



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/150389/>

Version: Published Version

Proceedings Paper:

Sneyd, A. and Stevenson, R. (2019) Modelling stopping criteria for search results using Poisson processes. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP). 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, 03-07 Nov 2019, Hong Kong, China. Association for Computational Linguistics (ACL), pp. 3484-3489. ISBN: 9781950737901.

<https://doi.org/10.18653/v1/D19-1351>

© 2019 The Association for Computational Linguistics. Licensed on a Creative Commons Attribution 4.0 International License (<https://creativecommons.org/licenses/by/4.0/>).

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Modelling Stopping Criteria for Search Results using Poisson Processes

Alison Sneyd and Mark Stevenson

Natural Language Processing Group

Department of Computer Science

University of Sheffield

a.sneyd,mark.stevenson@sheffield.ac.uk

Abstract

Text retrieval systems often return large sets of documents, particularly when applied to large collections. Stopping criteria can reduce the number of these documents that need to be manually evaluated for relevance by predicting when a suitable level of recall has been achieved. In this work, a novel method for determining a stopping criterion is proposed that models the rate at which relevant documents occur using a Poisson process. This method allows a user to specify both a minimum desired level of recall to achieve and a desired probability of having achieved it. We evaluate our method on a public dataset and compare it with previous techniques for determining stopping criteria.

1 Introduction

Document retrieval systems are widely used in information retrieval (IR) to return a set of documents in response to a user query. They often return large numbers of documents, particularly for large collections, which places a significant burden on a user who has to manually review them to determine which are relevant (i.e. of interest). The problem is particularly acute in scenarios where relevant documents are rare in the set of retrieved documents, or where there is a high cost associated with missing a relevant document.

Systematic reviews of medical literature are a good example of such a task. They aim to identify and summarise all available evidence for a specific research question such as ‘Is systemic inflammation present in stable chronic obstructive pulmonary disease?’ (Gan et al., 2004). Evidence included in systematic reviews is typically identified by executing a Boolean query optimised for high recall, with the results manually reviewed by domain experts. However, this process is extremely expensive and a single review may cost up to a

quarter of a million dollars and take over a year to complete (McGowan and Sampson, 2005; Karimi et al., 2010). A range of studies have demonstrated that ranking the results of the Boolean query can support the identification of relevant studies by placing them higher in the ranking, e.g. (Miwa et al., 2014; Wallace et al., 2010; O’Mara-Eves et al., 2015; Shemilt et al., 2016). However, ranking alone does not reduce the manual review workload since it does not inform reviewers about when they can stop looking at search results.

Search stopping criteria are methods for determining when manual examination of search results can stop. They essentially provide information about the risk of missing relevant documents that is associated with not examining some portion of the documents returned by a query. Stopping criteria have been proposed for technology assisted reviews (Cormack and Grossman, 2016) and IR test collection development (Losada et al., 2019).

Cormack and Grossman (2016) introduced the target and knee stopping methods. The former involves examining randomly chosen documents until a specified number of relevant documents have been found. The number of relevant documents required is determined using statistical inference and guarantees a minimum recall score with a specified probability. The knee method employs a knee detection algorithm (Satopaa et al., 2011) to estimate when the gain curve of recall versus effort has exceeded a certain tipping point, called the slope ratio. Cormack and Grossman (2016) found that this method was better able to predict the most efficient stopping point than the target method, but did not provide any probabilistic guarantees about the level of recall that would be achieved. Additionally, Losada et al. (2019) proposed simple baselines and used power curves to determine stopping criteria. A significant limi-

tation of previous approaches is that they are evaluated using a single ranking algorithm, but for a stopping approach to be generally useful it must be reliable for a range of ranking algorithms.

