

Query-Task Mapping

Michael Völske

Ehsan Fatehifar

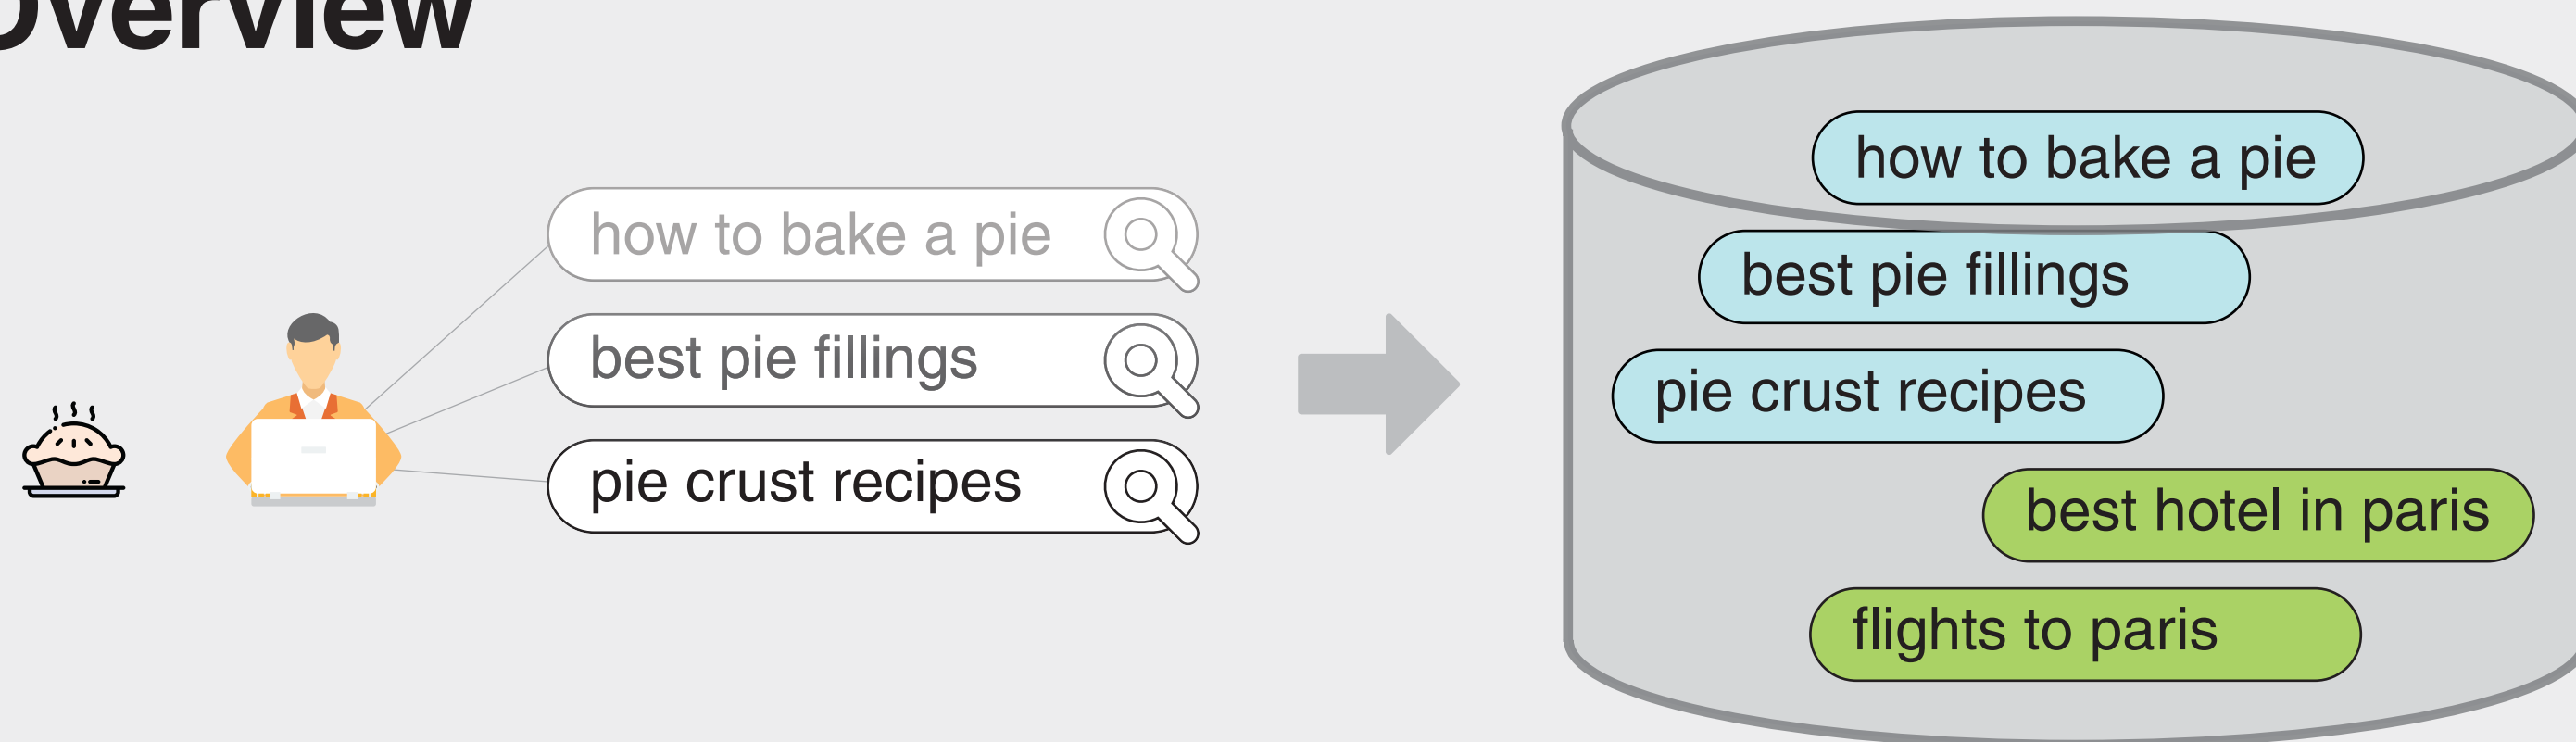
Benno Stein

Matthias Hagen

<firstname>.<lastname>@uni-weimar.de


matthias.hagen@informatik.uni-halle.de

Overview

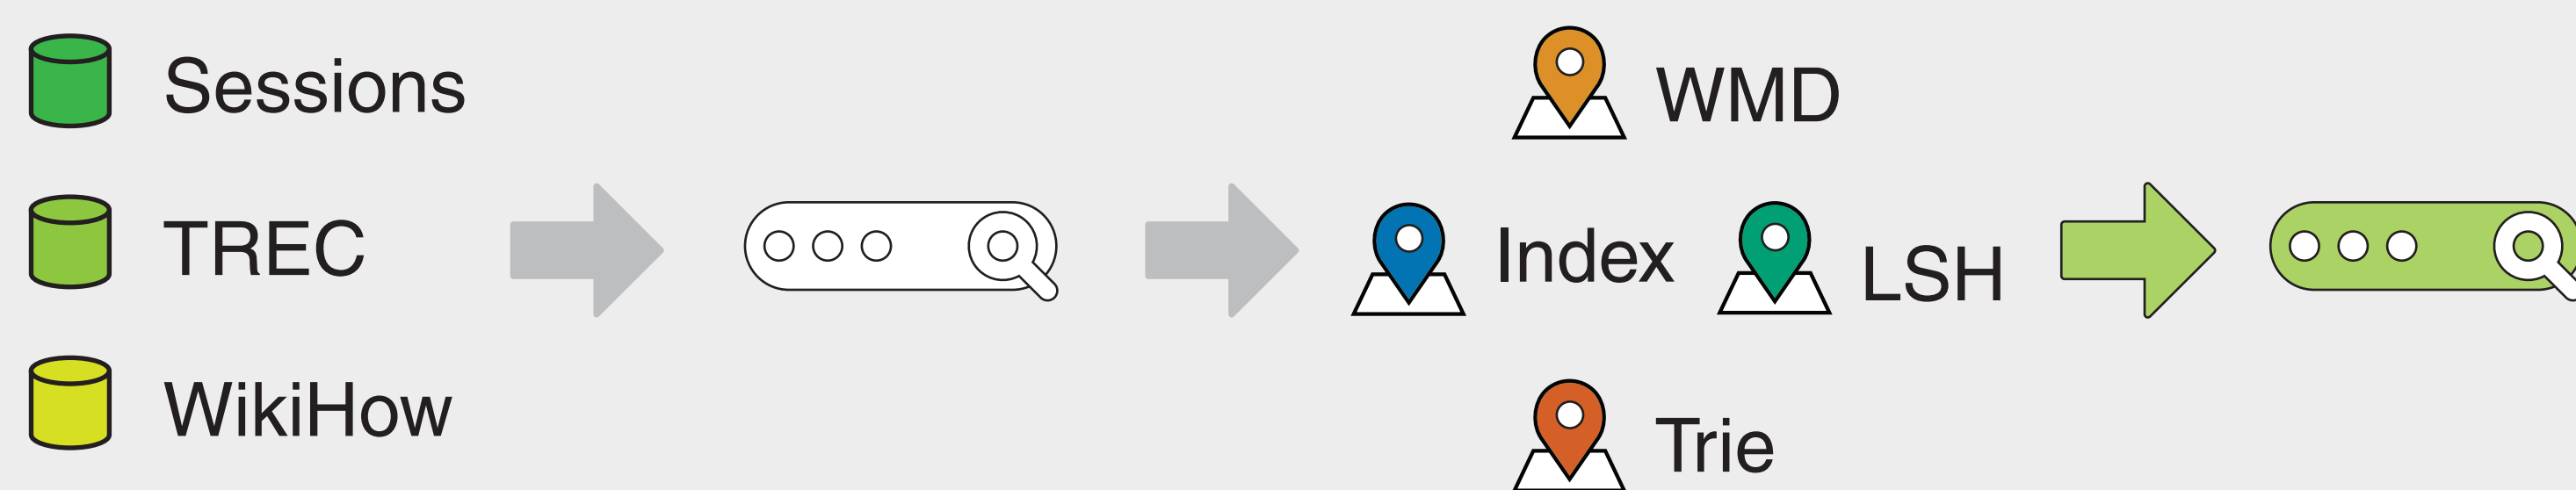


Search engine users submit their queries in order to accomplish some underlying task behind their search.

Previous work on task-based search has established methods for splitting a query log into sets of queries for the same task.



However, the related problem of **quickly and accurately mapping a new query to the best-fitting task** has received little attention so far.



We present three **benchmark datasets**, as well as four baseline **mapping approaches** for the query-task mapping problem.

Benchmark Datasets

