

Team OpenWebSearch at LongEval: Query and Document Rewriting

CLEF 2024, 9–12 September 2024, Grenoble, France




Daria
Alexander¹



**Maik
Fröbe**²



Gijs
Hendriksen¹



Ferdinand
Schlatt²



Martin
Potthast³



Matthias
Hagen²



Djoerd
Hiemstra¹



Arjen P.
de Vries¹

Radboud Universiteit



1



FRIEDRICH-SCHILLER-
UNIVERSITÄT
JENA

2




UNIVERSITÄT
LEIPZIG

3

webis.de

OWS at LongEval

Motivation: Imagine you Have Seen something Again and Again and Again...



OWS at LongEval

Motivation: Imagine you Have Seen something Again and Again and Again...



How should this change the behaviour of Bill Murray?

OWS at LongEval

Motivation: Imagine you Have Seen something Again and Again and Again...



How should this change the behaviour of a search engine?

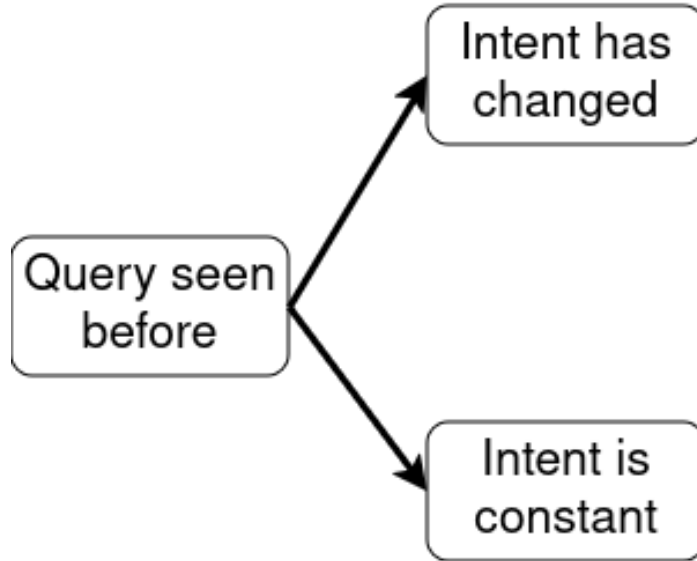
OWS at LongEval

How Should a Search Engine behave if it has seen the Query before?

Query seen
before


OWS at LongEval

How Should a Search Engine behave if it has seen the Query before?



