

NC STATE UNIVERSITY



Laboratory for
Analytic Sciences

Collaborate. Innovate. Transform.

2018 LAS Research Symposium

December 12, 2018

North Carolina State University



Laboratory for Analytic Sciences

COLLABORATE. INNOVATE.
TRANSFORM.

ANNUAL RESEARCH SYMPOSIUM JAMES B. HUNT LIBRARY

8:00 – 8:30	Registration		Lobby
8:30 – 8:45	Welcome	Mr. Mike Bender Director, LAS	Auditorium
8:45 – 9:00	Keynote Address	Dr. Anne Bowser Director of Innovation, Wilson Center	Auditorium
9:00 – 9:30	LAS Overview & 2019 Focus Areas	Dr. Alyson Wilson Principal Investigator, LAS	Auditorium
9:30 – 9:45	NCSU Welcome	Dr. Randy Woodson, Chancellor	Auditorium
9:45 – 10:00	BREAK		Auditorium
10:40 – 11:30	Panel Discussion: Human Machine Teaming		Auditorium
10:40 – 11:30	Panel Discussion: Advancing the Science of Analytics		Auditorium
11:30 – 1:00	NO-HOST LUNCH		
1:00 – 4:30	Poster Displays and Demonstrations		Duke Energy Hall, 2nd Floor
5:00 – 6:30	NO-HOST SOCIAL		State View Hotel Bar/Lounge



Laboratory for
Analytic Sciences

Collaborate. Innovate. Transform.

2018 Symposium Poster Session Abstracts

* indicates in day two session

1. OpenKE: Innovating Techniques for Open Source Tradecraft in Real-time*

OpenKE Team

Given the volume, variety, velocity and veracity of publicly available information (PAI), what proposed research solutions could assist analysts in effectively employing techniques for sense-making of information specific to a topic or domain of interest? How do we efficiently retrieve, organize and analyze big data/PAI instantaneously? The Laboratory for Analytic Sciences (LAS) has innovated the Open Source Knowledge Enrichment (OpenKE) prototype to manage the complexities of PAI. OpenKE hosts techniques designed to be semi-automated and allows for the configuration to scan a wide range of structured and unstructured publicly available data sources. OpenKE assists in identifying the 'right' information from the 'right' (reliable) sources and delivers capabilities to holistically analyze big data in an unclassified environment and integrate data of value into the mission environment. Simply stated, OpenKE blends art, science and technology to advance open source tradecraft, demonstrate novel capabilities, apply rigorous analytic standards and aid in the development of analytics required to meet the dynamic challenges presented by big data.

2. Textual Report Generation from Email Utilizing Temporal Topic Analysis*

**Colin Potts,
Sean Lynch,
Tracy Standafer**

We are building a system that automatically generates textual reports from large email datasets. This work has been tested extensively on the Enron Corpus and the Avocado Research Email Dataset. The system computes communication clusters, email topics and categorization, and aggregate statistics for each cluster (e.g., volume, typical topics). After computing this analysis, we developed a modular system to generate textual reports automatically. Our work this year has focused on improving topic analysis. In addition to labeling emails and clusters with topics, we also compute a topic flow graph. This allows us to reason about how emails influence each other and how topics change over time. All of this analysis is handed to the report generation/text realization subsystems which select and organize the content, including explanations and contexts for why information is included/excluded.

3. 2018 Cybersecurity Efforts*

Vince Streiff

In 2018, Cyber efforts at LAS centered around making anticipatory cybersecurity through Structured Analytic Techniques practical.

4. BEAST 2018 Breakthrough: Refactored & Analyst Tested Queryless Tradecraft*

Samantha Szymczak

The BEAST is an integration platform currently supporting 45 tradecraft experiments and associated technologies. Multiple 2018 research efforts holistically enabled queryless tradecraft including: A refactored auto-query engine, integrated NLP and other extraction services, refactored relevance engine, improved triage user experience, and the addition of observability of queryless tradecraft. These breakthroughs in opportunistic structured tradecraft improve analytic rigor. And anticipatory analysis focus is enabled by queryless tradecraft and automated reporting. These efforts were proven in 2018 via an Enterprise integrated prototype currently in use by early adopters.

5. WESTWOLF: Insights into the Fundamental Structure of Analyst Workflow and Mission Performance

**Mark Wilson,
Michele Kolb,
Sam Wilgus,
Ruth Tayloe,
Jody Coward,
Mariza Marrero**

WESTWOLF is a light-touch application of behavioral science to optimize the workflow and performance of teams through alignment activities. The WESTWOLF process involves identifying leverage points for managerial action across workflow activities (e.g., communicating with customers, planning, attending meetings) and performance dimensions (e.g., efficiency, customer relations, performance recognition). In 2018, WESTWOLF has seen wide-scale implementation across several field sites which has resulted in helping managers increase transparency and candor when it comes to shared understandings of work activities and performance goals. Aggregated data from WESTWOLF implementation has also led to insights into the fundamental structure of analyst workflow and mission performance, revealing underlying workflow dimensions (general analysis, intelligence analysis, and language analysis) and performance dimensions (mission goal performance, mission health, and customer relations).

6. Why WESTWOLF? A Hypothetical Use Case?*

**Mark Wilson,
Michele Kolb,
Sam Wilgus,
Ruth Tayloe,
Jody Coward,
Mariza Marrero**

WESTWOLF applies behavioral science to optimize team workflow and performance by aligning activities where it has the most impact. Here we demonstrate a hypothetical WESTWOLF use case in a storyboard format to show how WESTWOLF works in the team space. Be sure to also visit the WESTWOLF: Insights into the Fundamental Structure of Analyst Workflow and Mission Performance poster that walks through the technical aspects of WESTWOLF.