We mine tasks from query logs and similar sources, and **extend the task size** with query suggestions mined from commercial search engines.

Session-based
Combination of Lucchese et al. [1] and Hagen et al. [2] annotated AOL Query Log excerpts.

[1] C. Lucchese et al., *Identifying Task-based Sessions in Search Engine Query Logs*. (WSDM '11).
[2] M. Hagen et al., *From Search Session Detection to Search Mission Detection*. (OAIR '13)

TREC-based
Real query logs on Web Track topics from the Webis-TRC-12 dataset [3] combined with Session Track queries.


[3] M. Hagen et al., *How Writers Search: Analyzing the Search and Writing Logs of Non-fictional Essays*. (CHIIR '16)

WikiHow-based
Procedural-Task questions from article titles, considered as same-task if annotated as mutually related by WikiHow users.

	Tasks	Queries	Queries per Task		
			min	avg	max
Session-based dataset					
Lucchese et al. [1]	223	771	1	3.5	55
+ Hagen et al. [2]	1,423	4,502	1	3.2	147
+ Google suggestions	1,423	29,441	1	20.7	924
+ Bing suggestions	1,423	41,780	1	29.4	1,368
TREC-based dataset					
Webis-TRC-12 query log [3]	150	3,848	1	25.7	122
+ Additional TREC topics	276	7,771	1	28.2	144
+ Google suggestions	276	38,478	8	139.4	858
+ Bing suggestions	276	47,514	8	172.2	997
WikiHow-based dataset					
WikiHow	7,202	15,914	1	2.2	22
+ Google suggestions	7,202	119,283	1	16.6	197
+ Bing suggestions	7,202	119,292	1	16.6	197

Mapping Approaches

Construct an index of the task-split query log during **preprocessing** (⚙️). Later, when **mapping** (🔍) a new query, rank the queries in the log and assign the majority task from among the top k matches.



