15th ICCRTS
"The Evolution of C2"

Priority Intelligence Requirement Answering and Commercial Question-Answering: Identifying the Gaps

Topic 3: Information Sharing and Collaboration Processes and Behaviors
Topic 9: C2 Architectures and Technologies

Brian Ulicny, PhD
VIStology, Inc., 5 Mountainview Drive, Framingham, MA, U.S.A.

Gerald M. Powell, PhD
U.S. Army Research Laboratory HRED , Fort Monmouth, NJ, U.S.A

Christopher J. Matheus, PhD
VIStology, Inc., 5 Mountainview Drive, Framingham, MA, U.S.A.

Michael Coombs, PhD
VIStology, Inc., 5 Mountainview Drive, Framingham, MA, U.S.A.

Mieczyslaw M. Kokar, PhD
Northeastern University, Boston, MA, U.S.A.

Point of Contact:

Brian Ulicny, PhD
VIStology, Inc.
5 Mountainview Drive
Framingham, MA 01701 USA
508 788-5088
bulicny@vistology.com

Priority Intelligence Requirement Answering and Commercial Question-Answering: Identifying the Gaps

Brian Ulicny[a], Gerald M. Powell[b] Christopher J. Matheus[a], Michael Coombs[a] and Mieczyslaw M. Kokar[a,c]

[a] VIStology, Inc., 5 Mountainview Drive, Framingham, MA, U.S.A.
{bulicny, mcoombs, cmatheus}@vistology.com
[b] U.S. Army Research Laboratory HRED , Fort Monmouth, NJ, U.S.A
gerald.m.powell@us.army.mil
[c] Northeastern University, Boston, MA, U.S.A.
mkokar@ece.neu.edu

***Abstract**: Doctrinally, Priority Intelligence Requirements (PIRs) are information that the commander needs to know in order to make a decision or achieve a desired effect. Networked warfare provides the intelligence officer with access to multitudes of reports and sensor outputs. What technology can the intelligence officer use to find answers to PIRs in this sea of information? Recent developments in enterprise search technology have accelerated. To what extent can commercial search technologies assist with the task of PIR answering? In this paper, we outline doctrinal approaches to PIRs and PIR answering and explore how they have been adapted to contemporary warfare. We then explore five types of question-answering technologies: structured data technologies; unstructured text-based technologies; semi-structured or "tagged text"-based; logic or "semantic web" technologies; and social question-answering technologies. We identify gaps in the PIR answering process that cannot be filled by contemporary Commercial off-the-shelf (COTS) solutions.*

1. Introduction

Doctrinally, Priority Intelligence Requirements (PIRs) are information that the commander needs to know in order to make a decision or achieve a desired effect. Networked warfare provides the intelligence officer with access to multitudes of reports and sensor outputs. What technology can the intelligence officer use to find answers to PIRs in this sea of information? Recent developments in commercial search technology have accelerated and become ubiquitous in civilian life. To what extent can commercial search technologies assist with the task of PIR answering?[1]

In this paper, we outline the doctrine of PIRs and explore how it has been adapted to contemporary warfare. In traditional warfare, PIRs were requests for information about the disposition of the enemy and its assets; the related tasking, observations and assessment of answers were relatively straightforward. In the context of counterinsurgency and counterterrorism, the relationship between PIRs and the indicators that would answer them is less straightforward, and the tasking of assets and the collection of information to produce an answer are complex.

We then explore five types of question-answering technologies: structured data solutions; unstructured text-based or information retrieval solutions; semi-structured text or tagged

[1] This paper does not represent an endorsement by the Army Research Laboratory of any of the commercial products discussed.

text solutions; logic-based or "semantic web" solutions; and social question-answering solutions. We identify gaps in the PIR answering process that have not yet been filled by contemporary COTS solutions.

2. Priority Intelligence Requirements

Priority Intelligence Requirements (PIRs) drive the military intelligence collection process. PIRs are "those intelligence requirements for which a commander has an anticipated and stated priority in his task of planning and decision making" (FM 2-0 "Intelligence, section 1-32). PIRs are a subset of the whole spectrum of information requirements, broadly speaking, that a military intelligence officer (S2) and his staff are tasked with answering. Additional intelligence requirements aimed at filling gaps in commanders' understanding of the operating environment and requests for information may come from higher echelons, lower echelons, and lateral organizations, or from the intelligence staff itself, but it is the PIRs that an S2 has been tasked with that are most important. While military intelligence officers help in developing intelligence requirements, it is the commander who is responsible for designating an intelligence requirement as a priority. The intelligence staff regularly updates the commander on its progress toward answering each PIR.

Doctrinally, PIRs focus on the activities of objects and agents other than those controlled by friendly forces. Information requirements concerning friendly forces are called Friendly Force Information Requirements (FFIRs). Collectively, PIRs and FFIRs, are referred to as Commander's Critical Information Requirements (CCIRs) along with (optionally) Essential Elements of Friendly Information (EEFIs). EEFIs are essentially, PIRs from the threat's point of view for which the commander may want to understand the current state of the threat's knowledge or potential knowledge. (FM 2-0, 1-30, 1-32).

While emphasis shifts in various doctrinal publications, PIRs are generally supposed to[2]:

1. Ask a single question.
2. Be ranked in importance.
3. Be specific: Focus on a specific event, fact or activity.
4. Be tied to a single decision or planning task the commander has to make.
5. Provide a last time by which information is of value (LTIOV).
6. Be answerable using available assets and capabilities.

Usually, a commander only designates three to five intelligence requirements as PIRs at any one time.

The military intelligence staff is responsible for associating each PIR (or intelligence requirement, more generally), with one or more Indicators (Figure 1). Indicators are empirically observable variables about which information can be collected that would provide an answer to the overall PIR. Each Indicator is then associated with one or more Specific Information Requirements (SIRs) that detail how information is to be collected about an Indicator. Finally, these SIRs are assigned as concrete collection tasks to particular

[2] This characterization of PIRs is adopted from [11]. The authors survey the development of PIR doctrine since 1994 in FM 34-2 Collection Management and Synchronization Planning; FM 34-8-2, Intelligence Officer's Handbook, 1 May 1998, D-2 ; ST 2-50.4 ST 2-50.4 (FM 34-8), *Combat Commanders Handbook on Intelligence* B-2; and FMI 2-01 FMI 2-01, ISR Synchronization (2008)

personnel or sensors; assets are scheduled, units are deployed, and attempts to collect the information are made. Based on the collected information, the S2 produces answers to the assigned PIRs.

