Title: A Novel Model of Cyber Combat

Primary Author: Dr. Stephen Spey

Abstract: We have developed a novel model for cyber combat that allows us to begin exploring the complex decision space of cyber warfighting. Our model creates abstract representations of elements present in any cyber combat, such as defensive tools, offensive weapons, and aggressor teams, for two or more cyber combatants. Every combatant can attack and be attacked by every other combatant. The model sits in the middle ground between very broad high-level cyber-kill-chain models and very detailed models that simulate a single defended network against specific attacks. Each model timestep, each represented element takes actions as directed by the strategy in use by their combatant. The outcomes of elements interacting, such as when a defensive tool searches for an aggressor team that has penetrated its defended network terrain using a given offensive weapon, are resolved by comparing calculated outcome probabilities to a random number draw. Statistics of basic interactions between attackers and defenders are validated against real-world network intrusion data. Our model allows for rapid experimentation in force-level strategies and tactics in the cyber domain. We will present results for the types of targets aggressor teams prioritize and for when defenders act on detected intrusions. Other findings include how to allocate marginal additional defensive spending and relative skill levels and techniques necessary for successful long-duration network intrusions.

Title: From CNNs to Symbolic Explainable Models to Protection Against Adversarial Attacks to Automated Video Processing

Primary Author: Dr. Asim Roy

Abstract: Explainability and adversarial attacks are two major problems for deep learning CNNs. We resolve both with symbolic models created from CNNs. DARPA's Explainable AI (XAI) program conceptualized symbolic models. We present a method to decode a CNN to recognize parts of objects and extract a symbolic explainable model as desired by DARPA. For example, such a symbolic model might predict that a cat is in the image by verifying its parts - that it can see the cat's face, legs, body, and tail. We present experimental results for object detection in satellite images obtained from the xView dataset provided by the Defense Intelligence Unit (DIU). We show that it is easy to extract part information (e.g., the wings, fuselage and tail of an airplane) from a CNN model to produce a part-based symbolic explanation of objects in a scene. Overall, this method for scene analysis achieves the following: (1) improved accuracy because of identification of parts, (2) explainability based on symbolic variables corresponding to the parts, (3) protection against adversarial attacks, and (4) automated scene

analysis in real-time based on the symbolic information. Since we verify object parts, adversarial attacks cannot degrade these models. Thus, a school bus will never turn into an ostrich with a few pixel changes. Thus, we resolve DARPA's adversarial attack concerns. Also, real-time video processing - in drones, UAVs, satellites, airplanes - with transparent and trusted models is now feasible.

Title: Quantifying Non-functional Software User Requirements Using SNAP
Primary Author: Charles B. Tichenor
Abstract: QUANTIFYING SOFTWARE NON-FUNCTIONAL USER REQUIREMENTS USING SNAP

It is important to measure the size of software by quantifying the software's functional user requirements. This is at least in part because there is a statistical correlation between this software functional user requirements' size and the work effort to develop that software. The larger the application, the more it costs (and the longer it usually takes to complete). This sizing can be done using function points, which measures "what" the software delivers to its users. The function point sizing metric of the International Function Point Users Group (IFPUG) is an ISO standard. In addition, there are software non-functional user requirements. These represent not "what" the software will do but "how" it will do it. SNAP is the "Software Non-functional Assessment Process," and is sponsored by IFPUG. This metric, which is standardized by both IEEE and ISO, quantifies four categories and 14 subcategories of software non-functional user requirements using "SNAP points." The cost to develop software is too often underestimated. One reason may be that only functional, and not non-functional, software user requirements are considered. However, an extensive amount of work effort may be required to develop this non-functionality such as with algorithms, data encryption, help manuals, etc. The purpose of this presentation is to define and show how to size these categories and subcategories of software non-functional user requirements and demonstrate the value which SNAP can provide to the software cost estimation process.