7. Emerging Technologies: The (Continued) Rise of the Hobby Drones*

Peter Merrill

Technological advancements and manufacturing proficiencies in the consumer unmanned aerial vehicle (UAV aka drone) market have fueled rapid evolution of product specifications, flight capabilities, and streamlined user interfaces; significantly lowering the consumer's barrier to entry and pushing the boundaries for applications. Development of fully-autonomous consumer drones, integrating advanced AI and ML algorithms, could further enable independent operation of consumer drones. As users increase and applications and capabilities evolve, public safety, privacy, and national security concerns must be addressed.

8. How to Make a Magician*

**Suvodeep Majumder,
Timothy Menzies,
James Campbell,
Jared Stegall,
Aaron Wiechmann**

Most users of data mining technology are not magicians. They view those tools as complex black boxes controlled by many magic parameters. These magic parameters are difficult to set since there are so many of them. Due to a shortage of magicians, much data mining technology is used in inappropriate or sub-optimum ways. Clearly, we need to make more magicians that can suggest better ways to combine and configure the data mining tools. Since human magicians are in such short supply, we will use AI to

build automatic agents for helping users of data mining technology. Such tools will be able to convert LAS analysts into data mining “magicians”.

9. TUNINGWOLF*

**Felix Kim,
Jesse Miller,
Duncan Page,
Rob Tapp,
Sean Lynch,
Aaron Wiechmann,
James Campbell**

For many applications big data analytics (such as Map/Reduce and Spark jobs) are run on a cloud system over and over again, perhaps on a daily basis, to process new data which arrives continuously onto the system. The cloud systems running these analytics often implement resource allocation strategies which are only loosely informed from characteristics of the analytics themselves. For example, the number of reducers provided to an analytic may be chosen using a basic formula based on the amount of data read in by the analytic's mappers and the amount of memory reducers have available. Such strategies likely result in analytics being allocated more/less resources than they would in an optimal system. TUNINGWOLF is a prototype which employs mathematical techniques to attempt to determine an optimal set of resources for each analytic through an iterative and empirical process. Allocating resources to analytics using TUNINGWOLF is experimentally shown to improve runtimes of individual analytics, thereby increasing the overall throughput of analytics on the system as a whole.

10. Machine Learning Workflow Management*

**James Campbell,
Aaron Wiechmann,
Michael Green,
Jason Mooberry,
Natalie Smith**

This research project demonstrates an architecture designed to enhance analytic integrity across the full lifecycle of machine learning projects, including creating labeled datasets, model management, and ongoing validation. The architecture provides three key services--attribute-based data sampling, data labeling, and model validation--as well as two reusable repositories, for labeled data and models. By reusing common services, project owners can focus more on specific data and user requirements.

11. Domain-Specific Machine Translation*

**Natalie Smith,
James Campbell,
Richard Tait,
Chris Carr**

Our experiments show that domain-specific machine translation (MT) models can achieve adequate results with a training corpus much smaller than reported in other research when the domain is consistent and highly constrained. Several parameters affect the performance of statistical and neural MT models. Neural transformer models dramatically outperformed models created using Moses, an open source statistical machine translation system, but demonstrated more subtle patterns of error that may not be captured using traditional measures of translation quality.

12. Multifaceted Classification of Crowdsourced Futures to Support GEN1 Anticipatory Analysis*

Patrick Laughlin

While intelligence analysis conducted on an expanded future horizon is more likely to prevent cases of strategic surprise (e.g., Arab Spring), little to no tooling currently exists to support such anticipatory

analysis. By applying natural language processing (NLP) to public data sources, we can identify crowdsourced futures for analyst consideration. Such futures consist of a broad range of eventualities with possible intelligence implications. Furthermore, classifying crowdsourced futures along horizontal, modal, and attitudinal dimensions allows analysts to triage and explore large numbers of futures within domains of interest.

13. The Need for Better Cryptocurrency Tradecraft: Analytics and Visualizations* **Peter Merrill**

Promises of decentralization, cryptographic security, and enhanced anonymity make cryptocurrencies an attractive option for nefarious actors. Current forensic and investigative techniques employed against cryptocurrencies by financial and law enforcement analysts remain highly manual, time consuming, and frequently ineffective. Further, many techniques lack scalability to handle Big Data. To better identify illicit activities conducted with cryptocurrencies, track illicit transactions, and identify entities responsible, new forensic software, tools, and analytic capabilities need to be developed and deployed. In particular, capabilities must address analysts' needs for data volume, characterization, and visualization.

14. Deanonymizing ShapeShift: Linking Transactions across Multiple Blockchains*

**Nathan Borggren,
Gary Koplik,
Paul Bendich,
John Harer**

We identify ShapeShift transactions in the blockchains of Litecoin, Bitcoin, Dash, and Zcash as they are converted to many other cryptocurrencies, including the popular Monero. These transactions account for more than \$250 million of cryptocurrency traffic that has bypassed Anti-Money-Laundering (AML) and Know-Your-Customer (KYC) laws. Using machine learning, we can recall ~3/4 of ShapeShift transactions looking at only ~1/4 of all transactions.

15. A Workshop-Based Collaborative Reporting Investigation*

**Hongyang Liu,
Byungsoo Kim,
Ruth Tayloe,
Sharon Joines**

With new and evolving technology and ever-increasing remote collaboration, intelligence reporting processes are at a crossroads. Key questions have arisen about how best to collaboratively produce reports to get the right information to the right decision-makers in the right context are more important than ever. The LAS Collaborative Reporting project has attempted to address one part of this puzzle: Namely, how can reports be produced in a collaborative manner across two or more offices or agencies? As a result of the 3-day workshop and the report assessment from the reviewers, this investigation has gained valuable research data on how to develop analytic tradecraft to aid analysts and reporters with collaborative reporting in order to improve collaboration and advance writing strategy within the intelligence community.

