Mercator Research Institute on Global Commons and Climate Change, EUREF Campus 19, Torgauer Straße 12-15, 10829 Berlin, Germany

²Priestley International Centre for Climate, University of Leeds, Leeds, LS2 9JT, UK

Full list of author information is available at the end of the article



© The Author(s). 2020 **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

reviews for which inclusion or exclusion labels already exist. As all human labels are known a priori, it is possible to simulate the screening process, recording when a given recall target has been achieved. In live review settings, however, recall remains unknown until all documents have been screened. In order for work to really be saved, reviewers have to stop screening while uncertain about recall. This is particularly problematic in systematic reviews because low recall increases the risk of bias [9]. The lack of appropriate stopping criteria has therefore been identified as a research gap [10, 11], although some approaches have been suggested. These have most commonly fallen into the following categories:

- *Sampling criteria*: Reviewers estimate the number of relevant documents by taking a random sample at the start of the process. They stop when this number, or a given proportion of it, has been reached [12].
- *Heuristics*: Reviewers stop when a given number of irrelevant articles are seen in a row [6, 7].
- *Pragmatic criteria*: Reviewers stop when they run out of time [3].
- *Novel automatic stopping criteria*: Recent papers have proposed more complicated novel systems for automatically deciding when to stop screening [13–15].

We review the first three classes of these methods in the following section and discuss their theoretical limitations. They are then tested on several previous systematic review datasets. We demonstrate theoretically and with our experimental results that these three classes of methods cannot deliver consistent levels of work savings or recall—particularly across different domains, or datasets with different properties [2]. We also discuss the limitations of novel automatic stopping criteria, which have all demonstrated promising results, but do not achieve a given level of recall in a reliable or reportable way. Without the reliable or reportable achievement of a desired level of recall, deployment of AL systems in live reviews remains challenging.

This study proposes a system for estimating the recall based on random sampling of remaining documents. We use a simple statistical method to iteratively test a null hypothesis that the recall achieved is less than a given target recall. If the hypothesis can be rejected, we conclude that the recall target has been achieved with a given confidence level and screening can be stopped. This allows AL users to predefine a target in terms of uncertainty and recall, so that they can make transparent, easily communicable statements like “We reject the null hypothesis that we achieve a recall of less than 95% with a significance level of 5%”.

In the remainder of the paper, we first discuss in detail the shortcomings of existing stopping criteria. Then, we

introduce our new criteria based on a hypergeometric test. We evaluate our stopping criteria and compare their performance with heuristic- and sampling-based criteria on real-world systematic review datasets on which AL systems have previously been tested [13, 16–18].

Methods review

We start by explaining the sampling- and heuristic-based stopping criteria and discussing their methodological limitations.

Sampling-based stopping criteria

The stopping criterion suggested by Shemilt et al. [12] involves establishing the Baseline Inclusion Rate (BIR), by taking a random sample at the beginning of screening. The BIR is used to estimate the number of relevant documents in the whole dataset. Reviewers continue to screen until this number, or a proportion of it corresponding to the desired level of recall, is reached.

However, the estimation of the BIR fails to correctly take into account sampling uncertainty.¹ This uncertainty is crucial, as errors can have severe consequences. Let us assume that users will stop screening when they have identified 95% of the relevant number of documents. If the estimated number of relevant documents is more than the true number of relevant documents divided by 0.95, then the users will never see 95% of the estimated number. This means that they will keep screening until they have seen all documents, and no work savings will be achieved. Conversely, if the number of relevant documents is underestimated by even a single unit, then the recall achieved will be lower than the target.

The number of relevant documents drawn without replacement from a finite sample of documents follows the hypergeometric distribution. Figure 1a shows the distribution of the predicted number of documents after drawing 1000 documents from a total of 20,000 documents, where 500 documents (2.5%) are relevant. The left shaded portion of the graph shows all the cases where the recall will be less than 95%. This occurs 48% of the time. The right shaded portion of the graph shows the cases where the number of relevant documents is overestimated so much that no work savings could be made to achieve a target recall of 95%. This occurs 29% of the time. In only 23% of cases can work savings be achieved while still achieving a recall of at least 95%.

¹Although Shemilt et al. [12] employ a method to choose a sample size based on uncertainty, they fail to acknowledge the potential implications for recall of their choice. Their margin of error of 0.0025 and observed proportion of relevant studies of 0.0005 translate to estimates of 400 ± 451 relevant results. To reduce the margin of error to $\pm 5\%$ of estimated relevant studies, they would have had to screen 638,323 out of 804,919 results. See the notebook https://github.com/mcallaghan/rapid-screening/blob/master/analysis/bir_theory.ipynb that accompanies this paper for a detailed discussion.




Figure 1b shows the probability distribution of these errors according to the sample size. Even with very large samples, both types of error remain frequent. This shows how baseline estimation inevitably offers poor reliability, either in terms of recall or in work saved.

Heuristic stopping criteria

Some studies give the example of heuristic stopping criteria based on drawing a given number of irrelevant articles in a row [6, 7]. We take this as a proxy for estimating that the proportion of remaining documents that are relevant in the unseen documents is low, as the probability of observing 0 relevant documents in a given sample (analogous to a set of consecutive irrelevant results) is a decreasing function of the number of relevant documents in the population. We find this a promising intuition, but argue that (1) it ignores uncertainty, as discussed in relation to the previous method; (2) it lacks a formal description that would help to find a suitable threshold for the criterion; and (3) it misunderstands the significance of a low proportion of relevant documents in estimating the recall.