Title: Quantitative Intelligence Source Uncertainty Analysis Using Multi-Objective Decision Modeling
Primary Author: Adam Nesmith
Abstract: An all-source intelligence analyst's primary job is delivering timely, well-sourced assessments on relevant targets based on uncertain and incomplete information. Each assessment includes a likelihood that the assessment is true and a confidence level based on the uncertainty of the sources used. Quantitative all-source intelligence analysis is not currently implemented despite the acknowledged limitations of qualitative intelligence assessments and the existence of proposed quantitative methods. This is due to the challenge of quantitatively representing uncertainty in human-sourced intelligence reporting (i.e., HUMINT, OSINT, SIGINT), which limits the effectiveness and usability of previously suggested methods. This presentation proposes a novel quantitative approach for individual intelligence source uncertainty by adapting decision models used in multi-objective decision analysis. This allows analysts to easily identify and mathematically account for the underlying causes of a source's uncertainty, weight the importance of these causes, and output a single value in between 0 and 1 representing the source's overall uncertainty. The analyst can then use this numerical output as an input into the previously proposed quantitative intelligence analysis methods. Ultimately, this framework for quantifying source uncertainty facilitates the use of previously proposed methods and creates more traceable and defendable intelligence assessments.

Title: Digital Engineering Support to the Joint All Domain Command and Control (JADC2) Network
Primary Author: Mr. Charles D Burdick, CAP
Abstract: To potentially support JADC2, we have introduced the concept of a Network Digital Design Twin (NDDT) that employs both the existing technology of virtually cloning a network and the emulation capabilities of the Network Digital Twin (NDT), which can predict its physical twin's responses to any changes in hardware, software, configurations, etc. as well as to malware attacks and destruction of components.

The JADC2 Network Digital Design (NDD) starts as a Clone of the design built from periodic design updates from all participating JADC2 developers. In addition to providing an interrogatable 3D visualization of the current design, the NDD Clone would maintain a complete record of design changes, each with their source and date in much the same fashion as the that used for the NDT.

The NDDT would then be built from the current NDD Clone using the same digital artifacts as for an existing network e.g., virtual hardware, and existing network software. Where physical hardware or software does not yet exist to validate its virtual twin, the NDDT artifact constructs would be based on requirements documents, prototypes, etc.

As the NDT already does for existing networks, the NDDT would provide a predictive capability allowing for tests of the design, both end-to-end and in segments not yet connected to the larger network. Since much of JADC2 already exists or will exist in prototype networks, they (or a portion of them) would be represented as NDT emulations that can be connected to the JADC2 NDDT to test their interfaces and confirm interoperability without risks to the actual physical networks.

Updates of existing physical networks can be done automatically and passively to update their Digital Clones without being on the network or needing to secure an Authority to Operate. However, combining the submitted updates to the NDD Clones will require a common set of standards and artifacts. But if existing standards and artifacts can be expanded and their use maintained, there is every hope of automating these updates from the individual JADC2 developers. The conversion of Clones to emulations (both NDT and NDDT) is already automated.

This approach would both advance the use of Network Digital Engineering and meet the Testing Community's desire to "Shift Left" allowing much earlier testing in the development cycle, uncovering potential problems, and discovering potential network improvements long before the full system is fielded.

Title: Can you DIG it?: Transparency at a Glance: A Concept for Transparency of Semi-Autonomous and Fully Autonomous (SAFA) and Artificial Intelligence (AI) Systems through Dynamic Infographics (DIG)(Patent Pending)
Primary Author: Gina Hartnett
Abstract: Autonomous and Artificial Intelligence (AI) systems are beginning to inundate nearly all military operations and can overwhelm Soldiers with vast amounts of information (and slow decision-making) during missions. The rapid and complex actions taken by these systems are often difficult for Soldiers to understand the 'why' and the 'how' a system took a specific action. Transparency of the

system is crucial to understanding the 'why' and the 'how', which will in turn increase Soldier trust and ensure effective collaboration within the human-system team. Ensuring that the human understands the behavior of intelligent systems through eXplainable AI (XAI) (Gunning 2016) highlights the importance of transparency. A model of human trust in autonomous systems, Situation Awareness-based Agent Transparency (SAT) (Chen 2014) was developed to describe the issues of user trust in these systems. This presentation will introduce a concept, Dynamic Infographics (DIG), which seeks to provide user transparency of autonomous or AI systems thereby increasing user trust through increased Situation Awareness (SA). These DIGs, displayed in augmented reality, provide real-time battlefield information. For example, we developed an Army Aviation MEDEVAC scenario with Forward Arming and Refueling Points (FARPs), 9-Line, Patient, and Medical Treatment Facilities (MTFs) DIGs to increase aircrew and flight medic situational awareness and enhance decision-making during missions. Displaying a vast amount of information from an autonomous or AI system is nearly impossible. However, displaying the right information to the right Soldiers at the right time is key to mission success. The DIG concept will help provide timely and accurate information to Soldiers at the speed of battle.

