STRIDER: Toward an AI Collaborator for Intelligence Analysis

Scott Friedman, Mark Burstein, David McDonald
SIFT, Minneapolis, MN, USA
{friedman, burstein, dmcdonald}@sift.net

Roger Rosewall
George Washington University
rosewall@gwu.edu

J. Benton
NASA ARC & AAMU-RISE
j.benton@nasa.gov

Abstract

Intelligence analysts gather information from diverse sources and integrate it into a product that adheres to standards of quality and provenance, often under time pressures and information overload. The STRIDER system, which we describe in this paper, enables collaborative exploration, hypothesis formation, and information fusion from open-source text. STRIDER presents relevant information to the human analyst in an intuitive fashion and allows various forms of feedback through a diagrammatic interface to enhance its understanding of shared problem-solving objectives. The human analyst supports STRIDER's collaborative workflow by postulating new entities, events, associations, and hypotheses, to improve STRIDER's information extraction and relevance judgments. STRIDER models the analyst's objectives and focus in order to avoid presenting information to the analyst at the wrong time or in the wrong context, which could be distracting, or worse, misleading. The technology in STRIDER is motivated by known human cognitive biases and limitations, and compensates for these to improve the breadth, efficiency, and objectivity of intelligence analysis. We focus on two pillars of collaborative cognitive computing: (1) interfacing for bi-directional human-computer interaction that encodes the analyst's objectives and presents relevant information, and (2) support for mutual decision-making by the analyst and the system. We present preliminary empirical evidence to demonstrate STRIDER's effectiveness in extracting and identifying relevant information.

Introduction

This paper describes progress toward closing the cognitive loop in an intelligence analysis setting, where analysts face an information overload and require up-to-date information relevant to their intelligence objectives. We describe our progress toward this goal with our integrated system STRIDER (*Semantic Targeting of Relevant Individuals, Dispositions, Events, & Relationships*). STRIDER diagrammatically elicits the intelligence objectives of the analyst, automatically gathers relevant information from multiple sources of unstructured text, encodes necessary metadata, and presents information to analysts when relevant. This will facilitate compliance with quality and provenance policies and make analysts more efficient and effective. We describe STRIDER with respect to two primary pillars of cognitive computing: interfacing and decision support.

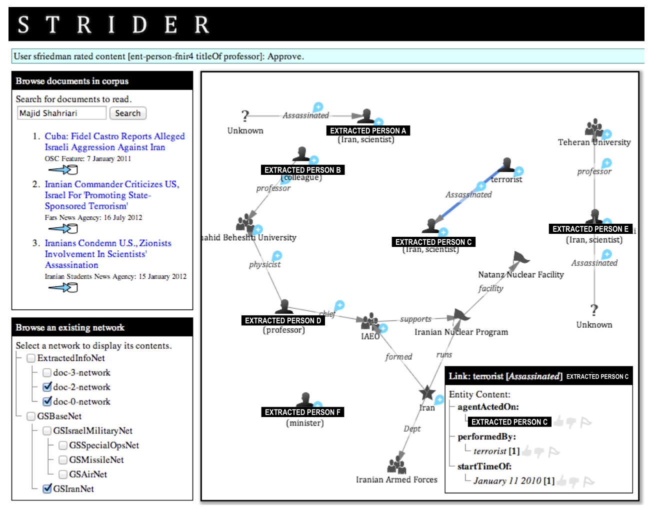

Figure 1: STRIDER's web-based interface.

Interfacing. STRIDER uses *link diagrams* to display individuals, events, organizations, and other entities as nodes in a network, connected by directed semantic links, as shown in the Figure 1 screenshot. Link diagrams are intuitive intelligence analysis interfaces: they do not require technical expertise with ontologies or knowledge representation, and other software systems use these representations for intelligence analysis (e.g., Carley et al., 2012, Stasko, Gorg, and Liu, 2008) and education. STRIDER exploits link diagrams as an interface for two purposes:

- **Soliciting objectives and queries from the analyst.** STRIDER's link diagrams have unambiguous semantics. This allows the analyst to extend and manipulate link diagrams to perform high-level fusion, specify objectives, and issue unambiguous directives to STRIDER.
- **Presenting relevant information to the analyst when appropriate.** While the analyst reads an article, STRIDER presents semantic information from that article— as well as semantically-related information from other sources— in a link diagram.