This paper proposes a novel method for generating search stopping criteria based on an inhomogeneous Poisson process. The method assumes the rate at which relevant documents occur in the ranking decreases exponentially and predicts how many relevant documents need to be found in order to achieve a defined level of recall. This approach, along with the knee and target methods, is evaluated using data from a shared task for which over 30 submissions are available. This allows approaches to be evaluated against a wide range of rankings of varying effectiveness. We found that, while each method returns a reliably high level of recall, the proposed method produces the most accurate predictions while also being robust to ranking quality.¹

2 Method

Our proposed method models the rate at which relevant documents occur in a ranked set of retrieved documents, using an exponential function learned from a sample of the data. It is then possible to predict the total number of relevant documents to a desired probability using a Poisson process, and from this, to predict the rank at which a desired level of recall is achieved.

2.1 Poisson Process

Poisson processes model the occurrence of random events that take place at a given rate (Kingman, 1993). It is assumed that the events are independent and that the number of occurrences in an interval follows a Poisson distribution. An inhomogeneous (or non-homogeneous) Poisson process has a variable rate function, λ , where the rate at which events occur is a function of the space over which the process is defined. Let:

$$\Lambda(a, b) = \int_a^b \lambda(x) dx, \quad (1)$$

and let $N(a, b)$ denote the number of events occurring in the interval $(a, b]$. Then:

$$P(N(a, b) = r) = \frac{[\Lambda(a, b)]^r}{r!} e^{-\Lambda(a, b)}. \quad (2)$$

¹Code available at: https://github.com/alisonsneyd/poisson_stopping_method

In particular, N is a Poisson random variable with expected value $\Lambda(a, b)$.

We model the occurrence of relevant documents in a set of ranked, retrieved documents as an inhomogeneous Poisson process, where λ is an exponential function (i.e. $\lambda(x) = de^{kx}$). So, from Equation 1:

$$\begin{aligned} \Lambda(a, b) &= \int_a^b de^{kx} dx \\ &= \frac{d}{k} (e^{kb} - e^{ka}). \end{aligned}$$

We are interested in knowing the total number of relevant documents and so restrict our attention to the interval $(0, n]$, where n is the number of documents in the set, yielding:

$$\Lambda(0, n) = \frac{d}{k} (e^{kn} - 1). \quad (3)$$

By Equations 2 and 3, the probability of there being r relevant documents in the document set is:

$$P(N(0, n) = r) = \frac{[\frac{d}{k}(e^{kn} - 1)]^r}{r!} e^{-\frac{d}{k}(e^{kn} - 1)}.$$

Since $N(0, n)$ is a Poisson random variable, the upper limit of the number of relevant documents in the collection (R), subject to a given probability, can be estimated using the cumulative distribution function of $N(0, n)$. The number of relevant documents required to achieve a desired level of recall (T) is then RT . The proposed stopping criteria is to cease relevance judgments when $\lceil RT \rceil$ relevant documents have been found. This method is flexible and can be adapted to whatever level of recall a user desires, and to the level of certainty they want associated with it.

2.2 Rate Estimation

The approach proceeds by examining the first α documents in the ranking, and then predicting the total number of relevant documents in the entire ranking. If the number of relevant documents observed in the first α documents ($rel(\alpha)$) is greater than or equal to the number of relevant documents needed to achieve the desired level of recall (i.e. $rel(\alpha) \geq \lceil RT \rceil$), the process stops and no more documents are examined. Otherwise, the next β documents are added to the pool of examined documents and the process repeated.

Accurate estimation of the rate at which relevant documents occur in the ranking is an impor-

tant component of the proposed approach. The parameters of the rate function are estimated as follows: Split the documents which have been examined so far into m sub-intervals and define a set of points (x_i, y_i) , $i = 1, \dots, m$, where x_i is the midpoint of the i^{th} interval and y_i is the number of relevant documents it contains. Fit a curve of the form de^{kx} , $d, k \in \mathbb{R}$, to the points.

This approach assumes that relevant documents occur frequently enough for their rate to be predicted and that this rate decreases approximately exponentially, thereby implying the shape of the gain curve (number of relevant documents versus rank) will be roughly convex. While these assumptions frequently hold, they are not guaranteed to do so and they are therefore tested in multiple ways.