Title: Bayesian adaptive design as an effective testing approach

Primary Author: Victoria Rose Carrillo Sieck

Abstract: When developing a system, it is important to consider system performance from a user perspective. This can be done through operational testing---assessing the ability of representative users to satisfactorily accomplish missions with the system in operationally-representative environments. This process can be expensive and time-consuming, but is critical for evaluating a system. We show how an existing design of experiments (DOE) process for testing can be leveraged to construct a Bayesian adaptive design. This novel approach, nested within the larger design created by the DOE process, allows interim analyses using predictive probabilities to stop testing early for success or futility. Representative simulations demonstrate how these interim analyses can be used in an operational test setting, and reductions in necessary test events are shown. The method allows for using priors when data from previous testing is not available, or for priors built using developmental testing data when it is available. The proposed method for creating priors using testing data allows for more flexibility than the current process, and demonstrates that it is possible to get more precise parameter estimates. Benefits and limitations derived from a case study will also be considered. Ultimately, this method will allow future testing to be conducted in less time and at less expense, on average, without compromising the ability of the existing process to verify the system meets the user's needs.

Title: Upcoming MORS Special Meeting: "Scoping the Analytical Implications of Climate Change and Extreme Events for National Security"

Primary Author: Mr. Donald H. Timian

Abstract: National security communities around the world are recognizing that climate change has significant consequences both at home and abroad and are now beginning to incorporate strategies for mitigating the impact of climate change into their security assessments. Thus, MORS is organizing a 6-8 December 2022 (Unclassified) Climate Change Mini-Symposium; hosted by Johns Hopkins University Applied Physics Laboratory. Its focus will be on information sharing from Centers of Excellence around the world, identifying ongoing and recent analyses, tools, and analytical approaches used to

address/manage the impact of climate change on defense and homeland security issues. The purpose of this presentation is to provide an update on our agenda and keynote slate, discuss the mini-symposium's goals, and envisioned outcomes.

Title: Safe Machine Learning Prediction and Optimization via Extrapolation Control
Primary Author: Dr. Thomas A. Donnelly
Abstract: Uncontrolled model extrapolation leads to two serious kinds of errors: (1) the model may be completely invalid far from the data, and (2) the combinations of variable values may not be physically realizable. Optimizing models that are fit to observational data can lead to extrapolated solutions that are of no practical use without any warning. In this presentation we introduce a general approach to identifying extrapolation based on a regularized Hotelling T-squared metric. The metric is robust to certain kinds of messy data and can handle models with both continuous and categorical inputs. The extrapolation model is intended to be used in parallel with a machine learning model to identify when the machine learning model is being applied to data that are not close to that model training set or as a non-extrapolation constraint when optimizing the model.

Title: Dynamically Tracking Space R&D Funding using Python and Machine Learning
Primary Author: Nicholas Wagner
Abstract: Maintaining awareness of the funding landscape for space technologies is a time-consuming struggle. We combine open source data engineering and machine learning Python libraries with federal funding data to curate a dataset of space-relevant contract opportunities and awarded contracts. We then serve this information to users via a dashboard interface to allow for user-driven analysis. In addition, we provide daily email alerts of new opportunities to reduce user attention burden. We discuss the potential for other technology areas which have needs for rapid analysis.

Title: Ethics, Rules of Engagement, and AI: Neural Narrative Mapping Using Large Transformer Language Models
Primary Author: Dr. Philip Gregory Feldman
Abstract: The problem of determining if a military unit has correctly understood an order and is properly executing on it is one that has bedeviled military planners throughout history. The advent of advanced language models such as OpenAI's GPT-series offers new possibilities for addressing this problem. We have developed an approach to harness the narrative output of large language models and produce diagrams or "maps" of the relationships that are latent in the weights of such models as the GPT-3. The resulting "Neural Narrative Maps" (NNMs), are intended to provide insight into the organization of information, opinion, and belief in the model, which in turn provide means to understand intent and response in the context of physical distance. This presentation will discuss the problem of mapping information spaces in general, and then will present a concrete implementation of this concept in the context of OpenAI's GPT-3 language model. In this scenario, we will determine if a subordinate is following a commander's intent in a high-risk situation. The subordinate's locations within the NNM allow a novel capability to evaluate the intent of the subordinate with respect to the commander. We show that it is possible not only to determine if they are nearby in narrative space, but also how they are

oriented, and what "trajectory" they are on. Our results show that our method is able to produce high-quality maps, and demonstrate new ways of evaluating intent more generally.

