



CALL FOR WHITE PAPERS FOR 2021 LAS ACTIVITIES

June 15, 2020

1.0 Executive Summary

The Laboratory for Analytic Sciences (LAS, <https://ncsu-las.org>) is an academic-industry-government partnership that works at the intersection of technology and tradecraft. LAS participants collaborate on projects that entail research, development, and validation of broadly applicable analytic approaches that have an operational impact in the U.S. Intelligence Community (IC).

LAS is now calling for white papers from potential academic and industry partners interested in working with LAS in 2021. Your white paper should reflect your individual interests and propose potential research and development efforts that could be funded by LAS in 2021. Decisions about funding will be based, in part, on how well your interests and proposed efforts can be integrated with those of LAS staff and IC stakeholders. See [Section 2: Projects and Project Teams at LAS](#) for more details about how projects and project teams work at LAS.

The main purpose of this document is to provide you with descriptions and examples to help you determine how your interests align with those of LAS. This list is not exhaustive, and you are encouraged to submit white papers with project ideas not explicitly called for in this document. In order to provide you with more opportunities to discover areas of mutual interest, we will describe LAS interests in both research and operational terms. [Section 3: Areas of Research Interest](#) describes research questions and technical areas that are applicable in many different IC mission contexts. [Section 4: Areas of Operational Interest](#) describes use cases and application areas that may leverage many different areas of research.

We will provide multiple opportunities to engage with LAS staff prior to the deadline for submission to discuss your interests in greater detail. The primary venue for engagement with LAS staff will be the LAS Virtual Collaborators Week, which will be held June 15-18, 2020. Due to COVID-19 precautions, this event will be held entirely online (see [Section 5: LAS Virtual Collaborators Week](#)). During the week, you will have the opportunity to learn more about the interests described in this document and engage with LAS staff to discuss how your interests might align.

This document provides guidance and important dates for submitting your white papers. **White papers describing your interests and potential projects will be accepted through July 17, 2020.** Please see [Section 6: Submitting a White Paper](#) for details on how to submit a white paper.

2.0 Projects and Project Teams at LAS

LAS interprets its role as a mission-oriented research lab as supporting the maturation of IC mission-relevant technology and tradecraft. We execute this role through projects that produce tangible deliverables that enable progress towards operational solutions. Funded efforts at LAS have a 12-month period-of-performance that runs from January 1 to December 31, with the potential for follow-on funded efforts in subsequent years. Each effort is expected to produce at least one (and hopefully more) tangible deliverable during the period of performance. A deliverable can take various forms, and include items such as research papers, documented workflows, storyboards, training aids, implemented technological components, and experimental prototype systems. For examples of efforts and outcomes that LAS has supported in the past, please visit our website at <https://ncsu-las.org/about/accomplishments/>.

Projects are the main vehicle through which collaboration among academia, industry, and government personnel occurs at LAS. At LAS, we pride ourselves on our model of “immersive collaboration,” where project teams, including funded performers and LAS staff members, work toward common mission-relevant outcomes. Project teams are formed after your white paper has been selected but before the start of the period of performance. While you will be ultimately responsible only for your own individual deliverables, you should be prepared to work closely with your project team to refine, execute, and integrate each other’s efforts.

All LAS projects are required to be mission-relevant, which means they consider questions and outcomes of interest to the IC. This requirement is motivated by the LAS objective of transitioning deliverables to our stakeholders for further development into operational solutions. However, mission relevance can take on different forms depending on the maturity level of an effort’s intended deliverables. For example, all of the following efforts would be considered mission-relevant:

- A research effort designed to answer fundamental questions raised by a particular use case of interest to the IC.
- A design effort to conceptualize a novel analytic approach for achieving an outcome of interest to the IC.
- A proof-of-concept effort aimed at implementing an analytic approach on proxy datasets that serve as stand-ins for IC analogs.
- An evaluation effort focused on investigating the expected costs and benefits of a novel analytic approach in an IC-proximate environment.

