

Appendices

Appendix A: Technical Glossary	601
Appendix B: Acronyms Found in This Report	617
Appendix C: Key Considerations for the Responsible Development and Fielding of Artificial Intelligence (Abridged)	633
Appendix D: Draft Legislative Language	663
Appendix E: Funding Recommendation Table	729
Appendix F: Commissioner Bios	741
Appendix G: Commission Staff and Contributors	749

Appendix A: Technical Glossary

3D Chip Stacking: The process of building integrated circuits with both horizontal and vertical interconnections between transistors. This brings elements of the chip physically closer together, increasing density and allowing for greater performance (i.e., speed) at lower power levels and at a smaller footprint than comparable two-dimensional devices, which only feature horizontal interconnects.

Additive Manufacturing: A computer-controlled process in which successive layers of material are deposited to create a part that matches a 3D design.

Adversarial Machine Learning: A broad collection of techniques used to exploit vulnerabilities across the entire machine learning stack and lifecycle. Adversaries may target the data sets, algorithms, or models that an ML system uses in order to deceive and manipulate their calculations, steal data appearing in training sets, compromise their operation, and render them ineffective.¹ Adversarial AI may be used as a phrase that broadens the considerations to attacks on AI systems, including approaches that are less dependent on data and machine learning.

Agile: A philosophy and methodology used to describe the continuous, iterative process to develop and deliver software and other digital technologies. User requirements and feedback inform incremental development and delivery by developers.²

AI Assurance: The defensive science of protecting AI applications from attack or malfunction.

AI Digital Ecosystem: A technology stack driving the development, testing, fielding, and continuous update of AI-powered applications. The ecosystem is managed as a multi-layer collection of shared AI essential building blocks (e.g., data, algorithms, tools, and trained AI models) accessed through common interfaces.

AI Governance: The actions to ensure stakeholder needs, conditions, and options are evaluated to determine balanced, agreed-upon enterprise objectives; setting direction through prioritization and decision-making; and monitoring performance and compliance against agreed-upon directions and objectives.³ AI governance may include policies on the nature of AI applications developed and deployed versus those limited or withheld.

AI Lifecycle: The steps for managing the lifespan of an AI system: 1) Specify the system's objective. 2) Build a model. 3) Test the AI system. 4) Deploy and maintain the AI system. 5) Engage in a feedback loop with continuous training and updates.⁴

AI Stack: AI can be envisioned as a stack of interrelated elements: talent, data, hardware, algorithms, applications, and integration.⁵

Algorithm: A series of step-by-step instructions or calculations to solve an instance of a problem. There are fundamentally two ways that algorithms are implemented by AI: explicit engineering of the algorithm (e.g., in symbolic reasoning and expert systems) or by machine learning, where the algorithm is derived from data or feedback from interactions.

Anonymization: Also referred to as data de-identification, this is the process of removing or replacing with synthetic values any identifiable information in data. This is intended to make it impossible to derive insights on any specific individual in the data while remaining useful for the intended use of the data.⁶ (See de-anonymization.)

Application Programming Interfaces (APIs): Programming tools for describing how one program can access the functionality of another⁷ while hiding the implementation details inside each program.

Application-Specific Integrated Circuit (ASIC): A chipset custom designed to perform a particular task. ASICs could provide significant performance gains over generic chips but are inflexible in their functions compared to central processing units.

Architecture: A set of values, constraints, guidance, and practices that support the active evolution of the planning, designing, and construction of a system. The approach evolves over time, while simultaneously supporting the needs of current customers.⁸ Architecture can refer to sets of components in a computing system and their operational interrelationships as well as other important configurations such as the architecture of a neural network, which captures the patterns of connectivity within and between layers of units in the network model.

Artificial General Intelligence (AGI): A phrase that has been used to capture the possibility of developing more general AI capabilities, in distinction to the typically narrow capabilities of AI systems that have been developed to date. Some use the term to refer to the prospect of achieving more human-like intelligence, developing AI systems with the ability to perform many of the intellectual tasks that humans are capable of doing, or developing systems that might employ a wide range of skills across multiple domains of expertise.

Artificial Intelligence (AI): The ability of a computer system to solve problems and to perform tasks that have traditionally required human intelligence to solve.

Auditability: A characteristic of an AI system in which its software and documentation can be interrogated and yield information at each stage of the AI lifecycle to determine compliance with policy, standards, or regulations.

Augmented Reality: Enhanced digital content, spanning visual, auditory, or tactile information, overlaid onto the physical world.⁹

Authorization to Operate (ATO): The official management decision given by a senior organizational official to authorize operation of an information system and to explicitly accept the risk to organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, and other organizations based on the implementation of an agreed-upon set of security controls.¹⁰

Automation Bias: An unjustified degree of reliance on automated systems or their outcomes.

Autonomous: A system with functions capable of operating without direct human control.

Biological Sensors (Biosensors): Devices used to detect the presence or concentration of a biological analyte, such as a biomolecule, a biological structure, or a microorganism. Biosensors consist of three parts: a component that recognizes the analyte and produces a signal, a signal transducer, and a reader device.¹¹

Biometric Technologies: Technologies that leverage physical or behavioral human characteristics that can be used to digitally identify a person and grant access to systems, devices, or data, such as face, voice, and gait recognition.¹²

Black Box: The nature of some AI techniques whereby the inferential operations are complex, hidden, or otherwise opaque to their developers and end users in terms of providing an understanding of how classifications, recommendations, or actions are generated and what overall performance will be.