STRIDER uses the same diagrammatic interface to present information and to elicit objectives and feedback. This provides a shared workspace for the analyst and the machine to collaborate and share progress toward the explicit— and potentially changing— objectives.

Decision support. Just as human collaborators mutually influence and support each others' decisions, the analyst's actions support STRIDER's decisions, and STRIDER's actions are designed to support the analyst's decisions. The influence is bi-directional:

- **Analyst influences STRIDER's decisions** by extending or annotating link diagrams. For instance, STRIDER labels *information gaps* in the analyst's diagram, which become goals for STRIDER's information extraction. Also, as the link diagram grows or otherwise changes, the set of relevant entities and relationships change, which affects the space of information that STRIDER will decide is relevant enough to present to the analyst.
- **STRIDER influences analyst decisions** by presenting relevant information. STRIDER may thereby influence the analyst's focus, e.g., toward relevant organizations or events. Further, since STRIDER maintains provenance according to IC directives (Office of the Director of National Intelligence, 2007a,b), the analyst may expand relevant entities in the link diagram and peruse other supporting documents from the corpus.

This mutual influence is desirable from a collaborative workflow perspective: it allows the analyst to drive the objectives and utilize their deep intuition and common sense, while exploiting the machine's broad parallel processing and book-keeping. However, if the machine collaborator extracts erroneous information or displays irrelevant information to the analyst, it will distract or mislead the analyst and thereby derail the workflow. We conducted a pilot study, described below, to estimate the precision and completeness of STRIDER's information extraction and relevance judgments, compared to a senior IC analyst, and we present encouraging results. This pilot study precedes more detailed workflow analyses of analysts using STRIDER, which is a central goal of future work, as we describe below.

We continue by outlining the tasks and cognitive biases relevant to intelligence analysis, which motivate STRIDER's complementary cognitive computation. We then describe the STRIDER architecture and the information flows that support the pillars of cognitive computation described above. We present results from a pilot study to demonstrate STRIDER's effectiveness on these tasks, and close with a discussion of relevant and future work.

Strategic Intelligence Analysis

STRIDER's design is guided by the following guidelines of strategic intelligence analysis, based on Intelligence Community Directives (ICDs) (e.g., Office of the Director of National Intelligence, 2007a,b, 2009). We describe each guideline and STRIDER's contribution.

Objectivity. Analysis should be free of emotional content, regard alternative/contrary reports, and acknowledge developments. STRIDER supports this ideal with objective information extraction: deep semantic parsing extracts the semantics reported by the source; its only interpretive bias is the ontology with which it represents information.

Based on all available sources. Analysis should be informed by all relevant information available, and information collectors should address critical gaps. STRIDER explicitly identifies *information gaps* (i.e., missing data about an individual or event) and labels the source coverage (i.e., sources of information that support each datum).

Describe quality & reliability of sources. Open-source references should include metadata such as reference type, author, publication, title/subject, date, and more. STRIDER tracks all of these data, down to the specific paragraphs and character offsets supporting the extracted information.

Distinguish between intelligence & assumptions. Assumptions are hypotheses or foundations on which conclusions are reached, so critical assumptions must be explicitly identified, and so should indicators that may validate or invalidate assumptions. STRIDER helps analysts identify assumptions in their diagrams by automatically identifying unsupported information and information gaps.

Incorporate alternative analyses & hypotheses. Analytic products should identify and qualify alternative hypotheses in light of available information and information gaps. STRIDER's hypothesis-based organization (described below) helps analysts segment and compare competing hypotheses, which helps compensate for known cognitive limitations (Heuer, 1999, Johnston, 2005).

Timeliness. Analytic products must be disseminated to customers with enough time to be actionable. The integrated STRIDER system— from information-gathering to provenance to reporting— aims to improve efficiency.

Approach

Here we outline the general STRIDER approach, starting with STRIDER's interfacing advances, then describing supporting technology, and finally describing STRIDER's decision-making and how it is affected by the analyst's objectives and directives.

