

Evaluation Measures for Relevance and Credibility in Ranked Lists

Christina Lioma, Jakob Grue Simonsen

Department of Computer Science
University of Copenhagen
Copenhagen, Denmark
{c.lioma,simonsen}@di.ku.dk

Birger Larsen

Department of Communication
University of Aalborg in Copenhagen
Copenhagen, Denmark
birger@hum.aau.dk

ABSTRACT

Recent discussions on *alternative facts*, *fake news*, and *post truth* politics have motivated research on creating technologies that allow people not only to *access* information, but also to *assess* the credibility of the information presented to them by information retrieval systems. Whereas technology is in place for filtering information according to relevance and/or credibility [15], no single measure currently exists for evaluating the accuracy or precision (and more generally *effectiveness*) of *both* the relevance *and* the credibility of retrieved results. One obvious way of doing so is to measure relevance and credibility effectiveness separately, and then consolidate the two measures into one. There at least two problems with such an approach: (I) it is not certain that the same criteria are applied to the evaluation of both relevance and credibility (and applying different criteria introduces bias to the evaluation); (II) many more and richer measures exist for assessing relevance effectiveness than for assessing credibility effectiveness (hence risking further bias).

Motivated by the above, we present two novel types of evaluation measures that are designed to measure the effectiveness of both relevance and credibility in ranked lists of retrieval results. Experimental evaluation on a small human-annotated dataset (that we make freely available to the research community) shows that our measures are expressive and intuitive in their interpretation.

KEYWORDS

relevance; credibility; evaluation measures

ACM Reference format:

Christina Lioma, Jakob Grue Simonsen and Birger Larsen. 2017. Evaluation Measures for Relevance and Credibility in Ranked Lists. In *Proceedings of ICTIR '17, Amsterdam, Netherlands, October 1–4, 2017*, 8 pages. DOI: <http://dx.doi.org/10.1145/3121050.3121072>

1 INTRODUCTION

Recent discussions on *alternative facts*, *fake news*, and *post truth* politics have motivated research on creating technologies that allow people, not only to *access* information, but also to *assess* the credibility of the information presented to them [8, 14]. In the broader area of information retrieval (IR), various methods for approximating [10, 23] or visualising [11, 17, 18, 21] information credibility have

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ICTIR '17, Amsterdam, Netherlands

© 2017 Copyright held by the owner/author(s). ISBN 978-1-4503-4490-6/17/10...\$15.00
DOI: <http://dx.doi.org/10.1145/3121050.3121072>

been presented, both stand-alone and in relation to relevance [15]. Collectively, these approaches can be seen as steps in the direction of building IR systems that retrieve information that is both relevant and credible. Given such a list of IR results, which are ranked decreasingly by both relevance and credibility, the question arises: how can we evaluate the quality of this ranked list?

One could measure retrieval effectiveness first, using any suitable existing relevance measure, such as NDCG or AP, and then measure separately credibility accuracy similarly, e.g. using the F-1 or the G-measure. This approach would output scores by two separate metrics¹, which would need to somehow be consolidated or considered together when optimising system performance. In such a case, and depending on the choice of relevance and credibility measures, it would not be always certain that the same criteria are applied to the evaluation of both relevance and credibility. For instance, whereas the state of the art metrics in relevance evaluation treat relevance as graded and consider it in relation to the rank position of the retrieved documents (we discuss these in Section 2), no metrics exist that consider graded credibility accuracy in relation to rank position. Hence, using two separate metrics for relevance and credibility may, in practice, bias the overall evaluation process in favour of relevance, for which more thorough evaluation metrics exist.

To provide a more principled approach that obviates this bias, we present two new types of evaluation measures that are designed to measure the effectiveness of both relevance and credibility in ranked lists of retrieval results *simultaneously* and *without bias* in favour of either relevance or credibility. Our measures take as input a ranked list of documents, and assume that assessments (or their approximations) exist both for the relevance and for the credibility of each document. Given this information, our Type I measures define different ways of measuring the effectiveness of both relevance and credibility based on differences in the *rank position* of the retrieved documents with respect to their ideal rank position (when ranked only by relevance or credibility). Unlike Type I, our Type II measures operate directly on *document scores* of relevance and credibility, instead of rank positions. We evaluate our measures both axiomatically (in terms of their properties) and empirically on a small human-annotated dataset that we build specifically for the purposes of this work. We find that our measures are expressive and intuitive in their interpretation.

¹In this paper, we use *metric* and *measure* interchangeably, as is common in the IR community, even though the terms are not synonymous. Strictly speaking, *measure* should be used for more concrete or objective attributes, and *metric* should be used for more abstract, higher-level, or somewhat subjective attributes [2]. When discussing effectiveness, which is generally hard to define objectively, but for which we have some consistent feel, Black et al. argue that the term *metric* should be used [2].

2 RELATED WORK

The aim of evaluation is to measure how well some method achieves its intended purpose. This allows to discover weaknesses in the given method, potentially leading to the development of improved approaches and generally more informed deployment decisions. For this reason, evaluation has been a strong driving force in IR, where, for instance, the literature of IR evaluation measures is rich and voluminous, spanning several decades. Generally speaking, relevance metrics for IR can be split into three high-level categories:

- (i) earlier metrics, assuming binary relevance assessments;
- (ii) later metrics, considering graded relevance assessments, and
- (iii) more recent metrics, approximating relevance assessments from user clicks.

We overview some among the main developments in each of these categories next.

