Characterizing Searches for Mathematical Concepts

Behrooz Mansouri†
Golisano College of Computing and Information Sciences
Rochester Institute of Technology
Rochester, NY, USA
bm3302@rit.edu

Richard Zanibbi
Golisano College of Computing and Information Sciences
Rochester Institute of Technology
Rochester, NY, USA
rlaz@rit.edu

Douglas W. Oard
iSchool and UMIACS
University of Maryland
College Park, MD 20742, USA
oard@umd.edu

ABSTRACT

Although there has been considerable interest in recent years in the development of specialized mathematical digital libraries that can index and search for mathematical documents, little is yet known about how people will search for mathematical concepts. To begin to gain some insight into that question, this paper examines the nature of queries for mathematical concepts that were created by users of a search engine. A total of 392,586 queries that contained at least one distinctive mathematical term (e.g., "Taylor Series") were identified in a two-year query log from the Parsijoo search engine; these queries were each issued in one of 69,014 search sessions. Descriptive and comparative analysis indicates that math search sessions are typically longer and less successful than general search sessions. Queries for mathematical concepts exhibit greater diversity than do queries in general – essentially the query distribution is nearly all "tail." The use of well-formed questions as queries (e.g., "How can I prove that …") is also surprisingly common, as is the creation of queries by copying text from a document. Among the implications of this study for the design of search engines with specialized functions for handling mathematical notation are that robust support for query refinement and reformulation could prove beneficial.

KEYWORDS

Math search, User behavior, Query log analysis

1 Introduction

One distinguishing feature of digital libraries is that their digital nature makes it possible to leverage a broad range of "signals" to characterize the elements of a collection, notably including metadata, content, and user behavior. The initial focus of content indexing was on text, but other content types (e.g., images and music) have received attention as well. Techniques have recently been developed to search based on mathematical notation, but because those capabilities are yet to be widely deployed that technology development has to date largely been driven more by fundamental questions such as how to recognize when one formula is similar to another than by a solid grounding in actual user tasks. Whether we are considering the proposed Global Library of the Mathematical Sciences [58] or the addition of new math-oriented search capabilities for existing digital libraries such as CiteSeerX that already contain mathematically rich materials, we see a need to better understand how users will interact with what we might refer to as Math-Enabled Digital Libraries (MEDL).

The foundation of math search is a growing body of information retrieval research in which the principal focus has been on matching mathematical notation [54,55]. Several challenges have been addressed, most notably mapping between surface forms and logical forms of mathematical notation [14], computing similarity based on logical forms [18], and linking mathematical notation with associated text [20]. This work has led to the deployment of specialized search engines designed to demonstrate the potential of that rapidly evolving search technology. For instance, "Search on Math"[1] searches Web pages, providing a keyboard that users can employ to include formulae in their queries. There are also specialized search engines for mathematical content, such as Mathematics Stack Exchange [2] which supports search over previously created content in a community question answering site focused on mathematics questions. Of course, users also employ general search engines when looking for information about mathematics, or in which mathematical concepts might be useful.

Too little is currently known about how searchers will use any of these capabilities. Information seeking behavior researchers have examined the information seeking practices of scientists [34], and collaboration behavior in Mathematics StackExchange has been studied [44], but in general, there is little in the way of studies of user behavior to guide MEDL design. To address this gap, in this paper we use search engine query log analysis. Our goals are to gain some insight into what people are trying to do when they are performing math search, into how well they are currently doing it, and into what design guidelines might be inferred based on those insights. To the best of our knowledge, this is the first work to use search engine query logs to study how users search for mathematical concepts, and specifically to examine how that behavior differs from more general search behavior.

In our study, we analyze a two-year query log of a Persian search engine, recognizing mathematical queries based on the use of specialized terminology, and assembling sessions around those queries in ways that leverage a combination of persistent HTTP connections and idle-time thresholds. Query logs are a valuable resource for this kind of study, but because we do not have access to the users posing those queries, log analysis ultimately relies on inferences drawn based on statistical aggregation. We should, therefore, think of this as the first step of a process that will

[1] www.searchonmath.com

[2] https://math.stackexchange.com/

ultimately grow to include other types of user studies and search engines with more specialized capabilities.

Our analysis is organized into three main parts. We begin our analysis by focusing on the queries, finding that math queries are considerably longer on average than is typical of Web queries generally. We find that one reason for this is that the use of fully formed questions is surprisingly common and that the use of what appears to be long runs of cut-and-paste text is common among the longest queries. Importantly, we also find that the query distribution is essentially "headless", with fewer than 1% of the query instances resulting from frequently repeated queries, and about 90% of the unique queries being issued only once in a collection of nearly 400,000 math queries.

We next focus on how searchers use math search results, through observing the clicks that follow the queries. We find that long queries are not just more common, but also less often successful. We infer this from a "sat click" analysis in which we infer that staying on a Web page long enough to read it is (in the aggregate) a useful indication of user satisfaction [19]. Our click analysis also indicates that frequently repeated queries are not merely rare, but also difficult for the search engine to respond to well. We base that analysis on measurements of click entropy showing a decline in entropy (meaning that the click pattern exhibits less scatter across the ranked results list) as queries become more common, but only to a point. Beyond that point (in our collection, for queries issued more than five times), click entropy increases. This suggests that the few commonly issued queries are generating click patterns that are more, not less, scattered. We also note that the clicks tend to accumulate on a relatively small number of specialized Websites, as might be expected for a specialized search task.

Finally, we group sequences of queries and clicks into sessions to study the process by which users iteratively refine their searches. We find that math search sessions are both notably longer on average, and notably less often successful than is typical of general searches on the same search engine.

Throughout our analysis, we suggest design implications based on our results. For example, we note that replacing query terms is often more successful than adding or deleting query terms (leading to a "sat click" that we infer to indicate user satisfaction more than three times as often); that result could have implications for the design of query suggestion techniques. The next section reviews related work on math search, user behavior analysis, and user behavior analysis for math search. Section 3 then describes the search engine and query log that we have used, how we identified math queries, and how we divided query and click sequences into sessions. Section 4 then presents our three types of analysis, focusing in turn on queries, clicks, and sessions. Finally, Section 5 concludes the paper with a few remarks on the next steps.

2 Related Work

Our study of math search behavior builds upon existing work in math-aware search engines and query log analysis techniques that observe users' information seeking behaviors and identify opportunities for improving search results. We summarize these studies and their relationship with our own work below.

