

Calibration of Google Trends Time Series

Robert West
EPFL
robert.west@epfl.ch

ABSTRACT

Google Trends is a tool that allows researchers to analyze the popularity of Google search queries across time and space. In a single request, users can obtain time series for up to 5 queries on a common scale, normalized to the range from 0 to 100 and rounded to integer precision. Despite the overall value of Google Trends, rounding causes major problems, to the extent that entirely uninformative, all-zero time series may be returned for unpopular queries when requested together with more popular queries. We address this issue by proposing *Google Trends Anchor Bank (G-TAB)*, an efficient solution for the calibration of Google Trends data. Our method expresses the popularity of an arbitrary number of queries on a common scale without being compromised by rounding errors. The method proceeds in two phases. In the offline preprocessing phase, an “anchor bank” is constructed, a set of queries spanning the full spectrum of popularity, all calibrated against a common reference query by carefully chaining together multiple Google Trends requests. In the online deployment phase, any given search query is calibrated by performing an efficient binary search in the anchor bank. Each search step requires one Google Trends request, but few steps suffice, as we demonstrate in an empirical evaluation. We make our code publicly available as an easy-to-use library at <https://github.com/epfl-dlab/GoogleTrendsAnchorBank>.

ACM Reference Format:

Robert West. 2020. Calibration of Google Trends Time Series. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*, October 19–23, 2020, Virtual Event, Ireland. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3340531.3412075>

1 INTRODUCTION

Google makes aggregate statistics about the popularity of search queries publicly available via Google Trends. The platform has become an important tool for researchers across disciplines (e.g., health [6] or economics [4]), journalists [7], and many others.

Google Trends is accessible via an official Web interface,¹ as well as via unofficial APIs wrapping the Web interface.² Users of Google Trends specify as input up to 5 search queries (or “topics”), a time range, and a geographic region. Search queries may

¹<https://www.google.com/trends>

²<https://github.com/GeneralMills/pytrends>, <https://github.com/PMassicotte/gtrendsR>

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '20, October 19–23, 2020, Virtual Event, Ireland

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6859-9/20/10...\$15.00

<https://doi.org/10.1145/3340531.3412075>

be specified as plain text (e.g., “Munich population”) or as entity identifiers from the Freebase knowledge base [3] (e.g., `/m/02h6_6p` for MUNICH). Freebase identifiers are particularly handy as they allow for grouping various surface forms relating to the same topic, across languages (the English queries “Munich” and “Munich population”, German “München”, Italian “Monaco di Baviera”, etc., are all counted toward the entity `/m/02h6_6p`).

As output, Google Trends returns, among other things, time series of *search interest* for the specified input queries. Importantly, search interest is not returned in terms of absolute search volume, but normalized as described in the Google Trends FAQ:³ “Each data point is divided by the total searches of the geography and time range it represents to compare relative popularity. [...] The resulting numbers are then scaled on a range of 0 to 100 based on a topic’s proportion to all searches on all topics.” Finally, the resulting search interest numbers are rounded to integer precision.

Whereas normalizing by geography and time is certainly useful, scaling and rounding all numbers to integers between 0 and 100 poses considerable problems. In particular, it is impossible to

- (1) directly compare more than 5 search queries, even for a fixed geographic region and time span (due to scaling),
- (2) directly compare queries with vastly different search interests (due to rounding errors).

To get around problem 1, one might be tempted to always include among the up to 5 input queries a fixed reference query, against which all other queries could be compared. But this solution will unfortunately fail for most queries due to problem 2. Fig. 1(a–b) illustrates with an example. We used Google Trends to obtain the worldwide search interest in 5 Bavarian towns over the course of 2019. The fact that the search interest in MUNICH vastly outweighs that in the 4 other towns makes the time series for the less popular towns close to useless due to rounding errors. In particular, the time series for the least popular town, ARNSTORF, is 0 everywhere.

We resolve these issues with *Google Trends Anchor Bank (G-TAB)*, a method for obtaining calibrated Google Trends time series that

- (1) are directly comparable across any number of queries (addressing problem 1) and
- (2) are not compromised by rounding errors (addressing problem 2).

In the above example, G-TAB returns the calibrated time series of Fig. 1(c), depicted here together with the time series for about 200 more Bavarian towns. As opposed to the raw, rounded Google Trends data shown in Fig. 1(a–b), G-TAB returns all time series with high precision across orders of magnitude, revealing details that remained hidden in the raw Google Trends results (cf. Sec. 3).

We make our code publicly available as an easy-to-use library at <https://github.com/epfl-dlab/GoogleTrendsAnchorBank>.

