# Context-Aware Search Principles in Automated Learning Environments

by

Eric Van Egmond

A Thesis Presented in Partial Fulfillment
of the Requirements for the Degree
Master of Science

Approved November 2014 by the
Graduate Supervisory Committee:

Winslow Burleson, Co-Chair
Violet Syrotiuk, Co-Chair
Brian Nelson

ARIZONA STATE UNIVERSITY

December 2014

# ABSTRACT

Many web search improvements have been developed since the advent of the modern search engine, but one underrepresented area is the application of specific customizations to search results for educational web sites. In order to address this issue and improve the relevance of search results in automated learning environments, this work has integrated context-aware search principles with applications of preference based re-ranking and query modifications. This research investigates several aspects of context-aware search principles, specifically context-sensitive and preference based re-ranking of results which take user inputs as to their preferred content, and combines this with search query modifications which automatically search for a variety of modified terms based on the given search query, integrating these results into the overall re-ranking for the context. The result of this work is a novel web search algorithm which could be applied to any online learning environment attempting to collect relevant resources for learning about a given topic. The algorithm has been evaluated through user studies comparing traditional search results to the context-aware results returned through the algorithm for a given topic. These studies explore how this integration of methods could provide improved relevance in the search results returned when compared against other modern search engines.

DEDICATION

To Kasey, for all the love and support you give me.

TABLE OF CONTENTS

Page

LIST OF FIGURES

# CHAPTER 1

# INTRODUCTION

## 1.1 Purpose

Whether one is simply communicating with another person, attempting to convey information through text, or trying to collect new information on a given subject, context can be quite important. If one were to try to ascertain the meaning of words connected in a given statement without looking at the context in which the words were connected, major misunderstandings could occur. In discussing “cranes” in construction or “cranes”, the long-necked, long-legged bird, the context of the word’s usage can be very revealing. Similarly, if commenting on a “sweet date”, one could become quite confused between a companion who was very kind and a piece of fruit that tasted sugary if not clarified by the context. In the same way, if the developers of a search engine attempt to define the meaning of a word or a set of words without taking into account the context in which those words were used, one cannot expect the engine to provide the most accurate results in many situations. In order to account for this, researchers in the area of context-aware search have attempted to address the idea of accounting for the environment in which a web search is performed [1].

The area of web search is one of the most far reaching and impactful subject areas involving the internet today. As the vast network of resources on the internet drastically expands, the job of parsing the massive amounts of data and retrieving it for consumption has fallen to search algorithms and the services they enable. Providers such as Google, Bing, and Yahoo all leverage various search techniques in order to index and retrieve as much of the web as possible. When a user then wants to find information about a given topic on the web, these algorithms parse both the search query provided and its relation to all the pages of which the search engine is aware. The developers of these search algorithms attempt to achieve the enormous task of understanding the intent of the search terms provided by the user, with an end goal of providing a relevant and helpful set of information to the user, where relevant information is defined as the information that the user needs and is searching for. In order to accomplish this, the designers of search algorithms must determine what information on the web is the most pertinent to the given search and return these results to the user in order of importance and relevance to the search.

The concept then emerges of setting the ranking of the pages found to have a relation to the given search. Generally, a user performing a search is seeking a set of information which will help them progress in their knowledge or understanding of the area of the search. Whether the user simply is looking for recent news, sports scores, or research papers on a given topic, the intent of the user is always to find a set of information. The purpose of the search algorithm then is to return a set of data that is as close as possible to the desired set of information. By ranking the pages known to the search algorithm in relation to a given search, the developer of the search engine hopes to provide the most relevant information first for the given query.

However, there are many possibilities for what information is the most relevant to a given person, and therefore generalized search engines tend to rank the pages returned based on their general importance and relevance on the web by evaluating the page's relevance to the given query through the number and placement of the appearances of the query in the page, and by evaluating the number of external pages linking to the result page to quantify the page's general importance [2]. These algorithms tend not to consider the exact preferences of a given user, since these preferences are not quantified in this context [1]. By focusing on the general relevance of the information to the world of the web at large, search engines can generically provide the most important information to the broadest number of users. Adaptations such as the Google Scholar portion of Google's search web site [3] have taken the approach of limiting the type of results returned by the search to specific content types, such as research papers and patents, in order to provide more relevant results for a certain set of users. However, this approach succeeds only through the exclusion of many other possible types of content that could relate to the context of the search and to the topic area. In this way, search engines such as Google Scholar do not attempt to take into account the context of the search, but instead simply limit the data set of the search.

In this way web search results have been advancing and improving since the invention of search engines through research into the ranking techniques applied to a given set of pages. This work has most often focused on generalized web search in which the search query is entered directly through a search engine's web page. These search results only progress out of the general context of the search engine's page and do not attempt to quantify the specific contexts through which the search is being performed [1]. Therefore, these standard search algorithms do not have any way to rank the pages returned beyond ranking the results based on their relevance to the search terms themselves and the determined importance of the results in the general context of the web [2]. By taking into account both the importance of the page relative to the search terms and the importance of the page relative to the scope of information on the web in a general sense, the algorithm can provide a broad audience with a broadly relevant set of data [2]. This attempt to ascertain the general relevance of information or pages on the web in comparison to the rest of the information on the internet through tracking characteristics such as backlinks and user bounce rates has led to improvements in the results returned by general search engines [4].

However, many contexts could benefit from improved search results by taking the details of that context into account when a search is being executed [1]. By taking into account the specific attributes of the given context and the desires of the users in that context, the relevance of the results returned for the search performed could be improved [1]. The area of context-aware search has arisen from this need in a generalized sense, but specific contexts have not received as great a level of attention and research [5, 6]. As such, context-aware search principles could potentially improve on these types of searches by including many types of content on the web while still taking into account the desires of the users in a specified context.

The area of automated learning environments, such as web sites where assistive technology helps the user to collect and learn from resources available on the internet, could benefit from enhanced search algorithms customized specifically for use in this area, as this would give the users of these sites access to more relevant resources [6]. In order to collect and curate educational content on the web, novice users must be enabled to access the most relevant content to a given subject. Since novice users may not have expertise in a given subject, they will not necessarily know where to find resources at higher levels of difficulty than they are accustomed to for a particular subject. As the goal of sites such as the Inventor’s Workshops [7] currently in development under the guidance of Dr. Winslow Burleson and Dr. Cecil Lozano is to provide modules of highly relevant content that can assist users at all experience levels in the given subject learn and progress in the topic, the assistance of a context-aware search algorithm can be beneficial to the curation of an automated learning environment by providing better access to relevant material.

In order to begin to address the lack of application of context-aware search improvements to an educational context, the work in this research has been towards the development of a context-aware search algorithm designed for assistance in content curation in an automated learning environment. By re-ranking search results based on context and filtering out content unrelated to the learning context, this research hopes to improve the process of collecting and curating educational materials on a given topic and enhance the ability of standard users to advance in science, technology, engineering, and mathematics (STEM) learning as well as passion-based learning areas. This research investigates several aspects of context-aware search principles, specifically context-sensitive and preference based re-ranking of results. These techniques take user inputs as to their preferred content types and combines this with search query modifications which automatically search for a variety of modified terms based on the given search query, combining these results into the overall re-ranking for the context. In this way, this work relates these search concepts specifically to automated learning environments, web sites where technology assists users in collecting and learning from resources, creating an integration of methods in which little previous research has been performed in this application area.

## 1.2 Scope

With the advancement of STEM learning in the field of education, various technological solutions are under development to assist in the progress of the 21st century learning model, through which education and technology combine to provide better access to passion-based and project-based learning environments [8]. These initiatives encourage users to explore subjects they are passionate about and to advance their knowledge in the area by working through projects that improve their understanding of the subject material [9]. Many initiatives, such as the IW website [7], work to develop this area in a meaningful and impactful manner in an attempt to promote a passion-based and project-based learning environment, through which learners in any subject can advance their knowledge on that subject by progressing through curated modules of content on the topic while interacting with an online community to support their efforts through mentorship and guidance.

Websites such as the Inventor's Workshops [7] accomplish this goal through the curation of learning content, bringing together various forms of information on a given subject at all difficulty levels in the topic in a cohesive manner. The content which could aid in learning can range from introductory materials, such as topic overviews or Wikipedia articles, to instructional content, such as Khan Academy lectures or tutorials such as those from Instructables.com, to advanced work in the field, such as research papers and journal articles. However, a standard web search would bring back all results on the given topic, not just those related to learning and advancing in the topic. The principles of context-aware search could be applied in this case by modifying the given query and then re-ranking all the results returned by giving additional weight to the desired content types in this context such as the overviews, tutorials, or research papers residing in the set of content.

In order to promote learning at all depths of knowledge, many various forms of content are intended to be included in a learning environment, and by giving these content types additional weight, the quality of the results returned can be improved. The collection and curation of learning content such as this is currently often performed by higher knowledge level users or experts in the field due to the lack of specific search tools to aid regular or novice users in collection of relevant content on a given topic. However, by using context-aware search principles this work hopes to give any user the ability to help collect and curate material on any topic that interests them. By creating this avenue for all users to advance the creation of educational content, this work hopes to promote passion-based learning at all levels of knowledge.

Through the combination of context-aware search principles, preference based user inputs, and query modifications, this work creates a novel approach to web search in learning environments. By taking into account the specific environment in which the given search is being performed and modifying the algorithm to fit this context's needs, this work is able to return more applicable and varied search results for educational environments. This work therefore attempts to improve the methods in which search results are returned in this context and hopes to assist in the advancement of automated learning environments throughout all subject areas.