Figure 2 illustrates this third point. We show two scenarios with identical low proportions of relevant documents observed in the unseen documents. In the top part of the figure, machine learning (ML) has performed well, and 74% of the screened documents were relevant. In the bottom part of the figure, ML has performed less well, and only 26% of the screened documents were relevant. In both cases, only 2% of unseen documents are relevant, but 2% of a larger number means more relevant documents are missed. Recall is not simply a function of the proportion of unseen documents that are relevant, but also of the number of unseen documents. This also means that where ML has performed well (as in the top figure), a low proportion of relevant documents in those not yet checked is indicative of lower recall than where ML has performed less well. Likewise, where the proportion of relevant documents in the whole corpus is low, a simi-

larly low proportion of relevant documents is likely to be observed, even when true recall is low. This shows us that even a perfect estimator of the proportion of unseen documents that are relevant is insufficient on its own to provide sufficient information about when to stop screening. To estimate recall reliably, it is necessary to take into account the total number of unseen relevant documents (or their proportion times the number of unseen documents).

Pragmatic stopping criteria

Wallace et al. [4] develop a “simple, operational stopping criterion”: stopping after half the documents have been screened. Although the criterion worked in their experiment, it is unclear how this could be generalised, and its development depended on knowledge of the true relevance values. Jonnalagadda and Petitti [6] note that “the reviewer can elect to end the process of classifying documents at any point, recognizing that stopping before reviewing all documents involves a trade-off of lower recall for reduced workload”, although clearly the reviewer lacks information about probable recall.

Novel automatic stopping criteria

Two examples come from the information retrieval literature. Di Nunzio [14] presents a novel automatic stopping criterion based on BM25, although recall reported is “often between 0.92 and 0.94 and consistently over 0.7”. Yu and Menzies [13] also present a stopping criterion based on BM25 which allows the user to target a specific level of recall. However, reviewers are not given the opportunity to specify a confidence level, and for two of the four datasets in which they tested their criteria, the median achieved recall at a stopping criteria targeting 95% recall was below 95%. In each case, the reliability of the estimate is dependent on the performance of the model.

Finally, Howard et al. [15] present a method to estimate recall based on the number of irrelevant documents D observed in a list of documents since the δ th previous rele-

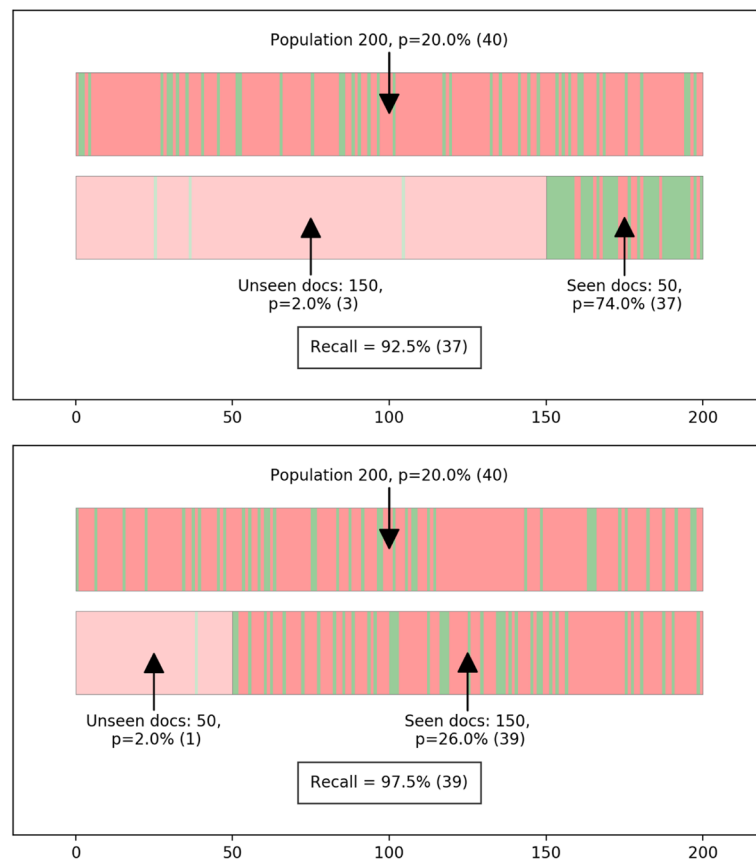


Fig. 2 Similar low proportions of relevant documents in unseen documents with different consequences for recall. The top bar shows a random distribution of relevant documents (green) and irrelevant documents (red) at a given proportion of relevance. The bottom bar shows distributions of relevant and irrelevant documents in hypothetical sets of seen (right) and unseen (left—transparent) documents

vant document. They reason that this should follow the negative binomial distribution based on the proportion of remaining relevant documents p , and use this information to estimate \hat{p} , and with this, the total number of relevant articles and the estimated recall.

However, their method does not quantify uncertainty, but can only claim that the method “tends to result in a conservative estimate of recall” (emphasis ours). This is not guaranteed by the criterion itself but rather a finding of the simulation with example datasets. Further, the authors do not give sufficient information to reproduce their results, providing neither code (they describe their own proprietary software) nor an equation for \hat{p} . Additionally, the criterion requires a tuning parameter δ , which users may have insufficient information to set optimally. Lastly, because screening is a form of sampling without replacement, the negative hypergeometric distribution should be preferred to the negative binomial, even though the latter can be a good approximation for cases with large numbers of documents.

These last examples are promising developments, but they all fail to take into account the needs of live sys-

tematic reviews, where the reliability of and ease of communication about recall are paramount, and the results must be independent of model performance. In the following, we explain our own method, which provides clearly communicable estimates of recall, and which manages uncertainty in a way robust to model performance.