³<https://support.google.com/trends/answer/4365533?hl=en>

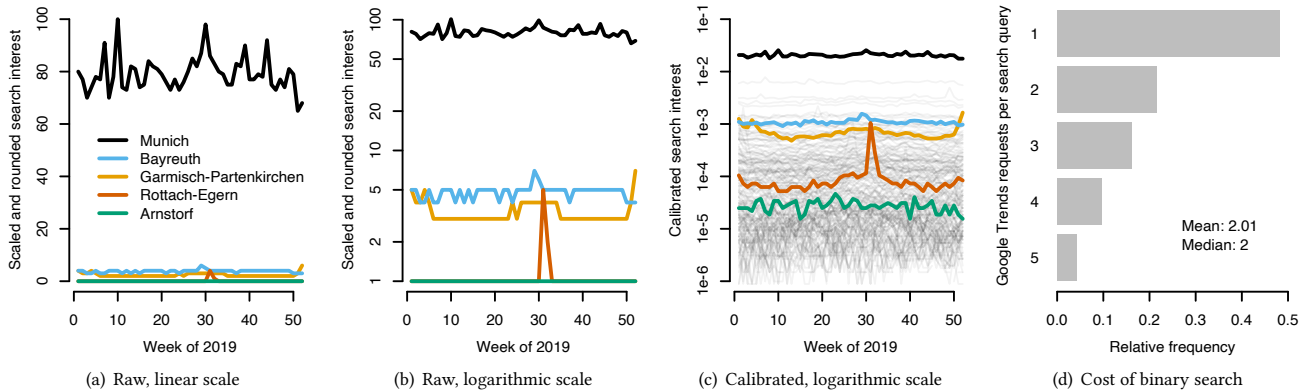


Figure 1: Google search interest in Bavarian towns. (a) Raw time series (scaled, rounded) for 5 towns as returned directly by Google Trends. (b) *Idem* on logarithmic scale (incremented by 1 in order to visualize values of 0). (c) Calibrated time series for 200 towns as returned by our method, G-TAB. (d) Number of Google Trends requests required for calibrating one search query.

2 GOOGLE TRENDS ANCHOR BANK (G-TAB)

Our calibration method, G-TAB, consists of two phases: offline preprocessing and online deployment. For a schematic description, see Fig. 2. All calibration is done for a given geographic region (worldwide or a specific country) and a given time span.

During **offline preprocessing** (Sec. 2.1; left panel of Fig. 2), we find a set of queries that collectively span all relevant orders of magnitude in terms of search interest. We refer to these queries as *anchor queries*, and to their entirety as the *anchor bank*. One of the anchor queries is selected as the *reference query* Q , and each anchor query x is calibrated against Q by calculating the *calibrated maximum search interest* in x , i.e., the maximum search interest attained by x , divided by that attained by Q .

During **online deployment** (Sec. 2.2; right panel of Fig. 2), we are given an arbitrary query q and return its *calibrated search interest* time series, which expresses the search interest in q as a fraction of the maximum search interest attained by the reference query Q . To achieve this, we first perform a binary search in the anchor bank in order to find an anchor query x whose search interest is on the same order of magnitude as that of q , such that q and x can be compared without harmful rounding errors. Then, since the anchor query x has already been calibrated against the reference query Q during offline preprocessing, the input query q , too, is readily calibrated against Q .

2.1 Offline phase: Building the anchor bank

In order to construct the anchor bank, the offline phase proceeds in three steps (cf. Fig. 2).

Step 1: Sample anchor queries. First, we identify an appropriate set of anchor queries, i.e., a set of queries that are likely to collectively cover a wide spectrum of search interest. Anchor queries need to be found heuristically, since prior to this step, we have no information yet related to the search interest in any queries.

As a rough proxy for search interest, we use the mention frequency of Freebase [3] entities in the large-scale ClueWeb corpus [5], as provided by the Freebase Easy project [2]. We fix an

entity type with a large number of instances (e.g., persons, towns, foods), sort the instances by ClueWeb mention frequency, discard all but the top N , and finally select a stratified sample of size n .

As mentioned, our method calibrates all queries against a reference query Q . The reference query may be chosen manually ahead of time and added to the set of anchor queries, or it may be chosen from among the anchor queries *post hoc*, e.g., as the most searched-for anchor query as determined later on, in step 3.

Step 2: Request time series from Google Trends. Next, we sort the n anchor queries by ClueWeb mention frequency and construct one Google Trends request for each set of k contiguous queries, in a shingling fashion, for a total of $n - k + 1$ requests. In Fig. 2, we illustrate with $k = 3$ for simplicity’s sake, although in practice we use the largest value allowed by Google Trends, $k = 5$.

The purpose of constructing overlapping Google Trends requests is to overcome the limitations posed by rounding errors that prevent us from directly comparing queries of vastly different search interest (cf. Fig. 1(a–b)). By later (in step 3) chaining together multiple

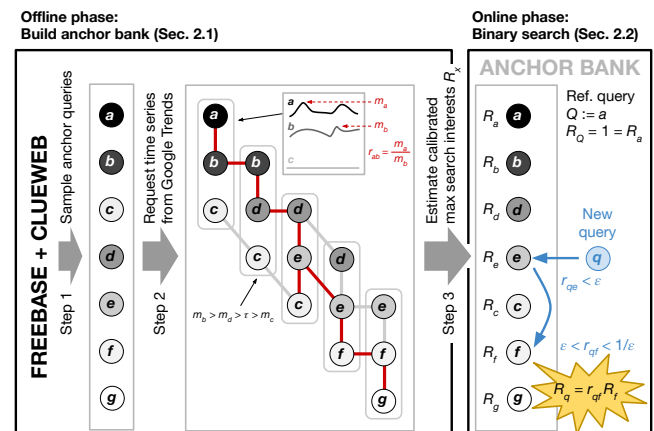


Figure 2: Overview of Google Trends Anchor Bank.

overlapping Google Trends requests, we can compare queries across orders of magnitude in search interest transitively.