## 1.3 Contribution

This research began by investigating the principles of passion-based and project-based learning as they are applicable here. The initial work was to take into account the approaches to the development of various learning processes [8, 10] and the application of web technologies in educational settings. From the advent of the web, educational technologies have been advancing and providing greater access to knowledge than ever before. Initiatives such as MIT's OpenCourseWare [11] and others like this have given access to educational materials in a way not previously possible. These initiatives have primarily focused on academic subjects, providing access to the content of university courses and content of other similar types. In order to continue this trend beyond academic materials and into all educational contexts, including all learning content which a person is passionate to discover, the technologies of the web must continue providing better and more efficient access to the wide spectrum of content currently available.

By first investigating how these technological advancements can lead to improvements in learning and educational environments [7, 12], this work developed new search technologies which can improve the retrieval and access of content available on the web. From this foundation this work progressed in its primary research goal of creating a search algorithm that ranks search results by integrating user preferences and ranking based on the search context while including additional results from the modifications of the search query. Through these methods this work adjusts the final calculation of the search result rankings in order to improve the quality of the content returned as reported by the users of the algorithm.

In pursuit of these advances, the questions this work attempts to answer are the following: when context-aware search principles are applied to an automated learning environment, can the perceived quality of the results be improved? In addition, can context-aware searches, preference based rankings, and query modifications improve the results of searches conducted for content collection and curation in learning websites, such that the results returned give a user more relevant results in a more efficient manner? In order to answer these questions this work has developed a context-aware search algorithm and has evaluated it with an exploratory user study in which the participants had a varied level of expertise in the given topic. This research used timed collection and result relevance rankings collected from test users to assess the progress the search algorithm has made towards improving the search results in this area.

The primary hypothesis tested through this research is the following: by applying context-aware search principles, user preference based re-ranking, and query modifications, the quality and applicability of educational search results returned is improved over generalized search engines as measured by user feedback regarding the relevance of the results returned. As the primary goal of this information retrieval system is to provide the most relevant information in the most efficient manner possible, the primary method of evaluating the success of this work is through the assessment of users from various experience levels attempting to collect information on a given topic. By collecting data as to the user's perceived relevance of results while the user is collecting materials on a given subject, this work hopes to explore the improvement that this system provides over the current generalized search engines.

In the investigation of the improvement this search algorithm provides to the web search landscape, the primary measurement used to assess the results of the work is the relevance ratings given by users to the results returned. The variable of the perceived relevance and usefulness of the results seen through user feedback gives a generally accepted measure by which to judge the results of the improvement in context-aware search explored here [13]. By answering these research questions and testing the given hypothesis, this research hopes to show a basis through which the search algorithm developed can provide improvement in to passion-based and project-based learning environments, in order to advance STEM learning and the educational landscape on the web.

In order to answer the questions proposed here, this work has conducted research into the area of context-aware search principles [14, 15, 16], with a focus on application to learning environments. The primary area of this research involves the computation of web search result rankings through an integration of methods which modifies the rank of various search results based on the needs of the given environment and the preferences generated by users dynamically during their interaction with the system. This is accomplished by taking user inputs as to the content desired through keyword input and the topic of the desired content through the search query itself. The results are then ranked based on static context attributes. This work extends the basic ideas of computing search result rankings [4], and incorporates the various techniques discussed here that improve upon these basic principles.

In order to further enhance the advancement of the searches being performed, this work has investigated search result improvements through query modifications [17] in which the query submitted by the user is modified into several related forms by adding clarifying phrases to the query to improve the accuracy of the search results. This further works has been included in an effort to ascertain the applicability of this technique to search algorithms used in learning environments similar to IW [7] and to integrate these findings into the test results. By adding phrases to the query based on the context of the website and modifying the query into similar alternate forms, this work creates a more enhanced and complete context-sensitive system in the resulting search algorithm, through which the general applicability of the conclusions of this research can be better evaluated. This also hopes to enhance the context-sensitivity of the resulting algorithm, adding to the improvement of the search results returned.

Through this set of techniques, this work creates a new, integrated approach to the improvement of relevance in search results for a given context. The improvement and integration of these techniques ultimately results in better collection of educational content for automated learning environments through which a broad range of people can advance in subjects they are passionate about. By increasing access to this content, this research hopes to encourage greater education as a whole and to promote life-long learning in any area in which a person has interests. From learning technical skills, to achieving new artistic endeavors, to building communities and promoting events surrounding technical challenges, this work hopes to achieve broader access and improved quality in content for all subjects in which users of the internet have a passion for accomplishment. This advancement of automated learning environments can continue to open the door for future generations to pursue the education that they deserve.

# CHAPTER 2

# BACKGROUND

## 2.1 Primary Theories

The idea of the relevance of a web page is a fairly relative concept, as the relevance of a search result is defined in this work as the amount to which a page provides beneficial content in the pursuit of gaining understanding of a given subject. Under this definition, the relevance of a page is dependent on how much that page provides information that advances a user's understanding of a topic. Pages that might be extremely relevant to a user performing a given search may have no relevance whatsoever to another user performing a similar search. For example, a person searching for "robotics" might be looking to get involved in a robotics club or other events in their area and therefore pages containing this information would be highly relevant. However, another person might search for "robotics" to try and find robotics kits such as the Lego kits for building introductory robots, and this person would want pages where they can read about and buy such kits. The basis of this concept of importance is that a page does not contain content inherently relevant to everyone at all times. Instead, any given page contains material that is relevant to a specific subset of people who are searching for that type of information about the given topic. From articles, to events, to biographical pages, to tutorials, to online auction results, the material contained in a page is only relevant to a user who is attempting to find and understand that subset of information on the given topic.

However, once a subject is defined by the user in the form of a search query, the idea of a page's relevance can be explored. Simply from the search term used, pages can now be analyzed for how many times that search term appears. This in turn leads to an initial understanding of how relevant that page is to the user performing the search. Other factors can then begin to be considered in calculating the relevance of the given page. Many theories of web search relevance have their roots in the original PageRank algorithm developed by Larry Page and Sergey Brin, which helped lay the foundations for calculating the importance of a given page based on the search performed and the number of links from other pages pointing to the page being evaluated [18]. By essentially evaluating the popularity of a given page, a form of relevance can begin to be expressed.

As search algorithms have advanced, many other factors have begun to be considered when evaluating the generic rank of a page for a given search, such as meta tags, the quantity of matching phrases, and the completeness of site mapping information on the site containing the given page. Consequently, most of these factors assume a generic context for the search, giving equal initial weight to all possible results. However, many searches are performed in a specific context, instead of from a generic Google or Bing search page. The specific context through which certain results are referenced allows for an environment where certain results are inherently more relevant than others. By quantifying the context of the query given, additional factors can be leveraged in the calculation of the rank of the result such as the types of content preferred in a given context [19].

This area of context-aware search has developed as the need arose for more specifically applicable search results. By taking into account the specific characteristics of a certain context in an application and re-ranking the search results based on these elements, the user can be presented with search results that are more applicable to their desired intent [19]. The context of a search is defined as the characteristics of the page from which the search is performed. For example, in a website about music, a search for "chili peppers" should return results about the band the Red Hot Chili Peppers, not the spice, chili pepper. In order to define this context, search algorithms began to include context-specific elements in their calculations for ranking search results, such as relevant content types like tutorials and lectures for educational contexts. In addition the algorithms began to include the relation of search queries to their topic areas by adding relevance to results including "music" or "band" when searching for "chili peppers" from a page for musicians. This area has shown good progress in developing searches that return more accurate results to what the specific user desires [19].

The primary work in context-aware search has been to calculate additional weighting vectors based on what the user has selected in a document or page, the area being viewed by the users, the previous queries submitted, or user preference inputs, in order to better calculate the weight of the links returned and rank them accordingly [13, 14, 20, 21, 22]. These approaches all stem from the idea that a query has various meanings depending on the context and previous user actions, and therefore the context information must be quantified in some way if the search results are to be improved [23, 24, 25]. The work in this area has taken a generalized approach that attempts to quantify any given context and the factors which this context may apply to its search results. While this general approach can be applied to the widest range of contexts, the specificity of a single context may not be fully quantified and applied.

Additionally, work has been done on the efficient computation of user preferences at run-time based on their interactions with a document [16, 26]. In this methodology the idea of the context of the search is expanded beyond the page and its given context to the specific user and the preferences that individual may have, such as specific sub-topics that user is interested in. By tracking user actions or allowing for specific user inputs to the algorithm through keywords, the search result rankings can be customized even further, allowing the relevance of the content returned to be further improved. Some of the most specific possible relevance can be obtained in this manner, as the user is directly inputting their desires in various ways, allowing the search algorithm to tailor the rankings of the results to each specific individual.

Further work has independently developed the technique of modifying the search query submitted by the user by adding to or modifying the terms entered by the user [17, 27]. These query modifications arise from a set of related words to the query based on the desired result of the user or the context of the search being performed. The techniques of query modifications include "changing terms (or making phrases), removing terms, or adding extra terms" [17]. In the end, "the goal is an internal query that is more representative of the user's intent" [17]. This technique is used to modify the search being performed, based on both the specific user and the context of the page. For example, a page focused on scholarly work might add the phrases "research" or "journal articles" to any given search term. The modified searches are then performed in addition to the original search, and all the results found are combined in the final result set. The addition or subtraction of terms can be set statically in this manner, or dynamically based off user input and actions. By modifying the search query internally, the user is not burdened with defining and applying these changes, but instead is simply given more relevant results for their intended purpose. This leads to a more complete and relevant data set returned from the search without requiring the user to have the knowledge required to tailor their search in this manner.