Though many of the missions of LAS stakeholders are classified, projects and teams are organized so that most of the work done by funded performers is unclassified. Cleared LAS staff members involved in the project teams bring with them a reservoir of knowledge and understanding of the IC mission context, which is useful for guiding projects to mission-relevant outcomes. For specific projects, funded performers with appropriate clearances (minimum TS//SCI) can be helpful, and active or eligible clearances should be noted in the submitted white papers.

3.0 Areas of Research Interest

One way to organize the interests at LAS is around broad themes that are applicable across a multitude of IC mission contexts. This section describes four research themes of interest to LAS in 2021.

3.1 Analytic Rigor and Performance

The IC aims to achieve analytic rigor in its products and processes ([ODNI, 2015](#)), but this standard can be challenging to realize in practice. This is in part due to the diverse and varied nature of intelligence work, which makes the possibility of a comprehensive evaluation criteria for rigor challenging. It is perhaps also a function of analytic rigor itself, which is a somewhat nebulous concept that can be difficult to pin down. Indeed, some have argued that analytic rigor is not a prescriptive set of rules to be followed, but rather a credible assurance that one has minimized the risk of reaching conclusions prematurely by probing to a sufficient depth of inquiry within a particular context ([Zelik et al., 2010](#)).

LAS is interested in understanding how IC analysts and analytic teams can obtain such “assurances of sufficiency” with confidence. What can be done to ensure that intelligence work products and processes are analytically rigorous? And on what basis should the IC assess the performance and effectiveness of analytic teams that are producing and contributing to these products? Building upon its previous work in this domain, LAS is interested in projects in the following areas:

- **Analytic production:** A number of new technologies are reshaping the landscape of contemporary journalism and news reporting.¹ We seek to understand and demonstrate the potential implications of these technologies for the publication, dissemination, and sharing of IC work products. We are also interested in exploring how contemporary journalism is grappling with the challenge of mitigating reporter and reader biases amidst its adoption of these technologies.
- **Analytic workflow:** Much of the IC emphasis on rigor focuses on published analytic work products. There is likewise a need to ensure that analytic processes and practices (i.e., the procedural “workflows” that analysts follow in searching for, filtering, and making sense of information) are rigorous. We seek to incorporate the evaluation of analytic rigor into intelligence and language analyst workflows, as well as explore how concepts like multi-factor optimization² can be productively applied to the creation and execution of workflows.

¹ This is seen, for example, in the increased use of audio and voice-activated technologies, as well as adoption of artificial intelligence and machine learning to provide a more personalized experience for readers, craft more compelling and engaging content, and support reporters who are suffering from information overload ([Newman, 2019](#)).

² Also known as multiobjective optimization, this is the problem of an end user (i.e., an analyst) pursuing a goal that entails minimizing and/or maximizing multiple (often conflicting) functions at the same time. We believe this concept has promising applications to analytic rigor, given the competing goals that analysts are expected to satisfy and balance (e.g., timeliness, thoroughness, accuracy, etc.).

- **Analytic performance:** LAS has previously leveraged insights from industrial and organizational psychology to develop mechanisms for empirically measuring the performance of analytic teams. We seek to extend this work into the area of *augmenting* analytic performance—specifically by testing whether proposed “improvements” work as expected for actual analytic teams. We are also interested in exploring the relationship between rigor and performance, to include establishing the degree to which the latter is dependent upon the former.
- **Rigor in language analysis:** The job of an IC language analyst entails using knowledge of foreign language and English to provide key decision makers an understanding of the nuance, context, cultural overtones and dialect in which a particular communication occurs ([NSA, n.d.](#)). In 2021 we seek to investigate what rigor means—and operationalize it—within a language analysis context. We specifically seek to explore alternatives for achieving rigor in language analysis beyond the conventional approaches of performing post-hoc quality control and quality review checks on foreign language transcriptions and translations.

3.2 Machine Learning Integrity

The IC is increasingly using artificial intelligence (AI) as a means of coping with the vast, disparate, and dynamic data that it collects and processes ([ODNI, 2019](#)). Machine learning (ML) technologies, in particular, show great promise in helping the IC understand and adjust to the changing nature of threats and adversaries. This constitutes a dramatic shift in how the IC develops analytics—from a reliance on static rules and logic that are supplied by humans to increased use of models that automatically generate and refine their own rules through continuous training ([Heckman, 2020](#)).