Carbon Nanotubes: Nano-scale structures that can be used to make transistors and could potentially replace silicon transistors in the future. Compared to existing silicon transistors, carbon nanotube transistors are both capable of being shrunk to a smaller size and more amenable to being stacked in three dimensions (see 3D chip stacking).

Cloud Computing: The act of running software within information technology environments that abstract, pool, and share scalable resources across a network.¹³

Cloud Infrastructure: The components needed for cloud computing, which include hardware, abstracted resources, storage, and network resources.¹⁴

Commonsense Reasoning: The process of forming a conclusion based on the basic ability to perceive, understand, and judge things that are shared by (“common to”) most people and can reasonably be expected without need for debate.¹⁵ Endowing computing systems with the commonsense knowledge of humans has been found to be a difficult and standing AI challenge.

Computational Thinking: The thought processes involved in formulating problems so their solutions can be represented as computational steps and algorithms.¹⁶

Computer Vision: The digital process of perceiving and learning visual tasks in order to interpret and understand the world through cameras and sensors.¹⁷

Continuous Delivery: A process that builds on continuous integration by taking the step of orchestrating multiple builds, coordinating different levels of automated testing, and moving the code into a production environment in a process that is as automated as possible.¹⁸

Continuous Integration: A process that aims to minimize the duration and effort required by “each” integration episode and deliver at any moment a product version suitable for release. In practice, this requires an integration procedure that is reproducible and mostly automated. This is achieved through version control tools, team policies, and conventions.¹⁹

Data Architecture: The structure of an organization’s logical and physical data assets and data management resources.²⁰

Data Privacy: The right of an individual or group to maintain control over, and the confidentiality of, information about themselves.²¹

Data Protection: The practice of safeguarding information from unauthorized access, use, disclosure, disruption, modification, or destruction, to provide confidentiality, integrity, and availability.²²

De-anonymization: Matching anonymous data (also known as de-identified data) with publicly available information, or auxiliary data, in order to discover the individual to whom the data belong.²³ (See anonymization.)

Deepfake: Computer-generated video or audio (particularly of humans) so sophisticated that it is difficult to distinguish from reality.²⁴ Deepfakes have also been referred to as synthetic media.

Deep Learning: A machine learning implementation technique that exploits large quantities of data, or feedback from interactions with a simulation or the environment, as training sets for a network with multiple hidden layers, called a deep neural network, often employing

an iterative optimization technique called gradient descent, to tune large numbers of parameters that describe weights given to connections among units.²⁵

Deep Neural Networks (DNN): A deep learning architecture that is trained on data or feedback, generating outputs, calculating errors, and adjusting its internal parameters. The process is repeated possibly hundreds of thousands of times until the network achieves an acceptable level of performance. It has proved to be an effective technique for image classification, object detection, speech recognition, some kinds of game-playing, and natural language processing—problems that challenged researchers for decades. By learning from data, DNNs can solve some problems much more effectively and also solve problems that were never solvable before.²⁶

Deployed AI: AI that has been fielded for its intended purpose within its relevant operational environment.

DevSecOps: Enhanced engineering practices that improve the lead time and frequency of delivery outcomes, promoting a more cohesive collaboration between development, security, and operations teams as they work toward continuous integration and delivery.²⁷

Differential Privacy: A criterion for a strong, mathematical definition of privacy in the context of statistical and machine learning analysis used to enable the collection, analysis, and sharing of a broad range of statistical estimates, such as averages, contingency tables, and synthetic data, based on personal data while protecting the privacy of the individuals in the data.²⁸

Digital Ecosystem: The stakeholders, systems, tools, and enabling environments that together empower people and communities to use digital technology to gain access to services, engage with each other, and pursue missional opportunities.²⁹

Digital Infrastructure: The foundational components that enable digital technologies and services. Examples of digital infrastructure include fiber-optic cables, cell towers, satellites, data centers, software platforms, and end-user devices.³⁰

Distributed System: A system whose components are located on different networked computers, which communicate and coordinate their actions by passing messages to one another in order to appear as a single system to the end user.³¹

Domain-Specific Hardware Architectures: Hardware that is specifically designed to fulfill certain narrow functions, seeking performance gains through specialization.

Edge Computing: A distributed-computing paradigm that brings computation and data storage closer to the location where it is needed (i.e., the network edge where smart sensors, devices, and systems reside along with points of human interaction) to improve response times and save bandwidth.³²

Expert System: A computer system emulating the decision-making ability of a human expert through the use of reasoning, leveraging an encoding of domain-specific knowledge most commonly represented by sets of if-then rules rather than procedural code.³³ The term “expert system” was used largely during the 1970s and '80s amidst great enthusiasm about the power and promise of rule-based systems that relied on a “knowledge base” of domain-specific rules and rule-chaining procedures that map observations to conclusions or recommendations.

Explainability: A characteristic of an AI system in which there is provision of accompanying evidence or reasons for system output in a manner that is meaningful or understandable to individual users (as well as to developers and auditors) and reflects the system's process for generating the output (e.g., what alternatives were considered, but not proposed, and why not).³⁴

False Negative: An example in which the predictive model mistakenly classifies an item as in the negative class. For example, a false negative describes the situation in which a junk-email model specifies that a particular email message is not spam (the negative class) when the email message actually is spam, leading to the frustration of the junk message appearing in an end user's inbox.³⁵ In a higher-stakes example, a false negative captures the case in which a medical diagnostic model misses identifying a disease that is present in a patient.