16. Design Thinking through Research Course

**Sharon Joines,
Byungsoo Kim,
Hongyang Liu,
Josephine Dorsett,
Catalina Salamanca,
Jennifer Peavey**

The Design Thinking through Design Research short course focused on introducing participants to design thinking through a weeklong immersive design experience. Participants engaged in a situated design challenge for which they applied a variety of design methods to develop solutions as a process of learning

and exploring design research process. Participants explored primary and secondary source data; analyzed and synthesized information; proposed and evaluated design solutions; and materialized and communicated design alternatives using a variety of tools for presentation and representation. This poster presents the 2018 short course including the course overview, course organization, changes for the 2018 short course, synthesis of students' evaluations, and insights and lessons learned.

**17. Attention Focus of LAS Participants:
How Boundary Objects Focus Diverse Perspectives**

**Beverly Tyler,
Jessica Jameson,
Brooke Lahneman**

A key problem for cross-sector interdisciplinary innovation teams, such as those at LAS, is that participants from government, academic, and industry focus their attention on different things. The Attention Based Theory posits that individuals' roles and organizational goals and values focus their attentional resources on one issue versus others, rather than being determined solely by training and experience. Individuals' roles focus their attention on certain issues, which inhibits their ability to shift their attention elsewhere. This can shape how individuals direct their focus, i.e., the specific spatial (global, regional, national) and temporal scales (i.e., long-term or short-term), and how they perceive the nature of a problem. It argues that formal organizational practices that encourage rich communication and the development of mutual understanding focuses diverse individuals' attention on the same problems at the same temporal scales, so they can perceive the issues at hand from similar perspectives. The Practice Based Perspective proposes a specific type of integrating device - boundary objects. Boundary objects establish a shared syntax or language for individuals to represent their knowledge and facilitates a process by which people can transfer their knowledge to one another. Common boundary objects would be standardized practices, a physical object, or shared concept (e.g., internet of things, smart cities). In this study based on 52 interviews with LAS participants, we consider how boundary objects help focus diverse team members' attention, creates a shared understanding, and facilitates cross-sector innovation.

**18. Facilitating Interdisciplinary Collaboration among the Intelligence
Community, Academy, and Industry**

**Jessica Jameson,
Beverly Tyler,
Kathleen Vogel,
Sharon Jones**

This poster provides information about our forthcoming book with Cambridge Scholars Publishing. It includes empirical data from our ethnographic study of collaboration at LAS. Data come from our analysis of several surveys, in-depth interviews, focus groups, and observations over a multi-year period (2014-2018). The book features case studies authored by members of the LAS community to model immersive collaboration and illustrate specific examples of research and discovery that is translated into new tools and techniques for analysis and tradecraft. Each chapter presents implications for collaboration theory and practice at the program, team, and individual levels of analysis.

19. Organizational Acculturation: Educate to Collaborate at LAS

Judith Johnston

The results of two LAS internal studies - one that addressed communication and collaboration at LAS, and one that addressed LAS members' understanding of translational research as the cornerstone of LAS work - arrived at similar conclusions regarding the challenges of immersive collaboration. These conclusions were translated to a description of the gap between what currently exists at LAS versus the vision of LAS and guide development of intervention that will help close that gap.

20. Enacting Immersive Collaboration*

**Jessica Jameson,
Mariza Marrero**

In 2018, a workshop was created and facilitated by Jessica Jameson and Mariza Marrero based on (i) sharing the idea that communication supports and creates immersive collaboration and (ii) identifying the behaviors and expectations that build and sustain a collaborative environment. These include listening, engaging, acknowledging expertise, building rapport, and nurturing relationships. Communication scenarios at LAS and in work teams were examined in order to develop constructive communication strategies and techniques. We also identified and practiced tips for dealing with especially difficult conversations, including cross-cultural interactions, negotiating expectations, holding each other accountable, and engaging in productive conflict management.

21. Building an Experimental Testbed for Structured Analytics & Collaboration

**Brent Younce,
Rob Johnston,
Judith Johnston**

Analysts' needs to employ structured analytic tradecraft (SAT) is surpassed by the existence of tools that actually facilitate the use of SATs. Initially, this project was conceived to address that need. It is simple, easy to use, scalable tool that guides the analysts' problem decomposition activity. It organizes and tracks data gathered to provide an audit trail, identify data gaps, and support report generation. As such, it was recognized as an environment that could form the framework for determining a tool's cognitive fitness for use in the IC. Based on our findings, we hope to determine experimental requirements for tools that are candidates for use in the IC. While there are technical requirements for software maturity and fitness for use in the IC, particularly in classified environments, there is a need to identify tests for cognitive/usefulness fitness for these tools.

22. The Leadership Effect: Optimal Team Performance*

**Mariza Marrero,
Carmen Vazquez**

What do the New England Patriots and the San Antonio Spurs have in common? No, not that! They are two of the most diverse sports teams in the U.S. They are also two of the most winning teams in American sports. Surprised? Well, their coaches and leaders on and off the field and court play an important role in the teams' success. Effective leaders have the ability to optimize a diverse talent pool. These leaders are able to communicate better and create an environment in which each team member or player feels valued, safe, and respected. Diverse teams perform better than homogenous teams, to include yielding positive business outcomes, or championships. Diverse teams are more creative as they draw from a wider range of knowledge and experiences to examine questions and develop solutions. Diverse teams are also more innovative and more willing to push the envelope when working hard-to-solve problems. When maximizing Diversity, team composition matters. The ideal diverse team is composed of culturally and cognitive/skills diverse individuals who are led by an effective leader who values a collaborative approach. While homogenous teams can form and can solve some problems more quickly, they are less creative or innovative than diverse teams. On the other hand, diverse teams may require more time to get to know each other, build trust, and learn how to communicate with one another. Once that is achieved, success is just a matter of time.