A continuing area of research for LAS in 2021 is Machine Learning Integrity (MLI)—an interrelated set of challenges that must be tackled in order for the IC to realize the promise and potential of ML technology. Building upon its previous work in this domain, LAS is interested in projects in the following areas:

- **Operational Machine Learning:** The growing adoption of ML technology in the IC has highlighted a need for more thoughtful and deliberate approaches to deploying, maintaining, monitoring, and safeguarding ML models. Against this backdrop, LAS has particular interest in the following challenges:
 - **ML DevOps:** Continuously training and transitioning ML models and workflows from development environments (e.g., in notebooks) into production ML systems.
 - **ML Model Maintenance:** Maintaining, monitoring, and ensuring the quality of operational ML models—for example, by enabling users to provide feedback on model predictions, identifying opportunities for improved performance, and detecting and mitigating common issues (e.g., errors made during model development, data drift, and changes in data availability).
 - **Model Assurance:** Ensuring the integrity of the data upon which operational ML models are trained—for example, by detecting when adversarial “poisoning” has occurred or when mislabeled data is affecting a ML model’s performance.

- ML Human Factors: Determining the performance metrics and thresholds required of a ML model in order for a human analyst to adopt a workflow that incorporates the model's outputs.³
- **Efficient ML Development**: As the IC accelerates its incorporation of ML models into more analyst workflows, it follows that these models will become increasingly tailored to niche use cases and datasets. Thus, while the overall impact of ML in the IC may increase over time, the impact of individual models could decrease. The IC must weigh the reduced marginal benefits of future ML models against the non-trivial costs of a typical ML model development cycle, such as selecting, training, and refining models created by data scientists. Against this backdrop, LAS seeks to explore ways of decreasing the cost of the ML development cycle with particular interest in the following challenges:
 - Model Selection Support: Reducing the complexity of model selection, which often entails choosing class and hyper-parameters of models and fitting them to the available training data, without sacrificing rigor—for example, through intuitive tools and visualizations, model recommendations, or full automation of model selection.
 - Label Reduction: Reducing the effort needed to generate sufficient labeled data to train a ML model and evaluating the effectiveness of these approaches.
- **Reinforcement Learning**: Reinforcement learning is an approach that has demonstrated value in automatically identifying good decision strategies in a variety of domains, including gaming, traffic signal control, and robotics ([garychl, 2018](#)). These successes are typically the result of a careful framing of the problem in terms of actions, environments, states, and rewards. LAS seeks to explore ways of framing IC problems, such as human workflow optimization and data / systems management, as reinforcement learning problems, thereby enabling the application of these state-of-the-art approaches to mission challenges.

3.3. Human-Machine Collaboration

Artificial intelligence and machine learning cannot be implemented in a vacuum. For such technologies to make a difference, their powerful capabilities and “formal rationality” must be paired thoughtfully with human domain expertise and “substantive rationality” ([Saenz et al., 2020](#)). Central to this vision is the idea of human analysts and machines working together as trusted teammates in service of shared analytic objectives ([ODNI, 2019a](#)). Within such teams, machine agents monitor and learn what their analyst teammates are trying to accomplish, identify potentially beneficial workflow interventions (e.g., by suggesting or providing relevant data, tools, or visualizations), and independently perform relevant work for the analyst. Building upon its previous work in this domain, LAS is interested in projects in the following areas:

³ For example, a tool that automatically identifies objects of interest in video may allow analysts to more efficiently process video content than manual-intensive approaches. At some point, however, correcting for inaccuracies and false positives / negatives produced by the object detection methods could cause the overall workflow to be less efficient for analysts than relying on the aforementioned manual-intensive methods. We are interested in methods that could help identify this threshold.