Methods

A statistical stopping criterion for active learning

In our screening setup, we start off with N_{tot} documents that are potentially relevant. ρ_{tot} of these documents are actually relevant, but we do not know this value a priori. As we screen relevant documents, we include them, so ρ_{seen} represents the number of relevant documents screened, and recall τ is given by:

$$\tau = \frac{\rho_{seen}}{\rho_{tot}} \tag{1}$$

We set a target recall τ_{tar} and a confidence level α . We want to keep screening until $\tau \geq \tau_{tar}$, and devise a hypothesis test to estimate whether this is the case with a given level of confidence (Fig. 3). We do this based on

interrupting the active-learning process and drawing a random sample from the remaining unseen documents. We first describe this test, before showing how a variation on the test can be used to decide when to begin drawing a random sample.

Random sampling

At the start of the sample, N_{AL} is the number of documents seen during the active learning process, and N is the number of documents remaining, so that:

$$N = N_{tot} - N_{AL} \tag{2}$$

We refer to the number of relevant documents seen during active learning as ρ_{AL} , and the number of remaining relevant documents as K . We do not know the value of K but know that it is given by the total number of relevant documents minus the number of relevant documents seen during active learning.

$$K = \rho_{tot} - \rho_{AL} \tag{3}$$

We now take random draws from the remaining N documents, and denote the number of documents drawn with n and the number of relevant documents drawn with k . The number of relevant documents seen is updated by

adding the number of relevant documents seen since sampling began to the number of relevant documents seen during active learning.

$$\rho_{seen} = \rho_{AL} + k \tag{4}$$

We proceed to form a null hypothesis that the true value of recall is less than our target recall:

$$H_0 : \tau < \tau_{tar} \tag{5}$$

Accordingly, the alternative hypothesis is that recall is equal to or greater than our target:

$$H_1 : \tau \geq \tau_{tar} \tag{6}$$

Because we are sampling without replacement, we can use the hypergeometric distribution to find out the probability of observing k relevant documents in a sample of n documents from a population of N documents of which K are relevant. We know that k is distributed hypergeometrically:

$$k \sim \text{Hypergeometric}(N, K, n) \tag{7}$$

We introduce a hypothetical value for K , which we call K_{tar} . This represents the minimum number of relevant documents remaining at the start of sampling compatible




Fig. 3 A workflow for active learning in screening with a statistical stopping criterion

with our null hypothesis that recall is below our target.

$$K_{tar} = \lfloor \frac{\rho_{seen}}{\tau_{tar}} - \rho_{AL} + 1 \rfloor \quad (8)$$

This equation is derived by combining Eqs. 1 and 4. Because k can only take integer values, K_{tar} is the smallest integer that satisfies the inequality in Eq. 5. With K_{tar} , we can reformulate our null hypothesis: the true number of relevant documents in the sample is greater than or equal to our hypothetical value.

$$H_0 : K \geq K_{tar} \quad (9)$$

We test this by calculating the probability of observing k or fewer relevant documents from the hypergeometric distribution given by K_{tar} , using the cumulative probability mass function.

$$p = P(X \leq k), \text{ where } X \sim \text{Hypergeometric}(N, K_{tar}, n) \quad (10)$$

Because the cumulative probability mass function $P(X \leq k)$ is decreasing with increasing K , this gives the maximum probability of observing k for all values of K compatible with our null hypothesis. Similar arguments have been made to derive confidence intervals for estimating the parameter K in the hypergeometric distribution function [19, 20], and the derivation of an equivalent criterion could use the upper limit of such a confidence interval of an estimated K from the observation of k .

We can reject our null hypothesis and stop screening if the maximum probability of obtaining our observed results given our null hypothesis p is below $1 - \alpha$.² To further investigate the accuracy of the test, we perform an experiment drawing 1 million random samples in 6 scenarios with different characteristics. We vary the value of ρ_{AL} to simulate starting random sampling with different levels of recall achieved.

Figure 4 shows that in each case, as long as recall is lower than the target recall when sampling begins, the percentage of trials in which the criterion is triggered to early is within two tenths of a percentage point of 5% and the 5th percentile of achieved recall values is within two tenths of a percentage point of the target recall 95%.

Ranked quasi-sampling

We now proceed to describe a special case of the method described above which we (1) use as a heuristic in order to decide when to begin random sampling and (2) test as an independent stopping criterion. The method works by treating batches of previously screened documents as if they were random samples.

We calculate p as above for subsets of the already screened documents. Concretely, we use subsets of documents A_i by looking back to the last i documents, $A_i = \{d_{N_{seen}-1}, \dots, d_{N_{seen}-i}\}$, where the documents d are indexed in the order in which they have been screened. For a specific i , this corresponds to random sampling beginning after seeing i documents in the section above. Thus, we set N_{AL} to i , n to $N_{seen} - i$, ρ_{AL} to the number of relevant documents seen when i documents had been seen, and k to the number of relevant documents seen since i documents had been seen, and calculate p according to Eq. 10. We compute p for all sets A_i with $i \in N_{seen} - 1 \dots 1$. This gives us a vector p , representing the values of p which would have been estimated at each point at which we could have stopped active learning and began random sampling. The point at which the p value for our null hypothesis is lowest is given by p_{min} . With the vectorised implementation included in our accompanying code, these calculations are completed in less than the time it would take a human to code the next document.

First, we use this method as a useful heuristic for deciding when to stop active learning, and switch to random sampling. For this, we choose a higher threshold for the likelihood, $p_{min} < 1 - \frac{\alpha}{2}$. Second, we use the same ranked quasi-sampling as an independent stopping criterion, by continuing screening with active learning until $p_{min} < 1 - \alpha$. We present the results of this second procedure separately below.