23. Lost in Transliteration: Making America a Place Where Everybody Knows Your Name

**Jared Stegall,
Jonathan Stallrich,
Richard Tait**

Have you ever come across someone from another country at either a business meeting or a social outing, and when you try and pronounce their name you completely mess it up? It leaves everyone in the scenario feeling bad and is the ultimate precursor to a poor first impression. For the past two years, Jared Stegall, Jonathan Stallrich, and Richard Tait have been working to combat this problem, along with

several others, by analyzing how well native English speaking American citizens pronounce common Korean names that have been transliterated into English. They combined their linguistic and technical expertise to scientifically assess the effectiveness of not only the current IC standard transliteration method, but also methods developed by the South Korean Government and academia to determine which one yielded the best pronunciation performance from native English speaking American citizens. This poster will take you through the origins of the project, experiments that were conducted, results found, supplemental research efforts, and where it's headed from here.

24. Veracity: Exploring the International Landscape of News Media Reliability

**Hector Rendon,
Alyson Wilson,
Jared Stegall,
Chris Nobblitt,
Drew Hollis**

Chances are, you've likely heard of the phrase "fake news" before. While in some contexts it is used to describe unfavorable reporting, it actually speaks to the larger concern of measuring the reliability of news sources. Have you ever stopped to wonder what makes a news source to be reliable? During the past several years, LAS has explored this issue through the partnership of media experts and statisticians to scientifically automate the process that the Intelligence Community uses to determine the reliability of media outlets. Having previously addressed the reliability of sources familiar to the United States context, in this year's Symposium our research examines the reliability of news outlets consumed all over the world and what exactly makes them reliable or unreliable. Hector Rendon, Alyson Wilson, Jared Stegall, Chris Nobblitt, and Drew Hollis created an expansive dataset of international news media sources that included their editorial characteristics, social media following, web traffic metrics, and country of origin attributes. With the help of the reliability ratings provided by several intelligence analysts working with LAS, we conducted a regression analysis and developed a predictive algorithm to determine the effects that each one of the variables had on the reliability of global news sources.

25. Analytic Frameworks

Ty Simmons

This research effort focused on understanding how existing advanced analytics frameworks can help inform the intelligence analysis life cycle. It surveys a variety of process frameworks associated with the disciplines of advanced analytics and intelligence analysis.

26. Research Transition*

**Steve Cook,
Jody Coward,
Dawn Hendricks,
Matthew Schmidt**

The objective of the various research transition efforts at LAS is to help build on the innovation at LAS to provide novel benefits to mission problems. Toward this goal, the research transition efforts help link LAS research with mission problems; inform current and potential partners, collaborators, and stakeholders about the progress of various LAS efforts; and create a shared set of expectations for what research transitions looks like. These research transition efforts are shared between LAS performers, NCSU's I2I team and government personnel responsible for the operations of applications in the mission space.

27. Improving LAS Technology Transfer

Eli Typhina

2018 Symposium Poster on my Translational Research/Technology Transfer study conducted with LAS members and review of LAS documentation. Main focal point is a discussion of the process currently in use and proposed.

28. Advancements in Anticipatory Thinking*

**Sarah Margaret Tulloss,
Tracy Standafer,
John Slankas,
Jamie Roseborough**

The LAS team defines Anticipatory Thinking (AT) as the ability to reason about future events. The team hypothesizes that this ability is comprised of cognitive skills that can be developed and improved, and as such has developed a trans-disciplinary three pillar approach to explore AT. 2018 work has focused on AT assessment measures, AT support platforms, and AT training and there has been progress made under each category.

29. IOT 2018*

Deb Crawford

This poster provides an overview of the IoT Exemplar's research goals and activities.

30. Leveraging the Internet in the Internet of Things*

Jody Coward

Internet of Things (IoT) devices are varied in type, function, capabilities, etc. and their use is growing at a rapid pace. With the availability of an endless number of devices, the overall understanding is not keeping up at the same pace. A general knowledge is needed to understand how IoT devices interact with a digital and physical ecosystem. It would be useful to determine what information is available on the Internet regarding an IoT device, in hopes that the information will be a part of a larger corpus of understanding.

31. IIoT Lab – Infrastructure*

**William Amass,
Sheila Bent,
Drew Lombardo,
Stephen Williamson**

LAS is developing a repeatable process to stand up an Internet of Things (IoT) laboratory environment. From the beginning, methodologies were developed to focus on compliant acquisition of data.

32. Clearer than MUD: Extending MUD for Securing IoT Systems

**Yannis Viniotis,
Mihail Sichitiu,
Simran Singh,
Ashlesha Atrey**

IoT devices are expected to increase in number exponentially in the next years. Many different manufacturers will produce a vast number of different models, all including custom software of highly variable quality and security. Especially the security is often overlooked when manufacturers have to deliver functionality on a tight schedule. In this poster, we focus on an approach we developed involving capturing and demodulating wireless packets from IoT devices, and analyzing the resulting packet trace for usage patterns. These patterns can then be described and enforced using the Manufacturer Description Protocol (MUD): originally MUD was developed to restrict the degrees of freedom that an IoT device (or any network connected device) has (e.g., limit the range of ports used). In this project we use it to further characterize (limit) the normal behavior of an IoT device, and thus allowing for detecting (and preventing) abnormal behavior.

33. Security Analysis of Consumer IoT Devices

**Junia Valente,
Alvaro Cardenas**

This poster summarizes our work on IoT security and privacy. We analyzed the security practices of 40 IoT devices and identify trends, common vulnerabilities, and new potential security solutions.

34. IoT Fingerprinting

**David Anderson,
Mary Beth Simmons**

In this project, we developed a method of identifying device types based on the internet traffic metadata of individual IoT devices. We classified IoT devices into eight broad categories and calculated the average and standard deviation values for each device in that category on 20 features related to the data transmission and reception of each device over a two-month training period, on about 140 devices. We then took the observed traffic from a new set of devices during a different time period and tried to classify which type of device created the recorded traffic data.