Title: Leveraging Neural Networks to Modernize Placement Methodology for the U.S. Army Recruiting Command's Enlisted Recruiting Force

Primary Author: Charles Lovejoy

Abstract: This research project seeks to simplify and innovate the how the United States Army Recruiting Command (USAREC) determines the number of enlisted recruiters it needs to achieve projected missions on a battalion level. Historically, USAREC issued surveys to gauge workload, but this methodology has significant drawbacks in terms of time, resources, and data accuracy. In this project, we rely on alternate methodologies determining manpower requirements and develop a neural network to determine the number of enlisted recruiters. We utilize data from the Department of Defense (DoD) Recruiter Quality of Life Survey (RQoLS) and USAREC's battalion-level Mission Achievement Plans (MAP) as input for the neural network to forecast USAREC's recruiting personnel requirements. The model was trained on the most recent set of RQoLS data (2018 and 2020) and tested on independent 2021 data. The final neural network contained two hidden layers and (9,3) nodes in each respectively. The model achieved an "accuracy" (1-MAPE) of 93% with RMSE of 23 recruiters on the testing data. The model indicates that USAREC needed approximately 13% more recruiters to successfully meet the mission in 2022.

Title: Simulation-Supported Wargaming at the Campaign Level

Primary Author: Mr. Charles D Burdick, CAP

Abstract: Simulation-Supported Wargaming at the Campaign Level

This presentation addresses agent-based simulation support to multi-domain Wargaming at the Campaign Level capable of gaming the detailed maneuver and firepower of large global scenarios including all C4ISR assets and logistics support with less manual effort, more analytical rigor, and verified repeatability.

The government-owned Joint Analysis System (JAS) was in use until ten years ago when it was archived for "budgetary" reasons. In its pure simulation mode JAS was used by OSD/CAPE to develop entire Strategic Scenarios that could be executed in a single full-domain model and distributed to government and government-sponsored industry and FFRDCs.

In its simulation-supported wargaming mode, JAS provided JFCOM (J9) with a means to simulate weapons effects, calculated and displayed movements on digital maps, and provided the wargamers with status reports (units, supplies, etc.) and a perception-based Common Operational Picture (COP) to assess the situation and make decisions under uncertainty.

JAS brings human decision making into the simulated environment to assess situations under the "Fog and Friction" of global competition and conflict. And offers opportunities to evaluate and improve JAS computer agent actions and decisions by demonstrating contributions both SME human evaluators and the application of AI in the perceptions, actions, and decisions of the over 150 types of agents simulated in JAS.

JAS uses common meta-data for both human and agent decisions allowing the human decisions, orders, priority changes, information requests, and similar wargaming inputs to be treated by JAS in the same manner as those of its agents, i.e., they will be executed by subordinate agents in a seamless fashion and recorded along with all agent actions and decisions. With the same initial random seed, JAS' verified repeatability can later bring its simulation-supported wargames back to any given point for review, discussion, and new inputs. Since JAS records ground truth as well as agent perceptions, comparisons can be made and improvements in sensors, processing, etc. could be examined to bring the perceptions-based assessments closer to Ground Truth.

Analysts could also modify the JAS initial random seed creating a Monte Carlo simulation and generate a distribution of outcomes from the wargame decisions. Or changes can be inserted in the scenario conditions, e.g., weather, unit deployments, tactics changes, etc. to observe their effects. The results of the simulation of the original wargame could be analyzed by small teams of analysts and/or SMEs and done with just a copy of the wargame scenario digital recording and a desktop computer.

JAS is a minimally aggregated model with all weapon types explicitly represented in organized units or in temporarily spawned configurations such as sorties, patrols, truck convoys, etc. All units have countable assets in operationally distributed locations/ formations and are subject to engagement and destruction. When this enumeration of assets and resources is combined with the detailed locations of units, vertical federation of the JAS Campaign model with Mission level models can occur and has been demonstrated using the DoD High Level Architecture (HLA) well before its archival.