Given that the documents seen during active learning are ranked according to predicted relevance, they do not in fact represent a random sample. This means that the test is unlikely to be accurate. It would be reasonable to assume that the proportion of relevant documents in each ranked quasi-sample is as high if not higher than the proportion of relevant documents in the unseen documents. This assumption would make this estimator conservative. As such, it works in a similar way to the criterion proposed by Howard et al. [15], although it makes use of more information and provides hypothesis testing rather than just a point estimate of recall.

Evaluation

We evaluate each of the criteria discussed on real-world test data, operationalising the heuristic stopping criteria with 50, 100, and 200 consecutive irrelevant records. We run 100 iterations on each dataset and record the following measures:

- *Actual Recall*: The recall when the stopping criteria were met.
- *WS-SC*: Work saved when the stopping criteria were met.
- *Additional Burden*: The work saved when the criterion was triggered subtracted from the work saved when the recall target was actually achieved.

²The notebook, https://github.com/mcallaghan/rapid-screening/blob/master/analysis/hyper_criteria_theory.ipynb, in the github repository accompanying this paper contains a step by step explanation of this method with code and examples.

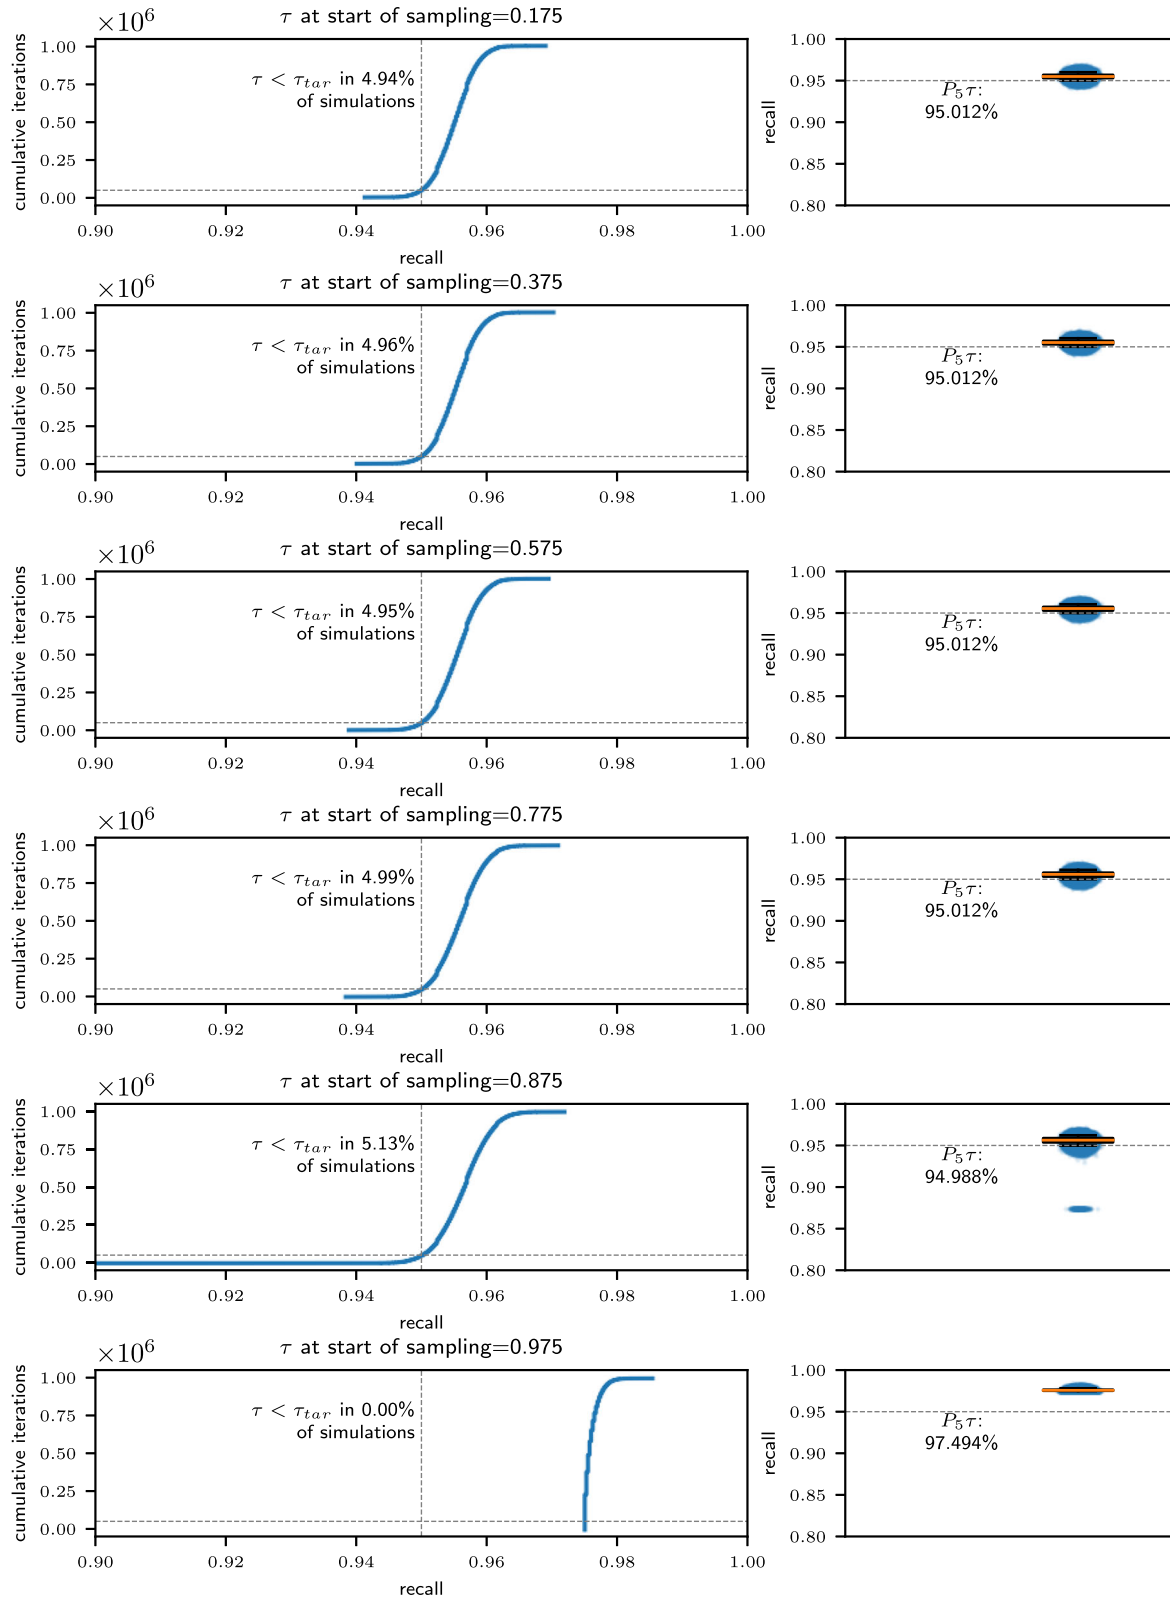


Fig. 4 The distribution of achieved recall values given our random sampling stopping criterion for 6 scenarios with different recall values at the start of sampling

For simplicity, we use a basic SVM model [21, 22], with 1–2 word n-grams taken from the document abstracts used as input data. We start with random samples of 200 documents (we do not employ Shemilt et al.'s methods for identifying the “optimal” sample size, as we showed these in the “Methods review” section to be unhelpful). Subsequently, we “screen”, that is, we reveal the labels of, batches of the 20 documents with the highest predicted relevance scores, retraining the model after each batch. Theoretically, using smaller batch sizes could mean that the recall target is achieved more quickly, but this is a trade-off between computational time spent training and the speed at which the algorithm can “learn”. However, this is a modelling choice which may affect work saved, but not recall. Each criterion is evaluated after each document is “screened”. For our criteria, we set the target recall value to 95% and the confidence level to 95%.

The systematic review datasets used for testing are described in Table 1. We use the seminal collection of systematic reviews used to develop machine learning applications for document screening by Cohen and co-authors in 2006 [16], along with the widely used Proton Beam [17] and COPD [18] datasets, and computer science datasets used to test FASTREAD [13]. Testing on datasets with different properties and from different domains is key to

establishing criteria appropriate for general use. Choosing as broad as possible data also prevents us from being able to “tune” our machine learning approach in ways that may work well for specific datasets but not generalise well. Work savings, even maximum work savings, are therefore below the state of the art recorded for each of these datasets. In this way, we can show how well the criteria perform even when the model performs badly.