35. Topological Data Analysis for the IoT: Discovery and Classification*

**Nathan Borggren,
Lihan Yao,
Gary Koplik,
Paul Bendich,
John Harer**

The persistogram, a topological signal analysis tool, is introduced and used to identify IoT devices from radio samples collected by a Software Defined Radio. We find that machine learning on persistogram-derived features is sufficient to distinguish devices without the need for demodulation and decryption, even in constrained environments with significant loss of packet captures. Machine learning is also performed with spectrogram-derived features which are also sufficient to classify devices. Persistograms, however, are shown to 1) train with less data, 2) concentrate information in fewer principle components, and most importantly, 3) are more resilient to the addition of noise. Persistograms may also prove to be useful in other classification tasks of time series where sampling rates are high.

36. Enhancing Information Discovery Workflows via Human-Machine Collaboration Interventions*

Kenneth Thompson

What if machines were more proactive during analysts' information discovery process? Instead of analysts doing all the work, including telling the machines what to do, the machine had agency to complete certain tasks (interventions). As a member of a human and machine collaborative team, the machine could help lessen the obstacles between the analysts and the discovery of actionable intelligence. However, to build a Human-Machine Collaboration (HMC) system, it is, in part, necessary to understand the complexities of the analyst workflow. We leveraged LAS and community analyst expertise; reviewed tradecraft taught to our analysts; analyzed what our analysts say they do; and, began analyzing workflow instrumentation data to see what they do. This allowed us to identify several steps/elements of the information discovery process, as well as a few pain-points experienced by analysts. As we identified pain-points, we worked with our academic and industry partners to begin research and development of HMC interventions that could alleviate those pain-points.

37. CyberMonkey*

**Brent Younce,
Matthew Schmidt,
Kenneth Thompson**

For analysts, the gathering and manipulation of data often requires the use of computerized tools such as spreadsheets, databases, and digital report archives or other lookup tools. However, each of these tools

requires first finding interesting inputs out of a large set of available data, choosing and configuring the tools themselves, analyzing outputs, and potentially restarting the cycle. With constantly growing computational power, computers are effectively waiting on users more than users wait on them. This presents an opportunity to take advantage of this excess computational capability to improve efficiency and improve the analytic process. In order to allow for a large number of analytic tools to be executed proactively in-browser (simplifying compliance and improving security), the CyberMonkey framework was created. CyberMonkey is a portable, modular system which provides several methods of user interaction and allows for an arbitrary set of JS- or web-based analytic tools to be run in-browser, automatically finding meaningful input and highlighting interesting information.

38. Automated Hypothesis Generation

**Shrey Anand,
Nirav Ajmeri,
Munindar Singh**

Effective analytics requires the consideration of multiple hypotheses to determine a suitable explanation or decision given the evidence. However, it is difficult to produce a sufficient variety of hypotheses and thus, analysts can be blinded by failing to consider certain possibilities. We present an approach for automated hypothesis generation that can help the analyst explore the space of possibilities more thoroughly than otherwise. Beginning from a textual prompt produced by an analyst or derived from an event description, we generate hypotheses using a statistically learned model. The hypotheses provide examples that stimulate creativity and present an adversarial perspective in the spirit of Devil's Advocacy.

39. Predicting View Changes in Online Argumentative Discussions

**Zhen Guo,
Munindar Singh**

Information overload complicates how an analyst may seek information and make sense of it. To assist analysts in information discovery and decision making from a human-machine collaboration perspective, we investigate the problem of understanding how analysts seek and are influenced by available information and develop a computational model to help analysts explore alternative hypotheses and evidence. To this end, we adopt the Reddit Change My View data and consider the problem of determining whether a posting can change the view of an opinion holder. We apply a sequential machine learning model to capture both adjacent and nonlocal dependencies during the information seeking process. We identify valuable features and patterns and reveal a latent state where an analyst may benefit from intelligent assistance. To our best knowledge, this is the first work on understanding how views change in response to arguments over time.

40. The Effects of Individual Differences on Visualization Use*

**Jordan Crouser,
Maddy Kulke,
You Jeen Ha**

Visual reasoning aides (e.g. maps, diagrams, data visualizations) support analysts in orienting themselves to a dataset, identifying patterns, exploring areas of interest, and ultimately building a mental model of the underlying phenomena. Because these tools are intended to support complex analytic tasks, their efficacy and use may be moderated by differences in individual analysts' cognitive style. In this survey, we provide a comprehensive overview of existing scholarship into how individual differences inform and influence visualization use.

41. Using Stories to Elicit Anticipatory Thinking*

**Sarah Margaret Tulloss,
Kelly Neville,
Robert Bechtel,
Robert Wray,
Tyree Cowell**

A software prototype was developed for the purpose of enriching intelligent analysts' anticipatory thinking. The tool, called Anticipatory Thinking (AT) Story Web, is designed to help analysts identify a broader and more varied set of possible futures. As analyses are developed, it presents users with information potentially relevant to their ongoing analysis work. Storyweb uses analogical reasoning to attempt to map analyst work products to a database of relevant, past cases, which are extracted from open-source news articles, case studies, and other analysis resources. A workshop was conducted to evaluate the usefulness of the Storyweb capabilities to analysts. Eight intelligence analysts volunteered to participate in two (n=4) or three (n=4) 1.25-hr analysis sessions. They were instructed to assess, across sessions, potential cyber threats to the 2019 Canadian federal election and to represent their work using concept-mapping software (IHMC's CMapTools). At the start of Sessions 2 and 3, the analysts were provided with Storyweb-generated recommendations to consider. Seven of the eight analysts reported that Story Web gave them ideas and suggested lines of inquiry that they otherwise may not have considered. An eighth analyst indicated that the background stories used by Storyweb were not supportive of the systems-focused analysis strategy the analyst prefers. Most (five) analysts report that Storyweb contributed to an enrichment of existing possibilities they were building out; three reported that Storyweb helped them think outside the box, i.e., divergently. All participants suggested improvements to the user interface, (which is not yet fully developed). Five analysts noted a need for additional context, e.g., meta-data about or access to the original source document, for the extracted information with which they were presented.