Diagrammatic Interfacing

STRIDER uses link diagram interfaces to display information and communicate objectives and queries. Figure 1 illustrates the link diagram display of STRIDER, which supports touch-based, pen-based, or mouse-based HTML5-enabled devices. Figure 2 illustrates how analysts manipulate STRIDER's link diagrams to express their intent.

STRIDER's interface provides an informative, intuitive, and domain-general *shared workspace* for human-machine collaboration, without requiring proficiency with ontologies or knowledge representation. To be sure, link diagrams are not as expressive as natural language, but as shown in Figure 2, annotating a link diagram offers significant flexibility for expressing objectives and queries. These annotations in Figure 2 include the following:

- **Connecting** annotations indicates that STRIDER should find direct or indirect relationships between existing entities or events in the link diagram.

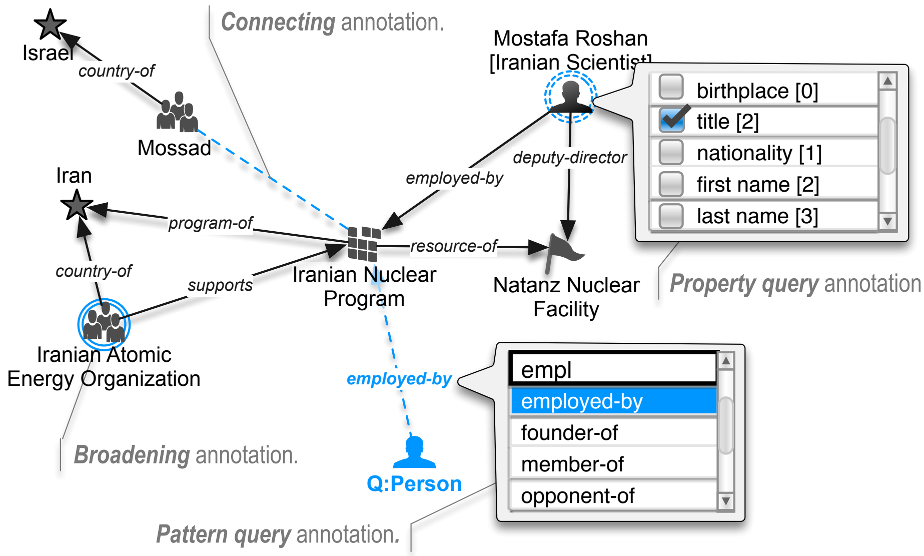

Figure 2: STRIDER's link diagram query interface.

- **Broadening** annotations indicate a broad interest in a specific entity or event in the diagram, and STRIDER should extract additional properties and events related to this.
- **Pattern query** annotations specify a pattern of interest—including a semantic relationship and one or more unknown entities— for STRIDER to match in the corpus.
- **Property query** annotations specify one or more properties of interest of an existing entity or event in the diagram, and STRIDER should extract additional evidence and/or values to fill that property.

STRIDER's use of link diagram manipulation to specify intent is inspired by Visual Query Systems (VQSs) for databases, web services, and other information repositories (e.g., Calvanese et al., 2010, Catarci et al., 1997). VQSs depict the domain of interest and express related requests, and aim to simplify complex query languages such as SQL and SPARQL. *Direct manipulation* (i.e., direct annotating or altering) of VQSs replaces the less-intuitive command language syntax, and benefits the user by reducing barrier of entry. This increases the ease of learning, providing high efficiency with experts, and reducing error rate (Ziegler and Fahnrich, 1988). We believe that STRIDER's direct diagram manipulations for querying and issuing directives, as shown in Figure 2, are novel interactions for specifying intent in a mixed-initiative information-gathering setting.