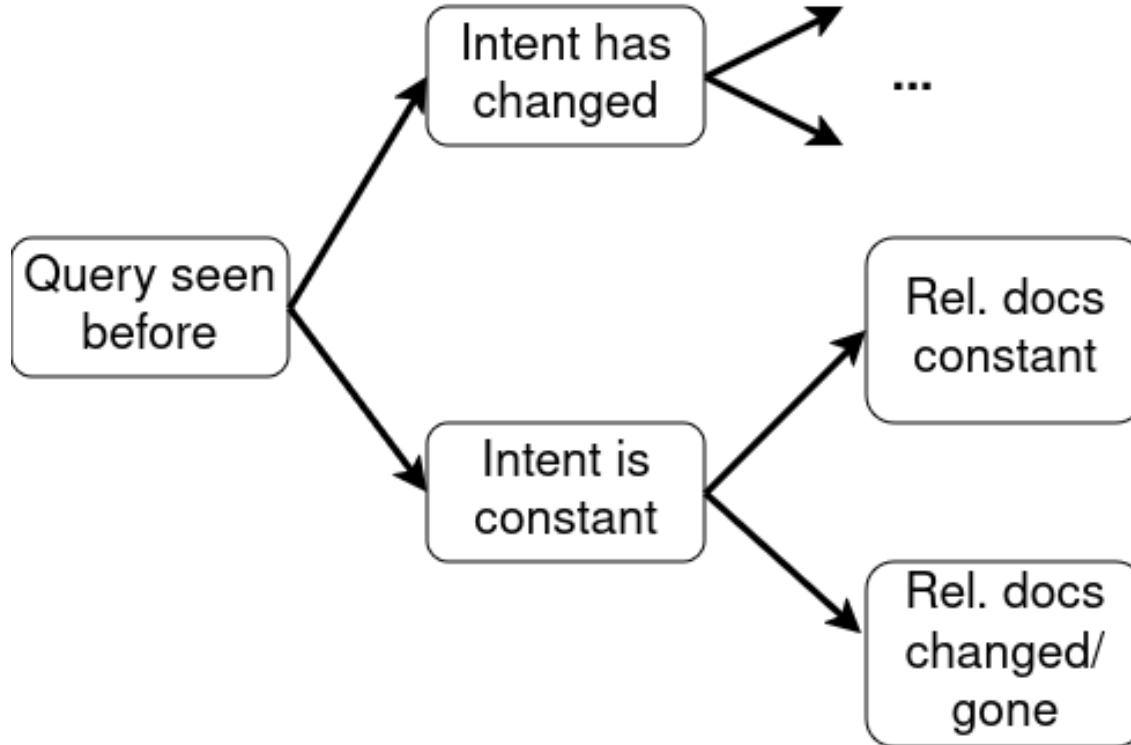
OWS at LongEval

How Should a Search Engine behave if it has seen the Query before?



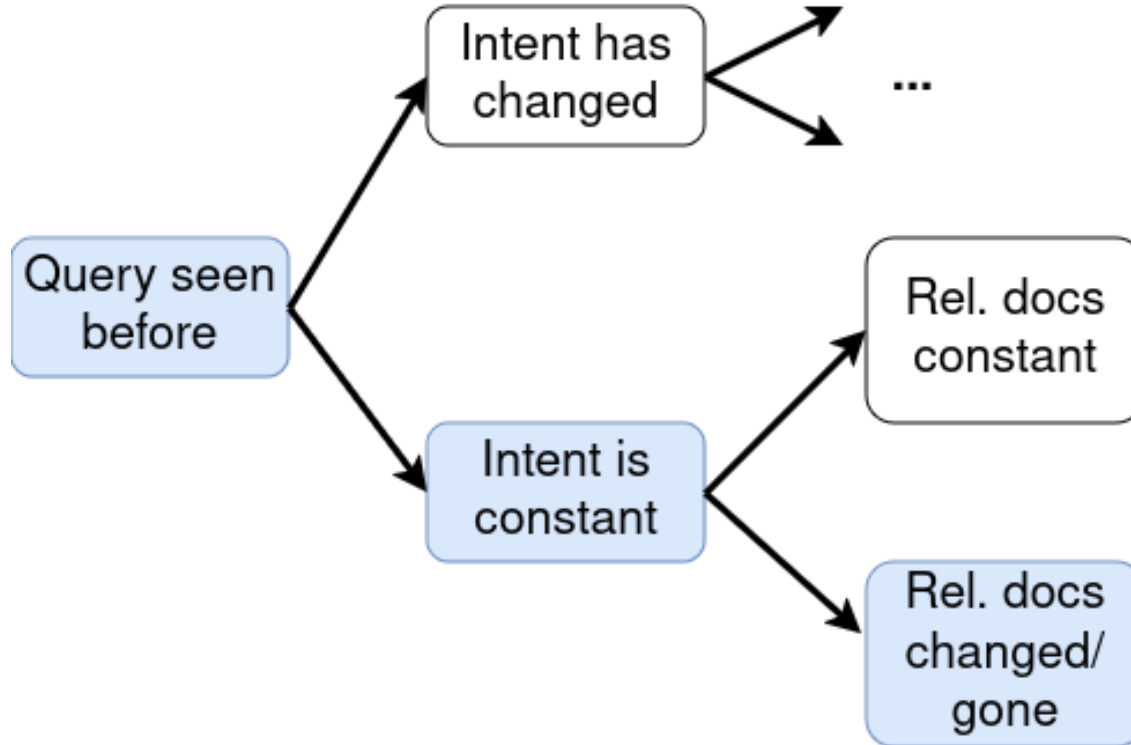
OWS at LongEval

How Should a Search Engine behave if it has seen the Query before?