False Positive: An example in which the predictive model mistakenly classifies an item as in the positive class. For example, the model inferred that a particular email message was spam (the positive class), but that email message was actually not spam, leading to delays in an end user reading a potentially important message.³⁶ In a higher-stakes situation, a false positive describes the situation in which a disease is diagnosed as present when the disease is not present, potentially leading to unnecessary and costly treatments.

Federated Data Repository: A virtual data repository that links data from distributed sources (e.g., other repositories), providing a common access portal for finding and accessing data.

Field-Programmable Gate Array (FPGA): An integrated circuit featuring reconfigurable interconnects that can be programmed by the user to be customized for specific functions after it is manufactured. FPGAs feature greater flexibility than ASICs, but at a cost to performance.

Gallium Nitride: An alternative material to silicon for transistors. Gallium nitride transistors feature higher electron mobility than silicon and are capable of faster switching speed, higher thermal conductivity, and lower on-resistance than comparable silicon solutions.

Generative Adversarial Networks (GANs): An approach to training AI models useful for applications like data synthesis, augmentation, and compression where two neural networks are trained in tandem: one is designed to be a generative network (the forger) and the other a discriminative network (the forgery detector). The objective is for each network to train and better itself off the other, reducing the need for big labeled training data.³⁷

Graphics Processing Unit (GPU): A specialized chip capable of highly parallel processing. GPUs are well-suited for running machine learning and deep learning algorithms. GPUs were first developed for efficient parallel processing of arrays of values used in computer graphics. Modern-day GPUs are designed to be optimized for machine learning.

High-Performance Computing (HPC): Developing, deploying, and operating very high-capacity computers (along with the requisite software, hardware, facilities, and underpinning infrastructure) to advance the computational upper limits of resolution, dimensionality, and complexity.³⁸

Homomorphic Encryption: A technique that allows computation to be performed directly on encrypted data without requiring access to a secret key. The result of such a computation remains in encrypted form and can at a later point be revealed by the owner of the secret key.³⁹

Human-Machine Teaming (or Human-AI Teaming): The ability of humans and AI systems to work together to undertake complex, evolving tasks in a variety of environments with seamless handoff both ways between human and AI team members. Areas of effort include developing effective policies for controlling human and machine initiatives,⁴⁰ computing methods that ideally complement people,⁴¹ methods that optimize goals of teamwork, and designs⁴² that enhance human-AI interaction.

Information Operations: The tactics, techniques, and procedures employed in both the offensive and defensive use of information to pursue a competitive advantage.⁴³

Internet of Things (IoT): A global infrastructure for the information society, enabling advanced services by interconnecting (physical and virtual) things based on existing and evolving interoperable information and communication technologies.⁴⁴

Intelligent Sensing: Utilizing advanced signal processing techniques, data fusion techniques, intelligent algorithms, and AI concepts to better understand sensor data for better integration of sensors and better feature extraction, leading to actionable knowledge that can be used in smart sensing applications.⁴⁵

Interpretability: The ability to understand the value and accuracy of system output. Interpretability refers to the extent to which a cause and effect can be observed within

a system or to which what is going to happen given a change in input or algorithmic parameters can be predicted. Interpretability complements explainability.⁴⁶

Legacy Systems: Outdated systems still in operation that are hard to maintain owing to shortage of skill sets and obsolete architecture.⁴⁷

Machine Learning (ML): The study or the application of computer algorithms that improve automatically through experience.⁴⁸ Machine learning algorithms build a model based on training data in order to perform a specific task, like aiding in prediction or decision-making processes, without necessarily being explicitly programmed to do so.

Microelectronics: A subfield of electronics involving small components such as transistors, capacitors, and resistors. These components are packaged together to form the integrated circuits that are used to perform computations.

MLOps: Enhanced engineering practices that combine ML model development and ML model operations technologies to support continuous integration and delivery of ML-based solutions.⁴⁹

Modeling and Simulation: Modeling the physical world to support the study, optimization, and testing of operations through simulation without interfering or interrupting ongoing processes. Modeling and simulation can be used to train AI systems, and AI technologies can be used to enhance modeling and simulation.

Multi-Party Federated Learning: An ML setting where many clients (e.g., mobile devices or whole organizations) collaboratively train a model under the orchestration of a central server (e.g., service provider) while keeping the training data decentralized. It can mitigate many of the systemic privacy risks and costs resulting from traditional, centralized ML and data science approaches.⁵⁰ However, it does introduce new attack vectors that must be addressed.⁵¹

Multi-Source Data: Data obtained and aggregated from different origins.

Multimodal Data: Data comprising several signal or communication types, such as speech and body gestures during human-to-human communication.

Natural Language Processing: The ability of a machine to process, analyze, and mimic human language, either spoken or written.

Natural Language Understanding: The ability of a machine to represent and act on the meaning that a language expresses utilizing language semantically rather than statistically.