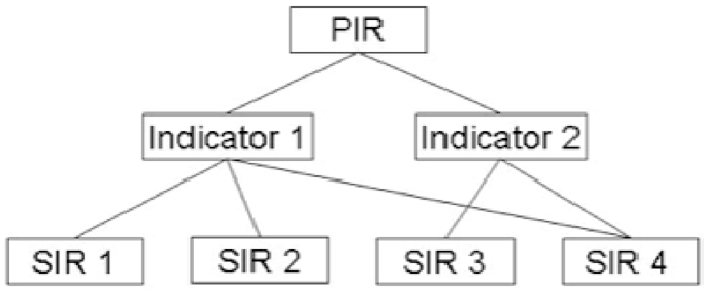

Figure 1 PIR Components

In recent years, there has been debate in military intelligence circles as to how to write PIRs that would be the most effective in the counterinsurgency, asymmetric warfare, and stability operations conflicts in which the U.S. military is currently engaged. Some have suggested that there should be more PIRs about the civilian or local government aspects of the operating environment [4]. Others have suggested that PIRs should focus more on 'Why' questions ("Why are counterinsurgency elements able to detonate IEDs in this Area of Operation?"), focusing on explanations rather than on more traditional questions about facts concerning the threat and other elements of the Area of Operations [10]. Many aspects of counterinsurgencies and asymmetric warfare have the effect of making the specification of Indicators that are relevant to a PIR less straightforward than it would be in conventional warfare, where PIRs are largely concerned with the location, courses of action, and disposition of enemy troops and equipment.

Outside the military, in recent years, the public has become comfortable with and increasingly dependent on Internet search engines as tools for answering questions both for their work and their personal lives. Search engines are incorporated directly into the web browsers by which users access the World Wide Web. Users increasingly turn to search engines to locate documents, phone numbers, driving directions, airline flight status information, product comparisons, breaking news stories, background information on potential employees, scholarly papers, patents, weather reports and a host of other information, both from their laptop and desktop computers as well as via handheld devices and over the phone.

To a certain extent, the PIR answering task is a question-answering task: queries are issued (PIRs, SIRs), and information is collated and presented as an answer to the commander in a briefing. To answer a PIR, an S2 must either identify relevant information that has already been collected or task the collection of information for an SIR. Then the S2 must identify all the relevant information that is being collected, locate relevant information that has already been produced, interpret it, and produce from it an answer to the commander's PIR based that can be briefed and justified to the commander. Information retrieval, broadly construed, and knowledge management form important elements of the PIR answering task.

Since information retrieval, question-answering, and search technologies have become ubiquitous in the civilian workplace as trusted and effective ways to navigate the huge amount of information accessible on the World Wide Web, it seems reasonable to ask (i) what aspects of the PIR Answering task can be automated by current technologies; (ii) where do gaps still exist in automating the PIR answering task, and (iii) what outstanding research issues do they present, if any?

3. Commercial Question-Answering Technologies

The idea of question-answering by computers has a long history in Artificial Intelligence and Information Retrieval [18, 23]. In this section, we classify contemporary information retrieval and question-answering systems by means of the way in which they internally represent information sources in order to produce an answer to a specified query. Our classification is based on published descriptions of the systems. In each case, we identify the representational scheme; provide examples of some familiar technologies using that representation; and provide examples of more recent, advanced applications that answer queries or retrieve information stored in that representation (Table 1).

Information Source Representation	**Common Application**	**Advanced Application**
Tables (Relational databases, spreadsheets, etc.)	Structured Query Language (SQL)	Wolfram Alpha
Text	Web search engines (Google, Ask, etc.)	TREC QA track; Aquaint (Advanced Question-Answering for Intelligence) systems
Tagged Text	Google Patent Search	Metacarta; Semantic MediaWiki; Palantir
Logic Statements	Prolog	

		Powerset (Microsoft Bing); Cyc
Trusted Teammates (Personal Knowledge)	Personal communication	Yahoo! Answers; Army Intelligence Knowledge Network Shoutbox

Table 1 Question-Answering Technology by Information Source Representation

3. 1 Structured Data

In computerized information systems, the use of relational databases has a long history. The vast majority of systems that store and retrieve data are based on representing the data in structured formats, in which the structure of the tables, and the significance of each column, is specified in advance. Structured Query Language (SQL) commands and queries are then used to insert and retrieve data elements in tabular form. While it has become increasingly sophisticated over the years, SQL was initially envisioned as a natural language interface to databases. In web-enabled database applications, the SQL queries and commands are mostly hidden from the user and are constructed and executed when a user fills out and submits a form on a web page.

The JC3IEDM (Joint Command, Control and Consultation Information Exchange Data Model) [13] is a vast relational data schema developed by the Multilateral Interoperability Programme (MIP) for encoding, storing and exchanging all information relevant to the contemporary battlespace. JC3IEDM 3.0 contains an "Intelligence Extension" that contains

contains data models for intelligence Requests and Request Answers. The model is designed to capture the most salient aspects of the military intelligence process, particularly as they are relevant to joint operations.

Wolfram Alpha represents a more sophisticated version of structured data querying. Wolfram Research is the producer of the major symbolic mathematical computation engine Mathematica. The Wolfram Alpha engine sits on top of quantitative data and other reference works that have been "curated" from authoritative sources [19]. When a user queries Wolfram Alpha (Figure 2) the engine attempts to interpret the query's intent so as to produce an output format that is the most likely to satisfy that query intention (here providing both a geospatial overlay and timeline as output, automatically defaulting to data within the last 30 years), without requiring the user to formulate the underlying Mathematica query him- or herself.