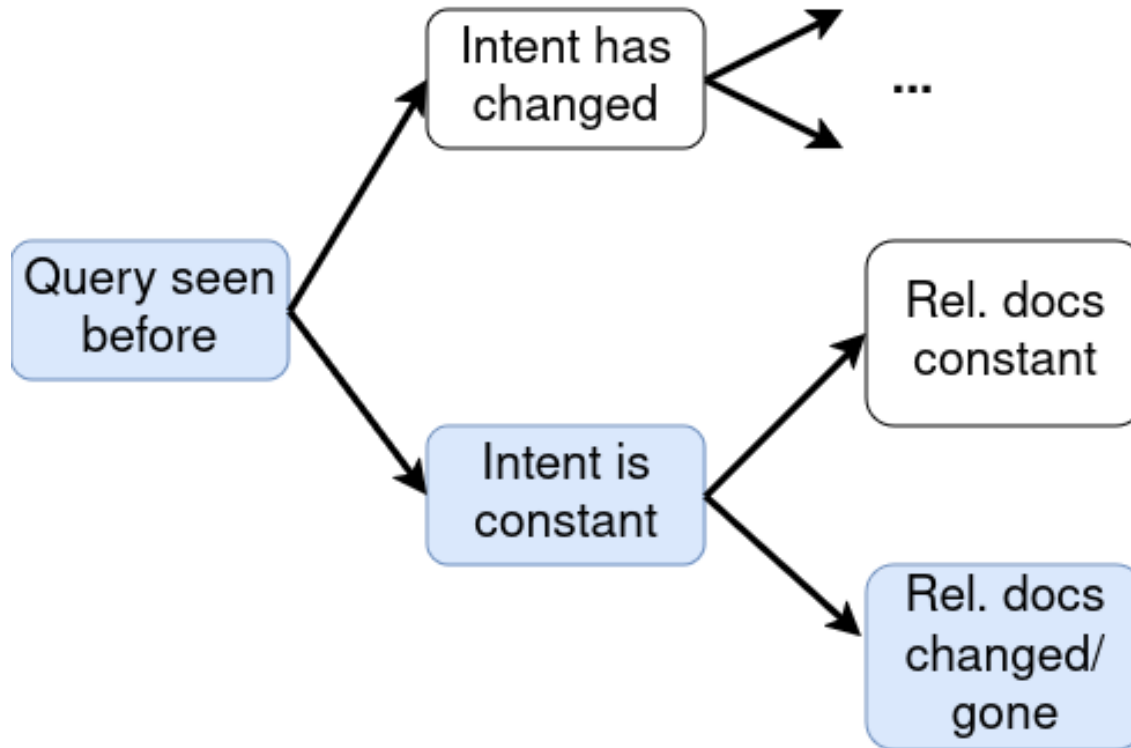
OWS at LongEval

How Should a Search Engine behave if it has seen the Query before?



OWS at LongEval

How Should a Search Engine behave if it has seen the Query before?



How can we incorporate past relevance information if documents might be gone?

- During indexing: rewrite documents via corpus graph and reverted index
- During retrieval: Rewriting the queries via keyqueries

OWS at LongEval

Past Relevance Information in the Index: Corpus Graph + Reverted Index

Query:

corpus to evaluate web search over time

OWS at LongEval

Past Relevance Information in the Index: Corpus Graph + Reverted Index

Query:

corpus to evaluate web search over time

Relevant document d :

LongEval-Retrieval: French-English Dynamic Test Collection for Continuous Web Search Evaluation

Petra Galuščáková*
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France
petra.galuscakova@univ-grenoble-
alpes.fr

Romain Deveaud*
Qwant
Paris, France
r.deveaud@qwant.com

Gabriela González Sáez
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Philippe Mulhem
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Lorraine Goeuriot
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Florina Piroi
Research Studios Austria,
Data Science Studio
Vienna, Austria