Neuromorphic Computing: Computing that mimics the human brain or neural network.⁵²

Object Recognition: The algorithmic process of finding objects in the real world from an image, typically using object models which are known a priori.⁵³

One Shot (or Few Shot) Learning: An approach to machine learning that leverages existing knowledge to enable learning in some applications (e.g., object recognition) on a few non-repeated examples, with the system rapidly learning similarities and dissimilarities between the training examples.⁵⁴

Open Knowledge Network (OKN): A vision to create an open knowledge graph of all known entities and their relationships, ranging from the macro (e.g., have there been unusual clusters of earthquakes in the U.S. in the past six months?) to the micro (e.g., what is the best combination of chemotherapeutic drugs for a 56-year-old female with stage 3 brain cancer?). OKN is meant to be an inclusive, open, community activity resulting in a knowledge infrastructure that could facilitate and empower a host of applications and open new research avenues, including how to create trustworthy knowledge networks/graphs.⁵⁵

Packaging: The final stage of the semiconductor fabrication process, in which a chip is placed in its protective case. For many years packaging was a low-value element of the semiconductor design process. However, advanced packaging techniques are enabling sophisticated new chip designs using processes such as 3D stacking, heterogeneous integration, and modular chiplets to create more complex and sophisticated semiconductors.

Pattern Recognition: The field concerned with the automatic discovery of regularities in data through the use of computer algorithms, with the use of these regularities to take actions such as classifying the data into different categories.⁵⁶

Planning and Optimization: Determining necessary steps to complete a series of tasks, which can save time and money and improve safety.

Platform Environment: Provides an application developer or user secured access to resources and tools (e.g., workflows, data, software tools, storage, and compute) on which applications can be developed or run.

Polymorphic Malware: A type of malware that constantly changes its identifiable features (i.e., signatures) in order to evade detection. Many of the common forms of malware can be polymorphic, including viruses, worms, bots, trojans, or keyloggers.⁵⁷

Precision: A metric for classification models. Precision identifies the frequency with which a model was correct when classifying the positive class. It answers the question “How many selected positive items are true positive?”—for example, the percentage of messages flagged as spam that actually are spam.⁵⁸

Prediction: Forecasting quantitative or qualitative outputs through function approximation, applied on input data or measurements.⁵⁹

Prior Art: The worldwide scientific and technical knowledge by which an invention is evaluated to determine if it is new.

Pseudonymization: A data management technique to strip identifiers linking data to an individual. Concern exists that such data could still be linked with other data that allows for a person's identity to be rediscovered.

PyTorch: A free and open-source software library for training neural networks and other machine learning architectures, initially developed by Facebook AI Research.

Quantum Computer: A machine that relies on the properties of quantum mechanics to perform computations. Quantum computers encode information in *qubits*, which can exist in a linear combination of two states. These states can be physically realized in a number of ways, such as superconducting circuits, trapped ions, optical lattices, and linear optics. Computation is performed by operating on the state of these qubits using quantum logic gates. For example, if the qubit is realized as an ion, the quantum logic gate might manipulate the ion's energy state with lasers.

Recall: A metric for classification models. Recall identifies the frequency with which a model correctly classifies the true positive items. It answers the question "How many true positive items were correctly classified"? For example, the percentage of spam messages that were flagged as spam.⁶⁰

Reinforcement Learning: A method of training algorithms to make suitable actions by maximizing rewarded behavior over the course of its actions.⁶¹ This type of learning can take place in simulated environments, such as game-playing, which reduces the need for real-world data.

Reliable AI: An AI system that performs in its intended manner within the intended domain of use.

Responsible AI: An AI system that aligns development and behavior to goals and values. This includes developing and fielding AI technology in a manner that is consistent with democratic values.⁶²

Robotics: A broad field of study including autonomous systems that exist in the physical world, sensing their environment and taking actions to achieve specific goals.⁶³

Robotic Process Automation (RPA): Software to help in the automation of tasks, especially those that are tedious and repetitive.

Robust AI: An AI system that is resilient in real-world settings, such as an object-recognition application that is robust to significant changes in lighting. The phrase also refers to resilience when it comes to adversarial attacks on AI components.

Self-Healing Robots: Robots that use structural materials to self-identify damage and initiate healing on their own, repeatedly.⁶⁴

Self-Replicating Robots: A means of manufacturing, so that fleets of autonomous rovers can extract water and metals from local terrain—say on the moon or Mars—to construct new industrial robots autonomously and continue the self-replication loop.

Self-Supervised Machine Learning: A collection of machine learning techniques that are used to train models or learn embedded representations without reliance on costly labeled data; rather, an approach is to withhold part of each data sample and require the algorithm to learn to predict the missing piece.⁶⁵ Self-supervision has been used to train some of the largest language models built to date by training on large amounts of natural language data.⁶⁶

Semi-Supervised Machine Learning: A process for training an algorithm on a combination of labeled and unlabeled data. Typically, this combination will contain a very small amount of labeled data and a very large amount of unlabeled data. One approach is to use the costly, smaller amount of labeled data to bootstrap a classification model, use that model to generate predicted labels across the larger, unlabeled data, and then use the outcome to retrain/refine the model and iterate until class label assignments stabilize.

Semiconductor Manufacturing Equipment (SME): The tools and equipment required to fabricate semiconductors (e.g., extreme ultraviolet and argon fluoride immersion lithography tools).

Semiconductor Photonics: As it relates to semiconductors, this refers to the use of light, rather than electricity, to transfer information on a chip. This allows for much faster data transfer speeds, resulting in significant performance improvements.

Semiconductors: The silicon-based integrated circuits that drive the operations and functioning of computers and most electronic devices.

Smart Sensors: Devices capable of pre-processing raw data and prioritizing the data to transmit and store, which is especially helpful in degraded or low-bandwidth environments.

Smart Systems: Information technology systems with autonomous functions enabled by AI.