These types of algorithms primarily discuss a generic application of an improved search technique to any given page or document. This approach has been shown to provide positive results such as improved relevance over a wide range of contexts, but this work proposes that those results could be further improved in user defined relevance by refining these techniques to a specific subset of pages, which in this case would be the subset of educational and learning related content and websites. In addition, these techniques of context-aware search, user preference inputs, and query modifications do not appear to have been combined in an integrated approach to improving the search results in a given area. This integration of methods has not been explored in previous work and therefore is the primary focus of this thesis. This research is implemented in a context-aware search algorithm and tested in developed web pages, through which this work conducted exploratory user studies as to the relevance of the search results provided in order to investigate the potential benefits provided towards obtaining more relevant search results when compared to generalized search engines.

## 2.2 Related Work

Beyond the initial PageRank algorithm, another researcher at Stanford University, Taher H. Haveliwala, contributed significantly to the ideas that underlie the area of context-aware search. His work extended the idea of the PageRank algorithm to initially become more efficient in its computation of the ranking vectors used [28]. His work in the optimization of the PageRank calculations provided advancement in the field and assisted in allowing on-the-fly re-ranking of results in an acceptable timeframe. Since the ranking of search results is usually performed on the fly in many current search approaches, his work contributed to making this computation possible in a manner that allowed for this ranking to occur while still maintaining usefulness to the user in typical scenarios.

His work progressed further into the idea of topic-sensitive ranking, introducing some of the first work on context-sensitive ranking algorithms for web search [14]. He proposed a set of ranking vectors based not only on the query, but also on the topic and context of the page from which the user searched [14]. By adding in the initial quantification of the context from which the search is performed, Haveliwala paved the way for future work to further explore this concept. This work provided the foundation of the context-sensitive search principles used extensively in this research. He additionally compared initial techniques for personalizing the PageRank system to individual users [15], providing the basis for future techniques such as query modifications and preference-based re-ranking. This contribution began to take into account specific user preferences, again opening the door to the customization of search results beyond the basic query entered. Haveliwala's work provided a major stepping stone for the principles used in this work. By combining techniques created by Haveliwala and others, this work was able to create the context-aware search algorithm described here.

# CHAPTER 3

# CONTEXT-AWARE SEARCH ALGORITHM

## 3.1 Context-Sensitive Search

The search algorithm developed in this work is a combination of approaches in which the integration of methods hopes to lead to advancement in the area of context-aware search for educational resource collection and curation. By focusing specifically on web search for educational-related resources, this work was tailored in its approach to use the most beneficial methods for collecting resources in this context. Through the fine-tuning of the approach to this area, this work hopes to bring greater improvement to search results than a generic context-aware search can provide. As this subject area can provide vast benefits to many under-represented groups, this work hopes that this focus on educational-related searches can bring meaningful benefit to the educational community as a whole.

There are three primary methodologies which are employed in this search algorithm: context-sensitive search, preference based re-ranking, and query modifications. These methods each bring specific improvements to the relevance of the search results returned and help to accurately re-rank these results, allowing the most relevant information to the user to be returned in a more efficient manner. The combination of these approaches results in a novel search algorithm and provides benefits to this research area that few other projects have previously explored. Each of these methods have been developed separately in other research [13, 17, 19] in an attempt to better quantify and include details of the intent of a given web search beyond the query itself. By combining these approaches and leveraging the benefits of all three, this work explores the additional improvement in the re-ranking of search results that none of the individual methodologies attain on their own.

The first approach employed in this search algorithm is the primary methodology of context-sensitive web search. The purpose of context-aware search at its foundational level is to incorporate the elements of the context in which the search was performed that are beneficial to determining the desired information of the search. The goal of this approach is to quantify the preferred types of content and key tags which beneficial results contain in their title and description. Since this research specifically has focused on the context of automated learning environments, the inclusion of this approach was primarily focused on incorporating content that would be the most beneficial to an educational context.

In order to accurately assess the types of data that should be given the most weight, several elements of educational data needed to be discussed. The first element to consider was the types of content that would be relevant in this context. The immediate content types such as textbooks and online lectures provided an initial set of pages that should be given weight, but various other content types are quite relevant in this case as well. In order to maintain a varied and accessible platform for all types of learners, other types of media needed to be weighted such as videos, online courses, local events, local people in the field that one could connect with, and research and development in the topic area. The inclusion of these various types at the basic context-awareness level was implemented in the algorithm by checking the set of results with tags quantifying the educational context which have been statically created based on the keywords of content found in the modules of the IW website [7]. By parsing the title, description, and URL of the results in the result set, this work compared these elements against statically defined context tags which would quantify the educational resources desired for this context. These context tags contained the context variables desired, such as checking for Khan Academy lectures, Coursera courses, iTunes U results, and many other desirable types of content.

This initial round of quantifying the context served to give the algorithm better rankings for these initially desirable types of content in an educational context. Although the types of content favored by the algorithm are statically set in this algorithm, these content types are based in the educational context on which this work has focused. In addition, these factors could easily be changed in the algorithm in order to accommodate changing factors in a given context, or to adjust to new contexts if the algorithm were to be repurposed for other uses. This context quantification simply allows for an initial level of awareness in the algorithm which can then be used in combination with the other approaches of this research to provide improvements to the relevance of the search results returned. By quantifying the context of an automated learning environment, this work adjusts the rankings of the results and attempt improve the quality of the results returned [16].

### 3.2 Preference Based Re-Ranking

The next methodology leveraged in the search algorithm is the dynamic collection of user preferences in order to personalize the ranking of the web search results returned by the search algorithm to the individual user at runtime. This approach involves the collection of user preferences through keyword input by the user to the search algorithm. This is performed by giving the user the option at runtime to include specific tags or keywords that they feel would be helpful or relevant in the given set of search results. These user preferences can then be applied to the set of search results obtained from the initial query and to expand the number of results by including new sets of results related both to the user preferences and to the original query. This new set of results can then be re-ranked based on the preferences collected and can also be passed to the rest of the search algorithm for inclusion in the other ranking techniques included in the algorithm.

The inclusion of additional sets of results based on dynamic user preferences allows for improvement in the overall result of the search algorithm. One of the greatest limitations of re-ranking systems is the lack of variety in the original set of results from the search query. By not giving a wide enough spectrum of results to the algorithm to re-rank, any potential benefit of the re-ranking process is undermined and the benefits are limited. The addition of preference based results also gives specific customization of the algorithm on a per-user basis at runtime. By taking into account each user's specific preferences within a broad topic, the relevance of the results for that particular user is directly improved. The collection of keyword preferences from the user allows for a greater variety of results to be included than the original query contained, and allows this personalization to occur. This directly increases the effectiveness of the other ranking techniques included, as the base data set of results has more accurate and varied content set available for processing [13].

## 3.3 Query Modifications

The final technique included in the search algorithm is the concept of using query modifications to adjust the search query entered and tailor it more appropriately to the given context of the search. This technique derives from the fact that any given search query is often under-defined, especially in relation to a given context. If a query is intended only for a generic, general search, then the query by itself may be sufficient. But when a specific context, such as an automated learning environment, is the intended context for a set of search results, then there are many modifications to a basic query that could be performed in order to more accurately relate to the intended information for the given search.

Query modifications can provide improvement in a search environment through both adding variety to the data set returned and by improving the specificity of the data returned in those results. When a context, such as an automated learning environment, has specific categories of content that are more relevant to the intended result than other categories of content, a given search term can be added to, modified, or subtracted from, in order to more accurately obtain the necessary set of data in the search results. The purpose of query modifications is to change the query submitted into several related queries which contain the original idea of the search but in a more targeted and specific manner [17]. Example query modifications are shown in Figure 3.1.

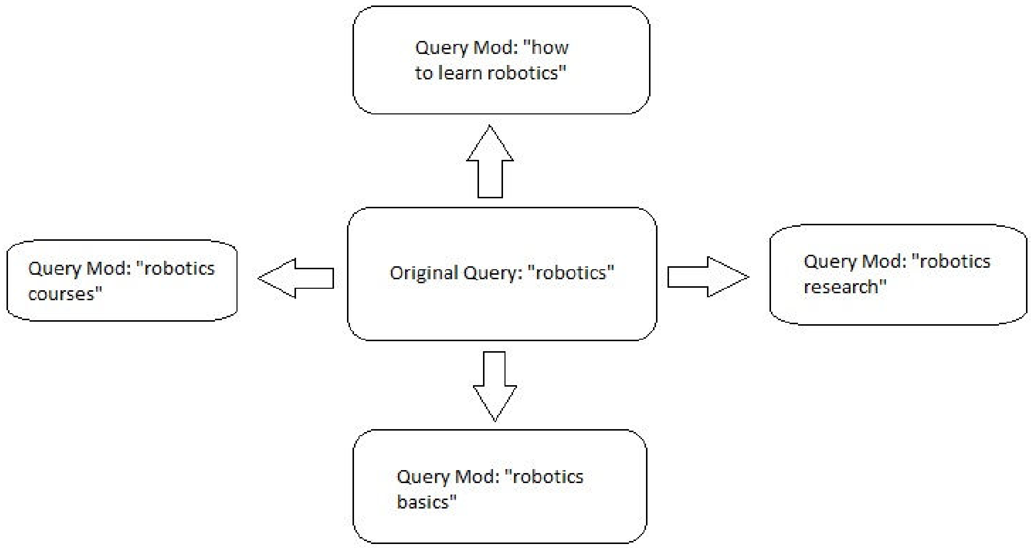


Figure 3.1: Query Modifications

Once the data sets from these modified queries have been obtained, they can be fed back into the rest of the context-aware algorithm for re-ranking. In this way, this work once again can improve the overall data set in both variety and quality.