JAS has simulated networks of communication connecting all units with a computer agent and other than observing local phenomena, this is the only way they receive information from sensors, receive and issue orders, request supplies, etc. Both kinetic and non-kinetic attacks including cyber and electronic warfare can disrupt those communications leading to impacts on operations as well.

Large scale JAS Scenarios should be built with collateral FVEY data and used by COCOMS and for the many DoD organizations that have little insight into the dynamics of combat with peer level opponents and the likely requirements for and impacts on their systems. These organizations can also modify their copies with higher classification data and proposed systems and to better reflect their current and potential future capabilities and observe their effects.

It is estimated by former JAS developers that the unclassified JAS model could be restored to operation in a matter of a few staff weeks and progress to a Microsoft OS 10 environment within a few staff months given that those same experienced JAS developers are used for the task. The benefits of leveraging tighter coordination between the simulation and wargaming communities would significantly improve our analytical capabilities. Let's revive JAS.

Title: Enterprise Risk Assessment and Strategic Security Measure Allocation for Optimal Aviation Security

Primary Author: Major Taylor Leonard

Abstract: This study establishes a generalizable quantitative construct for enterprise risk assessment and optimal portfolio investment to achieve the best aviation security and solve other risk and security

problems. We first analyze and model various aviation transportation risks and establish their interdependencies. Then, using the security measures and capabilities, we formulate a multi-objective portfolio investment model using mixed integer programming. The portfolio risk model determines the best security capabilities and resource allocation under a given budget while pinpointing potential capabilities impacts due to changes in the budget. Our analysis involves cascading and inter-dependency modeling multi-tier risk taxonomy and overlaying security measurements. The model determines an optimal security measure portfolio to distribute across airports nationwide with the objectives to minimize the probability of false clears, maximize the probability of threat detection, and maximize the ability to mitigate risks in aviation security. This work presents the first comprehensive model that links all resources across the 440 federally funded airports in the United States. To solve the resulting instances, we experiment and compare several computational strategies. At the same time, our solutions offer improved risk posture, lower false clears, and higher threat detection across all the airports, indicating a better risk enterprise strategy and decision obtained from our system.

Title: Resilient Position Data Fusion with a Generic Kalman Filter Method

Primary Author: Dr. Paul Edward Fanto

Abstract: Military and civil platforms and infrastructure rely on the Global Positioning System (GPS) to provide essential positioning, navigation, and timing (PNT) information. Given the risk of GPS denial or spoofing, the Department of Defense (DoD) is focused on providing resilient alternative PNT for military platforms. Fusing information from multiple alternative sources promises to increase the accuracy and resilience of the PNT solution. We will present a generic Kalman-filter-based approach for fusing position information from multiple sources for a notional trajectory. In particular, we assess the degradation in position accuracy and precision due to GPS denial, and we find that combining GPS with multiple other sources improves the GPS solution. We also model the effect of GPS spoofing on the fused solution and present a simple algorithm for detecting and rejecting spoofed GPS. Our method can be used as a baseline against which to evaluate other position sensor fusion approaches.

Title: Functions to Assets: A Holistic Risk Assessment of Critical Infrastructure

Primary Author: Dr. Lauren Wind

Abstract: As critical infrastructure (CI) becomes more complex and interdependent, there is a need to a help stakeholders and operators of CI assess, understand, and manage risk across the Nation. The number of assets (facilities, software, services etc.), which are managed and operated by many entities and organizations, continue to expand. While CI system models exists, they generally focus on specific sectors, specific subsets of assets, and often specific threats. There is a need for data and models to be integrated and aggregated to understand the cascading impacts to CI at the national level. One way to do this, is to develop a functions-based approach to assessing and managing risk. SPA is working for a government client to integrate the first ever functions-based risk assessment tool for CI that bridges the gap between the silos of sector and asset level analysis. This risk analysis tool integrates varying levels of data (functions, sectors, entities, assets) and risk models (threat, consequence, vulnerability) to provide a holistic risk assessment to CI across all critical functions based on repeatable and defensible methodology. This tool is the first to connect the national critical function intra- and inter-dependencies and identify the cascading consequences that exists across the expansive CI network. This tool is scalable

and extensible as new data, models, and capabilities are integrated to help support analysts, stakeholders, and operators in risk assessment to CI.