- **Visualization:** LAS seeks to devise, implement, and demonstrate novel visualizations that enable mission analysts to optimally view data, as well as present information more effectively to consumers of intelligence reports. We also seek methods to empirically validate whether novel visualizations provide an advantage over currently available and used visualizations.
- **Workflows and Interventions:** Workflow interventions require machine agents obtaining a nuanced “understanding” of analyst’s behavior, intent, and motivation. LAS seeks to demonstrate methods of inferring these from data sources that capture or proxy IC analyst activities⁴. We also seek to explore methods that use this understanding to optimize machine interventions and analysis outcomes.
- **Task and Tool Recommenders:** LAS seeks to design and develop prototypes that provide analysts with tailored recommendations for their workflow—for example, on the next step(s) they could take within a given workflow, or the optimal visualizations for the data with which they are currently working.

3.4 Triage

In assessing the contemporary strategic environment, the ODNI notes how “[a]dvances in communications and the democratization of other technologies” have fueled an explosion of information and data that the IC must “collect, process, evaluate, and analyze” with sufficient speed to “provide relevant and useful insight to its customers” (ODNI, 2019b). Lying at the heart of this challenge is the core function of data triage, which entails identifying the “right” data for a particular analysis. Strategies, methods, and tools are needed to explore data, to search, filter, sort, and prioritize massive data sets, and to find all information that is responsive to customers’ information needs. Building upon its previous work in this domain, LAS is interested in projects in the following areas:

- **Information Retrieval:** Search-based retrieval methods remain a vital tool for analysts to find relevant content. The value of these methods depends greatly on the ability of an analyst to use the indexed data features and metadata to find what they are looking for. LAS seeks to answer the related questions of (a) what the “right” structure is for enabling efficient search and (b) how that structure can be extrapolated at scale for different application domains. (For example application domains of interest, see the voice content / transcription analysis and handwriting OCR use cases in Section 4.2.)
- **Information Synthesis:** LAS is interested in methods that synthesize and infer information from multiple data points and organize the extrapolated information in a structure that enables higher-order analysis.⁵
- **Data Exploration:** Search-based tools for datasets are of limited use to analysts when they do not know what information is contained within the dataset—for example, the use case of captured enemy materials, where the analyst does not know *a priori* what types of files exist

⁴ Activity data sources can go beyond keystrokes and mouse movements. They could also include eye tracking, fNIRS, silent speech, etc. However, consideration must be given to the eventual feasibility of operationalizing and scaling any activity data collection on which the analysis depends.

⁵ As an example, LAS seeks to develop methods to distill information from text-based cyber vulnerability reports into a knowledge structure such as MITRE’s ATT&CK framework.

in the corpus or what information they may contain. LAS seeks to develop and demonstrate methods for surveying a given data corpus and helping analysts determine its potential information value.

- **Data Prioritization:** LAS is also generally interested in methods of prioritizing the data that the analyst might find most relevant within a large corpus. Specifically, LAS is interested in investigating how these approaches might apply to prioritization tasks that IC analysts face often (e.g., prioritization for language or voice content analysis). Specific data prioritization applications of interest to LAS are discussed in greater detail in Section 4.
- **Data Economization:** The purpose of retaining data is to address customers' current and anticipated future information needs. That said, some data may be redundant or otherwise immaterial to these needs. LAS seeks to develop and demonstrate scalable methods to understand the information needs of a given user base and apply that understanding to rigorously reduce the data retention requirements. Additionally, LAS is interested in applying the understanding of user information needs to identify new data sources that could be of value.

4.0 Operational Areas of Interest

The interests described in Section 3 cover the practice of intelligence analysis broadly and will lead to advances that are applicable across many domains. However, LAS also has interest in projects that use specific mission use cases to motivate, demonstrate, and validate these general advances. This section describes some of these use cases and application areas of interest for LAS in 2021.

4.1 Influence Campaigns

Strategic adversaries and competitors of the United States have come to view influence campaigns⁶ as an effective and low-cost means of advancing their agendas. Influence campaigns entail deception, persuasion, dissemination of misinformation, and amplification of conspiracy theories ([NCSC, 2020](#)). They are waged through a variety of mechanisms, ranging from the traditional tools of statecraft (e.g., economic and infrastructure development assistance) to the use of modern social media platforms and even the mainstream press.