Martin Popel
Charles University
Prague, Czechia

OWS at LongEval

Past Relevance Information in the Index: Corpus Graph + Reverted Index

Query:

corpus to evaluate web search over time

Relevant document d :

LongEval-Retrieval: French-English Dynamic Test Collection for Continuous Web Search Evaluation

Petra Galuščáková*
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France
petra.galuscakova@univ-grenoble-
alpes.fr

Romain Deveaud*
Qwant
Paris, France
r.deveaud@qwant.com

Gabriela González Sáez
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Philippe Mulhem
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Lorraine Goeuriot
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Florina Piroi
Research Studios Austria,
Data Science Studio
Vienna, Austria

Martin Popel
Charles University
Prague, Czechia

Reverted Index:

[Pickens'10]

□ corpus to evaluate web search over time $\Rightarrow d$

OWS at LongEval

Past Relevance Information in the Index: Corpus Graph + Reverted Index

We observe more queries:

continuous evaluate web search over time

Relevant document d :

LongEval-Retrieval: French-English Dynamic Test Collection for Continuous Web Search Evaluation

Petra Galuščáková*
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France
petra.galuscakova@univ-grenoble-
alpes.fr

Romain Deveaud*
Qwant
Paris, France
r.deveaud@qwant.com

Gabriela González Sáez
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Philippe Mulhem
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Lorraine Goeuriot
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Florina Piroi
Research Studios Austria,
Data Science Studio
Vienna, Austria

Martin Popel
Charles University
Prague, Czechia

Reverted Index:

[Pickens'10]

- corpus to evaluate web search over time $\Rightarrow d$
- continuous evaluate web search over time $\Rightarrow d$

OWS at LongEval

Past Relevance Information in the Index: Corpus Graph + Reverted Index

Transfer the reverted index via the corpus graph:

[MacAvaney'22]

- For each query–document pair q, d in the reverted index:
 - Submit document d as query to the current index
 - Add $10 \cdot q$ to the top-1 hit
 - Add $9 \cdot q$ to the top-2 hit
 - Add $8 \cdot q$ to the top-3 hit
 - ...
- Do normal retrieval on the resulting reverted index

OWS at LongEval

Past Relevance Information in the Index: Corpus Graph + Reverted Index

Transfer the reverted index via the corpus graph:

[MacAvaney'22]

- For each query–document pair q, d in the reverted index:
 - Submit document d as query to the current index
 - Add $10 \cdot q$ to the top-1 hit
 - Add $9 \cdot q$ to the top-2 hit
 - Add $8 \cdot q$ to the top-3 hit
 - ...
- Do normal retrieval on the resulting reverted index

Results on Overlapping Queries

	Condensed nDCG@10	
	June	August
BM25	0.407	0.346
BM25@Reverted Index	0.429	0.371

OWS at LongEval

Past Relevance Information for Retrieval: Keyqueries

How to re-write queries for past relevance judgments?

OWS at LongEval

Past Relevance Information for Retrieval: Keyqueries

How to re-write queries for past relevance judgments?

We use the concept of so-called Keyqueries

[Gollub'13]

Given a document d and a query q . The query q is a keyquery for d if and only if:

1. q retrieves d within its top-results
2. q does not overfit to d
3. q is minimal

OWS at LongEval

Examples for Keyqueries

Lets build a keyquery for Document d :

LongEval-Retrieval: French-English Dynamic Test Collection for Continuous Web Search Evaluation

Petra Galuščáková*
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France
petra.galuscakova@univ-grenoble-
alpes.fr

Romain Deveaud*
Qwant
Paris, France
r.deveaud@qwant.com

Gabriela González Sáez
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Philippe Mulhem
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Lorraine Goeuriot
Univ. Grenoble Alpes, CNRS,
Grenoble INP[†], LIG
Grenoble, France

Florina Piroi
Research Studios Austria,
Data Science Studio
Vienna, Austria

Martin Popel
Charles University
Prague, Czechia

OWS at LongEval

Examples for Keyqueries

Lets build a keyquery for Document d :