Speech Recognition: The algorithmic process of turning speech signals into text or commands.⁶⁷

Supervised Machine Learning: A process for training algorithms by example. The training data consists of inputs paired with the correct outputs. During training, the algorithm will search for patterns in the data that correlate with the desired outputs and learn to predict the correct output for newly presented input data over iterative training and model updates.

SWaP: Size, weight, and power, typically used in the context of reducing the overall dimensions of a device, increasing its efficiency, and lowering the overall footprint and cost—all contributing factors to viable edge computing.⁶⁸

Symbolic Logic: A tool for creating and reasoning with symbolic representations of objects and propositions based on clearly defined criteria for logical validity.⁶⁹

Synthetic Data Generation: The process of creating artificial data to mimic real sample data sets. It includes methods for data augmentation that automate the process for generating new example data from an existing data set. Synthetic data generation is increasingly utilized to overcome the burden of creating large labeled datasets for testing and at times training deep neural networks.

Technical Baseline: The government's capability to understand underlying technology well enough to make successful acquisition decisions independent of contractors.⁷⁰

TensorFlow: A free and open-source software library for training neural networks and other machine learning architectures, initially developed by Google Brain.

Test and Evaluation, Verification and Validation (TEVV) of AI Systems: A framework for assessing, incorporating methods and metrics to determine that a technology or system satisfactorily meets its design specifications and requirements, and that it is sufficient for its intended use.

Traceability: A characteristic of an AI system enabling a person to understand the technology, development processes, and operational capabilities (e.g., with transparent and auditable methodologies along with documented data sources and design procedures).

Unintended Bias: Ways in which algorithms might perform more poorly than expected (e.g., higher false positives or false negatives), particularly when disparate outcomes are produced (e.g. across categories, classes or groups).

Unsupervised Machine Learning: A process for training a model in which the model learns from the data itself without any data labels. Two common approaches are clustering (in which inherent groupings are discovered) and association (in which rules that describe large portions of the data are discovered).⁷¹

Virtual Reality: A simulated experience in a computer-generated synthetic, artificial world involving immersion, sensory feedback, and interactivity.⁷²

Appendix A - Endnotes

- ¹ See *Adversarial Machine Learning 101*, GitHub/MITRE (last accessed Feb. 18, 2021), <https://github.com/mitre/advm1threatmatrix/blob/master/pages/adversarial-ml-101.md#adversarial-machine-learning-101>; see also Ionut Arghire, *Microsoft, MITRE Release Adversarial Machine Learning Threat Matrix*, Security Week (last accessed Feb. 16, 2021), <https://www.securityweek.com/microsoft-mitre-release-adversarial-machine-learning-threat-matrix>.
- ² GAO-20-590G, *Agile Assessment Guide*, U.S. Government Accountability Office at 169 (Sept. 2020), <https://www.gao.gov/assets/710/709711.pdf>.
- ³ See *Glossary*, ISACA (last accessed Feb. 13, 2021), <https://www.isaca.org/resources/glossary>.
- ⁴ Note that for data-driven AI systems, step 2 is expanded and replaced with 2.a) Acquire data to meet the objective, and 2.b) Train the AI system on the data. These two steps are usually repeated, with data acquisition and training continuing until desired performance objectives are attained. For further discussion on the ML lifecycle, see Saleema Amershi, et al., *Software Engineering for Machine Learning: A Case Study*, IEEE Computer Society (May 2019), <https://www.microsoft.com/en-us/research/publication/software-engineering-for-machine-learning-a-case-study/>.
- ⁵ The stack of elements listed here is an adaptation from Andrew W. Moore, Martial Hebert, and Shane Shaneman. See Andrew Moore, et al., *The AI Stack: A Blueprint for Developing and Deploying Artificial Intelligence*, Proc. SPIE 10635 (May 4, 2018), <https://doi.org/10.1117/12.2309483>. For a graphical depiction of the AI stack, see *About*, Carnegie Mellon University Artificial Intelligence (last accessed Jan. 1, 2021), <https://ai.cs.cmu.edu/about>.
- ⁶ See **Recital 26 EU General Data Protection Regulation (EU-GDPR)**, PrivazyPlan (last accessed Feb. 17, 2021), <https://www.privacy-regulation.eu/en/recital-26-GDPR.htm>.
- ⁷ Vinton G. Cerf, *APIs, Standards, and Enabling Infrastructure*, Communications of the ACM, Vol. 62 No. 5, at 5 (May 2019), <https://m-cacm.acm.org/magazines/2019/5/236425-apis-standards-and-enabling-infrastructure/fulltext?mobile=true>.
- ⁸ GAO-20-590G, *Agile Assessment Guide*, U.S. Government Accountability Office at 169 (Sept. 2020), <https://www.gao.gov/assets/710/709711.pdf>.
- ⁹ See *Augmented Reality*, Google (last accessed Feb. 13, 2021), <https://arvr.google.com/ar/>.
- ¹⁰ See *Authorization to Operate*, NIST Computer Security Resource Center (last accessed Feb. 13, 2021), https://csrc.nist.gov/glossary/term/authorization_to_operate.
- ¹¹ See *Biosensors*, Nature (last accessed Feb. 13, 2021), <https://www.nature.com/subjects/biosensors>.
- ¹² Maria Korolov, *What Is Biometrics? 10 Physical and Behavioral Identifiers That Can Be Used for Authentication*, CSO (Feb. 12, 2019), <https://www.csoonline.com/article/3339565/what-is-biometrics-and-why-collecting-biometric-data-is-risky.html>.
- ¹³ See *Understanding Cloud Computing*, Red Hat (last accessed Feb. 13, 2021), <https://www.redhat.com/en/topics/cloud>.
- ¹⁴ See *What Is Cloud Infrastructure?*, Red Hat (last accessed Feb. 13, 2021), <https://www.redhat.com/en/topics/cloud-computing/what-is-cloud-infrastructure>.
- ¹⁵ See Matt Turek, *Machine Common Sense (MCS)*, DARPA (last accessed Feb. 13, 2021), <https://www.darpa.mil/program/machine-common-sense>.
- ¹⁶ See Alfred V. Aho, *Ubiquity Symposium: Computational and Computational Thinking*, ACM (January 2011), <https://ubiquity.acm.org/article.cfm?id=1922682>.
- ¹⁷ See *Computer Vision: What It Is and Why It Matters*, SAS (last accessed Feb. 13, 2021), https://sas.com/en_in/insights/analytics/computer-vision.html.
- ¹⁸ GAO-20-590G, *Agile Assessment Guide*, U.S. Government Accountability Office at 171 (Sept. 2020), <https://www.gao.gov/assets/710/709711.pdf>.
- ¹⁹ *Id.* at 172.