All computational steps required to reproduce this analysis are documented online at <https://github.com/mcallaghan/rapid-screening>.

Results

Figure 5 shows the actual recall and work savings achieved when each stopping criterion has been satisfied. For comparison, we also include the results that would have been achieved with a priori knowledge of the data, that is, the work saved when the 95% recall target was actually reached. In a live systematic review, reviewers would never know when this had been reached, but these are the work savings most often reported in machine learning for systematic review screening studies.

Both the random sampling and the ranked sampling criteria achieve the target threshold of 95% in more than 95% of cases. That this is greater than 95% is accounted

Table 1 Dataset properties

	Dataset	Data source	N documents	N relevant documents	Proportion relevant documents
0	UrinaryIncontinence	cohen	284	68	0.24
1	Antihistamines	cohen	287	90	0.31
2	Estrogens	cohen	349	79	0.23
3	NSAIDS	cohen	358	83	0.23
4	OralHypoglycemics	cohen	475	135	0.28
5	Triptans	cohen	594	205	0.35
6	ADHD	cohen	803	83	0.10
7	AtypicalAntipsychotics	cohen	1030	333	0.32
8	CalciumChannelBlockers	cohen	1103	257	0.23
9	ProtonPumpInhibitors	cohen	1210	227	0.19
10	SkeletalMuscleRelaxants	cohen	1348	30	0.02
11	COPD	copd_pb	1443	179	0.12
12	Kitchenham	fastread	1700	45	0.03
13	Opioids	cohen	1769	43	0.02
14	BetaBlockers	cohen	1872	270	0.14
15	ACEInhibitors	cohen	2234	168	0.08
16	Statins	cohen	2743	152	0.06
17	ProtonBeam	copd_pb	4108	240	0.06
18	Radjenovic	fastread	5999	47	0.01
19	Wahono	fastread	7002	62	0.01
20	Hall	fastread	8911	104	0.01




Fig. 5 Distribution of recall and work saved after each stopping criteria. Green dots show results for datasets with less than 1000 documents, orange dots show datasets with 1000–2000 documents, and blue dots show datasets with more than 2000 documents




Fig. 6 Distribution of recall and additional burden after each stopping criterion. Additional burden is the work saved when the criterion was triggered minus the work saved when the target was reached. Colouring of data points as in Fig. 5

for by the fact that random sampling sometimes begins after the target recall has been achieved, in which case the null hypothesis would be a priori impossible. The ranked quasi-sampling criterion outperforms the random sampling criterion with respect to both recall and work savings, saving a mean of 17% of the work compared to 15%, and missing the target in only 0.95% compared to 3.29% of cases. In theory, the ranked sampling criterion is conservative if the assumption holds that documents chosen by machine learning are not less likely to be relevant than those chosen at random. Based on our experiments, this assumption seems reasonable and accounts for the higher recall. Because the ranked quasi-sampling criterion can flexibly choose its sample, whereas the random criterion has to wait for a random sample to be triggered, the criterion is also triggered earlier, as it can make use of more data. This accounts for the higher work savings.