In this context-aware search algorithm, this research has implemented query modifications by statically adjusting any given query with a pre-computed set of modifications that relate directly to learning environments and educational materials. This work uses a variety of adjustments by adding words such as “events” or “research” to a given query. These searches are then executed by the search algorithm along with the original query and all the search results returned are included in the re-ranking procedures of the algorithm. The modifications also include phrases such as “learning” and “basics”, in order to obtain more specifically relevant educational content. These query modifications have been developed for learning environments, but could be modified to fit any other potential context in which this algorithm might be used. By setting these query modifications, this work has included more specific results than the basic query would have otherwise obtained and these can again drive the overall quality of results after processing the re-ranking portion of the algorithm. Additional query phrasings and modifications could be performed on a given query if the context specified a reason for which these modifications should occur. All of the adjustments made here are used to diversify the data set sent to the final ranking portion of the algorithm, allowing for improved context-awareness in the search functionality created here.

## 3.4 Search Algorithm

The final algorithm developed in this work combines these techniques by first executing the queries used to collect the data set and then re-ranking the results returned using the context elements and keywords mentioned. The initial step of the algorithm is to perform the search using the search query provided by the user. This search is executed by calling a Bing search API and returning the data set of results to the algorithm for processing. The next step of the algorithm is to call the API to perform additional searches for the original query with the user specified preference keywords added, bring in more results to the data set and including more personalized results in the data set being created. The original query is then adjusted using the query modification techniques described through the addition of varied search terms, in order to collect an even greater data set of more varied and relevant data. For each modified query, the Bing API is again called to return each specific set of results.

After these data sets have been obtained, the results are passed to the re-ranking portion of the algorithm. This portion of the logic takes the large and varied data set collected by the previous steps and begins to parse the results of each set. The algorithm assigns a base ranking to each result based on its current positioning in its original data set. It then analyzes the result's title, description, and URL to see if the results match any of the user provided keywords. If the algorithm finds the result contained any user preferred terms, it modifies the base ranking to move the given result closer to the top of the result list. The algorithm also analyzes the details of the result to see if it matches any of the given context tags and again modifies the rank of the result to move it higher up the list if a match is found. The final step of the algorithm is to interleave and sort the results obtained from all of the various data sets, sorting them by their ranking after all processing has been completed. The flow of the algorithm is shown in Figure 3.2, which also includes the possibility of future crowdsourcing and other feedback mechanisms which could be used to directly improve the user preference re-ranking ability of the algorithm.

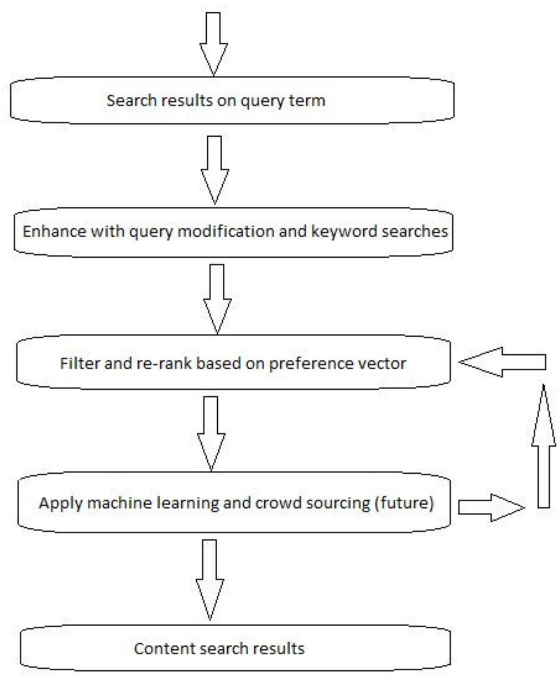


Figure 3.2: Algorithm Flow with Future Feedback Mechanisms

The pseudo code of the search algorithm is included in Figure 3.3 to further clarify the process this algorithm takes in creating the list of results provided.

```
read query from user;
read keywords from user;
read query modification terms;
read context tags;
connect to Bing search API;
execute search for query from user on search API;
foreach (keyword)
{
        execute search for query plus keyword;
        add results to new results list;
}
foreach(query modification)
{
        execute search for modified query;
        add results to new results list;
}
foreach(search result list)
{
        foreach(result in current list)
        {
                if(result not in hash table)
                {
                        set initial rank from result list;
                        foreach(context tag and keyword)
                        {
                                if(result matches)
                                {
                                        decrease rank value; //to move result up the list
                                }
                        }
                        add result to returnList;
                }
        }
}
sortByRank(returnList);
return returnList;
```

Figure 3.3: Algorithm Pseudo Code

Additionally, insight into the process of the algorithm can be seen through an example query. For this example, the query of “robotics” with the keywords “drones” and “flying” will be used as shown in Figure 3.4.

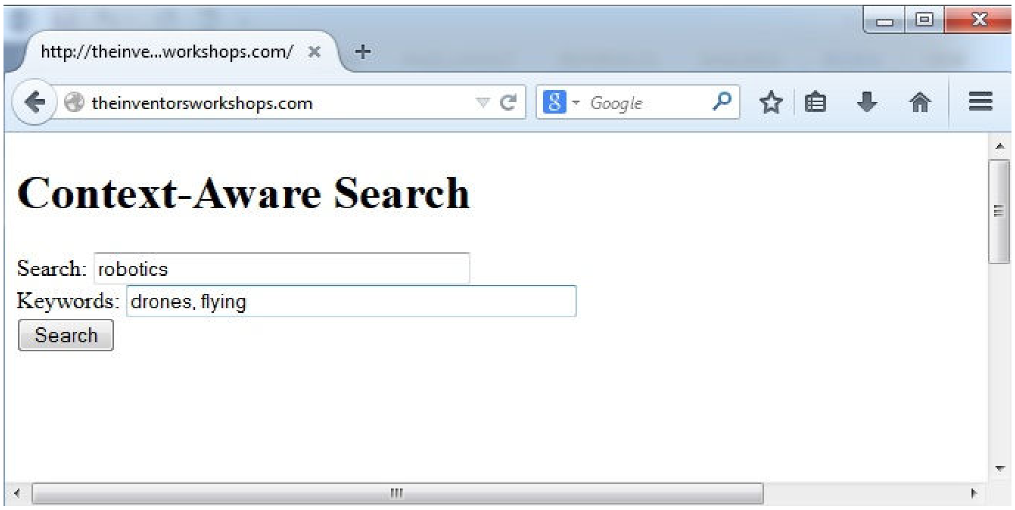


Figure 3.4: Example Search

When this search is executed, the algorithm will then begin the process of returning the context-aware search. It reads the query and keywords from the input and reads the lists of query modifications and context tags. It then performs the initial search through Bing for the search query and stores the list of results. Then for each keyword entered, the algorithm performs another Bing search for the query plus each keyword individually and stores the results of each search in a new list. Next, the algorithm performs a search for each of the modified queries created through the query modification logic and stores the results of each search in a new list. At this point, the algorithm has created a data set with more specific and varied content than the original search term returns. The algorithm then begins the interleaving and re-ranking process. Each list of results is scanned and the rank of each search result is set by starting with its initial rank in the results originally returned by the search API. This rank is then adjusted by checking the result against the keywords and context factors defined in the algorithm. When the title, description, or URL of the search result matches one of the desired tags, the rank of the result is decreased in order to move the result towards the top of the search result list. The result is also then checked against a hash table to remove any duplicate results. Once all the search result lists have been processed, all the results are interleaved and sorted based on their rank. The algorithm then returns the results to the user for viewing as shown in Figure 3.5. The GUI shown is only a test page and the search results returned could be integrated into any site as desired by the developer and is currently being integrated into the Inventor's Workshops website [7].