2.1 Binary relevance measures

Binary relevance metrics are numerous and widely used. Examples include:

Precision @ k ($P@k$): the proportion of retrieved documents that are relevant, up to and including position k in the ranking;

Average Precision (AP): the average of (un-interpolated) precision values (proportion of retrieved documents that are relevant) at all ranks where relevant documents are found;

Binary Preference (bPref): this is identical to AP except that bPref ignores non-assessed documents (whereas AP treats non-assessed documents as non-relevant). Because of this, bPref does not violate the *completeness assumption*, according to which “all relevant documents within a test collection have been identified and are present in the collection”) [5];

Mean Reciprocal Rank (MRR): the reciprocal of the position in the ranking of the first relevant document only;

Recall: the proportion of relevant documents that are retrieved;

F-score: the equally weighted harmonic mean of precision and recall.

2.2 Graded relevance measures

There exist noticeably fewer graded relevance metrics than binary ones. The two main graded relevance metrics are NDCG and ERR:

Normalised Discounted Cumulative Gain (NDCG): the cumulative gain a user obtains by examining the retrieval result up to a rank position, where the relevance scores of the retrieved documents are:

- *accumulated* over all the rank positions that are considered,
- *discounted* in order to devalue late-retrieved documents, and
- *normalised* in relation to the maximum score that this metric can possibly yield on an ideal reranking of the same documents.

Two useful properties of NDCG are that it rewards retrieved documents according to both (i) their degree (or grade) of relevance, and (ii) their rank position. Put simply, this means that the more relevant a document is *and* the closer to the top it is ranked, the higher the NDCG score will be [12].

Expected Reciprocal Rank (ERR): ERR operates on the same high-level idea as NDCG but differs from it in that it penalises documents that are shown below very relevant documents. That is, whereas NDCG makes the *independence assumption* that “a document in a given position has always the same gain and discount independently of the documents shown above it”, ERR does not make this assumption, and, instead, considers (implicitly) the immediate context of each document in the ranking. In addition, instead of the discounting of NDCG, ERR approximates the expected reciprocal length of time that a user will take to find a relevant document. Thus, ERR can be seen as an extension of (the binary) MRR for graded relevance assessments [6].

2.3 User click measures

The most recent type of evaluation measures are designed to operate, not on traditionally-constructed relevance assessments (defined by human assessors), but on approximations of relevance assessments from user clicks (actual or simulated). Most of these metrics have underlying user models, which capture how users interact with retrieval results. In this case, the quality of the evaluation measure is a direct function of the quality of its underlying user model [24].

The main advances in this area include the following:

Expected Browsing Utility (EBU): an evaluation measure whose underlying user click model has been tuned by observations over many thousands of real search sessions [24];

Converting click models to evaluation measures: a general method for converting any click model into an evaluation metric [7]; and

Online evaluation: various different algorithms for interleaving [19] or multileaving [3, 4, 20] multiple initial ranked lists into a single combined ranking, and by approximating clicks (through user click models) on the resulting combined ranking, assigning credit (hence evaluating) the methods that produced each initial ranked list [9].

In addition to the above three types of IR evaluation measures, there also exists further literature on IR measures that consider additional dimensions on top of relevance, such as query difficulty for instance [16]. To the best of our knowledge, none of these measures consider credibility. The closest to a credibility measure we could find is the work by Balakrishnan et al. [1] on source selection for deep web databases: their method considers the agreement between different sources in answering a query as an indication of the credibility of the sources. An adjusted version of this agreement is modeled as a graph with vertices representing sources. Given such a graph, the credibility (or quality) of each source is calculated as the stationary visit probability of a random walk on this graph.

The evaluation measures we present in Sections 4 - 5 are the only ones, to our knowledge, that are designed to operate both on relevance and credibility. Beyond these two particular dimensions, reasoning more generally about different dimensions of effectiveness, the F-score, and its predecessor, van Rijsbergen’s E-score [22], are early examples of a single evaluation measure combining two different aspects, namely precision and recall. We return to this discussion in Section 5, where we present a variant of the F-score for aggregating relevance and credibility.

3 EVALUATION DESIDERATA

Given a ranked list of documents, the aim is to produce a measure that reflects how effective this ranking is with respect to *both* the relevance of these documents to some query *and* also the credibility of these documents (irrespective of a query).

There are at least two basic ways to produce such a metric:

Either

- (I) gauge the difference in rank position(s) between an input ranking and “ideal” relevance and credibility rankings,

or

- (II) employ relevance and credibility *scores* to gauge how well the input ranking reflects high versus low scores.

Note that while (II) is reminiscent of existing measures for relevance ranking, the fact that two distinct kinds of scores (relevance and credibility) – perhaps having different ranges and behaviour – must be combined may lead to further complications.

Accordingly, in the remainder of the paper, we call measures *Type I* if they are based primarily on differences in rank position, and *Type II* if they are based primarily on relevance and credibility scores.

Regardless of whether it is Type I or Type II, we reason that any measure must be easily interpretable. Hence, its scores should be normalised between 0 and 1, where low scores should indicate poor rankings, and high scores should indicate good rankings. The extreme points (0 and 1) of the scale should preferably be attainable by particularly bad or particularly good rankings; as a minimum, if the ranking can be measured against an “ideal” ranking (as in, e.g. NDCG), the value 1 should be attainable by the ideal ranking.

In addition to the above, there also exist desiderata for evaluation measures that are more debatable (e.g., how the measure should act in case of identical ranking scores for distinct documents). Below, we list what we believe to be the most pertinent desiderata. The list encompasses desiderata tailored to evaluate measures that gauge ranking based on either *rank position* or on (relevance or credibility) *scores*. For the desiderata pertaining to rank position, we need the following ancillary definition:

Let D_i be a document at rank i . We then define an *error* as any instance where

- either (a) the relevance of a document at rank i is greater than the relevance of a document at rank $i - 1$,
- or (b) the credibility of a document at rank i is greater than the credibility of a document at rank $i - 1$.