²⁰ See Thor Olavsrud, *What Is Data Architecture? A Framework for Managing Data*, CIO (Nov. 4, 2020), <https://www.cio.com/article/3588155/what-is-data-architecture-a-framework-for-managing-data.html>.

²¹ *Digital Strategy 2020-2024*, USAID at 48 (June 2020), https://www.usaid.gov/sites/default/files/documents/15396/USAID_Digital_Strategy.pdf.

²² *Id.*

²³ See Jake Frankenfield, *De-Anonymization*, Investopedia (Dec. 27, 2020), <https://www.investopedia.com/terms/d/deanonymization.asp#:~:text=De%2Danonymization%20is%20a%20technique,person%2C%20group%2C%20or%20transaction>.

²⁴ *Interim Report*, NSCAI at 9 (Nov. 2019), https://www.nscai.gov/wp-content/uploads/2021/01/NSCAI-Interim-Report-for-Congress_201911.pdf.

²⁵ See Ian Goodfellow, et al., *Deep Learning*, MIT Press, (2016), <https://www.deeplearningbook.org/>.

²⁶ *Perspectives on Research in Artificial Intelligence and Artificial General Intelligence Relevant to DoD*, MITRE, at 9-25 (Jan. 2017), <https://fas.org/irp/agency/dod/jason/ai-dod.pdf>.

²⁷ See *Understanding the Differences Between Agile and DevSecOps—From a Business Perspective*, General Services Administration (last accessed Feb. 13, 2021), https://tech.gsa.gov/guides/understanding_differences_agile_devsecops/.

²⁸ Kobbi Nissim, et al., *Differential Privacy: A Primer for a Non-technical Audience*, Working Group of the Privacy Tools for Sharing Research Data Project, Harvard University (Feb. 14, 2018), https://privacytools.seas.harvard.edu/files/privacytools/files/pedagogical-document-dp_new.pdf.

²⁹ *Digital Strategy 2020-2024*, USAID at 4 (June 2020), https://www.usaid.gov/sites/default/files/documents/15396/USAID_Digital_Strategy.pdf.

³⁰ *Id.* at 49.

³¹ Maarten van Steen & Andrew Tanenbaum, *Distributed Systems* (3rd ed.), distributed-systems.net (2017), <https://www.distributed-systems.net/index.php/books/ds3/>.

³² See Eric Hamilton, *What Is Edge Computing: The Network Edge Explained*, Cloudwards (Dec. 27, 2018), <https://www.cloudwards.net/what-is-edge-computing>.

³³ Peter Jackson, *Introduction to Expert Systems* (3rd ed.), Addison Wesley at 2 (1998).

³⁴ For further discussion see P. Jonathon Phillips, et al., *Four Principles of Explainable Artificial Intelligence*, NIST (Aug. 2020), <https://www.nist.gov/system/files/documents/2020/08/17/NIST%20Explainable%20AI%20Draft%20NISTIR8312%20%281%29.pdf>.

³⁵ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.

³⁶ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.

³⁷ Ian Goodfellow, et al., *Generative Adversarial Nets*, Neural Information Processing Systems (2014), <https://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.

³⁸ *Fiscal Year 2019: Stockpile Stewardship and Management Plan—Biennial Plan Summary, Report to Congress*, U.S. Department of Energy at 3-7 (Oct. 2018), <https://www.energy.gov/sites/prod/files/2018/10/f57/FY2019%20SSMP.pdf>.

³⁹ See *Introduction*, Homomorphic Encryption Standardization (last accessed Feb. 13, 2021), <https://homomorphicencryption.org/introduction/>.

⁴⁰ Eric Horvitz, *Principles of Mixed-Initiative User Interfaces*, CHI '99: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems at 159-166 (May 1999), <https://dl.acm.org/doi/pdf/10.1145/302979.303030>.