Firstly, if fewer than a given threshold, γ , of relevant documents are found in the initial sample of α documents, it is assumed that either relevant documents are too rare for accurate estimates or that the ranking places them later, meaning their rate is not exponential. No stopping point is predicted and all documents in the ranking are examined (while this is costly in terms of document judgment, it preserves recall and is the default approach in systematic reviews).

Secondly, the Poisson process is dependent on $\lambda(x)$ reliably estimating the rate at which relevant documents occur. Therefore when a curve is fitted, its accuracy is tested as follows: Suppose the sample ends at rank n . The number of relevant documents estimated by $\lambda(x)$ at rank n , given by $\sum_{i=1}^n \lambda(i)$, is compared to the actual number of relevant documents in the sample ($rel(n)$). If the number of relevant documents found is less than a scalar multiple (δ) of the predicted number (i.e. if $rel(n) < \delta \times \sum_{i=1}^n \lambda(i)$) then the curve is assumed to be a bad predictor. No stopping point is predicted and the next β documents in the ranking are examined.

3 Experiments

3.1 Data

The dataset used for the experiments originated from the CLEF 2017 e-Health Lab Task 2 “Technology Assisted Reviews in Empirical Medicine” (Kanoulas et al., 2017). The task required participants to identify relevant studies for a set of 30 Diagnostic Test Accuracy systematic reviews

(known as “topics”) from the Cochrane Library.² Participants were provided with the (unordered) set of documents retrieved by a complex Boolean query and asked to rank them based on relevance to the review topic.

The number of documents retrieved by the Boolean queries for the 30 topics ranged from 64 to 12,807, with median 2,070, while the total number of documents across all topics was 117,562. The number of relevant documents ranged from 2 to 460, with median 38. On average, only 1.58% of documents were relevant across the 30 topics.³

A total of 33 runs ranking the entire set of articles were submitted by 11 groups of participants.⁴ The rankings produced by these runs were based on a wide range of techniques, including manual judgments of document relevance, and automated techniques based on (semi-)supervised classifiers and relevance feedback. These runs contain a diverse set of rankings of varying effectiveness, making them very suitable for testing the robustness of search stopping criteria.

3.2 Evaluation Metrics

The goal is to reduce the amount of manual judgment required, subject to maintaining a minimum level of recall. Following Cormack and Grossman (2016), the following metrics were used: For a stopping criteria S and topic T , the *effort* of S in T is defined as $|E_T|$, where E_T is the set of documents that were examined before stopping. For a desired level of recall R , the *acceptability* of S with respect to T is defined as:

$$acceptability(S_T) = \begin{cases} 1, & \text{if } recall(S_T) \geq R, \\ 0, & \text{else.} \end{cases}$$

As sets of runs of topics are being evaluated, additional metrics are used. For a stopping criteria S considered over a set of topics in a run \mathcal{R} , the *reliability* of S with respect to \mathcal{R} is:

$$\frac{|\{T \in \mathcal{R} : acceptability(S_T) = 1\}|}{|\mathcal{R}|}$$

S is considered *reliable* over a set of runs if its mean reliability over the runs is not less than a desired threshold p , which is equivalent to the proportion of all topics in all runs that are acceptable

²<https://www.cochranelibrary.com/>

³The CLEF data includes relevance judgments at two levels: abstract and content. The abstract level information is used for the experiments described in this paper.

⁴Available at <https://github.com/CLEF-TAR>

being greater than or equal to p . The *effort* of S with respect to \mathcal{R} is $\sum_{T \in \mathcal{R}} |E_T|$. The *percentage of effort saved* by S over R will denote the percentage of all documents in the set that do not have to be examined:

$$100 \times \sum_{T \in \mathcal{R}} \frac{|T| - |E_T|}{|T|}.$$

3.3 Approaches

The Poisson process approach (Section 2) and knee and target methods (Section 1) were implemented. For all methods, the minimum level of recall was set to 0.7 and the reliability threshold to 0.95, following Cormack and Grossman (2016).