Mathematical Information Retrieval (MIR) Systems. Currently, widely-used general search engines have very limited support for mathematical notation. LaTeX may be used in keyword search, but this generally produces low-quality search results. Neither the syntax nor semantics of mathematical notation are represented in these engines. Further, LaTeX is unfamiliar to most search engine users, often leaving them incapable of expressing unfamiliar math symbols and structures in their queries [48]. These limitations have motivated the development of specialized math-aware search engines. A typical MIR system has an interface where users can express their information need (as some combination of text and formulas), and an interface for presenting search results. Formula retrieval models can be grouped into two categories: text-based and tree-based [56]. Text-based approaches adapt traditional text retrieval models by linearizing formulas into a sequence of tokens and then treat them as text [38,21]. Tree-based methods represent formulas separately [15,56], using hierarchical (tree-based) representations for formula appearance (e.g., LaTeX or Presentation MathML) and formula semantics (e.g., Content MathML, representing operation sequences). Tree-based methods have proven to be much more effective for formula search, as demonstrated in the NTCIR [1,54] Math Information Retrieval tasks. NTCIR-12 used two corpora: one containing excerpts from technical articles in arXiv, and a second corpus containing English Wikipedia articles. For both corpora, there were two subtasks, one for combined math and text queries, and the other for isolated formula queries. In the NTCIR-12 MIR task and subsequent research, the intelligent use of context and embeddings have shown to be promising directions for math search [12,37,52].

User Search Behavior Studies. Currently, information retrieval researchers primarily use three approaches to study user search behaviors: analysis of search engine query logs, user studies, and user surveys. There are a number of advantages for query logs over the other two approaches, related to cost and ecological validity. Most importantly, query logs contain interactions between a search engine and users performing actual search tasks. A much larger number and a greater diversity of users (e.g., ages, geographic locations, and backgrounds) may be observed in a query log than in a survey sample or a user study -- it is not uncommon for a log to contain millions of records. Studying query logs is also cheaper: for surveys, collecting a meaningful sample may require paying participants, along with additional time for coding and analyzing open-ended responses, and surveys are subject to self-report bias. User studies can be designed to minimize the effect of self-report bias, but structured lab studies sacrifice some degree of ecological validity to focus on specific phenomena, whereas observational studies of situated users are expensive both to conduct and to analyze. For these reasons, we find query log analysis to be an attractive approach for this initial study.

Logs may be used to identify interaction patterns and places to improve the quality of services including search results, query suggestion, and query auto-completion. Query log studies may be divided into two categories; 1) *general search*, considering an information retrieval concept such as query reformulation for all types of search, and 2) *specialized search*, considering specific

Table 1. Parsijoo transaction logs for two search sessions with math queries.

Session ID	Query	Issue Time	Search Type	Clicked URL (Site shown)	Click Time	Session End Time
A33C14AC80CD4	double integral concept	2016-08-23 10:19:03	Web			
A33C14AC80CD4	double integral solving methods	2016-08-23 10:19:28	Web	fa.wikipedia.org	2016-08-23 10:19:44	
A33C14AC80CD4	double integral solving	2016-08-23 10:20:09	Web			
A33C14AC80CD4	double integral examples	2016-08-23 10:20:35	Web	riazisara.ir	2016-08-23 10:22:09	2016-08-23 10:57:14
A33C14ACB5698	cauchy schwarz inequality problems pdf	2017-02-05 16:49:13	Web			
A33C14ACB5698	cauchy schwarz inequality explained notes	2017-02-05 16:52:20	Web	kanoon.ir	2017-02-05 16:52:39	
A33C14ACB5698	cauchy schwarz inequality	2017-02-05 16:59:41	Web			
A33C14ACB5698	Tutorial on cauchy schwarz inequality	2017-02-05 17:19:43	Web			
A33C14ACB5698	Tutorial on cauchy schwarz inequality	2017-02-05 17:19:52	Video	kelasdars.org	2017-02-05 17:20:14	2017-02-05 18:01:30

domains such as medical, academic, and job search, and examining how they differ from general Web search. Our work falls within the specialized search category.

Related to log studies concerned with general search, Huang and Efthimiadis [15] consider how users modify queries to obtain better results in AOL query logs. They provide a taxonomy for query reformulation and identify which reformulations are most effective. This study was later extended to consider query reformulation and suggestions [17,31,41].

For log analysis in specialized search tasks, Spina et al. [42] extract queries related to job searches from a general search engine, and then compare job searches with general Web search. They show that job search characteristics are distinct from general search, and that different methods might be used to better satisfy job-related information needs. In another approach, some work considers specialized search tasks in vertical search engines. Li et al. [25], use a 5-month query log from the ScienceDirect search engine to study academic queries. They focus on *null queries*, where search engines fail to provide any results to users. In our own work, we examine extracted math queries and search sessions using logs from the Parsijoo[3] search engine, and include an analysis of success and failure in math search sessions.

User Search Behavior in Math Search. To date, there has not been a characterization of math search from general search engine logs. The most similar work to ours is a previous study by Zhao et al. [57], where 13 individuals affiliated with the Math Department at the National University of Singapore were interviewed. This was the first study of MIR user requirements, but the sample size is small, and all participants had a strong mathematical background: participants included two undergraduate students, seven graduate students, one professor, and three librarians. In contrast, here we present the first study of math search based on actual queries and search sessions for users of a real-world search engine.

3 The Parsijoo Query Log

For our analysis, we use the query logs of a Persian general-purpose search engine, Parsijoo. Parsijoo is the second most popular search engine in Persian language countries: in 2015 Parsijoo indexed one billion multilingual Web pages. The query log used in this study contains 27 million records of user interactions over two years (March 2015 to March 2017). Previously, Parsijoo query logs have been used to study temporally specific ("spiky") events and job search [29,30].

The Parsijoo transaction logs record communication between users and the search engine. As Table 1 shows, these transactions logs include Session IDs, submitted queries, query issue times, search type (e.g., Web or video), URLs for search results that users click on, click times, and session end times (*Note:* Session ID and Clicked URL fields are shortened for space).