This assumes that documents are ranked decreasingly by relevance and credibility, i.e. that the “best” document occurs at the lowest (i.e. first) rank.

We define the following eight desiderata (referred to as **D1-D8** henceforth):

- D1** Larger errors should be penalised more than smaller errors;
- D2** Errors high in the ranking should be penalised more than errors low in ranking;
- D3** Let δ^r be the difference in relevance score between D_i and D_{i-1} when D_{i-1} is more relevant than D_i . Similarly, let δ^c be the difference in credibility score between D_i and D_{i-1} when D_{i-1} is more credible than D_i . Then, larger δ^r and δ^c values should imply larger error;
- D4** *Ceteris paribus*, a credibility error on documents of high relevance should be penalised more than a credibility error on documents of low relevance;
- D5** The metric should be well-defined even if all documents have identical ranking/credibility scores;
- D6** Scaling the document scores used to produce the ranking by some constant should not affect the metric;
- D7** If all documents have the same relevance score, the metric should function as a credibility metric; and vice versa;
- D8** We should be able to adjust (by some easily interpretable parameter) how much we wish to penalise low credibility with respect to low relevance, if at all.

Next, we present two types of evaluation measures of relevance and credibility that satisfy (wholly or partially) the above desiderata: Type I measures (Section 4) operate solely on the *rank positions* of documents; Type II measures (Section 5) operate solely on *document scores*.

4 TYPE I: RANK POSITION MEASURES

Given a ranking of documents that we want to evaluate (let us call this *input ranking*), we reason in terms of two additional *ideal rankings*: one by relevance only, and one by credibility only (the two ideal rankings are entirely independent of each other). So, for each document, we have:

- (1) its rank position in the input ranking;
- (2) its rank position in the ideal relevance ranking; and
- (3) its rank position in the ideal credibility ranking.

The basic idea is then to take each adjacent pair of documents in the input ranking, check for errors in the input ranking compared to the ideal relevance and separately the ideal credibility ranking, and aggregate those errors. We explain next how we do this.

Let D_i be the rank position of document D in the input ranking. We then denote by $R_{D_i}^r$ the rank position of D_i in the ideal relevance ranking, and by $R_{D_i}^c$ the rank position of D_i in the ideal credibility ranking. Note that subscript i refers to the rank position of D in the input ranking *at all times*. That is, $R_{D_i}^r$ should be read as: the position in the ideal relevance ranking of the document that is at position i in the input ranking; similarly for $R_{D_i}^c$.

Let the *monus operator* $\dot{-}$ be defined on non-negative real numbers by:

$$a \dot{-} b = \begin{cases} 0 & \text{if } a \leq b \\ a - b & \text{if } a > b \end{cases} \quad (1)$$

That is, $a \dot{-} b$ is simply subtraction as long as $a > b$ and otherwise just returns 0. Then, using the monus operator and the notation introduced above, we define a “relevance error” (ϵ^r) and a “credibility error” (ϵ^c) as:

$$\epsilon^r = R_{D_i}^r \dot{-} R_{D_{i+1}}^r \quad (2)$$

$$\epsilon^c = R_{D_i}^c \dot{-} R_{D_{i+1}}^c \quad (3)$$

In the above, i and $i + 1$ are the rank positions of two documents in the input ranking. Given two such documents, a “relevance error” occurs iff the document that is ranked lower (at rank i) in the input ranking is ranked after the other document in the ideal relevance ranking. Otherwise, the error is zero. Similarly for the “credibility error”.

For example, if three documents A, B and C are ranked as C, A, B in the input ranking (i.e., $D_1 = C, D_2 = A, D_3 = B$), but ranked as $R^r = [A, B, C]$ in an ideal relevance ranking, there are two relevance errors, namely

$$(i) R_{D_1}^r \dot{-} R_{D_2}^r = R_C^r \dot{-} R_A^r = 3 \dot{-} 1 = 2, \text{ and}$$

$$(ii) R_{D_2}^r \dot{-} R_{D_3}^r = R_A^r \dot{-} R_B^r = 1 \dot{-} 2 = 0.$$

We use the above “relevance error” and “credibility error” to define the two evaluation measures, presented next.

4.1 Normalised Local Rank Error (NLRE)

Let n be the total number of documents in the ranked list. We define the Local Rank Error (LRE) evaluation measure as $LRE = 0$ if $n = 1$, and otherwise:

$$LRE = \sum_{i=1}^{n-1} \frac{1}{\log_2(1+i)} ((\mu + \epsilon^r)(\nu + \epsilon^c) - \mu\nu) \quad (4)$$

where ϵ^r, ϵ^c are the relevance error and credibility error defined in Equations 2 – 3, and μ, ν are non-negative real numbers (with $\mu + \nu > 0$) controlling how much we wish to penalise low relevance with respect to low credibility. For instance, a high ν weighs credibility more, whereas a high μ weighs relevance more. The reason for the term $-\mu\nu$ inside the summation at the end is to ensure that the value of the LRE measure is zero if no error occurs.

Because Equation 4 is *large* for *bad* rankings and *small* for *good* rankings, we invert and normalise it (Normalised LRE or NLRE) as follows:

$$NLRE = 1 - \frac{LRE}{C_{LRE}} \quad (5)$$

where C_{LRE} is the normalisation constant, defined as:

$$C_{LRE} = \sum_{j=0}^{\lfloor \frac{n}{2} - 1 \rfloor} \frac{(n - 2j - 1)^2 + (\mu + \nu)(n - 2j - 1)}{1 + \log_2(1 + j)} \quad (6)$$

ensuring that $LRE/C_{LRE} \leq 1$. Note the “floor” function of the angular brackets above \sum in Equation 6, which rounds the contents of the brackets down to the next (lowest) integer.