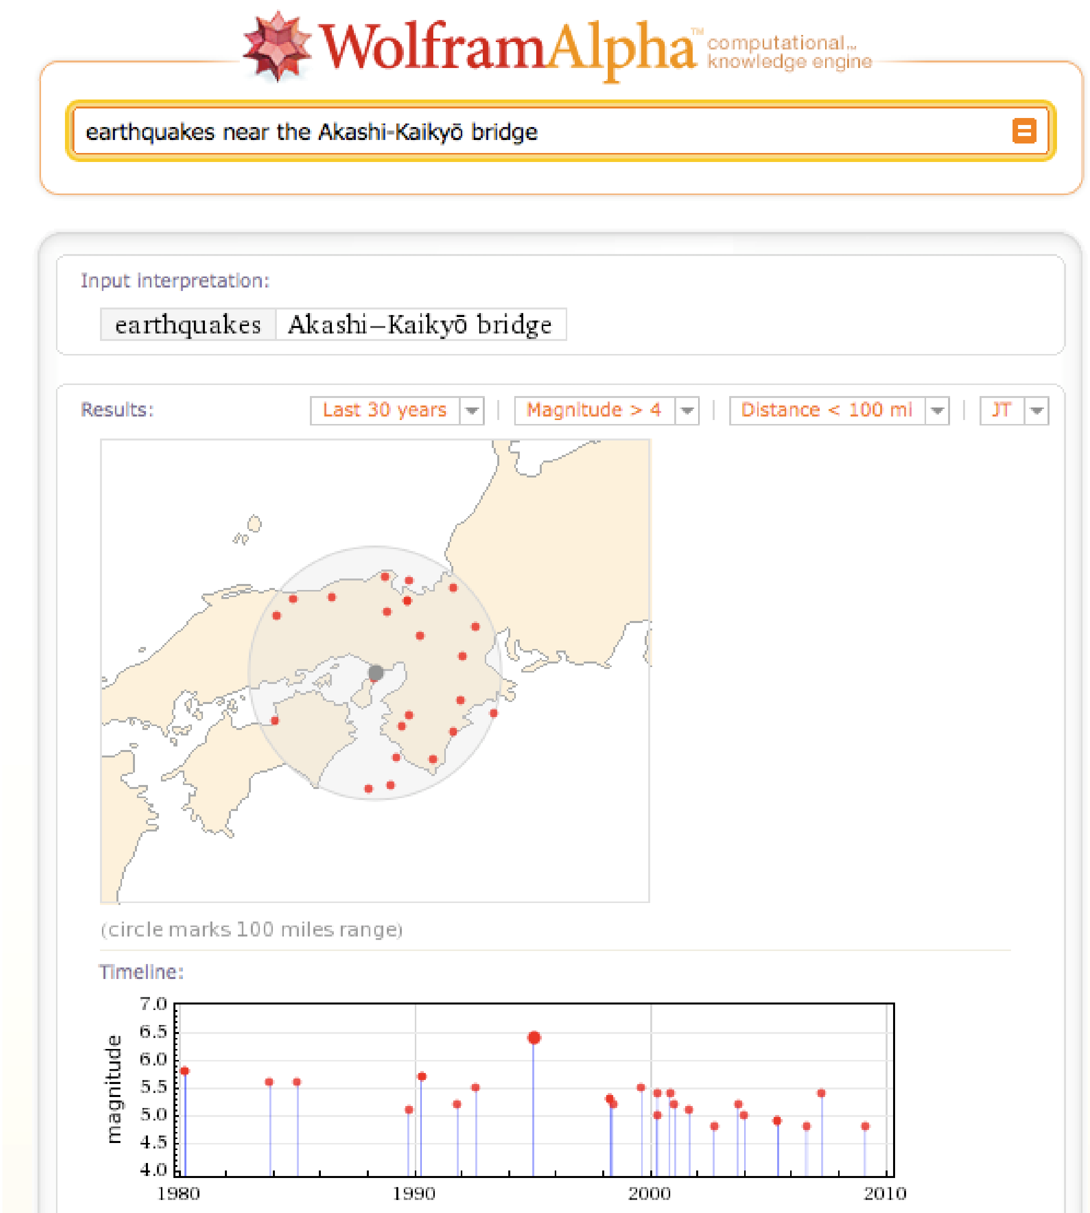

Figure 2 Wolfram Alpha output for query "Earthquakes near the Akashi-Kaikyō bridge"

3.2 Text

Text search engines are the most ubiquitous technology in Table 1. Users increasingly use text search engines such as Google, Yahoo!, Ask.com[3] and others many times a day in their work and personal lives. Search engines are not strictly speaking question-answering engines, because what they return is a ranked set of documents determined by the search engine to be relevant to the user query, not specific answers. The documents are ranked based on the frequency and position of the query terms in the document, as well as evaluations of the document's quality, as determined by Google's PageRank algorithm [15], or other measures of document quality. Document quality metrics like PageRank, which rely on properties of the hypertext network in which the document occurs, are not as applicable in a military context, in which information reports are not typically hyperlinked to other intelligence reports; rather, they are produced as standalone plain text documents.

Often, the user can determine the answer to a question posed as a search engine query simply by surveying the document snippets that the search engine returns, without clicking through to the documents themselves. Thus, from the snippets returned by Google below (Figure 3), one can gather that the consensus answer to the question "Where was Elvis born?" is Tupelo, MS. Notice that 'Elvis' is automatically interpreted as 'Elvis Presley' according to the results; no other Elvis is mentioned.

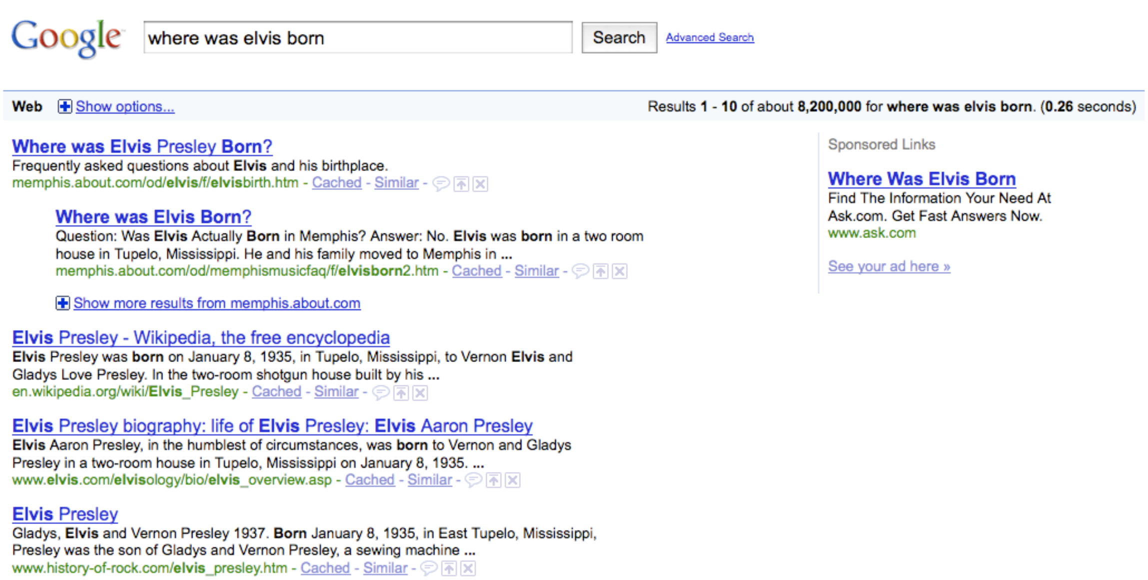

Figure 3 Google output for question "Where was Elvis born?"