We extract mathematical queries using the keyword-based approach of Mansouri et al. [29]. In that work, editors were asked to manually construct a list of phrases used for job searches, and queries containing those phrases were then identified as job search queries. In our work, three editors compiled a list of mathematical concepts in Persian from two widely used college-level math textbooks (Thomas' Calculus - Thirteenth Edition (Volumes I, II) and Essential Calculus with Applications (Volumes I, II, III)) along with textbooks from the elementary through high school levels. The editors created a list of 681 mathematical concepts such as "Gaussian distribution," "Taylor series" and "cosine" (in Persian). The length of these concept phrases was between 1 to 3 words, with an average of 1.8 words per phrase. Queries containing any complete math concept phrase are included in our initial pool of math queries. At the moment, Parsijoo, does not support user profiling, so there is no information available about the users' knowledge of mathematical concepts, but the keywords covered concepts from elementary school to undergraduate degree.

To have a query set with high precision, phrases such as "division" that have both mathematical and more general meanings were not considered. Also, some mathematical concepts are used in the titles of movies or as the name of music albums. So in our cleaning steps, we removed queries containing words such as "music", "album" or "cinematic movies." In total, 392,586 mathematical query log records were extracted.

Parsijoo assigns each unique browser a Session ID when it first connects that typically persists (using persistent HTTP) until that browser is closed or the connection is lost. Session IDs are not reused by Parsijoo, so the next time that same browser is opened

[3] www.parsijoo.ir

and connects it will receive a new Session ID. Some of these Parsijoo sessions are far longer than we would expect a user to be working on a specific task, so we subdivide these sessions into smaller math search session units for the purpose of our analysis. We do this by first partitioning Parsijoo sessions into math search sessions whenever there is no transaction for 60 minutes or more, and then removing all partitions that do not contain at least one math query. This results in 69,014 math search sessions. Note that math search sessions need not start or end with a math query; they simply must contain at least one math search query. In order to compare math searches with general searches, we partitioned every session in the same way and then selected partitions at random until we reached the number of math queries that we had started with (392K).[4]

4 Results

In this section, we present the results of our analysis of math queries, page visits (clicks) and math search sessions.

4.1 Math Queries

We begin by studying math queries in isolation.

4.1.1 Content Type

Math search is perhaps not just a single vertical since the technology to search for papers on arxiv.org that are available as PDF might be quite different from the technology needed to search for Khan Academy videos, even if the queries are similar. So the first question we might ask is what kinds of content is it that users are looking for? As with many search engines, Parsijoo allows users to specify a vertical to search (e.g., Web, News, Video). One easy source for the desired content type would therefore be the request type field in the query log. However, that field is not very informative because 95.8% of math queries were issued as general Web searches with no more specific vertical selected. However, users also will often specify something about the type of content that they hope to find as a part of their queries, and Web search engines will typically parse and interpret such terms as soft content type specifications. We therefore looked for content type terms in the queries. As Table 2 shows, a content type term was present surprisingly often in math queries. For resource type such as Video, list of video formats such as .mp4, .avi and .3gp was considered.

Table 2. Frequency of common content type terms in math queries. Results are cumulative, starting with the most frequent content type (e.g., "PDF" is counted only in queries that do not contain "Tutorial").

Resource Type	Math Queries
Tutorial	12.3%
PDF	9.7%
Video	8.4%
Download	6.1%
Notes	5.9%
Powerpoint	5.3%
Total	*47.7%*

[4] Note, however, that our math query partitions actually contain a total of 401,662 queries, 392,586 of which are math queries.

As can be seen from the total, nearly half of the math queries include a keyword that we might reasonably interpret as indicating the type of content that is being sought. Note that this is a rather eclectic mix of types, some of which specify file types and others of which describe a document genre. We selected these terms that we believe to be indicative of content type manually after inspecting enough of the query log to have confidence that we had seen the most frequent content type terms.

Tutorials (which might suggest text, slides, video, or any combination of that) were the most frequently requested content type (12.3%), while PDF and Video were the next most popular requests (9.7% and 8.4%, respectively). From this we can conclude that Math search engines will ultimately need to be able to do more than just search HTML; video search [8], and specialized handling for specific formats such as PDF and Powerpoint that might include math in either text of image form, would be useful. There may also be some scope for personalization [10] here since although user preferences are not always explicit in their queries, to the extent those preferences are persistent across sessions the search engine might learn to bias the results towards content types that have previously been requested.

4.1.2 Query Length

One challenge search engines face with Web queries is that they are typically quite short, thus requiring a good deal of guesswork on the part of the search engine to arrive at a useful interpretation of the searcher's intent [2]. For example, the set of general Parsijoo queries that we assembled as a basis for comparison has an average length of 3.4 words.

The average math query, by contrast, is nearly twice as long, at 6.7 words. Figure 1 compares the distribution of query lengths for math queries and general (non-math) queries. Math queries are much longer on average: indeed, the third quartile of the average query length for general queries is about where the first quartile is for math queries. 46% of the math queries are between 4 to 6 words long, producing a right-skewed distribution for math queries.

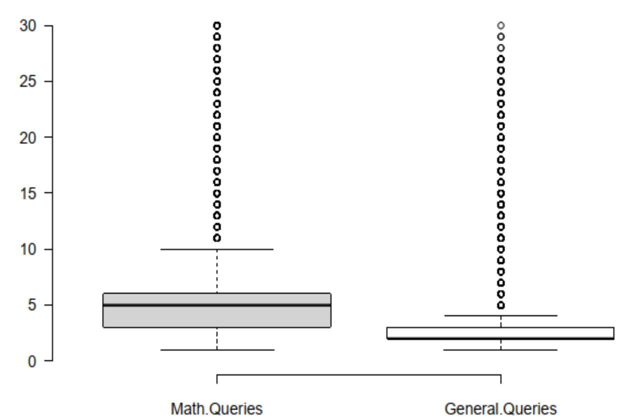

Figure 1. Query length distributions: math vs. general queries.

4.1.3 Cut-and-Paste Queries

One user behavior that can produce long queries is text reuse when searchers copy long passages of text into the query box. To check for this, we selected the top-2000 most frequent unique math queries with length ≥ 20 and issued them to Parsijoo. By using quotation marks to obtain exact matches on longer sequences of text, 49.8% of these queries found documents containing exactly the query text, making it likely that these queries were entered into the query field using cut-and-paste from some document in the indexed collection.

4.1.4 Verbosity

Another thing that can make queries longer is the presence of unnecessary words, a phenomenon referred to as *verbosity* [9]. Using an example from our query log, the query "How can I expand Taylor series?" is considered verbose. Bendersky and Croft [3] have shown a correlation between users' click behavior and the query length – generally, the performance of retrieval decreases as query length increases.

To check for verbosity, we considered the Top-2000 most frequent unique math queries by the frequency with the length greater than 4 (i.e., longer than the average query length over general search queries). We then asked 3 annotators to label the queries as verbose or not. The annotators were given the definition of verbose query as "long query with words that are less informative and can be ignored without changing the purpose of query", along with 30 verbose queries examples. To measure the inter-rater agreement between annotators we used Fleiss' kappa statistic [11]. Overall, the level of agreement was 0.79, which represents high agreement between annotators. 17.8% of the queries under study were tagged as verbose. There has been some work on converting long queries into better structured and more concise representations using segmentation, reformulation, weighting and reduction [22,28,35,39] that may prove useful for this part of the query stream.

4.1.5 Question Queries

One possible reason that math queries are longer and more verbose may be that they are issued as questions (i.e., are *question queries*). In this subsection, we explore the use of actual questions as math queries. To extract question queries, we used the approach of Zahedi et al. [53], where they studied how Persian question queries are posed to a search engine. We use their keyword-based approach, creating a list of Persian question words to identify math queries that are question queries.

Our study found that 18.4% of math query instances (and 19.8% of unique math queries) are issued as some form of questions. This number is much higher than the 1.8% reported in [53] for general queries, and close to the 17.8% of queries that our annotators marked as verbose. Indeed, 90% of question queries were marked as verbose, so a very substantial number of the queries marked as verbose were marked in that way because they contain question structure that is not related to the content of the query. Question type analysis can, however, be used to refine search engine results, for example returning different results for "Why" questions than for "How" questions [45].

Table 3. Frequent question words and accompanying terms in math question queries. 'Other' question words include *Which, Where, Whether,* and *Who*. '?' indicates no question word is used, but a question mark is included. Query percentages are of question queries; Accompanying term percentages are of that question type.

Queries	Cue	Accompanying Term
69.5%	What	Formula (*60%*), Equation (*11%*), Used for (*9%*)
13.8%	How	Prove (*51%*), Exact (*10%*), Accurate (*5%*)
6.2%	Why	Is (*20%*), Not (*15%*), Correct (*7%*)
3.9%	When	Use (*52%*), Apply (*21%*), Consider (*4%*)
4.3%	*Other*	Formula (*32%*), Equation (*12%*), Invent (*9%*)
2.3%	?	True (*14%*), Answer (*12%*), Principle (*6%*)

To get better insight into math question queries, the distribution of question cues and associated terms in math question queries are presented in Table 3. The table shows that some question words are frequently paired with other words. For instance, "What" is the most frequent question word, used in 69.5% of math question queries. "What" is often followed by the words "Formula" and "Equation," suggesting that users are looking for mathematical notation. Such information can be used in query suggestion, auto-completion, and query expansion. For instance, when a user issues the query "What is Poisson distribution …" probable words for auto-completion are "formula" or "equation". Another example is the question word "How," which was more than half of the time repeated with the word "Proof", indicating that users want to know how to prove a mathematical statement.

Question type analysis is a subtask studied in the broader field of Question Answering, a complex task that requires a retrieval system to correctly answer questions posed in natural language. This is an active research area at the intersection of information retrieval and natural language processing [43,50,51]. Math question answering, in particular, has also received some attention in recent years [33,46,47].

Overall, our analysis shows that math queries are longer than general search queries due to the use of copy-and-paste queries, formulation of a substantial number of queries as questions, and some degree of verbosity. These results suggest that multiple complementary strategies may be needed to produce the best results for longer queries.

4.1.6 Query Uniqueness

One advantage that search engines often seek to leverage is that a small number of identical queries are often issued by many different users. This phenomenon has important implications for both search quality (since the clicks of early users can be used to predict what later users are most likely to want to see) and efficiency (since responses to frequently issued queries can be cached, which can dramatically be reducing time-consuming disk and network activity). To understand how often math queries are

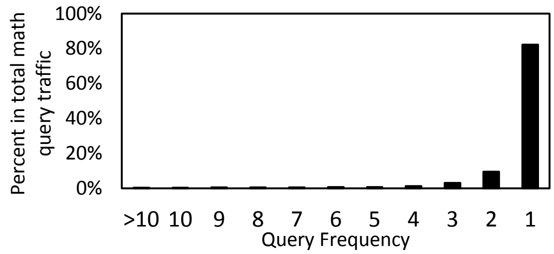

Figure 3. Distribution of math query frequencies.

repeated, we plotted the distribution of math query instances by frequency over the two-year period in Figure 3.

As can be seen from this figure, more than 80% of our math query instances appear in our query log only once. One implication of this result is that query suggestion techniques that rely on rich data about query frequency are not likely to work well in this setting. Recently, researchers have tried using side content (e.g., putatively relevant documents, or contemporaneous blogs) as a basis for query suggestion in a manner similar to pseudo-relevance feedback [23, 40]. Mitra and Craswell [32] have also proposed a vector representation approach for query auto-completion of rare queries. Similar methods might be tried to produce better query suggestions for our relatively infrequent math queries.

4.2 Page Visits (Clicks)

Queries tell only one part of the story – we can learn more by looking at queries and clicks together.

4.2.1 Frequently Visited Websites

First, we consider which Websites users view during math search. Patterns of this type can be leveraged to bias search results in favor of sites that are known to be useful. The top-10 Websites (i.e., the top-level domain of the Web page) clicked on in math search sessions. Together, these account for more than 25% of all page clicks. With the exception of Wikipedia, all of the top Websites focus on mathematical or scientific content. Wikipedia, of course, includes many pages with that type of content as well.

4.2.2 Click Entropy

Going further, we investigate the variability of Web pages clicked in math searches. There were two reasons to study this. First, the top-10 clicked sites account for more than a quarter of clicked pages, which suggests that users prefer certain Web sites for math search. The second reason is to check if the diversity of resources specified by users will lead to diversity in clicked pages. To measure the diversity of clicked pages we calculated the click entropy introduced in [10] as follows:

$$S\,(q) = -\sum_{URL\ u\,\in U(Q)} P(u|q)\log_2 P(u|q) \qquad (1)$$

Where $U(q)$ is set of URLs clicked by users after issuing query q and $P(u|q)$ is the probability that the URL u is clicked by users in when the search results for query q is shown. The results obtained show the entropy for the average math query is 3.33, which is quite close to the average click entropy of 3.18 for the general queries (selected as described in section 3). Averages obscure many details, however, so in Figure 2 we plot the click entropy for unique queries stratified by frequency.

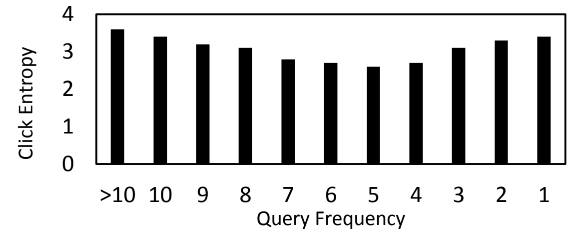

Figure 2. Average click entropy for math queries by frequency.

For unique queries issues just once we see an entropy of about 3.4. Converting entropy to perplexity by taking the antilog, this equates to an average of $2^{3.4}$=10.6 clicks per query which, for singleton queries, are by definition uniformly distributed. For instance, for the query "prove that the diagonals of a parallelogram bisect each other using vectors?", a user viewed 8 different pages to find the information they were looking for.

As queries are reused, position bias and perceived content quality effects result in greater probability mass accruing on some queries, and we see the expected consistent reduction in entropy. However, as queries become even more common – here, once unique queries have been issued more than five times – the entropy begins to climb. Common queries are often short, and short queries are often ambiguous, so one possible explanation for this observed effect is that as queries become more commonly issued, a point is reached where the ambiguity effect (which tends to increase entropy) overwhelms the quality and positional bias effects, thus resulting in a net increase in entropy.

For frequent queries, different users preferred different pages. For instance, for the query "triangle inequality", the click entropy was 4.7. Examining the most clicked-on URLs, there appear to be different user intentions: some of the users were looking at pages related to the definition, proof, equation, or application of this concept. Diversity ranking techniques that use result clustering to maximize the coverage of alternative interpretations or alternative facets of a query [7] can be a helpful response to ambiguity, and thus are potentially useful for more commonly issued queries.

4.2.3 No-Click Queries

One way to study failure is to analyze the characteristics of queries that led to no clicks before the next query was issued (or the end of the session), which we call *No-click queries*. Because we expect long queries to be challenging, we chose to look at the distribution of query length (in words) both for queries in general and for No-click queries in particular. We consider query lengths between 1 to 30 words and studied the distributions for general queries, all math queries, and No-click math. While the distribution of both general and all math queries fall off as the query length increases beyond some point, the distribution of no click math queries has a longer tail. So we see that math queries are long, and that long queries are often unsuccessful.

To get a clearer view, we randomly selected 100 No-click queries of length 5 or 6 (a range where there were also many successful queries) and examined them to see if we might guess the reason why there had been no click. We found three plausible explanations for why users are not clicking on results for queries of that length. One is that the user may be unsure about the mathematical concept (s)he is searching for, with the user's initial query being different from the final query, mostly moving from more general to more detailed and specific mathematical concepts. Consider one search session where the initial query was "examples for math series expansion," reformulated to "Series expansion methods" and ending the session with the query "Maclaurin Series." Another example is a search session starting with the query "area and perimeter of geometric shapes," where the top results are for general shapes such as circles, rectangles and squares. The user then made the query more specific by issuing the query "area and perimeter of shapes with multiple sides." The last two queries issued by this user were, "area and perimeter of Octagon" and "area and perimeter of Pentagon."

Another type of No-click query that we observed resulted from the user requesting exact matches (by using quotation marks). For instance, one user first issued the query "Cholesky decomposition and Newton's method comparison" and then reformulated the initial query to the new query "matrix inversion solutions" and finally posed the query "comparison of matrix inversion solutions" with the entire thing in quotation marks. That did not work well. Finally, a third apparent cause for No-click queries was invalid input, either because of spelling or because of deficiencies in vocabulary. For example, there was a search session where the initial query was "the median value theorem for integrals." However, the correct mathematical concept is "the average value theorem for integrals." Another example is a session where the user meant to issue the query "Kruskal's algorithm in graph theory," but instead of the word "Kruskal" used the word "Truskal." That user finished his/her search session with no click.

All three of these cases point to the value of query suggestion Considering the first scenario, suggesting related queries for more specific mathematical concepts could help guide users to desired pages faster. For exact-match queries and invalid input, perhaps query auto-completion might help users to avoid entering No-click queries in the first place.

4.3 Math Search Sessions

Queries and clicks merit attention individually, but usually they are part of some larger tapestry that we call a session. Users come to a search engine with some goal in mind, but their strategy for achieving the goal, and perhaps the goal itself, evolves over the course of their search. We call the sequence of queries and clicks for this often-complex search process a "search session." In this section, we study math search sessions from the perspectives of user effort, math query refinement, and failed searches. We then examine some long math search sessions to obtain additional insight into math search behavior.

4.3.1 User Effort

One simple way of quantifying user effort is to look at the number of queries in a session. We can also count the time between the first query and the last click of a session, which we call session duration, or the total number of clicks. As Table 4 shows, all of these measures are markedly greater for math sessions than for

Table 4. User effort metrics for math vs. general search. The mean and standard deviation for each measure is shown.

Session Type	Queries (μ, σ)	Duration (mins.) (μ, σ)	Clicks (μ, σ)
Math	5.82 (1.62)	10.42 (5.58)	3.28 (1.03)
General	1.86 (0.90)	3.11 (2.89)	1.32 (0.94)

general search sessions. On average, math sessions are more than three times longer, both in duration and in the number of queries, than general search sessions; the average number of clicks is about two and a half times greater. Some of this might be attributed to math searches being predominantly informational, using Broder's query type taxonomy [5], whereas the higher prevalence of navigational queries in general Web search would skew the average toward general searches toward shorter sessions. Indeed, it may be that exploratory search [49], a particularly challenging type of informational search, is more common for math since one common use case we have observed for math search is math learners of school age who are working on assignments.

User Satisfaction. Regardless of the cause, one thing we can say with some confidence is that searchers seem to be working pretty hard on their math searches. That naturally inspires us to ask about the degree to which that hard work is paying off. While we can't ask our users about how much they like their search results, we can see what they do after they get them. Essentially, we want to find what White has called "sat clicks" – interaction patterns that suggest (fairly reliably, when viewed in aggregate) user satisfaction. Specifically, we quantify three outcomes that we expect to be informative with regard to user satisfaction:

- *Zero-click sessions.* Sessions in which the user did not click on any search engine result. We interpret this as clear evidence of dissatisfaction.
- *Click-final sessions*. Sessions where the user ended the search session by visiting a Web page, suggesting that the user's information need may have been satisfied.
- *Sat-click sessions.* Sessions in which the user stays on a page for more than 30 seconds, suggesting that at least some part of the user's information need may have been satisfied. Note that many click-final sessions will also be Sat-click sessions since a Click-final session in which the browser remains open for 30 seconds is also a Sat-click session.

Table 5 compares these satisfaction measures between math and general search sessions, showing the percentage of sessions in

Table 5. User satisfaction metrics for math vs. general search.

Session Type	Zero-click	Click-final	Sat-click
Math	23.9%	25.2%	60.7%
General	10.5%	42.4%	82.2%

which each outcome occurs. From these results we can reasonably conclude that math search users are not just working harder, they are also getting worse outcomes. More than twice as many math search sessions result in no clicks at all, a bad outcome, markedly fewer math search sessions are Click-final, and even Sat-clicks are less common for math search sessions.

Example Sessions. What's so hard about math search? A few examples can help to illustrate the types of problems we see in our query log. Consider one session where the user was trying to answer the question query "what is derivative of ax power 2+bx+c?". The user first tries to input the formula "$ax^2 + bx + c$." After not obtaining a relevant result, the user then tried the query "how to calculate derivative for polynomial?" Another example is a search session where the user first submitted the query "C (10,5)." After not clicking on the result pages, the user reformulated the previous query to "C (n,k)". For the second query, the user clicked on three Web pages but did not stay on any of them more than 30 seconds. Then the user reformulated the query as a verbose text query, "How can I select 5 items from list of 10 items". For this query, the user did not click on any result Web pages, and changed the query to the more abstract query "Selecting k items from n items". After staying on a Web page for two minutes, the user returned to the search engine and issued the query "Combination k of n." As we see from these examples, users are hunting for ways of expressing themselves in ways that match what the search engine can use well.

Next, we consider an example session that clearly signals the importance of formula input for math search. The query "integral from 0 to 1 of 1/x^p", was submitted by a user searching for the formula "$\int_0^1 \frac{1}{x^p}$." After viewing four result pages, the user changed the query to "examples on Integral calculation". After issuing this query, the user spent 42 minutes on one of the result pages and then finished the search session.

Two common features of each of these examples are that general search engines have an impoverished vocabulary that formula input might help with, and that the key to success in the face of an impoverished vocabulary is to support effective query reformulation.

4.3.2 Query Reformulation

Query reformulation has been widely studied, both with the goal of improving retrieval effectiveness [13] and with the goal of supporting effective reformulation [17]. Here we consider seven types of query reformulation that are common to many previous studies [4,6,24,36]:

- *Reordering words.* Words from the previous query are re-ordered (e.g., Series Taylor → Taylor Series). This type of reformulation might be motivated by the user seeing the words in some other order in a document.
- *Removing words*. Removing at least one word from the previous query. This is often associated with a generalization strategy in which a user elects to search for some broader, and thus less full specified, concept (e.g., shortest path in graph with Prim→ shortest path in graph).

Table 6. Math query reformulation frequencies, and percentage of reformulations leading to user clicks by type.

Reformulation Type	Frequency	Led to click
Substituting words	32.1%	62.9%
New query	21.5%	44.7%
Multi-reformulation	15.2%	58.3%
Adding words	13.5%	17.4%
Removing words	9.6%	18.3%
Reordering words	4.3%	12.9%
Revisiting	3.8%	15.9%

- *Adding words*. At least one word is added to the previous query (e.g. optimization method → optimization newton method). In contrast to removing words, here the user makes the query more specific.
- *Substituting words*. Having the same number of words, with at least one word in common with the previous query (e.g. Standard normal distribution → general normal distribution). This behavior can occur when refining a previous query, or exploring related concepts.
- *New query*. There is no common word with the previous query (e.g. Taylor Series → Fourier Transform). This situation sometimes occurs where the user's goal changes [4], but it might equally well reflect an iterative strategy in which the user is working through a set of possibilities without having changed their goal.
- *Multi-reformulation*. Here more than one of the reformulation types defined above occur (e.g., Taylor Series Expansion Example → Taylor Series Formula: "Expansion" is removed, and "Example" substituted by "Formula").
- *Revisiting*. The user returns to using a query they previously issued in the same search session. This might occur when the user abandons a strategy and wants to start over from a known point with a different strategy, or it may simply reflect the user satisficing by using a previous query to return to some document that they had previously seen, perhaps after learning that no better document can be found.

Every consecutive pair of queries within the same session defines a reformulation, and these reformulation types are easily detected with simple string edit measures. A total of 83.6% of our math search sessions contain more than one query, and thus (by our definition) at least one reformulation. Table 6 shows the frequency of each reformulation type among sessions containing at least one reformulation, along with a simple measure or the quality of whether the reformulation was useful (for which we simply use how often a reformulation of that type led to a click).

As the table shows, the substitution was the most common reformulation type (accounting for nearly one-third of all reformulations), and the most effective reformulation type (leading to a click nearly two-thirds of the time). Multi-reformulations, 76% of which included at least one substitution, was also nearly as effective as substitution alone (at least at finding something to click on). Reordering words and revisiting queries submitted earlier were the least common reformulations, and also the least successful. These results suggest that as query interfaces for math search become more capable, users might benefit not just form the ability to enter equations, but also from the ability to easily reformulate equations that they have entered in earlier queries. These insights might also help with the design of query suggestion techniques, which today often focus more on adding words (e.g., for auto-completion) than on substitution. We also looked at how often users switch from using search queries to asking questions, given the high frequency of question queries in math searches compared to general searches. Liu et al. [27] had observed that users unsatisfied with their search results sometimes changed their queries into a question. However, our analysis of math searches shows that in 40.3% of cases where users posed a question, the question query was the initial query issued in the search session. For instance, in one search session, the user's initial query was "how to prove a relation is symmetric." As best we can guess the user seemed as if they may have been satisfied with the result, concluding their session after visiting two of the returned pages. From this relatively large prevalence of question-initial search sessions, we might speculate that users may be adopting this strategy because they expect the additional context to help the search engine to return better results. If so, perhaps we should try to do that!

4.3.3 Long Sessions

As we've discussed, in math search sessions users are submitting more queries and spending more time than for search sessions in general. To get a sense for the most extreme cases, we examined the ten math search sessions with the largest number of queries to try and determine why users are issuing more queries. Overall, in these sessions two types of behavior were noted: 1) the user is trying to solve a set of mathematical problems (perhaps homework), or 2) the user is trying to fully understand a mathematical concept.