For the Poisson process (**PP**), the initial sample size (α) and increases in batches (β) were set to $0.3|T|$ and $0.05|T|$ respectively. The value of α was chosen as it is large enough to predict the rate curve with reasonable reliability but without requiring too many documents to be examined. The γ and δ parameters were set to 20 and the 0.7 respectively.

For the target method (**TM**), we follow Cormack and Grossman (2016) in requiring that a set of 10 relevant documents be identified before the stopping point is predicted since they demonstrated that this produced a minimum recall of 0.7 with 0.95 probability.

The knee method requires a slope ratio to be set given by $\epsilon + 6 - \min(\text{relret}, \epsilon)$, where relret is the number of relevant documents retrieved for a given rank and ϵ is a variable that prevents the method stopping too soon. A version of the approach (**KM-default**) was implemented using $\epsilon = 150$, as recommended by Cormack and Grossman (2016). The value of ϵ was also tuned using 3-fold cross-validation with $\epsilon \in \{0, 25, 50, 100, 150, 200\}$. The smallest reliable value for ϵ was found to be 50 in each fold and consequently a second version (**KM-tuned**) was implemented using this value.

For comparison, an Oracle (**OR**) method was also implemented which stops at the point in the ranking where 70% recall has been achieved.

4 Results

Each method was found to be reliable (i.e. achieved the desired recall of 0.7 for at least 95% of topics). The top part of Table 1 shows the results over all topics in all 33 runs. The proposed

approach, PP, required less effort than other methods (TM, KM-tuned and KM-default). PP allows the desired recall to be achieved without examining 42.1% of the documents, thereby allowing 4.9% more of the documents to go unexamined than the next most efficient approach (KM-tuned).

Further analysis was carried out to explore the effect of the quality of the initial ranking on the performance of the various approaches. The 33 runs were ranked based on their Area Under Recall Curve (AURC) scores in the official evaluation (Kanoulas et al., 2017). (The AURC metric computes the area under the cumulative recall curve normalized by the optimal area.) The top, middle and bottom five runs in the rankings were selected. The top five had AURC scores in the range 0.92 - 0.93, indicating effective rankings in which most of the relevant documents appear early. The bottom five runs had scores in the range 0.48 - 0.60, indicating the rankings are not much different from random orderings of the documents.

PP was the most efficient in terms of effort for the middle and bottom five runs, and the second most efficient for the top five runs. KM-tuned required the least effort for the top five runs, but the second most for the middle five runs and the third most for the bottom five runs. As might be expected, each method required the least effort for the top five runs and the most effort for the bottom five.

A graph of the OR and PP effort scores for each run is shown in Figure 1.




Figure 1: Oracle and Poisson process effort versus AURC score.

The difference in performance between the tuned and untuned version of the knee method (KM-tuned and KM-default) demonstrates the sensitivity of this method to appropriate setting of parameter values. These results demonstrate that

		Mean %	
	Mean Eff.	Eff. Saved	
All 33 runs			
TM	79,111		32.7%
KM-default	102,681		12.7%
KM-tuned	73,826		37.2%
PP	68,122		42.1%
OR	33,760		71.3%
Top 5 runs			
TM	56,762		51.7%
KM-default	87,886		25.2%
KM-tuned	33,841		71.2%
PP	54,792		53.4%
OR	7,202		93.9%
Middle 5 runs			
TM	76,704		34.8%
KM-default	107,703		8.4%
KM-tuned	75,888		35.4%
PP	60,311		48.7%
OR	27,565		76.6%
Bottom 5 runs			
TM	104,620		11.0%
KM-default	111,055		5.5%
KM-tuned	110,778		5.8%
PP	100,622		14.4%
OR	74,389		36.7%

Table 1: Mean effort and mean percentage of effort saved for the 33 CLEF 2017 runs. Results are shown for all runs plus the top, middle and bottom five runs. Highlighted figures are the best non-Oracle results.

the approach performs strongly for effective rankings, such as the one it was developed for (Cormack and Grossman, 2015), but is not robust to less effective ones. The PP approach is less sensitive to ranking quality (which is unlikely to be known a priori).