From 1999 to 2007, the National Institute of Standards and Technology sponsored an automated question-answering competition as part of its Text REtrieval Conference (TREC), bringing together research teams from across the world to compete on identifying short

---

[3] http://www.google.com, http://www.yahoo.com, http://www.ask.com, respectively.

answers ("factoid" questions such as "How many calories are in a Big Mac?") and list answers (e.g. "Which countries don't have an extradition agreement with the US?") to questions in a shared evaluation context [14]. One of the techniques found to be effective in the TREC competitions was to identify consensus answers across document snippets identified using standard document retrieval techniques.

The Intelligence Community's Advanced Research and Development Activity (ARDA), now known as the Disruptive Technology Office, part of the Intelligence Advanced Research Projects Activity (IARPA), sponsored research on similar question-answering systems, from 2000 to 2006 [1].

3.3 Tagged Text

By "tagged text" systems, we mean question-answering systems that operate upon semi-structured data sources. By this we mean, textual data to which some information about objects, properties or relationships has been identified and marked up within the document, either as metadata separate from the body of the text, or marked up inline within the text, as in this bit of HTML, which indicates that the contents of this HTML element are of the kind "date-header".

```
<h2 class='date-header'>Thursday, January 21, 2010</h2>
```

While the identification and markup of such data elements can occur when a document is authored, many of the information retrieval systems using tagged data as their source data representation include automated text processing in which specific types of information are identified within a text using natural language processing techniques For example, a system might identify persons or organizations which are then inserted as markup into the system's representation of the document.

Examples of technologies making use of tagged data in question-answering or information retrieval include technology provided by companies including MetaCarta and Palantir. MetaCarta's technology [12] processes documents in order to identify any expression related to a location (e.g. location name or postal code), and marks up its representation of the document with geo-coordinates corresponding to that location expression. The system can then be queried for documents that contain some combination of keywords and have some geocoordinates within a specified bounding box or radius.

Palantir's analyst toolset allows users to ingest documents and tag specific elements according to a user-specified ontology [24] using an ontology framework that is equivalent to RDFS in allowing classes of objects, properties (RDFS dataypes) and relations (RDFS properties). Once the data is ingested and tagged, users can search and filter data according to objects, properties and events that have been identified in the unstructured or structured input data. Users can then filter this data according using the ontology, focusing only on data containing certain other objects properties and relations. Various "Helper" plug-ins then allow the user to visualize and analyze this data elements by means of histograms, social network analysis, geospatial overlays, timelines, and so on (Figure 4).

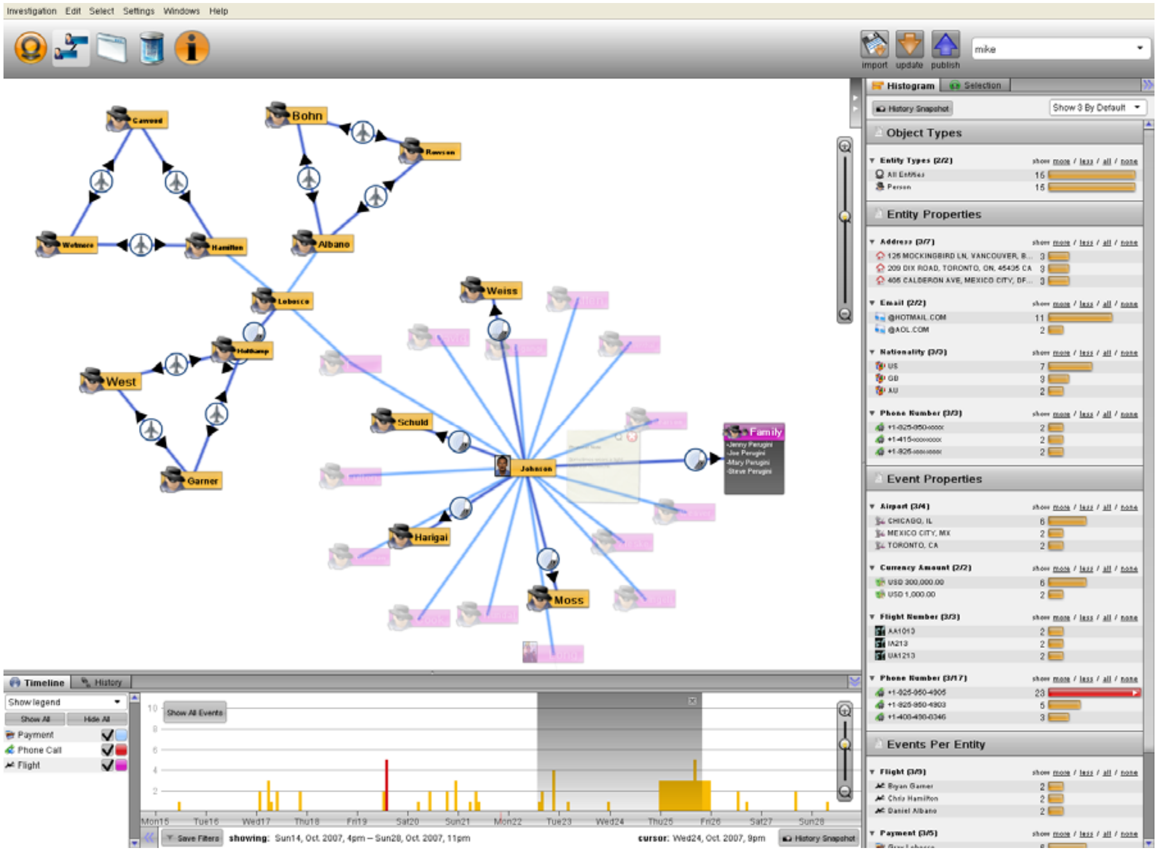

Figure 4 Palantir Screenshot from Palantir Tech Blog. The graph is linked to the histogram view that allows view filtering based on time ranges. (Source: http://blog.palantirtech.com/2008/07/04/palantir-screenshots-round-two/)

3.4 Logical Statements

Logic-based systems, such as Powerset, recently acquired by Microsoft and incorporated into its Bing search engine, parse all of a text into a logical representation, using sophisticated natural language processing [22, 20]. After analyzing free text and converting it into a logic-based representation, questions can formulated as queries over these logical clauses and returned as answers. Notice that in the example query below (Figure 5) clauses corresponding to various “Elvises” are highlighted, including Elvis Presley, skater Elvis Stojko, and Norwegian Elvis impersonator Kjell Elvis. Because Powerset is processing logical statements derived from a single source of data (Wikipedia), the answers do not converge on a single Elvis (Presley), as they do on the Web, where hyperlinking and network typology are used by Google to rank documents for return as well.