Title: Data Fabric: A Distributed Concept of Data Operations for Information Superiority in all Operational Environments, from the Enterprise to the Battlespace

Primary Author: Bhaarat Sharma

Abstract: To harness the power of data to achieve information superiority over adversaries and competitors, the DoD should adopt an open and distributed architecture – a data fabric – as its concept of data operations. The data fabric represents a significant departure from the traditional approach of standardizing, normalizing, and centralizing data – an approach that is expensive to implement and maintain, impractical to enforce, and too brittle for data "in use." Instead, a data fabric uses automated and self-orchestrating applications to continuously connect disparate and heterogeneous data sources to enable integrated discoverability and governance all data asset and enable users to access and connect to the data they need, when they need it, and regardless of where it is produced or stored. The data fabric architecture will also enable DoD to realize its vision of AI-enabled all-domain information fusion – even in contested, denied, and resource constrained environments. This presentation will provide a high-level technical overview of the data fabric design concept; review the challenges and shortcomings of traditional approaches; and present the case for the adoption of the data fabric as the DoD's concept for data in use. It will include examples from Raft's work to implement the Department of the Air Force's Data Fabric Minimum Viable Product, including use case-driven development for Space Force operational users. Finally, the presentation will conclude with a discussion of over-the-horizon challenges and opportunities – technical, bureaucratic, and cultural – to building and adopting the data fabric based on research and previous efforts to implement traditional and data fabric-based concepts across the U.S. Government.

Title: Operationalizing open-world probabilistic programming at the edge

Primary Author: Dr David Rushing Dewhurst

Abstract: Probabilistic programming (PP) elevates probabilistic modeling and statistical inference concepts to first-class programming language constructs, facilitating rapid development of capabilities that more closely meet operational constraints. Open-world PP further enables reasoning over an a priori unknown number of random events, objects, or scenarios, leading to further increased operational relevance. For example, intrusion detection systems (IDS) are vital tools on today's cyber battlefield, but human operators quickly cease to trust IDS that exhibit high false positive rates (FPR); open-world PP can be used to develop automated FPR reduction mechanisms that balance explainability with performance. Unfortunately, most implementations of open-world PP currently suffer from lack of scalability to the "edge" of low size, weight, and power (SWaP) hardware and real-time computation constraints. Low SWaP and real-time computation are realities of life in multiple operational application areas (e.g., critical cybersecurity systems, tactical sensor suites). We enumerate diverse issues involved with moving open-world PP to the edge, including: (a) choice of implementation language; (b) choice of memory model; (c) considerations arising under severe computational constraints (e.g., lack of heap memory); and (d) the interplay between inference algorithm choice and computational constraints. We

demonstrate several of the discussed concepts using a new open-source, open-world PP library hosted in modern C++.

Title: Community and Infrastructure Adaptation to Climate Change (CIACC): From Scientific Literature to Actionable Insights using Advanced Natural Language Processing

Primary Author: John K. Hutchison

Abstract: Critical infrastructure systems throughout the U.S. are increasingly at risk due to systemic underinvestment and intensifying natural hazards driven by anthropogenic climate change. Disruptions of critical infrastructure systems, such as drinking water and electric power systems, threaten the safety and well-being of people and communities, cause significant financial damage, and may inflict socioeconomic damages that span years or decades. Research on climate change, impacts on infrastructure, and infrastructure adaptation are constantly evolving and researchers are publishing at a blistering pace. This makes it nearly impossible for individuals or even teams across government, academia, and industry to review scientific and engineering advancements and use them to inform climate adaptation decision-making.

To address this need, our team of researchers at Argonne National Laboratory is developing the Community and Infrastructure Adaptation to Climate Change (CIACC) tool. It leverages advanced natural language processing (NLP) techniques to analyze large volumes of climate-change-related scientific literature. We developed an NLP workflow to generate categories for a large corpus of climate research articles, perform category-based topic modeling, search for relevant documents from the corpus given a search query, summarize the text from a large number of documents, and respond to specific questions in order to address the difficulty decision-makers face in comprehending the research on infrastructure impacts and climate change. The tool offers decision-makers cutting-edge and actionable information on climate hazards, threats to critical infrastructure, and climate adaptation best practices – helping them better safeguard systems and communities. In this presentation, we present 1) a summary of our AI system, architecture, and toolkits 2) how our system is capable of detecting significant climate change patterns, risk factors, and environmental impacts and 3) finally, a discussion of how this will enhance research on critical infrastructure systems and climate change.