Step 3: Estimate calibrated maximum search interests. For any query pair (x, y) , we define the *maximum ratio*

$$r_{xy}^* := \frac{M_x^*}{M_y^*} = \frac{m_x^*}{m_y^*}, \quad (1)$$

where $M_x^*, M_y^* \in \mathbb{R}_+$ denote the true, unobserved maximum search interests of x and y before scaling and rounding to $\{0, \dots, 100\}$, and $m_x^*, m_y^* \in [0, 100]$ denote the scaled—but still unrounded, and thus also unobserved—versions of M_x^* and M_y^* . The true maximum ratio r_{xy}^* is unobserved, but for anchor queries x, y co-occurring in the same Google Trends request, it can be estimated as

$$r_{xy} := \frac{m_x}{m_y}, \quad (2)$$

where $m_x, m_y \in \{0, \dots, 100\}$ denote the maximum values attained by the scaled and rounded time series for x and y as returned by Google Trends.

The estimate r_{xy} will be less reliable if m_x or m_y is small, due to larger relative rounding errors. We hence discard the query pair (x, y) if $\min\{m_x, m_y\} < \tau$, where $\tau \in \{0, \dots, 100\}$ is a predetermined threshold. For instance, choosing $\tau = 10$ ensures that the two anchor queries in each retained pair fall into the same order of magnitude with respect to search interest. One purpose of constructing Google Trends requests by grouping queries of similar ClueWeb mention frequency—a rough proxy for search interest—is to keep the number of discarded query pairs low.

Maximum ratios can be estimated directly only for anchor queries x, y that co-occurred in at least one of the Google Trends requests made in step 2. For other anchor query pairs, maximum ratios can be estimated inductively, as exemplified in Fig. 2 by the red tree connecting all anchor queries; e.g., although the queries a, d did not co-occur in any Google Trends request, their maximum ratio r_{ad}^* can still be estimated because a and d both co-occurred with b .

Formally, maximum ratios can be inferred based on the relation

$$r_{xz}^* = \frac{m_x^*}{m_z^*} = \frac{m_x^*}{m_y^*} \frac{m_y^*}{m_z^*} = r_{xy}^* r_{yz}^*, \quad (3)$$

which gives rise to a simple dynamic programming algorithm: if the estimates r_{xy} and r_{yz} have already been computed, but not so r_{xz} , then set $r_{xz} := r_{xy} r_{yz}$. Repeating this procedure iteratively allows us to compute r_{xz} for all anchor query pairs (x, z) , even if x and z did not co-occur in any of the Google Trends requests.

There are in general multiple chains for transitively relating two queries x, z via this procedure, and not all chains lead to equally tight estimates r_{xz} of r_{xz}^* : due to integer rounding,

$$m_x^* \in [m_x^-, m_x^+], \quad \text{where} \quad m_x^- = m_x - \frac{1}{2}, \quad m_x^+ = m_x + \frac{1}{2} \quad (4)$$

(except if $m_x = 100$, where there is no rounding, so $m_x^- = m_x^+ = 100$). Hence, for x, y that co-occurred in the same Google Trends request,

$$r_{xy}^* \in \left[\frac{m_x^-}{m_y^+}, \frac{m_x^+}{m_y^-} \right]. \quad (5)$$

We can capture the tightness of the estimate r_{xy} of r_{xy}^* as the ratio η_{xy} of the upper and lower bounds:

$$\eta_{xy} := \frac{m_x^+ / m_y^-}{m_x^- / m_y^+} = \frac{m_x^+}{m_x^-} \frac{m_y^+}{m_y^-}. \quad (6)$$

An indirect estimate $r_{xz} = r_{xy} r_{yz}$ comes with the corresponding bound ratio $\eta_{xz} = \eta_{xy} \eta_{yz}$.

For each query pair (x, z) , we are thus interested in finding a chain $\langle x = x_1, x_2, \dots, x_l = z \rangle$ of queries such that any two adjacent queries (x_i, x_{i+1}) co-occurred in the same Google Trends request and the product $\eta_{xz} = \prod_{i=1}^{l-1} \eta_{x_i x_{i+1}}$ is minimized. This problem is readily solved by finding a shortest path from x to z in a weighted directed graph G with queries as nodes, edges between queries that co-occurred in the same Google Trends request, and edge weights $w_{xy} := \log \eta_{xy}$.⁴ The product $r_{xz} = \prod_{i=1}^{l-1} r_{x_i x_{i+1}}$ along the shortest path from x to z yields the tightest estimate of r_{xz}^* that is possible given the Google Trends results at hand.