The somewhat involved definition of C_{LRE} is due to the fact that we wish the maximal possible error attainable (i.e., rankings that produce the largest possibly credibility *and* relevance errors) to correspond to a value of 1 for LRE/C_{LRE} . Observe that $NLRE$ is 1 if no errors of any kind occur (because, in that case, LRE is 0).

Our NLRE measure satisfies the desiderata presented in Section 3 as follows:

- **D1** holds if we interpret error size as the size of the rank differences;
- **D2** holds due to the discount factor of $1/\log_2(1+i)$;
- **D3** is satisfied in the sense that larger differences in *credibility* or *relevance* ranks mean larger error;
- **D4**: The credibility error is scaled by the relevance error, if there is any (i.e., they are multiplied). If there is no relevance error, the credibility error is still strictly greater than zero;
- **D5**: The measure is well-defined in all cases;
- **D6**: No scores occur explicitly, only rankings, so scaling makes no difference;
- **D7** is satisfied because if all documents have equal relevance, the relevance error will be zero. The resulting score will measure only credibility error. And vice versa;
- **D8** is satisfied through μ and ν .

We call NLRE a *local* measure because it is affected by differences in credibility and relevance between documents at each rank position in the input ranking. We present next a *global* evaluation metric that does not take such “local” effects at each rank into account (i.e., any differences in credibility and relevance between documents at rank i in the input ranking do not affect the global metric; only the total difference of credibility and relevance of the entire input ranking affects the global metric).

4.2 Normalised Global Rank Error (NGRE)

We define the Global Rank Error (GRE) evaluation measure as $GRE = 0$ if $n = 1$, and otherwise:

$$GRE = \left(1 + \mu \sum_{i=1}^{n-1} \frac{1}{\log_2(1+i)} \epsilon^r\right) \left(1 + \nu \sum_{i=1}^{n-1} \frac{1}{\log_2(1+i)} \epsilon^c\right) - 1 \quad (7)$$

The notation is the same as for LRE. Similarly to LRE, we invert and normalise GRE, to produce its normalised version (NGRE) as follows:

$$NGRE = 1 - \frac{GRE}{C_{GRE}} \quad (8)$$

where C_{GRE} is the normalisation constant, defined as:

$$C_{GRE} = \mu\nu \left(\sum_{j=0}^{\lfloor \frac{n}{2} - 1 \rfloor} \frac{n - 2j - 1}{1 + \log_2(1 + j)} \right)^2 + (\mu + \nu) \sum_{j=0}^{\lfloor \frac{n}{2} - 1 \rfloor} \frac{n - 2j - 1}{1 + \log_2(1 + j)} \quad (9)$$

C_{GRE} is chosen to ensure that $GRE/C_{GRE} \leq 1$ and that $GRE/C_{GRE} = 1$ is possible iff the ranking has the maximal possible errors compared to both the ideal relevance and ideal credibility rankings. The square brackets above both \sum s in Equation 9 also use the floor function, exactly like in Equation 6.

As with NLRE, NGRE is 1 if no errors of any kind occur. In spite of the differences in computation, NGRE satisfies all eight desiderata for the same reasons given for NLRE.

The main intuitive difference between NLRE and NGRE is that in NGRE the credibility errors and relevance errors are cumulated separately, and then multiplied at the end. Thus, there is no immediate connection between credibility and relevance errors at the same rank (*locally*), hence we say that the metric is *global*.

The advantage of such a global versus local measure is that, in the global case, it is more straightforward to perform mathematical manipulations to achieve, e.g., normalisation, and easier to intuitively grasp what the measure means. The disadvantage is that local information is lost, and this may, in theory, lead to poorly performing measures. As the notion of “error” defined earlier is inherently a local phenomenon, the desiderata concerning errors are harder to satisfy formally for global measures.

5 TYPE II: DOCUMENT SCORE MEASURES

The two evaluation measures presented above (NLRE and NGRE) operate on the rank positions of documents. We now present three evaluation measures that operate, not on the rank positions of documents, but directly on document scores.

5.1 Normalised Weighted Cumulative Score (NWCS)

Given a ranking of documents that we wish to evaluate, let $Z^r(i)$ denote the relevance score with respect to some query of the document ranked at position i , and let $Z^c(i)$ denote the credibility score of the document ranked at position i . Then, we define the Weighted Cumulative Score (WCS) measure as:

$$WCS = \sum_{i=1}^n \frac{1}{\log_2(1+i)} (\lambda Z^r(i) + (1-\lambda)Z^c(i)) \quad (10)$$

where n is the total number of documents in the ranking list, and λ is a real number in $[0, 1]$ controlling the impact of relevance versus credibility in the computation. We normalise WCS by dividing it by the value obtained by an “ideal” ranking maximizing the value of WCS (this is inspired by the normalisation of the NDCG evaluation measure [12]):

$$NWCS = \frac{WCS}{IWCS} \quad (11)$$

where $IWCS$ is the *ideal* WCS, i.e. the maximum WCS that can be obtained on an ideal ranking of the same documents.

NWCS uses a simple weighted combination of relevance or credibility scores in the same manner as the metric *NGRE*, but is applicable directly to relevance or credibility scores (instead of ranking positions).

Our NWCS measure satisfies the following of the desiderata presented in Section 3:

- **D1** is satisfied as both Z^r and Z^c occur linearly in WCS;
- **D2** is satisfied due to the logarithmic discounting for increasing rank positions;
- **D3** is satisfied by design as both Z^r and Z^c occur directly in the formula for WCS;
- **D5** is satisfied as the measure is well-defined in all cases;
- **D6** is satisfied due to normalization;
- **D7** is satisfied because the contribution of the credibility scores (if all are equal) is just a constant in each term (and vice versa if relevance scores are all equal);
- **D8** is satisfied due to the presence of λ .