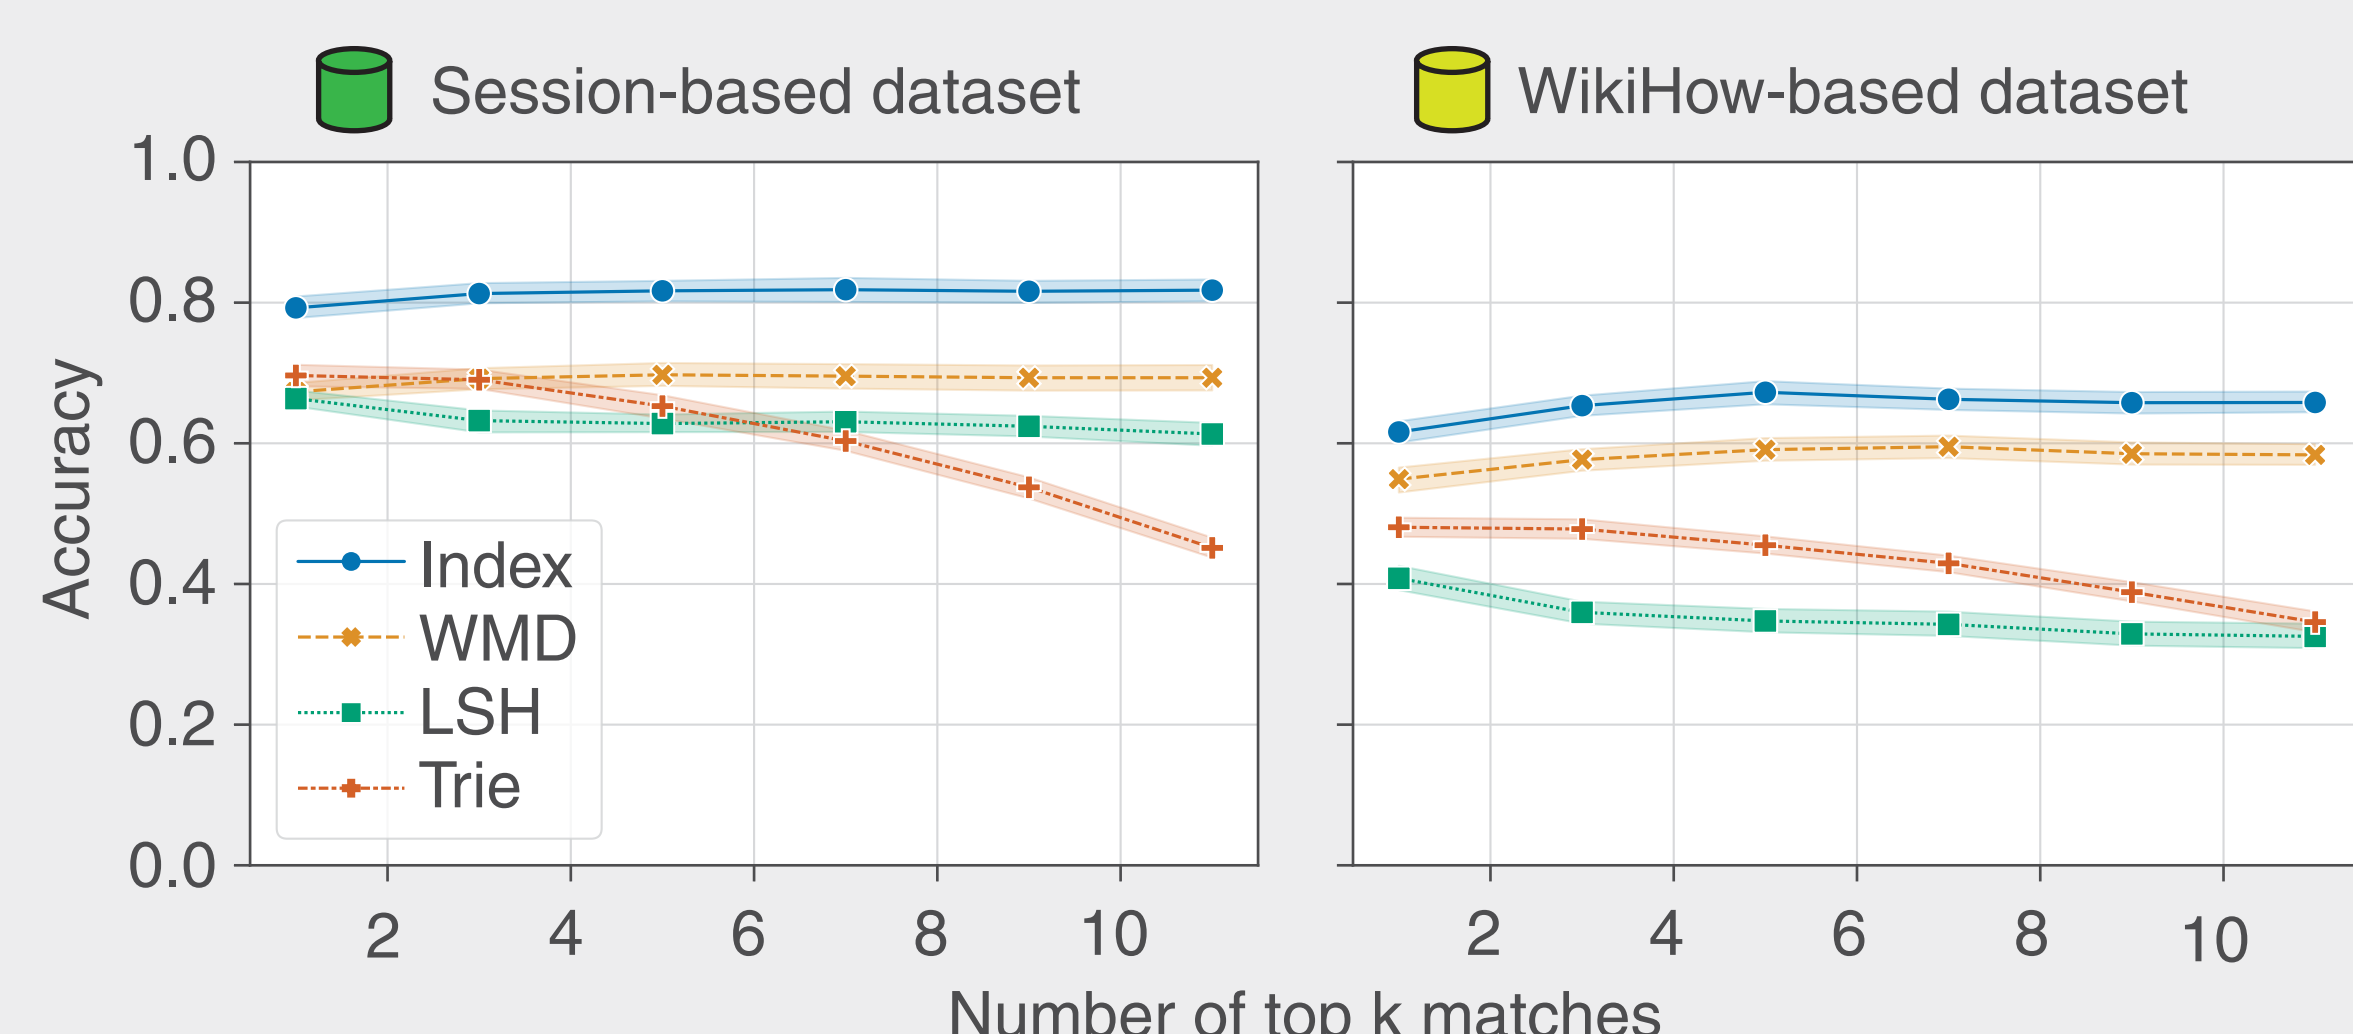
Index
⚙️ Inverted index of task-split query log
🔍 BM25 score
<https://www.elastic.co>

WMD
⚙️ Word2vec embeddings
🔍 Word mover's distance
<https://github.com/src-d/wmd-relax>

LSH
⚙️ MinHash-LSH of binary term vectors
🔍 Estimated Jaccard similarity
<https://github.com/ekzhu/datasketch>

Trie
⚙️ Trie of task-split query log
🔍 Length of matching prefix
<https://github.com/google/pygtrie>

	Index	WMD	LSH	Trie
Preprocessing time (entire dataset)				
Session-based	24.14s	9.60s	53.79s	10.03s
TREC-based	26.90s	11.14s	62.09s	13.26s
Wikihow-based	53.48s	26.50s	141.65s	28.00s
Query-task mapping time (per query)				
Session-based	2.80ms	7.16s	2.42ms	0.46ms
TREC-based	2.95ms	9.24s	2.50ms	0.51ms
Wikihow-based	4.21ms	22.65s	2.28ms	0.33ms
Query-task mapping accuracy				
Session-based	0.78	0.67	0.66	0.69
TREC-based	0.80	0.73	0.68	0.66
Wikihow-based	0.63	0.55	0.41	0.48




The procedural-task queries from the WikiHow dataset are noisier and harder to map; they benefit more from top-k smoothing.

Read the Paper



webis.de/downloads/publications/papers/stein_2019i.pdf

Get the Data



webis.de/data/webis-qtm-19.html