The goal of offline preprocessing is to calibrate each anchor query x against the reference query Q , as captured by the maximum ratio r_{xQ} , henceforth termed x 's *calibrated maximum search interest*

$$R_x := r_{xQ}. \quad (7)$$

The calibrated maximum search interest R_x expresses the maximum search interest in query x as a fraction of the maximum search interest in the reference query Q . The result of offline preprocessing—the anchor bank—consists of the list of anchor queries x sorted in increasing order of their calibrated maximum search interests R_x .

In the next section, we describe how to efficiently calibrate Google Trends results for any arbitrary query q , rather than for anchor queries only.

Obtaining an optimal anchor bank. In Appendix A,⁵ we show that, in order to obtain the most precisely calibrated maximum search interests (i.e., subject to the lowest rounding errors), the anchor bank should consist of a list of queries such that neighboring queries x, y have a constant maximum ratio $r_{xy} \approx 1/e \approx 0.37$. As shown in Appendix B, this ideal can be approximated in a second round of Google Trends requests, based on the results of the above-described first round.

2.2 Online phase: Binary search

In the online deployment phase, we are given any Google query q and return a calibrated time series of search interest, i.e., a time series of q 's search interest expressed as a fraction of the maximum search interest achieved by the reference query Q (in the specified geographic region and time span). This is achieved by multiplying q 's uncalibrated search interest time series, as obtained directly from Google Trends, by R_q/m_q . The task, thus, is to compute R_q .

In principle, $R_q := r_{qQ}$ could be measured by including q and Q in the same Google Trends request and estimating their maximum ratio $r_{qQ} = m_q/m_Q$. In practice, however, this direct estimate may be unusable due to errors incurred by rounding m_q and m_Q to

⁴To begin with, there are generally multiple edges per (x, y) , one for each Google Trends request in which both x and y occurred. As we are interested in shortest paths in G , we define the edge for (x, y) in G based on the request with the smallest η_{xy} .

⁵Appendices available online at <https://arxiv.org/abs/2007.13861>

integer precision. To overcome this issue, we observe that, for any anchor query x , the following equality holds:

$$R_q = r_{qQ} = r_{qx}r_{xQ} = r_{qx}R_x. \quad (8)$$

Here, R_x is already known from offline preprocessing, and the maximum ratio $r_{qx} = m_q/m_x$ can be reliably computed from the uncalibrated Google Trends results, provided that $m_q, m_x \in \{0, \dots, 100\}$ are both reasonably large, for then rounding errors will be small.

The main challenge of the online phase is therefore to search the anchor bank for an anchor query x^* such that neither m_q nor m_{x^*} is too small in the result returned by Google Trends for a joint request for q and x^* . To quickly find such an anchor query x^* , we apply binary search: Let A be the anchor bank, i.e., a list of all anchor queries x in increasing order of R_x . Compare q to the mid point x' of A . If $\epsilon < r_{qx'} < 1/\epsilon$ (where $\epsilon < 1$ is a fixed parameter), terminate the search with $x^* := x'$. Otherwise, if $r_{qx'} < \epsilon$, recursively search in the left half of A , and if $r_{qx'} > 1/\epsilon$, in the right half of A .

Computing each maximum ratio $r_{qx'} = m_q/m_{x'}$ requires sending a joint Google Trends request for the two queries q and x' . In practice, the search terminates after a very small number of steps (cf. Sec. 3), so the overhead incurred by calibration is generally low.

Quantifying uncertainty. Due to integer rounding, the estimate R_q is approximate. To quantify the uncertainty, we compute upper and lower bounds for R_q , as well as for q 's calibrated time series of search interest, based on the calculations of step 3 in Sec. 2.1.

3 EXAMPLE RESULTS

We now showcase the power of G-TAB empirically. In step 1 of the offline phase (Sec. 2.1), we consider as anchor queries all food entities from Freebase (types `FOOD` and `DISH`) and sample $n = 100$ from the top $N = 2000$ entities in a stratified manner.

Search interest generally follows a heavy-tailed distribution, headed by navigational queries for common websites [1]. In order to also cover these head queries, we manually add 6 common navigational queries to the anchor bank. Manual probing of Google Trends revealed `FACEBOOK` as probably the most popular Google query, which we use as the reference query Q .

In step 2, we group $k = 5$ queries per Google Trends request; in step 3, we use $\tau = 10$; and during the online phase, $\epsilon = 0.1$.

Example 1: Bavarian towns. Resuming the example from Sec. 1, we select the top 100 Bavarian towns mentioned most frequently in ClueWeb, unioned with 100 more sampled from the top 1000, and analyze their worldwide Google search interest during 2019. Queries are represented as Freebase IDs, rather than plain text.

Our calibration method can place an arbitrary number of queries on a common scale, as exemplified in Fig. 1(c), which shows the time series for all 200 towns in one single plot. We see that `MUNICH` is about 2%, and `ARNSTORF` about 0.003% as popular as the reference query `FACEBOOK`. Moreover, although search interest spans 5 orders of magnitude, all 200 time series are available at high precision, revealing insights that remain hidden in the results obtained directly from Google Trends (cf. Fig. 1(b)). For instance, `ROTTACH-EGERN` is consistently more popular than `ARNSTORF`, and its popularity grows gradually in summer, rather than as an impulse in week 32.