Appendix A - Endnotes

- ⁴¹ Bryan Wilder, et al., *Learning to Complement Humans*, Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20) at 1526-1533 (Jan. 2021), <https://www.ijcai.org/Proceedings/2020/0212.pdf>.
- ⁴² Saleema Amershi, et al., *Guidelines for Human-AI Interaction*, CHI '19: Proceedings of the CHI Conference on Human Factors in Computing Systems, at 1-13 (May 2019), <https://dl.acm.org/doi/pdf/10.1145/3290605.3300233>.
- ⁴³ Catherine Theohary, *Defense Primer: Information Operations*, Congressional Research Service (Dec. 15, 2020), <https://fas.org/sgp/crs/natsec/IF10771.pdf>.
- ⁴⁴ See interactive *ITU Terms and Definitions*, United Nations International Telecommunication Union (last accessed Feb. 15, 2021), <https://www.itu.int/net/ITU-R/asp/terminology-definition.asp?lang=en&rlink={42AA741E-A0A7-48C4-905B-AAAFDA29E5F2}>.
- ⁴⁵ See *Intelligent Sensors*, MDPI Sensors (last accessed Feb. 13, 2021), https://www.mdpi.com/journal/sensors/sections/Intelligent_Sensors.
- ⁴⁶ See Richard Gall, *Machine Learning vs Interpretability: Two Concepts That Could Help Restore Trust in AI*, KDnuggets (Dec. 2018), <https://www.kdnuggets.com/2018/12/machine-learning-explainability-interpretability-ai.html>.
- ⁴⁷ A. Sivagnana Ganesan & T. Chithralekha, *A Survey on Survey of Migration of Legacy Systems*, ICIA-16: Proceedings of the International Conference on Informatics and Analytics at 1-10 (Aug. 2016), <https://dl.acm.org/doi/10.1145/2980258.2980409>.
- ⁴⁸ Thomas M. Mitchell, *Machine Learning*, McGraw-Hill (1997).
- ⁴⁹ See *2021 Technology Spotlight: The Emergence of MLOps*, Booz Allen Hamilton (2021), https://www.boozallen.com/content/dam/boozallen_site/dig/pdf/white_paper/the-emergence-of-mlops.pdf.
- ⁵⁰ Peter Kairouz, et al., *Advances and Open Problems in Federated Learning*, arXiv (Dec. 10, 2019), <https://arxiv.org/pdf/1912.04977.pdf>.
- ⁵¹ See Vale Tolpegin et al., *Data Poisoning Attacks Against Federated Learning Systems*, ArXiv (Aug. 11, 2020), <https://arxiv.org/abs/2007.08432>; Arjun Nitin Bhagoji, et al., *Analyzing Federated Learning Through an Adversarial Lens*, arXiv (Nov. 25, 2019), <https://arxiv.org/abs/1811.12470>.
- ⁵² See *Beyond Today's AI: New Algorithmic Approaches Emulate the Human Brain's Interactions with the World*, Intel (last accessed Feb. 13, 2021), <https://www.intel.com/content/www/us/en/research/neuromorphic-computing.html>.
- ⁵³ Ramesh Jain, et al., *Machine Vision*, McGraw-Hill at 459 (1995), https://www.cse.usf.edu/~r1k/MachineVisionBook/MachineVision.files/MachineVision_Chapter15.pdf.
- ⁵⁴ Adam Santoro, et al., *One-Shot Learning with Memory-Augmented Neural Networks*, arXiv (May 19, 2016), <https://arxiv.org/pdf/1605.06065.pdf>.
- ⁵⁵ See *About Workshop, Open Knowledge Network* at National Institutes of Health, Subcommittee on Networking & Information Technology Research & Development, Big Data Interagency Working Group, (Oct. 4-5, 2017), https://www.nitrd.gov/nitrdgroups/index.php?title=Open_Knowledge_Network.
- ⁵⁶ Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer at 1 (2006), <https://www.microsoft.com/en-us/research/uploads/prod/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf>.

- ⁵⁷ See Nate Lord, *What Is Polymorphic Malware? A Definition and Best Practices for Defending Against Polymorphic Malware*, Digital Guardian (July 17, 2020), <https://digitalguardian.com/blog/what-polymorphic-malware-definition-and-best-practices-defending-against-polymorphic-malware#:~:text=Definition%20of%20Polymorphic%20Malware.bots%2C%20trojans%2C%20or%20keyloggers.>
- ⁵⁸ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.
- ⁵⁹ Trevor Hastie, et al., *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2nd ed.), Springer at 9-11 (Jan. 13, 2017), https://web.stanford.edu/~hastie/ElemStatLearn/printings/ESLII_print12_toc.pdf.
- ⁶⁰ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2021), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.
- ⁶¹ Richard S. Sutton & Andrew G. Barto, *Reinforcement Learning: An Introduction* (2nd ed.), MIT Press (2018).
- ⁶² *Key Considerations for Responsible Development and Fielding of Artificial Intelligence*, NSCAI (July 22, 2020), <https://www.nscai.gov/previous-reports/>.
- ⁶³ See Erico Guizzo, *What Is a Robot?*, IEEE (May 28, 2020), <https://robots.ieee.org/learn/what-is-a-robot/>.
- ⁶⁴ See Evan Ackerman, *Soft Self-Healing Materials for Robots That Cannot Be Destroyed*, IEEE (Sept. 5, 2019), <https://spectrum.ieee.org/automaton/robotics/robotics-hardware/soft-selfhealing-materials-for-robots-that-cannot-be-destroyed>.
- ⁶⁵ See Andrew Zisserman, *Self-Supervised Learning*, Google DeepMind (last accessed Feb. 17, 2021), <https://project.inria.fr/paiss/files/2018/07/zisserman-self-supervised.pdf>.
- ⁶⁶ Jacob Devlin, et al., *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*, arXiv (May 24, 2019), <https://arxiv.org/pdf/1810.04805.pdf>.
- ⁶⁷ Jianliang Meng, et al., *Overview of the Speech Recognition Technology*, 2012 Fourth International Conference on Computational and Information Sciences at 199-202 (2012), <https://ieeexplore.ieee.org/document/6300437/>.
- ⁶⁸ See *What Is Low-SWaP?*, REDCOM (last accessed Feb. 13, 2021), <https://www.redcom.com/what-is-low-swap-size-weight-and-power/>.
- ⁶⁹ Tony Roy, *Symbolic Logic: An Accessible Introduction to Serious Mathematical Logic* at 2-3 (Feb. 8, 2021), <https://tonyroypphilosophy.net/symbolic-logic/>.
- ⁷⁰ William LaPlante, *Owning the Technical Baseline*, Defense AT&L at 18-20, (July-Aug. 2015), <https://apps.dtic.mil/dtic/tr/fulltext/u2/1016084.pdf>.
- ⁷¹ See Jason Brownlee, *Supervised and Unsupervised Machine Learning Algorithms*, in *Machine Learning Algorithms*, Machine Learning Mastery (Aug. 20, 2020), <https://machinelearningmastery.com/supervised-and-unsupervised-machine-learning-algorithms/>.
- ⁷² See *What Is Virtual Reality?* Virtual Reality Society (last accessed Feb. 13, 2021), <https://www.vrs.org.uk/virtual-reality/what-is-virtual-reality.html>.