Deep Natural Language Understanding

Deep parsing allows STRIDER to extract precise semantics and determine entity types from local lexical context. STRIDER uses the SPARSER (McDonald, 1996) rule-based, type-driven semantic parser to read unstructured news articles. SPARSER'S rules succeed only if the types of the constituents to be composed satisfy the type constraints (i.e., value restrictions) specified by the rule. SPARSER compiles a semantic grammar from a semantic model of the information to be analyzed, including a specification of all the ways

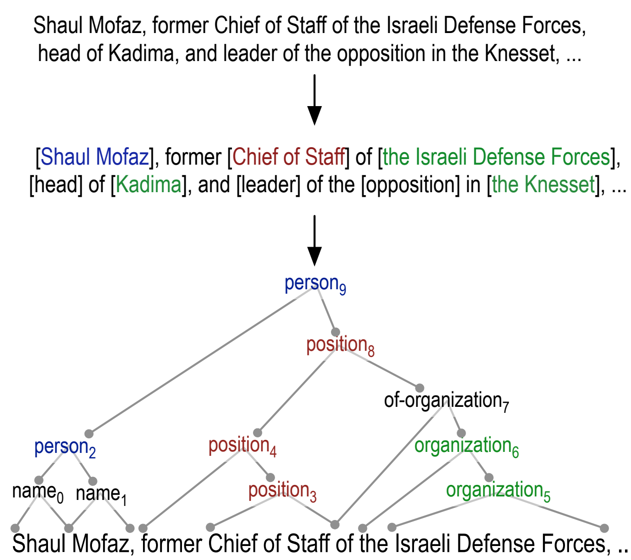

Figure 3: SPARSER efficiently analyzes text over multiple passes. Subscripts in the final semantic structure indicates the order in which SPARSER instantiated the instances.

each of the concepts can be realized in the language of the domain (e.g., open-source news articles). This ensures that everything SPARSER is able to parse it can model, and that every rule in the compiled grammar has an interpretation.

Figure 3 illustrates SPARSER'S scanning algorithm at a high level. In the first step, SPARSER segments or brackets the text into phrases, referred to as *segments*. Some of the segments include both known and unknown words, and some words are not included in any segment. SPARSER then detects instances of people, organizations, titles, times, locations, and more. It links these and other instances to exact locations in the document corpus to preserve the data source, in compliance with U.S. IC directives (e.g., Office of the Director of National Intelligence, 2007b).

Next, SPARSER uses rules and discourse heuristics to identify relations that connect the phrase segments. This is shown in the final SPARSER scan in Figure 3. Unless an established grammar and conceptualization applies, SPARSER relies on a set of textual relations drawn from its standard syntactic vocabulary. Along with the recognized entities, these segment-spanning relations represent the meaning of the text for STRIDER, and they are used to populate larger content models, as we describe below.

Organizing Competing Hypotheses

Tracking alternative hypotheses— and gathering evidence for each— is a recurring theme in the intelligence analysis literature, and is notably difficult for analysts due to cognitive biases and attentional limitations (Heuer, 1999, Johnston, 2005, Pirolli and Card, 2005).

Traditional link diagrams are *flat*, in that they conflate potentially disjoint or competing hypotheses into the same network structure. Consequently, they may contain entities and links that are unrelated, competing, or disjoint. We believe that competing hypotheses should be displayed and considered separately to preserve relevance and help the analyst weigh alternative outcomes.

We have integrated SPIRE inference and knowledge base (KB) technology— from our ongoing work on parsing biomedical texts (Friedman et al., 2016)— to house the STRIDER knowledge base. This allows us to store STRIDER sub-networks hierarchically in different *logical contexts*. Logical contexts support inheritance, as illustrated in Figure 4: the *Core* context contains the portion of the network shared by each competing scenario; and each competing scenario (e.g., *Covert* and *Missile*) inherits the *Core* network, but none of its sibling networks.

Importantly, the branching of hypotheses can continue beyond the single four-way split shown in Figure 4. For instance, SPIRE could support multiple sub-scenarios that inherit from the *Covert* scenario shown here, where each sub-*Covert* scenario contains mutually exclusive entities and relations. SPIRE also supports multiple inheritance of contexts, so a STRIDER context could inherit from both *SOF* and *Air* hypotheses to describe a joint strike.

Aggregating Semantic Content

STRIDER uses *content models* to organize, inherit, and prioritize knowledge about different types of entities, events, and relationships. Content models relate to the object-level ontology like a database view relates to a database. Each content model is associated with a category in STRIDER's ontology and specifies a partially-ordered list of properties that may be relevant for analysis. For example, the content model *Person CM* in Figure 5, associated with category *Person*, inherits all properties from the content model *Base CM*. STRIDER uses content models to aggregate presentable or queriable information, to support the following capabilities:

1. Detect gaps in information (i.e., unpopulated properties) or evidence (i.e., properties without support from the corpus). This helps the analyst manage uncertainty judgments (Heuer, 1999) and reduces the cognitive cost of monitoring for information gaps (Pirolli and Card, 2005).
2. Determine whether information should be rendered as a node (e.g., like the Person CM) or a link (e.g., like the Assassination CM) in the diagram.
3. Display relevant *drill-down* data for nodes and links.
4. Specify *equivalence classes* over categories, to help STRIDER detect equivalent entities (i.e., references to the same real-world entity or event) and data conflicts (i.e., multiple, inconsistent property values) within and across information sources.

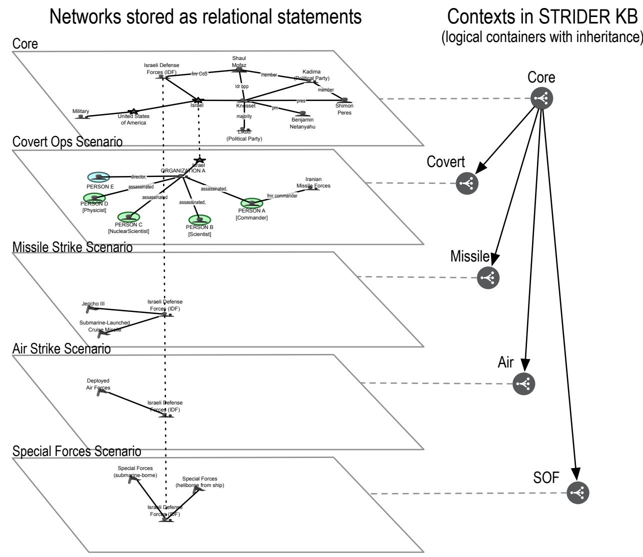

Figure 4: STRIDER records alternative hypotheses in separate hierarchical contexts.

Similarity-Based Reasoning

Given an entity, individual, or event of interest, STRIDER uses similarity-based retrieval to identify semantically similar *analogs* from its knowledge base. These retrieved analogs may help the analyst establish precedence and reason from previous examples to identify possible outcomes. STRIDER's similarity-based retrieval feature is motivated by well-known cognitive biases in memory retrieval and likelihood estimation; for instance, people use the sub-optimal *availability strategy* to estimate the probability of an event based on memory retrieval and imagination: assuming that if an event occurs frequently (and is therefore more probable), we can recall more instances of it (Heuer, 1999). Unfortunately, the ability to recall instances of an event is influenced by recency, context, vividness, and many other factors that are unrelated to the objective probability of an event.

STRIDER uses *structure-mapping*, a constrained graph-matching algorithm (e.g. Falkenhainer, Forbus, and Gentner, 1989, Friedman et al., 2016) to compute isomorphisms between semantic graphs and compute a numerical similarity rating. Given an entity (e.g., an Iranian nuclear scientist) in the network, STRIDER computes a subgraph of the entity

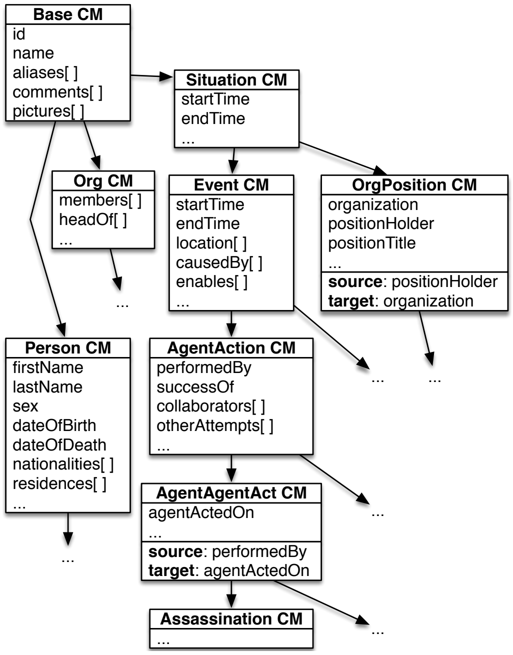

Figure 5: A portion of STRIDER's content model hierarchy.

from the content model and related events (e.g., the event describing the scientist's assassination), and matches it over the rest of the extracted semantic content in the KB to find similar analogues.