What price do we need to pay for calibration? Fig. 1(d) answers this question by showing the distribution of the number of Google

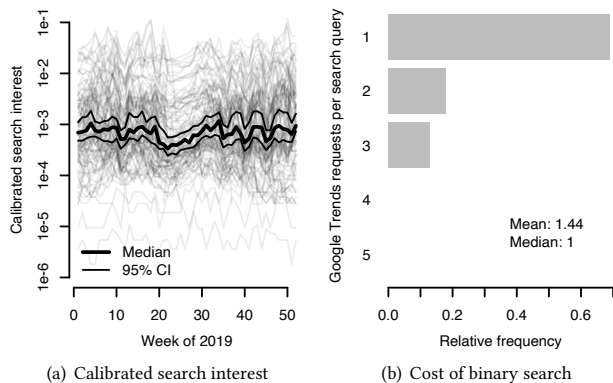


Figure 3: Search interest in top 100 soccer clubs.

Trends requests required during the binary search in order to process one query. We see that, on average, only two Google Trends requests are needed to calibrate one query.

Example 2: Soccer clubs. As a second example, Fig. 3(a) plots calibrated time series for the 100 soccer clubs mentioned most frequently in ClueWeb. The median over all clubs is shown as a thick black line (with bootstrapped 95% confidence intervals). It reveals that the search interest in soccer clubs drops in summer (mid May to mid July), when many soccer leagues break between seasons. Without calibration, such insights could not emerge.

Calibration is even cheaper here than in example 1. Most queries require only a single Google Trends request (mean 1.44; Fig. 3(b)). Generally, the better the distribution of R_x in the anchor bank matches the distribution seen in online deployment, the more efficient the binary search.

4 CONCLUSION

Google Trends has emerged as a Swiss army knife for data scientists. So far, users have been hampered by its imprecision, stemming from the fact that all results are scaled and rounded to integers from 0 to 100. With G-TAB, our simple, effective, and efficient calibration method, we hope to have sharpened the knife for all its users.

Acknowledgments. I thank Gorjan Popovski and Manoel Horta Ribeiro for help with Python coding, and Microsoft, Google, Facebook for funding. Dedicated to Barbara West and Johann Winkler.

REFERENCES

- [1] Ricardo Baeza-Yates and Berthier Ribeiro-Neto. 1999. *Modern Information Retrieval*. ACM Press.
- [2] Hannah Bast, Florian Bärle, Björn Buchhold, and Elmar Haußmann. 2014. Easy access to the Freebase dataset. In *Proc. WWW*.
- [3] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: A collaboratively created graph database for structuring human knowledge. In *Proc. SIGMOD*.
- [4] Hyunyoung Choi and Hal Varian. 2012. Predicting the present with Google Trends. *Economic Record* 88 (2012), 2–9.
- [5] Evgeniy Gabrilovich, Michael Ringgaard, and Amarnag Subramanya. 2013. FACC1: Freebase annotation of ClueWeb corpora, Version 1 (Release date 2013-06-26).
- [6] Jeremy Ginsberg, Matthew H Mohebbi, Rajan S Patel, Lynnette Brammer, Mark S Smolinski, and Larry Brilliant. 2009. Detecting influenza epidemics using search engine query data. *Nature* 457, 7232 (2009), 1012–1014.
- [7] Seth Stephens-Davidowitz. 2012. How racist are we? Ask Google. *New York Times* (9 June 2012).

A PROPERTIES OF AN OPTIMAL ANCHOR BANK

An anchor bank is a list $A = (x_1, x_2, \dots, x_n)$ of anchor queries sorted in increasing order of calibrated maximum search interest, i.e., $R_{x_i} < R_{x_{i+1}}$. In order to be able to compare Google queries of arbitrary search interest on the same scale using A , we assume that x_1 is the least frequent, and x_n the most frequent, Google query (in practice, the extreme queries can be chosen heuristically by manual probing). Without loss of generality, we assume the reference query $Q = x_n$. That is, in order to calibrate the anchor queries, we need to estimate $r_{x_i x_n}^*$ for all $1 \leq i < n$.

Here we ask the question: What are the properties of an optimal anchor bank?

First, let us define optimality: in an optimal anchor bank for fixed least and most frequent queries x_1 and x_n , the estimate $R_{x_1} = r_{x_1 x_n}$ (the estimated maximum ratio of the most distant query pair) should be as close as possible to the true, unobserved

$$r^* := r_{x_1 x_n}^*, \quad (9)$$

in the sense that the bound ratio $\eta_{x_1 x_n}$ (cf. (6)) should be as small, i.e., as close to 1, as possible. Additionally, if multiple anchor banks meet this criterion, an anchor bank is optimal if it does so with the smallest number n of anchor queries.