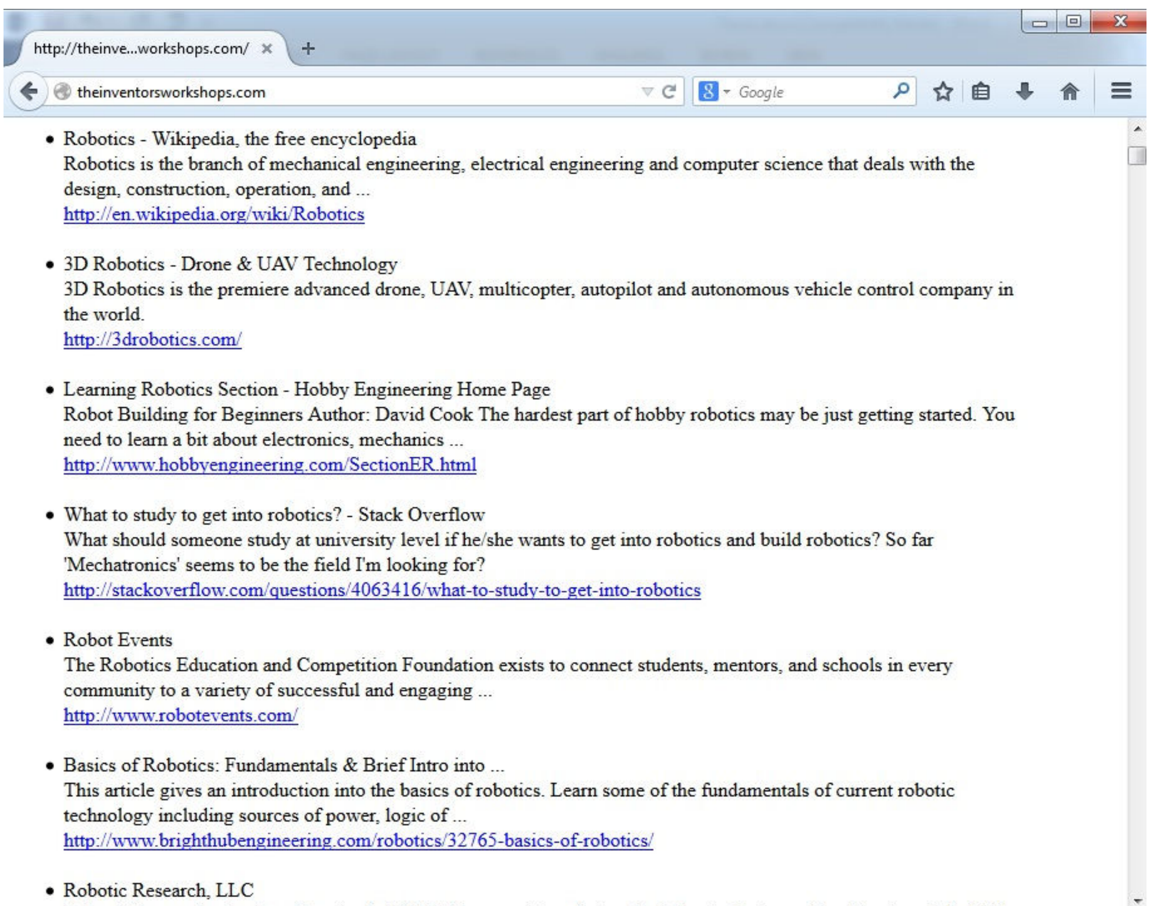


Figure 3.5: Example Search Results

Through these approaches and results, the algorithm calculates an entirely new set of rankings for the given content. The results returned hope to improve the relevance of the data set returned to the user over the results a generic search engine provides and give users at all experience levels a method to collect resources that are more relevant to an educational context. By this combination of approaches this works has created a novel context-aware search algorithm which hopes to bring relevance improvements to automated learning environments.

# CHAPTER 4

# TESTING AND RESULTS

## 4.1 Testing Methodology

In order to evaluate the improvement resulting from the context-aware search developed in this work, an exploratory user study has been conducted to compare the search results returned by this search algorithm to the search results for the same query using Google's search engine at www.google.com. In this study, the participants were asked to perform a search on the topic of robotics and to evaluate the relevance of the results returned by the search engine they were assigned to use. The definition of relevant educational content described to the participants was any content which the participant would find helpful if they were attempting to learn about the given subject. All sets of users entered the base query "robotics" and evaluated the results by recording on a form their rating of the relevance of each search result. Study participants were invited to the study from robotics groups at Arizona State and generic undergraduate computer science courses. The respondents came from a variety of expertise levels in the topic area of educational content on robotics used for this investigation, ranging from a number of completely novice users with little to no expertise in the area, to some intermediate users with moderate knowledge, and finally to a small group of expert users with substantial domain knowledge. This range was used in order to evaluate the value of the context-aware search to a general set of people who would use a site containing modules of information on a topic in which they had various levels of interest, but exact statistics on the participants backgrounds could not be recorded as no personal information was to be recorded under the protocol approved for the study.

The process used in the lab sessions for this user study was relatively straightforward. The 31 study participants were first divided into random sampling groups by counting out two out of every three of the participants for the experimental group which used the context-aware search developed in this research. The other one out of every three of the participants was placed in the control group which used Google's web search. The participants were then given result relevance forms on which they recorded whether or not the given web search result was relevant to educational content on the topic of robotics, and if the result was deemed relevant, the participant recorded their determination of the relevance on a scale of 1 to 10, with a rating of 1 being only barely relevant, to a rating of 10 being extremely relevant. If a result was deemed not relevant in any way, it was assigned a score of 0.

The lab sessions were capped at approximately 15 minutes of evaluation time with a total time cap of 30 minutes when including the time used for explaining the study, obtaining consent, and instructing the participants on the process of the study. During this time, the participants were allowed to evaluate as many links as they chose, and the number of links evaluated was recorded and averaged over each category of participants in order to report the average depth (average number of search results evaluated) for each group of participants. This was recorded in order to account for the potential decrease in relevance created for search results returned farther down the list of results by the ordering of the search results returned. In this way, the study performed evaluates the overall relevance of the content returned to the user as determined by the user performing the search. This provides a platform through which to evaluate the benefit of the techniques used by this context-aware search algorithm through search queries which model real-life usage.

## 4.2 Google Results

The first group of 11 participants was instructed to use Google's web search to attempt to find relevant educational content on the topic of robotics. The search performed by the participants using Google was "robotics" and did not include any further terms in the query so as to reflect a typical search that a user might perform when interested in learning about a given topic. In addition, if any further terms had been added to the Google search, the search results would have excluded the set of results that matched only the original query. This can only be overcome in Google by performing multiple searches or Boolean logic searches, and these techniques would consume more of the user's time than a simple keyword entry and could require the user to have advanced knowledge on how to formulate such a query in the given search engine. The participants then proceeded to evaluate the results returned as described in section 5.1. This control group of participants evaluated 277 web search results for relevance, creating an average depth of search results evaluated per user of 25.18 links. Since the relevance of the search results returned is expected to decrease as the depth to which the results are evaluated increases, the testing performed here maintained a relatively close depth between the general context-aware search results and the Google results, in order to maintain the relative relevance of the results evaluated between the control group and the experimental group.

The control group in the user study evaluated the relevance of the links returned from Google's search engine, and the individual relevance ratings for each link reported were averaged to find the overall relevance of the results returned. In this study, the participants reported a mean relevance rating of 4.79 out of 10, with 277 individual search results evaluated. In addition, by analyzing the data points at a confidence level of 95%, the confidence interval found was 0.42. This indicates that at this exploratory level the participants of the study found a moderate rating for the relevance of the results returned by Google's search engine. While this rating is not terrible when considering that Google is attempting to appeal to a broad range of audiences, it does suggest room for improvement. Through the initial results of the control group a baseline is established by which to evaluate the potential improvement in relevance gained by using the context-aware search developed here. In the following section, the relevance of the Google search results explored in this section is compared against the relevance results from the context-aware search.

### 4.3 Context-Aware Search Results

The evaluation of the experimental group testing the context-aware search algorithm developed here was performed in the same lab sessions as the control group, with the experimental group using the context-aware web search algorithm to obtain their search results. However, in order to more thoroughly test the contributions of this algorithm, the experimental group was broken into separate groups, each of which used a version of the context-aware search with specific techniques enabled. The first experimental group used the algorithm with only the context-sensitive elements enabled, which included preference-based and context specific re-ranking. The second group used the algorithm with only the query modifications technique enabled. Finally, the third and largest experimental group used the full context-aware search with all the techniques discussed in this work enabled. In this manner, the user study could more fully explore the individual benefits of each of the techniques used in this research and investigate the overall benefit of combining these techniques as performed in this work. The results from the first two experimental groups give only a partial picture of the improvements of the full search algorithm, but are beneficial for demonstrative purposes in understanding how the overall algorithm attempts to achieve its results. Thus the sample sizes for the partial algorithm groups are smaller than the group testing the full algorithm, and the results from these partial groups should not be assumed to be an indicative measure of the benefits of the partial algorithm but instead only a suggestion of the benefits provided by each technique. Larger studies would be necessary to determine more clearly the contribution of each portion of the algorithm. However, by testing each portion of the algorithm individually, the results seen in the full algorithm group are more clearly displayed and understandable.

In the first experimental group, only the preference based re-ranking and context specific factors were enabled in the algorithm. The participants of the study used the same testing structure and were not informed that they were only testing the context-aware search algorithm with some and not all of the techniques enabled. In this way, the results were not biased or influenced by this information. The experimental group here contained 4 users who evaluated 69 total search results, giving an average depth per user of 17.25 links. This depth is less than the control group and therefore the results are not directly equivalent, but as this experimental group is only used for demonstrative purposes and not quantitative judgments as to the effectiveness of the technique, the results still provide interesting information for general comparison. In this group's evaluation of the search results returned, the participants found an average relevance of 5.52 out of 10 for the 69 links assessed. This shows a small increase of approximately 15.2% in the relevance of the results returned from this partial context-aware search over the baseline Google search results.