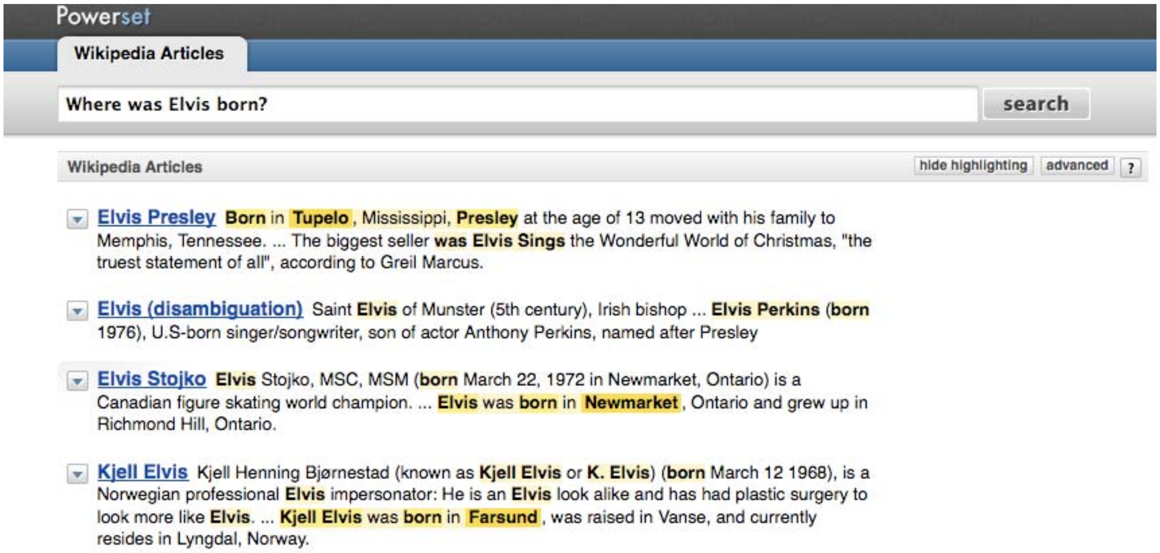

Figure 5 Powerset output for question "Where was Elvis born"

Many formal reasoning systems today have converged on a Subject, Predicate, Object representation of logical clauses (e.g. "John kissed Mary" has subject "John", predicate "kiss", and Object "Mary", ignoring tense) as being computationally more tractable than more flexible representations.[4] Standards for representing and sharing logic based information and reasoning about it are being developed under the rubric of the Semantic Web [2]. Resource Description Framework (RDF), Resource Description Framework Schema (RDFS), and the Web Ontology Language OWL are all triples-based standards adopted for such representations.

3.5 Trusted Teammates

Finally, the last representation of answer sources is knowledge contained within the heads of what we call "Trusted Teammates". Surely, the oldest technique for question-answering is simply to ask someone you trust who you believe knows the answer. In the context of the Web, this has been updated in services such as Yahoo! Answers[5] to allow users to pose answers to a community of online respondents, who provide answers asynchronously. The users can then use statistics compiled on the various respondents in order to assist in evaluating both the source and the content of the answer provided (the number of answers they have provided, their areas of expertise, the amount of positive feedback they have received and so on, Figure 6). A simpler question-broadcast service is incorporated as the "Shout Box" function in the Army's Intelligence Center Online Network (ICON) [3].

[4] Logical statements in triples or non-triples formats are equivalent, and they can be automatically transformed from one format to another, through a process known as 'currying' or 'Schoenfinkalization'.

[5] http://answers.yahoo.com

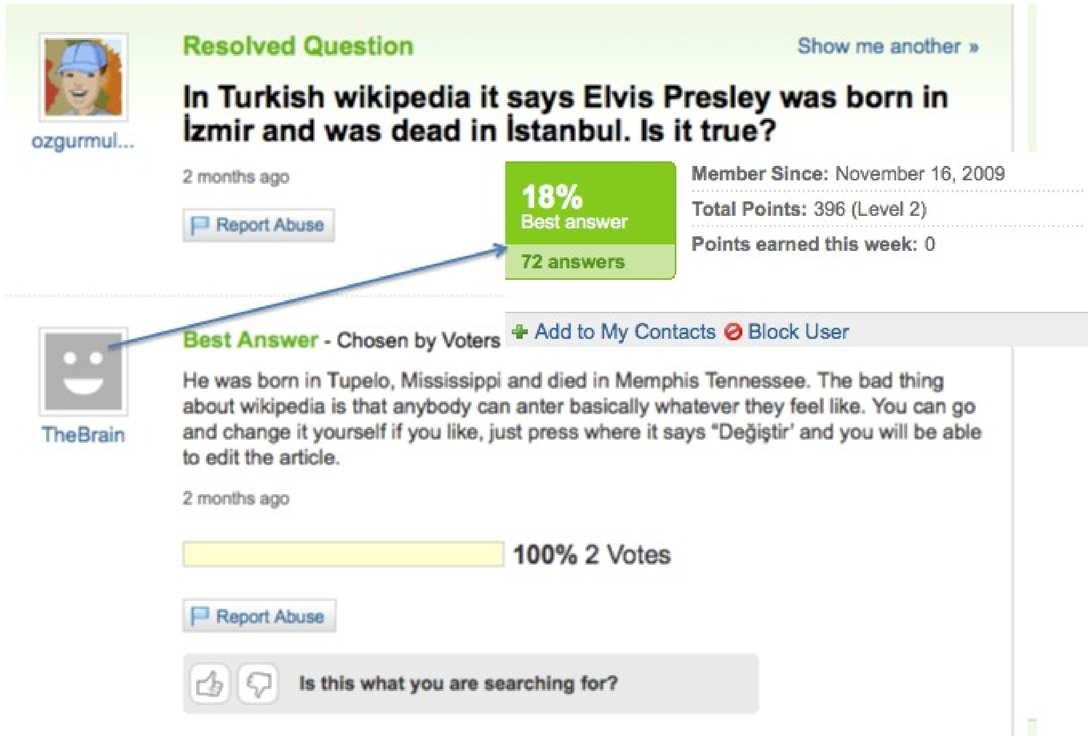

Figure 6 Question-answer example from Yahoo! Answers. The box indicated by the arrow displays metadata about the answer provider ("TheBrain") that the user could use to assess his or her credibility.