This definition entails that, in an optimal anchor bank A^* , the best way to estimate r^* is via the product

$$r_{x_1 x_n} = \prod_{i=1}^{n-1} r_{x_i x_{i+1}}, \quad (10)$$

where each $r_{x_i x_{i+1}} = m_{x_i} / m_{x_{i+1}}$ is obtained from a pairwise Google Trends request $\{x_i, x_{i+1}\}$. The reason for multiplying ratios $r_{x_i x_{i+1}}$ of neighboring queries, rather than directly comparing non-neighboring queries x_i, x_j (with $j > i + 1$) in a separate Google Trends request, is that, if non-neighboring queries could be compared directly without increasing the overall rounding error, then the intermediate queries x_{i+1}, \dots, x_{j-1} would be unnecessary, and A^* would not be optimal—a contradiction. The reason for including only $k = 2$, rather than the maximum allowed $k = 5$, queries per request is that $k = 2$ keeps rounding errors to a minimum, as only one of the maximum values $m_{x_i}, m_{x_{i+1}}$ will be rounded after scaling.

Next, we analyze the bound ratio $\eta_{x_1 x_n}$ that arises when constructing an anchor bank in the above-described way. Each pairwise Google Trends request $\{x_i, x_{i+1}\}$ returns $m_{x_{i+1}} = 100$ and $m_{x_i} = 100c_i < 100$, for $c_i = r_{x_i x_{i+1}} \in [0, 1]$. Recall from (4) that m_{x_i} is the rounded version of the true, unobserved $m_{x_i}^* \in [m_{x_i}^-, m_{x_i}^+]$ (the maximum search interest of x_i after scaling to $[0, 100]$, but before integer rounding), with $m_{x_i}^- = m_{x_i} - \frac{1}{2}$, and $m_{x_i}^+ = m_{x_i} + \frac{1}{2}$. Since $m_{x_{i+1}} = 100$, we have $m_{x_{i+1}}^- = m_{x_{i+1}}^+ = 100$ (only scaling, no rounding), so the true maximum ratio $r_{x_i x_{i+1}}^* \in [m_{x_i}^- / 100, m_{x_i}^+ / 100]$, and the ratio of the upper and lower bounds is

$$\eta_{x_i x_{i+1}} = \frac{m_{x_i}^+}{m_{x_i}^-} = \frac{100c_i + \frac{1}{2}}{100c_i - \frac{1}{2}} = \frac{c_i + \epsilon}{c_i - \epsilon}, \quad (11)$$

where $\epsilon = 1/200$. When multiplying local maximum ratios (cf. (10)), rounding errors accumulate, yielding the global bound ratio

$$\eta_{x_1 x_n} = \prod_{i=1}^{n-1} \frac{c_i + \epsilon}{c_i - \epsilon}. \quad (12)$$

In fact, in an optimal anchor bank, anchor queries are equidistant, in the sense that $c_i = c$ for a global constant c . To see this, consider the case $n = 3$ (the case $n > 3$ follows by induction):

$$\eta_{x_1 x_3} = \left(\frac{c_1 + \epsilon}{c_1 - \epsilon} \right) \left(\frac{c_2 + \epsilon}{c_2 - \epsilon} \right), \quad (13)$$

which is minimized for $c_1 = c_2$: The numerator,

$$(c_1 + \epsilon)(c_2 + \epsilon) = c_1 c_2 + \epsilon(c_1 + c_2) + \epsilon^2, \quad (14)$$

is minimized for $c_1 = c_2$ because $c_1 c_2 = r^*$ is fixed and the circumference $2(c_1 + c_2)$ of a rectangle of fixed area $c_1 c_2$ is minimized by a square. Analogously, the denominator is maximized for $c_1 = c_2$, so $\eta_{x_1 x_3}$ as a whole is minimized for $c_1 = c_2 =: c$.

The product of (12) thus simplifies to

$$\eta_{x_1 x_n} = \left(\frac{c + \epsilon}{c - \epsilon} \right)^{n-1} = \left(\frac{c + \epsilon}{c - \epsilon} \right)^{\log_c r^*}, \quad (15)$$

and finding an optimal anchor bank boils down to finding the optimal value of the constant c . Writing the bound ratio $\eta_{x_1 x_n} = \eta(c)$ as a function of c ,

$$\eta(c) = \left(\frac{c + \epsilon}{c - \epsilon} \right)^{\log_c r^*}, \quad (16)$$

our goal is to minimize $\eta(c)$ subject to the constraint that $r^* = c^{n-1}$ for some integer n , i.e., that c should be a root of r^* .

Since the maximum ratio r^* of the most and the least frequent queries is fixed, a smaller $c < 1$ will lead to an anchor bank of a smaller size n . This leads to a tradeoff: for smaller c (i.e., smaller n), we will on the one hand require fewer factors for computing $r_{x_1 x_n}$ (cf. (10)), but on the other hand each factor will be subject to larger rounding errors (cf. (15)). Finding the optimal constant c is to find the optimal tradeoff between the number of comparisons and the precision of comparisons.