The screenshot shows a search engine interface with a search bar containing the text "corpus to evaluate web search over time" and a magnifying glass icon. Below the search bar, it indicates "About 767,000 results (0.03 sec)". The results are listed as follows:

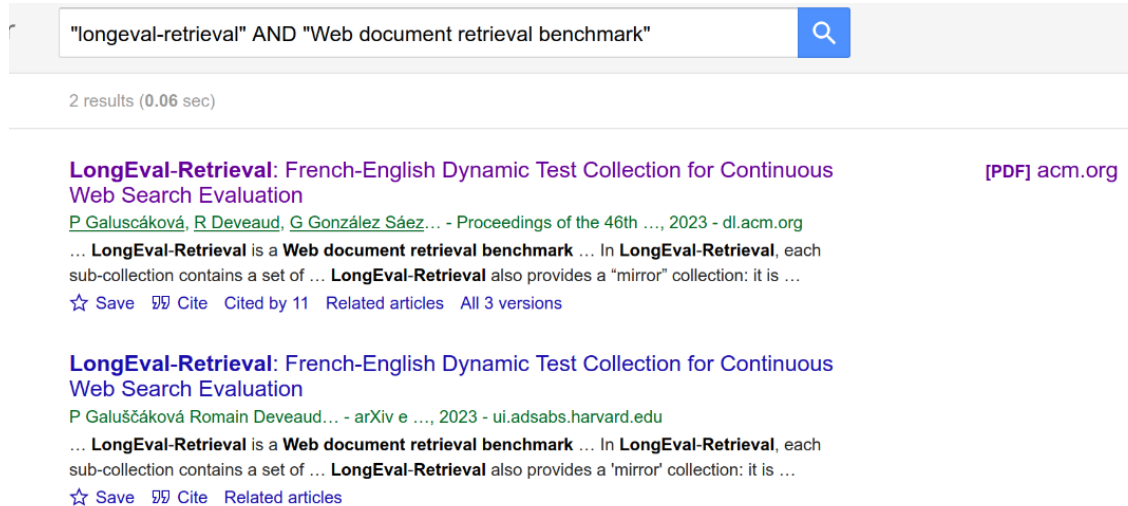
- [PDF] Corpus analysis of the world wide web** [PDF] psu.edu
WH Fletcher - The encyclopedia of applied linguistics, 2012 - Citeseer
... acquiring **Web** content and processing it into a static **corpus** or else accessing it directly as a ... **corpus**, a distinction captured in **Web for/as corpus** (De Schryver, 2002); here **Web corpus** ...
☆ Save 📄 Cite Cited by 112 Related articles All 3 versions 🔗
- [BOOK] Web corpus construction**
R Schäfer, F Bildhauer - 2013 - books.google.com
... Chapter 5 briefly touches upon the question of how we can **assess** the quality of a **web corpus** (... Now, since our goal is not **search** engine design but **web corpus** construction, we might ...
☆ Save 📄 Cite Cited by 108 Related articles All 7 versions 🔗
- [PDF] The 'body'and the 'web': The web as corpus ten years on** [PDF] uni.no
M Gatto - ICAME Journal, 2011 - clu.uni.no
... destroy its **value** as a source of linguistic information from a **corpus** linguistics perspective. ... control **over** the **corpus** itself, the **web**, but also **over** the **search** tools, ie the **search** engines, ...
☆ Save 📄 Cite Cited by 37 Related articles All 6 versions 🔗

Does not retrieve d in the top-results \Rightarrow no keyquery


OWS at LongEval

Examples for Keyqueries

Lets build a keyquery for Document d :



The screenshot shows a search engine interface with a search bar containing the query: "longeval-retrieval" AND "Web document retrieval benchmark". Below the search bar, it indicates "2 results (0.06 sec)". The first result is titled "LongEval-Retrieval: French-English Dynamic Test Collection for Continuous Web Search Evaluation" and is linked to "acm.org". The second result is titled "LongEval-Retrieval: French-English Dynamic Test Collection for Continuous Web Search Evaluation" and is linked to "arXiv e ...". Both results include a brief description of the collection and a list of actions: "Save", "Cite", and "Related articles".