Title: Micro Baselines for Operational Environments

Primary Author: Dr. Gabriel A. Weaver

Abstract: Critical infrastructure stakeholders need to baseline their networks to understand expected communications. Top-down approaches to baselining rely on observables that are generally available but lack properties upon which traditional statistical tools depend. We propose to construct micro-baselines: signatures within operational networks based on observables associated with specific events. Such observables are informed by precursor analysis reports of historical

cyber attacks on operational environments developed by Cybersecurity for Operational Technology Environments (CyOTE). Baseline measurements depend upon context beyond the cyber domain. An energy plant's baseline running in the summer may statistically differ from a similar facility in a colder region. Domain knowledge must be integrated to apply general micro-baselining algorithms to a facility-

specific context. Therefore, we propose to explore the feasibility of different micro baselining algorithms across different facilities. Facilities that implement the same processes in

different geographic locations will be compared relative to observable measurements used in micro-baselining for comparable events. One evaluation approach would condition or augment dynamic observables measured within a facility network testbed with additional observables derived from geographic context or infrastructure dependencies such as those provided by the All-Hazards Analysis tool.

Title: Recurrent Neural Networks to Streamline Data Interoperability

Primary Author: Dr. Donald Williams, Jr.

Abstract: This abstract represents an ongoing study exploring the use of artificial intelligence (AI) recurrent neural networks (RNNs) to streamline data interoperability between operations research models. The research question is: “How could RNNs improve statistical analysis, recognize data compatibility changes, and/or modify data sets to improve data compatibility between models?" This research is relevant to operations research methodology in three ways. First, it addresses how RNNs are especially suited for optimizing linear and non-linear statistical models. Second, it explores how RNNs may recognize data incompatibility between models and inform the researcher of the effects of this incompatibility. Third, it examines how an RNN may correct data incompatibility between models by converting data types and maintaining accuracy during the data transformation process. A neural network algorithm assigns weight to inputs and produces output based on the comparative weights of its inputs. Researchers may use the technology to recognize and correct situations in which data is lost or misrepresented as different models use a single data set. An RNN is uniquely suited for this research because it is agile enough to accommodate a wide range of modern optimization tools. It understands the context of its calculations, making it particularly useful for streamlining data interoperability between operations research models.

Title: Using Coupled Optimization to Solve Complex Systems Decision Analysis

Primary Author: Christina Bridges

Abstract: Conventional design decision analysis techniques applied to complex systems often wrongly assume the independence of system quality attributes, thereby risking inaccurate evaluation of design alternatives and delivery of suboptimal systems. Stakeholders need accurate evaluation of a system’s attributes ability to meet the requirements to make informed decisions. To optimize system design and manage increasingly complex systems, a decision analysis process starting in the design phase, that explicitly considers the inter-relationships of attributes, and works with flexible design methods is needed. The subject research is developing a process to analyze complex systems by assembling an ensemble of analysis models that incorporates the individual analysis of quality attributes and establishes the connections between coupled attributes. Genetic algorithms, a class of meta-heuristics, are being applied to the system-level optimization to find a suitable solution, because a pure mathematical solver does not exist. The deliverable is an industry-independent process that provides specific recommendations on the type of multi-object optimization models and techniques, along with a prescribed implementation translating the system performance and quality attributes into an ensemble

of coupled multi-object optimization models. The research outputs will provide evidence of attribute dependencies at the system analysis level and a validated process usable across industries.

Title: C-sUAS Analytic Model

Primary Author: Dr. Michael Yereniuk

Abstract: Army Futures Command (AFC) and Futures and Concept Center (FCC) sponsored a 6 month Counter-small Unmanned Aircraft Systems (C-sUAS) study to assess the operational effectiveness and to determine requirements of future C sUAS technologies by echelon and battlespace geometry. Study time and lack of existing models to represent and assess emerging technologies compelled the team to develop the Counter Information Seeker Computational Research (CISCR) model. The CISCR model is a stochastic, highly-abstracted, time-stepped adjudication model to analyze measures of operational effectiveness. Using a balance of explicit and implicit techniques, CISCR represents Friendly and Threat intent including sUAS behaviors to gain unit level knowledge. Combining global communications with individual entity-level knowledge causes emergent, swarm-like sUAS behavior. With every run, CISCR tracks the sUAS and C-sUAS identifications, engagements, and defeats, thus allowing the study team to complete quantitative analysis. The team used the CISCR model to simulate over 250 scenario situations that provided insights into the effectiveness of Threat sUAS and opportunities for future investment in Friendly C-sUAS technologies. This presentation will provide an example of how quantitative analysis of future systems, where limited data is available, are gained by using the CISCR model.