Having surveyed these classes of computer-aided question-answering systems, we turn to requirements for PIR answering.

4. PIR-Answering Requirements

The answer to a PIR must meet certain requirements. In Table 2 we list nearly all of the various doctrinal components of an intelligence report (INTREP) [8, p. 126]. Intuitively, the criteria for a satisfactory intelligence report should translate into requirements for a satisfactory answer to a PIR.

We have also gleaned requirements from the Intelligence Extension of the JC3IEDM 3.0 standard, which covers encoding information requests and their responses in a militarily satisfactory way (Table 3).

Table 2 INTREP Elements from FM6-99.2 (US Army Reports and Message Formats, April 2007).

Intel Report Element	Description

1. Date and Time	DTG
9. Sources	Reliability Rating of Source and Credibility Rating of Information

10. Evaluation	Evaluation of Source, Information, and BDA

Table 3 Intelligence Request and Request-Answer Requirements from JC3IEDM.

a. The degree of certainty of the data represented by each record must be modeled. This very general requirement is due to the fact that meta-information is important to the intelligence community.

b. An observation must be accompanied by information about the observer, the activity observed, the period of validity, the observation date, and a confidence factor.
c. A request for intelligence information must be modeled. A request for intelligence information identifies data that is currently not contained in the database. Such a request must be part of the data to be replicated so that some intelligence-collecting organization might be able to respond, collect the required information, update the database, and replicate the data to the requester.

d. The subject of a request for intelligence information may include any of the following topics: the identification of opposing force actions, capabilities, locations, as well as the identification of operational objects or their types.
f. A request for intelligence information can be permanent. Such requests (e.g., a standing request for information) always specify critical information needs.
g. A standing request for intelligence information can have an immediate impact. An example of such a request is the enemy's use of chemical, biological, or nuclear weapons.

h. A request for intelligence information can be linked to a resource cited in an operational plan. Such a request can be useful in validating a hypothesis but can also be necessary for an operational plan. For example, an attack in a given area can be executed under assurance that no friendly forces are present in this area.
i. It is highly desirable that a response to a request for intelligence information be explicitly positive or negative. This requirement is related to the "closed world" assumption, which states that, if data is missing in the database, then the data must also be missing in "reality." This assumption is not realistic from the intelligence point of view. Specifically, if the database indicates that no enemy forces are located in a given area, it must be ascertained that this is also true in "reality." An explicit negative response means that the given area was reconnoitered and no enemy forces were detected.

In response to the epistemic aspects of requirement (a) and (b), JC3IEDM Reporting Data elements contain the following codes (Table 4):

Table 4 JC3IEDM (Intelligence) Reporting Data Codes

reporting-data-accuracy-code—The specific value that represents, for intelligence purpose, the general appraisal of the subject matter in graded terms to indicate the extent or degree to which it has been judged to be free from mistake or error or to conform to truth or some recognized standard value
reporting-data-counting-indicator-code—The specific value that denotes whether the data referred to by a specific REPORTING-DATA is based on a count of objects
reporting-data-credibility-code—The specific value that represents, for normal operational use, the degree of trustworthiness of the data referenced by a specific REPORTING-DATA
reporting-data-reliability-code ↓The specific value that represents, for intelligence purpose,

the general appraisal of [6]the source in graded terms to indicate the extent to which it has been proven it can be counted on or trusted in to do as expected

Finally, based on the previous requirements and an understanding of the context in which it occurs, the PIR answering task requires:

Query Persistence: PIRs and SIRs are in force until they are retracted or the LTIOV deadline has passed. Therefore, attempts to answer them must be ongoing. (JC3IEDM requirement f).

Data Fusion: As suggested in Figure 1, PIRs can be complex, containing multiple indicators, and SIRs. Information corresponding to each SIR must be fused into an answer to the overall question. (FM 2-0, section 4-38)[7]

Partiality: Reports or other data sources may contain partial answers to an SIR or PIR. These must not be ignored, but must be retained and possibly conjoined with other partial answers to create a complete answer. This is what has become known as "connecting the dots" across documents, and even data repositories (FM 2-0, section 1-28).[8]

Incrementality: PIRs are often longstanding, and form the basis of a running assessment of the situation. As such, it is often important to determine how the answer to a PIR has changed over a preceding period. (JC3IEDM requirement g).

5. Evaluation of Current COTS Technologies with Respect to Intelligence Requirements

Given the requirements for intelligence reporting outlined in the previous section, to what extent can current technologies provide the required functionality? In this section, we aggregate the requirements described above and assess the ability of commercial technologies to address them.

5.1 Source

Requirement: Each element of information in a PIR answer should have a determinate source, time and location detailing who provided the information. (Table 2, item 9; Table 3, item b).

Structured data technologies can provide any required data element, such as identity of source or observer, when the activity was observed, and so on, as long as the data is already structured this way. The JC3IEDM standard provides a data model that can be implemented

[6] JC3IEDM MAIN - DMWG 20090514 Edition 3.0.2, p. 379, Section 13.2.1

[7] FM 2-0 Intelligence 2004. 4-38. In situation development, the intelligence staff analyzes information to determine its significance relative to predicted ECOAs and the CCIRs (PIRs and FFIRs). … They use the indicators developed for each ECOA and CCIRs (PIRs and FFIRs) during the MDMP as the basis for their analysis and conclusions. (Approved for public release, distribution is unlimited)

[8] The produce step [of the intelligence process] involves evaluating, analyzing, interpreting, **synthesizing, and combining information and intelligence from single or multiple sources** into intelligence or intelligence products in support of known or anticipated requirements. FM 2-0 Intelligence (Approved for public release; distribution is unlimited). Section 1-98. p. 1-28. (emphasis added)

by structured data solutions. We are not aware of any commercial technologies that implement this data model, however.

Text, tagged text and logic based technologies typically return only the URL of a source document, not the specific person, organization or sensor that provided the information.[9] In order to identify the source, a user must either decode the URL or identify the author of the page. Search engines based on text, tagged text, or logic statements do not provide any information about the location or time of the source's observation, or how long it is valid. Information retrieval systems do not usually identify who is responsible for the information in a document, by evaluating footnotes or attributing quoted materials. They only identify the source document. Extracting who is responsible for each element of data would require sophisticated text analysis; we are not aware of any system that implements this functionality. At best, some commercial text analysis software, such as OpenCalais[10], relates speakers to quotations in unstructured text.

In a social question answering setting, teammates provide only the information they choose to provide, so there is no guarantee as to what information about a source a question answerer will provide.

5.2 Reliability/Certainty/Accuracy/Confidence

Requirement: Each element of information in a PIR should have an assessment of the accuracy of the information provided, how certain it is, and how reliable the source is. The overall answer provided should be provided with an assessment of the analyst's confidence in it. (Table 2, items 9 & 10; Table 3 item a; Table 4).

Again, the JC3IEDM data model shows how such epistemic information could be incorporated in a structured data system, but we are not aware of any commercial application that implements it.

For systems that require extracting information from unstructured text, we are aware of no systems that attribute how certain or reliable information in a text is, nor represents uncertain information in a systematic way. For example, text analysis systems may identify "Mosul", say, as a location, and perhaps provide its geocoordinates, but they usually do not attempt to provide geocoordinates for an uncertain location, such as "near Mosul". Nor do they rank "Between five and ten men threw rocks at the convoy" as a less certain response than "Seven men threw rocks…" to the query "How many men threw rocks at the convoy?"

Information retrieval applications do use metrics such as Google PageRank or Technorati Authority (a measure of the number of inlinks a blogger receives over the last six months) to rank retrieved documents qualitatively as information sources. In this way, information retrieval systems use hyperlink topologies as a basis for document quality metrics. Such hyperlink-based metrics are not generally applicable in the military intelligence context,

[9] One exception is Technorati's blog search (http://www.technorati.com), which provides the title of a blog as the source. Twitter Search (search.twitter.com), on the other hand, provides only the username and content of Twitter status updates relevant to a search query, but it does not provide a URL for the status update itself as the default format of search results.
[10] http://www.opencalais.com

since reports do not typically hyperlink to other reports or data sources, nor are they typically linked to by other documents.

In social question answering systems, such as Yahoo! Answers, we have seen that statistics aggregated about other people's assessments of a question answerer allows a user to evaluate the quality of a respondent as well.

5.3 Activity/Unit/Size/Equipment/Time/Location

Requirement: PIR answers often involve analyzing what happened. This requires a representation of data in terms of events, when they happened, how they happened (by means of what equipment), who made them happen and so on (Table 2, items 3-8; Table 3, item b).

The JC3IEDM model provides a data model for such representations, but there are no commercial systems that implement it that we are aware of. Tagged data systems such as Palantir can represent event data and their participants in this way, particularly if structured data with the appropriate fields already exists.

In general, it is a challenge for systems based on processing unstructured text to identify the depiction of an event and identify what happened, who participated, who or what underwent the event, what things or methods were used, when it occurred, and so on. This applies to both tagged text and logic-based solutions that require some initial, automated semantic analysis of textual information. While commercial text analysis tools such as OpenCalais can identify some events and relations in free text, our experience has been that they do not identify most of the events and their participants in typical military intelligence reports. Moreover, the types of entities identified by commercial text analysis tools are not as specific as those needed for military intelligence purposes. Classes of items such as weapons or military units are typically not identified; only more general categories such as equipment or organization are identified.[11]

In social question-answering applications, again, there is no guarantee that the required information will be provided.

5.4 Data Fusion

Requirement: Answers to multiple SIRs must be fused into an answer to the overarching PIR. (FM 2-0, section 4-38.)

Structured data solutions and some text based or tagged text base solutions can do some simple data fusion across multiple documents. That is, structured data queries can calculate the average value of some data element across multiple database rows, for example. Similarly, we have seen that information retrieval solutions can produce answers

---

[11] The complete set of types of entities identified by Open Calais 3.0 is: Anniversary, City, Company, Continent, Country, Currency, EmailAddress, EntertainmentAwardEvent, Facility, FaxNumber, Holiday, IndustryTerm, MarketIndex, MedicalCondition, MedicalTreatment, Movie, MusicAlbum, MusicGroup, NaturalFeature, OperatingSystem, Organization, Person, PhoneNumber, Position, Product, ProgrammingLanguage, ProvinceOrState, PublishedMedium, RadioProgram, RadioStation, Region, SportsEvent, SportsGame, SportsLeague, Technology, TVShow, TVStation, URL

across multiple documents to list-type questions (e.g. "Which countries have extradition treaties with the U.S?").

However, none of the technologies we have surveyed make any attempt to fuse the answers for multiple subquestions (SIRs) into the answer to an overarching question, unless the question is posed this way in a social question-answering system. Transforming answers to component questions into answers to the parent question is left entirely to the analysts themselves. The analyst must decide how much to emphasize one indicator over another, or the reliability of one answer over another, and construct and evaluate an answer to the entire PIR. Automating this effectively would seem to be an open research question.

5.5 Query Persistence

Requirement: PIRs are standing queries, and any new relevant information must be considered in answering them (Table 3, item f).

The corpus of information that can be used to answer PIRs changes rapidly. Potential answers must always be constructed with the totality of relevant information available at the time. This requires standing queries for PIR answering. Standing queries can be implement in structured data, information retrieval (text and tagged text) and logic-based systems. In logic-based systems, this is frequently implemented by means of forward-chaining inference, by which new information is joined to partial answers to queries that have already been constructed. In the other system types, queries are simply re-issued periodically.

In social question-answering systems, query persistence is more problematic because once someone has answered a question, they are less likely to answer it again. Some social question-answering systems implement incentive systems in order to motivate behavior that is counter to user predilections.

5.6 Partiality

Requirement: Partial answer elements must be identified and combined across data sources (FM 2-0, section 1-28).

In structured data systems, partiality is addressed by joining information across database tables. This requires knowing where the partial answers are in various tables and understanding how the tables are structured so as to join them. In text based information retrieval systems, it is difficult to identify sets of documents that jointly are relevant to a query even when none individually satisfies it. We know of no commercial system that implements such document set retrieval, as opposed to returning individual documents that entirely match a query (cf. [16]).

Palantir, and other tagged text systems, can construct answer sets by linking extracted data elements across documents, for example, in social network diagrams. If the initial data source is structured, then complete accuracy can be achieved unless entity names are ambiguous, or name variants are not used in different tables.

If the initial data source is unstructured and text analysis is required in order to extract entities, then the cross-document answer sets are only as good as the text analysis that feeds them. Text analysis engines are often not very context sensitive. Geotaggers will normally tag every instance of "Boston cream pie" as associated with the city of Boston, Massachusetts, for example.