To make the problem easier (and because r^* , though fixed, is unobserved), we consider the continuous relaxation without the constraint $r^* = c^{n-1}$. (We shall see later, in Fig. 4, that adding the constraint does not change the optimal c by much.) This way, we can find the optimal c by setting the derivative $\eta'(c)$ to zero:

$$\eta'(c) = \eta(c) (\log r^*) \left(\frac{\frac{1}{c+\epsilon} - \frac{1}{c-\epsilon}}{\log c} - \frac{\log \frac{c+\epsilon}{c-\epsilon}}{c(\log c)^2} \right) = 0. \quad (17)$$

Since $\eta(c) (\log r^*) < 0$, we have $\eta'(c) = 0$ iff

$$\frac{\frac{1}{c+\epsilon} - \frac{1}{c-\epsilon}}{\log c} = \frac{\log \frac{c+\epsilon}{c-\epsilon}}{c(\log c)^2} \quad (18)$$

$$(\log c) \left(\frac{c}{c+\epsilon} - \frac{c}{c-\epsilon} \right) = \log \frac{c+\epsilon}{c-\epsilon} \quad (19)$$

$$(\log c) \left(\frac{c}{c+\epsilon} - \frac{c}{c-\epsilon} \right) = \log \left(\frac{c}{c-\epsilon} + \frac{\epsilon}{c-\epsilon} \right) \quad (20)$$

$$\log c = \frac{\log \left(\frac{1}{1-\epsilon/c} + \frac{\epsilon/c}{1-\epsilon/c} \right)}{\frac{1}{1+\epsilon/c} - \frac{1}{1-\epsilon/c}}. \quad (21)$$




Figure 4: Bound ratio $\eta(c)$ achieved by an anchor bank where neighboring queries x_i, x_{i+1} have a constant maximum ratio $m_{x_i}/m_{x_{i+1}} = c$, for various choices of r^* (the maximum ratio of the least and most frequent queries in the anchor bank). **Black dots: roots of r^* , i.e., values of c that meet the constraint $r^* = c^{n-1}$ for some integer n . **Red:** value $c = 1/e \approx 0.37$ for which $\eta(c)$ is approximately minimized.**

First-order Taylor expansions around $\epsilon/c = 0$ yield the following approximations:

$$\frac{1}{1 + \epsilon/c} \approx 1 - \epsilon/c \quad (22)$$

$$\frac{1}{1 - \epsilon/c} \approx 1 + \epsilon/c \quad (23)$$

$$\frac{\epsilon/c}{1 - \epsilon/c} \approx \epsilon/c, \quad (24)$$

so (21) can be approximated as

$$\log c \approx \frac{\log(1 + 2\epsilon/c)}{-2\epsilon/c}. \quad (25)$$

Another first-order Taylor expansion, around $2\epsilon/c = 0$, yields

$$\log(1 + 2\epsilon/c) \approx 2\epsilon/c, \quad (26)$$

so (25) simplifies to

$$\log c \approx -1 \quad (27)$$

$$c \approx 1/e \approx 0.37. \quad (28)$$

That is, in the ideal case, we would want to construct an anchor bank in which all neighboring queries x_i, x_{i+1} have the constant maximum ratio $m_{x_i}/m_{x_{i+1}} \approx 1/e$. In Fig. 4, we plot $\eta(c)$ for choices of r^* spanning all relevant orders of magnitude. The vertical red lines confirm visually that $\eta(c)$ is minimized for $c \approx 1/e$, regardless of r^* .




Figure 5: Theoretically optimal bound ratio $\bar{\eta}(r^*)$ (black) and empirical bound ratios (yellow, blue). **Yellow: empirical bound ratios $\eta_{x_i x_n}$ obtained for anchor queries x_i when using the optimized anchor bank. **Blue:** *Idem*, but when using the initial anchor bank before optimization (showing only the subset of anchor queries retained for the optimized anchor bank). As we see, the optimized anchor bank comes much closer to the theoretical optimum (black). Since the exact maximum ratios $r_{x_i x_n}^* = m_{x_i}^*/m_{x_n}^*$ are unobserved, we plot, on the horizontal axis, the intervals $[m_{x_i}^-/m_{x_n}^+, m_{x_i}^+/m_{x_n}^-]$ instead (cf. (5)).**

We now fix $c = 1/e$ and turn the bound ratio $\eta_{x_i x_n}$ into a function $\bar{\eta}(r^*)$ of r^* (rather than, as before, of c):

$$\bar{\eta}(r^*) = \left(\frac{\frac{1}{e} + \epsilon}{\frac{1}{e} - \epsilon} \right)^{-\log r^*} \approx 1.028^{-\log r^*} \approx (r^*)^{-0.027}. \quad (29)$$

In practice, $r^* > 10^{-7}$, so $\bar{\eta}(r^*) < 1.55$. That is, in the worst case, the upper bound on an optimally estimated maximum ratio is at most 55% higher than the lower bound. We plot $\bar{\eta}(r^*)$ in Fig. 5 as the black curve.

B CONSTRUCTING AND USING AN OPTIMAL ANCHOR BANK

The considerations of Appendix A where theoretical. Next, we describe how to implement them in order to construct and deploy an approximately optimal anchor bank in practice.

Offline phase: Building the anchor bank. Starting from an initial anchor bank A constructed as described in Sec. 2.1, we can build an approximately optimal anchor bank A^* (as characterized in Appendix A) as follows:

- (1) Select a subset of the anchor queries of the initial anchor bank A such that, in the subset, subsequent queries x_i, x_{i+1}

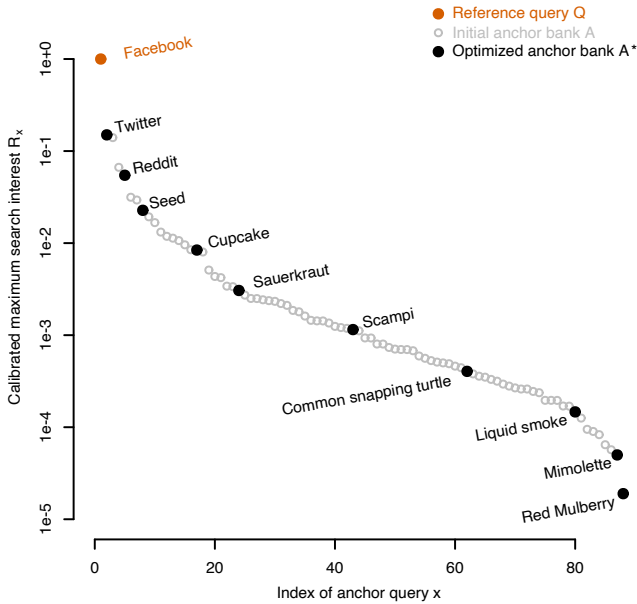


Figure 6: Anchor bank before and after optimization. *Gray circles: Initial anchor bank A constructed as described in Sec. 2.1. Black disks: Optimized anchor bank A^* . Note that the optimized anchor bank divides the (logarithmic) y -axis into segments of approximately equal length, corresponding to the optimal maximum ratio of $1/e$ for neighboring queries (cf. Appendix A).*

have an estimated maximum ratio $r_{x_i x_{i+1}}$ of about $1/e$ (cf. Appendix A).

- (2) Obtain more precise estimates of the true maximum ratios $r_{x_i x_{i+1}}^*$ of subsequent queries in a second round of Google Trends requests, where we include only $k = 2$ (rather than $k = 5$) subsequent queries $\{x_i, x_{i+1}\}$ per request.

To implement step 1, we first construct a complete, directed graph G' with edge weights $w_{xy} := |\log(e^{-1}/r_{xy})| = |1 + \log r_{xy}|$, where r_{xy} are the maximum-ratio estimates computed in step 3 of offline preprocessing (Sec. 2.1), and then find a shortest path from

the least frequent query x_1 to the most frequent query x_n in G' . The nodes along the shortest paths then serve as the anchor queries of the optimized anchor bank A^* .

In Fig. 6, we visualize the initial anchor bank A constructed as described in Sec. 2.1, as well as the optimized anchor bank A^* constructed as described above. Note that the optimized anchor bank divides the (logarithmic) y -axis into segments of approximately equal length, corresponding to the optimal maximum ratio of $1/e$ for neighboring queries, as derived in Appendix A.

We visually demonstrate the improvements afforded by optimizing the anchor bank in Fig. 5. Here we calibrated the anchor queries x_i retained for the optimized anchor bank A^* against the reference query $Q = x_n = \text{FACEBOOK}$ in two ways: first, using the initial anchor bank A , and second, using the optimized anchor bank A^* . (Recall that the anchor queries of A^* form a subset of those of A .) In blue, we plot, on the vertical axis, the bound ratios $\eta_{x_i x_n}$ (cf. (6)) obtained for the initial anchor bank A . In yellow, we plot the corresponding bound ratios obtained for the optimized anchor bank A^* . We clearly see that the optimized anchor bank (yellow) comes much closer to the theoretically optimal bound ratio $\bar{\eta}(r_{x_i x_n}^*)$ (black) than the initial anchor bank (blue) does.

Online phase: Binary search. At deployment time, during online binary search (Sec. 2.2), we start by comparing the input query q to the query x' from the optimized anchor bank A^* that is closest (in terms of calibrated maximum search interest $R_{x'}$) to the median query of the initial anchor bank A , based on the following rationale: the search-interest distribution in the initial anchor bank approximates the overall search-interest distribution, so by first comparing to a query that is close to the overall median, we will reduce the number of search steps, which keeps the number of Google Trends requests as well as rounding errors low.

As mentioned at the end of Sec. 2.2, in addition to the calibrated search-interest time series, we also return upper and lower bounds of the time series based on the largest possible (unknown) rounding errors that may have been encountered. In order to keep the bounds as tight as possible, we also use the close-to-median query x' (as defined above) as the reference query Q in practice (whereas, for ease of exposition, we have used the most frequent query x_n as the reference query throughout this paper).