Two examples of the actual gain curves of a ranking and the gain curve estimated by the rate function $\lambda(x)$ learned by the PP approach are shown in Figure 2. In the top graph, the rate at which relevant documents occur decreases expo-

entially and the high similarity between the two curves shows λ is fit well. In the bottom graph, the rate at which relevant documents occur does not decrease exponentially. However, PP still returns the desired level of recall because the probability of the Poisson process was set to a high value (0.95).




Figure 2: The gain curves of two rankings from CLEF 2017 runs, compared to the gain curve estimated by λ .

5 Conclusion

This work proposed a novel method for determining a stopping point in search results using a Poisson process. The method was applied to 33 varied CLEF 2017 runs and proved reliable in terms of recall, while also requiring substantially less effort and demonstrating less sensitivity to ranking quality than existing methods. On average, the Poisson process saved manual judgment of over 40% of the documents associated with a run. Although our method has been explored in the context of systematic reviews, it could be applied to a variety of other document retrieval problems.

References

- Gordon V. Cormack and Maura R. Grossman. 2015. [Autonomy and reliability of continuous active](#)

- learning for technology-assisted review. *CoRR*, abs/1504.06868.
- Gordon V. Cormack and Maura R. Grossman. 2016. [Engineering quality and reliability in technology-assisted review](#). In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16*, pages 75–84, New York, NY, USA. ACM.
- W Qi Gan, SFP Man, A Senthilselvan, and DD Sin. 2004. Association between chronic obstructive pulmonary disease and systemic inflammation: a systematic review and a meta-analysis. *Thorax*, 59(7):574–580.
- Evangelos Kanoulas, Dan Li, Leif Azzopardi, and Rene Spijker. 2017. CLEF technologically assisted reviews in empirical medicine overview. In *Working Notes of CLEF 2017 - Conference and Labs of the Evaluation forum*, CEUR Workshop Proceedings, Dublin, Ireland. CEUR-WS.org.
- Sarvnaz Karimi, Stefan Pohl, Falk Scholer, Lawrence Cavedon, and Justin Zobel. 2010. Boolean versus ranked querying for biomedical systematic reviews. In *BMC Med. Inf. & Decision Making*.
- J F C Kingman. 1993. *Poisson Processes*. Oxford studies in probability; 3. Oxford : Clarendon Press.
- David E. Losada, Javier Parapar, and Alvaro Barreiro. 2019. [When to stop making relevance judgments? a study of stopping methods for building information retrieval test collections](#). *Journal of the Association for Information Science and Technology*, 70(1):49–60.
- Jessie L McGowan and Margaret Sampson. 2005. Systematic reviews need systematic searchers. *Journal of the Medical Library Association : JMLA*, 93 1:74–80.
- Makoto Miwa, James Thomas, Alison O’Mara-Eves, and Sophia Ananiadou. 2014. Reducing systematic review workload through certainty-based screening. *Journal of Biomedical Informatics*, 51:242–253.
- Alison O’Mara-Eves, James Thomas, John McNaught, Makoto Miwa, and Sophia Ananiadou. 2015. Using text mining for study identification in systematic reviews: a systematic review of current approaches. *Systematic reviews*, 4(1):5.
- Ville Satopaa, Jeannie Albrecht, David Irwin, and Barath Raghavan. 2011. [Finding a “kneedle” in a haystack: Detecting knee points in system behavior](#). In *Proceedings of the 2011 31st International Conference on Distributed Computing Systems Workshops, ICDCSW '11*, pages 166–171, Washington, DC, USA. IEEE Computer Society.
- Ian Shemilt, Nada Khan, Sophie Park, and James Thomas. 2016. [Use of cost-effectiveness analysis to compare the efficiency of study identification methods in systematic reviews](#). *Systematic Reviews*, 5.
- Byron C Wallace, Thomas A Trikalinos, Joseph Lau, Carla E Brodley, and Christopher H Schmid. 2010. Semi-automated screening of biomedical citations for systematic reviews. *BMC Bioinformatics*, 11(1):55.