42. Scenario Explorer: An Imagination Support Platform for Anticipatory Thinking **Chris Argenta, Matt Lyle, Abigail Browning**

The Scenario Explorer platform is a combination of technology and tradecraft designed to help analysts systematically anticipate or imagine a wide range of diverse futures.

43. Operationalizing Social Reasoning with Theory of Mind

**Markus Eger,
Chris Martens**

Social reasoning is important for communication between humans and therefore plays an important role in effective communication between computers and humans. A key component of social reasoning is an understanding of how belief and intention drive action, which can be used to make sense of humans' actions as well as to produce more interpretable machine decision-making. We use a formal model of theory of mind based on dynamic epistemic logic to represent "nested beliefs" that simulated agents have about each other. We apply our model to the experiment domain of a social deduction card game, One Night Ultimate Werewolf, and develop an agent model that plays the game through social moves that convince other players of certain belief states. We evaluate different approaches to agents' commitments to their intentions, revealing variation in performance and believability correlated with different levels of commitment.

44. Towards Adaptive Support for Anticipatory Thinking*

**Michael Geden,
Andy Smith,
James Campbell,
Randall Spain,
Sarah Margaret Tulloss,
Bradford Mott,
Jing Feng,
James Lester**

Adaptive training and support technologies have been used to improve training and performance in a number of domains. However, limited work on adaptive training has examined how to successfully train and assess anticipatory thinking skills, which reflect an individual's ability to deliberately explore and generate divergent outcomes and consequences of events on future scenarios. Anticipatory thinking skills are particularly important for intelligence analysis, mission planning, and mission forecasting wherein individuals must identify a range of outcomes and consequences that can arise from actions, events, or decisions and account for these effects during their decision-making process. Although anticipatory thinking has been recognized as an important skillset for practitioners who are asked to engage in sense-making and scenario-generation, there is currently no underlying theory supporting these methodologies, no adaptive technologies to support their training, and no existing measures to assess their efficacy. We are engaged in an ongoing effort to design adaptive technologies to support the acquisition and measurement of anticipatory thinking. As a first step, we have developed an interactive task to measure anticipatory thinking wherein participants engage in a divergent thinking exercise to provide insights about the impact of certain events on the future. We present preliminary results from a study to examine the validity of this measure and discuss multiple factors that affect anticipatory thinking including attention, inhibitory control, need for cognition, need for closure, convergent thinking, divergent thinking, and creativity. We then introduce our initial prototype environment for supporting training, application, and assessment of anticipatory thinking.

45. Overcoming Failures of Imagination with Elaboration Prompts

**Jim Davies,
Eve-Marie Blouin-Hudon**

The problem that we investigated in this project is this: If people make systematic errors when thinking about the future, what interventions would contract them? The project targets fundamentally understanding cognitive mechanisms involved in anticipatory thinking (i.e., imagining the future), as well as developing technology that iterates on pen and paper exercises to enable imagination of future scenarios. As such, it supports understanding the efficacy of structured analytic techniques, which can be used to help government analysts make better predictions about future scenarios and threats. We recruited university students to read a scenario about the political climate in Iran. Participants were then asked to improve on the scenario in 'a more creative way', to re-write the scenario in the past-tense, or to re-write the scenario by changing the most common attributes. Then, original versus new scenarios were rated for creativity by Mechanical Turk workers. Results suggest that lay perceptions of creativity may influence how individuals chose to create scenarios, and how others chose to rate creativity.

46. Resisting Influence and Supporting Stability Campaign (RISSC)*

**Michelle Kolb,
Rob Johnston**

The RISSC exemplar seeks to model complex behaviors and decision making to create a more effective early warning system for fragile states, influence campaigns, illicit networks, and radicalization. The focus is three-fold:

- Understanding the relative importance of factors that hallmark the impending collapse or deterioration of a fragile or failing state.
- Identifying and countering the effects of key influencing measures that are used to motivate people to make decisions counter to their own best interests.

- Identifying the structure and impact of illicit networks underpinning failing states and influence campaigns.

47. A Qualitative Analysis of Drivers among Lone Actor Terrorists: Does Military Affiliation Matter?*

**Alexa Katon,
Christine Brugh,
Kaleb Rostmeyer,
Samantha Zottola,
Sarah Desmarais,
Joseph Simons-Rudolph,**

Using qualitative data from the Western Jihadism Project (Klausen, 2017), this presentation compares drivers of violent action among military-affiliated (n=5) and civilian (n=5) lone actor terrorists. Iterative content analysis yielded four parent codes (Religious Fervor, Action, Growing Jihad, and Grievance). Military lone actors spoke more about action-related concepts and their grievances while civilian lone actors showed more religious fervor. Growth in Jihad was discussed slightly more often among military lone actors. Findings suggest meaningful differences and similarities between the statements made by military-affiliated lone actors and their civilian counterparts. Results may help improve identification of individuals at risk of lone actor terrorism.

48. Presence of Mental Illness among Lone Actor Terrorists*

**Kaleb Rostmeyer,
Alexa Katon,
Samantha Zottola,
Sarah Desmarais,
Joseph Simons-Rudolph,
Christine Brugh**

Using the Western Jihadism Project (Klausen, 2017) dataset, we describe the presence of mental illness among lone actors (n=79). Less than half of lone actors displayed evidence of mental illness (n=33, 41.8%). Substance abuse disorders (n=10, 30.3%) were most frequent and comorbidity was common (n=16, 48.5%). Only eight lone actors (24.2%) received or sought treatment before perpetrating a terrorist act and two (6.1%) were hospitalized in a psychiatric institution after. Findings show that lone actors are heterogeneous with regard to mental illness, but also underscore the need for further research on the role of mental illness in lone actor terrorism.