The largest number of math queries issued by a user in a math search session was 26. In this session, our guess is that the user was a student working on geometry homework. Different math queries, each related to one of five geometry concepts discussed in high school, were issued. Interestingly, the user first apparently tried all of the questions using copy-and-paste. For two out of the five geometry concepts, (s)he clicked on Web pages after issuing the copied text (which was in the form of a question), and neither of those two queries were reformulated. This suggests that, the users may have found the answer to their question. However, for the three other queries, after not being satisfied with preliminary search results, the user tried to issue different queries to learn the geometry concept. By starting from the basic concept, the user tried different reformulations to locate more complex material on that concept. Overall, the user finished searching for the first four concepts with a click, but seems not to have been satisfied with results for the last concept as no clicks resulted.

Six more of the ten longest sessions seem to have been users were looking for answers to certain math problems. In the other three very long sessions users seemed to be trying to understand a mathematical concept in detail, continuing the search even after visiting Web pages. For instance, in the third-longest math search session, 20 queries were issued by a user concerning "trigonometric identities." In this search session, a 'new query' reformulation occurred 15 times as the user tried different queries included "sine formula", "tangent" and "cotangent."

Our analysis of these ten long math search sessions suggests that math searches can be long for different reasons, and that it might therefore be useful for math search engines to include functions for inferring user intent [16].

5 Conclusion

Our current understanding of how users search for mathematical information is limited. In this paper, we have presented the first query log analysis for math search, carried out using the general-purpose Parsijoo search engine. Our main finding is that we observe high user effort and low user satisfaction for math search.

Several properties of math-related queries were different from general ones such as query length, use of question queries and diversity of clicked page which indicates the importance of different search service design for mathematical searches. During math search session users spend more time and issue more queries compared to general searches, and user satisfaction measures were lower. Several of the longest search sessions appear to be searching for solutions to a set of homework problems, but others seemed to focus on trying to understand a mathematical concept in detail.

These findings have important implications for the design of future Math-Enabled Digital Libraries, although much remains to be done. In our MathSeer project we are building a vertical search engine for math. As a next step, we need to begin to integrate provisions for indexing mathematical content to enrich the text-based search capabilities that we have studied. After that, it would be natural to explore the potential for leveraging text together with mathematical notation found in documents and traces from user behavior to organize search results, and ultimately to support new interaction modalities. This study tends to confirm our belief that there are both user needs and technical opportunities that will ultimately come together in the creation of Math-Enabled Digital Libraries.

ACKNOWLEDGEMENT

We would like to acknowledge Parsijoo for providing data for this research. This material is based on upon work supported by the Alfred P. Sloan Foundation under Grant No. G-2017-9827 and the National Science Foundation (USA) under Grant No. IIS-1717997.

REFERENCES

[1] Aizawa, Akiko, et al. "NTCIR-11 Math-2 Task Overview." NTCIR. Vol. 11. 2014.
[2] Bailey, Peter, et al. "Mining historic query trails to label long and rare search engine queries." ACM Transactions on the Web (TWEB) 4.4 (2010): 15.
[3] Bendersky, Michael, and W. Bruce Croft. "Analysis of long queries in a large scale search log." Proceedings of the 2009 workshop on Web Search Click Data. ACM, 2009.
[4] Boldi, Paolo, et al. "Query reformulation mining: models, patterns, and applications." Information retrieval 14.3 (2011): 257-289.
[5] Broder, Andrei. "A taxonomy of web search." ACM Sigir forum. Vol. 36. No. 2. ACM, 2002.
[6] Bruza, Peter D., and Simon Dennis. "Query reformulation on the internet: Empirical data and the hyperindex search engine." Computer-Assisted Information Searching on Internet. LE CENTRE DE HAUTES ETUDES INTERNATIONALES D'INFORMATIQUE DOCUMENTAIRE, 1997.
[7] Carbonell, Jaime G., and Jade Goldstein. "The Use of MMR and Diversity-Based Reranking for Reodering Documents and Producing Summaries." (1998).

[8] Davila, Kenny, and Richard Zanibbi. "Visual Search Engine for Handwritten and Typeset Math in Lecture Videos and LATEX Notes." 2018 16th International Conference on Frontiers in Handwriting Recognition (ICFHR). IEEE, 2018.

[9] Di Buccio, Emanuele, Massimo Melucci, and Federica Moro. "Detecting verbose queries and improving information retrieval." Information Processing & Management 50.2 (2014): 342-360.

[10] Dou, Zhicheng, Ruihua Song, and Ji-Rong Wen. "A large-scale evaluation and analysis of personalized search strategies." Proceedings of the 16th international conference on World Wide Web. ACM, 2007.

[11] Fleiss, Joseph L. "Measuring nominal scale agreement among many raters." Psychological bulletin 76.5 (1971): 378.

[12] Gao, Liangcai, et al. "Preliminary Exploration of Formula Embedding for Mathematical Information Retrieval: can mathematical formulae be embedded like a natural language?." arXiv preprint arXiv:1707.05154 (2017).

[13] Guan, Dongyi, Sicong Zhang, and Hui Yang. "Utilizing query change for session search." Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval. ACM, 2013.

[14] Guidi, F., & Coen, C. S. (2016). A survey on retrieval of mathematical knowledge. Mathematics in Computer Science, 10(4), 409-427.

[15] Huang, Jeff, and Efthimis N. Efthimiadis. "Analyzing and evaluating query reformulation strategies in web search logs." Proceedings of the 18th ACM conference on Information and knowledge management. ACM, 2009.

[16] Jansen, Bernard J., Danielle L. Booth, and Amanda Spink. "Determining the user intent of web search engine queries." Proceedings of the 16th international conference on World Wide Web. ACM, 2007.

[17] Jiang, Jyun-Yu, et al. "Learning user reformulation behavior for query auto-completion." Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval. ACM, 2014.

[18] Kamali, Shahab, and Frank Wm Tompa. "Improving mathematics retrieval." Towards a Digital Mathematics Library. Grand Bend, Ontario, Canada, July 8-9th, 2009 (2009): 37-48.

[19] Kim, Youngho, et al. "Modeling dwell time to predict click-level satisfaction." Proceedings of the 7th ACM international conference on Web search and data mining. ACM, 2014.

[20] Kristianto, Giovanni Yoko, Goran Topic, and Akiko Aizawa. "MCAT Math Retrieval System for NTCIR-12 MathIR Task." NTCIR. 2016.

[21] Kumar, P. Pavan, Arun Agarwal, and Chakravarthy Bhagvati. "A structure based approach for mathematical expression retrieval." International Workshop on Multi-disciplinary Trends in Artificial Intelligence. Springer, Berlin, Heidelberg, 2012.