Of all desiderata, only **D4** is not satisfied: there is no scaling of credibility errors based on relevance. Despite this, the advantage of NWCS is that it is interpretable in much the same way as NDCG.

The main idea of the next two measures is that any two separate measures of *either* relevance *or* credibility, but not both, can be combined into a single aggregating measure of relevance and credibility. We next present two such aggregating measures.

5.2 Convex aggregating measure (CAM)

We define the convex aggregating measure (CAM) of relevance and credibility as:

$$CAM = \lambda M^r + (1-\lambda)M^c \quad (12)$$

where M^r and M^c denote respectively any valid relevance and credibility evaluation measure, and λ is a real number in $[0, 1]$ controlling the impact of the individual relevance or credibility measure in the overall computation. CAM is normalized if both M^r and M^c are normalised.

Our CAM measure satisfies the following desiderata:

- **D1** is satisfied for the same reasons as NWCS;
- **D2** is not satisfied in general;
- **D3** is satisfied for the same reasons as NWCS;
- **D4** is not satisfied in general;
- **D5** is satisfied for the same reasons as NWCS;
- **D6** is not satisfied in general; it is satisfied if both M^r and M^c are scale-free;
- **D7** is satisfied because the contribution of the credibility scores (if all are equal) is just a constant in each term (and vice versa if relevance scores are all equal);
- **D8** is satisfied for the same reasons as NWCS.

With respect to **D2**, **D4**, and **D6** not being satisfied in general: The tradeoff in this case is that as CAM is just a convex combination of existing measures, the scores are readily interpretable by anyone able to interpret M^r and M^c scores.

5.3 Weighted harmonic mean aggregating measure (WHAM) or “F-score for credibility and ranking”

We define the weighted harmonic mean aggregating measure (WHAM) as zero if either M^r or M^c is zero, and otherwise:

$$WHAM = \frac{1}{\lambda \frac{1}{M^r} + (1-\lambda) \frac{1}{M^c}} \quad (13)$$

where the notation is the same as for CAM in Equation 12 above. WHAM is the weighted harmonic mean of M^r and M^c . Observe that if $\lambda = 0.5$, WHAM is simply the F-1 scores of M^r and M^c . Note that WHAM is normalized if both M^r and M^c are normalised.

Similar definitions of metrics can be made that use other averages. For example, one can use the weighted arithmetic and geometric means instead of the harmonic mean.

Our WHAM measure satisfies the following desiderata:

- **D1** is satisfied for the same reasons as CAM;
- **D2** is not satisfied in general;
- **D3** is satisfied for the same reasons as CAM;
- **D4** is not satisfied in general;
- **D5** is satisfied for the same reasons as CAM;
- **D6** is not satisfied in general; it is satisfied if both M^r and M^c are scale-free;

Table 1: The 10 queries used in our experiments.

Query no.	Query
1	Smoking not bad for health
2	Princess Diana alive
3	Trump scientologist
4	UFO sightings
5	Loch Ness monster sightings
6	Vaccines bad for children
7	Time travel proof
8	Brexit illuminati
9	Climate change not dangerous
10	Digital tv surveillance

- **D7** is satisfied for the same reasons as CAM;
- **D8** is satisfied for the same reasons as CAM.

The primary advantage of CAM and WHAM is that their definitions appeal to simple concepts already known to larger audiences (convex combinations and averages), and hence the measures are simple to state and interpret. The consequent disadvantage is that this simplicity comes at the cost of not satisfying all desiderata.

We next present an empirical evaluation of all our measures.

6 EVALUATION

There are two main approaches for evaluating evaluation measures:

Axiomatic Define some general fundamental properties that a measure should adhere to, and then reflect on how many of these properties are satisfied by a new measure, and to what extent.

Empirical Present a drawback of existing standard and accepted measures, and illustrate how a new measure addresses this. Ideally, the new measure should generally correlate well with the existing measures, except for the problematic cases, where it should perform better [13].

We have already conducted the axiomatic evaluation of our measures, having presented 8 fundamental properties they should adhere to (Desiderata in Section 3), and having subsequently discussed each of our measures in relation to these fundamental properties in Sections 4 - 5. We now present the empirical evaluation. We first present our in-house dataset and experimental setup, and then our findings.

6.1 Empirical Evaluation

The goal is to determine how good our measures are at evaluating both relevance and credibility in ranked lists. We do this by comparing the scores of our measures to the scores of well-known relevance and separately credibility measures. This comparison is done on a small dataset that we create for the purposes of this work as follows². We formulated 10 queries that we thought were likely to fetch results of various levels of credibility if submitted to a web search engine. These queries are shown in Table 1. We then recruited 10 assessors (1 MSc student, 5 PhD students, 3 postdocs, and 1 assistant professor, all within Computer Science, but none

²Our dataset is freely available here: <https://github.com/diku-irlab/A66>

Table 2: Conversion of graded assessments to binary. The same conversion is applied to both relevance and credibility assessments.

Graded	Binary
1 (not at all)	0 (not at all)
2 (marginally)	0 (not at all)
3 (medium)	1 (completely)
4 (completely)	1 (completely)

working on this project; 1 female, 9 males). Assessors were asked to submit each query to Google, and to assign separately a score of relevance and a score of credibility to each of the top 5 results. Assessors were instructed to use the same graded scale of relevance and credibility shown in the first column of Table 2.

Assessors were asked to use their own understanding of relevance and credibility, and not to let relevance affect their assessment of credibility, or vice versa (relevance and credibility were to be treated as unrelated aspects). Assessors were instructed that, if they did not understand a query, or if they were unsure about the credibility of a result, they should open a separate browser and try to gather more information on the topic. Assessors received a nominal reward for their effort.