The baseline sampling criterion (Fig. 5c) misses the 95% recall target in 39.67% of cases, while the most common work saving is 0%. This is in line with our expectations that, due to random sampling error, the expected number of documents will often be overestimated or underestimated, resulting in zero work savings or poor recall.

The heuristic stopping criteria, both for 50 consecutive irrelevant results (Fig. 5d—IH50) and for 200 irrelevant results (Fig. 5e), also perform unreliably. Although the mean work saved for IH50 is 41%, the target is missed

in 39% of cases. The cases below the horizontal grey line indicate instances where work has been saved at the expense of achieving the recall target.

In Fig. 6, we rescale the x axis, calling it additional burden, which is simply the work saved when the criterion is triggered minus the work saved when the recall target was actually achieved. This measure indicates whether the stopping criterion was triggered too early (negative values) or too late (positive values). The figure directly highlights the trade-offs involved in deciding when to stop screening: For our criteria, there is mostly a small additional burden which comes with the necessity to make sure the desired recall target has been reached and reject the null hypothesis that this has not been the case. For the other criteria, there are many cases in which additional burden is negative, i.e. the criterion has been triggered too early. In these cases, however, the desired recall is missed.

To help explain the different work savings that were observed in our experiments, we show the distribution of work savings from our ranked quasi-sampling criterion for each dataset in Fig. 7. In general, higher work savings are possible when the total number of documents is larger. However, in datasets with a low proportion of relevant documents, many documents need to be screened to achieve a high confidence that there are only few relevant documents remaining in the unseen ones. Therefore, smaller work savings are possible.




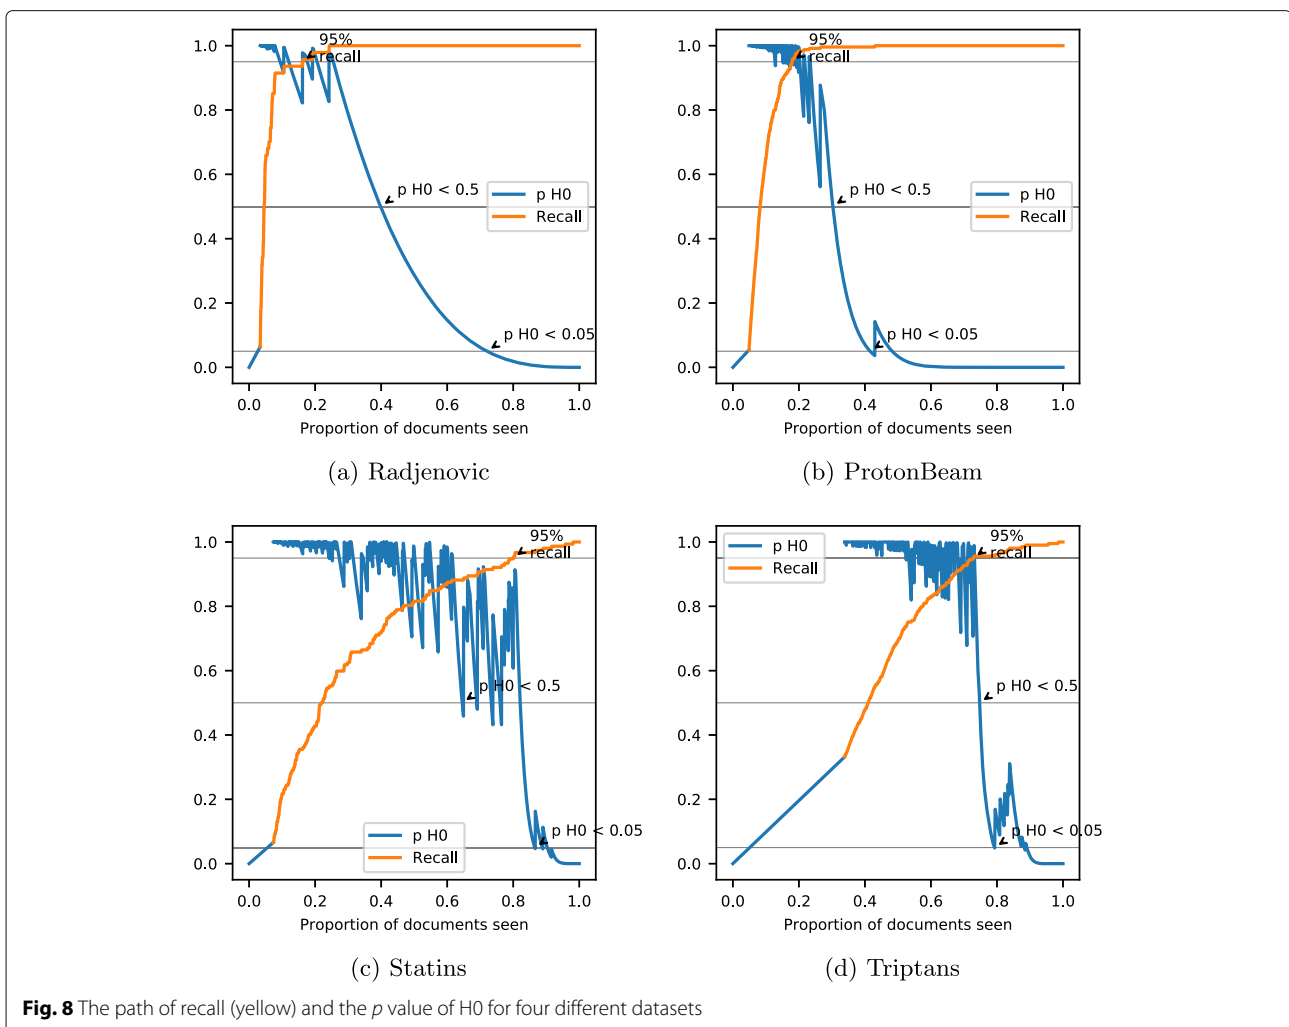
Figure 8 shows the recall and the p value for the null hypothesis for the iteration where the recall target is reached first for four datasets. Although the 95% recall target is achieved very quickly in the Radjenovic dataset, the null hypothesis cannot be excluded until much later. This is because the dataset has only 47 relevant documents out of a population of 5999. After the 95% recall target was achieved, 45 out of 47 relevant documents had been seen and 5029 documents remained. The null hypothesis was therefore that 3 or more of these 5029 documents were relevant, which requires a lot of evidence to disprove. The burden of proof was smaller in the case of the Proton Beam dataset: at the point that the 95% recall threshold was reached, the null hypothesis to disprove was that a minimum of 13 out of 3369 remaining documents were relevant.