Building upon its previous work in this domain, LAS is interested in projects that advance the state of the art in detecting, characterizing, and countering state-sponsored influence campaigns. Specific areas of interest include the following:

- **Indicators, origins, and provenance.** Projects will develop and test methods for identifying the origins, provenance, and key indicators of state-sponsored influence campaigns in publicly available information sources.
- **Message content.** Projects will understand how message content changes over time within influence campaigns. LAS is specifically interested in devising methods for extracting the

⁶ The RAND Corporation defines influence campaigns, also known as influence operations, as “the collection of tactical information about an adversary as well as the dissemination of propaganda in pursuit of a competitive advantage over an opponent” ([RAND, n.d.](#)).

core “talking points” and inflection points of campaigns, as well as determining and detecting how specific tactics (e.g., the use of memes, doctored media, and “source hacking”⁷) are deployed within them.

- **Impact and effectiveness.** Projects will develop ways to measure the relative impact and effectiveness of influence campaigns, as well as that of specific agents (i.e., influencers) for particular topics.
- **Countering malign influence.** Projects will identify and develop effective strategies and techniques for countering and neutralizing malign influence campaigns.

A common aim for all of LAS’s influence-related work is to develop intuitive, user-facing solutions (e.g., desktop visualizations) that directly enable IC analysts to pose and answer questions similar to those listed above.

4.2 Selected Cybersecurity Challenges

An additional application area of interest for LAS in 2021 relates to a subset of unclassified technical challenge problems raised by U.S. Cyber Command to “draw the attention of solution providers and focus them on [its] most pressing needs” ([U.S. Cyber Command, 2019](#)). LAS is interested in pursuing research and prototyping focused on any of the following topics drawn from the aforementioned document:

- Exploring fast, intuitive, and scalable ways to identify software security vulnerabilities and make evidence-based determinations of their exploitability
- Exploring state of the art methods for malware rapid triage, reverse engineering, correlation (e.g., grouping into families), and obfuscation
- Enabling defenders to recognize polymorphic malware in real-time at the perimeter
- Keeping code functionally persistent within network devices after measures have been taken to de-activate or disable it (e.g., system reboot)

4.3 Additional Use Cases

LAS is interested in a number of use cases that do not fall in a broad application domain like influence campaigns or cybersecurity. However, white papers that aligned with one of these specific use cases would still be of interest to LAS. These additional use cases include:

- **Prioritizing and Filtering Voice Data.** Prioritizing voice data for analysis remains a significant challenge for the IC, as intimated in Section 3.4 above. A specific use case of interest is exploring the merits and feasibility of leveraging novel filtering attributes (e.g., a speaker's tone, emotion, and sentiment) to address this challenge.⁸ As part of this investigation, we seek to obtain a corpora of voice data and work with expert language

⁷ Source hacking is “a set of techniques for hiding the sources of problematic information in order to permit its circulation in mainstream media.” Examples of source hacking include developing and testing false or inflammatory messages in one forum (e.g., an anonymized discussion group) and repackaging them for use in other forums with broader reach ([Donovan & Friedberg, 2019](#)).

⁸ Suppose, for example, that we could automatically identify all recordings in which one or more speakers expressed themselves intensely and forcefully. Could such an approach outperform conventional methods in quickly isolating high-value exchanges within large voice datasets?

analysts to prototype and apply voice filtering algorithms against it, with a view to prioritizing voice content for analysis. We are also interested in working with language analysts to evaluate the effectiveness of these prototype algorithms.

