

NASA/TM-2001-210913



Searching the ASRS Database Using QUORUM Keyword Search, Phrase Search, Phrase Generation, and Phrase Discovery

Michael W. McGreevy

April 2001

The methods described in this paper are patent pending.

Please contact David Lackner of the NASA Ames Commercial Technology Office, dlackner@mail.arc.nasa.gov, (650) 604-5761, for information about licensing the methods. A software implementation of the methods is available to licensees.

Table of Contents

Summary	1
Introduction	1
The Aviation Safety Reporting System (ASRS)	1
QUORUM.....	2
Enhancing QUORUM to support ASRS SRs and QRs	4
Core QUORUM methods	4
QUORUM analysis	4
QUORUM modeling	4
QUORUM relevance-ranking	4
New QUORUM methods.....	4
QUORUM keyword search	5
QUORUM phrase search.....	6
QUORUM phrase generation	8
QUORUM phrase discovery	8
Using the new QUORUM methods	9
Using QUORUM keyword search	9
Searching for "engage"	9
Searching for "rest"	10
Searching for "emergency"	12
Searching for "language", "English", or "phraseology"	12
Searching for "frequency congestion"	13
Using QUORUM phrase search.....	15
Searching for "conflict alert"	15
Searching for "frequency congestion"	15
Searching for "cockpit resource management"	16
Searching for "similar sounding callsign"	17
Searching for "flight crew fatigue"	19
Searching for a particular sentence that occurs only once in the database.....	19
QUORUM phrase generation and QUORUM phrase discovery	20
Using QUORUM phrase generation.....	21
Generating phrases containing "rain"	21
Searching for narratives containing "light moderate rain"	22
Searching for narratives containing common "rest" phrases	22
Using QUORUM phrase discovery	24

Related work	27
QUORUM keyword search and retrieval	27
QUORUM phrase search and retrieval	28
QUORUM phrase generation	28
QUORUM phrase discovery.....	28
Conclusion	29
References	29
Appendix 1. QUORUM relations and relational metrics	32
Calculating the metric values.....	32
An example of combining instance relations to produce a document relation	33
An alternative directionality metric	34
Creating additional metrics by removing some effects of frequency	34
Appendix 2. How QUORUM keyword search works	35
Appendix 3. How QUORUM phrase search works	38
Appendix 4. How QUORUM phrase generation works	41
Appendix 5. How QUORUM phrase discovery works	44

Searching the ASRS database using QUORUM keyword search, phrase search, phrase generation, and phrase discovery

MICHAEL W. MCGREEVY

NASA Ames Research Center

Summary

To support Search Requests and Quick Responses at the Aviation Safety Reporting System (ASRS), four new QUORUM methods have been developed: keyword search, phrase search, phrase generation, and phrase discovery. These methods build upon the core QUORUM methods of text analysis, modeling, and relevance-ranking. QUORUM keyword search retrieves ASRS incident narratives that contain one or more user-specified keywords in typical or selected contexts, and ranks the narratives on their relevance to the keywords in context. QUORUM phrase search retrieves narratives that contain one or more user-specified phrases, and ranks the narratives on their relevance to the phrases. QUORUM phrase generation produces a list of phrases from the ASRS database that contain a user-specified word or phrase. QUORUM phrase discovery finds phrases that are related to topics of interest. Phrase generation and phrase discovery are particularly useful for finding query phrases for input to QUORUM phrase search.

The presentation of the new QUORUM methods includes: a brief review of the underlying core QUORUM methods; an overview of the new methods; numerous, concrete examples of ASRS database searches using the new methods; discussion of related methods; and, in the appendices, detailed descriptions of the new methods.

Introduction

Avoiding accidents is a high priority for everyone involved in commercial aviation. Since commercial aviation accidents typically involve unique chains of events, prevention efforts must address the more common underlying factors that are the links in the chain. To reveal latent factors, and to measure their prevalence, the aviation safety community is increasingly relying on collection and analysis of incident reports. Leaders of the aviation safety community are concerned, however, that the implications of the many collected incident reports might not be fully appreciated. They want to make sure that the raw data provided by the incident reports is transformed into operationally useful information (Steenblik, 1999; Logan, 1998).

The Aviation Safety Reporting System (ASRS)

The Aviation Safety Reporting System (ASRS) is dedicated to the task of gathering aviation incident reports and helping to transform them into operationally useful information (ASRS, 1999; Connell, 1999). It has collected, distributed, analyzed, and interpreted aviation incident reports for over twenty years. According to the Federal Aviation Administration (FAA, 1999a), "The Aviation Safety Reporting System (ASRS) is a voluntary, confidential and anonymous incident reporting system. It is a cooperative program established under FAA Advisory Circular No. 00-46D, funded by the FAA and administered by NASA. Information collected by the ASRS is used to identify hazards and safety discrepancies in the National Aviation Airspace System. It is also used to formulate policy and to strengthen the foundation of aviation human factors safety research." To achieve its mandate, the ASRS annually processes the raw data of over 35,000 incident reports to add to its incident database, while producing formal safety alerts for aviation authorities, regular publications for the operational community, special studies, and customized collections of incident reports.

The ASRS strongly supports users who wish to access the information in their database. For example, the ASRS conducts many hundreds of database searches each year in response to Search Requests, to create collections for its Internet site, and to support its Quick Response studies.

As the leading U.S. repository of aviation incident reports, ASRS responds to 300-400 Search Requests per year from a wide variety of individuals and organizations. These include, for example, agencies of U.S. and foreign governments, aviation organizations, universities, researchers, and private citizens. Search Requests are also initiated by analysts within the ASRS and related organizations such as the Aviation Performance Measurement System office and the human factors research division at NASA Ames Research Center. These searches are typically done in association with special studies, presentations and publications, and "Alert Messages". Alert Messages are triggered by reports of significant aviation hazards, reports having significant accident prevention potential, and other incidents of particular concern.

In addition to responding to search requests, the ASRS generates topical collections of incident reports and posts them on its Internet site (ASRS, 1999). These collections represent a diversity of the most commonly requested topics, such as flight crew fatigue, controlled flight toward terrain, and runway incursions. Posting these reports on the Internet makes them readily available without the necessity of a formal Search Request.

The ASRS also performs a large number of carefully crafted searches for specially requested studies called Quick Responses. These studies are performed only for the Federal Aviation Administration, the National Transportation Safety Board, the Congress of the United States, or comparable organizations. The ASRS performs 5 to 10 Quick Responses each year. These analytic studies vary in scope, style, and format, but all of them require highly focused searches of the database. The retrieved incident narratives are thoroughly reviewed by ASRS analysts to ensure that they are relevant to the concerns of the study. The analysts then work to transform the raw data of the searches into operationally useful information that is published in the Quick Response reports.

Performing the large number of Search Requests and generating the detailed Quick Responses can be labor intensive. The many searches must be performed quickly without sacrificing quality. Yet the rapid retrieval of topically relevant reports is an art, and finding the most relevant reports among them can be time consuming. New methods that can enhance search and relevance-ranking capabilities in support of Search Requests and Quick Responses could improve productivity, giving analysts at ASRS more time for in-depth analysis. It is also possible that the quality of the search results could be improved, while reducing the time and effort to produce them.

The ASRS has consistently worked to improve its methods and to share them with the worldwide aviation safety community. As a result, the designs of confidential incident reporting and analysis systems have been beneficially influenced by the work of the ASRS. As airlines develop their internal incident reporting systems, improvements at the ASRS could suggest improvements in their systems. Given improved methods, the aviation community would be better able to transform their collections of incident reports into operationally useful information so that aviation safety can be further improved. Incident reporting systems in other domains, such as medicine, maritime operations, electric power generation, and Space Shuttle maintenance, could also benefit from demonstrated improvements at the ASRS.

QUORUM

QUORUM is a suite of NASA-developed text processing tools. Its core methods and software are capable of analyzing, modeling, and relevance-ranking incident narratives or other text (McGreevy & Statler, 1998; McGreevy, 1997; McGreevy, 1996; McGreevy, 1995). QUORUM measures the degree of contextual association of large numbers of word pairs in narratives or other text to produce models that capture the contextual structure of the text. It compares models to measure their degree of similarity. By ranking text items on their degree of similarity to a query model, QUORUM can retrieve the items that are most relevant to the query. As described in the present paper, these methods and software serve as the basis of new search and retrieval capabilities for ASRS Search Requests and Quick Responses.

QUORUM has already been shown to be effective in obtaining incident narratives from the ASRS database that are relevant to a variety of topics in aviation safety. For example, a recent study by McGreevy and Statler (1998) showed that QUORUM can find incidents that are relevant to the many factors involved in accidents. In that study, QUORUM found ASRS incident narratives that are relevant to the factors of the crash of American Airlines Flight 965 near Cali, Colombia in December 1995. That accident involved controlled flight into terrain, over-reliance on automation, confusion during descent and approach, problematic operations in foreign airspace, and a number of other factors. ASRS analysts independently judged that 84 of the top 100 incident narratives retrieved by QUORUM were relevant to the Cali accident. These excerpts are from some of the ASRS narratives that were judged to be relevant:

ATC CLRED US DIRECT TO THE CALI VOR AND DSND TO 5000 FT. ... FURTHER CHKING...SHOWED TERRAIN AT 14000 FT TO 11000 FT DIRECTLY ALONG OUR PATH. A SIMILAR ATC CLRNC HAPPENS VERY OFTEN FLYING INTO LIMA, PERU. MANY, MANY, MANY PLTS ARE NOT AWARE OF JUST HOW CRUCIAL IT IS NOT TO ACCEPT THESE DEADLY CLRNCs. PLEASE GET THE WORD OUT AGAIN. (ASRS report number 310143)

WHAT I FAILED TO NOTICE WAS THAT BY INSERTING THE ARR IN THE FMS, THE COMPUTER DUMPED THE XING RESTRICTION I HAD INSERTED JUST A FEW MOMENTS EARLIER. ... THE CAUSE, I BELIEVE, WAS A COMBINATION OF COCKPIT MGMNT OVERLOAD DURING THE APCH PHASE COUPLED WITH AN OVERCONFIDENCE IN THE FMS TO PRESENT VALID DSCNT PROFILE INFO. I ALLOWED MYSELF TO GET TOO BUSY DURING THE DSCNT TO MAKE ESSENTIAL XCHKS TO CONFIRM THE FMS WAS WORKING AS ADVERTISED. (272508)

AT THE LAST MIN, AFTER WE WERE VECTORED DIRECT TOWARD THE OUTER LOCATOR 'OC', WE WERE CLRED FOR A 'STRAIGHT IN LNDG ON RWY 11' AND TOLD TO RPT OVER 'OC.' ... THE FO INITIALLY SET UP HIS RADIO ON THE LOC 110.1, BUT THERE WAS NO LOC OR ANYTHING ON THAT FREQ. ... WE HAD BRIEFED BOTH THE ILS TO RWY 35 WITH A CIRCLE TO LAND AND THE LOC-VOR-DME RWY 11 APCH, BUT NOT A STRAIGHT IN APCH. THE ONLY STRAIGHT IN APCH WAS AN ADF LOCATOR APCH, WITH DME. ... MEANWHILE I WAS TRYING TO FIND AN APPROPRIATE APCH PAGE. WE SETTLED ON 11-2 CHART SINCE THE CTLR HAD CALLED THE APCH A 'STRAIGHT-IN APCH.' ... I SAID 'I AM CONFUSED.' I DIDN'T UNDERSTAND WHY WE WERE DSNDING AND THE FO HAD ALL FLAGS WITH HIS RADIO ON THE ILS FREQ. ... I COULDN'T FIGURE OUT WHICH APCH HE WAS USING, AND I HAD TROUBLE READING HIS CHART FROM ACROSS THE COCKPIT. THEN THE SO MENTIONED THAT WE HAD A 3000 FT MSA. WE WERE AT 2650 FT... THE APCH WE WERE FINALLY GIVEN, OR FLEW ANYWAY, DID NOT CONFORM TO ANY OF THE PLATES. ... I ACCEPTED THE CLRNC FOR A STRAIGHT-IN APCH, NOT KNOWING WHICH APCH. (310130)

THE CAPT STATED THAT WE WERE GOOD TO DSND NOW TO 4100 FT. I COMMENTED THAT THE AIRSPACE THAT PROTECTED US AT 4100 FT WAS ONLY VALID ONCE WE WERE ESTABLISHED ON FINAL AND OVER THE FIX INBOUND. HE SAID HE WAS SURE IT WAS SAFE... THE CAPT SEEMED VERY CONFIDENT AND NOTHING IN HIS MANNER SIGNALLED THAT I SHOULD BE AT ALL CONCERNED ABOUT HIS JUDGEMENT. I REMEMBERED THINKING EARLIER THAT HE SEEMED LIKE A REALLY GREAT GUY TO FLY WITH: VERY PROFESSIONAL AND SELF-ASSURED, WITH VERY GOOD PEOPLE SKILLS TOO. ... ALMOST IMMEDIATELY THE CAPT SAID SOMETHING ABOUT A MOUNTAIN BEING VISIBLE OUTSIDE THE WINDOW. I LOOKED OUT AND OUR LNDG LIGHTS WERE CLEARLY ILLUMINATING A LARGE PEAK BELOW OUR NOSE. THE CAPT SAID, 'THE RADIO ALT IS SHOWING 1000 FT, LET'S GET OUT OF HERE!' I DISCONNECTED THE AUTOPLT AND INITIATED A CLB AT TOGA THRUST BACK UP TO 6000 FT, SHORTLY AFTER I STARTED THE CLB, THE GPWS CALLED 'TERRAIN, TERRAIN!' (352618)

In an earlier study (McGreevy, 1997), QUORUM successfully relevance-ranked text from ASRS narratives, as well as wire service reports of the July 1996 crash of TWA Flight 800 near Long Island, New York. The study showed that QUORUM can effectively rank text items (e.g., sentences or narratives) according a variety of relevance criteria. For example, QUORUM can find the most typical text in a collection, the most topically relevant text in a collection, and text that is most similar to other text. As an illustration, here are the most typical sentences among the news stories, according to QUORUM:

"Mysterious explosion on TWA Flight 800 to Paris kills 230."

"The National Transportation Safety Board on Friday issued several urgent recommendations to the Federal Aviation Administration to protect fuel tanks from heat sources that could touch off the kind of explosion that occurred with TWA Flight 800."

The most typical sentence involving both the Federal Bureau of Investigation (FBI) and National Transportation Safety Board (NTSB) in the news stories was:

"The FBI and the National Transportation Safety Board are still investigating three theories: a missile, a bomb and mechanical failure."

The sentence most relevant to the topics of automation and training in a collection of ASRS narratives was:

"WITH REGARD TO TRAINING RECEIVED ON THIS ACFT, THE RPTR STATED THAT THE SIMULATOR DID NOT HAVE THIS CHARACTERISTIC OR PROB SO WAS NOT TRAINED IN THE EXACT AUTOPLT DESIGN THAT IS IN THE ACFT."

The sentence most relevant to the topic of "crew pressure" in a collection of ASRS narratives was:

"CAPT FELT SCHEDULE PRESSURE AND FELT RUSHED DURING SHORT TAXI."

In summary, previous studies of QUORUM's performance, particularly the Cali study, support the assertion that it is effective in analyzing, modeling, and relevance-ranking ASRS incident narratives and other text. The next step is to make QUORUM available for use by the ASRS so that it can be further evaluated and refined in an operational setting.

Enhancing QUORUM to support ASRS SRs and QRs

Given QUORUM's potential benefits to the ASRS (Connell, 1998; Rosenthal, 1998; Statler, 1998), the next step was to enhance QUORUM to support ASRS Search Requests and Quick Responses, while making it usable by analysts. This was accomplished in two phases. First, the core QUORUM software was redesigned and rewritten to make it portable, improve its performance, simplify its use, and extend its capabilities. Second, new QUORUM methods and software, which build upon the core methods, were developed to enable analysts to easily perform contextual keyword searches and flexible phrase searches.

The newly developed QUORUM methods are described in the section "New QUORUM methods", which follows the summary of the core methods. Demonstrations of the new methods are provided in the section, "Using the new QUORUM methods". QUORUM relations and relational metrics, which were enhanced to support the new methods, are described in appendix 1.

Core QUORUM methods

The core QUORUM methods include text analysis, modeling, and relevance-ranking. These have been thoroughly described elsewhere (McGreevy, 1999; McGreevy & Statler, 1998; McGreevy, 1997; McGreevy, 1996; McGreevy, 1995), and are summarized here.

QUORUM analysis — QUORUM analysis converts bodies of text to sequences of terms and measures the contextual associations among terms. QUORUM contextual analysis measures the structure of text as a way of measuring the structure of the domains and situations represented by the text. Terms that are contextually related in the structure of the text are considered to be contextually related in the world represented by the text.

QUORUM modeling — QUORUM modeling represents bodies of text, and the worlds they describe, as collections of QUORUM relations, that is, paired terms with measurements of the degree of their contextual association. If each term is considered to be a node, and each contextual association is an arc, then each QUORUM model consists of a network of contextually associated terms. A QUORUM model can represent any body of text, from the entire ASRS database taken as a whole to a short phrase, and it can represent any domain, sub-domain, situation, situational detail, topic, or subtopic.

QUORUM relevance-ranking — The similarity of any two QUORUM models can be quantified by comparing their features, that is, their paired terms and contextual measurements. By using one model's features as relevance-ranking criteria and comparing that model to a collection of models, the models in the collection can be ranked on their relevance to the criteria. For example, if the criterion model represents a narrative of interest, then a collection of narrative models can be relevance-ranked to find other narratives that are similar to it. Alternatively, if the criterion model represents a topic, then a collection of narrative models can be relevance-ranked according to that topic. As a further alternative, if the criterion model represents a collection of phrases of interest, then narratives or collections of phrases can be relevance-ranked according to those phrases. In addition, by ranking a collection of models on every model in the collection, an association matrix can be created to provide data for input to clustering methods. Given the diversity of text that can be represented by QUORUM models, QUORUM relevance-ranking has a great diversity of applications.

New QUORUM methods

To support ASRS Search Requests and Quick Responses, four new QUORUM methods have been developed: keyword search, phrase search, phrase generation, and phrase discovery.

- **QUORUM keyword search** retrieves from the ASRS database narratives that contain one or more user-specified keywords in typical or selected contexts, and ranks the narratives on their relevance to the keywords in context.
- **QUORUM phrase search** retrieves from the ASRS database narratives that contain one or more user-specified phrases, exactly or approximately, and ranks the narratives on their relevance to the phrases.
- **QUORUM phrase generation** produces a list of phrases from the ASRS database that contain a user-specified word or phrase.
- **QUORUM phrase discovery** finds phrases from the ASRS database that are related to topics of interest.

These methods build upon the core QUORUM methods of text analysis, modeling, and relevance-ranking. QUORUM phrase generation and phrase discovery are particularly useful for finding query phrases for input to QUORUM phrase search.

QUORUM keyword search— QUORUM keyword search is a user-oriented method of retrieving from the ASRS database narratives that contain one or more user-specified keywords in typical or selected contexts. The retrieved narratives are ranked according to their relevance to the keyword(s) in context. This differs from conventional keyword search methods, which detect the presence of keywords but do not consider the keywords in context. A step-by-step description of the keyword search method is shown in appendix 2. Its use and options are summarized here.

To initiate a keyword search, the user provides one or more keywords. QUORUM searches its models of the narratives of the ASRS database for the words in their typical or selected contexts and relevance-ranks the narratives accordingly. The most relevant narratives are automatically displayed to the user in a Netscape web browser window (fig. 1), allowing the user to scroll and review them. The relevant parts of each narrative are highlighted.

Each of the narratives in the browser window is accompanied by a list of the QUORUM relations that contributed to the relevance of the narrative. Casual users can ignore this information. Experienced users can refer to these relations to understand which features of the narrative were interpreted as contributing to the relevance of the narrative. In some cases, this can lead the user to modify the search strategy.

The automatic display of the most relevant narratives and relations might be sufficient to satisfy the needs of users in many cases. For others, the following information is available in three files:

- the highlighted narratives that were automatically displayed, along with the relations, in HyperText Markup Language (HTML) format;
- the accession numbers of all relevant narratives, sorted in decreasing order of relevance; and
- the QUORUM query model used in the keyword search.

Additional information is output to a window other than the narrative browser window. It shows the progress of the search, along with detailed documentation including parameters, options, and sub-processes. This information can be redirected to a file, if desired.

In using QUORUM keyword search, the user has a number of options. The most important option is to consider only a subset of the narratives in the ASRS database, rather than all of the narratives in the database. This allows QUORUM to work with the results of non-QUORUM search methods, such as those already in place at the ASRS. Those methods can find reports based on a wide variety of non-narrative, categorical information about the incidents, including: date, time of day, geographic location, type of anomaly, type of aircraft involved, etc. It is also possible to use QUORUM iteratively, so that incidents found to be relevant on one pass can be used as the subset of interest on a subsequent pass.

Other options of QUORUM keyword search instruct QUORUM to:

- Use either contained matching or exact matching. With contained matching (the default), a match is recognized if a query is contained in a word in a relation of a model. (See appendix 1 regarding relations.) So, for example, the word fragment "fatigu" would match "fatigue", "fatigued", and "fatiguing".
- Find narratives containing the logical union (that is, the logical OR) of the members of a set of keywords. This is the default. The search must match at least one of the keywords. Alternatively, QUORUM can find narratives containing the logical intersection (that is, the logical AND) of two sets of keywords. The search must match at least one of the keywords from each set. Each set can use either exact matching or contained matching.
- Disable mapping of input keywords to standard ASRS abbreviations and usage. The default is to do the mapping.
- Rank a user-specified database, rather than the default one representing the narratives of the ASRS database.
- Generate only the query model. This allows the query model to be refined, if desired.
- Apply a user-supplied query model, such as a refined model.

- Generate the query model and rank the narratives, but skip generation of the highlighted narratives.
- Include non-relevant narratives in the ranked list (at the end). The default is to list only those that are relevant.