The Statins and Triptans datasets show how the criterion performs when the machine learning model has performed poorly in predicting relevant results. In each

case, 95% recall is achieved with close to 20% of documents remaining. With fewer documents remaining, it takes fewer screening decisions to rule out the possibility that the number of relevant documents left is incompatible with the achievement of the recall target.

Discussion

Our results show that it is possible to use machine learning to achieve a given level of recall with a given level of confidence. The trade-off for achieving recall reliably is that the work saving achieved is less than the maximum possible work saving. However, for large datasets with a significant proportion of relevant documents, the additional effort required to satisfy the criterion will be small compared to the work saved by using machine learning. This makes the approach well suited to broad topics with lots of literature. In other words, it is precisely where machine learning will be most useful that the additional effort will be small.



Different use cases for machine learning enhanced screening may also carry different requirements for recall, or different tolerances for uncertainty. These can be flexibly accommodated within our stopping criterion. Importantly, the ability to make statements about the authors' confidence in achieving a given recall target makes it possible to clearly communicate the implications of using machine learning enhanced screening to readers and reviewers who are not machine learning specialists. This is extremely important in live systematic reviews.

Our criteria have the further advantage that they are independent of the choice or performance of the machine learning model. If a model performs badly at discerning relevant from irrelevant results, the only consequence will be that the work saved will be low. With other criteria, this may result in poor recall. When using machine learning for screening, poor recall can result in biased results, while low work savings represent no loss to the reviewer as compared to not using machine learning.

One caveat in the derivation of our criteria is that we did not address the problem of multiple testing formally. Such a derivation is mathematically challenging and beyond the scope of this paper. However, the performance of the criteria shows that this is of limited practical concern. Formally describing screening procedures with iterative testing should be a next step towards even more rigorous stopping criteria and should be fully worked out in future research.

So far, systematic review standards have no way of accommodating screening with machine learning. We hope that the reliability and clarity of reporting offered by our stopping criteria make them suitable for incorporation into standards, so that machine learning for systematic review screening can fulfil its promise of reducing workload and making more ambitious reviews tractable.

Conclusion

This paper demonstrates the drawbacks of existing stopping criteria for machine learning approaches to document screening, particularly with regard to reliability. We propose a simple method that delivers reliable recall, independent of machine learning approach or model performance. Our statistical stopping criteria allow users to easily communicate the implications of their use of machine learning, making machine learning enhanced screening ready for live reviews.

Acknowledgements

Max Callaghan is supported by a PhD scholarship from the Heinrich Böll Foundation. Finn Müller-Hansen acknowledges funding from the German Federal Ministry of Research and Education within the Strategic Scenario Analysis (START) project (grant reference: 03EK3046B).

Authors' contributions

MC designed the research and conducted the experiments. FMH contributed to the development of the statistical basis for the stopping criterion. Both authors wrote and edited the manuscript. The authors read and approved the final manuscript.

Availability of data and materials

All computational steps required to reproduce this analysis are documented online at <https://github.com/mcallaghan/rapid-screening>.

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Mercator Research Institute on Global Commons and Climate Change, EUREF Campus 19, Torgauer Straße 12-15, 10829 Berlin, Germany. ²Priestley International Centre for Climate, University of Leeds, Leeds, LS2 9JT, UK. ³Potsdam Institute for Climate Impact Research (PIK), Member of the Leibniz Association, P.O. Box 60 12 03, 14412, Potsdam, Germany.