Title: Understanding the Future Through Army Capability Threads

Primary Author: Mr. David DiCarlo

Abstract: The US Army delivers capabilities to the Joint Force that enables the Combatant Commands to fight and win our nation's wars. Capabilities are system-of-systems comprised of DOTMLPF-P solutions that must work in tandem to be effective. Traditionally, these solutions have been analyzed, evaluated, and tracked on an individual system or portfolio basis. Army Capability Threads (ACTs) are an emerging Army concept being developed by Army Futures Command (AFC) that encompasses the network of systems for the purpose of analysis that enables decision-making.

ACTs track the interdependencies between systems, portfolios, and DOTMLPF-P solutions, informed by DoD mission engineering threads, in a way that enables strategic decision making. The initial prototype of an ACT started development in August 2022. To realize their potential, ACTs will require collaboration from the Army Staff, the Secretariat, and Army Commands.

This presentation will describe what ACTs are, the process that has been envisioned to build them, and hypothetical use cases to highlight how the existence of ACTs can augment analyses.

Title: Assessing Ransomware Activity in Operational Technology Environments Using Bayesian Networks

Primary Author: Dr. Lee T Maccarone

Abstract: Critical infrastructure and other operational technology (OT) environments face increasing cybersecurity risks from adversarial behavior. The Cybersecurity for the Operational Technology Environment (CyOTE) program enables asset OT system operators to secure their OT environments through a range of resources and tools. The cornerstones of the CyOTE methodology are the perception of observable events and the comprehension of these observables in broad context including people, processes, and technologies. This research defines a risk-based approach to enhance comprehension of attack observables. The CyOTE team constructed a Bayesian network using the MITRE ATT&CK® for Industrial Control Systems (ICS) framework as a common lexicon for describing potential adversary behavior in the OT environment. As the adversary utilizes techniques and generates observables, the Bayesian network calculates the probability of adversarial behavior. Opportunities for improved comprehension of MITRE ATT&CK® for ICS techniques are also identified through analysis of each technique's observables. The results are provided in a graphical display with both quantitative data and plain language interpretations. Three historical case studies of ransomware attacks demonstrate how the Bayesian network assessment of MITRE ATT&CK® for ICS techniques provides actionable information to assist investigators with determining the cause of an anomaly.

Title: Generative Methods and Meta-learning for Cybersecurity

Primary Author: Capt Marc Winczer Chale

Abstract: The internet of battlefield things is vast, growing, and contested. Network intrusion detection systems (NIDS) seek to identify malicious data before it reaches the target location within a network. Most modern NIDS use machine learning. Adversarial attack is a strategy to fool machine learning classifiers and has been demonstrated in domains such as computer vision. Adversarial attacks in the cyber domain are more complex due to the risk of corrupting encoded data. Previous works speculate that adversarial attacks are possible in the cyber domain but fail to provide a viable attack vector. This study algorithmically generates an adversarial attack that evades a typical target NIDS with 69% success rate. An adversarially trained stacked ensemble meta-learner is presented that defeats all tested adversarial attacks. While our community should be weary of adversarial attacks, the demonstrated meta-learning classifiers may offer protection.

Title: Using Transfer Learning to Accelerate Model Development

Primary Author: Andy Block

Abstract: Transfer learning is a rapidly expanding field within data science that lowers the barrier to entry for applying deep learning models to your domain-specific data. These neural networks have already been trained by the scientific and academic communities on massive, generic datasets, such that their weight and biases matrices optimize the given training loss function. The intention is for these models to be "smart" enough to transfer their learning from the training phase directly to your particular real-world data, with the exception of training the bespoke output layers configured to the scope of your technical problem. These techniques have been successfully applied across a broad range of problem scopes, from image recognition to natural language processing and beyond -- this discussion will focus on how to design a transfer learning approach to solving your technical challenges and how it can accelerate the implementation of your solutions.