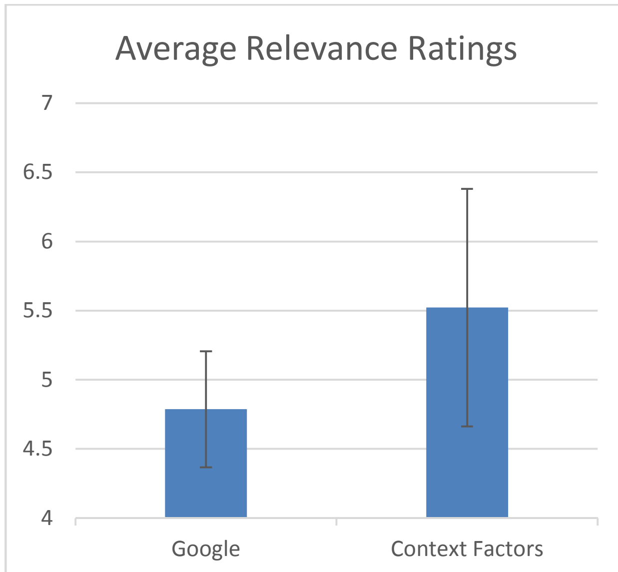


Figure 4.1: Relevance Ratings for Context Factors

Figure 4.1 shows the comparison between the average relevance for the control group results and the results for this context factor portion of the search algorithm evaluated by the first experimental group. In addition, the figure includes the confidence interval of 0.86 found by analyzing the context factors data points at a 95% confidence level and the 0.42 confidence interval found with the Google results. These values do show a small overlap indicating the potential at this confidence level that, given the possible error in the context factor results, the increase in the relevance noted is encouraging but should be viewed only as exploratory. A larger study group for the context factors would be necessary to determine more accurately the individual benefit of this technique.

The second experimental group evaluated the results returned by the search algorithm with only query modifications enabled. As with the first experimental group, the participants were not informed that they were only testing one element of the context-aware search algorithm's techniques. This experimental group contained 5 participants with a total number of 95 search results evaluated, creating an average depth per user of 19.00 links. This experimental group was closer in depth to the Google control group, allowing for a better comparison between this technique and the baseline results. However, the depth still differs to a degree that the results are only demonstrative and not a direct comparison. In this group, the users of the partial context-aware search algorithm recorded an overall average relevance rating of 5.78 out of 10 for the 95 results evaluated with a confidence interval of 0.70 at a confidence level of 95%. This result suggested a slightly larger improvement over the context specific factors and a potential improvement over the baseline Google search results of approximately 20.7%. This improvement over the baseline Google results is shown in Figure 4.2.

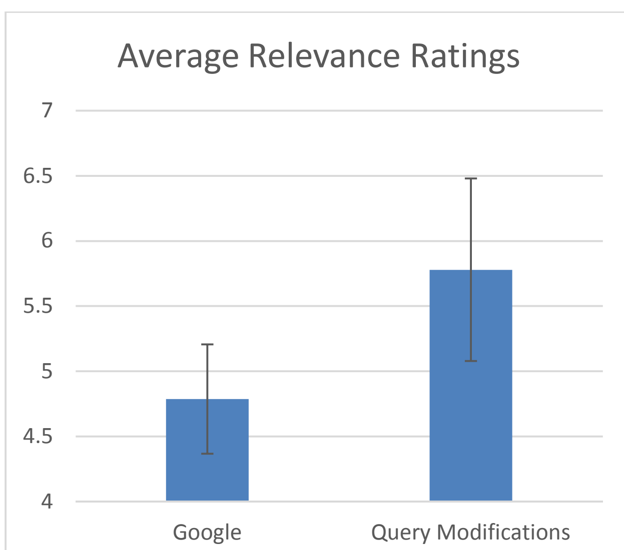


Figure 4.2: Relevance Ratings for Query Modifications

In Figure 4.2, the potential increase of relevance over the baseline Google search results is shown with the confidence intervals included. In addition, after the slight improvement in relevance and the slightly tighter confidence interval, the results showed only very slight overlap at 95% confidence. The evaluation of the query modifications portion of the context-aware search algorithm suggests an increase in relevance similar to the increase in relevance from the context factors portion of the search algorithm, but the decrease in overlap of the confidence intervals does suggest slightly more strongly that this technique provides an improvement over the baseline Google results.

In the third and final experimental group, participants of the study used the complete context-aware search algorithm developed through this research. This algorithm includes all the factors of context-aware search with preference based re-ranking and query modifications. This experimental group evaluated the total relevance change available in the complete context-aware search algorithm and shows the overall benefit of the work performed here. The group contained 11 participants who evaluated a total of 248 search results. This resulted in an average depth of per user of 22.55 links, with a difference of depth between this group and the baseline group of only approximately 2.6 links per user. This experimental group is more directly comparable to the baseline Google search results, as the number of results evaluated and the number of participants in the group is fairly closely aligned. The data from this experimental group recorded an average overall relevance rating of 6.16 out of 10 over the 248 search results evaluated with a confidence interval of just 0.38 at a 95% confidence level. Therefore, the total improvement in relevance of the full context-aware search algorithm suggested by this study is approximately 27.0% over the control group's Google search relevance data.

This result suggests that the full context-aware search provides a potential improvement in overall relevance of the search results returned over the results returned from a similar Google search on the given topic of robotics. Figure 4.3 shows the relevance data from all four study groups including their confidence intervals.

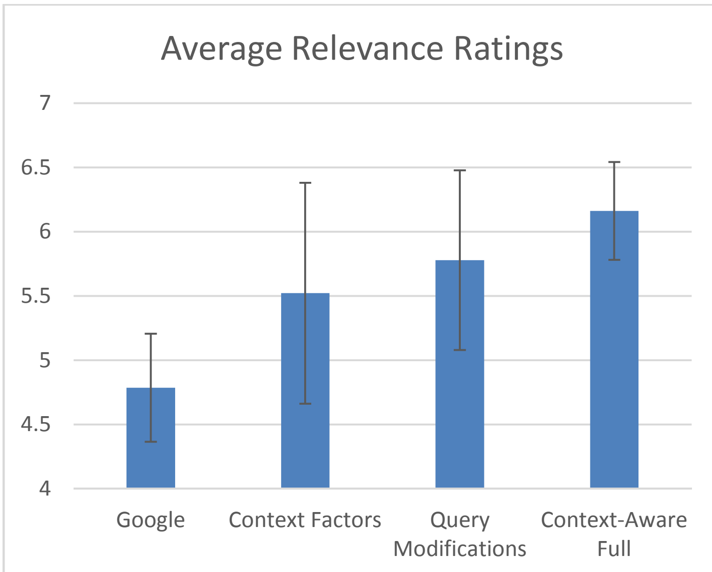


Figure 4.3: Relevance Ratings for Full Context-Aware Search

This comparison shows that the combination of both the context factor techniques and query modifications does suggest an improvement in relevance over either individual technique on its own, and that the overall context-aware search provides potentially increased relevance in the search results returned for educational content as defined in this user study as shown by the increased relevance recorded and lack of overlap of the confidence intervals at a 95% confidence level. Although these results should be viewed as exploratory due to the size of the study conducted, the data shown is quite encouraging and could be an interesting area of further study.

# CHAPTER 5

# DISCUSSION

## 5.1 Analysis

In order to more completely investigate the increases in relevance suggested by the user study, the data collected for the Google results and the full context-aware search results was statistically analyzed using a single factor ANOVA (analysis of variance) test. In this test, the null hypothesis is that the two groups being evaluated do not vary significantly in their group mean values. For this test $\alpha = 0.05$ was used to indicate a significance level of 95%. The results of this test are shown in Figure 5.1.

| SUMMARY | | | | | | |
|---|---|---|---|---|---|---|
| *Groups* | *Count* | *Sum* | *Average* | *Variance* | | |
| Google | 277 | 1326 | 4.79 | 12.65 | | |
| Context-Aware Full | 248 | 1528 | 6.16 | 9.39 | | |
| | | | | | | |
| ANOVA | | | | | | |
| *Source of Variation* | *SS* | *df* | *MS* | *F* | *p-value* | *F crit* |
| Between Groups | 247.13 | 1 | 247.13 | 22.24 | 3.09E-06 | 3.86 |
| Within Groups | 5811.9816 | 523 | 11.11 | | | |
| | | | | | | |
| Total | 6059.11 | 524 | | | | |

Figure 5.1: ANOVA Analysis

The particularly interesting value found in this result is the p-value at $3.09 \times 10^{-6}$, as this value is well below the $\alpha$ value of 0.05. Since $p = 3.09 \times 10^{-6} < 0.05 = \alpha$, the ANOVA test dictates that the null hypothesis be rejected, meaning that statistically the two groups are found to not be equal in their mean relevance ratings at a 95% confidence level. Again, the study performed here was small in size and therefore the results must be characterized as exploratory only. But this statistical result for the study performed does show positive and encouraging results at this exploratory level.

The user studies conducted in this research show several interesting aspects of the techniques incorporated here. Initially the query modifications show some benefit in increasing the relevance of the search results returned by the algorithm. This element of the algorithm provides a way to expand the data set included in the search results and thereby allows for more specific results to be included without the user providing any additional input. By automatically adjusting the query entered by the user, this method provides a straightforward avenue for improving the relevance in a given context. Since these query modifications are easily adjustable and could potentially be crowd-sourced and dynamically created in the future, this technique could be used to apply to any specific context and improve the search results. The results of testing this technique suggest that query modifications are a helpful basis for expanding the data set of results returned in context-aware search.

In addition, the context factors that incorporate the context-awareness for this search algorithm also suggest moderate benefit for increased relevance. These factors let the algorithm to re-rank all the results in the data set provided, allowing this improvement to extend to the data included through the query modifications technique as well. As the user testing for the complete search algorithm suggests, the benefit of the combined techniques did continue the upward trend in relevance of results when these techniques were combined with the query modifications.

Interestingly, the total improvement did not show a linear correlation with the increases found from each partial algorithm user study. If the improvement of both techniques had combined in a linear manner, the overall algorithm data would have shown approximately 35% more relevant results over the baseline Google results.

However, this non-linear increase was somewhat expected as the most relevant results in either technique were seen to be duplicate results of the other technique in several instances. Therefore the increase in each partial algorithm's results would not be linearly combined upon the combination of the techniques into the full algorithm. However, the overall increase in the combined algorithm does indicate that there are results increased in rank in each technique that are not found in the others, and therefore the combination of these approaches does result in a more relevant set of results than any one approach on its own.

Through the user studies conducted to evaluate this work, an interesting increase in the relevance of the search results returned is explored. The increase in relevance as recorded by the participants of the user study of approximately 27.0% at a high confidence level suggests that by combining query modifications with preference based re-ranking and context-sensitive search principles, the quality of results returned by the context-aware search algorithm could potentially be improved. Through this increased relevance in results, the time and effort needed to collect educational resources in automated learning environments can be reduced by presenting the user with more relevant resources. This directly benefits the advancement and collection of online educational resources by giving even novice users the ability to collect resources on a given topic of interest more easily, with these resources ranging across all levels of difficulty in the subject. By further developing these advances and adding additional techniques to further increase overall relevance in search results, even greater improvements in this area could potentially be achieved.

## 5.2 Future Work

This research in context-aware web search has several additional avenues by which it could be further extended for even greater improvement. These enhancements and future work could further improve the relevance of the search results returned by taking into account additional context factors and other approaches to collecting user preferences. By more completely quantifying the context in which the searches are being performed, these techniques could provide additional frames of reference by which the search results returned could be improved. These techniques span related research areas of computer science which did not fall under the scope of the work being carried out here, but could be beneficial after further investigation and future research.

One technique which could yield further improvements in this area would be to dynamically obtain relevant keywords to the initial search term by lexically analyzing the content of the pages returned in the initial search query and determining the most popular terms that occur in the results found by the initial search topic [29, 30]. Since many subjects can fall under broader search terms and only certain sub-topics are usually relevant to a given user, analyzing the pages returned for potential sub-topics keywords would allow the algorithm to return these sub-topics to the user, allowing the user to adjust the ranking and weight given to pages relating to each sub-topic. In this manner, the personalization of the search to each individual user could be further enhanced. The resulting search set could then be re-ranked based on the feedback obtained from the user, allowing for more relevant results to be given additional weight. By allowing the user more specific control over the functionality of the search algorithm, the results returned could be personalized to a level not previously achieved.

In addition, the ability to present a user with a set of sub-topics would allow for greater accessibility for novice users. If the sub-topics obtained were to be returned to the user and graphically displayed in a word cloud [31] or similar representation, this would allow users with less initial knowledge of the subject area to quickly identify the various categories contained within the parameters of their initial search. They could then use the feedback mechanisms of the search algorithm to narrow the scope of the search they are performing. In this way, they would be able to collect educational content on more specific topics in their subject area, without having expert knowledge of the subject domain in advance. This would add to the accessibility of the context-aware search algorithm for users with smaller amounts of domain knowledge.

The challenge with achieving this additional customization is the efficiency with which the pages in the search results are processed and analyzed for keywords. The lexical analysis must be completed and returned in an efficient manner, which could be a challenge when hundreds or thousands of pages are being processed in a given search result set. In addition, the lexical analysis has the challenge of maintaining the accuracy of the counts of the keywords returned. The analysis must leave out language constructs, such as connectives, articles, and verbs, and only return topical keywords related to the subject of the original search. It also must recognize when various forms of a word or phrase are used to mean the same topic and to correctly count these occurrences towards the overall frequency of that sub-topic. If the accuracy of the results returned is not kept to a high enough standard, then the feedback submitted to the search algorithm will incorrectly re-rank the search results based on incorrect weighting data. Challenges such as these must be overcome in future work in order to allow this approach to provide substantial benefit to the overall re-ranking of the search result set.

Another potential area for future research in context-aware search principles would be the incorporation of machine learning techniques into the re-ranking structure of the algorithm [32, 33]. If the type of the content, the subject matter, or the level of difficulty of a page could be quantified in the search results, the algorithm could then record the behavior of the user and identify which types of content, subjects, etc., that the user is naturally gravitating towards in their current collection of resources. By tracking the resources the user is collecting, the search algorithm can then better rank the results returned in future searches. If this mechanism were added to the context-aware search algorithm, it would provide an avenue to quantify more specifically what results are the most relevant to a given user and would allow for greater user-specific personalization of the search results provided.

The primary challenge facing the addition of machine learning components to the context-aware search algorithm is again maintaining acceptable running times for the execution of the search. The classification of web pages in the search result set would need to occur dynamically in this case, as the search algorithm does not have access to the data stored by the original crawler about all the pages that might be returned. This would necessitate that the algorithm must evaluate the pages in the base result set for their various characteristics, such as content type and subject matter. There are various methodologies available to accomplish this [34, 35, 36, 37], however many of these methods have not been tested over a large scale result set or over a wide range of web pages. In addition, these methods only attempt to categorize the pages by topic and do not take into account the content type, such as text, video, or audio, which can be very beneficial to record when attempting to track and understand user preferences.

One final approach that could be applied to the context-aware search algorithm to great effect would be to implement crowdsourcing techniques which could allow users to rank many various aspects of the search results returned. If each result found by the search algorithm was stored locally in a cache of results, many revealing statistics could then be recorded about the individual results by leveraging user ratings of these aspects. Some of the potentially beneficial ratings that could be applied are the content type of the page, the difficulty level of the information contained in the page, and the overall quality of the information provided by the page. By allowing the users of the search algorithm to record their opinions on these characteristics, the algorithm could then use these ratings to categorize and rank the search results in a more comprehensive manner. By using crowdsourcing instead of various other forms of page analysis, the efficiency of the algorithm would not be reduced, while the functionality of the algorithm could potentially be improved.

Applying crowdsourcing to the search results collected by the algorithm would open a wide range of possible improvements [38]. By taking into account the rated quality of the content provided in each search result, the algorithm could filter out content that users feel to be lower quality. In addition, the classification of the pages returned could be augmented and possibly corrected by the user ratings on the content tags for the page. This could then be incorporated back into the classification engine to improve the results of future page classifications. Additionally, this avenue would be the most feasible option for recording the type of content provided in a web page. Certain content types that would be beneficial to determine, such as event listings or profiles of local experts and activists in a topic, could best be determined by crowdsourcing this information, as attempts to automate this categorization could be more inefficient than other aspects of the improvements considered here.

The greatest challenge to integrating crowdsourcing into the current search algorithm is the necessity of recording and storing all the search result data being used and rated. Currently the algorithm only retrieves results temporarily and does not store them permanently. Adding a crowdsourcing aspect to the results would require storing some amount of identifying information about the search result in order to allow for its identification in future searches and to maintain the relation of user ratings to the result. This would necessitate a much larger storage system than the algorithm uses currently and would have a direct effect on the efficiency of the algorithm by adding database calls or other storage access functions. However, these efficiency concerns might be outweighed by the efficiency savings produced by not requiring dynamic classification of the content of the search results. By only attempting to categorize the content of a search result during the first access of that result by the algorithm, all future access to that result would be made more efficient. In addition, the classification of the result should over time become more accurate as the principles of crowdsourcing come into play and the rating and categorization of the result is reviewed by each user and gradually made more accurate.

These potential improvements are areas in which future research could reveal the extent of the possible benefits to the context-aware search algorithm provided by these techniques. Concerns of efficiency and implementation would need to be overcome, but the improvement in the relevance of the search results returned could be quite interesting. By leveraging these techniques in cooperation with the methods already employed in this work, a synergistic effect might be achieved which could further improve the relevance of the results returned. This future work would continue to build on the foundation provided in this research and could achieve an even greater improvement in context-aware web search.

REFERENCES

[1] Lawrence, Steve. "Context in web search." IEEE Data Eng. Bull. 23.3 (2000): 25-32.

[2] Langville, Amy N., and Carl D. Meyer. "Deeper inside pagerank." Internet Mathematics 1.3 (2004): 335-380.

[3] “Google Scholar.” Google. Web. 14 Nov. 2014.

[4] Franceschet, Massimo. "PageRank: Standing on the shoulders of giants." Communications of the ACM 54.6 (2011): 92-101.

[5] Truong, Hong-Linh, and Schahram Dustdar. "A survey on context-aware web service systems." International Journal of Web Information Systems 5.1 (2009): 5-31.

[6] Brusilovsky, Peter, and Mark T. Maybury. "From adaptive hypermedia to the adaptive web." Communications of the ACM 45.5 (2002): 30-33.

[7] Burleson, Winslow. “IW Motivational Environments.” ASU Ira A. Fulton Schools of Engineering. Web. 9 Apr. 2014.

[8] Brown, John Seely, and Richard P. Adler. "Open education, the long tail, and learning 2.0." Educause review 43.1 (2008): 16-20.

[9] Greenhow, Christine, Beth Robelia, and Joan E. Hughes. "Learning, teaching, and scholarship in a digital age Web 2.0 and classroom research: What path should we take now?." Educational Researcher 38.4 (2009): 246-259.

[10] Karampiperis, Pythagoras, and Demetrios Sampson. "Adaptive learning resources sequencing in educational hypermedia systems." Educational Technology & Society 8.4 (2005): 128-147.

[11] “MIT OpenCourseWare.” Massachusetts Institute of Technology. Web. 14 Nov. 2014.

[12] Alexander, Bryan. "Web 2.0: A new wave of innovation for teaching and learning?." Educause review 41.2 (2006): 32.

[13] Radlinski, Filip, and Susan Dumais. "Improving personalized web search using result diversification." Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval. ACM, 2006.

[14] Haveliwala, Taher H. "Topic-sensitive pagerank: A context-sensitive ranking algorithm for web search." Knowledge and Data Engineering, IEEE Transactions on 15.4 (2003): 784-796.

[15] Haveliwala, Taher, Sepandar Kamvar, and Glen Jeh. "An analytical comparison of approaches to personalizing PageRank." Technical Report. Stanford. (2003).

[16] Chen, Liang Jeff, and Yannis Papakonstantinou. "Context-sensitive ranking for document retrieval." Proceedings of the 2011 ACM SIGMOD International Conference on Management of data. ACM, 2011.

[17] Glover, Eric J., et al. "Improving category specific web search by learning query modifications." Applications and the Internet, 2001. Proceedings. 2001 Symposium on. IEEE, 2001.

[18] Page, Lawrence, et al. "The PageRank citation ranking: Bringing order to the web." (1999).

[19] Xiang, Biao, et al. "Context-aware ranking in web search." Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval. ACM, 2010.

[20] Cao, Huanhuan, et al. "Context-aware query suggestion by mining click-through and session data." Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining. ACM, 2008.

[21] Cao, Huanhuan, et al. "Context-aware query classification." Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval. ACM, 2009.

[22] Shen, Xuehua, Bin Tan, and ChengXiang Zhai. "Context-sensitive information retrieval using implicit feedback." Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval. ACM, 2005.

[23] Agrawal, Rakesh, Ralf Rantzau, and Evimaria Terzi. "Context-sensitive ranking." Proceedings of the 2006 ACM SIGMOD international conference on Management of data. ACM, 2006.

[24] Liu, Tie-Yan. "Learning to rank for information retrieval." Foundations and Trends in Information Retrieval 3.3 (2009): 225-331.

[25] Mitra, Mandar, Amit Singhal, and Chris Buckley. "Improving automatic query expansion." Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval. ACM, 1998.

[26] Chakrabarti, Soumen. "Dynamic personalized pagerank in entity-relation graphs." Proceedings of the 16th international conference on World Wide Web. ACM, 2007.

[27] Jones, Rosie, et al. "Generating query substitutions." Proceedings of the 15th international conference on World Wide Web. ACM, 2006.

[28] Haveliwala, Taher. "Efficient computation of PageRank." Technical Report. Stanford. (1999).

[29] Matsuo, Yutaka, and Mitsuru Ishizuka. "Keyword extraction from a single document using word co-occurrence statistical information." International Journal on Artificial Intelligence Tools 13.01 (2004): 157-169.

[30] Chekuri, Chandra, et al. "Web search using automatic classification." Proceedings of the Sixth International Conference on the World Wide Web. 1997.

[31] Cui, Weiwei, et al. "Context preserving dynamic word cloud visualization." Pacific Visualization Symposium (PacificVis), 2010 IEEE. IEEE, 2010.

[32] Agichtein, Eugene, Eric Brill, and Susan Dumais. "Improving web search ranking by incorporating user behavior information." Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval. ACM, 2006.

[33] Agichtein, Eugene, et al. "Learning user interaction models for predicting web search result preferences." Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval. ACM, 2006.

[34] Shen, Dou, et al. "Web-page classification through summarization." Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval. ACM, 2004.

[35] Qi, Xiaoguang, and Brian D. Davison. "Web page classification: Features and algorithms." ACM Computing Surveys (CSUR) 41.2 (2009): 12.

[36] Kwon, Oh-Woog, and Jong-Hyeok Lee. "Web page classification based on k-nearest neighbor approach." Proceedings of the fifth international workshop on Information retrieval with Asian languages. ACM, 2000.

[37] Marath, Sathi T., et al. "Large-scale web page classification." System Sciences (HICSS), 2014 47th Hawaii International Conference on. IEEE, 2014.

[38] Alonso, Omar, Daniel E. Rose, and Benjamin Stewart. "Crowdsourcing for relevance evaluation." ACM SIGIR Forum. Vol. 42. No. 2. ACM, 2008.

# APPENDIX A

# IRB APPROVAL

ASU Knowledge Enterprise Development

EXEMPTION GRANTED

Brian Nelson
CIDSE: Computing, Informatics and Decision Systems Engineering, School of
480/965-0383
Brian.Nelson@asu.edu

Dear Brian Nelson:

On 10/17/2014 the ASU IRB reviewed the following protocol:

| | |
|---|---|
| Type of Review: | Initial Study |
| Title: | Evaluating the Improvement of Web Search Results using Context-Aware Search |
| Investigator: | Brian Nelson |
| IRB ID: | STUDY00001643 |
| Funding: | None |
| Grant Title: | None |
| Grant ID: | None |
| Documents Reviewed: | • nelsonConsent Form.docx, Category: Consent Form;<br>• HRP-503a.docx, Category: IRB Protocol;<br>• Result Relevance Form.docx, Category: Measures (Survey questions/Interview questions /interview guides/focus group questions);<br>• nelsonRecruitment (1).docx, Category: Recruitment Materials; |

The IRB determined that the protocol is considered exempt pursuant to Federal Regulations 45CFR46 (2) Tests, surveys, interviews, or observation on 10/17/2014.

In conducting this protocol you are required to follow the requirements listed in the INVESTIGATOR MANUAL (HRP-103).

Sincerely,

IRB Administrator

cc: Eric Van Egmond
Eric Van Egmond
Brian Nelson
Cecil Anahi Lozano Fortun

# APPENDIX B

# RECRUITMENT EMAIL

Dear potential participants,

I am a graduate student under the direction of Professor Brian Nelson in the Ira A. Fulton Schools of Engineering, at Arizona State University. I am conducting a research study on the benefits of context-aware web search in learning environments.

I am recruiting individuals to use a search engine for finding educational content and record whether or not that result is relevant to the given context, which will take approximately 30 minutes. You must be 18 years or older to participate.

While there may be no immediate benefits to you, your participation will help in the improvement of web search technology for specific contexts, in particular aiding in the advancement in learning environments through which anyone interested in a topic can find resources to learn and interact with others interested in that topic.

Your participation in this study is voluntary. If you are interested in participating please fill out and send me the form below. If you have questions concerning the research study, please contact me at (602) 456-1899, or evanegmo@asu.edu. Thank you very much for your time.

Thanks,
Eric Van Egmond

Name: ________________________
Email address: __________________

# APPENDIX C

# PARTICIPANT CONSENT FORM

CONSENT FORM

*Evaluating the Improvement of Web Search Results using Context-Aware Search*

I am a graduate student working under the direction of Professor Brian Nelson in the Ira A. Fulton Schools of Engineering, at Arizona State University. I am conducting a research study to evaluate the benefit of context-aware web search in learning environments. Specifically, we will evaluate the improvements made through the context-aware search algorithm we have developed, which is designed for assistance in content curation in an automated learning environment.

I am inviting your participation which will involve one session where you will be asked to use either www.google.com or our context-aware search algorithm to search on the topic of robotics for the purpose of collecting web links into a cohesive module of educational material on the subject of robotics. For each search result returned, you will record whether or not that result is relevant to the given context. Your participation will last for a maximum of 30 minutes. You have the right not to answer any questions, and to stop participation at any time.

Your participation in this study is voluntary. If you choose not to participate or to withdraw from the study at any time, there will be no penalty.

While there may be no immediate benefits to you, your participation will help in the improvement of web search technology for specific contexts, in particular aiding in the advancement in learning environments through which anyone interested in a topic can find resources to learn and interact with others interested in that topic. There are no foreseeable risks or discomforts to your participation.

All information obtained in this study is strictly confidential. The results of this research study may be used in reports, presentations, and publications, but your name will not be known.

Any questions you have concerning the research study or your participation in the study, before or after your consent, will be answered by the Eric Van Egmond, (602) 456-1899, evanegmo@asu.edu or Brian Nelson at Brian.Nelson@asu.edu. If you have questions about your rights as a subject/participant in this research, or if you feel you have been placed at risk; you can contact the Chair of the Human Subjects Institutional Review Board, through the ASU Office of Research Integrity and Assurance, at (480) 965-6788.

Your participation in this study is your consent to participate.

# APPENDIX D

# RESULT RELEVANCE FORM

For each web search result returned on the web page, please circle whether the result is relevant as educational content for the search topic.
If the result is relevant, please rate the relevance of the result on a scale from 1 to 10, where 1 is only barely relevant and 10 is extremely relevant.

| | | |
|---|---|---|
| 1. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 2. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 3. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 4. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 5. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 6. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 7. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 8. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 9. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 10. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 11. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 12. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 13. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 14. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 15. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 16. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 17. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 18. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 19. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 20. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 21. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 22. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 23. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 24. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 25. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 26. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 27. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 28. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 29. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 30. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 31. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 32. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 33. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 34. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 35. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 36. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 37. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 38. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 39. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 40. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |

| | | |
|---|---|---|
| 41. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 42. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 43. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 44. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 45. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 46. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 47. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 48. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 49. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 50. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 51. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 52. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 53. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 54. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 55. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 56. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 57. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 58. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 59. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 60. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 61. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 62. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 63. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 64. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 65. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 66. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 67. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 68. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 69. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 70. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 71. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 72. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 73. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 74. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 75. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 76. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 77. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 78. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 79. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 80. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 81. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 82. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 83. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 84. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 85. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |

| | | |
|---|---|---|
| 86. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 87. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 88. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 89. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 90. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 91. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 92. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 93. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 94. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 95. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 96. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 97. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 98. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 99. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |
| 100. | Relevant? Yes or No | If yes, relevance rating (1 to 10): ____ |