Even though assessors used the same queries, the top 5 results retrieved from Google per query were not always identical. Consequently, we compute our measures separately on each assessed ranking, and we report the arithmetic average. For NLRE and NGRE, we set $\mu = \nu = 0.5$, meaning that relevance and credibility are weighted equally. Similarly, for NWCS, CAM, and WHAM, we set $\lambda = 0.5$.

As no measures of *both* relevance and credibility exist, we compare the score of our measures on the above dataset to the scores of:

- NDCG (for graded relevance), AP (for binary relevance);
- F-1, G-measure (for binary credibility).

F-1 was introduced in Section 2 for relevance. We use it here to assess credibility, by defining its constituent precision and recall in terms of true/false positives/negatives (as is standard in classification evaluation). The G-measure is the geometric mean of precision and recall, which are defined as for F-1.

To render our graded assessments binary (for AP, F-1, G-measure), we use the conversion shown in Table 2.

6.2 Findings

Table 3 displays the scores of all evaluation measures on our dataset. We see that relevance-only measures (NDCG, AP) give overall higher scores than credibility-only measures (F-1, G). It is not surprising to see such high NDCG and AP scores, considering that we assess only the top 5 ranks of Google. What is however interesting, is the comparatively lower scores of credibility (F-1 and G). This practically means that even the top ranks of a high-traffic web search engine like Google can be occupied by information that is not entirely credible (at least for this specially selected set of queries).

Table 3: Our evaluation measures compared to NDCG, AP, F-1 and G. For NDCG we see our graded assessments. For the rest, we convert our graded assessments to binary as follows: 1 or 2 = not relevant/credible; 3 or 4 = relevant or credible. All measures are computed on the top 5 results returned for each query shown in Table 1. We report the average across all assessors.

RELEVANCE	
NDCG	0.9329
AP	0.7842
CREDIBILITY	
F-1	0.4786
G	0.5475
RELEVANCE and CREDIBILITY	
NLRE	0.8262
NGRE	0.6919
NWCS	0.9413
$CAM_{NDCG, F-1}$	0.7058
$CAM_{NDCG, G}$	0.7402
$CAM_{AP, F-1}$	0.6311
$CAM_{AP, G}$	0.6659
$WHAM_{NDCG, F-1}$	0.6326
$WHAM_{NDCG, G}$	0.6900
$WHAM_{AP, F-1}$	0.6089
$WHAM_{AP, G}$	0.6448

Looking at our measures of evaluation and credibility, we see that they range from roughly 0.6 to 0.9. This coincides with the range between the score of credibility-only measures and relevance-only measures. All of our measures are strongly and positively correlated to NDCG, AP, F-1, and G (from Spearman’s $\rho = 0.79$ for NDCG and F-1, up to $\rho = 0.97$ for NDCG and NLRE).

Table 4 shows examples of high divergence between the relevance and credibility of the retrieved documents, for three of our measures (the scores of our remaining metrics can be easily deduced from the respective relevance-only and credibility-only scores, as our omitted measures – CAM and WHAM – aggregate the existing relevance-only and credibility-only metrics shown in Table 4). Note that, whereas we found several examples of max relevance and min credibility in our data, there were (understandably) significantly fewer examples of max credibility and min relevance (this distribution is reflected in Table 4). We see that NWCS gives higher scores for queries 2 and 4-10 than NLRE and NGRE. For the first five examples (of max relevance and min credibility), this is likely because NWCS does not satisfy **D4**, namely that credibility errors should be penalised more on high relevance versus low relevance documents. We also see that NGRE gives consistently lower scores than NLRE and NWCS. This is due to its *global* aspect discussed earlier: NGRE accumulates credibility and relevance errors separately and then multiplies them at the end, meaning that local errors in each rank do not impact as much the final score (unlike NLRE and NWCS, which are both *local* in that sense, the first using document ranks, the second using document scores).

7 CONCLUSIONS

The credibility of search results is important in many retrieval tasks, and should be, we reason, integrated into IR evaluation measures that are, as of now, targeting mostly relevance. We have presented several measures and types of measures that can be used to gauge the effectiveness of a ranking, taking into account both credibility and relevance. The measures are both axiomatically and empirically sound, the latter illustrated in a small user study.

There are at least two natural extensions of our approach: First, the combination of rankings based on different criteria goes beyond the combination of relevance and credibility, and several such combinations are used in practice based on different criteria (e.g., combinations of relevance and upvotes on social media sites); we believe that much of our work can be encompassed in more general approaches, suitably axiomatised, that do not necessarily have to satisfy the same desiderata as those of this paper (e.g., do not have to scale credibility error by relevance errors as in our **D4**). Second, while we have chosen to devise measures that are both theoretically principled and conceptually simple using simple criteria (satisfaction of desiderata, local versus global, amenable to principled interpretation), there are many more measures that can be defined within the same limits. For example, our Type II measures are primarily built on simple combinations of scores or pre-existing measures that can easily be understood by the community, but at the price that some desiderata are hard or impossible to satisfy; however, there is no theoretical reason why one could not create Type II measures that incorporate some of the ideas from Type I metrics. We intend to investigate these two extensions in the future, and invite the community to do so as well.

Lastly, while the notion of credibility, in particular in news media, is subject to intense public discussion, very few empirical studies exist that contain user preferences, credibility rankings, or information needs related to credibility. The small study included in this paper, while informative, is a very small step in this direction. We believe that future substantial discussion of *practically relevant* research involving credibility in information retrieval would greatly benefit from having access to larger-scale empirical user studies.