QUORUM phrase search— A phrase is a sequence of two or more words that convey a single thought. QUORUM phrase search is a user-oriented method of retrieving from the ASRS database narratives that contain one or more user-specified phrases. A step-by-step description of the phrase search method is shown in appendix 3. Its use and options are summarized here.

To initiate a phrase search, the user provides one or more phrases. QUORUM searches its phrase-oriented models of the narratives of the ASRS database for the phrase(s) and relevance-ranks the narratives accordingly. The most relevant narratives are automatically displayed to the user in a Netscape browser window. The relevant parts of each narrative are highlighted. The more relevant narratives typically contain more instances of the exact phrase. Less relevant narratives contain only instances of close matches to the phrase, or fragments of the phrase. This differs from conventional phrase search methods, which can only detect exact matches.

Each of the narratives in the browser window is accompanied by a list of the QUORUM relations that contributed to the relevance of the narrative. Casual users can ignore this information. Experienced users can refer to these relations to understand which features of the narrative were interpreted as contributing to the relevance of the narrative. In some cases, this can lead the user to modify the search strategy.

The automatic display of the most relevant narratives and relations might be sufficient to satisfy the needs of users in many cases. For others, the following information is available in three files:

- the highlighted narratives that were automatically displayed, along with the relations, in HTML format;
- the accession numbers of all relevant narratives, sorted in decreasing order of relevance; and
- the QUORUM query model used in the phrase search.

Additional information is output to a window other than the narrative browser window. It shows the progress of the search, along with detailed documentation including parameters, options, and sub-processes. This information can be redirected to a file, if desired.

In using QUORUM phrase search, the user has a number of options. As with QUORUM keyword search, the most important option for QUORUM phrase search is the ability to consider a subset of the ASRS database, allowing the phrase search to work with the results of non-QUORUM search methods. The subset capability also allows earlier QUORUM searches to provide subsets to later ones. Another phrase search option improves results when the query phrase or phrases contain common stopwords such as "the", "and", "to", etc., by treating non-stopwords as more important than stopwords. A third option slightly loosens the definition of a phrase match so that a greater number of narratives is retrieved. This option accepts narratives containing only some fragments of the query phrase, but ranks them after those that match the whole phrase and those that match all fragments of the phrase.

Most users will be satisfied with results obtained using the default behaviors of QUORUM phrase search, or the options just described. Other phrase search options are available that instruct QUORUM to:

- Emphasize one or more particular words in the query phrase(s).
- Emphasize contextual relations that appear in multiple query phrases.
- Ignore (rather than de-emphasize, as described earlier) any stopwords in the query phrases.
- Ignore contextual relations involving a particular word or pair of words. Multiple instances are allowed.
- Output the QUORUM phrase query model without performing the search.
- Include non-relevant narratives in the file of ranked accession numbers (at the end).
- Rank a user-supplied database of narrative phrase models, rather than the default database.




Figure 1. View of a browser window containing output of a QUORUM keyword search. The window contains highlighted narratives and associated QUORUM relations, and may be scrolled to view the other narratives and relations. This excerpt is taken from the output of a search on the keywords: English, language, and phraseology. The contents and format of this display are described in the section "Using the new QUORUM methods" in the subsection "Using QUORUM keyword search."

QUORUM phrase generation— QUORUM phrase generation produces a list of phrases from the ASRS database that contain a user-specified word or phrase. This output can be helpful in suggesting queries for phrase searches.

To generate phrases, the user provides a word or phrase that is to be contained in each of the output phrases. QUORUM builds phrases around this input, based on its phrase models. The resulting phrases are displayed on the computer screen, or can be redirected to a file. The phrases are sorted on an estimate of their prominence in the ASRS database.

Options of QUORUM phrase generation instruct QUORUM to:

- Generate the N most prominent phrases. This is the default. The default value of N is 10.
- Generate phrases having a prominence metric value of at least M. Default is to generate N phrases (see above).
- Allow up to S stopwords in each of the output phrases. The default value of S is zero.
- Use a file of query words and/or phrases (one per line). The default is to use a single query entered by the user.
- Apply a given stoplist rather than the default one, or temporarily append some stop words to the default stoplist.
- Use a particular phrase database, rather than the one representing phrases in the narratives of the ASRS database.

A step-by-step description of the QUORUM phrase generation method is shown in appendix 4.

QUORUM phrase discovery— QUORUM phrase discovery finds phrases that are related to topics of interest. Given a topic such as "fatigue", it can discover related phrases such as "rest period", "reduced rest", "duty period", "crew scheduling", "continuous duty overnight", "crew fatigue", and "compensatory rest". The topical phrases can assist in understanding the variations and scope of topics in the ASRS database. They can also be used selectively as input to QUORUM phrase search so that narratives containing particular topical variations can be retrieved.

While QUORUM keyword search, phrase search, and phrase generation can each be done in one easy step, QUORUM phrase discovery involves a sequence of automated and manual steps. The first, fully automated part of the process produces many topical and situational phrases, and these results can be very useful. To achieve a more comprehensive list of topical phrases, however, further manual and automated processing is done.

The first step of QUORUM phrase discovery is to perform a QUORUM keyword or phrase search. Next, phrases are automatically extracted from the most relevant narratives. The phrases produced at this point can be useful, but further processing improves the results. From these phrases, topical phrases are distilled by a combination of manual and automated methods. The refined set of topical phrases is then used to query the database using QUORUM phrase search. The cycle of phrase extraction and search is repeated to produce a final set of narratives. If narratives relevant to the topic are available in the database, this final collection of narratives will be highly relevant to the topic.

The main product is a list of the topical phrases that are extracted from the final collection of narratives. Other products include:

- a display of the narratives that are most relevant to the whole topic of interest, with the relevant sections highlighted, along with the corresponding HTML file;
- a list of the accession numbers of all relevant narratives, sorted in decreasing order of relevance; and
- the QUORUM query model used to rank the database of narratives.

Narratives among the final collection retrieved by the QUORUM phrase discovery method need not contain the initial keywords or phrases used as the query. For example, given the topic "fatigue", the retrieved narratives are highly relevant to "fatigue" but do not necessarily contain the word "fatigue". Relevant narratives contain a variety of fatigue-related phrases such as "rest period", "reduced rest", "duty period", "crew fatigue", etc. Narratives containing the greatest number of the most prominent phrases are considered to be the most relevant.

The current implementation of this method is more of a research tool than a user-oriented method. It is presented here because it utilizes QUORUM keyword and phrase search methods, and it could be very helpful for ASRS Search Requests and Quick Responses once it is packaged as an easy-to-use search tool. Example results are presented in the section, "Using QUORUM phrase discovery". A detailed, step-by-step description of the method is shown in appendix 5.

Using the new QUORUM methods

In this section, the new QUORUM methods are used in illustrative examples. QUORUM keyword search and QUORUM phrase search are easy to use, and can be used to perform searches for ASRS Search Requests and Quick Responses. QUORUM phrase generation is also easy to use, and can be used to suggest query phrases likely to be found in ASRS narratives. The complete QUORUM phrase discovery method is currently too complex for casual users, but it is a powerful method for finding phrases that are related to topics of interest. This could be of significant benefit in ASRS searches, and for elaboration of taxonomies.

Using QUORUM keyword search

Using QUORUM keyword search is easy. All the user needs to do is provide the keyword or keywords of interest. QUORUM then sorts the narratives of the ASRS database according to their relevance to the query, and displays the most relevant narratives with the relevant sections highlighted. Despite this simplicity, there are some details that must be understood. The examples of QUORUM keyword search that are shown below illustrate the most important details.

Searching for "engage"— To find narratives relevant to "engage", the user provides the word "engage" to QUORUM keyword search. QUORUM displays the most relevant narratives, with their relevant sections highlighted. Also output are three files and documentation of the search, as described in section "New QUORUM methods" in the subsection "QUORUM keyword search". The files contain the highlighted narratives, a complete list of relevant narratives, and the QUORUM model used to search the ASRS database.

Here is an example of a relevant narrative:

ON FEB / XX / 95 AT ABOUT XA00 PM SAN JUAN TIME WE DEPARTED RWY 8 ENRTE TO MIAMI . WE INTERCEPTED THE JAAWS 9 DEP , AND SHORTLY AFTER PASSING THROUGH 10000 FT WE WERE **CLRED** DIRECT (RNAV) TO JUNUR , WHICH IS A **POINT** IN THE CLAMI 1 ARR INTO MIAMI . I THEN **ENGAGED** THE **AUTOPLT** AND TURNED THE ACFT IN THE DIRECTION OF THE WAYPOINT (JUNUR) WE WERE **CLRED** TO . AT THIS **POINT** I AM **NOT SURE** IF I **ENGAGED** THE **AUX NAV PORTION** OF THE **AUTOPLT** . THE REASON I SAY THIS IS BECAUSE APPROX 1 HR LATER WE DISCOVERED THAT THE **AUX NAV PORTION** OF THE **AUTOPLT** WAS **NOT ENGAGED** AND WE HAD **DRIFTED** ABOUT 45 NM OFF **COURSE** . IT IS UNKNOWN WHETHER THE **AUX NAV** WAS **NEVER ENGAGED** OR IF THE KNOB WAS SOMEHOW KNOCKED OFF DURING THE FLT . I DO REMEMBER PASSING ALMOST DIRECTLY OVER GTK VOR WHICH IS ALONG THE NORMAL **RTE** THE ACFT WOULD **TAKE** IF THE **OMEGA** WERE **ENGAGED** . 2 SCENARIOS ARE POSSIBLE . THE **OMEGA** WAS **NEVER ENGAGED** , AND DUE TO **LIGHT HIGH ALT** WINDS , THE ACFT AFTER INITIALLY BEING POINTED IN THE CORRECT DIRECTION , ONLY BEGAN TO DRIFT DRAMATICALLY AFTER PASSING GTK VOR . OR , THE **AUX NAV KNOB** WAS ACCIDENTLY **DISENGAGED** AND WAS **NOT NOTICED** . THERE IS NO **AURAL** OR OTHER **TYPE WARNING** WHEN THE **OMEGA** BECOMES **DISENGAGED** . THERE IS A GREEN ' **AUX NAV ' LIGHT** THAT IS **ILLUMINATED** WHEN **ENGAGED** , **BUT** THE **LIGHT** IS **NOT** VERY OBVIOUS TO THE CREW . SOME TYPE OF OBVIOUS WARNING (HAD IT BEEN AVAILABLE) WOULD HAVE ALERTED THE CREW IN THE EVENT OF AN INADVERTENT DISCONNECT . ONE THING WE FOUND UNUSUAL DURING OUR FLT WAS THAT ATC NEVER SAID A WORD TO US DURING OUR SMALL DETOUR . (300563)

The default pattern-matching behavior of QUORUM keyword search is to do a "contained match". This means that any word that contains the string of characters "engage" is considered to be a match. So, narratives containing the following words are retrieved:

- engage
- disengage
- reengage
- engagement
- engaged
- disengaged
- reengaged
- disengagement

In the example narrative, the word "engaged" appears 7 times, "disengaged" appears twice, and "engage" does not appear. This shows the value of allowing the "contained match" as the default. The user need not know the various forms of the word that appear in the narratives, but can find the narratives that are clearly relevant to their query word.

Not only are the various forms of the word "engage" highlighted in the example narrative, but other words are also highlighted. These other words are often found in the context of "engage" in the ASRS database. Highlighting is currently limited to a user-selectable number of the most prominent contextual associations of the keyword in the database. The default number is 1000. New user-selectable options to keyword search could limit highlighting to just the keyword(s), or to contextual associations that have some fraction of the prominence of the most prominent association in the database or the particular narrative.

The automatic display of the most relevant narratives will satisfy the needs of many users, but others might be interested in a deeper understanding of which contextual associations contribute to the relevance of each narrative. By referring to the data table that is displayed after each narrative, it is possible to identify the words in the narrative that are most often found in the context of the query word(s). Here is the top part of the data table for the example narrative:

W1	W2	A	B	C
ENGAGED	AUTOPLT	17905	70	41.6048
NOT	ENGAGED	2484	72	33.4334
NAV	ENGAGED	898	94	30.8952
ENGAGED	ALT	6015	27	28.6804
ENGAGED	LIGHT	508	74	26.8164
OMEGA	ENGAGED	386	87	26.5982
DISENGAGED	NOT	896	39	24.9047
ENGAGED	BUT	984	24	21.902
NEVER	ENGAGED	159	73	21.7479
AUX	ENGAGED	117	94	21.636
CLRED	ENGAGED	364	26	19.2135
ENGAGED	COURSE	239	32	18.98
OMEGA	DISENGAGED	202	34	18.7189
WARNING	DISENGAGED	202	34	18.7189

Each line in the table represents a contextual association between two words (i.e., the words in columns W1 and W2). Column A is a measure of the strength of the contextual association of the word pair in the whole ASRS database. Column B is a measure of the strength of the same contextual association in this narrative. Column C is a combination of these two metrics and represents a measure of the contextual association of the paired words. In this table, C is the product of the natural logarithms of A and B. The value of C is large when the values of both A and B are large. The relations are sorted on column C.

Word pairs toward the top of the list have stronger contextual associations. The top relation, for example, is between ENGAGED and AUTOPLT (i.e., autopilot). This relation is at the top of the list because AUTOPLT is very often found in the context of ENGAGED in the ASRS database (as indicated by 17905 in column A) and that relationship is also relatively prominent in this narrative (as indicated by 70 in column B). The word ENGAGED is in column W1, and the word AUTOPLT is in W2 because ENGAGED tends to precede AUTOPLT in the narratives of the ASRS database. In general, each pair of words appears in the more typical order.

The contextual relationship between ENGAGED and AUTOPLT can be seen in these excerpts from the example narrative:

- I THEN *ENGAGED* THE *AUTOPLT*
- IF I *ENGAGED* THE *AUX NAV PORTION* OF THE *AUTOPLT*
- THE *AUX NAV PORTION* OF THE *AUTOPLT* WAS *NOT ENGAGED*

The user could take further advantage of the default "contained match" rule by using "engag" as the query. This would match several forms of "engage", including not only those listed earlier, but also "engaging" and "disengaging". Alternatively, the user has the option of requiring an exact match, so that only narratives containing the word "engage" would be retrieved.

Searching for "rest"— A search for narratives relevant to "rest" requires the use of the "exact match" option. That is because the default "contained match" option that worked so well in the previous example becomes a liability when the query is contained in too many words. "Rest" is such a query, as indicated by this long list of words from the ASRS database that contain "rest":

RESTR	INTERESTED	RESTRICTING	RESTARTING	RESTRAINED
REST	INTERESTING	RESTRICTIVE	CREST	RESTRAINTS
RESTRICTION	RESTATED	UNRESTR	INTERESTS	BREST
RESTRICTIONS	ARRESTED	RESTING	RESTATE	OVERESTIMATED
NEAREST	RESTED	RESTAURANT	RESTRICTS	RESTATING
RESTART	ARREST	ARRESTING	PRESTART	RESTORATION
RESTRS	RESTORE	RESTROOM	UNDERESTIMATED	RESTRAINING
INTEREST	UNRESTRICTED	RESTRICTED	INTERESTINGLY	ARMREST
RESTARTED	RESTRICT	RESTS	RESTORING	RESTLESS
RESTORED	FOREST	CRESTVIEW	RESTRAINT	

To find narratives relevant to "rest", the user can provide the query "rest" to QUORUM keyword search and select the "exact match" option. QUORUM displays the most relevant narratives, with their relevant sections highlighted. Here is one of the most relevant narratives:

CREW REST REGS : UNFORTUNATELY , EVERY ONCE IN A WHILE FOR A VARIETY OF REASONS , THIS REG (DESIGNED TO ENSURE PROPERLY RESTED PLTS) GETS FORGOTTEN ! TRY AND FIGURE THIS ONE . 2 DAY PAIRING SCHEDULE FOR 10 PLUS 09 , THE FIRST DAY SHOW TIME IS LATE EVENING AND FLT TIME IS SCHEDULED FOR 3 PLUS 44 . DUE TO MECHANICAL PROBLEM WE PUSHED : 20 LATE , WX IN THE AREA DELAYED OUR TKOF . WITH AN UNSCHEDULED FUEL STOP WE LANDED AND PARKED AT THE DEST GATE 1 PLUS 51 LATE . ORIGINALLY WE WERE **SCHEDULED** FOR 10 PLUS 16 LAYOVER . OUR COMPANY 'S STD RESPONSE WHEN CALLED TO **CHK CREW REST** IS 8 PLUS 44 **BLOCK TO BLOCK** (XX AND 8 PLUS 44 = A PUSH TIME OF XXY) SINCE OUR PUSH TIME WAS SCHEDULED FOR XXY THERE WAS NOT A CONFLICT IN OUR THINKING . AT EARLY SCHEDULING AWOKED THE **CAPT** , INFORMING HIM THAT THE FO AND SO ' **REQUIRED** 9 PLUS 45 ' **BLOCK TO BLOCK CREW REST** . WE ALL SHOWED AS PLANNED THE PREVIOUS EVENING FOR **SCHEDULED** VAN . THE **CAPT** INFORMED FO AND I ABOUT CALL FROM SCHEDULES , IT JUST DID NOT MAKE SENSE . WE FLEW 4 PLUS 13 THE NIGHT BEFORE AND WERE SCHEDULED TO FLY 6 PLUS 25 THIS DAY . WHAT WERE WE TO DO ? GO BACK TO OUR ROOMS AND SLEEP FOR ANOTHER 45 MINS ? WE SHOWED ON THE ACFT (8 PLUS 51 FROM BLOCK IN) ACFT WAS BOARDED NORMALLY AND WE SAT WITH THE PARKING BRAKE SET SO AS NOT TO TRIP ACARS UNTIL SCHEDULING GOT THEIR IMPOSED 9 PLUS 45 **BLOCK TO BLOCK** , HOWEVER , I **SEE** THAT 1) THEY **INTERRUPTED CAPT CREW REST** . 2) THEIR **REST INTERPRETATION** WAS SOMEHOW FLAWED (ALTHOUGH APPRECIATED WHEN WE GET ' MORE ' **REST**) . 3) ' MORE ' **REST** I DO **NOT NEED SPENT SITTING** 54 MINS WITH PARKING **BRAKE SET - WAITING** TO BE **LEGAL** . MY AIRLINE **USES FAR** MIN **REST** AS NORMAL **PRACTICE** AND **ROUTINELY VIOLATES CREW REST** FOR **PERHAPS** MISINTERPRETED **REST REGS REQUIRED** . I **FEEL** 1) **FAA** MUST **MAKE** BOTH FLT TIME AND **DUTY** TIME **HENCE REST** TIMES **EASIER TO UNDERSTAND** (THROW OUT **INTERPRETATIONS**) ! 2) **HOLD CREW SCHEDULERS** ACCOUNTABLE FOR VIOLATIONS OF **CREW REST** , A GOOD **SCHEDULE PRACTICE** WOULD HAVE BEEN TO **INFORM US** ON **ARR** THE PREVIOUS **NIGHT** OF **REST REQUIRED** . (183457)

The words CREW, REQUIRED, BLOCK, NOT, DUTY, CAPT (i.e., captain), FAR (i.e., Federal Aviation Regulations), REGS (i.e., regulations), LEGAL, FAA (i.e., Federal Aviation Administration), NIGHT, FEEL, SCHEDULED, and others are highlighted in the narrative because they are often found in the context of REST in the narratives of the ASRS database.

The needs of many users will be satisfied by the display of the most relevant narratives, but others might wish to better understand the relevance of each narrative. By referring to the data table that is automatically displayed after each narrative, the user can see the relative association of REST with the words found most often found in the context of REST. Here is the top part of the data table for the example narrative:

word1	word2	A	B	C
CREW	REST	9241	264	50.9163
REST	REQUIRED	2281	115	36.6896
BLOCK	REST	1181	124	34.0992
REST	NOT	4639	44	31.9471
DUTY	REST	4595	43	31.7172
CAPT	REST	1302	66	30.0468
FAR	REST	1534	56	29.5285
REST	REGS	643	93	29.3084
LEGAL	REST	1606	47	28.4199
REST	FAA	1207	54	28.3054
NIGHT	REST	2375	34	27.4095
REST	FEEL	462	60	25.1211
REST	SCHEDULED	2372	24	24.6982
REST	NEED	693	42	24.4482
REST	SCHEDULE	852	35	23.99

The format of this table was described in the previous example. In this case the table indicates, for example, that CREW is often found in the context of REST in both the database and in this narrative, and CREW typically precedes REST in the database. Further, since the value in column C for that word pair is greater than that for any of the other word pairs, the contextual association of CREW and REST is stronger than that of any of the other word pairs. The other contextual associations can be interpreted in a similar fashion.

Should the user be interested in searching for all forms of "rest", the best approach is to use the exact match and to provide the words "rest", "resting", and "rested" as query terms.

Searching for "emergency"— To find narratives relevant to "emergency", the user provides that keyword, and QUORUM then retrieves and displays the most relevant narratives, with their most relevant sections highlighted. Here is an example narrative:

A FEW MINS AFTER REACHING *FL350 CABIN RAPIDLY* DEPRESSURIZED . *COCKPIT CREW VERIFIED RAPID DECOMPRESSION* , BEGAN *EMER DSCNT* , *DECLARED AN EMER CONDITION* WITH *ARTCC* AND *SIMULTANEOUSLY REQUESTED A DIRECT VECTOR* TO THE *NEAREST SUITABLE ARPT* WHICH WAS *DETERMINED* BY *CAPT* TO BE *STL 110 MI AWAY* . ALL *EMER CHECKLISTS* AND NORMAL *CHECKLISTS COMPLETED* AND AN *UNEVENTFUL APCH* AND *LNDG* WAS *MADE* . NO *INJURIES* . I HAVE UNFORTUNATELY DONE 2 *EMER DSCNTS* IN THE LAST 18 MONTHS DUE TO THE SAME COMPUTER FAILURE OF THE PRESSURIZATION SYS . THE ODDS AGAINST THAT ARE STAGGERING . I BELIEVE THIS ACFT 'S AUTO *CABIN CTRLRS* SHOULD BE *LOOKED* AT CAREFULLY . ALSO , *EMER PROC TRAINING* AT MY *COMPANY* FOR *EMER DSCNTS NEEDS* TO BE *REVIEWED* AND MODIFIED AS WELL AS *THOUGHT* GIVEN TO MANY FACTORS NEVER DISCUSSED DURING *TRAINING* . (110788)

The word "emergency" does not appear in the narrative because the ASRS abbreviates the word "emergency" as "emer". QUORUM keyword search automatically maps the user's input words to the ASRS abbreviations, as long as those mappings are contained in the mapping file used by QUORUM. The user has access to the mapping file, and also has the option of turning off the automatic mapping.

The highlighted words include the query word (as abbreviated by the ASRS) and those words that are often found in the context of the query in the narratives of the ASRS database. The user might, in some cases, wish to exercise an option of highlighting only the original query words, or an option of highlighting them in a manner different from that of the other words.

Searching for "language", "English", or "phraseology"— A realistic search example can be drawn from ASRS Quick Response (QR) 290, "An analysis of international airspace incidents". Among the factors being investigated in that report were "ATC language barrier factors". One way to obtain incident reports containing those factors is to retrieve narratives containing the words "language", "English", and/or "phraseology". To accomplish this search, the user enters these keywords, and then QUORUM ranks the narratives of the ASRS database according to their relevance to the typical or selected contexts of these words in the database.

In order to match the search done in QR 290, the search needs to be limited to a subset of the ASRS database. That subset consists of reports occurring in non-U.S. airspace, from January 1993 through December 1996, in which the reporter was a member of a U.S. flight crew. This subset was obtained, using the current ASRS retrieval methods, by searching the information coded in the non-narrative data fields that are associated with each report in the ASRS database. The subset is described to QUORUM by providing a list of accession numbers of reports in the subset.

Here is an example of one of the most relevant narratives retrieved and displayed by QUORUM keyword search:

TKOF *CLRNC* WAS MISUNDERSTOOD BY CREW . *TWR CTRL* 'S *ENGLISH* WAS *NOT* VERY CLR AND HE *USED INCORRECT PHRASEOLOGY* WHICH *CAUSED* AN APPARENT *ALT* ' BUST . ' ATC *CLRNC* WAS TO 9000 FT , WHICH IS NORMAL FOR THEM . WE WERE USING RWY 21 . TKOF CLRNC WAS ' CLRED FOR TKOF , RWY HDG 210 DEGS , CONTACT DEP . ' DEP SAID WE WERE CLRED TO 2100 FT (AS WE WERE PASSING 3000 FT) . EVIDENTLY THE ' 21 ' AFTER ' RWY HDG ' WAS MEANT AS AN AMENDED *ALT CLRNC* . IF *PROPER PHRASEOLOGY* HAD BEEN *USED* , I AM *SURE* WE WOULD HAVE EITHER *UNDERSTOOD* OR *ASKED* FOR A *CLARIFICATION* . *PROPER PHRASEOLOGY* IS EVEN MORE IMPORTANT WHEN SPEAKING TO *PEOPLE* WHOSE *PRIMARY LANGUAGE* IS *NOT ENGLISH* . *PLTS* SHOULD UNDERSTAND THIS BECAUSE OF TRYING TO GIVE POS RPTS , ETC , TO SO MANY DIFFERENT *PEOPLE* . (236336)

Here are some relevant sentences from other highly relevant narratives:

EXTREMELY DIFFICULT TO COPY *CLRNC* BECAUSE OF *POOR ENGLISH* OF *CTRL* AND NO *SPANISH* BY *PLTS* . (306637)

I THINK AN IMMEDIATE REVIEW OF RELATED FIX *NAMES* FOR SIMILAR SOUNDING *NAMES* AS PRONOUNCED BY THE *LCL SPEAKER* 'S *LANGUAGE* IS ESSENTIAL . (242971)

THE COM BTWN THE *FRENCH CTLRS* AND *ENGLISH SPEAKING PLTS* HAS BEEN POOR FOR SOME TIME , AND IS GETTING WORSE . (301205)

FLYING A LOT OF TIME IN CENTRAL AND S AMERICA , I EXPERIENCE THAT ATC *CTLRS* DON'T HAVE *FLUENT* TALKING AND *UNDERSTANDING* OF THE *ENGLISH LANGUAGE* , AS THE WAY HAS TO BE CONSIDERING THAT *ENGLISH* IS THE *UNIVERSAL* AND INTL *LANGUAGE* IN *AVIATION* . (302310)

THE RPTR *SAID* THAT HE OFTEN HEARS *IMPROPER PHRASEOLOGY* DURING HIS *FOREIGN OPS* . (352400)

AFTER MUCH DISCUSSION AND *PROBS* WITH *LANGUAGE BARRIER* AND *PHRASEOLOGY* , *CREW* WAS GIVEN *PERMISSION* TO START *TAXI* AT *CREW* 'S OWN RESPONSIBILITY . (295947)

ALTHOUGH *ENGLISH* IS THE *OFFICIAL LANGUAGE* OF TRINIDAD , *LCL DIALECT* MAKES IT *DIFFICULT* TO *UNDERSTAND CTLRS* . (294060)

BETTER ENGLISH SPEAKING FOREIGN CTLRS AND *USE OF STD PHRASEOLOGY* IS *NEEDED* . (268223)

SITUATIONAL AWARENESS IS NONEXISTENT WHEN *CTLRS* SPEAK TO EVERYONE ELSE IN A *FOREIGN LANGUAGE* AND TO YOU IN *BROKEN ENGLISH* ! (344832)

TWR PHRASEOLOGY WAS NON *STD* AND HIS *COMMAND* OF *ENGLISH* WAS *LIMITED* , BUT WE WERE *CLRED* TO *LAND* . (332620)

Given the keywords used in this search, the top-ranked narratives typically describe incidents involving miscommunication between air traffic controllers and flight crews due to language barriers, including poor use of the English language and the use of non-standard phraseology. For each search term, here are some of the typical contexts, as indicated by the QUORUM query models and reflected in some of the excerpts above:

- "Language" is often found in the context of *barrier*, *English* and *Spanish*; *clearances*; *air traffic controllers* and *ATC*; and *problems*, *differences* and *difficulties*.
- "English" is often found in the context of the verbs *speaking* and *understanding*; the attributes *poor*, *broken*, and *limited*; *Spanish* and *French*; and *air traffic controllers* and *pilots*.
- "Phraseology" is often found in the context of the attributes *standard* and *proper*; the verbs *use*, *used*, and *using*; *ATC*, *air traffic controllers*, and *towers*; and *clearances* and *runways*.

While the top narratives retrieved in this search all involve "ATC language barrier factors" it should be noted that there was no requirement that the narratives should involve ATC. Since the typical contexts of language barrier factors do, in fact, involve ATC, the top narratives also involved ATC. As a consequence, however, as one goes farther down the list of relevant reports, at some point reports will be found that involve language barrier factors but not ATC.

The most reliable way to address this is to use a different subset of reports. In the search just demonstrated, as in QR 290, the subset consisted of reports occurring in non-U.S. airspace, from January 1993 through December 1996, in which the reporter was a member of a U.S. flight crew. This subset was obtained using conventional ASRS methods to search the non-narrative fields of the database. These same methods should be used to search the non-narrative field "PERSONS FUNCTIONS" for ATC roles. The subset would then become: reports occurring in non-U.S. airspace, from January 1993 through December 1996, in which the reporter was a member of a U.S. flight crew, and in which an ATC person played a role. The QUORUM search shown above would then be more strictly limited to ATC language barrier factors.

Searching for "frequency congestion"— QUORUM keyword search will take any number of keywords as queries, as in the above examples, but the user must understand that each word is treated individually. A search on the keywords "frequency congestion" will return narratives that contain either one or both of these keywords and their contexts. There is no guarantee, however, that both of the keywords will appear in the top-ranked narratives because the search treats each query word as an independent item.

To address this kind of situation, QUORUM keyword search has an option to require the logical intersection of two searches. The query for each search can be specified by one or more words. In this example, the "frequency" search uses the query "freq freqs" and requires an exact match. This query avoids matches on words such as "frequently". The "congestion" search uses the query "congestion congested" and requires an exact match. This query avoids matches on "uncongested".

To perform the search, the user provides the two search queries, "freq freqs" and "congestion congested", and indicates that both matches are to be exact, and that the logical intersection of these searches is required. QUORUM keyword search then retrieves and relevance-ranks narratives that contain both "frequency" in context and "congestion" in context.

The following are excerpts from some of the most relevant narratives:

SEVERAL ATTEMPTS WERE *MADE TO CONTACT TWR* , BUT DUE TO *EXTREME CONGESTION* ON THIS *FREQ* NO *LNDG CLRNC* WAS OBTAINED *FREQ* 124.15 WAS SO *CONGESTED* THAT NO ACFT COULD *XMIT* ON THIS *FREQ* CORRECTIVE ACTIONS : ... NOTAM *FREQ* 124.75 AS AN *ALTERNATE FREQ* ON *ATIS* [.] DECREASE *CONGESTION* OF *TWR FREQ* . (151711)

I *FINALLY SWITCHED* BACK TO THE ORIGINAL *CTLR FREQ* BUT , DUE TO *CONGESTED FREQ* , I *SWITCHED* TO THE *TWR FREQ* TO GET THROUGH , WHICH I *FINALLY* DID MAYBE ON SUBSEQUENT FLTS , IF THIS PROB SHOULD COME ABOUT , IT MIGHT BE A GOOD IDEA TO ALWAYS *LEAVE* ONE OF THE *RADIOS SET* TO THE LAST *FREQ* TO GO BACK TO WHEN THE *FREQ* GETS *BUSY* OR WHEN NOBODY SEEMS TO BE *WORKING* THAT *FREQ* . (237353)

AFTER CLRING RWY 33L , WE WERE UNABLE TO CONTACT *GND CTL* DUE TO *FREQ CONGESTION* TAXIING INBND WITHOUT FIRST RECEIVING A *CLRNC* IS *NOT* AT ALL UNUSUAL AT *FREQ CONGESTED ARPTS* . IN SIMILAR SITS AT *BWI* AND ELSEWHERE , IF THE *FREQ* IS *BLOCKED* AND A CUSTOMARY *TAXI RTE* IS KNOWN AND *CLR* OF TFC , NEARLY AL[L] CAPTS I HAVE OBSERVED WOULD PROCEED SLOWLY , AS WE DID . WE PROGRESSED FARTHER THAN MOST ONLY BECAUSE THE *FREQ* WAS *CONGESTED LONGER* , IN PART BECAUSE THE *CTLR* WOULD *NOT* UNKEY HIS MIC WHILE *MAKING* MULTIPLE *XMISSIONS* . (173324)

BECAUSE OF *EXTREME FREQ CONGESTION* , ABBREVIATED *TAXI INSTRUCTIONS* ARE *GIVEN* AT *ORD* THE *FREQ CONGESTION* AND *CTLR WORKLOAD* AT *ORD* MAKE IT HARD TO VERIFY *INSTRUCTIONS* THAT ARE UNCLR . WE ATTEMPTED CONTACT A FEW TIMES BEFORE BEING *TOLD* TO *TURN NEAR* THE BARRICADES , BUT WERE THEN *GIVEN* AN IMMEDIATE *FREQ CHANGE* WHICH *PREVENTED* PROMPT FEEDBACK FROM THE *CTLR* WHO *GAVE* US THE *INSTRUCTIONS* . TO THEIR CREDIT , THEY DID SPOT THE *ERROR* QUICKLY AND *CALLED* ON *TWR FREQ* WITH *NEW INSTRUCTIONS* . (WE MAY NOT HAVE *HEARD* SOME CALLS DUE TO RECEPTION *PROBS* .) THE *CONGESTION* AT *ORD* WOULD BE TOUGH TO FIX , BUT BETTER ARPT SIGNS SHOWING *TAXI RTES* THROUGH THE CONSTRUCTION AREAS WILL DEFINITELY CUT DOWN ON FUTURE *PROBS* . (252779)

These and other relevant narratives indicate that the topics "frequency" and "congestion" are often found in the same contexts, but that the exact phrase "frequency congestion" is not always present. Instead, many forms are found, such as:

- *CONGESTION* ON THIS *FREQ*
- *FREQ* 124.15 WAS SO *CONGESTED*
- *CONGESTION* OF *TWR FREQ*
- *CONGESTED FREQ*
- *FREQ CONGESTION*
- *FREQ CONGESTED*
- *FREQ* WAS *CONGESTED*

A QUORUM phrase search would also be useful for finding narratives relevant to "frequency congestion". The preceding phrases suggest that an effective search would use a variety of phrase forms as queries, including:

- FREQ CONGESTION
- FREQ CONGESTED
- CONGESTION FREQ
- CONGESTED FREQ

One might also want to include the plural form, "freqs".

- FREQS CONGESTION
- FREQS CONGESTED
- CONGESTION FREQS
- CONGESTED FREQS

Using QUORUM phrase search

QUORUM phrase search makes it easy to find ASRS incident narratives that contain phrases of interest. As examples, and to illustrate some important considerations, several phrase searches are presented here, including: "conflict alert", "frequency congestion", "cockpit resource management", "similar sounding callsign(s)", and "flight crew fatigue". These examples are representative of phrase searches that would be useful to the ASRS (Frank, 1998).

Searching for "conflict alert"— The simplest phrase search uses a single phrase as the query. This can be helpful when looking for a thing, concept, or action that is expressed using multiple words, such as "conflict alert." A "conflict alert" is "A function of certain air traffic control automated systems designed to alert radar controllers to existing or pending situations recognized by the program parameters that require his immediate attention/action." (DOT, 1982).

A search for the narratives that contain the phrase "conflict alert" is simple. The user merely enters the phrase. QUORUM then displays the most relevant narratives in a Netscape browser window, with instances of the phrase highlighted. Also output are three files and documentation of the search, as described in section "New QUORUM methods", in the subsection "QUORUM phrase search". The files contain the highlighted narratives, a complete list of relevant narratives, and the QUORUM model used to search the phrase database.

Here is one of the most relevant narratives found by QUORUM phrase search:

THIS ASRS RPT IS ADDRESSED TO THE ARTS IIA **CONFLICT ALERT** FEATURE USED IN MANY TRACONS IN THE COUNTRY . THIS FEATURE IS DESIGNED TO BE AN AID TO CTLRS IN PREDICTING IMPENDING CONFLICTIONS OF AIR TFC . THE ACTUAL OP OF THE **CONFLICT ALERT** IS THAT IT DOES NOT ACTIVATE , IN THE MAJORITY OF CASES , UNTIL THE ACFT ARE IN VERY CLOSE PROX OR HAVE ALREADY PASSED EACH OTHER . THE LATEST VERSION (A2.07) BECAME OPERATIONAL LAST MONTH AND THE PROB STILL EXISTS . THE SOFTWARE PROGRAM MUST BE IMMENSE AND I'M SURE THAT IT MUST BE A MONUMENTAL TASK TO DEBUG , HOWEVER , IT MUST BE DONE TO MAKE THE **CONFLICT ALERT** FEATURE A USABLE TOOL FOR CTLRS . A UCR RPT HAS BEEN SUBMITTED TO THE FAA . THE **CONFLICT ALERT** IS SUPPOSED TO PROJECT ACFT COURSES AND RATES OF CLB AND ALARM WHEN AN IMMINENT CONFLICT IS DETECTED . MY PAST EXPERIENCES WITH ARTS III AND ARTS IIIA PROVED THIS TO BE THE CASE . UNFORTUNATELY THE ARTS IIA SYS HAS NEVER FUNCTIONED AS WELL FROM THE ONSET TO THE PRESENT DAY . ARTS IIA VERSION A2.07 IS CURRENTLY IN USE AND THE **CONFLICT ALERT** HAS , IN MY ESTIMATION , LIMITED USE TO THE CTLR AS AN AID IN PREDICTING CONFLICTS . IT FUNCTIONS MORE AS AN IMMINENT COLLISION ALERT OR AN ' AFTER THE FACT ALERT ' (YOU JUST HAD A DEAL) . THE AURAL / VISUAL ALARM DOES NOT ACTIVATE UNTIL THE ACFT ARE IN VERY CLOSE PROX AND IMMEDIATE ACTION IS REQUIRED TO PREVENT A COLLISION , OR THE ACFT HAVE ALREADY PASSED EACH OTHER AND NOTHING CAN BE DONE (EXCEPT TURN YOURSELF IN) !! THE MAJORITY OF DATA CONCERNING **CONFLICT ALERT** ALARMS WAS RECEIVED ON ACFT UTILIZING VISUAL SEPARATION METHODS (WHEN THE SEPARATION IS VASTLY REDUCED) . THE **CONFLICT ALERT** FEATURE COULD BE A VALUABLE SEPARATION TOOL FOR THE CTLR IF IT WERE TO OPERATE AS DESIRED . THIS SHORTCOMING MUST HAVE SURFACED IN THE TESTING OF ARTS IIA BEFORE GOING OPERATIONAL . I ASSUME ' DEBUGGING ' A PROGRAM OF THIS SIZE MUST BE A MONUMENTAL TASK AND THIS IS WHY I HAVE WAITED THIS LONG TO INITIATE THE PAPERWORK . VERSION A2.07 WAS JUST RELEASED IN AUG AND THERE WAS NO CHANGE IN THE OP OF THE **CONFLICT ALERT** FEATURE .
(251367)

Since the phrase "conflict alert" is found in exactly the form of the query, and since there are many occurrences of the phrase, this narrative is considered to be highly relevant.

Searching for "frequency congestion"— A search for the narratives that contain the phrase "frequency congestion" is simple. The user merely enters the phrase "frequency congestion". In the keyword search done earlier on "frequency" and "congestion", however, multiple forms of the phrase "frequency congestion" were found in the ASRS database and others are possible. The forms include:

- FREQ CONGESTION
- FREQ CONGESTED
- CONGESTION FREQ
- CONGESTED FREQ
- FREQS CONGESTION
- FREQS CONGESTED
- CONGESTION FREQS
- CONGESTED FREQS

If the user provides these phrases as the query, QUORUM phrase search finds the narratives that contain one or more of them. QUORUM then displays the most relevant narratives in a Netscape browser window, with instances of the phrase highlighted. Here is one of the highly relevant narratives retrieved by QUORUM:

WE WERE CLRED A CIVET 1 ARR TO LAX . THE ARR ENDS AT ARNES AT 10000 FT WITH THE NOTE ' EXPECT ILS APCH . ' WE WERE SWITCHED TO APCH CTL AROUND ARNES . THERE WAS AN ACFT COMING BACK TO LAND AFTER TKOF AND THUS THE **FREQ** WAS **CONGESTED** . WE WERE BLOCKED ON SEVERAL ATTEMPTS TO CONTACT APCH CTL AND WERE UNABLE TO CHK IN . WE CONTINUED OUR DSCNT MEETING THE ALT CONSTRAINTS FOR ILS RWY 25L . SOMEWHERE AFTER ' FUELRL , ' APCH CTL CALLED US AND TOLD US TO LEVELOFF AT 7000 FT AND THAT WE WERE ONLY CLRED TO 10000 FT . THE QUESTION IS , ' IF YOU ARE UNABLE TO CONTACT APCH CTL , ARE YOU IN A LOST COM SIT ? ' IF YOU LEVELOFF AT ARNES , YOU VERY QUICKLY FIND YOURSELF TOO HIGH TO LAND . DO YOU FLY ALL THE WAY TO THE ARPT AT 10000 FT OR DO YOU FLY THE ILS APCH ? IS **FREQ CONGESTION** A LEGITIMATE LOST COM SIT ? CALLBACK CONVERSATION WITH RPTR REVEALED THE FOLLOWING INFO : RPTR SENT 2 CAPT RPTS TO HIS COMPANY QUESTIONING THE PROC , BUT AS YET , NO ANSWER . HE WAS NOT SURE WHAT WAS HIS CLRNC LIMIT BECAUSE THE CIVET 1 ARR ENDS AT ARNES WITH A NOTE TO ' EXPECT ILS APCH . ' THE RPTR THOUGHT THAT PERHAPS WHEN UNABLE TO OBTAIN APCH CLRNC PRIOR TO ARNES AND IF IT WAS A CLRNC LIMIT , THEN HE SHOULD ENTER HOLDING AS DEPICTED ON THE CHART . TO CLARIFY , THE SOCAL APCH CTLR SUPVR WAS CONTACTED AND HE SAID THAT THE ACFT WAS CLRED TO THE ARPT AS PART OF THE ORIGINAL CLRNC AND THAT THE ARR IS NOT A CLRNC LIMIT . ALSO , THAT THE ACFT MUST MAINTAIN THE LAST ASSIGNED ALT AND , IF APCH CTLR MESSES UP AND DOESN'T GIVE THE APCH CLRNC , THEN THE ACFT IS EXPECTED TO MAINTAIN ALT AND CONTINUE INBOUND ON THE LOC COURSE . THE SUPVR SAID THAT THE ACFT DEFINITELY SHOULD NOT ENTER HOLDING , BUT CONTINUE INBOUND AT THE LAST ASSIGNED ALT . (306082)

This narrative is relevant because it contains two of the query phrases. One is in exact form ("**FREQ CONGESTION**") and one is nearly in exact form ("**FREQ WAS CONGESTED**").

Searching for "cockpit resource management"— A search for the narratives that contain the phrase "cockpit resource management" is simple, but it raises two issues. First, the ASRS uses many abbreviations, and the word "management" is one of the words abbreviated. To save the user from having to know the abbreviations, QUORUM phrase search maps words to ASRS abbreviations. The second issue raised by a search for narratives containing the phrase "cockpit resource management" is the fact that the phrase has more than 2 words. As a consequence, the user has the option to accept narratives containing only part of the phrase. The default, however, is to require that the whole phrase be present in each retrieved narrative.