[22] Kumaran, Giridhar, and Vitor R. Carvalho. "Reducing long queries using query quality predictors." Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval. ACM, 2009.

[23] Li, Lin, et al. "Enhancing Keyword Suggestion of Web Search by Leveraging Microblog Data." J. Web Eng. 15.3&4 (2016): 181-202.

[24] Li, Xinyi, and Maarten de Rijke. "Do topic shift and query reformulation patterns correlate in academic search?." European Conference on Information Retrieval. Springer, Cham, 2017.

[25] Li, Xinyi, Bob JA Schijvenaars, and Maarten de Rijke. "Investigating queries and search failures in academic search." Information processing & management 53.3 (2017): 666-683.

[26] Lin, X., Gao, L., Hu, X., Tang, Z., Xiao, Y., & Liu, X. (2014). A mathematics retrieval system for formulae in layout presentations. In Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval (pp. 697-706). ACM.

[27] Liu, Qiaoling, et al. "When web search fails, searchers become askers: understanding the transition." Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval. ACM, 2012.

[28] Ma, Chao. "Methods and systems for query segmentation in a search." U.S. Patent Application No. 10/061,844.

[29] Mansouri, Behrooz, et al. "Online job search: Study of users' search behavior using search engine query logs." The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. ACM, 2018.

[30] Mansouri, Behrooz, et al. "Understanding User's Search Behavior towards Spiky Events." Companion of the The Web Conference 2018 on The Web Conference 2018. International World Wide Web Conferences Steering Committee, 2018.

[31] Mitra, Bhaskar. "Exploring session context using distributed representations of queries and reformulations." Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval. ACM, 2015.

[32] Mitra, Bhaskar, and Nick Craswell. "Query auto-completion for rare prefixes." Proceedings of the 24th ACM international on conference on information and knowledge management. ACM, 2015.

[33] Nguyen, Tam T., Kuiyu Chang, and Siu Cheung Hui. "A math-aware search engine for math question answering system." Proceedings of the 21st ACM international conference on Information and knowledge management. ACM, 2012.

[34] Niu, Xi, and Bradley M. Hemminger. "A study of factors that affect the information-seeking behavior of academic scientists." Journal of the American Society for Information Science and Technology 63.2 (2012): 336-353.

[35] Paik, Jiaul H. "A novel TF-IDF weighting scheme for effective ranking." Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval. ACM, 2013.

[36] Rieh, Soo Young. "Analysis of multiple query reformulations on the web: The interactive information retrieval context." Information Processing & Management 42.3 (2006): 751-768.

[37] Schubotz, Moritz, et al. "Evaluating and improving the extraction of mathematical identifier definitions." International Conference of the Cross-Language Evaluation Forum for European Languages. Springer, Cham, 2017.

[38] Sojka, Petr, and Martin Líška. "Indexing and searching mathematics in digital libraries." International Conference on Intelligent Computer Mathematics. Springer, Berlin, Heidelberg, 2011.

[39] Soldaini, Luca, Andrew Yates, and Nazli Goharian. "Learning to reformulate long queries for clinical decision support." Journal of the Association for Information Science and Technology 68.11 (2017): 2602-2619.

[40] Song, Yang, and Li-wei He. "Optimal rare query suggestion with implicit user feedback." Proceedings of the 19th international conference on World wide web. ACM, 2010.

[41] Sordoni, Alessandro, et al. "A hierarchical recurrent encoder-decoder for generative context-aware query suggestion." Proceedings of the 24th ACM International on Conference on Information and Knowledge Management. ACM, 2015.

[42] Spina, Damiano, et al. "Understanding user behavior in job and talent search: an initial investigation." Proc. SIGIR Wrkshp. eCommerce (2017).

[43] Tan, Ming, et al. "Improved representation learning for question answer matching." Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Vol. 1. 2016.

[44] Tausczik, Yla R., Aniket Kittur, and Robert E. Kraut. "Collaborative problem solving: A study of mathoverflow." Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing. ACM, 2014.

[45] Tomuro, Noriko. "Question terminology and representation for question type classification." Terminology. International Journal of Theoretical and Applied Issues in Specialized Communication 10.1 (2004): 153-168.

[46] Wang, Lei, et al. "Mathdqn: Solving arithmetic word problems via deep reinforcement learning." Thirty-Second AAAI Conference on Artificial Intelligence. 2018.

[47] Wang, Yan, Xiaojiang Liu, and Shuming Shi. "Deep neural solver for math word problems." Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing. 2017.

[48] Wangari, Keita Del Valle, Richard Zanibbi, and Anurag Agarwal. "Discovering real-world use cases for a multimodal math search interface." Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval. ACM, 2014.

[49] White, Ryen W., and Resa A. Roth. "Exploratory search: Beyond the query-response paradigm." Synthesis lectures on information concepts, retrieval, and services 1.1 (2009): 1-98.

[50] Xu, Huijuan, and Kate Saenko. "Ask, attend and answer: Exploring question-guided spatial attention for visual question answering." European Conference on Computer Vision. Springer, Cham, 2016.

[51] Yih, Wen-tau, Xiaodong He, and Christopher Meek. "Semantic parsing for single-relation question answering." Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics. Vol. 2. 2014.

[52] Youssef, Abdou, and Bruce R. Miller. "Deep Learning for Math Knowledge Processing." International Conference on Intelligent Computer Mathematics. Springer, Cham, 2018.

[53] Zahedi, Mohammad Sadegh, et al. "How questions are posed to a search engine? An empiricial analysis of question queries in a large scale Persian search engine log." 2017 3th International Conference on Web Research (ICWR). IEEE, 2017.

[54] Zanibbi, Richard, et al. "NTCIR-12 MathIR Task Overview." NTCIR. 2016.

[55] Zanibbi, Richard, and Dorothea Blostein. "Recognition and retrieval of mathematical expressions." International Journal on Document Analysis and Recognition (IJDAR) 15.4 (2012): 331-357.

[56] Zanibbi, Richard, et al. "Multi-stage math formula search: Using appearance-based similarity metrics at scale." Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval. ACM, 2016.

[57] Zhao, Jin, Min-Yen Kan, and Yin Leng Theng. "Math information retrieval: user requirements and prototype implementation." Proceedings of the 8th ACM/IEEE-CS joint conference on Digital libraries. ACM, 2008.

[58] National Research Council. Developing a 21st century global library for mathematics research. National Academies Press, 2014.