



**US Army Corps
of Engineers®**
Engineer Research and
Development Center



Automated Man-Machine Learning and Processes For 3D Terrain (AMP3D)

Foundations of Mission Analysis Storytelling (FOMAS)

Timothy K. Perkins, Harland Yu, and Sean Sovine

July 2020

The U.S. Army Engineer Research and Development Center (ERDC) solves the nation's toughest engineering and environmental challenges. ERDC develops innovative solutions in civil and military engineering, geospatial sciences, water resources, and environmental sciences for the Army, the Department of Defense, civilian agencies, and our nation's public good. Find out more at www.erdcl.usace.army.mil.

To search for other technical reports published by ERDC, visit the ERDC online library at <http://acwc.sdp.sirsi.net/client/default>.

Foundations of Mission Analysis Storytelling (FOMAS)

Timothy K. Perkins

*U.S. Army Engineer Research and Development Center (ERDC)
Construction Engineering Research Laboratory (CERL)
2902 Newmark Drive
Champaign, IL 61824*

Harland Yu and Sean Sovine

*U.S. Army Engineer Research and Development Center (ERDC)
Geospatial Research Laboratory (GRL)
7701 Telegraph Road
Alexandria, VA 22315*

Final Report

Approved for public release; distribution is unlimited.

Prepared for U.S. Engineer Research and Development Center
Geospatial Research Laboratory
7701 Telegraph Road
Alexandria, VA 22315

Under Program Element Number – 0602784A
Project Number – 855
Task Number - 24

Abstract

Mission analysis is a critical step in military planning and decision-making. It is currently time-consuming for analysts, who have few automated tools. The Foundations of Mission Analysis Storytelling (FOMAS) project developed algorithms, tools, and methods to automate sensemaking for mission analysis, which reduces the time and increases the effectiveness of the process.

This report describes the FOMAS research, specifically as it relates to storytelling and link analysis. It includes descriptions of storytelling and a related prototype implementation, “Spatio-temporal Retrieval and Introspection of Data and Embedded Relationships, (STRIDER).” It also describes user engagements involving STRIDER and a prototype information collection and processing tool, the Big Open Source Social Science (BOSSS).

The focus for this project was placed on capability development and adaptation. Both STRIDER and BOSSS underwent limited testing that demonstrated promise and received positive user feedback. However, project time did not allow for detailed rigorous technical evaluations of the prototypes. The team recommends further investigation and development of these automation tools to improve mission analysis for analysts.

DISCLAIMER: The contents of this report are not to be used for advertising, publication, or promotional purposes. Citation of trade names does not constitute an official endorsement or approval of the use of such commercial products. All product names and trademarks cited are the property of their respective owners. The findings of this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

DESTROY THIS REPORT WHEN NO LONGER NEEDED. DO NOT RETURN IT TO THE ORIGINATOR.

Contents

Abstract.....	ii
Figures and Tables.....	v
Preface.....	vi
1 Introduction.....	1
1.1 Background.....	1
1.2 Objectives.....	2
1.3 Scope.....	3
1.4 Approach	3
1.5 Technology transfer	3
2 Mission Analysis.....	5
2.1 Background.....	5
2.2 Army doctrine	5
2.3 Improving mission analysis	7
2.3.1 Initial intelligence preparation of the battlefield (IPB).....	8
2.3.2 Other processes of mission analysis.....	9
2.3.3 Mission analysis technology.....	10
3 Technology: Storytelling and Link Analysis	11
3.1 Introduction.....	11
3.2 Related work	12
3.3 Storytelling approach.....	13
3.3.1 General approach.....	13
3.3.2 Detailed approach.....	17
3.3.3 Prototype implementation.....	22
3.3.4 Evaluation	25
3.4 Link analysis.....	27
3.4.1 Approach.....	27
3.4.2 Evaluation	28
3.5 Storytelling prototype conclusions.....	30
4 Technology: Big Open Source Social Science.....	32
4.1 BOSSS user engagements	32
4.1.1 Enterprise Challenge 19	32
4.2 BOSSS products.....	33
4.2.1 Civil Affairs Proponent Office	34
5 Summary and Recommendations	35
References.....	36
Acronyms	40

Report Documentation Page

Figures

Figure 1. FM 6-0 - Mission Analysis - Figure 9-3, page 9-7.....	6
Figure 2. Architecture diagram.....	22
Figure 3. Data flow overview.	23
Figure 4. Data flow diagram, storytelling.	23
Figure 5. Example of a story.	24
Figure 6. Summary statistics for SYNCOIN dataset.	24
Figure 7. Story summary.	25
Figure 8. Summary statistics for BFB dataset.	26
Figure 9. Summary statistics for Atlantic Storm dataset.	26
Figure 10. SRW link prediction results.	29
Figure 11. Abridged SRW predictions (al-Hardi).	30

Preface

This study was conducted for the Engineer Research and Development Center - Geospatial Research Laboratory (ERDC-GRL) under Automated Man-Machine Learning and Processes For 3D Terrain (AMP3D). The funding was provided by Program Element Number 0602784A, Project Number 855, and Task Number 24. The technical monitor was Mr. Nathan Frantz.

The work was performed by the Information Generation and Management Branch, ERDC-GRL. At the time of publication, Mr. Jeffrey Murphy was Branch Chief; Ms. Martha Kiene was the Division Chief; and Mr. Ritchie Rodebaugh was the Technical Director for Geospatial Research and Engineering. The Deputy Director of ERDC-GRL was Ms. Valerie Carney and the Director was Dr. Gary Blohm.

COL Teresa A. Schlosser was Commander of ERDC, and Dr. David W. Pittman was the Director.

1 Introduction

1.1 Background

Mission analysis, the second step of the Military Decision Making Process (MDMP), is central to Army planning and critical to effective operations (Army Doctrine Publication (ADP) 5-0, May 2012). Mission analysis occurs after the Receipt of Mission with guidance from higher headquarters and other inputs. It interacts with the Army Design Process (ADP) and yields critical outputs such as problem and mission statements, initial Commander's intent, planning guidance, and information requirements.

Planners conduct the initial Intelligence Preparation of the Battlefield (IPB) during mission analysis, along with an assessment of the higher headquarters plan or order; determination of specified, implied and essential tasks; and other tasks (Field Manual (FM) 6-0, Figure 9-3, May 2014).

“...It is imperative the commander and staff 'get it right' concerning mission analysis. The ramifications of a failure of understanding in this step of the MDMP will be magnified as the commander and staff continue through the decision making process” (Loffert, 2002).

Improved understanding requires better sensemaking. Sensemaking (Weick 1995; Pirolli 2005; Klein et al. 2006a; Klein et al. 2006b; Klein et al. 2007; Weick et al. 2005; Weick 2009; Weick 2012; Sieck et al. 2007; Moore 2013; Egeth et al. 2014) is an iterative process that describes how people, such as analysts or first responders, actively work to understand a situation (Tolone et al. 2015) or how people fit data into a frame in order to interpret context (Klein et al. 2007; Whalley et al. 2014). Requirements relevant to sensemaking were expressed in the 2018 Army Center of Excellence Science & Technology needs statements from the Combined Arms Center - Training (Synthetic Training Environment (STE) Cross Functional Team (CFT)), Combined Arms Support Command (CASCOM), U.S. Army Intelligence Center of Excellence, and Maneuver Center of Excellence, and Soldier Lethality CFT.

Recent concepts such as Multi-Domain Operations (U.S. Army 2018) and the Joint Concept for Operating in the Information Environment (Joint Chiefs 2018) express the critical importance of understanding the interaction between the physical and information environments. To address Mission Analysis requirements at the intersection of these environments, technical capabilities were developed to provide automated support to the process.

A key capability of the Foundations of Mission Analysis Storytelling (FOMAS) that directly relates to sensemaking is "storytelling, i.e., automatically 'connecting the dots' between disparate entities, such as people and organizations, in an effort to form hypotheses and suggest non-obvious relationships" (Hossain et al. 2012). Storytelling is relevant to intelligence analysts (Wang et al., 2012; Hossain et al. 2012; Hossain et al. 2011; Dos Santos et al. 2016a; Dos Santos et al. 2016b), medical and scholarly research (Kumar et al. 2008; Hossain et al. 2012), media and event analysis (Barranco et al. 2017; Dos Santos et al. 2014; Dos Santos et al. 2016b; Kader et al. 2016), as well as broader applications.

Development of the Big Open Source Social Science (BOSSS) prototype was also guided during this research. BOSSS was designed to support IPB, incorporating publicly available information and user provided text data. The tools help the analyst understand groups, relationships, key influencers, the cognitive environment, and the information environment. The Boston Fusion Corporation developed BOSSS in response to the broad agency announcement topic, "A16A-To13: Beliefs, Values, Interests and Practices of Identities, Networks, Groups for Planning & Analysis" (Perkins 2016) of the Small Business Technology Transfer (STTR) program.

This report summarizes the accomplishments during the Fiscal Year (FY) 2019 FOMAS research.

1.2 Objectives

The objective of FOMAS was to experiment with, and improve, algorithms, tools, and methods to enable Army personnel to iterate more rapidly and comprehensively through stages of sensemaking and storytelling. The research contributes to mission analysis by minimizing the time and effort spent on basic and easily automated activities.

1.3 Scope

FOMAS was part of the Automated Man-Machine Learning and Processes for 3D Terrain (AMP3D) program.

FOMAS included the four tasks listed below.