"longeval-retrieval" AND "Web document retrieval benchmark" 

2 results (0.06 sec)

LongEval-Retrieval: French-English Dynamic Test Collection for Continuous Web Search Evaluation [\[PDF\] acm.org](#)
[P Galuscáková, R Deveaud, G González Sáez...](#) - Proceedings of the 46th ..., 2023 - dl.acm.org
... LongEval-Retrieval is a **Web document retrieval benchmark** ... In LongEval-Retrieval, each sub-collection contains a set of ... LongEval-Retrieval also provides a "mirror" collection: it is ...
[☆ Save](#) [🔗 Cite](#) [Cited by 11](#) [Related articles](#) [All 3 versions](#)

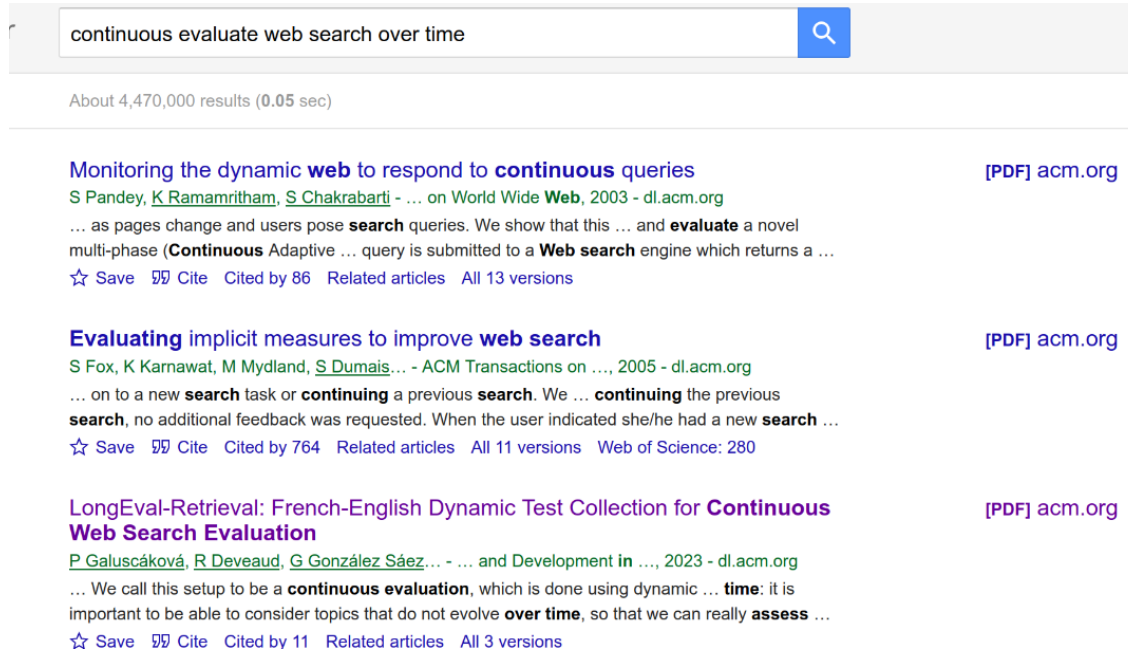
LongEval-Retrieval: French-English Dynamic Test Collection for Continuous Web Search Evaluation
[P Galuščáková Romain Deveaud...](#) - arXiv e ..., 2023 - ui.adsabs.harvard.edu
... LongEval-Retrieval is a **Web document retrieval benchmark** ... In LongEval-Retrieval, each sub-collection contains a set of ... LongEval-Retrieval also provides a 'mirror' collection: it is ...
[☆ Save](#) [🔗 Cite](#) [Related articles](#)

Overfits to $d \Rightarrow$ no keyquery

OWS at LongEval

Examples for Keyqueries

Lets build a keyquery for Document d :



The screenshot shows a search engine interface with a search bar containing the text "continuous evaluate web search over time" and a magnifying glass icon. Below the search bar, it indicates "About 4,470,000 results (0.05 sec)". Three search results are displayed, each with a title, a snippet, and a "[PDF] acm.org" link.