- **User-Centric Document Categorization:** Many IC analytic workflows entail distinguishing between documents with "relevant" and "not relevant" content, where relevance is a function of domain-specific and idiosyncratic criteria. While use cases such as these could likely be addressed to some degree with judiciously applied ML models, the costs of assembling a team of data scientists to create a model for each bespoke scenario would be prohibitively expensive. LAS is therefore interested in developing and evaluating methods that could enable an analyst without programming experience to easily create, train, and tune ML models to address their particular document classification needs.
- **Processing of Structured Forms:** A document's structure can reveal important information about its contents, particularly when the structure is standardized. LAS is interested in prototyping and evaluating automated methods for processing structured form documents and images. LAS would be particularly interested in methods that provide one or more of the following capabilities:
 - Identify structured form documents (and images thereof) within a large collection of heterogeneous data objects
 - Categorize identified structured form documents and images by their document type (e.g., tax form vs. voter registration form)
 - Localize and index all fields within a form to support rapid information retrieval by analysts.
- **Handwriting Recognition in Scanned Documents:** The challenge of processing document images (including, but not limited to, structured forms) is further compounded by the prevalence of handwritten notes that appear within them. Such notes are generally more difficult to read and parse automatically (e.g., via Optical Character Recognition) than typed text. Nevertheless, the handwritten notes are in many cases of greater information value than the typed text itself.⁹ LAS seeks to address this problem in 2021 by prototyping and testing methods that provide one or more of the following capabilities:
 - Identify document images within a given set that contain handwritten text
 - Locate each occurrence of handwriting within a document page
 - Use these capabilities to prioritize and filter document images for analysis.

5.0 LAS Virtual Collaborators Week

To learn more about the 2021 LAS interest areas and how to work with LAS, you are encouraged to participate in the events of the 2020 LAS Collaborators Week from June 15-18, 2020. Due to COVID-19 social distancing guidelines, this year's event will be 100% virtual and will run in segments across four days. All information about the 2020 LAS Collaborators Week, including detailed agendas for the days

⁹ Consider, for example, a senior executive's handwritten feedback appearing in the margins of a draft document, or an employee's explanatory notation on a personal expense report.

and registration forms for the webinars and office hours, will be accessible on the 2020 LAS Collaborators Week website at: <https://ncsu-las.org/2020-las-collaborators-day/>.

At 2:30p on **Monday, June 15**, the LAS Principal Investigator and Director of Programs will host a Zoom webinar for all interested potential performers. This webinar will provide an overview of LAS, further details on the research and application interest areas in this call, and logistics for submitting white papers for 2021 activities. The presentation will be given as a Zoom Webinar and recorded for later viewing. *Attendance is open to all, but you will need to register through the 2020 LAS Collaborators Week website to receive the Zoom webinar information.*

On **Tuesday, June 16**, members of the LAS Leadership team will host “office hour” sessions via Zoom to answer questions about the logistics of the white paper process, how to work with LAS, and provide general suggestions about how your research interests might align with the different LAS interest areas. *The sessions will be divided into time slots of 10 minutes each, and you will need to reserve a time slot through the 2020 LAS Collaborators Week website.*

On **Wednesday and Thursday, June 17-18**, LAS staff members and selected stakeholders will be available via Zoom to discuss each of the research and application themes described in Sections 3 and 4. These office hour sessions will provide you the opportunity to speak with LAS staff who have related interests about your project ideas and potential collaborations. *The sessions will be divided into time slots of 10 minutes each, and you will need to reserve a time slot through the 2020 LAS Collaborators Week website.*

Should you be unable to attend the webinar on June 15, we will record the presentation and make it available on the 2020 LAS Collaborators Week website shortly after the event. If you are unable to schedule an appropriate office hour time slot, please send an e-mail to lasoutreach@ncsu.edu, and NC State staff will work with you to arrange discussions after Collaborators Week.

6.0 Submitting a White Paper

If you would like to be considered for a funded project, you will need to submit a white paper to LAS for review no later than 11:59 pm EDT, July 17, 2020.

Your submission should NOT contain classified, proprietary, or sensitive information of any kind.

6.1 Scope, Schedule, And Budget

The period of performance for which we will make awards is January 1 – December 31, 2021. However, we understand that some projects may require multiple years to be fully realized. In your white paper, we would like you to submit a detailed description of the work that you would propose to undertake in 2021. You should describe your work as a project by defining a specific question or goal, an approach, and deliverables or outcomes that you would expect to be achievable by the end of 2021. For white papers whose full extent is unlikely to be realizable within a single year, you may optionally submit an additional section briefly describing potential follow-on work for 2022.