- Task 1) Analyzing gap statements, concepts, and doctrine to identify mission analysis automation requirements.
- Task 2) Investigating link prediction approaches for enhanced storytelling.
- Task 3) Overseeing BOSSS development.
- Task 4) Exercising BOSSS and Spatio-temporal Retrieval and Introspection of Data and Embedded Relationships (STRIDER) prototypes to support to mission analysis.

1.4 Approach

The tasks were executed using a combination of waterfall and agile project management practices. The scope and priorities were established at the start of the project, along with the activities associated with each task. As research and stakeholder outreach proceeded, additional activities and revised priorities were defined based on findings and stakeholder input.

Task 1 involved searching Army and Joint libraries, doctrine, task lists, and training documents for specific Army tasks associated with mission analysis. After collecting essential information, workflows relevant to the project technologies, Storytelling and BOSSS, were identified.

Task 2 implemented the link prediction algorithms within the STRIDER prototype and compared link prediction results with storytelling results.

Tasks 3 and 4 was executing the Contracting Officer Representative (COR) responsibilities, which included verifying technical performance, maintaining communication and coordination with Boston Fusion, and inspecting and testing deliverables and reports.

1.5 Technology transfer

As part of FOMAS task 4, the BOSSS and STRIDER prototypes were demonstrated at Warfighter Exercise 19-04. BOSSS was demonstrated at

Enterprise Challenge 19. These events provided exposure to potential end-users and opportunities to collect feedback on usability, current and potential workflows, as well as user priorities. This report includes select findings from these events, with additional details from Enterprise Challenge 19 available in the final BOSSS Phase II Option report.

Stakeholders were also engaged to identify opportunities for technology integration. As part of this interaction, the team established a Knowledge Transition Agreement (KTA) with the Civil Affairs Proponent Office. The final BOSSS report will serve as one of the FY19 KTA deliverables.

The BOSSS and STRIDER source code libraries have been retained, using the Defense Intelligence Information Enterprise (DI2E) platform to enable sharing and exchange of code to authorized, appropriate partners. Functional prototypes of BOSSS and STRIDER were also established on the Army Geospatial Enterprise (AGE) Node – a computing capability on the Defense Research and Engineering Network (DREN) that provides an ability to mimic Army computing environments, such as the Command Post Computing Environment (CPCE).

2 Mission Analysis

2.1 Background

For Task 1, gap statements, concepts, and doctrine were analyzed by searching Army and Joint libraries. Doctrine, task lists, and training documents were examined to identify specific Army tasks associated with Mission Analysis. This involved assessing the inputs, outputs, and personnel. The analysis was used to identify workflows for automation by Storytelling and BOSSS technologies.

2.2 Army doctrine

Mission Analysis is introduced in several doctrinal publications, including ADP 3-0 (2019) – Operations and ADP 5-0 (2019) – The Operations Process. FM 6-0 Commander and Staff Organization and Operations (May 2014) contains additional detail.

ADP 3-0 provides a foundational description of Operational Art and Operational Processes in the land operations context. It differentiates design from planning and introduces operational environment (OE) variables (i.e., Political, Military, Economic, Social, Information, Infrastructure, Physical Environment, and Time - 'PMESII-PT'); mission variables (Mission, Enemy, Terrain and Weather, Troops and Support Available, Time Available, and Civil Considerations – 'METT-TC'); and domains (Land, Maritime, Air, Space, and Cyberspace).

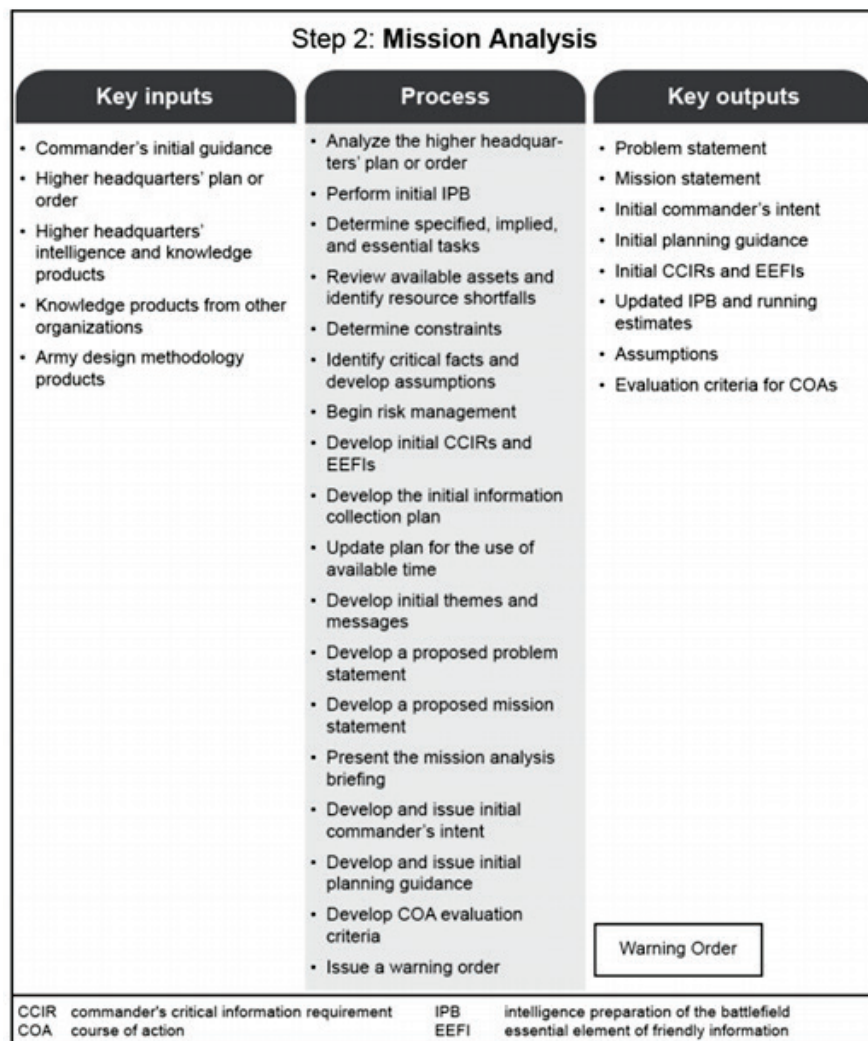
ADP 5-0 differentiates and characterizes three planning processes used by the Army – Army Design Methodology (ADM), MDMP, and troop leading procedures (TLP). This publication also describes the interactive iteration during planning between operational art and planning.

The Army Doctrine Reference Policy (ADRP) 5-0 – The Operations Process (2012) and ATP 5-0.1 - Army Design Methodology (2015) describe how, “Effective planning requires integration of the conceptual and detailed components... Army leaders employ and integrate the following planning methodologies throughout an operation – ADM, the Military Decision Making Process (MDMP), Troop Leading Procedures.” ATP 5-0.1 describes several techniques to enable interaction between ADM and MDMP, including having teams each work concurrently (section 2-4), or

embedding aspects of ADM in the Mission Analysis step of MDMP (section 2-5). This interactive work between ADM and MDMP supports situational understanding and framing operational environments (Chapter 3).

FM 6-0 (2014) states, “Since no amount of subsequent planning can solve an insufficiently understood problem, Mission analysis is the most important step in the MDMP” (9-26). As such, it is critical that time and effort spent on mission analysis be used effectively. FM 6-0, Figure 9-3 (replicated below as Figure 1), summarizes the inputs, processes and outputs of Mission analysis.

Figure 1. FM 6-0 - Mission Analysis - Figure 9-3, page 9-7.



2.3 Improving mission analysis

Mission analysis includes both top-down and bottom-up processes. From a top-down perspective, it requires receiving and analyzing higher headquarters' plans, guidance, intelligence, external knowledge products, specified tasks, critical facts, and determining constraints. From a bottom-up perspective, it requires performing an initial IPB, developing the commander's critical information requirements (CCIR), and initial commander's intent.

Previous articles and books have examined the process of sensemaking (Pirolli 2005; Klein et al. 2006a; Klein et al. 2006b; Klein et al. 2007; Weick et al. 2005; Weick 2009; Weick 2012; Sieck et al. 2007; Moore 2013). Some (Pirolli, Moore) focus on the sensemaking work of analysts and decision-makers. Similar to the Army's design methodology and mission analysis, sensemaking is an iterative top-down and bottom-up process with feedback loops, as illustrated by Pirolli and Card (2005). It includes foraging, encoding, and reasoning.

Pirolli and Card (2005) describe a parallel between foraging results and precision and recall in information retrieval. In information retrieval, the analyst makes the tradeoff between obtaining more information that includes less relevant or irrelevant results versus receiving less information, but highly relevant information that may miss other relevant information. They also describe techniques for highlighting potentially relevant information in advance, or summarizing documents as possibilities for improving costs imposed in lower levels of sensemaking. The key to efficiency is the speed an analyst can reject unimportant messages. Challenges for higher-level sensemaking processes include problem structuring (generating, exploring, and managing hypotheses) and evidentiary reasoning.

- *Some of the challenges identified for efficient, useful analysis by Pirolli and Card (2005) include:*
 - *Highlighting potentially relevant information*
 - *Summarizing documents*
 - *Reject unimportant information*
 - *Generate, explore, manage hypotheses*
 - *Challenge biases (esp. confirmation bias)*

2.3.1 Initial intelligence preparation of the battlefield (IPB)

The team worked with the STRIDER storytelling software and the BOSSS prototype to characterize how these tools would mitigate the Mission Analysis challenges, especially with respect to initial IPB. In summary:

- *Highlighting potentially relevant information*
 - Both BOSSS and STRIDER utilize named entity recognition (NER) techniques to automatically identify and highlight persons, organizations, and locations within blocks of text, presenting highlighted elements to the user.
 - BOSSS automates keyword search of datasets (e.g., Google News) to aid analysts in finding relevant documents about a location of interest.
 - BOSSS uses vector space models to organize the collected and extracted information into doctrinal formats, such as the PMESII-PT/ASCOPE crosswalk (ATP 2-01.3, 2019, pg 4-23).
- *Summarizing documents*
 - Both BOSSS and STRIDER use natural language processing techniques to extract entities referenced in a common context and then automatically build a network graph to depict these relationships.
 - * STRIDER uses storytelling techniques to build the graph and present found and potential relationships between entities. The user has options to adjust the thresholds when building actual relationships and additional flexibility when building potential relationships.
 - * BOSSS uses natural language processing and builds connections based on document co-occurrence alone.
 - * BOSSS also infers and includes sentiment edges between entities when quoted material of one entity references another entity.
- *Reject unimportant information*
 - Both BOSSS and STRIDER allow users to correct network relationships when an analyst discovers flaws in the relationships.
 - With BOSSS, analysts have the option to prepare PMESII-PT/ASCOPE crosswalk output charts using the automatically

extracted information. The analyst can adjust the crosswalk to include or exclude information.

- *Generate, explore, manage hypotheses*
 - STRIDER explicitly allows analysts to generate stories that are hypothetical relationships between entities. The relationships in these stories are supported by evidence in the documents. There is a configurable threshold of uncertainty for these relationships. The analyst may retain or reject the relationship diagrams and may export them to other tools.
 - The automated processing of BOSSS currently allows users to assess multiple information sources in the same context, such as reports about economic capabilities (the 'E' in PMESII-PT, and 'C' in ASCOPE). Users can also compare the reports.
 - Additionally, both BOSSS and STRIDER allow users to find documents that reference the same entity (e.g., named individual, organization, etc.) and review those documents.
- *Challenge biases (esp. confirmation bias)*
 - Using the STRIDER storytelling techniques, analysts can look for alternative explanations or alternative approaches.

2.3.2 Other processes of mission analysis

Other Mission Analysis processes relevant to the BOSSS and STRIDER technologies may include:

- *Analyze the higher headquarters' plan or order.*
- *Determine specified, implied, and essential tasks.*
- *Review available assets and identify resource shortfalls.*
- *Determine constraints.*
- *Identify critical facts and develop assumptions.*
- *Develop initial CCIRs and Essential Element of Friendly Information (EEFIs).*
- *Develop the initial information collection plan.*
- *Develop initial themes and messages.*
- *Iteration with Army Design methodology.*

FOMAS did not prototype these processes, as they were beyond the scope.

2.3.3 Mission analysis technology

Detailed information on storytelling techniques and STRIDER are detailed in Chapter 3 of this report. Full details about BOSSS are available in a separate BOSSS technical report and briefing, each available as Distribution B materials through the Defense Technical Information Center (DTIC) library (Palladino et al. 2019a and Palladino et al. 2019b).

3 Technology: Storytelling and Link Analysis

3.1 Introduction

The following scenario motivates this work

Suppose we are given a large collection of natural language text documents in digital format with which we have little familiarity. Further, suppose we are interested in two entities and that at least one of the documents mentions these entities. We wish to find out whether these two entities are connected, and, if so, we wish to discover the nature of the connection. Such a task has significant practical relevance for intelligence analysts, as well as IPB and mission analysis as described in section 2.3.

The most direct approach to answering these questions is to read the documents in the collection and search for evidence of a connection between the two entities of interest. The problem with this approach is the document collection might be large, so it could take a substantial amount of time and effort to obtain a sufficient grasp of its contents. In operational conditions, time and resource restrictions make such an approach impractical, costly or altogether infeasible. Yet, in some cases, this document collection may be the sole source of information about the target entities, and collecting further evidence is too expensive to be feasible. Storytelling and link analysis can improve efficiency in working through such a document collection.

Key processes of Mission Analysis include the IPB and development of the initial CCIRs. Outputs from Mission Analysis include updates to the IPB and running estimates. Link analysis (also known as ‘network analysis’) is an analytic method that is part of IPB and is used to understand organizational diagrams, relationships between “personalities,” and their associations to various groups or activities. The method is useful for analyzing events, understanding population characteristics, assessing Improvised Explosive Device (IED) activities, and evaluating some economic systems.

An automated capability was developed to aid the discovery of connections between entities of interest based on evidence contained in collections of digital, natural language documents. The team believes this capability can

discover connections quickly and with little effort. It is also believed that, in some cases, it can discover connections that might be missed in a manual examination of the documents.

3.2 Related work

This work builds directly on the work presented in Hossain et al. (2012). The authors automatically extracted and connected entities mentioned in a set of text documents to form an entity network. They applied a weighting scheme to the edges in the network and used a shortest-path algorithm to find an optimal path between two entities.

This framework was built upon by enforcing explicit constraints on connections between entities, using additional entity features to define edge weights, and introducing an extension to the shortest-path algorithm used for finding paths connecting pairs of entities. Hossain et al. (2012) also introduced an approximation of the A* algorithm that was designed to enhance efficiency on extremely large data sets. For this research, existing pathfinding algorithms were used in unmodified forms.

Several other works describe similar approaches to connecting entities in automatically generated networks, a task referred to as "connecting the dots." Jin et al. (2007a) and Jin et al. (2007b) considered shortest paths between concepts, where concepts include both named entities and more general noun-phrases that were automatically extracted from source documents. A1-Zaidy et al. (2012) considered chains of co-occurring terms connecting prominent communities, which are groups of one or more associated entities. Faloutsos et al. (2004) describe a method for efficiently generating connection subgraphs, which is a subgraph of an entity network that best describes the relationship between two nodes in that network.

Storytelling and link prediction assume a direct connection between two indirectly connected entities based on the statistical properties of reaching one entity from another via a random walk on an entity graph. However, link prediction does not determine the most likely path within the entity graph that can substantiate the inferred connection. Consequently, these approaches can be mutually reinforcing or complementary within a greater analysis task.

3.3 Storytelling approach

Following the approach of Hossain et al. (2012), it is proposed to find a potential connection between two entities of interest by finding an optimal path in an extracted entity network. To start, automatically generate a set of hypothesized connections between entities tagged in the input document corpus; then apply weights to those connections using several evidence features calculated from the corpus text; finally, find an optimal path in this weighted, automatically generated network.

3.3.1 General approach

3.3.1.1 *Connecting entities*

Modeling potential connections between entities using paths in a generated entity network is proposed, thus the direct connections that comprise that network are fundamental to this approach. If the direct connections that make up a path are not valid and strong, it is unlikely that the connection hypothesized by that path will be valid. Thus, a reliable method is needed for inferring direct connections between entities.

Since this capability will automatically generate inferences from potentially very large data sources, an efficient algorithm is also needed. The common sentence criterion was used to connect entities in the extracted network. This criterion says that two entities will connect in the extracted network if, and only if, they occur together in at least one sentence in a source document. This team hypothesizes that when a common sentence includes two entities, there is a substantial likelihood that these two entities are somehow connected.

Notice that a "connection" between two entities is not defined in a real-world sense. The team simply defines a hypothetical connection to be an edge in a generated entity network. It is left to the user to define what a real-world connection between two entities means. However, the approach was designed with the intent of capturing a useful notion of connection. The hope is this capability will help users to identify pairs of entities connected in a real-world sense and aid in investigating the nature of such connections.

3.3.1.2 *Applying weights to entity connections*

It is believed that certain aspects of the corpus text can provide additional evidence for the existence of connections between entities beyond the evidence captured by the common sentence criterion. Several characterizations of entities were developed using information from the source document text. These characterizations were used to define edge weight functions for the connections established using the common sentence criterion.

The edge weight functions are intuitively measures of distance (not necessarily metric) between the characterizations of entities. The weight of an edge between two entities with similar characterizations is small and the weight of an edge between entities with dissimilar characterizations is large. The value of this edge weighting scheme for applying edge weights will be seen when the optimal path connecting two entities is discussed. However, the intuition behind this approach is that an edge between two entities that has a low weight (i.e., short distance) is one that represents an inferred connection with stronger evidence.

The three sources of information that are used for characterizing an entity are: the documents containing the entity, the distribution of words occurring near mentions of the entity, and the connections of an entity to other entities based on the common sentence criterion. These entities are referred to as document characterization, word-context characterization, and network characterization.

The primary hypothesis motivating the document characterization is that entities, which appear in similar sets of documents, are more likely related than entities that appear in dissimilar sets of documents. A related hypothesis is that if there is evidence supporting a possible direct connection between two entities, it is most likely in a document that contains both entities.

The following hypotheses motivate the word-context characterization: entities that tend to occur with the same words are likely discussed in similar contexts; and, similarity between the contexts of two entities gives supporting evidence for a possible connection between them. The inspiration for this characterization comes from language models, such as Hyperspace Analogue to Language (HAL) as outlined by Lund and Burgess (1996). This characterizes a word in terms of the frequencies with which

other words occur near that word. Given a large amount of input data, the team believes the word-context characterization could be used to meaningfully categorize and compare entities. It is hypothesized that this information will discriminate between pairs of entities that are likely to be connected and pairs that are not.

Finally, entities are characterized in terms of their connections to other entities. The motivation for this characterization is the hypothesis that two entities connected to similar sets of other entities are more likely connected to one another.

3.3.1.3 Finding optimal paths

Once a network of inferred connections between entities and applied weights to those connections have been generated, the network is used to discover evidence for connections between entities. This evidence is looked for in the paths through the generated network, from one entity of interest to another. The interest here is only in paths that contain no entity more than once, i.e., paths that contain no loops. The simplest example is a path with a single edge connecting two entities that share a common sentence. However, a more complex case is a path with more than one edge, which corresponds to an indirect connection between entities formed by a chain of direct connections.

Finding the path that is most likely to represent evidence for a real-world connection is the goal. This means that some way to estimate the relative levels of evidence supporting each path was needed. In order to map evidence onto paths in this network, the simplifying assumption was made that the strength of evidence supporting a path is a simple function of the strength evidence supporting each of the edges that makes up the path.

In the most basic model of path strength, it was assumed that the length of a path can quantify the strength of evidence for the path connection. The length of a path is the sum of the weights of edges along the path. This model is referred to as the shortest path model. Recall that, in the extracted network, edges are weighted using a weight function based on descriptions of entities in terms of a number of features.

Because larger weights represent lower degrees of evidence, the edge weight functions were referred to informally as distances. Thus, the path length was interpreted informally as the total distance along the path.

Alternatively, the edge weight could be interpreted as a measure of uncertainty. The shortest path model assumes that uncertainty accumulates additively along the path as we move from one node to the next along the edges of the path.

Notice that the shortest path model involves a certain independence assumption: it is assumed that the contribution of one edge to the overall quality of the path is independent of the other edges in the path. The advantage of this method is that shortest path finding is a well-studied problem, and there are efficient shortest-path algorithms available.

A variation in the shortest path model was also considered. Intuition here suggests that a path is only as good as its weakest link. A minimax path is a path whose weakest link is as strong as possible. In the shortest-minimax path model, we first discover the set *S* of paths connecting the two entities of interest such that the longest edge of each of these paths is shorter than the longest edge of any other path not in the set *S*. Then, among these "minimax" paths, the shortest path was found and referred to as the shortest-minimax path. Among all of those paths that share the same weakest link, the best path is the one that is the shortest.

Both the shortest path and shortest-minimax path approaches were implemented in a prototype system (see section 3.3.3).

3.3.1.4 *Clique constraint*

One further restriction was considered on connections between entities beyond the common sentence criterion, which was the common clique constraint. It states that, in order for an inferred connection to exist between two entities in the extracted network, those two entities must share a common *k*-clique. A *k*-clique is a fully connected subgraph containing *k* vertices. In this context, that means a set of *k* entities in which each entity shares a common sentence with each of the *k* - 1 others. The hypothesis motivating the common clique constraint is that entities that are part of a highly connected group, based on inferred connections, are likely connected to one another in a real-world sense.

3.3.2 Detailed approach

This section will describe more of the specific algorithmic and implementation details of the approach to finding connections between entities.

3.3.2.1 *Extracting entities and forming connections*

The vertices in the entity network were automatically extracted from source text documents using the NER tool from Stanford's CoreNLP suite as described by Manning et al. (2014). Stanford's coreference tool was also used to discover coreference chains. A coreference chain is a set of n-grams in the text of a document, such that each of these n-grams is a mention of some single entity. An n-gram is just a sequence of n contiguous words.

Within a document, the longest n-gram in a coreference chain that is tagged as a named entity was chosen to represent the entity referred to within that chain. This means that all n-grams tagged as mentions within the same chain will be treated as instances of that longest named entity n-gram. After consolidating mentions within documents, entities were merged across documents by considering any two entities with the same name to be instances of the same entity.

Connections between entities are initially formed using the common sentence criterion. After applying entity tagging and intra- and inter-document coreference processing to the text, a straightforward and reasonably efficient procedure was applied to iterate over sentences and determine which pairs of entities occur together within a common sentence. Sentence boundaries were obtained from the Stanford CoreNLP output.

3.3.2.2 *Calculating edge weights*

As described above, the weights of edges in the extracted entity network were determined by using a combination of three distance functions where each was defined using a different characterization of entities. These characterizations are based on, respectively, occurrence of entities within documents, distributions of words occurring near entity mentions, and patterns of connectivity to other entities. The respective distance functions for these entities are referred to as document distance, context distance,

and network distance, abbreviated as D_{doc} , D_{cont} , and D_{net} . The final edge weight function is a simple linear combination:

$$D(e_i, e_j) = \alpha_d * D_{doc}(e_i, e_j) + \alpha_c * D_{cont}(e_i, e_j) + \alpha_n * D_{net}(e_i, e_j)$$

A further restriction was placed on the form of the distance functions by requiring that each produced values in the interval $[0,1]$. This provided some normalization when combining the function and facilitates the use of an upper-bound threshold for restricting entity connections. It has the following interpretation: if a distance function is zero, then the corresponding entity descriptions are identical; if the function is one, then the corresponding descriptions are completely different. Here, "completely different" means simply the greatest difference recognized by the particular distance function. The coefficients α_d , α_c , and α_n are required to be non-negative and sum to one. The result is that the function D is essentially a composite distance function that also takes values in the interval $[0,1]$.

3.3.2.2.1 Document characterization

The document characterization for an entity is the set of documents containing the entity. This set was denoted by S_q for entity q . Given two entities, e_i and e_j , the document distance between these is calculated as one minus the Jaccard index:

$$D_{doc}(e_i, e_j) = 1 - \frac{|S_i \cap S_j|}{|S_i \cup S_j|}$$

3.3.2.2.2 Context characterization

The "context" characterization of entities is calculated using a sliding window algorithm. A sliding window is a contiguous sequence of words of maximum length $2n + 1$, centered at some entity e in the source text. Here, a word is defined as a white-space delimited token. The sliding window contains the entity e plus the (at most) n words preceding e and the (at most) n words following e in a source document.

The center of the window is moved from the first entity to the last entity in each document — hence the name "sliding window." During this process, a window-count vector c_i is maintained for each entity e_i . Whenever the window is centered at an entity, e_i , for each word, w , in the window

(excluding e_i), the component in c_i corresponding to w is incremented by one. Once this process is complete, we have a complete set of window-count vectors c_i , one for each entity e_i .

Once the window-count vectors c_i have been obtained, two additional steps are performed to obtain the "context vector," which characterizes entities in terms of word context. First, a weight is applied to each word w in window-count vector c_i . This weight is a variation on the inverse document frequency (IDF) weight that is commonly used in information retrieval. Thus, the inverse context frequency (ICF) for the weight for word w is calculated as

$$ICF(w) = \log \frac{|E|}{|\{e_i \in E: c_i(w) > 0\}|}$$

where $c_i(w)$ is a component of the window-count vector for entity e_i corresponding to word w and E is the set of entities contained within the document set. These weights are applied to context vectors by setting $c'_i(w) = ICF(w) * c_i(w)$, for each entity e_i and each word w in c_i . Finally, each vector is normalized using the 1-norm by setting

$$k_i(w) = \frac{c'_i(w)}{\sum_{w \in V} c'_i(w)}$$

for each entity e_i and each component of c_i . V is used to denote the set of all unique words in the document set (e.g. the vocabulary). The result of this normalization is a vector whose components are non-negative numbers that sum to one. These vectors are referred to as context vectors.

The context distance between the characterizations of two entities e_i, e_j is now defined as a multiple of the Euclidean distance between the context vectors k_i, k_j

$$D_{cont}(e_i, e_j) = \left[\frac{1}{2} \sum_{w \in V} (k_i(w) - k_j(w))^2 \right]^{\frac{1}{2}}$$

where the factor $\frac{1}{2}$ is added to ensure that the function D_{cont} has range $[0,1]$.

3.3.2.2.3 Network characterization

The network characterization of an entity is the row corresponding to that entity in the adjacency matrix of the extracted network generated by the common sentence criterion. The adjacency matrix for a network with vertices $\{e_1, e_2, \dots, e_n\}$ is an $n \times n$ binary matrix, A , whose entry a_{ij} is equal to 1 if e_i and e_j are connected and equal to 0 otherwise. Thus, the i th row of A , denoted a_i , is a binary vector of length n whose components describe the connectivity of e_i to other entities in the network. The network distance between two entities e_i, e_j is defined as one minus the cosine similarity of their corresponding rows in A : $a_i \cdot a_j$ where the dot operation above represents the usual vector dot product.

3.3.2.3 *Clique constraint*

Recall the hypothesis in section 3.3.1.4 that states k -cliques provide ‘evidence’ that the clique members are most likely strongly related in the real world. A greedy algorithm was used to find an approximately optimal clique for each pair of entities e_i, e_j , whenever a clique exists that contains those two entities. This algorithm was essentially a depth-first search algorithm, which proceeds as follows. A clique is constructed stepwise, initially containing only e_i, e_j .

At each step, an optimal candidate entity was added to the clique. A candidate entity is an entity that is not already in the clique and connects to each entity already in the clique. An optimal candidate entity is one whose addition to the clique increases the total weight of edges in the clique no more than the addition of any other entity. If the last candidate added to the clique results in a clique of size k , the procedure halts and the clique is returned. However, it may be that, at some point, there are less than k entities in the clique and there are no candidate entities. In this case, the last entity added to the clique would be removed and the next-best candidate added. The "backtracking" continues until only two entities e_i, e_j are in the clique and there are no more candidates. The procedure then halts without returning a clique.

3.3.2.4 *Finding optimal paths*

Two approaches have been discussed to defining an optimal path in the extracted network: shortest path and shortest-minimax path. To find shortest paths, the Dijkstra's algorithm is used. To implement shortest-

minimax paths, the approach in Malpani and Chen (2002) was followed. In this approach, Kruskal's algorithm was first used to generate a minimum spanning tree (MST) for the extracted network. Because the network is undirected, the unique path connecting any two nodes, e_i, e_j in this MST, is guaranteed to be a minimax path*. Next, to determine whether any path is a minimax path, it is only necessary to ensure that the weakest link along that path is at least as strong as the weakest link of the minimax path contained in the MST (in fact, it can't be any stronger). Hence, all edges could be removed in the network with weight greater than the weight of the weakest link in the minimax path in the MST, and this ensures that all paths connecting e_i and e_j in the resulting network are minimax paths. To find the shortest-minimax path, one needs only to run a shortest-path algorithm on the resulting, pruned network.

As an alternative to Dijkstra's algorithm, it is possible to use the A* algorithm to find shortest paths in the extracted network. A* is an extension of Dijkstra's algorithm that uses a heuristic to guide the order in which nodes are visited, often resulting in improved computational performance (Hart et al. 1968; Hetland 2010)[†]. In graphs embedded in metric spaces, it is possible to employ the straight-line distance between vertices as a consistent A* heuristic. The document distance and context distance functions described above are both metrics on the set of entities, and the function

$$D_{net}(e_i, e_j) = \arccos(1 - D_{net}(e_i, e_j))$$

provides an alternative network distance function that is also a metric (see Leskovec et al. 2014)[‡]. A non-zero linear combination of distance metrics is also a distance metric. Hence, with the function D'_{net} in place of D_{net} , the distance function D becomes a metric, meaning that the vertices of this network can be viewed as points in a metric space. Thus, if this modified version of D is used for defining edge weights, then D could also be employed as a straight-line distance heuristic for the A* algorithm. This approach entails a trade-off between the efficiency of A* and the cost of calculating distance functions that require calculating arccosine values.

* For a proof, see supplemental document "MST Minimax Path Property"

† It is straightforward to show that, if the A* heuristic is the constant 0 function, then A* is equivalent to Dijkstra's algorithm.

‡ For a reference regarding the fact that "one minus Jaccard index" and "cosine distance" are metrics, see Leskovec, Rajaraman, and Ullman, Mining of Massive Datasets, chapter 3.

3.3.3 Prototype implementation

The team developed a prototype capability implementing the storytelling algorithm as part of GRL's STRIDER applied research project to investigate the applicability of storytelling on military decision-making and intelligence analysis. Aspects of the project continued under the Geo-enabled Augmented Intelligence for Decisive Engagement (GAIDEN) and FOMAS projects. For FOMAS, the storytelling algorithm was compared against the output of several link prediction algorithms to understand the relationship between the two approaches. The prototype used the results of a previous project, Grapevine, as a foundation for further development. Grapevine integrated information retrieval and spatial indexing to drive spatial pattern analysis and derive insight from unstructured and structured sources of data. Figures 2 and 3 below provide an overview of the prototype architecture and data flow.

Figure 2. Architecture diagram.

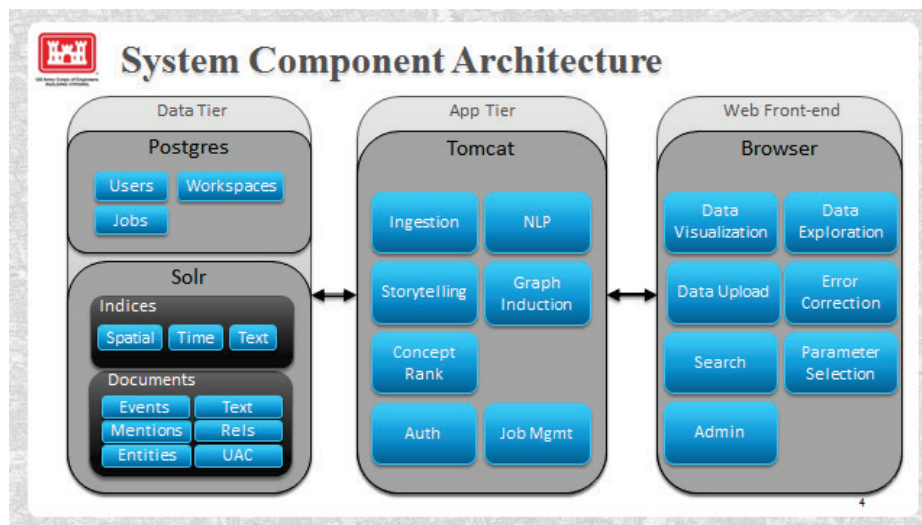
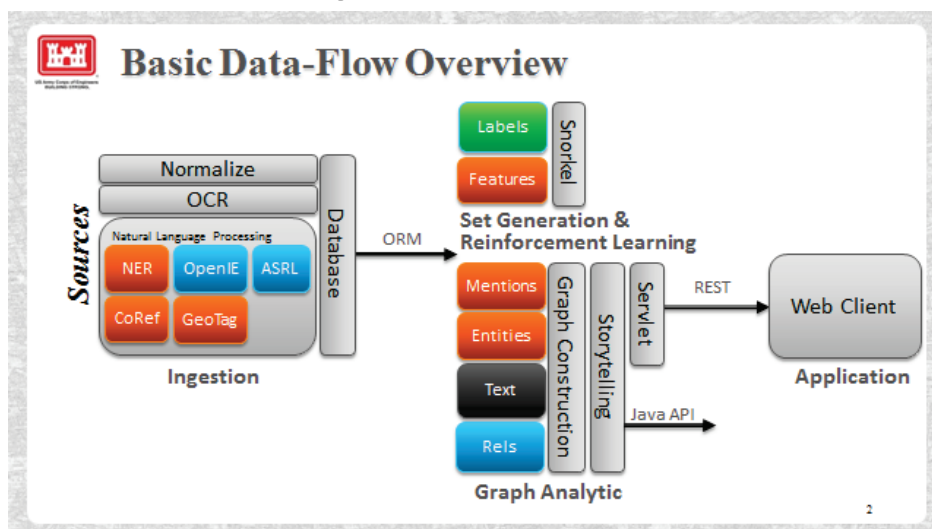
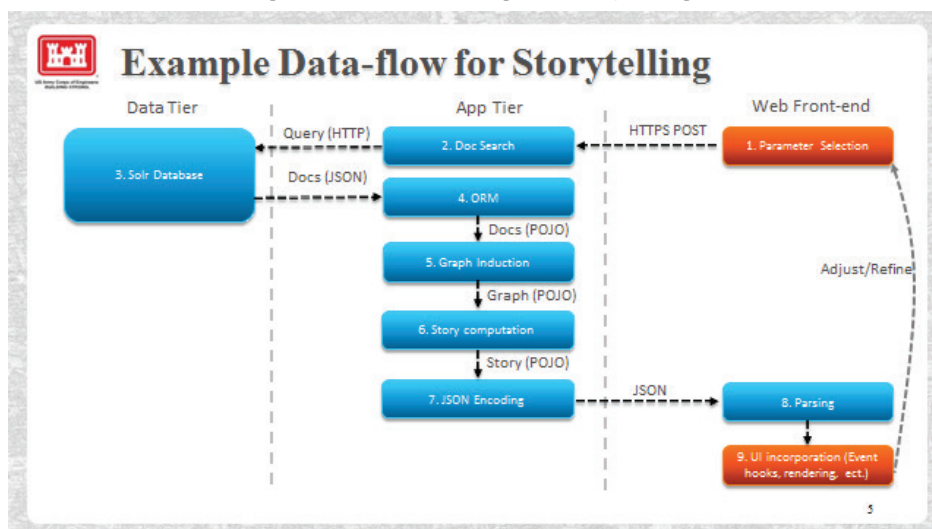


Figure 3. Data flow overview.



Storytelling is one of the analytical functions within the prototype, executed after Solr returns the search results for a query. Figure 4 shows the process diagram for storytelling.

Figure 4. Data flow diagram, storytelling.



An example of a story returned by the algorithm can be seen in Figure 5. This story was obtained by ingesting documents from the SYNCOIN (synthetic counterinsurgency) dataset (Graham et al. 2011) and querying for entities of “Al-Hardi” and “Tarik Mashal.” Figure 6 displays summary statistics for the dataset. The SYNCOIN dataset is used to evaluate emerging tools and exercise emerging capabilities and methods with military analysts in their occupational specific tasks and includes “ground truth” explanations for the threats represented within the dataset.

Figure 5. Example of a story.

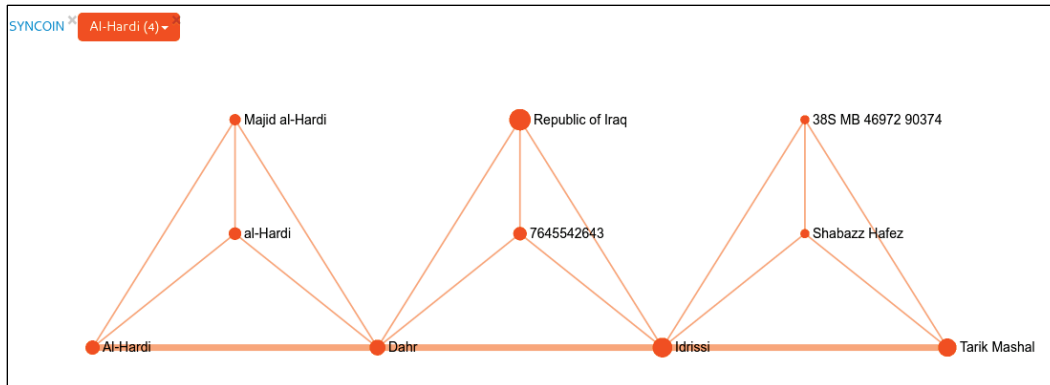


Figure 6. Summary statistics for SYNCOIN dataset.

Number of Documents:	1332
Number of Entity Occurrences:	5347
Number of nodes:	1578
Number of same-doc edges:	5084
Number of same-sentence edges:	4234
Last Updated:	Thu Aug 29 2019 16:41:43 GMT-0400 (Eastern Daylight Time)

From Figure 7, the sequence of entities connected in a straight line at the bottom of the graph from the starting entity on the left to the ending entity on the right denotes the story. Each entity connected to the story entities provides the context with which the algorithm used to score the path through the original base entity graph. A more human readable outline of the story is shown in the bottom pane of the user interface directly below the story graph pane (Figure 7).

Figure 7. Story summary.



If there is a sentence co-occurrence between two entities within the story, that sentence is selected for display to the user. Otherwise, the document(s) that contain that particular entity in the story are displayed.

3.3.4 Evaluation

3.3.4.1 TRADOC Operational Environment Center – May 2017

As part of the earlier STRIDER project, a limited user test was conducted of the storytelling algorithm at the Operational Environment Center (OEC), Training and Doctrine Command (TRADOC), Fort Eustis in May 2017. The results are documented here since they have not been published elsewhere.

Five employees from the OEC with a background in intelligence analysis participated in the test. Two employees used the prototype capability with the Atlantic Storm dataset while the other three used the prototype capability with the Bilasuvar Freedom Brigade (BFB) dataset. Both datasets are synthetic in nature and used to train intelligence analysts to construct entity networks and deduce possible threats contained with the message corpus. Each participant was given a worksheet with questions pertaining to those networks and threats in order to assess the accuracy of the algorithm as well as overall usability of the prototype capability. Each group of participants had four hours to complete the worksheet, not including time to offer critical feedback on storytelling and prototype mechanics. Figures 8 and 9 show summary statistics for each dataset.

Figure 8. Summary statistics for BFB dataset.

Number of Documents:	53
Number of Entity Occurrences:	778
Number of nodes:	257
Number of same-doc edges:	1748
Number of same-sentence edges:	699
Last Updated:	Wed Aug 07 2019 14:48:01 GMT-0400 (Eastern Daylight Time)

Figure 9. Summary statistics for Atlantic Storm dataset.

Number of Documents:	112
Number of Entity Occurrences:	1831
Number of nodes:	507
Number of same-doc edges:	4888
Number of same-sentence edges:	1404
Last Updated:	Wed Aug 07 2019 14:50:26 GMT-0400 (Eastern Daylight Time)

Overall, the Atlantic Storm dataset was about twice as large as the BFB dataset. The performance on the worksheet somewhat reflects this difference. The two participants using the BFB dataset achieved a mean score of ~89%, while the three participants on the Atlantic Storm dataset achieved a mean score of ~36%. For the latter, the mean score increases to ~48% if the unanswered questions are removed from the tally.

Although there were not enough personnel to serve as a control group for this limited user test, additional insight into the effectiveness of storytelling could be gleaned by surveying the anecdotal feedback given by the participants. Below are a few excerpts:

One analyst stated, "... the Storytelling was spot on. It was terrific as a verification tool" and "I would certainly want STRIDER if I was receiving a corpus of reports like this."

Another analyst answered whether it saves time by replying, "Yes. [For] General questions [,] it displayed relationships between organizations for additional analysis. [For] Specific questions [,] it focused analysis really quickly."

3.3.4.2 Warfighter Exercise 19-04 – April 2019

As part of the FOMAS project, the team was invited to participate during Warfighter Exercise 19-04, attending with a unit at Fort Bragg. The STRIDER prototype was demonstrated to participating personnel from a diverse range of Military Occupational Specialty (MOS). The Atlantic Storm dataset served as the primary demonstration dataset. However, working directly with unit personnel, exercise data was collected and ingested into STRIDER, and visualized and explored the data. Feedback from the users across multiple MOSs indicated that STRIDER would provide the users with useful capabilities unavailable within their existing tools. One analyst noted that STRIDER enabled the visualization of network data in 10 min versus the current time of hours or days.

3.3.4.3 Evaluation conclusions

Given the worksheet scores during the limited TRADOC user test, the limited amount of time to complete the worksheet, and the anecdotal responses from the participants at both evaluations, the results are encouraging and warrant further, more detailed testing. Such further testing is anticipated to demonstrate that storytelling is indeed effective at aiding intelligence analysts in performing their tasks.

3.4 Link analysis

3.4.1 Approach

To understand the relationship between storytelling and link prediction, three link prediction algorithms were implemented in the prototype capability - random walk with restart (RWR), superposed random walk (SRW), and PropFlow.

The link prediction algorithms, like storytelling, use the data returned from a search query. The random walk concept forms the basis for all three algorithms. Predicting a link between any two entities a and b is a function of the probability that a ‘walker’ on the graph will traverse from a to b . The algorithms primarily differ in the computation of similarity between nodes. All of the algorithms output a list of predictions along with a score.

RWR is similar to the PageRank algorithm (Page et al. 1999) - it models a walker, on the graph, that has a choice of continuing on to connected nodes or jumping back to the original starting node. The software iteratively calculates the steady state transition matrix until the change in probabilities between each iteration is less than a set threshold. When complete, the transition matrix is sorted through as an adjacency list, finding node pairs that do not currently have edges between them and that have a transition probability above a set threshold. These edges are used as predictions.

SRW has two major differences with RWR. It does not have a restart probability – instead it models multiple random walkers starting from each node. The similarities along the paths of walkers from the same starting node are added – *superposed* – in order to derive the final similarity score between nodes. The iterations are set at a fixed number and the resulting transition matrix list is sorted the same as in RWR.

PropFlow limits the length of the random walk and then adds the similarity along the length of the path, propagating it to each node. The software calculates the score added to each node on the path by dividing the score of the current node by the sum of the scores of all connected nodes one hop away.

3.4.2 Evaluation

Link prediction and storytelling have different characteristics. Link prediction can only rank pairs of starting and ending entities based on the likelihood that they should be directly connected. Storytelling performs the same function and can also specify the likeliest paths on an entity graph, providing evidence for a direct connection. In other words, storytelling can provide multiple hypothetical explanations for a predicted link in addition to the prediction itself.

Consequently, a reasonable way to combine the outputs of both approaches might be to use link prediction as an initial ‘filtering’ function in order to produce a prioritized list of predicted links. Follow-on analysis would use storytelling to investigate whether or not the predicted links are valid by examining the stories produced when using the same starting and ending entity.

Figure 10 shows an example output of link prediction using the SRW algorithm. The other two algorithms also output a similar list, sorted in descending order based on score by default.

Figure 10. SRW link prediction results.

SYNCOIN > LinkPrediction

Model Links

<ul style="list-style-type: none"> Dora Dhanun Ahmad 	0.0027
<ul style="list-style-type: none"> Dora Ali Hakem 	0.0025
<ul style="list-style-type: none"> Dora Hosni 	0.0025
<ul style="list-style-type: none"> Kaleed Dahr Idrissi 	0.0024
<ul style="list-style-type: none"> Dora Qazi 	0.0021
<ul style="list-style-type: none"> Dora Saiar al Habib 	0.0021
<ul style="list-style-type: none"> 	0.0020

If the results in Figure 10 are sorted based on the starting entity used in the story shown in Figure 5 (Al-Hardi), the list shown in Figure 11 is obtained.

Figure 11. Abridged SRW predictions (al-Hardi).

Starting entity	Ending entity	Score
Dora	Majid al-Hardi	0.0007
Mandari	Majid al-Hardi	0.0003
Majid al-Hardi	Idrissi	0.0003
Abdul	Majid al-Hardi	0.0002
Hassan Al-Buredi	Majid al-Hardi	0.0002
Ramses Al-Sabah	Majid al-Hardi	0.0002
Omar Khrayesh	Majid al-Hardi	0.0002
Majid al-Hardi	Tarik	0.0002

Note that the last link prediction captures the story from the example. However, the analyst must still sift through documents in order to understand why the algorithm predicted this link. Storytelling offers a quick way to calculate and evaluate multiple hypothetical explanations.

3.5 Storytelling prototype conclusions

In this report, a storytelling capability was introduced for automatically discovering hypothetical connections between pairs of entities that are mentioned within a large corpus of digital text documents. This capability first establishes inferred direct connections between pairs of entities that occur within a common sentence. Then, it applied weights to these inferred connections based on distance functions calculated from (1) the sets of documents containing the entities, (2) the distributions of words near mentions of each entity, and (3) the patterns of connection of entities to other entities. Finally, the capability used a shortest-path algorithm to discover the chain of inferred connections from one entity to another that gives the best evidence for a hypothetical connection between the two entities, if such a chain exists.

While a rigorous, formal evaluation of the effectiveness of this capability has not yet been performed, evidence from hands-on, informal use of the capability suggests that it will allow users to discover connections between entities more quickly, easily, and effectively.

In addition to storytelling work, three approaches to link analysis were developed: random walk with restart (RWR), superposed random walk (SRW), and PropFlow. Link analysis is faster than storytelling, but does

not provide explanations for the shortest path. Link analysis can be combined with storytelling, identifying potential results to explore in depth. Due to project scope constraints, no formal testing was conducted in FY 19.

4 Technology: Big Open Source Social Science

The FOMAS team provided contract oversight for the BOSSS prototype. COR responsibilities were executed, including verifying technical performance, maintaining communication and coordination with Boston Fusion, and inspecting and testing deliverables and reports, along with engaging military stakeholders. Initial BOSSS design and implementation work is reported in the conference paper, “Big open-source social science: capabilities and methodology for automating social science analytics” (Palladino et al. 2018). A complete, separate BOSSS technical report and briefing are each available as Distribution B materials through the DTIC library (Palladino et al. 2019a and Palladino et al. 2019b).

4.1 BOSSS user engagements

ERDC was invited to demonstrate BOSSS to users as part of the Warfighter Exercise 19-04 at Fort Bragg, North Carolina in April 2019, and Enterprise Challenge 19 at Fort Huachuca in May 2019. Feedback was collected about BOSSS features such as usability and IPB support. Due to network and prototype limitations, the team worked with BOSSS using previously collected and processed data at these events.

4.1.1 Enterprise Challenge 19

At Enterprise Challenge 19, ERDC-CERL and Boston Fusion Corporation, conducted a limited objective experiment of the BOSSS prototype software with 35-series intelligence analysts, DCGS-A trainers, requirements staff, and others. Nine personnel participated and started by completing a ‘before’ questionnaire. They then received an overview, instructions, and hands-on use of the BOSSS software (with three sample datasets included). It was concluded with an ‘after’ questionnaire and discussion.

Eight personnel fully completed ‘before’ questionnaires and all nine completed ‘after’ questionnaires. Ten people participated in the discussion sessions. Full details on the written responses are available in the BOSSS report (Palladino et al. 2019a).

User feedback indicated that the users believed BOSSS has relevance throughout the military decision making process and as an aid to IPB

before and during operations, as part of targeting, and for battle damage assessment for kinetic and nonkinetic actions. BOSSS was seen as an enabler for analysis in support of competition within Multi-Domain Operations (MDO). Participants indicated that, with continued development, BOSSS could save analysts significant time and greatly expand what analysts can make sense of and consider when analyzing the OE variables, especially with respect to the information environment. This suggests that BOSSS enables sensemaking.

In the written feedback, all analysts that used the prototype software indicated that the current software would make IPB, analysis of entity networks, and PMESII-PT/ASCOPE analysis more efficient. Analysts noted that PMESII-PT/ASCOPE analysis is a significant portion of the workload of Series '35F' (All Source) analysts. It requires many hours each week and the analysts have few specialized tools available to collect, analyze, understand, and present data.

BOSSS offers a unique, valuable, and workflow-relevant design. Analysts appreciated that the inclusion of news and other sources helps save significant time on the early sensemaking stages (i.e., collection, prioritization, organization). With minor modifications and enhancements, the Soldiers envision BOSSS useful at many echelons. One respondent succinctly stated, "I do love the tool. I also believe it could be used from the Company level to Corps because it's such an easy tool to use. Simply put it saves tons of time."

4.2 BOSSS products

The Army's STTR program funded the development of the BOSSS prototype software. As a product of STTR investment, the BOSSS prototype is subject to 'SBIR/STTR' rights, which are detailed in the Defense Federal Acquisition Regulation Supplement (DFARS). Under the contract, the BOSSS source code was retained using the Defense Intelligence Information Enterprise (DI2E) platform to enable sharing and exchange of code to authorized, appropriate partners. Additionally, a functional prototype of BOSSS is configured on the AGE Node – a computing capability on the DREN. The AGE Node provides a capability to mimic Army computing environments, such as the CPCE, enabling demonstrations, testing and evaluation.

4.2.1 Civil Affairs Proponent Office

The FOMAS team also worked with the Civil Affairs Proponent Office and established a KTA pertaining to BOSSS. The final BOSSS report will serve as one of the KTA deliverables.

5 Summary and Recommendations

This report documents the FOMAS work and the team's efforts to experiment with and improve algorithms, tools, and methods. These algorithms enable Army personnel to more rapidly and comprehensively iterate through the stages of sensemaking and storytelling.

The four tasks included: (1) analyzing gap statements, concepts and doctrine, and identifying workflows relevant to the technologies being investigated – BOSSS and storytelling; (2) investigating link prediction for enhanced storytelling; (3) overseeing BOSSS development; and, (4) exercising the BOSSS and STRIDER prototypes.

Through user engagement, it was determined that most analysts perceive opportunities for automation support during Mission Analysis and especially IPB. Under this applied research, the team did not have the opportunity to perform rigorous, formal evaluations of the prototypes. However, both STRIDER and BOSSS demonstrated promise.

Further investigation and development is recommended of automation tools to improve analyst capabilities. STRIDER and BOSSS should be integrated with other analytical tools and aligned with user training and workflows. These enhancements would improve sensemaking, IPB, and Mission Analysis by enabling analysts to minimize time spent on menial tasks and focus on tasks requiring human insight.

References

- Al-Zaidy, R., B. C. M. Fun, A. M. Youssef, and F. Fortin. 2012. Mining criminal networks from unstructured text documents. *Digital Investigation* 8(3-4): 147-160.
- Barranco, R. C., A. P. Boedihardjo, and M. S. Hossain. 2017. Analyzing evolving stories in news articles. *International Journal of Data Science and Analytics*, 1-16.
- Dos Santos, R. F., S. Shah, A. Boedihardjo, F. Chen, C. T. Lu, P. Butler, and N. Ramakrishnan. 2016a. A framework for intelligence analysis using spatio-temporal storytelling. *GeoInformatica*, 20(2): 285-326.
- Dos Santos, R. G., A. Boedihardjo, S. Shah, F. Chen, C. T. Lu, and N. Ramakrishnan. 2016b. The big data of violent events: algorithms for association analysis using spatio-temporal storytelling. *GeoInformatica*, 20(4): 879-921.
- Dos Santos, R., S. Shah, F. Chen, A. Boedihardjo, C. T. Lu, and N. Ramakrishnan. 2014. Forecasting Location-based Events with Spatio-temporal Storytelling. In *Proceedings of the 7th ACM SIGSPATIAL International Workshop on Location-Based Social Networks*. New York, NY: ACM.
- Egeth, J., Klein, G. L., & D. Schmorow. (2014). Sociocultural Behavior Sensemaking: State of the Art in Understanding the Operational Environment. No. MITRE-14-2487. McLean, VA: MITRE Corp.
- Faloutsos, C., K. S. McCurley, and A. Tomkins. 2004. Fast discovery of connection subgraphs. In *Proceedings of the tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York, NY: ACM.
- Graham, J. L., D. L. Hall, and J. Rimland. 2011. A COIN-inspired synthetic dataset for qualitative evaluation of hard and soft fusion systems. In *14th International Conference on Information Fusion* (pp. 1-8).
- Hart, P.E., N. J. Nilsson, and B. Raphael. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics* 4(2): 100-107.
- Headquarters, Department of the Army (HQDA). July, 2019. Intelligence. Army Doctrine Publication (ADP) 2-0. Washington, DC.
- Headquarters, Department of the Army (HQDA). July, 2019. Operations. Army Doctrine Publication (ADP) 3-0. Washington, DC.
- Headquarters, Department of the Army (HQDA). July, 2019. The Operations Process. Army Doctrine Publication (ADP) 5-0. Washington, DC.
- Headquarters, Department of the Army (HQDA). July 31, 2019. Mission Command: Command and Control of Army Forces. Army Doctrine Publication (ADP) 6-0. Washington, DC.
- Headquarters, Department of the Army (HQDA). March 1, 2019. Intelligence Preparation of the Battlefield. Army Techniques Publication (ATP) 2-01.3. Washington, DC.