In this fashion, STRIDER not only builds a semantic network describing an event or topic of interest; it also relates different event descriptions using structural similarity, finding similar, or related, events and people, which helps to broaden the analysis.

Influences on STRIDER's Decision-Making

STRIDER's decision-making is influenced by the analyst's objectives (i.e., the analyst's annotated link diagram) and the analyst's focus (i.e., the article they are reading).

When STRIDER extracts information from text, it uses its content models to fuse the information into new or existing nodes or edges in its link diagrams. If the analyst has manually created nodes or links for entities and events (e.g., an individual or an organization), and then these are subsequently mentioned in an article, STRIDER will automatically extend the existing nodes and edges with the new information or evidence from the article; otherwise, STRIDER will generate new nodes and edges in the diagram.

To establish link diagram portions that are relevant to an article that the analyst is presently reading, STRIDER (1) parses the article, (2) grounds the entities and events referenced by the article within its link diagram(s), and then (3) displays entities and events within a specified link distance from the mentioned events or entities. STRIDER's distance-based relevance metric is effective for our current means, but we believe other methods, such as token-passing spreading activation, will yield better results as STRIDER accrues dense diagrams, as we describe in future work.

As mentioned above, if STRIDER incorrectly extracts information and then decides to present it— or if STRIDER decides to present otherwise irrelevant data— then STRIDER could distract or mislead the analyst. We next describe a pilot study to evaluate STRIDER's information extraction and fusion, compared to a senior intelligence analyst.

Information Extraction Pilot Experiment

We conducted a pilot experiment to assess STRIDER's ability to produce report-ready diagrams to aide a human analyst collaborator. In this experiment, a U.S. IC analyst consultant used a third-party diagram tool (i.e., *not* STRIDER) to build a link diagram from news articles, and we compare STRIDER's ability to extract, aggregate, and present the same information in a diagram, completely autonomously.

We tested STRIDER's information extraction on three articles from the Open Source Center (OSC, http://www.opensource.gov) in order to reproduce real link diagrams created by the analyst on the same material. The IC analyst used the third-party tool to build a *gold-standard* link diagram from many articles, citing individual sources. We used the a subset of the gold-standard diagram, including one of STRIDER's three articles, that described four assassinated individuals, an attempted assassination, and more. From the single-article portion of the analyst's gold-standard diagram, we counted 34 *data fragments*, including names, relationships, categories, events, dates, titles, nationalities, affiliations, organizations, and more.

STRIDER extracted all individuals correctly, most organization affiliations, two of four assassination events with dates intact, and more. However, due to gaps in parsing coverage, it missed two assassination events, it missed one unsuccessful assassination attempt, and it did not label a person's nationality:

- 33 total data fragments were extracted (e.g., names, relationships, events, dates, and titles)
- 30 fragments were in the analyst's diagram.
- 4 fragments (three events and one relation) were missing.
- 3 fragments (about an individual) were correct but deemed irrelevant by the analyst.
- 0 fragments were incorrect.

In total, the precision was 1.0, the recall was 0.88, for an F1 score of 0.94. STRIDER analyzed all three articles and recorded the sources of information in under one minute, and the human expert analyzed and recorded this information in one hour. Four elements were not extracted due to complex grammatical constructions that makes inter-sentence references to "bombings" and "attacks."

We subsequently gave STRIDER the remaining two OSC articles about the same events. From these, STRIDER extracted some consistent and some additional information and used its content models to fuse the new information into the diagram. The additional documents contained two previously-missing assassination events, and three university affiliations (two were diagrammed by the analyst, one is novel). Figure 6 illustrates the diagram STRIDER deemed relevant to the the initial article, with names hidden. It contains relevant information from the two subsequent articles,

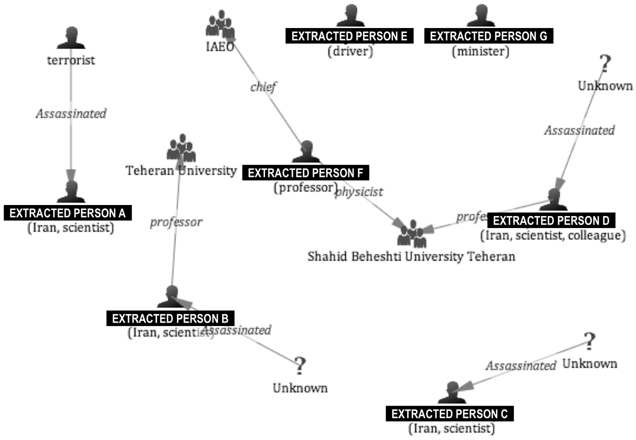

Figure 6: STRIDER screenshot (names hidden) showing information extracted and merged from three OSC articles.

hiding data that was determined irrelevant due to diagram distance from nodes mentioned in the article.

Overall, we demonstrated STRIDER extracting information, identifying gaps, presenting information and incorporating feedback using one consistent interface. These preliminary results suggest that STRIDER can:

- Extract information about individuals (e.g., their names, careers, nationalities, and titles), assassinations, organization affiliations, and role-based slot-fillers (e.g., an article mentions a"terrorist" assassinating somebody instead of a specific organization or individual).
- Extract partial information when complete information is not available.
- Utilize its content model and equivalence classes to merge and organize data across multiple sources.
- Display extracted information for non-technical users.

Information fusion is a time-intensive process of reading documents, extracting events and tying people and entities to related events. We demonstrated how a system like STRIDER can both help with the reading and information extraction process and also support the analyst in interactively tying relevant people and events together from multiple documents, improving the speed the process dramatically.

Related Work

STRIDER's interfacing exploits link diagram interfaces for information presentation and visual queries. Previous work in *network analysis* and *link analysis* utilize link diagrams as information displays. Organizational Risk Analysis (ORA) (Carley et al., 2012), a network analysis tool, automatically analyzes dynamic networks, social networks, geo-spatial data, and workflows. Jigsaw (Stasko, Gorg, and Liu, 2008) represents documents and their entities visually to help analysts examine them more efficiently, with an emphasis on illustrating connections between entities across documents. This reduces the cognitive load of data analysis. CRAFT (Gruen et al., 2008) supports wiki-like analyst collaboration with link diagram and form-based interfaces, to help analysts extend and share hypotheses, inquiries, and ontologies. These tools analyze link diagrams and support collaboration, but to our knowledge, they do not encode and autonomously react to analysts' changing objectives.

Semantic targeting in online advertising incorporates a semantic representation alongside an optional syntactic (i.e., keyword or bag-of-words) representation. Ad and page semantics can be evaluated for proximity (i.e., relevance) by calculating their taxonomic distance, allowing for semantic matches when no exact lexical matches are found. For instance, if a webpage describes a curling event, but no curling ads exist, ads belonging to the semantic class "skiing" (a sibling of class "curling" under the parent "winter sports") could be retrieved and delivered. STRIDER's relevance criteria is inspired, in part, by web-based semantic targeting.

Conclusion & Future Work

We described the cognitive computing technology of the STRIDER system for collaborative intelligence analysis, with a focus on (1) diagrammatic interfaces for eliciting the analyst's objectives and presenting relevant information, and (2) the influence of the analyst's objectives on STRIDER's decision-making. We presented preliminary results of STRIDER's information extraction and fusion, and we compared STRIDER's product to a senior intelligence analyst's diagram, with encouraging results (F1 = 0.94).

Our pilot experiment demonstrates that STRIDER can gather and present information, but it does not evaluate all of STRIDER's features as a collaborative closed-loop system. Evaluating the analyst's workflow with and without STRIDER will help qualify (1) how STRIDER addresses its user's cognitive biases and limitations, (2) the analyst time and effort required for inputting and maintaining information in STRIDER, and (3) STRIDER's overall impact on the breadth, objectivity, and efficiency of intelligence analysis.