If you are submitting a white paper as a potential individual academic collaborator, please scope your base effort at the level of a month of summer salary support or academic release and a 12-month graduate student. This is our most common award. You may also choose to fund different personnel (research staff, post-docs, etc.); however, you should plan on receiving only the amount of funding that would cover a graduate student. In rare cases, we may fund additional students; you may submit up to three additional scope options at the level of one additional graduate student each. If you are submitting a white paper as part of an integrated team, each individual faculty performer can propose salary support and a student.

If you are submitting a white paper as a potential industry collaborator, please scope your base effort at \$250K or less. This should include all personnel costs, fees, other direct costs, overhead, etc. You may submit up to three additional scope options at \$100K each. This figure is meant to be a ROM estimate to allow you to scope effort. For the white paper, we do not require a discussion of specific personnel at specific rates, nor do we require a cost proposal.

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

6.2 Outline and Content of White Paper

White papers describing potential 2021 work should be no more than *two* (2) pages long. However, if the white paper contains descriptions of possible follow-on work or a capabilities statement, as described below, the white paper may contain an additional page for each. White papers should address one or more of the areas discussed in Sections 3 and 4.

LAS intends to pursue projects as integrated project teams. You are encouraged to submit team white papers that include more than one participant from academia or industry. If you have discussed your ideas with LAS government staff, you may indicate this in your white paper, but government staff should not be included as performers. There is no limit to the number of white papers you may submit, and we request that you submit multiple white papers if you are proposing multiple projects.

For each white paper, please include the following information:

- **Name, affiliation, e-mail, phone number for the Principal Investigator (PI) or co-PIs of the proposed work.** If you have more than one team member, please list information for only funded PIs.
- **Website.** Any websites provided for individual performers, labs, or organizations will be reviewed by LAS to better understand a submitter's capabilities. If you do not have a current website or if you feel it does not adequately capture the capabilities of your group, you may submit an additional brief capabilities description, as described below.
- **Primary point of contact for the work proposed,** if different from the principal investigator.
- **Title and abstract of proposed project.** Please include a 2000 character or less description of the proposed project. *The title and abstract will be widely distributed to potential government collaborators during the 2021 program development process. Do NOT include proprietary or sensitive information in either the title or abstract.*

- **Description of proposed 2021 project.** Please define a specific question or goal, an approach, how it supports key challenges and/or applications of interest, and the specific outcomes that you expect. This description may include brief scope options for additional students (academics) or funds (industry), as described previously.
- **Description of potential follow-on work (optional, additional page).** You may add an additional page with a brief description of related work that could be performed in 2022. The primary intent of this optional section is to allow potential performers to appropriately contextualize projects that are unlikely to be fully realizable within a single year. There is no requirement to describe potential follow-on work, and selection of a white paper for 2021 does not imply selection of the follow-on work.
- **Description of capabilities (optional, additional page).** You may add an additional page with a brief, paragraph-long description of the capabilities of the submitting faculty member, lab, or organization. The intent of this optional section is to allow performers to describe the broader scope of their interests and capacity, apart from the specific proposed project. There is no need to submit a separate capabilities description if this information is already adequately captured on a website provided in the main white paper body.

The white papers will be uploaded using an online submission tool (Section 6.3). Please include all of the above information in your white paper. Some of the information from the white paper (e.g., the title and abstract) will also be entered into separate text boxes in the online tool.

Please do not include budget information in the white paper that you upload. There is a text box in the online tool where information about the total budget requested should be submitted. See Section 8.1 for additional details on budgets. For academics, it is sufficient to describe the level of support, e.g., “a month of faculty support and a graduate student.” For industry submissions, a rough order of magnitude (ROM) bottom line will suffice; please do not submit a detailed budget at this time.

6.3 White Paper Submission Process

Please submit white papers to LAS for consideration by July 17, 2020 at 11:59 pm EDT. You will submit your white papers at <https://whitepapers.ncsu-las.net>, where you can find detailed instructions on the submission process. We ask that you upload your white paper to the tool as an Adobe PDF.

6.4 Review and Notification

Your abstracts and white paper will be reviewed by North Carolina State University and government domain experts, as well as by selected partners from industry and academia hired to make recommendations to LAS on which approaches will best help us achieve our goals. Decisions about funding will be based on a variety of criteria, to include the technical quality of the proposed work, the relevance of your white paper to the LAS stakeholder interests, the capabilities of the proposed performers or organizations, how well the proposed work can be integrated into project teams, and the availability of funding. While information and reviews will be gathered from a variety of stakeholders and domain experts, the decision on who receives funding rests with the LAS Principal Investigator, Dr. Alyson Wilson, and the LAS Director of Programs, Dr. Matthew Schmidt.

We will make preliminary notification of decisions on white papers by September 15, 2020. Due to the volume of white papers, we will not be able to provide feedback on white papers that are not selected. If your white paper is selected, we will work with you to integrate your work into NC State's overall portfolio for 2021. Based on the capabilities of the proposed performers and organizations, we may also suggest ways your project might be revised to align more closely with LAS interests. Final confirmation of project funding is dependent on government budget decisions, but our goal is to make final notifications of funding by November 1, 2020.

7.0 Contacts

If you have additional questions, please contact us at lasoutreach@ncsu.edu.

If you received a forwarded copy of this document, but would like to ensure you are notified about future calls and events at LAS please contact Maggie Epps (mebeasle@ncsu.edu) to be added to the distribution list.

8.0 References

- Donovan, J. and Friedberg, B. (2019). *Source Hacking: Media Manipulation in Practice*. Data & Society. https://datasociety.net/wp-content/uploads/2019/09/DataSociety_Source_Hacking.pdf
- garychl (2018, August 1). Applications of Reinforcement Learning in Real World. *Towards Data Science*. <https://towardsdatascience.com/applications-of-reinforcement-learning-in-real-world-1a94955bcd12>
- Heckman, J. (2020, March 5). *Despite 'consensus' with DoD, ODNI moving ahead with its own AI principles*. Federal News Network (FNN). <https://federalnewsnetwork.com/artificial-intelligence/2020/03/despite-consensus-with-dod-odni-moving-ahead-with-its-own-ai-principles/>
- National Counterintelligence and Security Center [NCSC]. (2020). National Counterintelligence Strategy of the United States of America, 2020-2022. https://www.dni.gov/files/NCSC/documents/features/20200205-National_CI_Strategy_2020_2022.pdf
- Newman, N. (2019). *Journalism, Media and Technology Trends and Predictions 2019*. Reuters Institute for the Study of Journalism. <http://www.digitalnewsreport.org/publications/2019/journalism-media-technology-trends-predictions-2019/#2-3-ai-and-the-newsroom>
- National Security Agency [NSA] (n.d.). "NSA Career Fields: Foreign Language Analysis." Retrieved from: <https://www.intelligencecareers.gov/nsa/nsacareers.html>
- Office of the Director of National Intelligence [ODNI] (2015). *Intelligence Community Directive 203: Analytic Standards*. <https://www.dni.gov/files/documents/ICD/ICD%20203%20Analytic%20Standards.pdf>
- ODNI (2019). *The AIM Initiative: A Strategy for Augmenting Intelligence Using Machines*. <https://www.dni.gov/files/ODNI/documents/AIM-Strategy.pdf>
- ODNI (2019). *National Intelligence Strategy of the United States of America*. https://www.dni.gov/files/ODNI/documents/National_Intelligence_Strategy_2019.pdf
- RAND Corporation [RAND]. (n.d.). *Information Operations*. <https://www.rand.org/topics/information-operations.html>
- Saenz, M., Revilla, E., and Simón, C. (2020 March 18). *Designing AI Systems With Human-Machine Teams*. MIT Sloan Management Review. <https://sloanreview.mit.edu/article/designing-ai-systems-with-human-machine-teams>
- U.S. Cyber Command (2019). *Technical Challenge Problems Guidance*. Retrieved from: <https://www.cybercom.mil/Portals/56/Documents/Technical%20Outreach/Technical%20Challenge%20Problems.pdf?ver=2019-07-02-151118-497>
- Zelik, D. J., Patterson, E. S., & Woods, D. D. (2010). Measuring attributes of rigor in information analysis. In E. S. Patterson & J. Miller (Eds.), *Macro-cognition metrics and scenarios: Design and evaluation for real-world teams*. (pp. 65-83). CRC Press.