- Headquarters, Department of the Army (HQDA). July 2015. Army Design Methodology. Army Techniques Publication (ATP) 5-0.1. Washington, DC.
- Headquarters, Department of the Army (HQDA). May 5, 2014. Commander and Staff Organization and Operations. Field Manual (FM) 6-0. Washington, DC.
- Hetland, M. L. 2010. Python algorithms: Mastering basic algorithms in the Python language. New York, NY: Apress.
- Hossain, M. S., C. Andrews, N. Ramakrishnan, and C. North. 2011. Helping intelligence analysts make connections. In *Workshops at the Twenty-Fifth AAAI Conference on Artificial Intelligence*.
- Hossain, M.S., P. Butler, A. P. Boedihardjo, and N. Ramakrishnan. 2012. Storytelling in entity networks to support intelligence analysis. In *Proceedings from KDD '12: The 18 h ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York, NY: ACM.
- Jin, W., R. K. Srihari, and H. H. Ho. 2007a. A text mining model for hypothesis generation. In *Proceedings from ICTAI '07: IYh IEEE International Conference on Tools for Artificial Intelligence*.
- Jin, W., R. K. Srihari, H. H. Ho, and X. Wu. 2007b. Improving knowledge discovery in document collections through combining text retrieval and link analysis techniques. In *Proceedings from ICDM '07: IEEE International Conference on Data Mining*.
- Joint Chiefs of Staff. 2018. Joint Concept for Operating in the Information Environment. https://www.jcs.mil/Portals/36/Documents/Doctrine/concepts/joint_concept_s_jcoie.pdf?ver=2018-08-01-142119-830
- Kader, M. A., S. M. Naim, A. P. Boedihardjo, and M. S. Hossain. 2016. Connecting the dots using contextual information hidden in text and images. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- Klein, G., B. Moon, and R. R. Hoffman. 2006a. Making sense of sensemaking 1: Alternative perspectives. *IEEE Intelligent systems*, 21(4): 70-73.
- Klein, G., B. Moon, and R. R. Hoffman. 2006b. Making sense of sensemaking 2: A macrocognitive model. *IEEE Intelligent systems*, 21(5): 88-92.
- Klein, G., Phillips, J. K., Rall, E. L., and D. A. Peluso. (2007). A data-frame theory of sensemaking. In *Expertise out of context* (pp. 118-160). Psychology Press.
- Kumar, D., N. Ramakrishnan, R. F. Helm, and M. Potts. 2008. Algorithms for storytelling. *IEEE Transactions on Knowledge and Data Engineering*, 20(6): 736-751.
- Leskovec, J., A. Rajaraman, and J. D. Ullman. 2014. Mining of massive datasets. Cambridge, UK: Cambridge University Press.
- Loffert, J. M. 2002. Mission analysis: Giving commanders what they need. Army Command and General Staff College, Fort Leavenworth, KS, School of Military Studies.

- Malpani, N., and J. Chen. 2002. A note on practical construction of maximum bandwidth paths. *Journal of Information Processing Letters* 83(3): 175-180.
- Manning, C. D., M. Surdeanu, J. Bauer, J. Finkel, S. J. Bethard, and D. McClosky. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 55-60.
- Moore, D. T. 2013. Sensemaking: A structure for an Intelligence Revolution. Government Printing Office.
- Page, L., Brin, S., Motwani, R. and T. Winograd. 1999. The pagerank citation ranking: Bringing order to the web. Stanford InfoLab.
- Palladino, A., Bienenstock, E. J., George, C. A., & Moore, K. E. 2018. Big open-source social science: capabilities and methodology for automating social science analytics. In *Next-Generation Analyst VI, Vol. 10653, p. 106530D*. International Society for Optics and Photonics.
- Palladino, A., E. Bienenstock. 2019a. Big Open-Source Social Science (BOSSS) Phase II Option Final Report. Lexington, MA, Boston Fusion Corp.
- Palladino, A., E. Bienenstock, K. Moore, M. Salwen, S. Flanagan, C. Weir, R. Arantes, F. Diedrich. 2019b. Big Open-Source Social Science (BOSSS) Phase II Option Final Briefing. Lexington, MA, Boston Fusion Corp.
- Perkins, T. K., and M. K. Strand. 2016. Beliefs, values, interests and practices of identities, networks, groups for planning & analysis. <https://www.sbir.gov/sbirsearch/detail/871663>.
- Pirolli, P., and S. Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of International Conference on Intelligence Analysis* (Vol. 5, pp. 2-4).
- Sieck, W. R., G. Klein, D. A. Peluso, J. L. Smith, D. Harris-Thompson, and P. A. Gade. 2007. FOCUS: A model of sensemaking. Klein Associates Inc., Fairborn, OH.
- Tolone, W. J., J. Whitmeyer, J. Walsh, M. Hadzikadic, T. Carmichael, M. Armstrong, T. K. Perkins, and C. C. Rewerts. 2015. A Design for Computationally Enabled Analyses Supporting the Pre-Intervention Analytical Framework (PIAF) ERDC/CERL-TR-15-14. Champaign, IL: U.S. Army Engineer Research and Development Center, Construction Engineering Research Laboratory.
- U.S. Army Training and Doctrine Command. 2018. The U.S. Army in Multi-Domain Operations 2028. https://www.tradoc.army.mil/Portals/14/Documents/MDO/TP525-3-1_30Nov2018.pdf
- Weick, K. E. 1995. Sensemaking in organizations (Vol. 3). Sage.
- Wang, B., H. Dong, A. P. Boedihardjo, C. T. Lu, H. Yu, I. Chen, and J. Dai. 2012. An integrated framework for spatio-temporal-textual search and mining. In *Proceedings of the 20th international conference on advances in geographic information systems*, pp. 570-573. ACM.

- Weick, K. E. 2012. Making Sense of the organization, Volume 2: The impermanent organization. John Wiley & Sons.
- Weick, K. E., Sutcliffe, K. M., & Obstfeld, D. (2005). Organizing and the process of sensemaking. *Organization science*, 16(4): 409-421.
- Whalley, L. A., N. R. Myers, J. L. Burton, D. A. Morrison, A. M. Rhodes, I. Scarborough, J. Vendryzyk, W. Zhang, N. Vander-Most, D. R. Taylor, D. A. Krooks, M. K. Chawla, C. C. Rewerts, and T. K. Perkins. 2014. Framing an Understanding of Sociocultural Dynamics for Civil-Military Operations. Champaign, IL: U.S. Army Engineer Research and Development Center, Construction Engineering Research Laboratory.

Acronyms

Abbreviation	Term
ADM	Army Design Methodology
ADP	Army Doctrine Publication
ADRP	Army Doctrine Reference Publication
AGE	Army Geospatial Enterprise
AMP3D	Automated Man-Machine Learning and Processes for 3D Terrain
ASCOPE	Areas, Structures, Capabilities, Organization, People, Events
BFB	Bilasuvar Freedom Brigade
BOSSS	Big Open Source Social Science
CASCOM	Combined Arms Support Command
CCIR	Commander's Critical Information Requirements
CERL	Construction Engineer Research Laboratory
CFT	Cross Functional Team
COR	Contracting Officer's Representative
CPCE	Command Post Computing Environment
DFARS	Defense Federal Acquisition Regulation Supplement
DI2E	Defense Intelligence Information Enterprise
DREN	Defense Research and Engineering Network
DTIC	Defense Technical Information Center
EEFI	Essential Element of Friendly Information
ERDC	Engineer Research and Development Center
FM	Field Manual
FOMAS	Foundations for Mission Analysis Storytelling
FY	Fiscal Year
GAIDEN	Geo-Enabled Augmented Intelligence for Decisive Engagement
GRL	Geospatial Research Laboratory
HAL	Hyperspace Analogue to Language
IDF	Inverse Document Frequency
IED	Improvised Explosive Device
IPB	Intelligence Preparation of the Battlefield
KTA	Knowledge Transition Agreement
MDMP	Military Decision-Making Process

MDO	Multi-Domain Operations
METT-TC	Mission, Enemy, Terrain and Weather, Troops and Support Available, Time Available, and Civil Considerations
MOS	Military Occupational Specialty
MST	Minimum Spanning Tree
NER	Named Entity Recognition
OE	Operational Environment
OEC	Operational Environment Center
PMESII-PT	Political, Military, Economic, Social, Information, Infrastructure, Physical Environment, and Time
RWR	Random Walk with Restart
STE	Synthetic Training Environment
SRW	Superposed Random Walk
STRIDER	Spatio-temporal Retrieval and Introspection of Data and Embedded Relationships
STTR	Small Business Technology Transfer
SYNCOIN	Synthetic Counterinsurgency
TLP	Troop Leading Procedures
TRADOC	Training and Doctrine Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) July 2020		2. REPORT TYPE Final		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE Foundations of Mission Analysis Storytelling (FOMAS)				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 0602784A	
6. AUTHOR(S) Timothy K. Perkins, Harland Yu, and Sean Sovine				5d. PROJECT NUMBER 855	
				5e. TASK NUMBER 24	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Engineer Research and Development Center (ERDC) Construction Engineering Research Laboratory (CERL) 2902 Newmark Drive, Champaign, IL 61824 U.S. Army Engineer Research and Development Center (ERDC) Geospatial Research Laboratory (GRL) 7701 Telegraph Road, Alexandria, VA 22315				8. PERFORMING ORGANIZATION REPORT NUMBER ERDC TR-20-12	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Engineer Research and Development Center Geospatial Research Laboratory 7701 Telegraph Road Alexandria, VA 22315				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT <p>Mission analysis is a critical step in military planning and decision-making. It is currently time-consuming for analysts, who have few automated tools. The Foundations of Mission Analysis Storytelling (FOMAS) project developed algorithms, tools, and methods to automate sensemaking for mission analysis, which reduces the time and increases the effectiveness of the process.</p> <p>This report describes the FOMAS research, specifically as it relates to storytelling and link analysis. It includes descriptions of storytelling and a related prototype implementation, "Spatio-temporal Retrieval and Introspection of Data and Embedded Relationships, (STRIDER)." It also describes user engagements involving STRIDER and a prototype information collection and processing tool, the Big Open Source Social Science (BOSSS).</p> <p>The focus for this project was placed on capability development and adaptation. Both STRIDER and BOSSS underwent limited testing that demonstrated promise and received positive user feedback. However, project time did not allow for detailed rigorous technical evaluations of the prototypes. The team recommends further investigation and development of these automation tools to improve mission analysis for analysts.</p>					
15. SUBJECT TERMS Military planning – United States Military planning – Computer programs			Military intelligence – Analysis Automation Algorithms		
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code)