49. Interactions between Individual Characteristics and the Fragile State Index on Terrorism Outcomes*

**Samantha Zottola,
Christine Brugh,
Alexa Katon,
Kaleb Rostmeyer,
Sarah Desmarais,
Joseph Simmons-Rudolph**

Using data from the Western Jihadism Project (Klausen, 2017) and the Fragile States Index, this poster presents one of the only studies to analyze terrorism outcomes utilizing a multilevel design. Of the terrorists with terrorism-related arrests, involvement in plots or organization involvements, and domestic plot involvement in domestic plots specifically, those who were non-converts came from less stable countries (higher FSI). Further, of the terrorists involved in domestic plots, non-asylum seekers and asylum seekers came from countries with similar levels of stability. Results suggest that individual characteristics may differentially predict terrorism involvement based on outcomes of interest and stability of the country in which the individuals reside.

50. **Measuring Influence***

**Dhrubajyoti Ghosh,
James Robertson,
Soumendra Lahiri,
Rob Johnston,
William Boettcher,
Michele Kolb**

We analyze Twitter data on the 36th Amendment of the Constitution of Ireland. This amendment repealed the 8th amendment and allowed the government to legislate on abortion. The proposed legislation brought Ireland into line with the majority of European countries, allowing for abortion on request up to the 12th week of pregnancy (subject to medical regulation). The goal was to develop an initial predictive model using Twitter, analyze influencers, and identify indicators of changes in sentiment.

51. **Arabic OCR**

**Brent Younce,
Akram Khater**

The Moise A. Khayrallah Center for Lebanese Diaspora Studies at NC State maintains a growing dataset of over 250,000 historical Arabic documents, in the form of images. Researchers cannot make full use of these documents, however, as the set of images cannot be queried in any meaningful way due to a lack of digital text for these documents. Existing commercial OCR (Optical Character Recognition) systems perform extremely poorly on this dataset, correctly recognizing less than 5% of the visible characters. Through this project, several open source OCR solutions were tested, compared, and trained using both hand-transcribed text and over 100,000 artificially generated lines of Arabic text, producing several OCR solutions which performed with high accuracy (75% for one system, 95% for another). Additionally, a simple and an advanced web-based search system were developed, allowing querying of the entire dataset at once.

52. **Agent-Based Models for Illicit Network Simulations***

**Conor Artman,
Li Zhen,
Eric Laber,
Rob Johnston**

Modeling the complexity of human decision making and group behavior is inherently an interdisciplinary and difficult task. Understanding and describing human behavior crosses multiple disciplines from anthropology, cognitive science, economics, political science, psychology and other social sciences. Modeling those behaviors requires computational statistics, computer science, and mathematics. A large part of undertaking a task like this is developing a core interdisciplinary team that can speak and understand each other's professional jargon and come to a shared mental model of the work and a common language for the shared effort. Developing a simulation platform for a highly complex modeling environment to accommodate the scale of human behavior and decision-making led the team to the creation of an OpenABM gym. In tandem, we developed methodology for multi-agent reinforcement learning to automatically optimize agent behavior.

53. **On the Edge of Failure: Forecasting State Stability**

**William Boettcher,
Luis Esteves,
Dawn Hendricks,
Rob Johnston,
Michele Kolb,
Samantha Schultz**

Mali, Cameroon, and Senegal are three West African nations with sharply differing outlooks for future state stability, despite being broadly similar in many geographic and other characteristics. Case studies were developed for each of these countries, highlighting the axes on which the nations differ, in order to

better understand the possible causal mechanisms that may contribute to state instability. These results are used to inform upcoming work on state fragility in other regions of the world, as well as to guide the development of predictive statistical and machine learning models of state failure.

54. Explaining State Instability: Merging Quantitative and Qualitative Efforts* Arnab Chakraborty,
Samantha Schultz,
Luis Esteves,
Soumendhra Lahiri,
Rob Johnston,
William Boettcher,
Jascha Swisher

This poster describes an initial case study, focused on Cameroon, to support the development of a machine learning model to forecast state stability. The case study offers a check on the validity of the variables in the machine learning model and allows the comparison of real-life events with model forecasts. Through the combined qualitative and quantitative efforts, we hope to gain greater understanding of the underlying causal mechanisms that affect state stability.

55. Forecasting State Instability Arnab Chakraborty,
Soumendhra Lahiri,
Rob Johnston,
William Boettcher,
Michele Kolb,
Jascha Swisher

Policy makers have long sought early warning of the negative events associated with state fragility/failure. One well-known public effort regarding this is Fragile State Index (FSI). But this index is hampered by Western, Educated, Industrialized, Rich and Democratic bias; it is yearly-lagged and not sensitive to forecast major negative events or state instability. In this poster we present our efforts to develop an alternative model that eliminates the biases; relies on better and timely data; and utilizes more sophisticated statistical techniques and hence more sensitive to events and time.

56. Radicalization: A Meta-Cognitive Analytics Approach through Sentiment/Emotions Analysis and Deep Learning* William Agosto-Padilla,
Carlos Gaztambide,
Mariza Marrero,
Jason Mooberry

The work focused on identifying the Radicalizer's Cognitive Modus Operandi in order to understand the workflow for analytic tradecraft. This approach aimed to assist the intelligence analysts to effectively identify evidence of radicalization attempts. Cognitive Behavioral-Emotive theories, Sentiment and Affect Analysis methodologies and Machine Learning/Deep Learning algorithms were implemented in this study.

57. Understanding Analyst Workflow through Baseball Analytics Justin Middleton,
Kathryn Stolee,
Emerson Murphy-Hill

Modeling the processes and tools by which analysts collect, clean, process, and interpret their data is an important step to addressing, documenting, and transferring successful strategies and correcting unhelpful ones. To begin constructing a theory that describes such an analyst workflow, we selected an accessible community of analysts working in public---baseball analytics, also known as sabermetrics--- and recruited 10 analysts who use baseball as their subject on inquiry. With each participant, we interviewed them about the resources by which they learned how to practice analytics and about how,

from start to finish, they approach the typical analytic problem. After this, we applied the methods of grounded theory to code and categorize themes in the interviews and recorded the prominent themes that describe common elements of their processes.