In addition to evaluation, we have significant development remaining to realize our vision for STRIDER. For instance, STRIDER computes relevance using a diagram distance metric from the analyst's manually-created or annotated elements of the link diagram. We believe that a semantics-directed spreading activation algorithm will yield more complete and precise results. Also, we plan to have STRIDER learn from analysts' feedback on extracted data: if the analyst resolves a data conflict by choosing data from one source over another, STRIDER can generate or revise a topical model of source credibility.

Finally, STRIDER must support noisy data, e.g., as news reports are revised, as assumptions are violated, and as situations develop. This is crucial for cognitive aides in the intelligence analysis domain, to support human decision-makers in a partially-observable, uncertain, and changing world. We believe that metadata-based approaches for conflict resultion (e.g., Bleiholder and Naumann, 2006), in conjunction with human feedback and collaborative filtering, will help STRIDER semi-automatically prioritize conflicting data, but this is an empirical question and an area of future work.

References

Bleiholder, J., and Naumann, F. 2006. Conflict handling strategies in an integrated information system. In *Proceedings of the IJCAI Workshop on Information on the Web*.

Calvanese, D.; Keet, C. M.; Nutt, W.; Rodriguez-Muro, M.; and Stefanoni, G. 2010. Web-based graphical querying of databases through an ontology: the wonder system. In *SAC '10*.

Carley, K. M.; Pfeffer, J.; Reminga, J.; Storrick, J.; and Columbus, D. 2012. Ora user's guide 2012. Technical Report CMU-ISR-12-105, Carnegie Mellon University, School of Computer Science, Institute for Software Research.

Catarci, T.; Costabile, M. F.; Levialdi, S.; and Batini, C. 1997. Visual query systems for databases: A survey. *Journal of Visual Languages & Computing* 8(2):215–260.

Falkenhainer, B.; Forbus, K. D.; and Gentner, D. 1989. The structure-mapping engine: Algorithm and examples. *Artificial Intelligence* 41(1):1 – 63.

Friedman, S.; Burstein, M.; McDonald, D.; Paullada, A.; Plotnick, A.; Bobrow, R.; Cochran, B.; Pustejovsky, J.; and Anick, P. 2016. Learning by reading: Extending & localizing against a model. In *Proceedings of the Fourth Annual Conference on Advances in Cognitive Systems*.

Gruen, D. M.; Rasmussen, J. C.; Liu, J.; Hupfer, S.; and Ross, S. I. 2008. Collaborative reasoning and collaborative ontology development in craft. In *AAAI Spring Symposium: Symbiotic Relationships between Semantic Web and Knowledge Engineering*, 51–58.

Heuer, R. 1999. *Psychology of Intelligence Analysis*. Center for the Study of Intelligence, Central Intelligence Agency.

Johnston, R. 2005. *Analytic Culture in the United States Intelligence Community: An Ethnographic Study*. Central Intelligence Agency: Washington, DC.

McDonald, D. D. 1996. The interplay of syntactic and semantic node labels in partial parsing. In Bunt, H., and Tomita, M., eds., *Recent Advances in Parsing Technology*. Kluwer Academic Publishers. 295323.

Office of the Director of National Intelligence. 2007a. Intelligence community directive 203: Analytic standards.

Office of the Director of National Intelligence. 2007b. Intelligence community directive 206: Sourcing requirements for disseminated analytic products.

Office of the Director of National Intelligence. 2009. Intelligence community directive 501: Discovery and dissemination or retrieval of information within the intelligence community.

Pirolli, P., and Card, S. 2005. Sensemaking processes of intelligence analysts and possible leverage points as identified through cognitive task analysis. In *Proceedings of the 2005 International Conference on Intelligence Analysis*.

Stasko, J.; Gorg, C.; and Liu, Z. 2008. Jigsaw: supporting investigative analysis through interactive visualization. *Information Visualization* 7:118–132.

Ziegler, J. E., and Fahnrich, K. P. 1988. Direct manipulation. In Helander, M., ed., *Handbook of Human-Computer Interaction*. North-Holland, Amsterdam. 123–133.