REFERENCES

- [1] Raju Balakrishnan and Subbarao Kambhampati. 2011. SourceRank: relevance and trust assessment for deep web sources based on inter-source agreement. In *Proceedings of the 20th International Conference on World Wide Web, WWW 2011, Hyderabad, India, March 28 - April 1, 2011*, Sadagopan Srinivasan, Krithi Ramamritham, Arun Kumar, M. P. Ravindra, Elisa Bertino, and Ravi Kumar (Eds.). ACM, 227–236. DOI : <https://doi.org/10.1145/1963405.1963440>
- [2] Paul E. Black, Karen A. Scarfone, and Murugiah P. Souppaya (Eds.). 2008. *Cyber Security Metrics and Measures*. Wiley Handbook of Science and Technology for Homeland Security.
- [3] Brian Brost, Ingemar J. Cox, Yevgeny Seldin, and Christina Lioma. 2016. An Improved Multileaving Algorithm for Online Ranker Evaluation. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval, SIGIR 2016, Pisa, Italy, July 17-21, 2016*, Raffaele Perego, Fabrizio Sebastiani, Javed A. Aslam, Ian Ruthven, and Justin Zobel (Eds.). ACM, 745–748. DOI : <https://doi.org/10.1145/2911451.2914706>
- [4] Brian Brost, Yevgeny Seldin, Ingemar J. Cox, and Christina Lioma. 2016. Multi-Dueling Bandits and Their Application to Online Ranker Evaluation. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management, CIKM 2016, Indianapolis, IN, USA, October 24-28, 2016*, Snehasis Mukhopadhyay, ChengXiang Zhai, Elisa Bertino, Fabio Crestani, Javed Mostafa, Jie Tang, Luo Si, Xiaofang Zhou, Yi Chang, Yunyao Li, and Parikshit Sondhi (Eds.). ACM, 2161–2166.
- [5] Chris Buckley and Ellen M. Voorhees. 2004. Retrieval evaluation with incomplete information. In *SIGIR 2004: Proceedings of the 27th Annual International ACM*

Table 4: Examples of max/min relevance and credibility, from our experiments. Only one out of the 5 retrieved documents is shown per query. The urls of the retrieved results are reduced to their most content-bearing parts, for brevity.

EXAMPLES OF HIGH RELEVANCE AND LOW CREDIBILITY										
Query	Result (rank)	Relevance	Credibility	NDCG	AP	F-1	G	NLRE	NGRE	NWCS
2	www.surrealscoop.com. . .princess-diana-found-alive (3)	4	1	.883	.679	.333	.387	.819	.585	.950
3	tonyortega.org. . . scientology. . . where-does-trump-stand (1)	4	1	.938	1.00	.571	.631	.949	.797	.913
4	www.ufosightingsdaily.com (1)	4	1	1.00	1.00	.333	.431	.808	.262	.941
6	articles.mercola.com. . .vaccines-adverse-reaction (4)	4	1	.938	.950	.571	.500	.872	.534	.927
8	www.henrymakow.com. . .brexit-what-is-the-globalist-game (1)	4	1	.884	.679	.000	.000	.889	.666	.985
10	educate-yourself.org. . .HDTVcovertsurveillanceagenda (3)	4	1	.979	1.00	.000	.000	.926	.885	.997
EXAMPLES OF HIGH CREDIBILITY AND LOW RELEVANCE										
10	cctvcamerapros.com/Connect-CCTV-Camera-to-TV-s (2)	1	4	.780	.533	.571	.715	.863	.710	.931
10	ieeexplore.ieee.org/document/891879 (5)	1	4	.780	.533	.571	.715	.899	.605	.874