58. Detecting and Combating Social Media Influencers

**Anthony Weishampel,
William Rand**

Social Media provides a popular platform for marketers and organizations to diffuse content. The potential fear of malicious influencers (e.g., bots, trolls, and extremists) spreading falsehoods or attempting to radicalize the general public on social media has been validated in recent years. Some of this influence is perpetrated not directly by humans but by social media bots, which are artificial agents whose behavior is controlled by predetermined algorithms, often in an attempt to imitate authentic users, and have a particular effect on social media public discourse. This project consists of two parts. The first part of this project is to identify methods to detect the social media bots. The second part consists of an experiment to determine how legitimate user's behaviors change when informed that they have been interacting with a bot. The experiment involves reaching out to a select number of users, who have interacted with one of the detected bots. We will use a technique called causal state modeling (CSM) (sometimes called computational mechanics) to model the users' behaviors. By modeling the behavior of each user from individual social data traces, the core CSM approach is well-equipped to capture heterogeneity among individuals. These models have the advantage that they are efficient to construct and maintain and can scale to tens or hundreds of thousands of social traces, and our previous work has shown that this approach is quite accurate at predicting future behavior. Comparing the CSMs of the users before and after our interaction will allow us to determine any changes in the behavior of the users. Functional Data Analyses (FDA) methods were also used to detect the Bots. FDA models the users' behaviors through a linear combination of functions. Both of these methods proved to be successful in detecting the automated users.

59. Hybrid Methods for Estimating Regression Coefficients in Networks with Latent Community Structure

**Heather Mathews,
Alexander Volfovsky**

Networks allow us to investigate connections between people and what drives those connections. For example, we might be interested in what impacts a students' ability to make friends. To adequately model impacts of covariates on a network, we must account for the dependencies in our data and potential latent structures (here, specifically latent community structure). Given a network, Y , and covariates, X , can we make inference about effects of X on connections in Y ? We aim to efficiently estimate latent community structure that may be influencing effects of covariates. We condition on these structures since nodes belonging to different communities may be impacted by a particular covariate in different ways. Here, we introduce a model in which we allow for estimation of coefficient effects based on estimated latent community structure. When community structure is present, allowing for community dependent covariate coefficient estimation leads to more interpretable and meaningful inference as we can determine specific effects of covariates on particular communities.

60. The Geometry of Community Detection

**Vaishakhi Mayya,
Heather Mathews,
Ricardo Batista,
Jingwen Zhang,
Alexander Volfovsky,
Galen Reeves**

The problem of community detection is to partition a network into clusters of nodes (communities) with similar connection patterns. Specific examples include finding like-minded people in a social network and discovering the hierarchical relationships in organizations from observed behavior. A major limitation of

the current analysis of community detection is that it is relevant only to networks exhibiting high levels of homogeneity or symmetry. While the theory provides initial guidelines for how much data one needs to collect, it fails to describe the performance one expects to see in practice. Particularly in settings where individuals belong to multiple communities, there is high variability in the size of the communities, and there is additional covariate information. The contribution of this work is to study a much broader class of network models in which there can be high variability in the sizes and behaviors of the different communities. Our analysis shows that the performance in these models can be described in terms of a matrix of the effective signal-to-noise ratios (SNRs) that provides a geometrical representation of relationships between the communities. This analysis motivates new methodology for a variety of state-of-the-art algorithms, including spectral clustering, belief propagation (BP), and approximate message passing

61. Serious Fun: Can We Tune an Agent Based Model to Mimic Human Decisions?*

**Michele Kolb
Sarah Margaret Tulloss**

The LAS Anticipatory Thinking and RISSC exemplars teamed up in the latter part of 2018 to consider two questions: How might an agent based model perform vis-a-vis humans playing a game with the same parameters? And can human decision processes during a game be recorded and used to tune an agent based model? To that end, the team has developed an early prototype of a tabletop game to use to capture individual and group decisions within a set of parameters. The agent based model will be designed with the same parameters the game uses. You can see the alpha prototype game and even play-test it at our table today!

62. Propaganda Detection on Twitter

**Khuzaima Hameed,
Paul Thompson,
Eric Laber**

We perform exploratory analysis to inform features to be used in a classifier of propaganda on Twitter. We perform clustering with various Tweet and User attributes, then compare cluster text contents as well as their sentiment and polarity distributions. In addition to exploratory analysis, we have built a website to facilitate analyzing labeled data.

63. Intelligence Augmentation

**Michael Kowolenko,
Kenneth Thompson,
Felecia Vega**

The focus of our activities has been to leverage the use of Unstructured Text Analytics to capture facts from text as they relate to domains under investigation by a given analyst. We posit that this approach will allow the analyst to perform higher level tasks of abstract relationship mapping while the machine does the lower level tasks of data abstraction and isolation. The system is designed to leverage rules-based data extraction coupled with machine learning based on Part-of-Speech phrase extraction leading to both enhancement of the rules-based system as well as classification criteria. We are currently testing the system on curated data available through the Western Jihadism Project at Brandeis University, Waltham, MA, provided by Jytte Klausen. This database of terrorist activity was derived from open-source news articles and databases. The rules-based algorithms were derived from the database codebook and have been executed against free text fields in the database and will be compared to specific entries coded in the database for accuracy and precision. Once suitable results are obtained, the algorithms will be deployed against the original news articles for similar evaluation. In addition, the classification system will be evaluated to determine if additional information can be extracted from the news articles that was not captured in the database.

64. LAS Overview

LAS Staff

This poster contains an overview of the LAS and details on how to work with LAS.

Notes: