2018 Collaborators' Day

Dr. Alyson Wilson, Principal Investigator
agwilso2@ncsu.edu

Dr. Matt Schmidt, Director of Programs
mcschmid@ncsu.edu

June 20, 2018

Who We Are

What is LAS?

LAS is a mission-oriented research partnership between the Intelligence Community and NC State University to work at the intersection of technology and tradecraft.

What We Do

- We invest in core research challenges ("challenges") that are applied in mission-relevant contexts ("applications") to advance the science and practice of intelligence analysis.

- We develop scalable models, methodologies, technologies, and tradecraft that can be broadly applied.

How We Do It

- Mission-relevant projects
- ~90% of our work is unclassified
- Integrated, team-based approach
- Guidance is intentionally open-ended, as we are expecting you to help shape the direction of the research

LAS By The Numbers

- 40 government staff from research, mission, and technology

- 29 faculty and 40 students from 9 universities (60% STEM, 40% humanities and social sciences)

- 9 industry partners

2019 Challenges

- Human-Machine Collaboration
- Integrity for Analytic Methods
- Forecasting and Anticipation
- Analytics, Artificial Intelligence, and Machine Learning

2019 Application Areas

- Cybersecurity and Insider Threat
- Industrial Internet of Things and Critical Infrastructure
- Computational Social Science
- Emerging Technology

Call For Projects – Process Overview

- In order to propose work, you must submit a whitepaper
- If selected, you will be a subcontractor to NC State University
- Period of Performance will be January 1, 2019 – December 31, 2019

Whitepaper Submission Guidelines

- You may submit more than one
- You are encouraged to submit team whitepapers with more than one performer
- Each whitepaper should be no more than 3 pages
- Should NOT contain classified, proprietary, or sensitive information of any kind
- Submit to **https://las-whitepapers.oscar.ncsu.edu/**

Due NLT July 20, 2018

Whitepaper Submission Guidelines

Your Whitepaper Must Contain the Following:

- Name, affiliation, e-mail, phone number, and website for Principal Investigator and any additional team members
- Primary point of contact for the work proposed, if different from the principal investigator
- Total budget requested
- Indication of challenge and/or application where your work is relevant
- *Description of proposed effort, approach, how it supports areas of interest, and the specific deliverables you expect from your work*

Whitepapers – University Budgets

For Academic Partners:

- One month of summer salary support or academic release and a 12-month graduate student, plus $3,600 in other direct costs
- Can be used for post-docs, undergrads, etc… but must stay within total budget
- You may submit up to three additional scope options at the level of one additional graduate student each

Whitepapers – Industry Budgets

For Industry Partners:

- $250k or less (includes all direct, indirect, and ODCs)
- Can submit up to three additional $100k options, with appropriate scope delineated

If these levels of effort do not seem appropriate to the work you would like to propose, please contact Dr. Matt Schmidt, LAS Director of Programs, mcschmid@ncsu.edu, to discuss other options

Call for Projects – Key Dates

Key Planning Dates:

- July 20 – White Papers Due
- September 14 – Preliminary Notification of Accepted Whitepapers
- October 15 – (Tentative) Final Confirmation
- January 1, 2019 – Project Launch

Q&A

We invite your questions about challenges and applications.

2:30pm-4:00pm Duke Energy Hall

Duke Energy Hall

1. Working with LAS

<u>Challenges</u>

2. Human Machine Collaboration
3. Integrity for Analytic Models
4. Forecasting and Anticipation
5. Analytics, Artificial Intelligence, Machine Learning

<u>Applications</u>

6. Cybersecurity and Insider Threat
7. Industrial Internet of Things and Critical Infrastructure
8. Computational Social Science
9. Emerging Technology

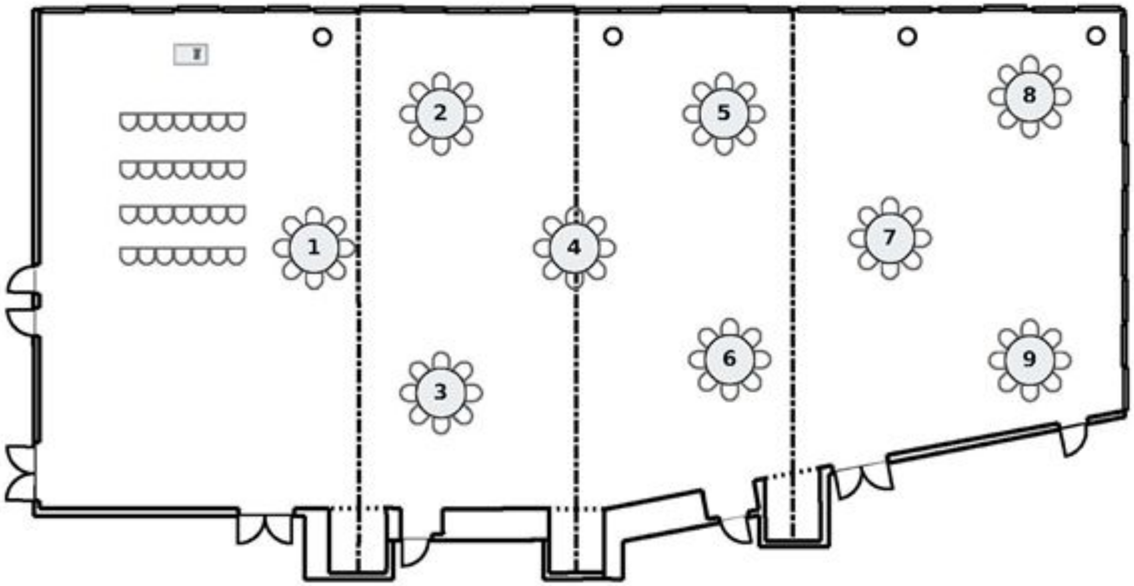

Human-Machine Collaboration

- Analytic workflows are a hybrid of manual analysis and automated analytics
- Typically, analysts view automated analytics as tools for their manual analysis
- This leads to a delegative model of coordination, which can be limited and inefficient

Human-Machine Collaboration

- We want analysts to view automated analytics more as collaborators then as tools.

- Features of collaborators
  - Proactive
  - Additive
  - Shared Understanding

- We want to enable automated analytics to proactively compute and integrate their results into the analyst's workflow

Human-Machine Collaboration

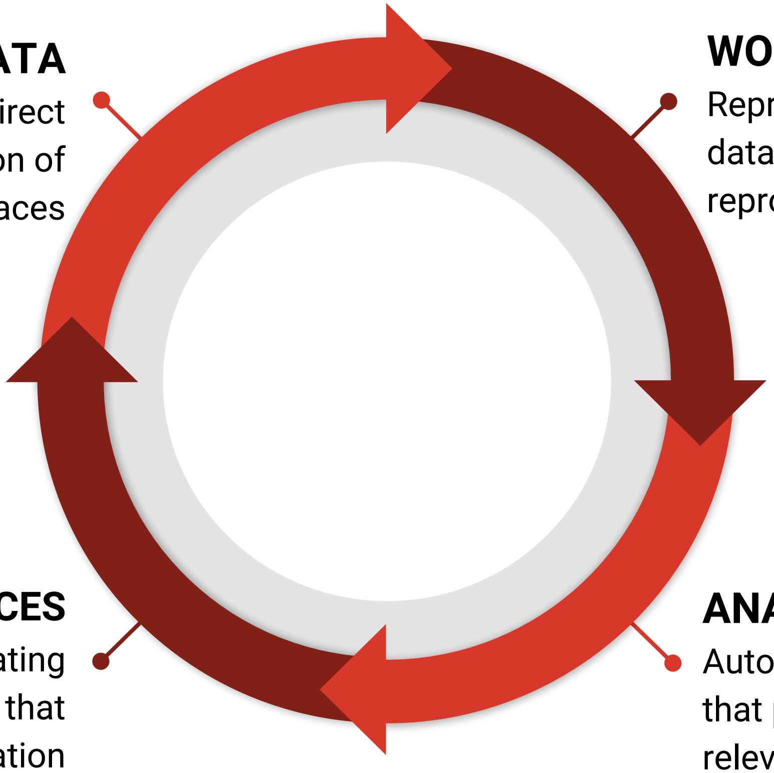

DATA

Information collected through direct observation or instrumentation of analyst tools and interfaces

WORKFLOWS

Representations of analytic workflow data used to assess, analyze, and/or reproduce the workflow

ANALYTICS

Automated workflow components that perform computational tasks relevant to the analyst’s workflow

INTERFACES

Methods of triggering, communicating and integrating analytic results that enable human-machine collaboration

Human-Machine Collaboration

- *Generating Data.* Can instrumented tools, and interfaces automatically gather data on analyst workflows and objectives?

- *Representing Workflows.* How should the we represent workflow data? What information is being communicated between manual and automated components?

- *Developing Analytics.* What are the right tasks for the machine to take on? Can we identify points in those workflows that are appropriate for technological interventions?
    - See also the Analytics, Artificial Intelligence, and Machine Learning Challenge.

Human-Machine Collaboration

- *Personalized Analytics.* Can we better understand the “science of personalization”? Can we better design analytics and models by tailoring them to an analyst and/or their task?

- *Collaborative Human-Machine Interfaces.* How does the human-machine interface support human-machine collaboration? Do collaboration methods between humans translate to collaboration between machines and humans?

Integrity of Analytic Methods

- The term *analytic methods* encompasses both both manual techniques and automated analytics

- The manner and context in which analytic methods are applied can introduce complications that negatively affect the resulting product

Integrity of Analytic Methods

- We want to enable and develop analytic tradecraft that provides high-confidence in analytic products in the face of these potential complications

- This is a challenge for all analytic methods, but has increased relevancy as methods from AI/ML are further integrated into analyst’s workflows

Integrity for Analytic Methods

- *Machine Learning / Statistics.* How do we ensure that the mathematical models learned through machine learning techniques are accurate, unbiased, and up to date? How can we ensure that adversaries are not able to manipulate or defeat the learned models?

- *Knowledge Infrastructure.* How can we scale the development, training, and evaluation of accurate and relevant models?

- *Privacy.* How can we provide privacy protections while enabling the large-scale application of machine learning techniques?

Integrity for Analytic Methods

- *Cognitive Science* - How do we protect against the bias of subject matter experts' mental models?
- *Critical Thinking* - What tools, techniques or training could be effective at integrating critical thinking and analytic rigor into an analyst's tradecraft?
- *Analytic Rigor* - How can we measure the effectiveness of tradecraft at ensuring or encouraging analytic integrity?
- *Psychology* - How do we build trust in the predictions of machine learned models?
- *Communication* - Are there ways to provide clear, accurate assessments of the confidence of the output of both manual and automated analytic methods?

Forecasting and Anticipation

The challenge of forecasting and anticipation is to improve analysts' abilities to both *anticipate* future scenarios and *forecast* their likelihood.

This can take the form of *foresight* (identifying emerging issues), *forecast* (developing potential scenarios), or *warning* (looking out to various time horizons).

Forecasting and Anticipation

- Develop quantitative methods and data analysis techniques
    - Can we develop tools and methodologies to support future-oriented analysis?
    - How should data and analytic results be integrated into these tools and methodologies?
    - What automation and analytics can enable future-oriented analysis?
    - How can we use available data to generate timely forecasts for well-defined future events and their characteristics (i.e. who, what, when, where, and how)?
    - How should we update these forecasts based on either newly discovered or streaming data?

Forecasting and Anticipation

- Define and improve anticipatory thinking skills
    - How can we train and develop analysts' skills to think about the future?
    - What are the underlying cognitive mechanisms used to think about the future?
    - How can decision aids support creating mental models of the future?
    - How can we measure an individual's ability to think about and anticipate future scenarios?
    - How do we measure the success of future-oriented tools and methodologies?

- How should we communicate the output of future-oriented analyses to decision-makers and other analysts?
    - How can we communicate uncertainty in these forecasts?

Analytics, Artificial Intelligence, and Machine Learning

LAS is interested in developing and deploying new methods, techniques, and tools for analyzing data in the context of mission-relevant applications.

Of particular interest are questions that arise in the context of “4 Vs of big data”: volume, velocity, variety, and veracity.

Analytics, Artificial Intelligence, and Machine Learning

- *Triage of an unstructured information corpus*, to include topic identification, document clustering, and automatic summarization. How can analysts understand increasingly large sets of data and rapidly triage the data to find the information they need?

- *Sensemaking* is the collection and organization of information for deeper understanding to facilitate insight and subsequent action.

Analytics, Artificial Intelligence, and Machine Learning

- *Inference and uncertainty quantification for heterogeneous information*. How do we combine heterogeneous data to characterize behavior and understand the uncertainties of our inferences?

- *Visualization*. What are the most effective visualization strategies for particular analytic tasks? What individual user characteristics make visualizations effective for particular analysts?

- *Anomaly detection*. What methods are effective for anomaly detection in heterogeneous data, spatio-temporal data, and network data?

Cybersecurity and Insider Threat

- Great progress has been achieved in the areas of detecting cyber attacks, designing defensible networks, and maintaining situational awareness of network activity.

- However, at its core cybersecurity is still essentially *reactive*

Cybersecurity and Insider Threat

- We are interested in moving beyond the reactive to enable scalable cyber defense without increasing the number of analysts.

- Motivating Challenges:
  - How do we fundamentally change the playing field, to make it even – or even give defenders the advantage?

  - How do we increase the odds of detecting malicious activity, while decreasing the time and effort required to do so?

Cybersecurity and Insider Threat

- ***Intersection with Forecasting and Anticipation***
  - *Moving beyond anomaly detection:* Can we anticipate or predict a cyber event or adversary action through the analysis of data sources not typically used in cybersecurity?

- ***Intersection with Human-Machine Collaboration***
  - *Scaling through technology-enabled tradecraft:* Can we develop techniques that make use of “big data” to improve performance and accuracy -- absent expectation that an analyst can sift through it -- to optimize human-machine interactions bounded by time?
  - *Technology-enabled structured analytic tradecraft environment:* Can we make use of structured cyber knowledge bases to enhance automated and human collaborative analysis?
  - *Managing and modeling cyber analytics:* Can an analyst request network sensor data related to a given intrusion, vulnerability, or technique and get everything relevant?

Cybersecurity and Insider Threat

- ***Intersection with Integrity for Analytic Methods***
  - *Incorporating disparate sources:* Can we significantly reduce the manual processes involved in correlating intrusion detection system alerts with text reports of intrusion activity? Can open source knowledge be used to correlate seemingly unassociated events? Is it possible to enhance cyber defense success without examining packets on a “wire”?

- ***Intersection with Analytics, Artificial Intelligence and Machine Learning***
  - *Robust cyber indications and warning:* Are there better techniques to triage, prioritize and discover threats and indications in time-bound, context sensitive environments? Can we provide related (possibly “non-cyber”) context to alerts?
  - *Insider threat detection:* Can we model individual user behaviors, identifying both outliers and significant changes? Are there techniques to find behavioral evidence that will lead to discovery of a prior compromise in a system or network? Can we anticipate threat behavior with enough confidence to take action?
  - *High-confidence attribution:* Are there reliable methods to attribute and track the phylogeny of malicious code? What contextual information is of most beneficial to attribution?

Critical Infrastructure and Industrial Internet of Things

- Critical infrastructures affect all areas of daily life, and it is a national priority to keep this critical infrastructure robust against disruptions from either unexpected situations or malicious attacks.

- We are focused on understanding two particular aspects of the critical infrastructure problem: the *interdependencies within the nation's critical infrastructure* and the security implications of integrating *Industrial Internet of Things* (IIoT) devices.

Critical Infrastructure and Industrial Internet of Things

- *Interdependency of Critical Infrastructures.* The nation's critical infrastructure is highly interconnected and mutually dependent in complex ways. Identifying, understanding, and analyzing such interdependencies is essential to assessing the robustness of the overall infrastructure.

- *Industrial Internet of Things (IIoT).* Many Industrial Control Systems users are transitioning to architectures that integrate cyber physical systems with the Internet of Things and cloud computing services. These IIoT architectures provide better visibility and insight, but have the potential to introduce numerous security challenges.

Critical Infrastructure and Industrial Internet of Things

- Example Questions: (potential intersections with Challenges in **RED**)
  - **[F&A]** How do we reduce the security risks to critical infrastructure that migrates to an IIoT architecture? What data/information is needed to better anticipate and avoid surprises?
  - **[HMC]** How can we provide a common operating picture between the cyber and physical environments? What about across interdependent sectors?
  - **[ANALYTICS]** How can we help decision makers quickly identify threats and define problems in interdependent critical infrastructure systems?
  - **[ANALYTICS]** How can we analyze and understand the large volumes of data/information generated by IIoT networks to provide a common operating picture of the environment?
  - **[ANALYTICS]** Can we use the communication behavior of devices to understand more about the physical environment?

Computational Social Science

LAS is interested in studying social phenomena of national security interest. Of particular interest are those projects where there are social science questions that can be effectively studied by applying computational methods to model, simulate, and analyze phenomena of interest.

Computational Social Science

***Impact of Instability*.**
Fragile and failing states frequently engage in armed conflict for power, and this conflict is often underpinned by illicit economies connected to global illicit networks.

Analytics, Artificial Intelligence, and Machine Learning

- Can we understand the relative importance of factors that predict the impending collapse or deterioration of a fragile or failing state?
- How well do machine learning models trained on historical events predict failing states?
- Can critical paths be identified in clandestine (non-attributable) and covert (mis-attributable) networks?

Computational Social Science

Adverse Influence.
Human beings are influenced every day from every angle. Examples abound, from Russian activities in Ukraine to the influence mass media, social media-based “fake news” or cyber campaigns had on recent US and European elections.

3 Types of “Fake News”

Fabrication
An intentional lie that doesn’t usually go beyond one source. The source is probably aware that the story is false. Depends heavily on clickbait. Think of these like an evolution of fake tabloid stories.

Hoax
Uses more sophisticated methods of fooling an audience, like forged evidence or social media manipulation. Often spread by multiple sources, some of which may believe the story is true.

Satire
A false news story that the source presents as true as a joke. When satire is shared with people that aren’t familiar with the source, there’s always a chance someone will think it’s real!

Source: Rubin, Victoria L., et al. “Deception Detection for News: Three Types of Fakes.” Proceedings of the Association for Information Science and Technology, vol. 52, no. 1, Jan. 2015, pp. 1–4.

Integrity of Analytic Methods

- What are scientifically rigorous methods to measure the effects of direct and indirect influence?
- What methods and tools are effective for identifying key direct and indirect influencing strategies for high-stakes decision-making such as those used in marketing, lobbying, fundraising, recruiting, propaganda and information campaigns?

Emerging Technology

- The emergence of a new technology can have an unforeseen impact in seemingly unrelated areas

- Examples of emerging technologies
  - Anonymization services and encryption
  - Blockchains, cryptocurrencies, and smart contracts
  - Data brokerages and data aggregation services
  - Drones
  - Some application areas discussed previously (Cyber, IIoT, AI/ML, etc.)

Emerging Technology

- Understanding the potential impact of an emerging technology is a typical intelligence need

- Example questions: (potential intersections with Challenges in **RED**)
  - What is the current state of the field? **[ANALYTICS, I4AM]**
  - What are some of the risk and opportunities that may result? **[F&A]**
  - What are the security implications and second-order effects that may result? **[F&A]**
  - What might have to adapt to these new technologies? **[F&A]**

LAS is interested in approaches that can help answer these questions an emerging technology domain

Emerging Technology

- Regardless of the specific technology, many of the difficulties in answering the questions are similar across technologies
  - Lack of consensus / authorities in the data
  - Rapidly changing state of the art
  - Lack of initial analyst familiarity and expertise in the domain

- Example general questions: (potential intersections with Challenges in **RED**)
  - What does a repeatable process look like for analysts and machine learning methods to bootstrap learning about a new domain? How can automation help? **[HMC, ANALYTICS]**
  - Can we quickly build a robust knowledge base for an emerging technology domain? How do we keep this knowledge current? **[ANALYTICS, I4AM]**
  - How should we handle the bias and veracity issues that arise in domains where there is a lack of consensus or authority? **[I4AM]**

LAS is interested in approaches that can address these questions for emerging technology domains, *in general*

Q&A

We invite your questions about challenges and applications.

2:30pm-4:00pm Duke Energy Hall