



Generation of Synthetic Query Auto Completion Logs

Unni Krishnan¹ , Alistair Moffat¹  , Justin Zobel¹ ,
and Bodo Billerbeck² 

¹ The University of Melbourne, Melbourne, Australia
ammoffat@unimelb.edu.au

² Microsoft Australia and RMIT University, Melbourne, Australia

Abstract. Privacy concerns can prohibit research access to large-scale commercial query logs. Here we focus on generation of a synthetic log from a publicly available dataset, suitable for evaluation of query auto completion (QAC) systems. The synthetic log contains plausible string sequences reflecting how users enter their queries in a QAC interface. Properties that would influence experimental outcomes are compared between a synthetic log and a real QAC log through a set of side-by-side experiments, and confirm the applicability of the generated log for benchmarking the performance of QAC methods.

1 Introduction

Query auto completion (QAC) systems offer a list of completions while users enter queries in a search interface. Users can either submit one of the completions as their *query*, or *advance* their *partial query* by selecting a completion and then continuing to type [33]. A detailed QAC log capturing the sequence of partial queries, along with the completions presented and the user interactions with them, is required in order to evaluate a QAC system [37, 38]. However, concerns about the privacy of query logs and regulatory requirements such as GDPR mean that there is a need for alternative ways of obtaining logs for academic purposes.

Here we explore a framework for generating synthetic QAC logs, extending the work of Krishnan et al. [33], who suggest converting a QAC log to an abstracted format (an *abstract QAC log*) that records only the length of each partial query and the lengths of words used, minimizing privacy concerns but removing the possibility of performing evaluations on actual strings. Synthetic QAC log generation seeks to produce a list of *plausible* synthetic partial query sequences by mapping the word lengths from the abstract QAC log to strings from a publicly available dataset. An example of the proposed process is shown in Fig. 1. On the left are partial queries typed by a user. The abstracting process converts these to the digit sequences shown in the middle column, describing the strings but not their characters; and then the corresponding synthesized strings are shown at the right. Note that it is neither necessary nor sufficient for the synthetic log to contain semantically valid phrases. Comparison between the




Fig. 1. A sample QAC log entry, the corresponding abstract QAC log entry, and a synthetic QAC log entry generated from the string “john of gaunt”.

original and synthetic logs across a range of properties show that the synthetic log can be used to evaluate QAC system performance. Moreover, the synthetic log eliminates the privacy concerns associated with the original.

2 Background

A QAC system retrieves a *candidate set* matching the partial query P , drawing from a target string collection, with strings in the target collection having an associated *score*. Query Auto Completion systems typically match P against past queries from a log; or, in the absence of logs, they can also be synthesized [14, 40]. Methods of ranking the candidates include static popularity [9], search context [9, 29], forecast popularity [15], personalized ranking parameters [15, 29, 42], and diversity [16]. It is also possible to choose an initial candidate set based on popularity and then apply a second ranking criteria to obtain the final strings [15, 16]. QAC implementation strategies vary based on how the partial query P is matched against the target strings [32]. A common approach is to use a trie [3, 4, 25, 27] to retrieve candidates that have P as a prefix; or inverted index-based approaches [10, 11, 23] that offer completions independent of the ordering of the words in the partial query. The functionality of a QAC system can be extended beyond character level matches by including contextual cues [11] or synonyms [12, 28]. Error-tolerant QAC approaches [17, 28, 36, 47] allow up to a fixed number of character mismatches to account for possible typing errors.

User interactions are a key factor in implementation and evaluation of QAC systems [26, 33, 38] and have been captured using a wide range of models [31–33, 37, 38, 43, 44]. In particular, users are not limited to entering single characters, and can alter the partial query by selecting a completion or deleting characters already entered. Until now, the *test collections* used to evaluate traditional search systems have been anonymized commercial logs [1, 19, 34, 48] or synthesized logs [5, 30, 45, 49]. QAC system evaluations have typically been performed over large publicly available string collections [10, 14, 23], with strings taken sequentially from left-to-right to generate partial query sequences [10]. However this approach does not account for the full range of possible interactions [32, 33]. In this work, we explore an approach to generation of synthetic partial query sequences that addresses this gap.



Fig. 2. The synthetic pattern generation process. Signatures of FinalP from the abstract QAC log are used to form the seedsig list, and the signatures from surrogate log form targsig list. A match between the signature (6, 3) in the seedsig list and in the targsig list might correspond to the target string “coffee mug”.

Terminology. A partial query P is the string currently displayed in the search box. An *interaction* updates P and results in loggable changes. A new *conversation* starts when the user begins a query, and continues until either explicitly terminated by the user or as a result of a session timeout. The last partial query from a conversation is referred to as FinalP. A QAC log records a set of conversations as a sequence of partial queries. A *surrogate log* is a dictionary of strings that can be used as substitutes for final partial queries, with each such string having an associated score reflecting its popularity. For a string T, the ordered tuple $\langle |w_1|, |w_2| \dots |w_k| \rangle$ representing the lengths of its words $w_1, w_2, \dots w_k$ in T is referred to as its *signature*. For each partial query P, an abstract QAC log records only its signature and its length |P|, including whitespace. The signature of each FinalP in the abstract QAC log is referred to as seedsig. The signature of a string in the surrogate log is its targsig.

Problem Definition. For each conversation in the abstract QAC log, find a *target string* in the surrogate log with the same signature as the final partial query FinalP. Then, starting from the first interaction in the conversation, apply the word lengths from the abstract QAC log to the target string in order to obtain a plausible partial query sequence. For example, consider the last conversation in Fig. 2, with its signature sequence $\langle 1 \rangle, \langle 2 \rangle, \langle 6 \rangle, \langle 6, 3 \rangle$. The seedsig for this conversation is (6, 3). The string “coffee mug” in the surrogate log has the same signature and hence might be selected as a target string. Mapping the signature sequence from the conversation, we get the synthetic partial query sequence “c”, “co”, “cof”, “coffee”, and then “coffee mug”.

3 Generation Process

Depending on the distribution of word lengths in the surrogate log, for every final partial query in the abstract QAC log, there may not be a string having the same signature. Moreover, strings are not entered by the users in the word order of the target collection. For instance, a user looking for the Wikipedia

main page might enter the queries “wikipedia”, “main page wikipedia”, or “wiki”. We narrow down the possible ways of matching a *seedsig* with the list of signatures in *targsig* list to the following hierarchical *modes*:

1. *Exact*. The *targsig* is equal to the *seedsig*. For example, the *seedsig* $\langle 3, 3, 2 \rangle$ only matches with the target signature $\langle 3, 3, 2 \rangle$. There might be zero, or multiple strings in the surrogate log that match.
2. *Prefix*. The *seedsig* is a prefix of the *targsig*. For example, *seedsig* $\langle 3, 3, 2 \rangle$ matches $\langle 3, 3, 2, 4 \rangle$ and $\langle 3, 3, 2, 4, 7 \rangle$, but not $\langle 9, 3, 3, 2 \rangle$.
3. *Match by Drop* (MbD). The *seedsig* is an ordered subset of the *targsig*, For example, $\langle 3, 3, 2 \rangle$ matches $\langle 3, 4, 3, 2 \rangle$ and $\langle 4, 3, 1, 3, 6, 2 \rangle$, but not $\langle 3, 4, 2, 3 \rangle$.
4. *Bag of Numbers* (BoN). Relaxing the ordering requirement, a BoN match occurs if the sets of word lengths in *seedsig* are a subset of those in *targsig*. For example, *seedsig* $\langle 3, 3, 2 \rangle$ matches $\langle 2, 3, 3, 4 \rangle$ and $\langle 5, 2, 3, 6, 3 \rangle$, but not $\langle 3, 4, 2, 6 \rangle$. In a BoN match the target string’s words are reordered to match the seed signature ordering.

Locating Target Strings. The first step locates, for a given *seedsig*, a set of matching signatures in the *targsig* list. The set of target candidates is maintained in lexicographically sorted order, so that the exact-match signatures for a given *seedsig* can be found via two binary searches, establishing a range $[r_{beg}, r_{end}]$. A similar process can be used to find the prefix-match range, which is a larger contiguous block in the lexicographically sorted array of target candidates.

If a target signature S_t matches with *seedsig* using MbD, any *targsig* having S_t as a prefix will also have an MbD match with *seedsig*. For example, under MbD the *seedsig* $\langle 3, 5 \rangle$ matches $\langle 3, 2, 5 \rangle$. Then signatures $\langle 3, 2, 5, 4 \rangle$ and $\langle 3, 2, 5, 9, 16 \rangle$ will also be a match because they have $\langle 3, 2, 5 \rangle$ as a prefix. Using this property, once a matching signature S_t for MbD is found, we can add the prefix match range for S_t to the MbD range for the current *seedsig*. In contrast to exact and prefix match ranges, MbD ranges may not form a continuous range over SigList.

For BoN matching, the tokens in the signature are sorted to form a canonical representation. For example, $\langle 3, 2, 4 \rangle$ becomes $\langle 2, 3, 4 \rangle$. This list is sorted into lexicographical order. A BoN match between the canonical representations of a *seedsig* and a target signature S_{tc} can be verified by a linear scan over the tokens. The matching ranges for BoN in the modified SigList are calculated by finding prefix match ranges for each S_{tc} that matches with *seedsig* using BoN.

Handling Deletions. To include deletions, we assume that something different was initially typed (the *replacement* word) and was converted to the word from the target string after the deletions. In particular, suppose that the current conversation contains the deletion of a word w_k starting from the i th interaction and ending with the j th interaction. Then, for every interaction preceding the j th interaction, w_k is substituted by a *replacement* word constructed via a set of deletion heuristics. If character sequences are deleted and re-entered several times during the conversation, several replacement words will be required. The last replacement word should be close to the original word, and the penultimate

Table 1. Deriving synthetic partial query sequence for a conversation containing multiple word deletions. The fourth column shows three target strings used.

Sequence	P	Signature	Modified target string	Synthetic P
1	12	$\langle 5, 6 \rangle$	“black mirror”	“black mirror”
2	11	$\langle 5, 5 \rangle$	“black mirror”	“black mirro”
3	5	$\langle 5 \rangle$	“black mirror”	“black”
4	6	$\langle 5 \rangle$	“black hole”	“black”
5	7	$\langle 5, 1 \rangle$	“black hole”	“black h”
6	8	$\langle 5, 2 \rangle$	“black hole”	“black ho”
7	7	$\langle 5, 1 \rangle$	“black hole”	“black h”
8	8	$\langle 5, 2 \rangle$	“black hawk down”	“black ha”
9	15	$\langle 5, 4, 4 \rangle$	“black hawk down”	“black hawk down”

replacement prior to it should similarly be a modification of the last replacement word. Table 1 gives an example of the partial query sequence generated by re-writing target strings when the conversation contains multiple deletions.

Users delete characters either due to typing errors [6, 20] made within that word, or to switch to an entirely different word. Among various typing errors that can occur in any character level entry systems [22], the following error categories that are predominantly discussed in past research [6, 20, 21, 46] are included in the generation process.

Deletion. A deletion error occurs when the user initially missed out one of the characters in the word, for example, in the correction “acount” → “account”. Deletion errors are more frequent at character repetitions, and occur more commonly at the beginning of a word [6]. To simulate a deletion error, if there is a deletion sequence from i th interaction to the j th interaction in which more than two characters of w_k are removed, then w_k from the $j + 1$ th interaction (after the deletions) is examined to see if the last two characters in w_k are repeating. The error is then simulated by deleting one of the repeating characters of w_k in the interactions prior to $j + 1$.

Insertion. These arise when an extra letter is initially typed, for example, the correction “sherriff” → “sheriff”. Our experimental results show that only a low fraction of insertion errors are observed in a QAC log and for that reason, they are not included in the synthetic QAC log generation process.

Substitution. These arise when the user enters one of the neighboring keys instead of the intended key, for example, the correction “disturv” → “disturb”. Single character deletions are assumed to be substitution errors. The substituted character is found via a probability distribution for mistyped keys around the current key.

Transposition. This is when the user swaps two characters in a word either not knowing the correct spelling or because of mis-ordered keystrokes, for

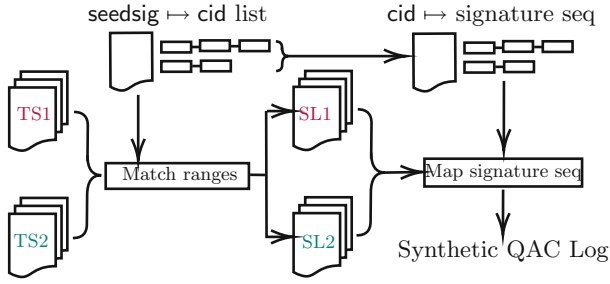


Fig. 3. Framework for generating a synthetic QAC log. See the text for details.

example, the correction “wierd” → “weird”. Deletions of length two are considered to be transposition errors and the last two characters of the word are swapped to get the replacement word.

In a QAC interface, typing errors may not be the only reason why users delete characters. Sometimes users replace certain parts of a well-formed partial query, for instance to get a different set of completions. In such cases, referred as *multichar* deletions, the corrected word will differ from the original word in two or more character positions. To allow for such cases, the deleted part of the word is replaced with another string starting with the word’s remaining prefix, but differing in the next character. For example, consider the word “hawk” in Table 1. A deletion chain removed all of the characters except the “h”; and hence (working backwards), “hawk” was replaced by an alternate word that starts with “h” but not “ha”, such as “hole”.

Generation Framework. The overall framework for generation of a synthetic QAC log is illustrated in Fig. 3. The seed signatures are stored with a mapping from signatures to the list of cids having the same signature. Signatures from the surrogate log are precomputed and stored in lexicographically sorted order (TS1). Additionally, the strings in the surrogate log are re-ordered based on TS1 ordering to obtain the permutation SL1. The permutations TS2 stores a sorted list of within sorted signatures to support BoN matching, and SL2 is the corresponding permutation of the surrogate log. To generate a synthetic QAC log entry, a signature is selected from *seedsig* → *cid list* and a list of target signatures retrieved based on the four modes. Using the mappings SL1 and SL2, strings from the surrogate log corresponding to the target signatures are then obtained.

The strings are converted to target strings by aligning the word lengths with the *seedsig*. A target string can be generated from multiple originals. If, say, *seedsig* is (9,4), then “wikipedia main page” gives the target string “wikipedia page” using MbD; and another string “personal web page wikipedia” gives the same target using BoN mode. If the same target string is produced by more than one string from the surrogate log, then the one with lower score is discarded.

The signature sequences obtained from the file providing the $\text{cid} \mapsto$ signature sequence mapping are then applied to the list of target strings, to obtain a synthetic partial query sequence for each conversation. Partial query sequences are generated for each cid from the current $\text{seedsig} \mapsto \text{cid}$ list until either the cid list or the target strings are exhausted. Finally, after the generation process is completed, the synthetic QAC log is re-ordered so that the conversations in the synthetic QAC log have one-to-one correspondence with the abstract QAC log.

4 Experiments

Datasets Used. The Wikipedia clickstream dataset,¹ generated from Wikipedia request logs containing tuples of the form (*referrer*, *resource*, *frequency*), is used to generate the surrogate log. The number of requests for an article (*resource*) is the *frequency*. Data dumps from January to March 2019 were aggregated by updating the *frequency* of each *resource* by the mean *frequency* over the dumps. We refer to the resulting dataset as **Wiki-Clickstream** and the synthetic QAC log generated from **Wiki-Clickstream** as **Wiki-Synth**.

A QAC log was formed by randomly sampling the logs from Bing² QAC system over a period of one week from 13 August 2018. This **Bing-QAC-2018** records the partial queries entered by the users and a unique cid for each conversation. The abstracted version of **Bing-QAC-2018**, referred to as **Bing-Abs-QAC-2018**, is generated by recording the signature and total length of each partial query along with a unique cid for each conversation. Note that **Wiki-Synth** was generated from **Bing-Abs-QAC-2018** only (with no use of **Bing-QAC-2018**), and that the subsequent comparisons between **Wiki-Synth** and **Bing-QAC-2018** were performed on secure Microsoft servers and in accordance with Microsoft privacy requirements.

Preprocessing. Conversations in **Bing-Abs-QAC-2018** where an intermediate signature contained more words than the signature of the final partial query, or contained words that were longer than the corresponding words in the final partial query, were removed. The strings from **Bing-QAC-2018** were *transliterated* from UTF-8 to ASCII encoding using **Unidecode**³ Python package. Conversations containing non-converting characters in any partial query were discarded. The strings from **Wiki-Clickstream** were transliterated in the same way and non-converting strings were removed. The conversion from UTF-8 to ASCII resulted in the loss of 0.26 million conversations from **Bing-QAC-2018** and 0.02 million strings from **Wiki-Clickstream**. After the preprocessing, there were 1.44 million conversations in **Bing-Abs-QAC-2018** and 5.11 million strings in **Wiki-Clickstream**. In **Bing-Abs-QAC-2018**, 7.53% of signatures were unique, of which 2.26% had no matches in **Wiki-Clickstream**. A total of 0.23 million conversations from

¹ https://meta.wikimedia.org/wiki/Research:Wikipedia_clickstream, accessed 29th October, 2019.

² <https://www.bing.com>, accessed 29th October, 2019.

³ <https://pypi.org/project/Unidecode>, accessed 29th October, 2019.

Bing-Abs-QAC-2018 were not included in Wiki-Synth because there were not enough matching strings in Wiki-Clickstream to map their signature sequence.

Sampling of Target Strings. The frequency of search queries has been found to exhibit a power law distribution [8] with the probability distribution $p(x) \propto x^{-\alpha}$. Therefore, while generating a synthlog, the target strings need to be sampled to obtain a power law distribution based on their frequencies. Strings in the Wiki-Clickstream log have a score that tends to follow a power law distribution with $\alpha = 3.09$. Using this observation we sample target strings based on a weighted probability distribution over their scores, so that the probability of a string T given a seed signature `seedsig` is

$$\text{Prob}(T \mid \text{seedsig}) = \frac{\text{score}(T)}{\sum_{i=1}^n \text{score}(T_i)} \quad (1)$$

where T_1, T_2, \dots, T_n are the set of target strings retrieved for `seedsig`. The frequency distribution of the resulting target strings used as the FinalP in Wiki-Synth is analyzed below, where we compare partial query frequencies.

Language Model for Finding Replacement Words. For deletion types other than multichar, the replacement word will be a modification of the original word based on the error category. For multichar deletions, the heuristics discussed above are used to find replacement words. The *best* replacement word based on contextual information among the candidate words that passes the heuristics is selected using a 4-gram language model (LM) trained on the title strings from Wiki-Clickstream. The language model is generated using KenLM [24],⁴ which is based on modified Kneser-Ney smoothing and provides fast model construction and querying. For example, an LM-based replacement for the target string “live queen”, is “live together” while a random replacement gives “live teufelshorner”; for the target string “web server”, a random replacement yields “web castelvetere” and an LM-based replacement gives “web content”. A similar character bigram model trained on the Microsoft spelling-correction dataset⁵ is used to find the most likely next mistyped character, given the previous character, to simulate substitution errors.

Comparison of Synthetic QAC Logs with QAC Logs. A synthetic QAC log should have some properties similar to the QAC log, and these can be used to validate the generation process. A comparison of these properties is given in Table 2. Other properties are desirable if a synthetic QAC log is to be a substitute in experiments for a QAC log. While comparing these properties, the partial queries from both the logs are treated as the *test queries* that will be queried against a collection of strings acting as the *test collection*. Extending our assumption that a user’s goal was to enter FinalP, we claim that these strings could get

⁴ <https://github.com/kpu/kenlm>, accessed 29th October, 2019.

⁵ <https://www.microsoft.com/en-us/download/details.aspx?id=52418>, accessed 29th October, 2019.

Table 2. Basic statistics of Bing-QAC-2018 and Wiki-Synth.

Attribute	Bing-QAC-2018	Wiki-Synth
Number of conversations (millions)	1.74	1.21
Number of interactions (millions)	11.75	7.90
Percentage of unique partial queries	38.99	39.41
Percentage of unique final partial queries	52.30	63.32
Lengths of partial queries – mean (characters)	10.37	9.67
Lengths of partial queries – std.dev. (characters)	9.68	8.53

Table 3. Goodness of power law fit against other distributions.

Frequency distribution	Exponential		Log-normal	
	R	p	R	p
Bing-QAC-2018, P	11885.4	<0.01	0.04	0.8
Wiki-Synth, P	10895.3	<0.01	0.16	0.15
Bing-QAC-2018, FinalP	5903.6	<0.01	-2.42	0.04
Wiki-Synth, FinalP	1211.9	0.01	0.12	0.04

eventually indexed by the QAC system and subsequently be queried. Thus, in our experiments, the list of FinalP strings are considered as a representative sample of a larger hypothetical test collection.

Heap’s Law. Heap’s law gives an estimate of the number of unique terms V in a collection as a function of the number of terms N . The relationship is given by $V = kN^\beta$ and the typical values of the parameters are $30 \leq k \leq 100$, with $\beta \approx 0.50$ for English text [39] and $\beta \approx 0.60$ for web documents [7]. The values for Heap’s law parameters estimated from Bing-QAC-2018 are $k = 14.50$, $\beta = 0.67$ and the parameters from Wiki-Synth are $k = 32.61$, $\beta = 0.57$. The difference in growth rate can be explained by the nature of the logs. The growth rate of V is expected to be higher in a collection containing numbers and spelling errors [39]. Each of the FinalP strings from Wiki-Synth is a Wikipedia title that comes from a curated collection whereas FinalP strings in Bing-QAC-2018 are strings entered by users. Considering, all partial queries, the parameters estimated are $k = 5.54$, $\beta = 0.71$ for Bing-QAC-2018 and $k = 20.94$, $\beta = 0.60$ for Wiki-Synth.

Frequency of Partial Query Strings. The likelihood of a query being repeated over time as a result of the power-law distribution affects the performance of various query processing strategies [13,35]. The frequency distributions of partial queries and final queries from Wiki-Synth and Bing-QAC-2018 are given in Fig. 4 (left) with dashed lines showing the power-law fits for the corresponding distributions. The distributions give similar exponents except for the distribution of FinalP from Wiki-Synth, which has $\alpha = 3.02$, indicating a steeper decay in the

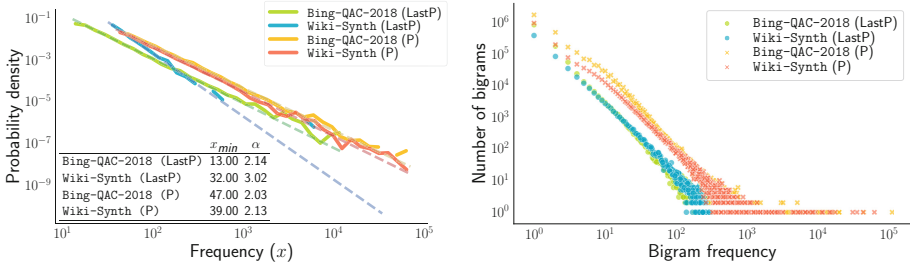


Fig. 4. Power law fit (left) and bigram frequency distribution (right).

probability of higher frequency partial query strings. The goodness of power-law fit against exponential and log-normal fits is compared using likelihood ratio R [2, 18] along with the significance level p , as reported in Table 3. Except for the frequency distribution of FinalP from Bing-QAC-2018, which is better explained with a log-normal distribution ($R = -2.42, 0.04$), the other three distributions suit a power-law distribution, perhaps as a result of the sampling process followed. For some datasets, it is not surprising to get a better log-normal fit [18].

N-Gram Frequency. The distribution of terms in a collection can be modeled using a decay law which estimates the collection frequency of the i th most common term as $F_i = Ci^{-k}$, where $k = -1$ and C is constant. Similarity in term distribution with a real log is considered to be a desirable property of a synthetic log [45]. Figure 4 (right) shows the bigram frequency distribution from partial queries and final partial queries. We find close correspondence between the bigram frequency distributions (Kolmogorov-Smirnov statistic, $D = 0.05, 0.63$ for FinalP and $D = 0.06, 0.04$ for all partial queries). Unigram frequencies follow a similar trend. While the bulk of the distribution can be explained with the decay law, the tail of the distributions coming from rare n -grams show deviations.

Empirical Entropy. The empirical entropy of the FinalP strings gives a measure of their compressibility. The k th-order empirical entropy of a string $T[1 \dots n]$ over an alphabet set Σ is given by

$$H_k(T) = \frac{1}{n} \sum_{w \in \Sigma^k} |T_w| \cdot H_0(T_w), \text{ with } H_0(T) = \frac{1}{n} \sum_{c \in \Sigma} n_c \cdot \log \frac{n}{n_c}, \quad (2)$$

where T_w is formed by collecting the characters that immediately follow *context* w in T , and where n_c is the frequency of character c in T . A lower bound on the number of bits required to encode T is given by $nH_k(T)$ [41]. For a list of strings T_1, T_2, \dots, T_l , it can be observed that $\sum_i n_i H_k(T_i) \leq n_c H_k(T_c)$ where T_c is the concatenation of strings T_i . This can be shown by extending the proof given by Navarro [41, Chapter 2], applying Jensen’s inequality to $\sum_i n_i H_k(T_i)$. Therefore, we consider $n_c H_k(T_c)$ as the worst case lower bound for the space required to encode the FinalP strings from Bing-QAC-2018 and Wiki-Synth. The values of $H_k(T_c)$ for $k = 0 \dots 4$ computed from Bing-QAC-2018 and Wiki-Synth are given

Table 4. Values of H_k and number of contexts C_t per length n ($\times 10^{-3}$) of the concatenated string formed by the FinalP strings of Bing-QAC-2018 and Wiki-Synth.

Dataset	$k = 0$		$k = 1$		$k = 2$		$k = 3$		$k = 4$	
	H_0	C_t/n	H_1	C_t/n	H_2	C_t/n	H_3	C_t/n	H_4	C_t/n
Bing-QAC-2018	4.39	0.00	3.64	0.00	2.50	0.17	1.41	2.51	0.64	11.27
Wiki-Synth	4.28	0.00	3.43	0.00	2.32	0.08	1.27	0.97	0.54	4.95

Table 5. Example for pre-correction and post-correction strings. In the table “pal” is taken as the pre-correction string and “pla” is taken as the post-correction string. This correction pair will be classified as a transposition error.

Inter.	Partial query	Word len.	Edit dist.
2	coffee pa	2	
3	coffee pal	3	
4	coffee pa	2	
5	coffee p	1	
6	cofee pl	2	1
6	coffee pla	3	1
6	cofee places	6	4

in Table 4. The values for H_k tend to be similar between the two logs. Slightly higher values of H_k and increased C_t/n for Bing-QAC-2018 can be explained by the higher degree of randomness expected from web search queries, compared to the results obtained for Heap’s law coefficients.

Identification of Typing Errors. Similarities in typing patterns between the synthetic QAC log and the original QAC log are verified by analyzing the typing errors present in both logs. Typing errors are identified by extracting and comparing a set of *pre-correction* and *post-correction* strings from the conversations, extending the method proposed by Baba and Suzuki [6]. When the k th word w_k in P gets deleted from the i th interaction, w_k from the $i - 1$ th interaction is taken as the pre-correction string. To find the post-correction string, the edit distance between the pre-correction string and the words w_k from the interactions after the end of current deletion sequence are calculated until $|w_k|$ from interaction j is less than that from interaction $j + 1$. From this list, the word having minimum edit distance with the pre-correction string is taken as the post correction string. Table 5 gives an example for how pre-correction and post-correction strings are calculated from a deletion sequence. The edit distance between two strings is calculated as a Damerau-Levenshtein distance, so transposition of two characters is given a cost of 1. The pre-correction and post-correction pairs are then compared to classify the possible typing errors as one of the types discussed above. The results obtained by analyzing the typing errors from Bing-QAC-2018

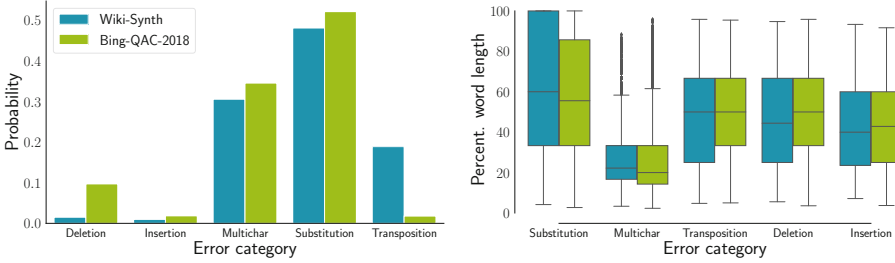


Fig. 5. Comparison between frequency (left) and position (right) of typing errors.

and the synthetic QAC log are in Fig. 5. The fraction of transposition errors was found to be higher in Wiki-Synth compared to Bing-QAC-2018 while the latter contains more deletion errors than Wiki-Synth. As a result of the method of simulating substitution errors, they tend to occur more towards the end of words in Wiki-Synth. The remainder of the error categories are distributed across similar word positions in both the logs, suggesting similarities in typing patterns.

5 Conclusions

We have explored a method for generating a synthetic QAC log from an abstract QAC log, by mapping the word lengths of the abstract QAC log to those of a publicly available string collection, and applying a range of corrective techniques. Synthetic QAC formation can also be posed as a language generation problem relying on various models of QAC systems [31, 33].

We have demonstrated that the synthetic log generated from a pre-existing string collection encompasses many of the properties found in the original QAC log from which the abstract QAC log was derived. In particular, an analysis of typing errors found in actual QAC logs is reported, along with a description of how they were introduced into the synthetic log. As a result, there is a close correspondence between the real QAC log and the synthetic QAC log across a range of properties, each of which might influence the computational cost of providing the completion strings. That is, experiments using the synthetic QAC log can be expected to provide close approximations to behaviors that would be observed on a real QAC log. A particular example is efficiency, which we plan to examine in future work, comparing the performance of a range of QAC implementations using both synthetic and real QAC logs.

Acknowledgments. This work was supported by the Microsoft Research Centre for Social Natural User Interfaces (SocialNUI) at The University of Melbourne. We are grateful to Peter Bailey (Microsoft Australia) for his support of that project, and for additionally facilitating the work that is reported here.

References

1. Adar, E.: User 4xxxxx9: anonymizing query logs. In: Proceedings of the WWW Query Log Analysis Workshop (2007). <http://www.cond.org/anonlogs.pdf>
2. Alstott, J., Bullmore, E., Plenz, D.: Powerlaw: a Python package for analysis of heavy-tailed distributions. *PLoS One* **9**(1), 1–11 (2014)
3. Askitis, N., Sinha, R.: HAT-trie: a cache-conscious trie-based data structure for strings. In: Proceedings of the Australasian Conference on Computer Science, pp. 97–105 (2007)
4. Askitis, N., Zobel, J.: Redesigning the string hash table, burst trie, and BST to exploit cache. *ACM J. Exp. Algorithmics* **15**, 1–7 (2010)
5. Azzopardi, L., de Rijke, M., Balog, K.: Building simulated queries for known-item topics: an analysis using six European languages. In: Proceedings of the SIGIR, pp. 455–462 (2007)
6. Baba, Y., Suzuki, H.: How are spelling errors generated and corrected? A study of corrected and uncorrected spelling errors using keystroke logs. In: Proceedings of the ACL, pp. 373–377 (2012)
7. Baeza-Yates, R., Saint-Jean, F.: A three level search engine index based in query log distribution. In: Nascimento, M.A., de Moura, E.S., Oliveira, A.L. (eds.) SPIRE 2003. LNCS, vol. 2857, pp. 56–65. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-39984-1_5
8. Baeza-Yates, R., Tiberi, A.: Extracting semantic relations from query logs. In: Proceedings of the KDD, pp. 76–85 (2007)
9. Bar-Yossef, Z., Kraus, N.: Context-sensitive query auto-completion. In: Proceedings of the WWW, pp. 107–116 (2011)
10. Bast, H., Weber, I.: Type less, find more: fast autocompletion search with a succinct index. In: Proceedings of the SIGIR, pp. 364–371 (2006)
11. Bast, H., Weber, I.: The CompleteSearch engine: interactive, efficient, and towards IR & DB integration. In: Proceedings of the CIDR, pp. 88–95 (2007)
12. Bast, H., Majumdar, D., Weber, I.: Efficient interactive query expansion with complete search. In: Proceedings of the CIKM, pp. 857–860 (2007)
13. Beitzel, S.M., Jensen, E.C., Chowdhury, A., Grossman, D., Frieder, O.: Hourly analysis of a very large topically categorized web query log. In: Proceedings of the SIGIR, pp. 321–328 (2004)
14. Bhatia, S., Majumdar, D., Mitra, P.: Query suggestions in the absence of query logs. In: Proceedings of the SIGIR, pp. 795–804 (2011)
15. Cai, F., Liang, S., de Rijke, M.: Time-sensitive personalized query auto-completion. In: Proceedings of the CIKM, pp. 1599–1608 (2014)
16. Cai, F., Reinanda, R., de Rijke, M.: Diversifying query auto-completion. *ACM Trans. Inf. Syst.* **34**(4), 25:1–25:33 (2016)
17. Chaudhuri, S., Kaushik, R.: Extending autocompletion to tolerate errors. In: Proceedings of the SIGMOD, pp. 707–718 (2009)
18. Clauset, A., Shalizi, C.R., Newman, M.E.J.: Power-law distributions in empirical data. *SIAM Rev.* **51**(4), 661–703 (2009)
19. Cooper, A.: A survey of query log privacy-enhancing techniques from a policy perspective. *ACM Trans. Web* **2**(4), 19:1–19:27 (2008)
20. Damerau, F.J.: A technique for computer detection and correction of spelling errors. *Commun. ACM* **7**(3), 171–176 (1964)
21. Dhakal, V., Feit, A.M., Kristensson, P.O., Oulasvirta, A.: Observations on typing from 136 million keystrokes. In: Proceedings of the CHI, pp. 646:1–646:12 (2018)

22. Gentner, D.R., Grudin, J.T., Larochele, S., Norman, D.A., Rumelhart, D.E.: A glossary of terms including a classification of typing errors. In: Cooper, W.E. (ed.) *Cognitive Aspects of Skilled Typewriting*, pp. 39–43. Springer, New York (1983). https://doi.org/10.1007/978-1-4612-5470-6_2
23. Hawking, D., Billerbeck, B.: Efficient in-memory, list-based text inversion. In: *Proceedings of the Australasian Document Computing Symposium*, pp. 5.1–5.8 (2017)
24. Heafield, K.: KenLM: faster and smaller language model queries. In: *Proceedings of the Workshop on Statistical Machine Translation*, pp. 187–197 (2011)
25. Heinz, S., Zobel, J., Williams, H.: Burst tries: a fast, efficient data structure for string keys. *ACM Trans. Inf. Syst.* **20**(2), 192–223 (2002)
26. Hofmann, K., Mitra, B., Radlinski, F., Shokouhi, M.: An eye-tracking study of user interactions with query auto completion. In: *Proceedings of the CIKM*, pp. 549–558 (2014)
27. Hsu, B.-J.P., Ottaviano, G.: Space-efficient data structures for top-*k* completion. In: *Proceedings of the WWW*, pp. 583–594 (2013)
28. Ji, S., Li, G., Li, C., Feng, J.: Efficient interactive fuzzy keyword search. In: *Proceedings of the WWW*, pp. 371–380 (2009)
29. Jiang, J., Ke, Y., Chien, P., Cheng, P.: Learning user reformulation behavior for query auto-completion. In: *Proceedings of the SIGIR*, pp. 445–454 (2014)
30. Jordan, C., Watters, C., Gao, Q.: Using controlled query generation to evaluate blind relevance feedback algorithms. In: *Proceedings of the JCDL*, pp. 286–295 (2006)
31. Kharitonov, E., Macdonald, C., Serdyukov, P., Ounis, I.: User model-based metrics for offline query suggestion evaluation. In: *Proceedings of the SIGIR*, pp. 633–642 (2013)
32. Krishnan, U., Moffat, A., Zobel, J.: A taxonomy of query auto completion modes. In: *Proceedings of the Australasian Document Computing Symposium*, pp. 6:1–6:8 (2017)
33. Krishnan, U., Billerbeck, B., Moffat, A., Zobel, J.: Abstraction of query auto completion logs for anonymity-preserving analysis. *Inf. Retrieval J.* **22**(5), 499–524 (2019). <https://doi.org/10.1007/s10791-019-09359-8>
34. Kumar, R., Novak, J., Pang, B., Tomkins, A.: On anonymizing query logs via token-based hashing. In: *Proceedings of the WWW*, pp. 629–638 (2007)
35. Lempel, R., Moran, S.: Predictive caching and prefetching of query results in search engines. In: *Proceedings of the WWW*, pp. 19–28 (2003)
36. Li, G., Ji, S., Li, C., Feng, J.: Efficient fuzzy full-text type-ahead search. *VLDB J.* **20**(4), 617–640 (2011). <https://doi.org/10.1007/s00778-011-0218-x>
37. Li, L., Deng, H., Dong, A., Chang, Y., Zha, H., Baeza-Yates, R.: Analyzing user’s sequential behavior in query auto-completion via Markov processes. In: *Proceedings of the SIGIR*, pp. 123–132 (2015)
38. Li, Y., Dong, A., Wang, H., Deng, H., Chang, Y., Zhai, C.: A two-dimensional click model for query auto-completion. In: *Proceedings of the SIGIR*, pp. 455–464 (2014)
39. Manning, C., Raghavan, P., Schütze, H.: *Introduction to Information Retrieval*. Cambridge University Press, Cambridge (2008)
40. Maxwell, D., Bailey, P., Hawking, D.: Large-scale generative query autocompletion. In: *Proceedings of the Australasian Document Computing Symposium*, pp. 9:1–9:8 (2017)
41. Navarro, G.: *Compact Data Structures: A Practical Approach*. Cambridge University Press, Cambridge (2016)

42. Shokouhi, M.: Learning to personalize query auto-completion. In: Proceedings of the SIGIR, pp. 103–112 (2013)
43. Smith, C.L., Gwizdka, J., Feild, H.: Exploring the use of query auto completion: search behavior and query entry profiles. In: Proceedings of the CHIIR, pp. 101–110 (2016)
44. Smith, C.L., Gwizdka, J., Feild, H.: The use of query auto-completion over the course of search sessions with multifaceted information needs. *Inf. Process. Manag.* **53**(5), 1139–1155 (2017)
45. Webber, W., Moffat, A.: In search of reliable retrieval experiments. In: Proceedings of the Australasian Document Computing Symposium, pp. 26–33 (2005)
46. Wobbrock, J.O., Myers, B.A.: Analyzing the input stream for character- level errors in unconstrained text entry evaluations. *ACM Trans. Comput.-Hum. Interact.* **13**(4), 458–489 (2006)
47. Xiao, C., Qin, J., Wang, W., Ishikawa, Y., Tsuda, K., Sadakane, K.: Efficient error-tolerant query autocompletion. *Proc. VLDB* **6**(6), 373–384 (2013)
48. Xiong, L., Agichtein, E.: Towards privacy-preserving query log publishing. In: Proceedings of the WWW Query Log Analysis Workshop (2007)
49. Zobel, J., Moffat, A., Ramamohanarao, K.: Inverted files versus signature files for text indexing. *ACM Trans. Database Syst.* **23**(4), 453–490 (1998)