In logic based systems, it is possible to join logical clauses extracted from different documents into a single answer, but cross-document co-reference becomes a problem, as it is in structured and tagged text solutions as well. For example, one document might say "*Hound Dog* was recorded in 1952" and another might say "Elvis Presley recorded *Hound Dog*", from which one might infer "Elvis Presley recorded *Hound Dog* in 1952", but this inference would be wrong, since the first clause refers to "Big Mama" Thornton's original recording of Hound Dog and the second refers to Elvis' remake.

In social question-answering systems, respondents naturally provide partial answers when they don't know a complete answer.

5.7 Incrementality

Requirement: PIR answers may require an assessment of how things have changed with respect to an information request over the last specified interval (Table 3, items e and g).

In any system in which data has been tagged temporally, it is possible to construct timelines or histograms of documents found matching a specified query criterion over time. For example, one can construct a timeline of incidents of a particular type stored as database rows, or documents containing certain phrases or tags, and bin them by time period. This can certainly be used to provide quantitative assessments of change over the specified period. Structured data or tagged text solutions can output graphs or other visualizations to represent change in a certain data variable over time. For the other types of systems, providing an assessment of change for a certain data variable over time is not functionality that is not generally attempted; at most, systems might provide 'term clouds' to show the most frequent or the most rapidly accelerating terms being used in a document set to provide some qualitative sense of what is going on in a recent set of documents, but it would be difficult to incorporate this data into a decision-making process.

In social question-answering systems, requests for incremental assessments are likely to be met with anecdotal evidence, unless the respondent already has the data analysis at hand.

Conclusion

In this paper, we have described Priority Intelligence Requirements and their place in the military intelligence process. We described five types of question-answering technologies based on their internal data representation. In the next section, we identified various doctrinal requirements for PIR answering. Finally, we surveyed how well the various types of question-answering technologies address the requirements of PIR answering outlined previously. Many significant gaps remain between the required functionality and the capabilities of existing systems. Several of the gaps identified, such as the fusion of multiple SIR answers into a PIR answer, the automatic analysis of intelligence reports into representations of militarily-significant events and their participants, and the automatic evaluation of how the answer to a query has changed over time, would seem to require significant research efforts.

References

1. ARDA (Advanced Research and Development Activity). AQUAINT (Advanced Question Answering for Intelligence) Bibliography. http://www-nlpir.nist.gov/projects/aquaint/aquaint.bib.082608.html

2. Berners-Lee, T., Hendler, J., Lassila, O., and Others (2001). The Semantic Web. *Scientific American*, 284(5):34-43.

3. Chunn, S. The Intelligence Center Online Network: walkthrough of a KM system. Military Intelligence Professional Bulletin. January,2008 http://www.thefreelibrary.com/The+Intelligence+Center+Online+Network:+walkthrough+of+a+KM+system.-a0193450188

4. Flynn, MG M. T., Pottinger, Capt. M., USMC, Batchelor, P. D., "Fixing Intel: A Blueprint for Making Intelligence Relevant in Afghanistan", Center for a New American Security (CNAS) Working Paper. January 4, 2010.

5. FM 2-0. Intelligence. 2004.

6. FM 34-2 Collection Management and Synchronization Planning, (8 March 1994).

7. FM 34-8-2, Intelligence Officer's Handbook, 1 May 1998, D-2 http://www.fas.org/irp/doddir/army/fm34-8-2.pdf. (1998).

8. FM 6-99.2 (FM 101-5-2), U.S. ARMY REPORT AND MESSAGE FORMATS April 2007. P, 126. Approved for public release; distribution unlimited.

9. FMI 2-01 FMI 2-01, ISR Synchronisation, 1 November 2008, (2008)

10. Guvendiren, Z.T., & Downey, S. (2009). Putting the "priority" back into PIR: PIR development in a COIN environment. *Small Wars Journal*, 1-14, retrieved on September 3, 2009 from http://smallwarsjournal.com/mag/docs-temp/216-guvendiren.pdf.

11. McDonough, LTC W. G., Conway, LTC J. A., "Understanding Priority Intelligence Requirements", Military Intelligence Professionals Bulleting, April-June 2009.

12. Metacarta http://www.metacarta.com/products-platform-information-retrieval.htm (Retrieved January, 2010).

13. Multilateral Interoperability Programme. THE JOINT C3 INFORMATION EXCHANGE DATA MODEL (JC3IEDM Main). Version 3.0.2. May, 2009. http://www.mip-site.org/

14. National Institute of Standards and Technology (NIST). Text Retrieval Conference Proceedings. http://trec.nist.gov/proceedings/proceedings.html

15. Page, L. Brin, S., Motwani, R., Winograd, T., (1998) The pagerank citation ranking: Bringing order to the web. Technical report. Stanford Digital Library Technologies Project.

16. Srihari, R.K., Lamkhede, S., Bhasin, A., "Unapparent Information Revelation: A Concept Chain Graph Approach," Proc. ACM Conf. Information and Knowledge Management (CIKM '05), ACM Press, Bremen, Ger., Nov. 2005, pp. 329-330.]

17. ST 2-50.4 ST 2-50.4 (FM 34-8), *Combat Commanders Handbook on Intelligence* (Fort Huachuca, AZ: United States Army Intelligence Center and Fort Huachuca, AZ, 6 September 2001), B-2. https://rdl.train.army. mil/soldierPortal/atia/adlsc/view/public/8684-1/st/2-50.4/st2-50.4.htm (2001).

18. Strzalkowski, T., Harabigu, S., Preface. Advances in open domain question answering. Springer. 2006

19. Talbot, D., "Search Me: Inside the launch of Stephen Wolfram's new "computational knowledge engine"." Technology Review. July/August 2009

20. Udell, J. (interview). Scott Prevost explains Powerset's hybrid approach to semantic search. http://channel9.msdn.com/posts/JonUdell/Scott-Prevost-explains-Powersets-hybrid-approach-to-semantic-search/

21. Ulicny. B. Lycos retriever: An information fusion engine. In Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers, pages 177-180, New York City, USA, June 2006. Association for Computational Linguistics.

22. Van den Berg, M., et al, FACT-BASED INDEXING FOR NATURAL LANGUAGE SEARCH. et al, US Patent Application. Application number: 12/201,596 Publication number: US 2009/0063550 A1 Filing date: Aug 29, 2008

23. Waltz, David L. (1975) Natural-Language Question-Answering Systems. Clinic on Library Applications of Data Processing, April 27-30, 1975, ed. F.W. Lancaster. Urbana, Il: Graduate School of Library Science: 137-144.

24. Wright, B. Payne, J. Steckman, M. Stevson, S. Palantir: A visualization platform for real-world analysis" IEEE Symposium on Visual Analytics Science and Technology, 2009. VAST 2009. Pp. 249-250