To find narratives containing the phrase: "cockpit resource management", the user enters that phrase. QUORUM maps the vocabulary of the phrase to that used in the ASRS narratives. In this case, the result is "cockpit resource mgmnt", and this phrase is used as the actual query phrase. QUORUM finds the narratives containing the phrase "cockpit resource mgmnt", and the most relevant narratives are displayed in a Netscape browser window with all instances of the phrase highlighted. Here is an example:

COPLT 'S BRASH ATTITUDE HAD BEEN A SORE SPOT WITH ME ALL MONTH AND REPEATED DISCUSSION WITH HIM HAD FAILED TO ACHIEVE ANY RESULTS . ALTHOUGH I NOTICED EARLY ON THAT HIS PLTING SKILLS DIDN'T JUSTIFY HIS CONFIDENCE LEVEL AND I HAD RECOGNIZED THE NEED TO CONTINUALLY MONITOR HIS PERF , I HAD TO TAKE MY EYES OFF OF HIM FOR ABOUT 2 MINS (2 MINS ! !) . IN THAT PERIOD OF TIME HE DEVIATED OFF OUR RTING BY ABOUT 8 MI PROMPTING AN INQUIRY FROM ZAU . THE FO 'S ATTITUDE WAS ' OK , I MADE A MISTAKE - - SO WHAT ? ' I BELIEVE (DUE TO INTERACTING WITH THIS INDIVIDUAL ON PREVIOUS TRIPS) THAT HE FELT HIS ROLE IN THE COCKPIT WAS ONE OF DECISION MAKER . ALTHOUGH I EXPLAINED TO HIM THAT WE WERE A TEAM , AND EACH MEMBER OF THE TEAM WAS ESSENTIAL TO OUR SAFETY , IT IS IN THE CAPT 'S JOB DESCRIPTION AS BEING THE FINAL AUTHORITY AS TO THE OP OF THE FLT . WITH THE ADVENT OF **COCKPIT RESOURCE MGMNT** I'VE NOTICED A TENDENCY WITH SOME FO 'S TO IGNORE THE FACT THAT THERE IS A HIERARCHY WITHIN THE COCKPIT , TO THE POINT OF CONSIDERING THEMSELVES AUTONOMOUS (AS IN THIS EXTREME CASE) . WHILE THE INTENT OF **COCKPIT RESOURCE MGMNT** IS OK , I MUST SAY THAT THE CREW 'S RELATIONSHIP WITH THE CAPT IS ONE OF ORDINATE - SUBORDINATE , AND **COCKPIT RESOURCE MGMNT** TENDS TO OVERLOOK OR MINIMIZE THIS CONCEPT . IF MY ASSESSMENT IS CORRECT , **COCKPIT RESOURCE MGMNT** SHOULD BE MODIFIED TO REFLECT THE REALITIES OF LINE OPS . (222230)

The narratives considered to be the most relevant are the ones that have the best and the most matches to the query phrase. If the user were also interested in narratives that might contain only a fragment of the phrase, such as "resource management", an option can be invoked to allow it. In that case, narratives containing only fragments of the phrase would be added at the bottom of the list of relevant narratives.

Here are some example excerpts from narratives containing only fragments of the phrase "cockpit resource management":

THIS AIRLINE HAS EXERTED A LOT OF ENERGY TO PROMOTE CREW **RESOURCE MGMNT** , BUT ALL OF MY EFFORT TO PROVIDE USEFUL INPUT FAILED . ALL DURING THIS INCIDENT I WAS WELL AWARE OF PREVIOUS ACCIDENTS IN WHICH NO ONE CHALLENGED THE CAPT AS HE MADE IMPROPER DECISIONS . I WANTED TO MAKE SURE THAT THIS WOULD NOT HAPPEN DUE TO MY INACTION . I DISCOVERED MY LIMITATIONS IN THE FACE OF A CAPT WHO MADE IMPROPER DECISIONS . (279099)

FO IS LOW TIME AND [CAPT] ADMITS HE EXERCISED POOR **COCKPIT MGMNT** . SHOULD HAVE INSISTED THAT FO HELP WITH TAXI VIGILANCE . (202096)

... NEW HIRES OFTEN BITE THEIR TONGUES RATHER THAN CONFRONT CAPTS ABOUT **COCKPIT CREW MGMNT** PROBS , BECAUSE OF THE POSSIBILITY OF A NEGATIVE EVALUATION BEING SENT TO THE COMPANY , WHICH COULD EFFECT YOUR BEING KEPT ON THE JOB BEYOND PROBATION . MY RELUCTANCE TO WORK THIS OUT CAUSED ME TO PUT UP WITH A COCKPIT ENVIRONMENT THAT WAS LESS THAN SATISFACTORY . (143981)

LACK OF TRAINING COVERING **COCKPIT MGMNT** RESOURCES . (206734)

COCKPIT RESOURCES **MGMNT** HAS HELPED IN THE ACFT ; MAYBE MORE PERSONAL CONTACT BTWN ATC AND PLTS WOULD DO THE SAME . (141625)

The benefit of matching phrase fragments is that a greater number of relevant reports can be found, even when the author of the narrative didn't get some standard phrase exactly right. Some of these reports can be highly relevant to the topics of interest.

Searching for "similar sounding callsign"— A search for the narratives that contain the phrase "similar sounding callsign" is easy for the user to accomplish, but it raises three issues.

The first issue is that the ASRS uses various forms of some words and phrases. Sometimes "call sign" is used, while other times "callsign" is used. Similarly, "descent" is sometimes abbreviated as "dscnt" while other times it is "dsnt". And there are other such examples. To achieve consistency, QUORUM standardizes usage. This is accomplished using the same mapping technique that is applied to handle ASRS abbreviations. That is, the various forms of some terms are mapped to standard forms. Since "call sign" is more common, that is the form used consistently by QUORUM. Thus, "callsign" is mapped to "call sign". Similarly, "callsigns" is mapped to "call signs".

The second issue involves singular and plural forms of phrases. Specifically, if a user is interested in the singular form, the plural form is often of interest as well, and vice versa. In this case, the user might want to find not only the narratives containing the phrase "similar sounding call sign" (singular), but also those containing "similar sounding call signs" (plural). Rather than make any assumptions, QUORUM requires the user to specify each of the forms to be found.

The third issue raised by this search involves QUORUM's ranking of narratives when searching for long and/or multiple phrases. In the case of "similar sounding call sign(s)", some narratives will contain both singular and plural forms of the phrase. Some narratives will contain only one of the forms. Some narratives will contain only fragments, such as "similar call sign", or "call signs". QUORUM's rank ordering of narratives containing these various forms is done in the order just described, as will be shown. This seems to be a useful order, as it is in accordance with an intuitive sense of what constitutes a good match to the query phrases.

In order to perform the search, the user enters the phrases "similar sounding callsign" and "similar sounding callsigns". QUORUM then displays the most relevant narratives in a Netscape browser window, with instances of the phrases highlighted. Here are excerpts from some of the most relevant narratives:

BECAUSE WE HAD BEEN ON TWR FREQ FOR SO LONG , WE HAD NO AWARENESS OF THE OTHER ACFT WITH A **SIMILAR CALL SIGN** THE FOLLOWING ARE CONTRIBUTING FACTORS . **SIMILAR SOUNDING CALL SIGNS** DURING SIMULTANEOUS INTERSECTING RY DEPS , EXTREME CARE SHOULD BE TAKEN WITH ACFT HAVING LIKE **CALL SIGNS** THEY HAD MISUNDERSTOOD TKOF CLRNC FOR AN ACFT WITH A **SIMILAR SOUNDING CALL SIGN** , ON ANOTHER RWY . (198106)

WHILE INBOUND TO DTW METRO ARPT FROM KALAMAZOO , MI , ON COMPANY XX50 THERE WERE 2 OTHER COMPANY FLTS : COMPANY XX53 AND COMPANY X50 WITH **SIMILAR SOUNDING CALL SIGNS** AS OURS APPARENTLY WE WERE FOLLOWING A CLRNC FOR AN ACFT OF A **SIMILAR SOUNDING CALL SIGN** . I DID READ BACK THE ORIGINAL CLRNC WITH OUR OWN **CALL SIGN** , HOWEVER . THERE WAS MUCH CONFUSION WITH **SIMILAR CALL SIGNS** . (192640)

I VERIFIED THE ALT AND FREQ AS BEING CORRECT BUT DID NOT CATCH THE **CALL SIGN** ALTHOUGH I DID NOT CLARIFY THE CORRECT **CALL SIGN** ... I CANNOT IMAGINE WHY ANY PLT WOULD CLB WITHOUT QUESTION WHEN HE HAD JUST BEEN ISSUED 2 CONVERGING TARGETS AT ALTS ABOVE HIM WE WERE INFORMED BY OUR UNION SAFETY CHAIRMAN THAT WE HAD ACCEPTED THE 13000 FT CLB AND FREQ CHANGE FOR ANOTHER FLT , ACR X , WITH A **SIMILAR SOUNDING CALL SIGN**
CORRECTIVE ACTION : REDUCE , IF NOT ELIMINATE , **SIMILAR SOUNDING CALL SIGNS** . (255236)

HE THEN STATED HE HAD ANOTHER COMPANY WITH A **SIMILAR SOUNDING CALL SIGN** ON THE FREQ ... THIS SAME CTLR WAS ALSO WORKING 2 OTHER PAIRS OF OUR COMPANY FLTS WITH **SIMILAR CALL SIGNS** ... MULTIPLE FLTS WITH **SIMILAR SOUNDING SIGNS** IN TODAY 'S CONGESTED ATC ENVIRONMENT IS DANGEROUS , AND OUR COMPANY HAS A BAD PRACTICE OF DOING THIS . I BELIEVE THEY DO IT FOR MARKETING REASONS , BUT RUNNING BANKS OF FLTS INTO A HUB AT PEAK HRS WITH **SIMILAR SOUNDING CALL SIGNS** IS NOT A GOOD PRACTICE , AND SHOULD BE STOPPED , THUS HELPING TO AVOID SOMEONE FROM MISUNDERSTANDING AND TAKING SOME OTHER FLT 'S CLRNC . THIS HAS THE POTENTIAL TO CREATE A VERY SERIOUS SIT . THIS **CALL SIGN** USAGE BY OUR COMPANY HAS RAISED THE IRE OF MANY PLTS , BUT OUR COMMENTS AND COMPLAINTS HAVE FALLEN ON DEAF EARS AT THE COMPANY . (236716)

THIS WAS A **SIMILAR ENOUGH SOUNDING CALL SIGN** THAT I BELIEVE SOME EFFORT SHOULD BE MADE TO DISTINGUISH BTWN THEM FLT # S SHOULD BE READ READ DIGIT BY DIGIT AND WARNINGS SHOULD BE ISSUED FOR **SIMILAR SOUNDING CALL SIGNS** . (173196)

PROBS THAT NEED TO BE IDENTED : TOO MANY **SIMILAR SOUNDING CALL SIGNS** BY SAME COMPANY IN SAME VICINITY AT THE SAME TIME NO ONE HAD SAID THERE WAS AN ACFT ON FREQ WITH A **SIMILAR CALL SIGN** AND WE HAD HEARD NO CALLS TO COMPANY ACR . WHEN THE FIRST CALL WAS MADE , THE FO WAS DISTR BY A FLT ATTENDANT IN THE COCKPIT ASKING ABOUT THE TEMP OF THE CABIN AND HE DID NOT HEAR THE **CALL SIGN** READ BY CTR . SUPPLEMENTAL INFO FROM ACN 224896 : OUR **CALL SIGN** SAME COMPANY ACR SIMILAR TO ACR X ... (224992)

The narratives considered to be the most relevant to multiple query phrases are the ones that best match, in whole or in part, the query phrases. The following observations illustrate the quality of the phrase matches relative to the rank ordering of the narratives.

- The narratives ranked 1-4 contain both of the query phrases: "similar sounding call sign" and "similar sounding call signs". Phrase fragments are also found in these narratives, including one or more of: "similar call sign(s)", "similar sounding sign(s)", or "call sign(s)".
- Narratives ranked 5-86 contain one or the other of the query phrases: "similar sounding call sign" or "similar sounding call signs". Narratives in this group usually also contain one or more of the phrase fragments: "similar call sign(s)" or "call sign(s)". Less common additions include: "similar enough sounding call sign", "similar to the call signs", "similar acft call signs", "similar-sounding but incorrect ident", and "like sounding call signs".
- Narratives ranked 87-91 contain one of the following: "similar sounding call sign", "similar sounding call signs", one of those phrases but with inclusions, or a collection of phrase fragments that, taken together, conveys the notion of "similar sounding call sign(s)". For example, the 87th narrative contains only "similar sounding acft call signs", and the 88th contains only "similar sounding flt numbers", "wrong call sign", and "similar call signs".
- Narratives 92-181 do not contain the whole phrase. Most of them (83) contain the fragment "similar call sign(s)", usually with some other fragments such as "call sign(s)" or "similar sign(s)". The other seven narratives include fragments containing "sounding" but not "similar", e.g., "close sounding or transposable call signs".
- Narratives 182-200 contain only the fragments "similar call sign(s)" or "call sign(s)". Narrative 182 is the highest ranking narrative that contains only the fragment "call sign(s)".
- Most of the many narratives beyond the 200th in rank contain only "call sign(s)".

In summary, the rank ordering of the narratives provided by QUORUM phrase search for long, multiple query phrases is appropriate. The highest ranked narratives (1-86) contain one or more instances of the query phrases "similar sounding call sign" and "similar sounding call signs", while a transition group (87-91) at least conveys the notion of the query. The next large group (92-181) mostly contains "similar call sign(s)", which is more general than "similar *sounding* call sign(s)", but represents the next best match to the query. These are followed by a large group of narratives (increasingly common beginning with 182) that contain only "call sign(s)", which is more general than "*similar* call sign(s)", but represents the next best match to the query.

Here are the accession numbers of the 91 ASRS incident reports that are most relevant to the phrase "similar sounding callsign(s)":

1. 236716	17. 104418	33. 142569	49. 342497	65. 273212	81. 192059
2. 192640	18. 333433	34. 144569	50. 94979	66. 286220	82. 160883
3. 198106	19. 246229	35. 89654	51. 339600	67. 173641	83. 262477
4. 255236	20. 361796	36. 139469	52. 90769	68. 298130	84. 105298
5. 173196	21. 364467	37. 136784	53. 152083	69. 299673	85. 133520
6. 144720	22. 259010	38. 334890	54. 142766	70. 120463	86. 266870
7. 273139	23. 337485	39. 332500	55. 217142	71. 304066	87. 108119
8. 269000	24. 268344	40. 210935	56. 230971	72. 304370	88. 85247
9. 95030	25. 165761	41. 146441	57. 160848	73. 178788	89. 92664
10. 310278	26. 93653	42. 206733	58. 308996	74. 82543	90. 217637
11. 224992	27. 202997	43. 86887	59. 307837	75. 325390	91. 266124
12. 249451	28. 150627	44. 158878	60. 306664	76. 249352	
13. 370586	29. 374529	45. 246471	61. 282179	77. 328055	
14. 143173	30. 347810	46. 201843	62. 112496	78. 248464	
15. 366360	31. 351689	47. 343091	63. 276472	79. 135501	
16. 139993	32. 343860	48. 342960	64. 109765	80. 330230	

Searching for "flight crew fatigue"— A user might consider searching for the phrase "flight crew fatigue", but the results would be less than satisfactory due to the small number of matched narratives. Only 8 of 67821 ASRS reports contain the phrase "flt crew fatigue". This small number does not, however, reflect the true prevalence of narratives involving flight crew fatigue.

The user might then consider limiting the search to the phrase "crew fatigue". A larger number of narratives contain that phrase. Among 67821 ASRS reports, a total of 102 narratives contain "crew fatigue", and an additional 9 contain phrases such as "crew's fatigue", "crew member fatigue", or "crew mental fatigue". This does not, however, reflect the true number of narratives on the subject.

Rather than doing a phrase search in this case, a keyword search on "fatigue" would be more effective. Even better would be a search on "fatigu", which would match "fatigue", "fatigued", and "fatiguing". To increase the probability that the retrieved narratives involve flight crew fatigue, the search should be limited to the subset of the reports that were submitted by flight crews. In a QUORUM keyword search on "fatigu" among 36361 reports submitted by the flight crews of large aircraft there were 743 relevant narratives. A search among 67821 ASRS reports of all kinds found 1364 narratives relevant to "fatigue", "fatigued", or "fatiguing".

Narratives that contain the topic of fatigue do not necessarily contain the words "fatigue", "fatigued", or "fatiguing". A method is shown in the section "Using QUORUM phrase discovery" that addresses this issue. That method finds a large number of fatigue-related phrases such as "duty time", "crew rest", etc. The process of finding these phrases also finds ASRS reports that contain the topic of fatigue even if no forms of the word "fatigue" are present in the narratives.

Searching for a particular sentence that occurs only once in the database— Since QUORUM represents phrases implicitly among the contextual relations of the documents, rather than explicitly as a pre-computed list, it is possible to find any phrase, even if it occurs only once. In addition, even though contextual relations in the phrase database are limited to spans of 4 words, indirect chains of relations allow longer phrases to be found (see appendix 3, step 6).

As an example, a sentence was selected from one of the incident narrative excerpts in McGreevy & Statler (1998), page 68, in the section, "Distracted by automation." The selected sentence is from report number 368360:

THE ENTIRE CREW WAS DISTR, AND WE BOTH FAILED TO MONITOR THE PERF OF THE ACFT.

A QUORUM phrase search was done using the whole sentence as the query. The search was repeated using the query:

The entire crew was distracted, and we both failed to monitor the performance of the aircraft.

Given either query, QUORUM phrase search identifies the relevant narrative and launches Netscape to display it with the relevant sections highlighted. Shown below is an excerpt. Notice that the query sentence is highlighted, as are additional fragments of the sentence.

I BELIEVE THAT THE COMPLEXITY OF FMS PROGRAMMING IS NOT ADDRESSED IN INITIAL TRAINING AT SCHOOL BECAUSE EACH ACFT HAS DIFFERENT EQUIP . HOWEVER , THIS LEAVES **THE FLT CREW** TO ' LEARN AS THEY FLY . ' THIS EFFECTIVELY TOOK MY FO OUT OF THE LOOP IN THAT IF HE WAS PROGRAMMING THE FMS , I COULD HAVE CONCENTRATED MORE ON MONITORING **THE ACFT** . I SHOULD HAVE LET THE FO FLY **THE ACFT** WITH THE AUTOPLT RATHER THAN ME DO ALL THE TASKS . **THE ENTIRE CREW WAS DISTR , AND WE BOTH FAILED TO MONITOR THE PERF OF THE ACFT** . I SHOULD HAVE JUST PUT MY HSI IN THE VOR MODE RATHER THAN DISPLAY FMS COURSE INFO . THIS WOULD HAVE ALLOWED US TO FOCUS MORE ON **THE ACFT** . (368360)

By doing the search using the option to include narratives containing only some of the fragments of the sentence, some near-matches can also be found. These are ranked as less relevant than the one containing the whole sentence. Here are excerpts from narratives containing only fragments of the sentence:

I WAS **DISTR** BY THE CAPT 'S CONVERSATION **AND WE BOTH FAILED TO MONITOR THE ACFT** 'S DSCNT . (265142)

WHILE **WE** CONTINUED **TO** WONDER WHY **THE** DSCNT DID NOT OCCUR AS PROGRAMMED , IT WAS OBVIOUS THAT **WE** HAD **BOTH FAILED TO MONITOR THE** DSCNT AS WE SHOULD HAVE . (253696)

WE WERE CLRD FOR THE OXI 2 ARR , FWA TRANSITION TO ORD , FO FLYING **THE ACFT** ALTHOUGH WE HAD TUNED THE OXI 095 DEG RADIAL FOR THE TURN AT SPANN INTXN , **WE FAILED TO** TURN BECAUSE OF OUR DISTR THE FO AND I DO NOT BELIEVE THAT WE MISSED A RADIO CALL , EVEN THOUGH WE WERE **DISTR AND** WERE OFF COURSE I BELIEVE THAT MY FAILURE **TO MONITOR THE** FO 'S NAV WHILE I INVESTIGATED POSSIBLE ACFT ABNORMALITIES WAS THE MOST IMPORTANT CONSIDERATION IN THIS OCCURRENCE . (201659)

This example shows the ability of QUORUM phrase search to find long or rare phrases, while also finding similar text if desired.

(Note: In the final excerpt above, the first occurrence of DISTR is not highlighted because it is not sufficiently phrasally related, within the narrative, to any of the words to which it is related in the query sentence. This lack of sufficient relatedness in the narrative is also true for some other words found in both the query and the excerpts.)

QUORUM phrase generation and QUORUM phrase discovery

The use of any phrase search tool requires the user to know or guess what phrases are likely to be in the database being searched. QUORUM phrase generation and QUORUM phrase discovery can show the user what phrases are likely to be useful queries. In addition, these tools can help the user to explore and understand the particular nuances of topics in the database.

QUORUM phrase generation differs from QUORUM phrase discovery. QUORUM phrase generation assembles phrases from word pairs that are often found in a particular order and close together in the narratives of the ASRS database. That is, the phrases are assembled from QUORUM phrase models. Many of the generated phrases are present in the narratives, and phrases are listed in order of their estimated frequency in the whole ASRS database. QUORUM phrase generation is a useful way of building phrases that are typically present, without actually storing and retrieving the phrases themselves. In contrast, QUORUM phrase discovery scans narratives for all possible phrases and distills them down to those which are topically relevant, as will be shown in the section "Using QUORUM phrase discovery".

Using QUORUM phrase generation

QUORUM phrase generation is used to show typical phrases that contain words or phrases of interest. The default is to produce the 10 most typical phrases, but the user has the option of specifying a different number. The output phrases can be used as query phrases for input to QUORUM phrase search.

Generating phrases containing "rain"— As a simple example, suppose a user wished to know what phrases contain the word "rain". Given the word "rain", and using the option to specify the number of generated phrases (30 in this case), QUORUM phrase generation produces the following list:

- LIGHT RAIN
- HVY RAIN
- RAIN SHOWERS
- FREEZING RAIN
- MODERATE RAIN
- LIGHT MODERATE RAIN
- HEAVY RAIN
- RAIN SHOWER
- RAIN FOG
- MODERATE HVY RAIN
- ENCOUNTERED RAIN
- ENCOUNTERED MODERATE RAIN
- ENCOUNTERED LIGHT RAIN
- ENCOUNTERED LIGHT MODERATE RAIN
- RAIN TURB
- MODERATE RAIN TURB
- LIGHT RAIN TURB
- ENCOUNTERED RAIN TURB
- LIGHT MODERATE RAIN TURB
- ENCOUNTERED MODERATE RAIN TURB
- ENCOUNTERED LIGHT RAIN TURB
- ENCOUNTERED LIGHT MODERATE RAIN TURB
- VISIBILITY RAIN
- VISIBILITY RAIN FOG
- VISIBILITY LIGHT RAIN
- TURB RAIN
- TURB ENCOUNTERED RAIN
- MODERATE TURB RAIN
- LIGHT TURB RAIN
- ENCOUNTERED TURB RAIN

The phrases toward the beginning of the list are the ones that appear more often in the narratives of the ASRS database. So, for example, "light rain" is more common than "moderate rain". Similarly, "hvy rain" is more common than "heavy rain". Some of the listed phrases, such as "light rain", typically appear in narratives exactly as shown. Other listed phrases, such as "light moderate rain", typically appear in narratives with other words intermixed. For example, the most common appearance of "light moderate rain" is "light to moderate rain".

The user has the option of eliminating phrases containing words that are not of interest at the moment. This is done by identifying such words as additions to the default stoplist. For example, the user could add the words LIGHT, MODERATE, ENCOUNTERED, TURB (i.e., turbulence), and CONDITIONS to eliminate the many variations on these themes. When re-running phrase generation with the expanded stoplist, a revised list of phrases is generated.

The user also has the option of allowing a number of stopwords within each phrase. To avoid generating an excessive number of similar phrases, however, the default is to display only those phrases that contain no stopwords. Otherwise, given the query word "rain", many phrases like the following would be output:

- THE LIGHT RAIN
- A LIGHT RAIN
- SOME LIGHT RAIN
- WAS LIGHT RAIN
- ANY LIGHT RAIN
- THE HVY RAIN
- A HVY RAIN
- SOME HVY RAIN

QUORUM phrase generation can also find phrases that contain other phrases. For example, given the query "freezing rain", these and other phrases would be generated:

- FREEZING RAIN
- LIGHT FREEZING RAIN
- FREEZING RAIN CONDITIONS
- LIGHT FREEZING RAIN CONDITIONS
- MODERATE FREEZING RAIN
- MODERATE FREEZING RAIN CONDITIONS
- LIGHT MODERATE FREEZING RAIN
- MODERATE LIGHT FREEZING RAIN
- MODERATE LIGHT FREEZING RAIN CONDITIONS
- LIGHT MODERATE FREEZING RAIN CONDITIONS
- FREEZING RAIN DRIZZLE
- LIGHT FREEZING RAIN DRIZZLE

When using QUORUM phrase generation, user query words are mapped (if necessary) to ASRS abbreviations and usage. For example, "runway" is mapped to "rwy". Most words, however, do not need to be mapped.

Searching for narratives containing "light moderate rain"— Any phrase can be used as input to QUORUM phrase search, including those produced by QUORUM phrase generation. For example, one might search for the phrase "light moderate rain". Here are excerpts from some of the most relevant narratives:

CONTRIBUTING FACTORS - *LIGHT TO MODERATE RAIN* WAS FALLING IN THE JFK AREA WITH STANDING WATER ON RAMP SURFACES - THIS COUPLED WITH LIGHTING ON THE CONCOURSE CAUSED A GLARE ON THE RAMP MAKING VIEW OF THE LEAD - IN LINE DIFFICULT . (86853)

THERE WERE LARGE AREAS OF *LIGHT TO MODERATE RAIN* SHOWERS AROUND THE LAX AREA THE GPWS SOUNDED ... I SUSPECT THIS WAS CAUSED BY THE EFFECT OF THE RAIN SHOWER ON THE GPWS . (233843)

JUST PRIOR TO FLYING INTO THE HAIL , ATC ASKED WHAT MY CONDITIONS WERE AND I RPTED *LIGHT TO MODERATE RAIN* . (373915)

The exact phrase "light moderate rain" never appears, but the phrase "light to moderate rain" is common. This shows the value of the flexible phrase matching available with QUORUM phrase search. Of course, the phrase "light to moderate rain" could itself be used as a query phrase.

Searching for narratives containing common "rest" phrases— It is often helpful to use multiple phrases from the list produced by QUORUM phrase generation as input to QUORUM phrase search. For example, if the user were unsure of what phrases typically contain the word "rest" as it relates to fatigue, the phrase generation program could be used to list the most common phrases containing the word "rest". These would include, in order of estimated prominence in the ASRS database:

- REST FLT (e.g., "rest of the flight")
- REDUCED REST
- CREW REST
- REST PERIOD
- CAME REST (e.g., "came to rest")
- MINIMUM REST
- REST REQUIREMENTS
- REST PERIODS
- REST APCH (e.g., "rest of the approach")
- MINIMUM REST APCH
- REST APCH FLT
- ACFT REST
- ACFT REST FLT
- ACFT REST APCH
- ACFT CAME REST
- ACFT REST APCH FLT
- REST TRIP
- CREW ACFT REST
- ADEQUATE REST
- etc.

Given an interest in "rest" as it relates to fatigue, the user would ignore "rest flt", "came rest", and other phrases unrelated to fatigue, and would select the fatigue-related phrases. To simplify the selection task, the user could list the words ACFT, CAME, APCH, TRIP, and perhaps others as extra stopwords and then re-run the phrase generation program. The fatigue-related phrases, such as those shown below, could be used as input to QUORUM phrase search:

- REDUCED REST
- CREW REST
- REST PERIOD
- MINIMUM REST
- REST REQUIREMENTS
- REST PERIODS
- ADEQUATE REST
- REQUIRED REST
- MINIMUM REQUIRED REST
- REST OVERNIGHT
- REQUIRED CREW REST
- PROPER REST
- REST PRIOR
- CREW REST PRIOR
- SCHEDULED REST
- REST PRIOR FLT
- LEGAL REST
- MINIMUM REST REQUIREMENTS
- COMPENSATORY REST
- REST NIGHT
- REST BREAK
- MINIMUM CREW REST
- REQUIRED REST PRIOR
- MINIMUM REQUIRED CREW REST
- REQUIRED REST PRIOR FLT
- REQUIRED CREW REST PRIOR
- LACK REST
- REST NIGHT PRIOR
- LACK PROPER REST
- LACK CREW REST
- LACK ADEQUATE REST

A QUORUM phrase search on these phrases retrieves narratives containing one or more of them. The most relevant narratives contain a greater variety of the more common phrases. Since QUORUM phrase generation was used to suggest the list of phrases, the user is assured that there are, in fact, narratives containing one or more of them. Here are excerpts from some of the narratives that are most relevant to the "rest" phrases:

AFTER A NUMBER OF YRS AS BOTH A MIL AND COMMERCIAL CARRIER PLT I'VE FOUND THAT EVERYONE 'S BODY NEEDS A ROUTINE , AND RADICAL CHANGES CAN ADVERSELY AFFECT ONE 'S PERF AND ABILITY TO GET ADEQUATE SLEEP DURING THE SUPPOSED *REST PERIOD* . OUR AIRLINE 'S SCHEDULING DEPT OPERATES UNDER CRISIS MGMNT DUE TO OUR MGMNT 'S ' STAFFING STRATEGY , ' AND THUS REQUIRES MANY RESERVE CREW MEMBERS TO COVER MORE THAN 1

SCHEDULED TRIP IN A CALENDAR DAY AND THUS WE HAVE A LARGE NUMBER OF ' **SCHEDULED REDUCED REST PERIODS** ' WHICH ARE 8 HRS , WHICH DOES NOT INCLUDE TRANSPORTATION LCL IN NATURE , WHICH , IN REALITY , REDUCES YOUR TIME AT A REST FACILITY WELL BELOW 8 HRS , PROVIDED YOU FALL TO SLEEP AS SOON AS YOU ARRIVE AT THE HOTEL . MY TRIP / RERTE FROM HELL STARTED AS A 3 DAY WITH AN 8 HR **REST** THE FIRST **NIGHT** WITH AN EARLY RPT . I HAPPENED TO BE COMING OFF A COUPLE OF NIGHT TRIPS AND THE EARLY MORNING RPT HAD ME A LITTLE OUT OF SYNC . WHEN WE ARRIVED AT OUR NEXT OVERNIGHT STATION , WHICH WE WERE **SCHEDULED COMPENSATORY REST** , I FELL ASLEEP EARLY NOT BEING ACCUSTOMED TO EARLY MORNING RPTS AND THUS WOKE VERY EARLY ON THE THE THIRD DAY THE FAA NEEDS TO RECOGNIZE THE IMPORTANCE OF QUALITY **CREW REST** AND IMPLEMENT GUIDELINES TO PREVENT SUCH SCHEDULING PRACTICES . (254345)

CREW HAD A LEGAL DUTY DAY , BUT LAST 2 DAYS CREW HAD BEEN ON **REDUCED REST** WITH **COMPENSATORY REST** TO **MINIMUM** ALLOWED . **CREW** WAS EXTREMELY FATIGUED DUE TO MIN **LEGAL REST** AND RATHER LENGTHY DUTY DAY . CREW HAD BEEN ON DUTY OVER 12 HRS . SUGGESTIVE ACTION : INCREASE **REST PERIODS** . MIN **REST PERIODS** ARE ADEQUATE PROVIDED YOU AREN'T FLOWN TO THOSE MINS 6 DAYS IN A ROW . IT 'S SIMPLY TOO FATIGUING . THERE WERE MANY SIMPLY MISTAKES MADE THIS FLT , ETC . MISSED CALLS , MISUNDERSTANDING HDG / ALT ASSIGNMENT / FREQ CHANGES . MOST OF THESE ERRORS WERE CAUGHT BY ONE OF THE CREW , THE ALT DEVIATION ON THE LAST LEG OF A 13.2 HR DUTY DAY WITH **MINIMUM REQUIRED REST** WAS JUST UNAVOIDABLE . PLEASE RESEARCH INCREASED **REQUIRED REST PERIODS** . (123335)

PRIOR TO DEPARTING ON THE LAST FLT OF DAY 2 , I BECAME CONCERNED ABOUT THE **REQUIRED CREW REST** , SINCE WE WERE BEING DELAYED BY MAINT . I KNEW THAT , THOUGH WE HAD 9 HRS **REST** THE PREVIOUS **NIGHT** , ONCE WE EXCEEDED 15 HRS DUTY TIME OUR REST FOR THE 24 HR " LOOKBACK " WOULD BE LESS THAN NORMAL . MY QUESTION WAS THIS : COULD I ACCEPT **REDUCED REST** ON THE SECOND NIGHT , SINCE I WAS STILL FLYING WHAT WAS SCHEDULED , OR DID WE NEED **COMPENSATORY REST** BECAUSE OF WHAT WAS ACTUALLY FLOWN ? I CALLED OUR COMPANY 'S HEAD OF (MY ACFT) TRNING AND EXPLAINED ABOUT MY SIT . HE STATED THAT , WHILE HE FELT I NEEDED **COMPENSATORY REST** , REPEATED DISCUSSIONS WITH OUR VP OF OPS INDICATED THAT THE COMPANY 'S POS WAS THAT **REDUCED REST** WAS LEGAL . BASED ON THAT , I WENT WITH **REDUCED REST** . ON COMPLETION OF THE TRIP I TALKED TO OUR DIRECTOR OF OPS , WHO PRODUCED A MEMO FROM OUR VP OF OPS . THE MEMO SUMMARIZED AN FAA RULING DATED 7 / 89 STATING (AGAIN , AS I UNDERSTAND IT) THAT **REQUIRED REST** IS BASED ON ACTUAL FLT TIME AND DUTY TIME DURING THE PREVIOUS 24 HRS . COMMUTER AIRLINES ROUTINELY USE THE DUTY TIME REGS AS A GOAL TO ACHIEVE MAX UTILIZATION OF PLTS . YET , I HAVE NOT MET A SINGLE LINE PLT THAT FULLY UNDERSTANDS THIS REG . AS AN EXAMPLE , NO LINE PLT I ASKED KNEW THE ANSWER TO MY QUESTION . WHY IS THIS REG SO UNNECESSARILY SUBTLE ? (145545)

These narratives contain a variety of the more prominent "rest" phrases, such as "reduced rest", "crew rest", and "rest periods". In the first of these narratives (254345), the phrases "scheduled reduced rest periods" and "scheduled compensatory rest" are also among the highlighted "rest" phrases, despite the fact that these phrases do not appear in their entirety among the query phrases. Instead, they match several of the query phrases, including "scheduled rest", "reduced rest", "rest periods", and "compensatory rest". This indicates the flexibility of QUORUM phrase search in highlighting larger phrases of interest built up from smaller ones.

The combination of QUORUM phrase generation and QUORUM phrase search provides the ability to avoid ambiguities in searches. An advantage of this method with a topic like "rest" is that it can focus on the uses of the word "rest" that involve fatigue, while avoiding others. A keyword search would sometimes retrieve narratives involving only "rest of the flight", "came to rest", etc. Without QUORUM phrase generation, the user would not know what phrases contained the word "rest", and so could not effectively use QUORUM phrase search to focus on the kinds of "rest" that are of interest. Using QUORUM phrase generation, however, the user can find topical phrases for use as queries in QUORUM phrase search, and thus find narratives that are focused on the topic of interest. In even more refined searches, the user can select just those phrases that represent particular nuances of the topic of interest, and can use that selection as a query to QUORUM phrase search. The retrieved narratives will reflect the desired nuances of the topic of interest.

QUORUM phrase generation also supports domain analysis and taxonomy development by showing prominent variations among topically related phrases. The "rest" phrases, for example, provide the analyst with a variety of variations on the concept of "rest", such as "reduced rest" and "compensatory rest", which, as the third narrative shows, have very particular meanings. With that insight, an analyst could then use QUORUM phrase search to find other narratives containing "reduced rest" and/or "compensatory rest" to further explore the implications of these issues on crew performance and operational safety.

Using QUORUM phrase discovery

QUORUM phrase discovery scans narratives to find phrases that are related to topics of interest. This is very different from QUORUM phrase generation, which uses phrase models to build likely phrases on a given word or phrase.

In the example shown here, phrases related to "fatigue" are discovered. These include, for example: "rest period", "continuous duty", "crew scheduling", "reserve or standby", "crew fatigue", and "continuous duty overnight". Unlike generated phrases, discovered phrases are not required to contain any of the query words.

For this example, the phrase discovery process began with a keyword search on the words: "fatigue", "fatigued", "fatiguing", "tired", "tiredness", "sleep", "asleep", "sleeping", "sleepy", and "circadian". The particular forms of these words were suggested by reviewing the vocabulary used in the narratives of the ASRS database. The phrase discovery process ultimately produced a collection of relevance-ranked narratives and a list of phrases that are topically-related to "fatigue".

The essence of the QUORUM phrase discovery method is described in the section "Overview of QUORUM methods" in the subsection "QUORUM phrase discovery". A step-by-step description is shown in appendix 5.

The following table shows 50 of 420 phrases related to the topic of fatigue. The 420 phrases were extracted from three sets of 200 narratives that were found to be most relevant to the topic of fatigue. The frequency of each phrase within a set of 200 narratives is shown in the first column. This list shows, for example, that in the context of fatigue, "rest period(s)", "reduced rest", and "crew rest" are the most prominent concerns. Further, these are greater concerns than "continuous duty", "duty period", and "crew duty". The list also shows that "crew scheduling" ranks high among the concerns of the reporters in the context of fatigue. Other prominent concerns include: "reserve or standby", "rest requirements", "crew fatigue", "continuous duty overnight(s)", "adequate rest", "minimum rest", "required rest", "plt fatigue" (i.e., pilot fatigue), and "compensatory rest". The prominence of these fatigue-related phrases parallels the prominence of these concerns in the industry (NTSB, 1999; FAA, 1999b; Mattick, 1999; ALPA, 1999).

<u>freq</u>	<u>phrase</u>		
		12	24 HR REST PERIOD
152	REST PERIOD	12	CREW SCHEDULER
109	REDUCED REST	12	FELL ASLEEP
79	CREW REST	12	LACK OF SLEEP
57	CONTINUOUS DUTY	12	SCHEDULING PRACTICES
46	CREW SCHEDULING	11	ENTIRE CREW
37	DUTY PERIOD	10	FATIGUE AND STRESS
36	REST PERIODS	10	REDUCED REST OVERNIGHT
34	RESERVE OR STANDBY	9	DUTY PERIODS
30	REST REQUIREMENTS	9	EARLY AM
28	CREW FATIGUE	9	FALL ASLEEP
22	CREW DUTY	9	FIRST NIGHT
20	CONTINUOUS DUTY OVERNIGHT	8	CIRCADIAN RHYTHMS
19	ADEQUATE REST	8	NOT SLEEP
18	MINIMUM REST	8	PROPER REST
18	REQUIRED REST	8	SCHEDULING DEPT
17	PLT FATIGUE	8	SHORT REST
16	COMPENSATORY REST	8	STANDBY PLT
16	STANDBY STATUS	7	14 HR DUTY
15	REDUCED REST PERIOD	7	BODY CLOCK
15	SLEEP THE NIGHT	7	CIRCADIAN RHYTHM
13	CONTINUOUS DUTY OVERNIGHTS	7	CONTEXT OF REST PERIOD
13	EARLY MORNING	7	DEFINITION OF DUTY
13	LONG DUTY	7	DUTY AND REST
13	NIGHT 'S SLEEP	7	DUTY REGS
13	RESERVE OR STANDBY STATUS		

It can be useful to subdivide the list of topical phrases into groups. One approach, shown here, is based on the prominence of words in the phrases. To find the prominence of each word among all 420 of the fatigue-related phrases, the frequencies of the word groups containing each word were summed. The top 10 of 304 phrase words are shown in the following table. This shows, for example, that "rest" is the most prominent word among the phrases.

<u>sum</u>	<u>phrase word</u>	291	CREW	135	SCHEDULING
855	REST	163	REDUCED	109	NIGHT
370	DUTY	151	FATIGUE	102	RESERVE
304	PERIOD	147	SLEEP		

These words can be used to group the prominent fatigue-related phrases. For example, one can find all of the phrases containing the prominent word "rest". Using this approach, the following 10 tables show prominent subtopics within the fatigue-related narratives. The frequency of each phrase within 200 fatigue-related narratives is shown in the first column.

These groupings show, for example, that "rest period" and "reduced rest" are the most prominent "rest" phrases. Similarly, "continuous duty" and "duty period" are the most prominent "duty" phrases. Among "period" phrases, "rest period" is far more common than "duty period", indicating that rest periods are a greater concern than duty periods among the sampled narratives.

<u>freq</u>	<u>REST phrases</u>	5	14 HR CREW DUTY	<u>freq</u>	<u>REDUCED phrases</u>
152	REST PERIOD	5	CALL FROM CREW SCHEDULING	109	REDUCED REST
109	REDUCED REST	5	CALLED CREW SCHEDULING	15	REDUCED REST PERIOD
79	CREW REST			10	REDUCED REST OVERNIGHT
36	REST PERIODS			7	SCHEDULED REDUCED REST
30	REST REQUIREMENTS			4	REDUCED REST PERIODS
19	ADEQUATE REST			3	REDUCED REST SCHEDULES
18	MINIMUM REST			3	REDUCED REST TRIPS
18	REQUIRED REST			2	BLOCK - TO - BLOCK REDUCED REST
16	COMPENSATORY REST			2	BLOCK REDUCED REST
15	REDUCED REST PERIOD			2	GIVEN A REDUCED REST PERIOD
<u>freq</u>	<u>DUTY phrases</u>	<u>freq</u>	<u>FATIGUE phrases</u>		
57	CONTINUOUS DUTY	28	CREW FATIGUE		
37	DUTY PERIOD	17	PLT FATIGUE		
22	CREW DUTY	10	FATIGUE AND STRESS		
20	CONTINUOUS DUTY OVERNIGHT	7	FATIGUE AND STRESS INDUCED FATIGUE		
13	CONTINUOUS DUTY OVERNIGHTS	5	EXTREMELY FATIGUED		
13	LONG DUTY	5	FATIGUE CAUSED		
9	DUTY PERIODS	4	CAUSED BY PLT FATIGUE		
7	14 HR DUTY	4	CHRONIC FATIGUE		
7	DEFINITION OF DUTY	4	LEVEL OF FATIGUE		
7	DUTY AND REST	4	SIGNS OF FATIGUE		
<u>freq</u>	<u>PERIOD phrases</u>	<u>freq</u>	<u>SLEEP phrases</u>		
152	REST PERIOD	15	SLEEP THE NIGHT		
37	DUTY PERIOD	13	NIGHT 'S SLEEP		
36	REST PERIODS	12	FELL ASLEEP		
15	REDUCED REST PERIOD	12	LACK OF SLEEP		
12	24 HR REST PERIOD	9	FALL ASLEEP		
9	DUTY PERIODS	8	NOT SLEEP		
7	CONTEXT OF REST PERIOD	7	SLEEP PATTERNS		
7	REQUIRED REST PERIOD	6	FALLING ASLEEP		
7	REST PERIOD EXISTS	6	SLEEP PRIOR		
7	SAID FOR REST PERIODS	5	ENOUGH SLEEP		
<u>freq</u>	<u>CREW phrases</u>	<u>freq</u>	<u>SCHEDULING phrases</u>		
79	CREW REST	46	CREW SCHEDULING		
46	CREW SCHEDULING	12	SCHEDULING PRACTICES		
28	CREW FATIGUE	8	SCHEDULING DEPT		
22	CREW DUTY				
12	CREW SCHEDULER				
11	ENTIRE CREW				
7	MINIMUM CREW REST				

5	CALL FROM CREW SCHEDULING	6	REST OVERNIGHT
5	CALLED CREW SCHEDULING	4	REST THE NIGHT
5	TYPE OF SCHEDULING	3	LATE AT NIGHT
3	CALL SCHEDULING		
3	CALLED SCHEDULING	<u>freq</u>	<u>RESERVE phrases</u>
3	SCHEDULING ASKED	34	RESERVE OR STANDBY
3	SCHEDULING CALLED	13	RESERVE OR STANDBY STATUS
		7	RESERVE ' OR ' STANDBY ' PLT
<u>freq</u>	<u>NIGHT phrases</u>	7	RESERVE OR STANDBY DUTY
20	CONTINUOUS DUTY OVERNIGHT	7	RESERVE OR STANDBY PLT
15	SLEEP THE NIGHT	6	RESERVE OR STANDBY FALLS
13	CONTINUOUS DUTY OVERNIGHTS	5	CONSISTENT INTERP OF RESERVE
13	NIGHT 'S SLEEP	4	RESERVE CREW
10	REDUCED REST OVERNIGHT	4	RESERVE PLT
9	FIRST NIGHT	3	AM A RESERVE CAPT
7	LATE NIGHT		

Two very useful by-products of the method used to produce the topically relevant phrases are a display of the most relevant narratives with their matching phrases highlighted, and a relevance-ranked list of the narratives that are relevant to the topic. Here is the most relevant narrative, in its entirety. Although it does not contain any form of the word "fatigue", it does contain a diversity of fatigue-related topics.

I WORK FOR A LARGE REGIONAL / NATIONAL CARRIER AND CURRENTLY **AM A RESERVE CAPT** . OUR CURRENT WORKING AGREEMENT HAS VERY LITTLE IN THE WAY OF WORK RULES REGARDING SCHEDULING AND HRS OF SVC , AND THUS , WE ARE SCHEDULED AND FLOWN TO THE MAX ALLOWED BY THE FARs , WHICH WE ALL KNOW LEAVES MUCH TO BE DESIRED WITH THE REALITY OF OUR **CIRCADIAN RHYTHMS** . MANY PEOPLE THINK THAT **CIRCADIAN RHYTHMS** ONLY APPLY TO LONG HAUL INTL PLTS . HOWEVER , AFTER A NUMBER OF YRS AS BOTH A MIL AND COMMERCIAL CARRIER PLT I'VE FOUND THAT EVERYONE 'S BODY NEEDS A ROUTINE , AND RADICAL CHANGES CAN ADVERSELY AFFECT ONE 'S PERF AND ABILITY TO GET **ADEQUATE SLEEP** DURING THE SUPPOSED **REST PERIOD** . OUR AIRLINE 'S **SCHEDULING DEPT** OPERATES UNDER CRISIS MGMNT DUE TO OUR MGMNT 'S ' STAFFING STRATEGY , ' AND THUS REQUIRES MANY **RESERVE CREW MEMBERS** TO COVER MORE THAN 1 **SCHEDULED TRIP** IN A CALENDAR DAY AND THUS WE HAVE A LARGE NUMBER OF ' **SCHEDULED REDUCED REST PERIODS** ' WHICH ARE 8 HRS , WHICH DOES NOT INCLUDE TRANSPORTATION LCL IN NATURE , WHICH , IN REALITY , REDUCES YOUR TIME AT A **REST FACILITY** WELL BELOW 8 HRS , PROVIDED YOU FALL TO SLEEP AS SOON AS YOU ARRIVE AT THE HOTEL . MY TRIP / RERTE FROM HELL STARTED AS A 3 DAY WITH AN 8 HR **REST THE FIRST NIGHT** WITH AN **EARLY RPT** . I HAPPENED TO BE COMING OFF A COUPLE OF **NIGHT TRIPS** AND THE **EARLY MORNING RPT** HAD ME A LITTLE OUT OF SYNC . WHEN WE ARRIVED AT OUR NEXT **OVERNIGHT STATION** , WHICH WE WERE **SCHEDULED COMPENSATORY REST** , I **FELL ASLEEP** EARLY NOT BEING ACCUSTOMED TO **EARLY MORNING RPTS** AND THUS WOKE VERY EARLY ON THE THE THIRD DAY . OUR DAY WAS SCHEDULED TO START AT 0450 AND END AT 1358 LCL . WHEN I WENT TO CHKOUT , **CREW SCHEDULER INFORMED** ME I HAD BEEN REROUTED AND I NOW HAD ADDITIONAL FLTS WITH ANOTHER OVERNIGHT AND MY DUTY DAY NOW WAS GOING TO BE 15:30 , LEGAL BUT SAFE ? LATER , AS I WAITED TO MAKE THE LAST FLT TO THE **OVERNIGHT STATION** THEY HAD ME DO AN ADDITIONAL 2 LEGS , WHICH BROUGHT ME UP TO 8 LEGS . AFTER CHKING THE TRIP ON THE SCHEDULING COMPUTER , I FOUND THE SCHEDULER HAD CHANGED THE TRIP TO SHOW A COMBINATION OF ACTUAL TIME FLOWN , AND MARKETING TIMES TO MAKE THE TRIP LEGAL (I.E . , UNDER 8 HRS SCHEDULED) AS OPPOSED TO USING THE HISTORIC BLOCK TIMES AS IS CALLED FOR BY BOTH OUR OPS MANUAL AND FAA POI . THE REMAINDER OF THE TRIP WAS MUCH THE SAME . THE FAA NEEDS TO RECOGNIZE THE IMPORTANCE OF **QUALITY CREW REST** AND IMPLEMENT GUIDELINES TO PREVENT SUCH **SCHEDULING PRACTICES** . ON THE THIRD AND FOURTH DAY , I WAS FAR FROM BEING AT PEAK PERF AND HAD THERE BEEN A SERIOUS EMER THE OUTCOME MAY HAVE BEEN QUESTIONABLE . THE FAA IS MANDATING MANY ITEMS TO ENHANCE SAFETY SUCH AS TCASII AND GPWS , HOWEVER , THEY SEEM TO FORGET THE MOST CRITICAL AND COMPLEX PIECE OF EQUIP ON THE ACFT : THE PLT ! (254345)

Numerous fatigue-related phrases are highlighted in this narrative, and most of these appear in the list of 420 fatigue-related phrases produced by QUORUM phrase discovery. Some phrases that are not on the list are also highlighted. The phrase "scheduled compensatory rest", for example, is highlighted because the phrases "scheduled rest" and "compensatory rest" are on the list. This approach aids the user in recognizing compound topical phrases in the narratives.

Here are the accession numbers of the 100 narratives that are most relevant to the fatigue-related phrases. The more relevant narratives appear closer to the top of the list.

1. 254345	18. 281704	35. 294130	52. 80148	69. 275586	86. 375952
2. 288683	19. 257793	36. 309408	53. 307314	70. 102754	87. 134612
3. 288893	20. 219810	37. 82286	54. 118537	71. 218676	88. 280233
4. 288846	21. 360800	38. 145545	55. 302099	72. 123335	89. 373770
5. 317360	22. 96245	39. 311602	56. 245026	73. 168334	90. 185044
6. 344664	23. 273938	40. 296275	57. 294430	74. 301360	91. 261246
7. 295352	24. 245003	41. 205528	58. 281395	75. 112090	92. 123033
8. 289770	25. 324660	42. 319125	59. 142582	76. 190632	93. 360420
9. 290921	26. 340923	43. 262904	60. 270256	77. 96789	94. 345560
10. 299489	27. 256799	44. 367822	61. 364640	78. 358723	95. 189506
11. 362160	28. 261075	45. 314510	62. 146711	79. 147013	96. 108189
12. 188837	29. 123541	46. 164061	63. 140005	80. 298219	97. 356959
13. 96242	30. 206207	47. 184813	64. 337600	81. 302300	98. 306800
14. 277949	31. 193131	48. 348901	65. 258759	82. 223012	99. 270930
15. 233057	32. 276356	49. 176651	66. 246248	83. 172229	100. 151142
16. 255852	33. 367856	50. 143879	67. 206734	84. 368250	
17. 297614	34. 254267	51. 244901	68. 254490	85. 206269	

This example shows that QUORUM phrase discovery is useful for finding topically-related phrases and narratives that do not necessarily contain the original query words or phrases.

Related work

Previous work involving QUORUM explored the nature of QUORUM models (McGreevy, 1995; McGreevy, 1996), showed how QUORUM models can represent and relevance-rank documents and topics (McGreevy, 1997), and evaluated the precision of QUORUM's relevance judgments (McGreevy & Statler, 1998). Discussion of related work by others in these areas is found in those papers.

The present paper shows how QUORUM can be used to accomplish keyword search, phrase search, and phrase extraction. Related work in these areas is discussed below, and the QUORUM method is shown to have certain advantages.

QUORUM keyword search and retrieval

Most keyword search methods use term indexing (Salton, 1981), where a word list represents each document and internal query. As a consequence, given a keyword as a user query, these methods use the presence of the keyword in documents as the main criterion of relevance. In contrast, QUORUM keyword search uses indexing by term association, where a list of contextually associated word pairs represents each document and internal query. Given a keyword as a user query, QUORUM uses not only the presence of the keyword in documents but also the contexts of the keyword as the criteria of relevance. This allows retrieved documents to be sorted on their relevance to the keyword in context. In fact, the QUORUM method is related to the Keyword In Context (KWIC) method developed by Luhn (1960), and described by Fischer (1964). QUORUM automates analysis of the contexts of keywords, such as those produced by the KWIC method and other concordance methods, to accomplish indexing by term association.

Some methods utilize term associations to identify or display additional query keywords that are associated with the user-supplied keywords (Jing & Croft, 1994; Gauch & Wang, 1996; Xu & Croft, 1996; McDonald, Ogden, & Foltz, 1997). These methods do not use term association to represent documents and queries, however, and instead rely on term indexing. As a consequence, "query drift" occurs when the additional query keywords retrieve documents that are poorly related or unrelated to the original keywords. Further, term index methods are ineffective in ranking documents on the basis of keywords in context.

Unlike QUORUM keyword search and retrieval, the search method of Hawking and Thistlewaite (1996) creates no model of the query and no models of the documents of the database. Instead, the query words are compared with the words that appear in the documents of the database, and documents containing greater numbers of query words in shorter sequences of words are considered to have greater relevance. This clumping of multiple query words is substantially different from the QUORUM method. Further, as with conventional term indexing schemes, the method of Hawking and Thistlewaite allows a single query term to retrieve documents containing the term, but unlike the QUORUM method, their method cannot rank the retrieved documents according to their relevance to the contexts of the single query term.

QUORUM phrase search and retrieval

Most phrase search and retrieval methods treat query phrases as single terms, and typically rely on pre-existing lists of key phrases (Fagan, 1987; Croft, Turtle, and Lewis, 1991; Gey & Chen, 1997; Jing & Croft, 1994; Gutwin, Paynter, Witten, Nevill-Manning, & Frank, 1998; Jones & Staveley, 1999). This approach allows little flexibility in matching query phrases with similar phrases in the text, and it requires that all possible phrases be identified in advance, typically using statistical or "natural language processing" (NLP) methods. In contrast, the QUORUM method represents phrases implicitly among contextual associations representing each document. This allows both exact matching of phrases and the option of flexible matching of phrases. In addition, the QUORUM method eliminates the need for explicit and inevitably incomplete lists of phrases.

Since QUORUM phrase search does not depend on phrase frequency, it is not hampered by the infrequency of most phrases (Turpin & Moffat, 1999) which reduces the effectiveness of statistical phrase search methods. Since QUORUM phrase search does not use NLP methods, it is not subject to problems such as mistagging (Fagan, 1987).

Croft, Turtle, and Lewis (1991) dismiss the notion of implicitly representing phrases as term associations, but the association metric they tested is not as definitive as that used by QUORUM. Unlike QUORUM, their pairwise associations do not include a measurement of degree of proximity. Further, while QUORUM restricts the scope of acceptable contexts to a few words and enforces word order, the association method of Croft et al. uses entire documents as the contextual scope, and uses no directional information.

Finally, unlike typical Internet search tools, QUORUM phrase search makes it easy to use large numbers of phrases as query phrases.

QUORUM phrase generation

QUORUM phrase generation is one of several methods that display phrases contained in collections of text as a way to assist users in domain analysis or query formulation and refinement (Godby, 1994; Gutwin, et al., 1998; Normore, Bendig, & Godby, 1999; Zamir & Etzioni, 1999; Jones & Staveley, 1999). While the other methods maintain explicit and incomplete lists of phrases, QUORUM's implicit phrase representation can provide all possible phrases. In addition, QUORUM phrase generation can provide the essence of multiple, similar phrases, which can be used as queries in QUORUM phrase search. The option of using the flexible matching of QUORUM phrase search allows the generated query phrases to match both identical and nearly identical phrases in the text. This ensures that inconsequential differences do not spoil the match.

Some phrase generation methods use contextual association to identify important word pairs, but do not identify longer phrases, or do not use the same associative method to identify phrases having more than two words (Church, Gale, Hanks, & Hindle, 1991; Gey and Chen, 1997; Godby, 1994). In contrast, QUORUM phrase generation treats phrases uniformly regardless of their size.

Some methods rely on manual identification of phrases at a critical point in the process (Gelbart & Smith, 1991; Gutwin, et al., 1998; Jones & Staveley, 1999), while QUORUM phrase generation is fully automatic.

QUORUM phrase discovery

QUORUM phrase discovery is similar to the so-called "natural language processing" (NLP) methods of phrase-finding in that it classifies words and requires that candidate word sequences match particular patterns. Most methods, however, classify words by part of speech using grammatical taggers and apply a grammar-based set of allowable patterns (Godby, 1994; Jing & Croft, 1994; Gutwin, et al., 1998; de Lima & Pedersen, 1999; Jones & Staveley, 1999). These methods typically remove all punctuation and stopwords as a preliminary step, and most then discover only simple or compound nouns leaving all other phrases unrecognizable. In contrast, QUORUM uses the full text, and applies a simple classification scheme where the main categorical distinction is between stopterms (stopwords and punctuation) and non-stopterms. In addition, QUORUM uses a simple, procedurally defined set of acceptable patterns that requires phrases to begin and end with non-stopterms, limits the interior stopterms, and allows the "-" (dash) character to be an interior term.

Like Keyphind (Gutwin, et al., 1998) and Phrasier (Jones & Staveley, 1999), QUORUM phrase discovery identifies phrases in sets of documents. In contrast to these methods, however, QUORUM phrase discovery requires no grammatical tagging, no training phrases, no manual categorization of phrases, and no pre-existing lists of identifiable phrases. Further, QUORUM phrase discovery identifies a far greater number of the phrases that occur within sets of

documents because its method of phrase identification is more powerful. The larger number of phrases identified by QUORUM phrase discovery also provides much more information for determining the degree of relevance of each document containing one or more of the phrases.

Conclusion

The new QUORUM methods of keyword search, phrase search, phrase generation, and phrase discovery significantly extend beyond the core QUORUM methods to address search tasks of importance to the ASRS. In particular, these methods have the potential to directly support and enhance ASRS Search Requests, Quick Responses, and special topical studies. Further, the current implementation of these methods, combined with a graphical user interface, can be used by the ASRS, commercial airlines, and other organizations that process narrative incident reports. As a result, these new methods can enhance commercial aviation safety by supporting incident analysis for better understanding of operational incidents.

The new QUORUM methods extend beyond the current state of the art by applying QUORUM's unique contextual associative indexing to enable: 1) keyword-in-context search, 2) flexible phrase search, 3) generation of phrases contained in collections of documents, using implicit phrase models of those documents, and 4) discovery of topically associated phrases, and the documents in which they are prominent. In addition, a new method of phrase extraction from text was introduced as part of the phrase discovery method.

In the immediate future, the core and new QUORUM methods and software described in this paper will be transferred to the ASRS for incorporation into their systems and processes. In parallel, the methods and software will be made available for further research and development, and possible commercial applications.

Continued research involving ASRS incident narratives, as well as preliminary applications and extensions to Space Shuttle maintenance incidents (McGreevy, Kanki, Stephenson, and Patankar, 1999), incidents in the electric power industry, and business intelligence applications (Keijola, 1998) will lead to refinement and solidification of QUORUM's functionality, while offering challenging new opportunities for its evolution.

References

ALPA: Flight time/duty time rules: Safety in the balance. <http://www.alpa.org/internet/ftdt/index.htm>, 1999.

ASRS: Aviation Safety Reporting System home page. <http://asrs.arc.nasa.gov>, 1999.

ASRS: An Analysis of International Airspace Incidents, Aviation Safety Reporting System, Quick Response No. 290, 1997.

Church, K., Gale, W., Hanks, P., and Hindle, D.: Using statistics in lexical analysis. In U. Zernik (ed.), *Lexical Acquisition: Using On-Line Resources To Build A Lexicon* (pp. 115-164). Lawrence Earlbaum, Hillsdale, New Jersey, 1991.

Connell, L.: Incident Reporting: The NASA Aviation Safety Reporting System. *GSE Today*, vol. 7, no. 2, <http://www.gsetoday.com/back-is/2-99/NASA-Safety.htm>, 1999.

Connell, L.: NASA Director, ASRS, personal communication, 1998.

Croft W. B., Turtle H. R., and Lewis D. D.: The use of phrases and structure queries in information retrieval. *Proc. ACM SIGIR*, 1991, pp. 32-45.

de Lima, E. F., and Pedersen, J. O.: Phrase recognition and expansion for short, precision-biased queries based on a query log. *Proc. ACM SIGIR*, 1999, pp. 145-152.

DOT: Air Traffic Control. Air Traffic Service, U.S. Dept. of Transportation, 7110.65C, 1982.

FAA: Search Aviation Safety Data. http://www.asy.faa.gov/safety_data/, 1999a.

- FAA: Flight crewmember flight time limitations and rest requirements; Final Rule, Federal Register, vol. 64, no. 114, 1999b, pp. 32175-32177.
- Fagan, J. L.: Experiments in automatic phrase indexing for document retrieval: A comparison of syntactic and non-syntactic methods. Ph.D. thesis TR87-868, Dept. of Comp. Sci., Cornell University, 1987.
- Fischer, M. F.: History and use of the KWIC index concept. Thesis, San Jose State College, 1964.
- Frank, S.: Database Manager, ASRS/Battelle, personal communication, 1998.
- Gauch, S., and Wang, J.: Corpus analysis for TREC 5 query expansion. Proc. TREC 5, NIST SP 500-238, 1996, pp. 537-547.
- Gelbart, D., and Smith, J. C.: Beyond boolean search: FLEXICON, a legal text-based intelligent system. Proc. ACM Artif. Intell. & Law , 1991, pp. 225-234.
- Gey, F. C., and Chen, A.: Phrase discovery for English and cross-language retrieval at TREC 6 . Proc. TREC 6, NIST SP 500-240, 1997, pp. 637-644.
- Godby, C. J.: Two techniques for the identification of phrases in full text. Annual Review of OCLC Research. Online Computer Library Center, Dublin, Ohio, 1994.
- Gutwin, C., Paynter, G., Witten, I. H., Nevill-Manning, C., and Frank, E.: Improving browsing in digital libraries with keyphrase indexes. TR 98-1, Computer Sci. Dept., Univ. of Saskatchewan, 1998.
- Hawking, D., and Thistlewaite, P: Proximity operators - So near and yet so far. Proc. TREC 4, NIST SP 500-236, 1996, pp. 131-143.
- Jing, Y., and Croft, W. B.: An association thesaurus for information retrieval, CIIR TR IR-47, Univ. of Mass., 1994.
- Jones, S., and Staveley, M.: Phrasier: A system for interactive document retrieval using keyphrases. Proc. ACM SIGIR, 1999, pp. 160-167.
- Keijola, M.: Researcher, Helsinki University of Technology, personal communication, 1998.
- Logan, T. J.: Safety data: Preserve it or lose it. <http://www.aviationweek.com/safety/library/safdat.htm>, adapted from Safety Data: An Airline Safety Director's Perspective, ISASI Seminar presentation, October 1998.
- Luhn, H. P.: Keyword-In-Context Index for Technical Literature (KWIC Index). American Documentation, vol. 11, no. 4, 1960, pp. 288-295.
- Mattick, T.: Reserve rest requirements. Air Line Pilot, vol. 68, no. 7, 1999, pp. 17-19, 47.
- McDonald, J., Ogden, W., and Foltz, P.: Interactive information retrieval using term relationship networks. Proc. TREC 6, NIST SP 500-240, 1997, pp. 379-383.
- McGreevy, M. W.: QUORUM installation guide and user reference. Off-line HTML-based QUORUM tutorial, 1999.
- McGreevy, M. W., Kanki, B., Stephenson, H., and Patankar, H.: Initial computer-aided analysis of the KSC Shuttle Processing Events Database. Unpublished progress report, 1999.
- McGreevy, M. W., and Statler, I. C.: Rating the relevance of QUORUM-selected ASRS incident narratives to a "Controlled Flight Into Terrain" accident. NASA TM-1998-208749. Ames Research Center, Moffett Field, Calif., 1998.
- McGreevy, M. W.: A practical guide to interpretation of large collections of incident narratives using the QUORUM method. NASA TM-112190. Ames Research Center, Moffett Field, Calif., 1997.

- McGreevy, M. W.: Reporter concerns in 300 mode-related incident reports from NASA's Aviation Safety Reporting System. NASA TM-110413. Ames Research Center, Moffett Field, Calif., 1996.
- McGreevy, M. W.: A relational metric, its application to domain analysis, and an example analysis and model of a remote sensing domain. NASA TM-110358. Ames Research Center, Moffett Field, Calif., 1995.
- Normore, L., Bendig, M., and Godby, C. J.: WordView: Understanding words in context. Proc. Intell. User Interf., 1999, pg. 194.
- NTSB: Testimony of Vernon S. Ellingstad, Director, Office of Research and Engineering, National Transportation Safety Board before the Subcommittee on Aviation, Committee on Transportation and Infrastructure, House of Representatives, Regarding Pilot Fatigue, August 3, 1999.
- Rosenthal, L.: Battelle Manager, ASRS, personal communication, 1998.
- Salton, G.: A blueprint for automatic indexing. ACM SIGIR Forum, vol. 16, no. 2, 1981. Reprinted in ACM SIGIR Forum, vol. 31, no. 1, 1997, pp. 23-36.
- Statler, I.: Program Manager, APMS, personal communication, 1998.
- Steenblik, Jan W.: 1999 Air Safety Forum: Schedule with Safety. Air Line Pilot, vol. 68, no. 9, 1999, pp. 10-14, 53.
- Turpin, A, and Moffat, A.: Statistical phrases for vector-space information retrieval. Proc. ACM SIGIR, 1999, pp. 309-310.
- Xu, J., and Croft, W.: Query expansion using local and global document analysis. Proc. ACM SIGIR, 1996, pp. 4-11.
- Zamir, E. and Etzioni, E. (1999). Grouper: A dynamic clustering interface to web search results. Proc. 8th Int'l World Wide Web Conf. (WWW8), 1999.

Appendix 1. QUORUM relations and relational metrics.

Previously published work on QUORUM utilized relations having a single contextual metric. To support the new methods presented in this paper, new metrics were developed. This extends QUORUM relations to include multiple metrics.

QUORUM relations are paired terms with one or more measurements of the degree of their contextual associations in text. Terms are usually words, but are sometimes numbers, punctuation marks, word fragments, or other strings of characters that are separated by spaces or other white-space characters. The general form of a QUORUM relation is:

term_1	term_2	metric_i	metric_j	...	metric_N
--------	--------	----------	----------	-----	----------

For example:

CREW	FATIGUE	150	103	...	500
------	---------	-----	-----	-----	-----

Each QUORUM relation can represent the contextual association of a word pair in any body of text, from a small excerpt to all of the text in a database. A QUORUM model is a collection of such relations.

The metrics that are used for QUORUM keyword and phrase methods include:

- 1) the standard relational metric value (RMV or stdRMV)
- 2) the left relational metric value (LRMV or leftRMV)
- 3) the right relational metric value (RRMV or rightRMV)

These metrics are further described in the next section. Given these metrics, the format of QUORUM relations for keyword and phrase methods is:

term_1	term_2	stdRMV	leftRMV	rightRMV
--------	--------	--------	---------	----------

For example, here are some of the QUORUM relations from the keyword query model of "engage". These relations represent contextual associations in all of the narratives of the ASRS database.

AUTOPLT	ENGAGED	17905	5442	12463
DISENGAGED	AUTOPLT	12805	4833	7972
ENGAGED	ALT	6015	2642	3373
ENGAGED	HOLD	2992	1439	1553
ALT	DISENGAGED	2937	1340	1597
NOT	ENGAGE	2839	441	2398
ENGAGE	AUTOPLT	2572	1282	1290
NOT	ENGAGED	2484	786	1698
DISENGAGE	AUTOPLT	2194	1087	1107

Calculating the metric values

Calculation of the standard RMV is presented in detail in McGreevy (1995), and is described in McGreevy (1996) and appendix 1 of McGreevy & Statler (1998). In essence, for each instance of a pair of words A and B appearing within a context window of C words in a text, their degree of association is measured. The measurements are summed across all occurrences of the word pair in the text to provide an overall measurement of their degree of contextual association.

Degree of association is measured with respect to the number (N) of terms that come between the words A and B. When they are immediately adjacent, no terms are between them and their degree of association is C minus 1. When A and B are at the limits of being in the same context, they are separated by C minus 2 terms and their degree of association is 1. In general, the degree of association between A and B, for each instance of the word pair, is C minus 1 minus N. This provides larger metric values for word pairs that have a greater degree of contextual association.

The standard RMV does not measure whether the word A precedes or follows the word B, so it cannot serve as a directional metric. Such a metric is required, however, for QUORUM phrase methods. To provide directional metrics, the left and right RMVs were developed. They are calculated in the same way the standard RMV is calculated, except

that the left RMV only considers the context to the left of (preceding) a particular word, while the right RMV only considers the context to the right of (following) a particular word. The standard RMV is equal to the sum of the left and right RMVs.

For example, in this sentence, which contains an instance of word ENGLISH followed by the word PHRASEOLOGY, the word PHRASEOLOGY is in the right context of the word ENGLISH, and the word ENGLISH is in the left context of the word PHRASEOLOGY.

BETTER ENGLISH SPEAKING FOREIGN CTRLRS AND USE OF STD PHRASEOLOGY IS NEEDED .

Using a context window of size C=10, treating the sentence in isolation from any other text, and noting that there are N=7 words between ENGLISH and PHRASEOLOGY, the metrics have these interpretations and values:

- stdRMV(ENGLISH, PHRASEOLOGY), the measure of the extent that ENGLISH and PHRASEOLOGY are in the same context, is $C-1-N = 10-1-7 = 2$; This is the same as stdRMV(PHRASEOLOGY,ENGLISH).
- rightRMV(ENGLISH, PHRASEOLOGY), the measure of the contextual association of ENGLISH followed by PHRASEOLOGY, is $C-1-N = 10-1-7 = 2$;
- leftRMV(ENGLISH, PHRASEOLOGY), the measure of the contextual association of ENGLISH preceded by PHRASEOLOGY, is 0;
- rightRMV(PHRASEOLOGY, ENGLISH). the measure of the contextual association of PHRASEOLOGY followed by ENGLISH, is 0;
- leftRMV(PHRASEOLOGY, ENGLISH), the measure of the contextual association of PHRASEOLOGY preceded by ENGLISH, is $C-1-N = 10-1-7 = 2$.

An example of combining instance relations to produce a document relation

Here is an example of combining QUORUM relations across all of the shared contexts in a text. The following are three schematic lines of text representing excerpts from text being modeled, where the items t are terms that are not A or B, and the contextual relationship between terms A and B is of interest. Assume that no other instances of A and B occur.

1.	...	t	t	t	A	B	t	t	t	...
2.	...	t	t	A	t	B	A	t	t	...
3.	...	t	t	t	B	B	A	t	t	...

Here are the QUORUM relations of each instance of the paired words A and B, assuming a context window of C=3 words. The format of each relation is: term_1, term_2, stdRMV, leftRMV, rightRMV, as discussed earlier. The line numbering indicates the line number containing the relation. For example, "2.1" is the first relation from line 2 above, and "2.2" is the second relation from that line. Each relation can take one of two forms, as shown, which are equivalent.

1.0.	A	B	2	0	2	same as	B	A	2	2	0
2.1.	A	B	1	0	1	same as	B	A	1	1	0
2.2.	A	B	2	2	0	same as	B	A	2	0	2
3.1.	A	B	1	1	0	same as	B	A	1	0	1
3.2.	A	B	2	2	0	same as	B	A	2	0	2

If lines 1-3 were the only ones containing A and B, the relations above would be summed to produce a single relation representing the overall contextual association of A and B. That relation can take one of two forms, which are equivalent:

A	B	8	5	3	same as	B	A	8	3	5
---	---	---	---	---	---------	---	---	---	---	---

The order of word pairs in the relations of QUORUM models is always shown in the same order as the typical reading order in the text. When the rightRMV is greater than the leftRMV of a relation, that relation is in the typical reading order. Accordingly, this form of the relationship between A and B is the one that would be used in the model:

B	A	8	3	5
---	---	---	---	---

This relation could be interpreted as saying that when A and B are contextually associated, A tends to follow B, with the degree of contextual association indicated by the metrics. This relationship can be observed in text lines 1-3, shown earlier. A QUORUM model consists of a collection of such relations for all word pairs of interest.

Here is a realistic example of a relation representing the overall contextual association of two words in a collection of text. It represents the degree of contextual association of the words DISENGAGED and AUTOPLT in the narratives of the ASRS database.

term_1	term_2	stdRMV	leftRMV	rightRMV
DISENGAGED	AUTOPLT	12805	4833	7972

This relation indicates that when reading the ASRS narratives and the words DISENGAGED and AUTOPLT (i.e., autopilot) are found in the same context (in this case, with no more than 24 words between the two words), the word AUTOPLT is typically encountered after the word DISENGAGED, and less typically can occur before the word DISENGAGED. The words are contextually associated to the degree indicated by the metrics (relative to those in other relations). A typical excerpt from the ASRS narratives is:

I DISENGAGED THE AUTOPLT AND BEGAN THE DSCNT .

An alternative directionality metric

The left and right RMVs can be used to compute a single directionality metric value, DMV, such as:

$$DMV = \text{rightRMV} - \text{leftRMV}$$

where $\text{rightRMV} \geq 0$, $\text{leftRMV} \geq 0$, and $\text{rightRMV} - \text{leftRMV}$, and $DMV > 0$. Applying this metric to the "engage" model relations shown earlier, and sorting on DMV, this table is the result:

term_1	term_2	stdRMV	leftRMV	rightRMV	DMV
AUTOPLT	ENGAGED	17905	5442	12463	7021
DISENGAGED	AUTOPLT	12805	4833	7972	3139
NOT	ENGAGE	2839	441	2398	1957
NOT	ENGAGED	2484	786	1698	912
ENGAGED	ALT	6015	2642	3373	731
ALT	DISENGAGED	2937	1340	1597	257
ENGAGED	HOLD	2992	1439	1553	114
DISENGAGE	AUTOPLT	2194	1087	1107	20
ENGAGE	AUTOPLT	2572	1282	1290	8

This table indicates, for example, that AUTOPLT and ENGAGED are often found in close proximity, and are usually found in that order in the text. In contrast, ENGAGE and AUTOPLT are less often found in close proximity, and are about equally found in that order and in the order AUTOPLT ENGAGE.

Use of the DMV allows relations to be weighted on a combination of: 1) overall degree of contextual association of the word pair in the text and, 2) the tendency of the word pair to appear in one order more than the opposite order. In contrast, using the rightRMV and leftRMV weights relations on their degree of contextual association in each direction. Accordingly, the rightRMV and leftRMV are used when doing phrase operations.

Creating additional metrics by removing some effects of frequency

For some applications, it can be useful to create additional metrics by removing some effects of frequency. If M is one of the metrics stdRMV, leftRMV, rightRMV, or DMV, and if F0 is the frequency of the less frequently occurring word in the pairwise relation, F1 is the frequency of the more frequently occurring word, Fp is the frequency of the probe term (see McGreevy, 1995, pg. 6), and Fmax is the maximum frequency of all words in the model, then four useful metrics are: M times Fmax divided by F0; M times Fmax divided by F1; M times the square of Fmax divided by the product of F0 and F1; and M times Fmax divided by Fp. Preliminary tests have shown that the resulting metrics are not useful for relevance-ranking but can be informative when developing context-based object-oriented models.

Appendix 2. How QUORUM keyword search works.

This appendix describes how the new QUORUM keyword search method works. All of the steps are performed automatically once the user invokes the program and provides one or more keywords. This method is built upon the core methods of text analysis, modeling, and relevance-ranking, which are described in great detail in appendix 1 of McGreevy & Statler (1998), as well as in McGreevy (1997), McGreevy (1996), and McGreevy (1995).

1. Read user-selected options (described on page 5), if any, and set parameters accordingly. If the user selected the option to provide the query model rather than to provide keywords, input the model and skip to step 7
2. Read the user's query keyword or keywords. (Word fragments are acceptable when using "contained match". See 5b.)
3. Capitalize the keyword(s), since text is capitalized in the narratives of the ASRS database.
4. Map the keyword(s) to ASRS abbreviations and usage. The user has the option of overriding this.
5. Gather relations for the model to be used in ranking the database.
- 5a. If the user invoked the exact match option, scan a collection of QUORUM models, which typically represent the narratives of the database to be searched, and extract the QUORUM relations that exactly match at least one of the query keywords. For an exact match, a query keyword and one of the two words in the QUORUM relation must be identical.
- 5b. If the user invoked the "contained match" option, scan a collection of QUORUM models, which typically represent the narratives of the database to be searched, and extract the QUORUM relations that contain at least one of the query keywords. For a contained match, the contiguous sequence of characters in a query keyword must also be found in one of the two words in the QUORUM relation.

Note for 5a and 5b: Currently, relations are gleaned from the entire ASRS database, even if only a subset of narratives is to be ranked. In future, the user should have the option of gleaning relations from the subset to be ranked, the entire database to be searched, or any user-selected collection of QUORUM models. Gleaning relations from the whole database captures the more broadly typical contexts of the keyword(s). In contrast, gleaned relations from a subset of the database (or some other collection of QUORUM models) could focus the subsequent search on some particular context or contexts of the keyword(s).

6. Reduce the relations gleaned in steps 5a or 5b so that there is only one relation for every pair of words, regardless of word order in the relation. For each unique relation, compute new metrics by finding the sum of the corresponding metrics in the redundant relations. Adjust the word order within each relation, and the left and right contextual metrics, so that the word order is the same as that which is typically found in the database. (See appendix 1, "An example of combining instance relations to produce a document relation", which involves a similar combination of relations.) Sort the resulting relations on their standard contextual metrics, so that the more prominent relations are closer to the top of the list. This list of relations is a QUORUM query model containing the ranking criteria. Create a file containing this model.

If the user invoked the option of only generating the query model, exit the process here.

7. Rank the QUORUM models that represent the narratives of the ASRS database using the relations in the query model produced in steps 5 and 6 (or provided by the user in step 1) as relevance criteria. That is, compare the query model with each narrative model and rank the narrative models on their degree of similarity to the query model. If the subset option was selected, rank only that subset of the ASRS database.

QUORUM keyword search uses large-context models of the narratives in the ASRS database. These models are similar to the ones used to represent the ASRS database in the Cali project (McGreevy & Statler, 1998). The models consider the context of a term to include 25 terms on either side of the term in question. In a narrative model, the paired terms in each QUORUM relation are put into their most typical order in the narrative. This makes it easier to visualize their relationship in the text. When ranking for keyword search, term order is not utilized.

In the comparison of the query model and a narrative model, each similarity value (one for each type of metric: standard RMV, left RMV, right RMV, see appendix 1) is found by taking the corresponding inner product of the two models and scaling the result. The inner product favors narratives that share prominent features with the query model. (See McGreevy (1997) for details of relevance-ranking calculations involving a single metric value.)

The scale factors are intended to favor narratives whose overall emphases are closer to the overall emphases of the query model, and to favor narratives whose lengths are closer to being at least average.

The narrative emphasis factor is equal to the sum of the standard relational metric values (stdRMVs) of the shared relations, as measured in the narrative, divided by the sum of all relations in the narrative model. This measures the fraction of the narrative model that is shared with the query model.

The query emphasis factor is equal to the sum of the standard relational metric values (stdRMVs) of the shared relations, as measured in the query model, divided by the sum of all the relations in the query model. This measures the fraction of the query model that is shared with the narrative model.

The length factor is computed by dividing either the number of terms in the narrative or the average number of terms across all narratives, whichever is smaller, by a number larger than the largest number of terms in any ASRS narrative (e.g., 2000). This factor somewhat disfavors overly terse reports, which tend to lack elaboration and detail, but adds no further reward for length that is beyond average. Using only relevance density would tend to favor only terse reports.

The similarity value for each metric is the product of the corresponding inner product and these three scale factors. A combination of the similarity values serves as a single relevance-ranking value.

8. Create a file containing a list of the identifiers of all relevant narratives and their relevance-ranking values (one for each type of metric), sorted in decreasing order of relevance to the query model. For keyword search, that sorting is applied to the relevance-ranking values associated with the standard RMV.
9. If the user selected the "and" (logical intersection) option, steps 2-8 are applied for each of the two sets of keywords. The two relevance-ranked lists, resulting from the two applications of step 8, are then combined. To do this, the corresponding relevance-ranking values for each narrative are combined, essentially by finding their product, as described below.

Each ranked list contains one line for every narrative in the database. For example, here is the line for ASRS report number 168127 from the file containing the ranking on the keywords "frequency" or "frequencies" (produced in the first application of steps 2-8):

accession number	RRV (std)	RRV (left)	RRV (right)
163127	5576	7224	2846

The header defining the columns is not part of the file. RRV is relevance-ranking value, whose derivation is described in McGreevy (1997), and is here applied to the three QUORUM relational metrics. The metrics are described in appendix 1 of the present paper.

Here is the line for the same report (163127) from the file containing the ranking of narratives on the keywords "congestion" or "congested" (produced in the second application of steps 2-8):

163127	2508	1771	1691
--------	------	------	------

To combine the two rankings of each narrative, the corresponding RRVs are combined by finding the product of their square roots, scaled by a constant, C, and divided by the square root of the maximum RRV ranking value for each metric. For example, to combine the standard relevance ranking values (stdRRVs), these values are used for all narratives:

The maximum stdRRV in the "frequency" ranking (the first file produced in step 8) is 143275.
The maximum stdRRV in the "congestion" ranking (the second file produced in step 8) is 6049.

The maximum of these two stdRRVs is 143275.
A useful value of C is 10000.

To combine the standard RRV rankings for narrative 163127, these values are used:

The stdRRV of report 163127 in the "frequency" ranking, as shown above, is 5576.
The stdRRV of report 163127 in the "congestion" ranking, as shown above, is 2508.

So, the combined standard relevance ranking value in this case is:

$$\text{sqrt}(5576) * \text{sqrt}(2508) * 10000 / \text{sqrt}(143275) = 98796$$

After combining each of the two other metrics in the same way, the following data represent the relevance of narrative 163127 to both frequency and congestion.

accession number	RRV (std)	RRV (left)	RRV (right)
163127	98796	83478	80140

The same process is applied to all other narratives to find their relevance to both topics. The list is then sorted on the standard RRV, in descending order, so that the most relevant narratives appear toward the top of the list. The other computed metrics (left RRV and right RRV) are provided for comparison, but are not essential at this point.

10. Output the list of all relevant narratives and their relevance metrics, sorted in decreasing order of relevance. If the user invoked the option of skipping the generation and display of the highlighted narratives, skip steps 11 and 12. Otherwise, continue to the next steps.
11. Generate a HyperText Markup Language (HTML) file containing the top N narratives (where N is selectable, but defaults to 20), with relevant sections highlighted. Show the rank order number of each narrative, and its ASRS accession number. Map the narrative text to standard ASRS abbreviations and usage. To do the highlighting, first select the top R query relations (where R is selectable, but defaults to 1000) as the highlighting relations. (In the case of the "and" option, take half of the relations from each of the two query models.) Then scan each narrative and determine which word pairs not only match the highlighting relations, but also are within the same context. The context size used in this determination, typically 25 terms on either side of a particular term (where terms are words, stand-alone punctuation, etc.), is the same one used to model the narratives of the ASRS database.

After each narrative in the file, list the relations that are shared by the narrative and the query model.

12. Run Netscape to display the highlighted narratives and shared relations.

Note: At each step of the process, information to document the process is produced and can be saved for later reference.

Appendix 3. How QUORUM phrase search works.

This appendix describes how the new QUORUM phrase search method works. All of the steps are performed automatically once the user invokes the program and provides one or more phrases. This method is built upon the core methods of text analysis, modeling, and relevance-ranking, which are described in great detail in appendix 1 of McGreevy & Statler (1998), as well as in McGreevy (1997), McGreevy(1996), and McGreevy(1995).

1. Read user-selected options (described on page 6), if any, and set parameters accordingly.
2. Read the user's query phrase or phrases.
3. If a filename string is specified by the user, use it as part of the output filenames. If no filename string is specified by the user, use "temp" as part of the output filenames.
4. Capitalize the phrase(s), since text is capitalized in the narratives of the ASRS database.
5. Map the words in the phrase(s) to ASRS abbreviations and usage. The user has the option of overriding this.
6. If the user enters a single query phrase, generate a QUORUM model of that phrase, using a context of three words on either side of every word, that is, a left context window of 4 terms and a right context window of four terms. Do not apply a stoplist, and allow numbers to be terms in relations. In contrast, the keyword models use a context of 25 words on either side of each non-rare word that is not a stopword.

Here is an example of words C_i in the context of word W , and words C_4 - C_6 in the context of word X :

C1 C2 C3 W C4 C5 C6 X

Note that the words C_1 and C_6 are the most distant words from word W that are explicitly considered to be in the same context as W . So, the relation $R(C_1, W)$ and the relation $R(W, C_6)$ are the most distant relations that are explicitly represented in the context of W . One could characterize this by saying that no more than two terms are allowed between related terms. Accordingly, as a further example, words C_2 and X are the most distant words from word C_4 that are explicitly considered to be in the same context as C_4 . Despite the small context applied here, indirect chains of associations allow phrases of any length to be found. For example, W is indirectly related to X because W is directly related to C_4 , C_5 and C_6 , while C_4 , C_5 , and C_6 are directly related to X .

See the example of finding a whole sentence in the section of this paper, "Searching for a particular sentence that occurs only once in the database," on pages 19 and 20.

7. If the user provides one or more phrases in a file, generate a QUORUM model of each phrase, as in step 6.

Given a context window of 4 terms, each two-word phrase produces a single QUORUM relation. Each N -word phrase, for N greater than 2, produces $(N - 2) * 3$ relations.

If each query phrase is accompanied by a numerical weight (usually frequency), multiply the relational metrics of each phrase's relations by the weight of the phrase.
8. Create a single query model.
 - 8a. If the user invoked the option of summing the redundant relations, then for each set of relations representing the same word pair in the same order in the text, find the sum of the corresponding relational metric values.
 - 8b. If the user did not invoke the option of summing the redundant relations, then for each set of relations representing the same word pair in the same order in the text, eliminate all but the one with the largest contextual association in the given order.
- 9a. If the user invoked the option of disfavoring stopwords, then decrease the weights of relations in which both terms are stopwords, increase the weights of relations in which neither term is a stopword, and leave unchanged relations containing only one stopword.

- 9b. If the user invoked the option of favoring "emphasis words", then increase the weights of relations in which both terms are emphasis words, decrease the weights of relations in which neither term is an emphasis word, and leave unchanged relations containing only one emphasis word.
10. If the user invoked the option of ignoring one or more relations, delete those relations.
11. If the user invoked the option of ignoring one or more words, delete all relations that exactly match the word(s).
12. Output the phrase query model. If the user invoked the option to produce the query model but skip the search, exit the process here.
13. Determine which version of the ASRS database narrative models will be ranked. If the user specified one, use it. Otherwise, decide between the one with stopwords and the one without stopwords. If any one of the query phrases contains a stopword, and if the user did not choose the option of ignoring stopwords, use the database models that include relations containing stopwords. Otherwise, use the database models that only include relations not containing stopwords. The database models containing no stopwords allow the process to run faster, so they are used when possible, for the sake of efficiency.
14. Rank the QUORUM models that represent the narratives of the ASRS database (selected in step 13) on the phrase query model (completed after step 11). That is, compare the query model with each narrative model, and rank the narrative models on their degree of similarity to the query model. If the subset option was selected, rank only that subset of the ASRS database.

QUORUM phrase search uses small-context, directional models of the narratives in the ASRS database, which is not the same kind of model used for keyword search, or to represent the ASRS database in the Cali project (McGreevy & Statler, 1998). The narrative models used by phrase search consider the context of a term to include 3 terms on either side of the term in question (see step 6), use the leftRMV and rightRMV as the relational metrics (see appendix 1), and include all contextual relations that have a leftRMV or rightRMV value of at least 1. The latter constraint ensures that even a single occurrence of word X followed by word Y in a narrative, with no more than two other terms (words or stand-alone punctuation) between X and Y, will be represented in the database. One database of phrase models of the narratives does not allow X or Y to be stopwords, and one does. The one used for a phrase search depends upon whether there are any stopwords in the query (see step 13).

In the phrase-oriented narrative models, the paired terms in each QUORUM relation are put into their most typical order. This makes it easier to visualize their relationship in the text. Regardless of the typical order, when matching narrative relations to query relations, directionality metrics must be aligned, as illustrated in the following example:

If a query relation R_q is A B x_0 x_1 x_2 , that is,

	word 1	word2	std RMV	left RMV	right RMV
$R_q(A,B)=$	A	B	x_0	x_1	x_2

and a matching narrative relation is

$R_n(A,B)=$	A	B	y_0	y_1	y_2
-------------	---	---	-------	-------	-------

then the order of A and B is the same, so the directional metrics are already aligned.

Given the same query relation, but a narrative relation of

$R_x(B,A)=$	B	A	z_0	z_1	z_2
-------------	---	---	-------	-------	-------

the relations must be aligned by, in effect, converting the narrative relation to

$R_x'(A,B)=$	A	B	z_0	z_2	z_1
--------------	---	---	-------	-------	-------

which is accomplished by swapping A and B, and swapping the values of the leftRMV and the rightRMV.

The query relation $R_q(A,B)$ matches $R_x(B,A)$ and it also matches $R_x'(A,B)$, but the alignment of directional metrics is done so that the correct sense of direction is used in calculating the strength of the match.

- 14a. If there is only one query phrase, and the user does not override the default behavior, all query relations must be found in a narrative if it is to be considered as matching the query. Given a match, each similarity value (one for each type of metric: standard RMV, left RMV, right RMV, see appendix 1) is found by computing the corresponding inner product of the two models and scaling it. The inner product favors narratives that share prominent features with the query model. (See McGreevy (1997) for details of relevance-ranking calculations involving a single metric value.)

The scale factors are intended to favor narratives whose overall emphases are closer to the overall emphases of the query model, and to favor narratives whose lengths are closer to being at least average.

The narrative emphasis factor is equal to the sum of the relational metrics of the shared relations, as measured in the narrative, divided by the sum of all relations in the narrative model. This measures the fraction of the narrative model that is shared with the query model.

The query emphasis factor is equal to the sum of the relational metrics of the shared relations, as measured in the query model, divided by the sum of all the relations in the query model. This measures the fraction of the query model that is shared with the narrative model.

The length factor is computed by dividing either the number of terms in the narrative or the average number of terms across all narratives, whichever is smaller, by a number larger than the largest number of terms in any ASRS narrative (e.g., 2000). This factor somewhat disfavors overly terse reports, which tend to lack elaboration and detail, but adds no further reward for length that is beyond average. Using only relevance density would tend to favor only terse reports.

The similarity value for each metric is the product of the corresponding inner product and these three scale factors. A combination of the similarity values serves as a single relevance-ranking value.

With single query phrases, the user has the option of ranking instead by summing the relational metrics of the query relations that are found in the narrative, as with multiple phrase ranking, discussed below.

- 14b. If there are multiple query phrases, and the user does not override the default behavior, the value representing the degree of similarity between the query and a narrative is found by summing the relational metrics of the query relations that are found in the narrative. This is intended to favor narratives that contain a greater breadth of the more prominent relations of the query model.

With multiple query phrases, the user has the option of instead requiring that all of the relations match, and calculating the ranking as with the single-phrase case, described earlier.

15. Output the list of all relevant narratives and their relevance metrics, sorted in decreasing order of relevance. If the user invoked the option of skipping the generation and display of the highlighted narratives, skip steps 16 and 17. Otherwise, continue to the next steps.
16. Generate an HTML file containing the top N narratives (where N is selectable, but defaults to 20), with relevant sections highlighted. Show the rank order number of each narrative, and its ASRS accession number. Map the narrative text to standard ASRS abbreviations and usage. To do the highlighting, first select the top R query relations (where R is selectable, but defaults to 1000) as the highlighting relations. Then scan each narrative and determine which word pairs not only match the highlighting relations, but also are within the same context. The context window used in this determination is 4 terms (i.e., words, stand-alone punctuation, etc.).

After each narrative in the file, list the relations that are shared by the narrative and the query model.

17. Run Netscape to display the highlighted narratives and shared relations.

Note: At each step of the process, information to document the process is produced and can be saved for later reference.

Appendix 4. How QUORUM phrase generation works.

This appendix describes how the new QUORUM phrase generation method works. The method constructs phrases that contain a user-provided word or phrase. It also has a batch mode allowing the user to supply a list of words and/or phrases, each of which will be used to generate a list of phrases. The generated phrases are based on QUORUM's implicit representations of the phrases in the narratives of the ASRS database. All of the steps are performed automatically once the user invokes the program and provides input.

The essence of the process is to repeatedly scan a phrase-oriented model of the narratives of the ASRS database and to build up phrases using a selection of the contextual relations contained in the model. On each pass, a copy of each phrase obtained so far is subject to gaining a new word at the beginning or at the end of the phrase. The rules for adding words are based on the directional contextual associations among the words in each candidate phrase. When the desired number of phrases is generated, or when a prominence metric reaches a certain threshold, the resulting phrases are sorted on their prominence in the collection and displayed to the user.

1. Read user-selected options (described on page 8), if any, and set parameters accordingly.
 - a. Specify the output requirement in terms of either the number of phrases to output or the minimum acceptable phrase weight ("threshold weight"). The default is to generate 10 phrases, but the user can select some other number, or specify a threshold weight. If both are specified, the threshold weight is ignored. There are not always enough phrases available to meet the output requirement. In such cases, the available phrases will be output after the search is exhausted.
 - b. Specify the number of stopwords that may appear in each of the output phrases. The default number is zero. The user has the option to allow up to N stopwords in each output phrase. If the user includes a stopword as a query word or in a query phrase, but does not allow at least one stopword per output phrase, no phrases will be generated for that query.
 - c. Specify the stoplist. Use the default stoplist, the default stoplist plus user-specified stopwords, or a user-specified stoplist.
 - d. Specify the phrase database, a phrase-oriented QUORUM model of the collection of narratives in the ASRS database. Use one of the default phrase databases or a user-specified one. When using one of the default databases, if any stopwords are allowed, use the phrase database that includes stopwords, otherwise use the database without stopwords.

A phrase-oriented QUORUM model is a set of QUORUM phrase relations having standard, left, and right relational metric values derived in phrase-sized contexts. (See appendix 1 and also appendix 3, step 14, paragraphs 2 and 3.) Each relation represents measured contextual associations between two words. A phrase-oriented model of a narrative represents the contents of the entire narrative, and it implicitly represents all of the phrases in the narrative. A phrase-oriented model of a collection represents the contents of all of the narratives in the collection, and implicitly represents the phrases in all of the narratives of the collection. Any particular set of two words must co-occur within a phrase-sized context (see appendix 3, step 6) in order for their relationship to be represented in the model. The two words need not be immediately adjacent to one another.

The phrase database used by QUORUM phrase generation is a single phrase-oriented model of the collection of narratives in the ASRS database. It is derived from the tens of thousands of phrase-oriented models that each represent one narrative. This is done by accumulating the sums of the standard, left, and right relational metric values of relations having the same word pairs, adjusting the metrics to account for word order in the text. Given the very large number of relations, a culling process is performed as the sums are accumulated. Culling eliminates particularly rare and weak contextual associations so as not to overflow computer memory. The parameters of that culling are the number of narrative models to be taken as a block, and the minimum accumulated standard relational metric value required of each relation after accounting for all relations in each block of models.

2. Read the user's query input: a keyword, a phrase, or a file of keywords and/or phrases (one per line). Capitalize the query text, since text is capitalized in the narratives of the ASRS database. Map the query text to ASRS abbreviations and usage.

3. Process each query word and phrase.
- 3a. If no query words or phrases remain to be processed, then phrase generation is complete. Otherwise, begin to process the next query word or phrase. Clear the output buffer. If the user did not specify a threshold weight (i.e., minimum acceptable phrase weight), start with an initially somewhat high threshold weight that will gradually be lowered (see step 3h) until the required number of phrases is available for output or the search is exhausted.
- 3b. If the user did not supply a weight for this query word or phrase, supply one. Given a word that is not a stopword, use the frequency of that word in the ASRS database. If it is not in the vocabulary list (which contains no stopwords), set the weight to an arbitrarily high number such as 1000000 so that it will appear at the top of the sorted list. Given a phrase, use the least frequency of the non-stopwords in the phrase.
- 3c. Add the current list of phrases to the output buffer. On the first pass of steps 3c and 3d, the current list consists of a query word or phrase provided by the user. On subsequent passes for this word or phrase (looping back from step 3e), the current list of phrases consists of phrases built upon the previous list. Each pass attempts to create phrases that contain an additional word. Each listed phrase can generate zero or more new phrases.
- 3d. Attempt to build phrases upon the current list of phrases. Compare each QUORUM phrase relation in the phrase database with the words in the phrases on the current list. (Initially, the current list typically contains a single query word, which is treated as a one-word phrase.)
 - a. If every word in a phrase has a non-zero phrase relation with a following word W , then add the word W following the words in the phrase to create a new phrase. The number of words separating each word of the initial phrase from the word W is not considered.

For example, given the two-word phrase $A B$, and a following word W , create the phrase $A B W$ if the phrase relation $R(A,W)$ has a non-zero rightRMV or $R(W,A)$ has a non-zero leftRMV, and $R(B,W)$ has non-zero rightRMV or $R(W,B)$ has a non-zero leftRMV. (See appendix 1.)
 - b. If every word in a phrase has a non-zero phrase relation with a preceding word W , then add the word W preceding the words in the phrase to create a new phrase. The number of words separating each word of the initial phrase from the word W is not considered.

For example, given the two-word phrase $A B$, and a preceding word W , create the phrase $W A B$ if the phrase relation $R(W,A)$ has a non-zero rightRMV or $R(A,W)$ has a non-zero leftRMV, and $R(W,B)$ has a non-zero rightRMV or $R(B,W)$ has a non-zero leftRMV. (See appendix 1.)
 - c. The basic weight of a new phrase is equal to the weakest of the phrase relations that associate the words in the phrase.
 - d. The basic weight is adjusted by subtracting from it the number of words in the phrase. When phrases have identical basic weights, this ensures that the ones having fewer words are slightly higher on the list. This is done because the shorter phrases are more likely to be found in actual narratives.
 - e. Retain newly-created phrases that have weights at least as large as the threshold weight.
 - f. Eliminate redundant phrases. Redundancy occurs, for example, when phrases $X Y$ and $Y Z$ are present in the list of current phrases and each is extended to $X Y Z$: from $X Y$ by adding a following Z , and from $Y Z$ by adding a preceding X , producing two copies of $X Y Z$.
 - g. Retain phrases that have no more than the specified number of stopwords.
- 3e. If any new phrases are produced by step 3d, return to step 3c, using them as the current list. Otherwise, continue to step 3f.
- 3f. If the output requirement is based on the threshold weight, sort the phrases in the output buffer in descending order of their weights and output them, then go to step 3a to process the remaining query words and/or phrases. Otherwise, the output requirement is a specified number of output phrases, so go to step 3g.

- 3g. Count the number of phrases in the current output buffer. If the number is less than the specified number of output phrases (N), then go to step 3h. Otherwise, sort the phrases in the current output buffer in descending order of their weights, print the N phrases having the largest weights, then go to step 3a to process the remaining query words and/or phrases.
- 3h. Lower the threshold weight.
- a. Adjust the initial threshold to reflect the frequency of the query. If the database frequency of the current query word or phrase is less than the current threshold weight, use it as the new current threshold weight. This can only occur during the first lowering of the threshold weight for the current query when the output requirement is a specified number of output phrases. It is necessary because the initial threshold weight is set to a high value that prevents very common words from initiating a time-consuming depth-first search. The adjustment in this step lowers the current threshold weight to an upper bound of the actual database frequency of the particular query. As long as the threshold weight is greater than the database frequency of the query, no phrases will be generated. Unfortunately for maximum efficiency, it is currently much easier to check this after the first attempt to generate phrases based on the current query.
 - b. Adjust the rate at which the threshold is lowered. If the number of phrases generated during this pass is the same as that generated during the last pass or if the number found so far is less than or equal to half the number sought, multiply the current threshold weight by $1/2$ and truncate to the nearest integer. Otherwise, multiply the current threshold weight by $3/4$ and truncate to the nearest integer.
 - c. If the threshold reaches the minimum possible value, print the N or fewer phrases having the largest weights, then go to step 3a. Otherwise, start over using the new, lower threshold: clear the output buffer, replace the current list of phrases with the query word or phrase, and go to step 3c. The lower threshold makes more QUORUM phrase relations available for phrase building in step 3d.

Note: Picking the best threshold value on the first try, before any processing, would improve performance. The best threshold value is the one that enables generation of the desired number of phrases in one pass. That value ultimately depends on the metric values in the phrase database used in step 3d.

Appendix 5. How QUORUM phrase discovery works

This appendix describes the new QUORUM phrase discovery method, which is used to discover phrases that are topically associated with user-specified words or phrases. Some steps are automated and some are manual.

The process is one of distillation. The initial search in step 1 provides narratives containing the original query terms or phrases. Phrases extracted from these narratives broaden the next search (the first pass of step 6), but since they are derived from narratives that contain the original terms, they are still somewhat limited. The narratives retrieved by using the extracted phrases as query terms do not necessarily contain the original query terms, so it is from them that the last set of query phrases is derived. Using these in a final search of the database (the second pass of step 6), a final collection of narratives is retrieved. The most unrestricted set of topical phrases is derived from these narratives. The three lists of derived phrases are combined in step 7b in order to weight the phrases according to the collection of narratives in which each phrase is most prominent. The result of the process is a frequency-ranked list of topically related phrases, along with a final collection of narratives which are ranked on their relevance to those phrases. Neither the phrases nor the narratives necessarily contain the original query terms.

1. First, use QUORUM keyword search, QUORUM phrase search, or any other means to find the initial set of narratives that are relevant to the topic of interest.

To illustrate the method, keyword search will be used. First, identify topical words of interest, for example: fatigue, tired, sleep, circadian. Review the vocabulary of the text in the database to see what forms of these words are present among the narratives in the database. In this case, the forms present are: fatigue, fatigued, fatiguing, tired, tiredness, sleep, asleep, sleeping, sleepy, circadian. Do a keyword search on this collection of words.

2. Extract phrases from the top N narratives. A useful value of N is 200.

Scan each narrative term by term, in reading order, in one pass. At each term, if it is not a stopword or punctuation, then consider it as the first word in a possible phrase of 2 words, then one of three words, and so on, to W words, typically 8. Meanwhile, however, if the last term in any possible phrase is a stopword or punctuation, ignore that possible phrase. If more than S interior terms of the phrase are stopwords, or if any is punctuation other than "-" (dash), ignore the possible phrase. The value of S and the list of stopwords are selectable by the user.

For example, given the text: "A B C D E F G H I J", where each letter represents a term (e.g., a word, number, or punctuation), if C, E, F, and I are stopwords or punctuation other than "-" (dash), and if the value of S is 1 (allowing up to one interior stopword), the following table shows all possible phrases. The underlined phrases in this table would be retained as phrase candidates, while the rest would be ignored.

<u>A B</u>	C D	E F G H
A B C	C D E	E F G H I
<u>A B C D</u>	C D E F	E F G H I J
A B C D E	C D E F G	F G
A B C D E F	C D E F G H	F G H
A B C D E F G	C D E F G H I	F G H I
A B C D E F G H	C D E F G H I J	F G H I J
B C	D E	<u>G H</u>
<u>B C D</u>	D E F	G H I
B C D E	D E F G	<u>G H I J</u>
B C D E F	D E F G H	H I
B C D E F G	D E F G H I	<u>H I J</u>
B C D E F G H	D E F G H I J	I J
B C D E F G H I	E F	
B C D E F G H I J	E F G	

For each unique phrase among those retained so far as phrase candidates, count the number of times they appear in the whole collection of text (e.g., in 200 narratives).

Next, process phrases contained within other phrases. If phrase X is contained in phrase Y and the frequency of X is no greater than that of Y, eliminate phrase X from further consideration. This means that if a sub-phrase only occurs as part of a larger phrase, and never stands alone, it is deleted from the set of retained phrases. In the above example, only phrases A B C D and G H I J avoid deletion and are output as the extracted phrases.

Output each remaining phrase and its frequency, in descending order of frequency. Here are the top 10 phrases from the first pass of the fatigue example:

34	RESERVE OR STANDBY	23	HOLD SHORT
32	REST PERIOD	22	CONTINUOUS DUTY
25	APCH CTL	21	ASSIGNED ALT
25	CREW FATIGUE	21	CREW REST
24	CREW MEMBERS	18	ALT ALERTER

These phrases are all relevant to fatigue. Some phrases are situationally relevant, while others are topically relevant. The non-automated part of this process further distills the phrases, but it is important to note that even the fully automated first pass produces this useful set of topically and situationally relevant phrases.

- Determine which words are most prominent among the phrases. To do this, count the number of occurrences of each unique non-stopword among the phrases, and sort those words in descending order of frequency. Here are the top 10 words from the phrases derived in the first pass of step 2:

214	REST	142	FATIGUE	89	PERIOD
189	CREW	140	APCH	88	STANDBY
185	RWY	138	SLEEP		
163	DUTY	120	ALT		

From this list, and referring to the phrases, manually select words that are related to the topic of interest. Of the ten words shown, for example, retain REST, DUTY, FATIGUE, SLEEP, and STANDBY.

- From the phrases gathered in step 2, extract those that contain any of the topical words selected in step 3. Here are the top 10 phrases that are retained in the first pass:

34	RESERVE OR STANDBY	18	REDUCED REST
32	REST PERIOD	18	REST PERIODS
25	CREW FATIGUE	17	DUTY PERIOD
22	CONTINUOUS DUTY	17	PLT FATIGUE
21	CREW REST	15	SLEEP THE NIGHT

- (This step is for fine-tuning and is not always necessary.) Manually review the resulting phrases, and delete any that are not on the topic of interest. Subtract the resulting list from the whole list generated in step 2 in order to list the phrases that were not selected. Manually review them. If any of these are on the topic of interest, manually add them to the list of topical phrases. Adjustments, if any, are usually among rarely occurring phrases.
- Use the accumulated topical phrases as input to QUORUM phrase search to find topically related narratives.
- Repeat steps 2-5 (i.e., extract and filter phrases), using the output of step 6. Then combine the resulting list of topical phrases with the first one. If any phrases appear more than once, use the one with the highest frequency. Repeat step 6 then use the output in step 7b.
- Repeat step 2 (i.e., extract phrases from a collection of narratives). Then combine the resulting list of topical phrases with the first two. If any phrases appear more than once, use the one with the highest frequency. At this point, a list of phrases that are topically and situationally relevant to the original query has been produced. To differentiate among these two categories of relevance, continue to step 8, otherwise skip to step 10.
- Manually label the phrases to assign them to two broad categories: 1) those that are topically relevant, and 2) those that are situationally relevant.

Here are the top 20 phrases related to the topic of fatigue/tired/sleepy/circadian. The topically relevant ones are labeled "topical". The others are labeled "situational":

152	REST PERIOD	topical	37	DUTY PERIOD	topical
109	REDUCED REST	topical	36	REST PERIODS	topical
79	CREW REST	topical	34	RESERVE OR STANDBY	topical
57	CONTINUOUS DUTY	topical	30	REST REQUIREMENTS	topical
46	CREW SCHEDULING	topical	28	CREW FATIGUE	topical

24	24 HR PERIOD	topical	19	CREW MEMBERS	situational
24	CHIEF PLT	situational	18	MINIMUM REST	topical
22	CREW DUTY	topical	18	REQUIRED REST	topical
20	CONTINUOUS DUTY OVERNIGHT	topical	17	PLT FATIGUE	topical
19	ADEQUATE REST	topical	16	COMPENSATORY REST	topical

9. Collect the set of topically relevant phrases as one of the products. The situationally relevant phrases might be useful for subsequent analysis of the situational contexts of the topic of interest, so collect them in a separate list.
10. Use the output of the final phrase search (second pass of step 6) as additional products. These include:
 - a. the most relevant narratives, with relevant sections highlighted, in HTML format;
 - b. the ranked list of all relevant narratives; and
 - c. the QUORUM query model used in the final database ranking.

Note: A fully automated version of this method is currently being developed.