continuous evaluate web search over time

About 4,470,000 results (0.05 sec)

Monitoring the dynamic web to respond to continuous queries [PDF] acm.org
S Pandey, K Ramamritham, S Chakrabarti - ... on World Wide Web, 2003 - dl.acm.org
... as pages change and users pose **search** queries. We show that this ... and **evaluate** a novel multi-phase (**Continuous** Adaptive ... query is submitted to a **Web search** engine which returns a ...
☆ Save 📄 Cite Cited by 86 Related articles All 13 versions

Evaluating implicit measures to improve web search [PDF] acm.org
S Fox, K Karnawat, M Mydland, S Dumais... - ACM Transactions on ..., 2005 - dl.acm.org
... on to a new **search** task or **continuing** a previous **search**. We ... **continuing** the previous **search**, no additional feedback was requested. When the user indicated she/he had a new **search** ...
☆ Save 📄 Cite Cited by 764 Related articles All 11 versions Web of Science: 280

LongEval-Retrieval: French-English Dynamic Test Collection for Continuous Web Search Evaluation [PDF] acm.org
P Galuscáková, R Deveaud, G González Sáez... - ... and Development in ..., 2023 - dl.acm.org
... We call this setup to be a **continuous evaluation**, which is done using dynamic ... **time**: it is important to be able to consider topics that do not evolve **over time**, so that we can really **assess** ...
☆ Save 📄 Cite Cited by 11 Related articles All 3 versions

No overfitting + d in top-results \Rightarrow We have a keyquery

OWS at LongEval

Keyqueries on LongEval

To apply keyqueries, we:

- ❑ Counterfactually assume past relevant documents would still exist
- ❑ Remove all counterfactual documents from retrieval results

OWS at LongEval

Keyqueries on LongEval

To apply keyqueries, we:

- ❑ Counterfactually assume past relevant documents would still exist
- ❑ Remove all counterfactual documents from retrieval results

Results on Overlapping Queries

	Condensed nDCG@10	
	June	August
BM25	0.407	0.346
BM25@Reverted Index	0.429	0.371
KeyQueries	0.572	0.488

OWS at LongEval

Additionally, we also applied LTR

- ❑ We used LambdaMART
- ❑ Over 100 features (including reverted index + Keyqueries)
- ❑ Derived from submissions to the Workshop on Open Web Search

	Condensed nDCG@10	
	June	August
BM25	0.407	0.346
BM25@Reverted Index	0.429	0.371
KeyQueries	0.572	0.488
LTR	0.414	0.355

Conclusion

Focus on overlapping queries with the same intent but changed/gone documents

- ❑ Reverted index with corpus graph during indexing time
- ❑ Keyqueries during retrieval time
- ❑ Both effective stand-alone: still, keyqueries worked much better
- ❑ Are not really effective combined via LTR

Implementation available online: github.com/OpenWebSearch/LONGEVAL-24

Conclusion

Focus on overlapping queries with the same intent but changed/gone documents

- ❑ Reverted index with corpus graph during indexing time
- ❑ Keyqueries during retrieval time
- ❑ Both effective stand-alone: still, keyqueries worked much better
- ❑ Are not really effective combined via LTR

Implementation available online: github.com/OpenWebSearch/LONGEVAL-24

Future work

- ❑ Implement exhaustive search over possible keyqueries
- ❑ Try to bring the reverted index to the same effectiveness as keyqueries

Conclusion

Focus on overlapping queries with the same intent but changed/gone documents

- ❑ Reverted index with corpus graph during indexing time
- ❑ Keyqueries during retrieval time
- ❑ Both effective stand-alone: still, keyqueries worked much better
- ❑ Are not really effective combined via LTR

Implementation available online: github.com/OpenWebSearch/LONGEVAL-24

Future work

- ❑ Implement exhaustive search over possible keyqueries
- ❑ Try to bring the reverted index to the same effectiveness as keyqueries

Thank You!