Received: 2 December 2019 Accepted: 5 November 2020

Published online: 28 November 2020

References

- Westgate M, Haddaway N, Cheng S, McIntosh E, Marshall C, Lindenmayer D. Software support for environmental evidence synthesis. *Nat Ecol Evol.* 2018;2:588–90. <https://doi.org/10.1038/s41559-018-0502-x>.
- O'Mara-Eves A, Thomas J, McNaught J, Miwa M, Ananiadou S. Using text mining for study identification in systematic reviews: a systematic review of current approaches. *Syst Rev.* 2015;4(1):1–22. <https://doi.org/10.1186/2046-4053-4-5>.
- Miwa M, Thomas J, O'Mara-Eves A, Ananiadou S. Reducing systematic review workload through certainty-based screening. *J Biomed Inform.* 2014;51:242–53. <https://doi.org/10.1016/j.jbi.2014.06.005>.
- Wallace B, Trikalinos T, Lau J, Brodley C, Schmid C. Semi-automated screening of biomedical citations for systematic reviews. *BMC Bioinforma.* 2010;11: <https://doi.org/10.1186/1471-2105-11-55>.
- Wallace BC, Small K, Brodley CE, Trikalinos TA. Active learning for biomedical citation screening. In: Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '10). New York: Association for Computing Machinery; 2010. p. 173–82. <https://doi.org/10.1145/1835804.1835829>.
- Jonnalagadda S, Petitti D. A new iterative method to reduce workload in systematic review process. *Int J Comput Biol Drug Des.* 2013;6((1/2)):5. <https://doi.org/10.1504/ijcbdd.2013.052198>.
- Przybyła P, Brockmeier A, Kontonatsios G, Le Pogam M, McNaught J, von Elm E, Nolan K, Ananiadou S. Prioritising references for systematic reviews with RobotAnalyst: a user study. *Res Synth Methods.* 2018;9(3): 470–88. <https://doi.org/10.1002/jrsm.1311>.
- Settles B. Active learning literature survey technical report. University of Wisconsin-Madison. 2009. <https://doi.org/10.1016/j.matlet.2010.11.072>.
- Lefebvre C, Glanville J, Briscoe S, Littlewood A, Marshall C, Metzendorf M-I, Noel-Storr A, Rader T, Shokraneh F, Thomas J, Wieland LS. Chapter 4: Searching for and selecting studies. In: Higgins JPT, Thomas J, Chandler J, Cumpston M, Li T, Page MJ, Welch VA, editors. *Cochrane Handbook for Systematic Reviews of Interventions* version 6.1 (updated September 2020). Cochrane; 2020. www.training.cochrane.org/handbook.
- Bannach-Brown A, Przybyła P, Thomas J, Rice A, Ananiadou S, Liao J, Macleod M. Machine learning algorithms for systematic review: reducing workload in a preclinical review of animal studies and reducing human screening error. *Syst Rev.* 2019;8(1):1–12. <https://doi.org/10.1186/s13643-019-0942-7>.
- Marshall I, Wallace B. Toward systematic review automation: a practical guide to using machine learning tools in research synthesis. *Syst Rev.* 2019;8(1):1–10. <https://doi.org/10.1186/s13643-019-1074-9>.
- Shemilt I, Siron A, Hollands G, Marteau T, Ogilvie D, O'Mara-Eves A, Kelly M, Thomas J. Pinpointing needles in giant haystacks: use of text mining to reduce impractical screening workload in extremely large scoping reviews. *Res Synt Methods.* 2014;5(1):31–49. <https://doi.org/10.1002/jrsm.1093>.
- Yu Z, Menzies T. FAST 2: an intelligent assistant for finding relevant papers. *Expert Syst Appl.* 2019;120:57–71. <https://doi.org/10.1016/j.eswa.2018.11.021>.

14. Di Nunzio GM. A Study of an Automatic Stopping Strategy for Technologically Assisted Medical Reviews. In: Pasi G, Piwowarski B, Azzopardi L, Hanbury A, editors. *Advances in Information Retrieval. ECIR 2018. Lecture Notes in Computer Science*, vol 10772. Cham: Springer; 2018. https://doi.org/10.1007/978-3-319-76941-7_61.
15. Howard BE, Phillips J, Tandon A, Maharana A, Elmore R, Mav D, Sedykh A, Thayer K, Merrick BA, Walker V, Rooney A, Shah RR. SWIFT-Active Screener: Accelerated document screening through active learning and integrated recall estimation. *Environ Int*. 2020;138:105623. <https://doi.org/10.1016/j.envint.2020.105623>.
16. Cohen AM, Hersh WR, Peterson K, Yen P-Y. Reducing workload in systematic review preparation using automated citation classification. *J Am Med Inform Assoc*. 2006;13(2):206–19. <https://doi.org/10.1197/jamia.M1929>.
17. Terasawa T, Dvorak T, Ip S, Raman G, Lau J, Trikalinos TA. Systematic review: charged-particle radiation therapy for cancer. *Ann Intern Med*. 2009;151(8):556–65. <https://doi.org/10.7326/0003-4819-151-8-200910200-00145>.
18. Castaldi P, Cho M, Cohn M, Langerman F, Moran S, Tarragona N, Moukhachen H, Venugopal R, Hasimja D, Kao E, Wallace B, Hersh C, Bagade S, Bertram L, Silverman E, Trikalinos T. The COPD genetic association compendium: a comprehensive online database of COPD genetic associations. *Hum Mol Genet*. 2009;19(3):526–34. <http://dx.doi.org/10.1093/hmg/ddp519>.
19. Buonaccorsi J. A note on confidence intervals for proportions in finite populations. *Am Stat*. 1987;41(3):215–8. <http://dx.doi.org/10.2307/2685108>.
20. Sahai H, Khurshid A. A note on confidence intervals for the hypergeometric parameter in analyzing biomedical data. *Comput Biol Med*. 1995;25(1):35–8. [http://dx.doi.org/10.1016/0010-4825\(95\)98883-f](http://dx.doi.org/10.1016/0010-4825(95)98883-f).
21. Cortes C, Vapnik V. Support-vector networks. *Mach Learn*. 1995;20:273–97. <https://doi.org/10.1007/BF00994018>.
22. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay É. Scikit-learn: Machine Learning in Python. 2011.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions

