

Topical result caching in web search engines


Ida Mele, Nicola Tonellotto, Ophir Frieder, Raffaele Perego

Presented by: Rushikesh Padia

Overview

- Along with Static Dynamic cache, another layer for topic-based caching, where the entries are allocated to different topics (e.g., weather, education).
- Queries that are neither sufficiently popular to be in the static portion nor requested within short-time intervals to be in the dynamic portion.
- Captures moderately popular queries not captured by the static cache; additionally captures requests repeated within a large time interval that are likewise not captured by the dynamic cache.

Overview




Overview

- Query classification and query-topic distillation
- One possible approach for understanding the topic of a query is enriching its search keywords with the content (from snippet or page) of its top results.
- Another solution consists of using the content of only the clicked results (Used by author).
- Latent Dirichlet Allocation (LDA) is completely unsupervised and domain independent; additionally, it leverages the word co-occurrences, providing better results

Overview

- A topic cache of size $|T| = f_t \cdot N$ entries, which is in turn partitioned in k topic-based sections $\{t_1, t_2, \dots, t_k\}$, where k is the number of distinct topics.
- Each section is considered as an independent cache, managed with some caching policy (e.g., LRU or SDC).
- Each captures the specific temporal locality of the queries belonging to a given topic, i.e., queries more frequent in specific time intervals or with periodic “burstiness”. (e.g., queries on weather forecasting, typically issued in the morning, or queries on sport events, typically issued in the weekend);

Results



Reference

[1] I. Mele, N. Tonello, O. Frieder, and R. Perego, "Topical result caching in web search engines", *Information Processing & Management*, vol. 57, no. 3, p. 102193, 2020, doi: <https://doi.org/10.1016/j.ipm.2019.102193>.