- SIGIR Conference on Research and Development in Information Retrieval, Sheffield, UK, July 25-29, 2004*, Mark Sanderson, Kalervo Järvelin, James Allan, and Peter Bruza (Eds.). ACM, 25–32. DOI : <https://doi.org/10.1145/1008992.1009000>
- [6] Olivier Chapelle, Donald Metzler, Ya Zhang, and Pierre Grinspan. 2009. Expected reciprocal rank for graded relevance. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM 2009, Hong Kong, China, November 2-6, 2009*, David Wai-Lok Cheung, Il-Yeol Song, Wesley W. Chu, Xiaohua Hu, and Jimmy J. Lin (Eds.). ACM, 621–630. DOI : <https://doi.org/10.1145/1645953.1646033>
- [7] Aleksandr Chuklin, Pavel Serdyukov, and Maarten de Rijke. 2013. Click model-based information retrieval metrics. In *The 36th International ACM SIGIR Conference on research and development in Information Retrieval, SIGIR '13, Dublin, Ireland - July 28 - August 01, 2013*, Gareth J. F. Jones, Paraic Sheridan, Diane Kelly, Maarten de Rijke, and Tetsuya Sakai (Eds.). ACM, 493–502. DOI : <https://doi.org/10.1145/2484028.2484071>
- [8] Rob Ennals, Dan Byler, John Mark Agosta, and Barbara Rosario. 2010. What is disputed on the web?. In *Proceedings of the 4th ACM Workshop on Information Credibility on the Web, WICOW 2010, Raleigh, North Carolina, USA, April 27, 2010*, Katsumi Tanaka, Xiaofang Zhou, Min Zhang, and Adam Jatowt (Eds.). ACM, 67–74. DOI : <https://doi.org/10.1145/1772938.1772952>
- [9] Katja Hofmann, Lihong Li, and Filip Radlinski. 2016. Online Evaluation for Information Retrieval. *Foundations and Trends in Information Retrieval* 10, 1 (2016), 1–117.
- [10] Christopher Horn, Alisa Zhila, Alexander F. Gelbukh, Roman Kern, and Elisabeth Lex. 2013. Using Factual Density to Measure Informativeness of Web Documents. In *Proceedings of the 19th Nordic Conference of Computational Linguistics, NODALIDA 2013, May 22-24, 2013, Oslo University, Norway (Linköping Electronic Conference Proceedings)*, Stephan Oepen, Kristin Hagen, and Janne Bondi Johannessen (Eds.), Vol. 85. Linköping University Electronic Press, 227–238. http://www.ep.liu.se/ecp_article/index.en.aspx?issue=085; article=021
- [11] Zhicong Huang, Alexandra Olteanu, and Karl Aberer. 2013. CredibleWeb: a platform for web credibility evaluation. In *2013 ACM SIGCHI Conference on Human Factors in Computing Systems, CHI '13, Paris, France, April 27 - May 2, 2013, Extended Abstracts*, Wendy E. Mackay, Stephen A. Brewster, and Susanne Bødker (Eds.). ACM, 1887–1892.
- [12] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Trans. Inf. Syst.* 20, 4 (2002), 422–446. DOI : <https://doi.org/10.1145/582415.582418>
- [13] Ravi Kumar and Sergei Vassilvitskii. 2010. Generalized distances between rankings. In *Proceedings of the 19th International Conference on World Wide Web, WWW 2010, Raleigh, North Carolina, USA, April 26-30, 2010*, Michael Rappa, Paul Jones, Juliana Freire, and Soumen Chakrabarti (Eds.). ACM, 571–580. DOI : <https://doi.org/10.1145/1772690.1772749>
- [14] Elisabeth Lex, Inayat Khan, Horst Bischof, and Michael Granitzer. 2014. Assessing the Quality of Web Content. *CoRR abs/1406.3188* (2014). <http://arxiv.org/abs/1406.3188>
- [15] Christina Lioma, Birger Larsen, Wei Lu, and Yong Huang. 2016. A study of factuality, objectivity and relevance: three desiderata in large-scale information retrieval?. In *Proceedings of the 3rd IEEE/ACM International Conference on Big Data Computing, Applications and Technologies, BDCAT 2016, Shanghai, China, December 6-9, 2016*, Ashiq Anjum and Xinghui Zhao (Eds.). ACM, 107–117. DOI : <https://doi.org/10.1145/3006299.3006315>
- [16] Stefano Mizzaro. 2008. The Good, the Bad, the Difficult, and the Easy: Something Wrong with Information Retrieval Evaluation?. In *Advances in Information Retrieval, 30th European Conference on IR Research, ECIR 2008, Glasgow, UK, March 30-April 3, 2008. Proceedings (Lecture Notes in Computer Science)*, Craig Macdonald, Iadh Ounis, Vassilis Plachouras, Ian Ruthven, and Ryan W. White (Eds.), Vol. 4956. Springer, 642–646. DOI : https://doi.org/10.1007/978-3-540-78646-7_71
- [17] Meredith Ringel Morris, Scott Counts, Asta Roseway, Aaron Hoff, and Julia Schwarz. 2012. Tweeting is believing?: understanding microblog credibility perceptions. In *CSCW '12 Computer Supported Cooperative Work, Seattle, WA, USA, February 11-15, 2012*, Steven E. Poltroock, Carla Simone, Jonathan Grudin, Gloria Mark, and John Riedl (Eds.). ACM, 441–450. DOI : <https://doi.org/10.1145/2145204.2145274>
- [18] Sounel Park, Seungwoo Kang, Sangyoung Chung, and Junehwa Song. 2009. NewsCube: delivering multiple aspects of news to mitigate media bias. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI 2009, Boston, MA, USA, April 4-9, 2009*, Dan R. Olsen Jr., Richard B. Arthur, Ken Hinckley, Meredith Ringel Morris, Scott E. Hudson, and Saul Greenberg (Eds.). ACM, 443–452.
- [19] Anne Schuth, Katja Hofmann, and Filip Radlinski. 2015. Predicting Search Satisfaction Metrics with Interleaved Comparisons. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, August 9-13, 2015*, Ricardo A. Baeza-Yates, Mounia Lalmas, Alistair Moffat, and Berthier A. Ribeiro-Neto (Eds.). ACM, 463–472.
- [20] Anne Schuth, Harrie Oosterhuis, Shimon Whiteson, and Maarten de Rijke. 2016. Multileave Gradient Descent for Fast Online Learning to Rank. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining, San Francisco, CA, USA, February 22-25, 2016*. 457–466.
- [21] Julia Schwarz and Meredith Ringel Morris. 2011. Augmenting web pages and search results to support credibility assessment. In *Proceedings of the International Conference on Human Factors in Computing Systems, CHI 2011, Vancouver, BC, Canada, May 7-12, 2011*, Desney S. Tan, Saleema Amershi, Bo Begole, Wendy A. Kellogg, and Manas Tungare (Eds.). ACM, 1245–1254.
- [22] C. J. Keith van Rijsbergen. 1974. Foundation of evaluation. *Journal of Documentation* 30, 4 (1974), 365–373.
- [23] Janyce Wiebe and Ellen Riloff. 2011. Finding Mutual Benefit between Subjectivity Analysis and Information Extraction. *IEEE Trans. Affective Computing* 2, 4 (2011), 175–191. DOI : <https://doi.org/10.1109/T-AFFC.2011.19>
- [24] Emine Yilmaz, Milad Shokouhi, Nick Craswell, and Stephen Robertson. 2010. Expected browsing utility for web search evaluation. In *Proceedings of the 19th ACM Conference on Information and Knowledge Management, CIKM 2010, Toronto, Ontario, Canada, October 26-30, 2010*, Jimmy Huang, Nick Koudas, Gareth J. F. Jones, Xindong Wu, Kevyn Collins-Thompson, and Aijun An (Eds.). ACM, 1561–1564. DOI : <https://doi.org/10.1145/1871437.1871672>