Appendix B: Acronyms Found in This Report

(alphabetical order)

A	
AAAI	Association for the Advancement of Artificial Intelligence
AaaS	applications as a service
AAMAS	Autonomous Agents and Multiagent Systems
AISC	AI Strategic Challenge
AAF	Adaptive Acquisition Framework
AAL	Army Applications Laboratory
ABMS	Advanced Battle Management System
ADL	Advanced Distributed Learning
AFC	Army Futures Command
AFOSR	Air Force Office of Scientific Research
AFRL	Air Force Research Lab
AGI	artificial general intelligence
ACM SIGKDD	Association for Computing Machinery's Special Interest Group on Knowledge Discovery and Data Mining
AI	artificial intelligence
AI CoE	AI Center of Excellence
AIM	Augmenting Intelligence using Machines
AIPfd	AI Partnership for Defense

Amii	Alberta Machine Intelligence Institute
ANPRM	Advance Notice of Proposed Rulemaking
API	Application Programming Interface
ArF	Argon fluoride
ARO	Army Research Office
ARPA	Academic Research Protection Act
ASC	Alternative Simplified Credit
ASIC	application-specific integrated circuit
ATO	Authorization (or Authority) to Operate
AVC	Bureau of Arms Control, Verification and Compliance

B

BA	Budget Activity
BARDA	Biomedical Advanced Research and Development Authority
BioMADE	Bioindustrial Manufacturing and Design Ecosystem
BIRD	Binational Industrial Research & Development Foundation
BIS	Bureau of Industry and Security
BSF	Binational Science Foundation

C

C2	command and control
C&ET	critical and emerging technologies
CBP	U.S. Customs and Border Protection
CBRS	Citizens Broadband Radio Service
CCMD	combatant command
CCP	Chinese Communist Party

CCW	Convention on Certain Conventional Weapons
CD	cardiovascular disease
CDC	Centers for Disease Control and Prevention
CDO	chief data officer
CFIUS	Committee on Foreign Investment in the United States
CFR	Code of Federal Regulations
CHIPS	Creating Helpful Incentives to Produce Semiconductors
CIA	Central Intelligence Agency
CIO	chief information officer
CISA	Cybersecurity and Infrastructure Security Agency
CMI	Component Mission Initiative
CNIPA	China National Intellectual Property Administration
COE	Center of Excellence
CONOPS	concept(s) of operations(s)
COTS	commercial off-the-shelf
COVID-19	coronavirus disease 2019
CReATE	Coding Repository and Transformation Environment
CRISPR	clustered regularly interspaced short palindromic repeats
CRCL	civil rights and civil liberties
CS	computer science
CSC	U.S. Cyberspace Solarium Commission
CSET Bureau	Bureau of Cyberspace Security and Emerging Technologies
CSIS	Center for Strategic and International Studies
CSTD	Comprehensive Science and Technology Dialogue
CTO	chief technology officer

D

DA	Decision Authority
DAC	Development Assistance Committee
DARPA	Defense Advanced Research Projects Agency
DDI	Bureau for Development, Democracy, and Innovation at USAID
DEXCOM	Deputies Executive Committee
DFC	U.S. International Development Finance Corporation
DFFT	data free flow with trust
DFI	development finance institution
DIA	Defense Intelligence Agency
DIB	Defense Innovation Board
DIU	Defense Innovation Unit
DHS	Department of Homeland Security
D/MR	Deputy Secretary of State for Management and Resources
DNA	deoxyribonucleic acid
DNI	Director of National Intelligence
DoD	Department of Defense
DoDD	Department of Defense Directive
DoE	Department of Energy
DOI	Department of the Interior
DOJ	Department of Justice
DOT	Department of Transportation
DOT&E	Director, Operational Test and Evaluation
DOTMLPF-P	doctrine, organization, training, materiel, leadership and education, personnel, facilities, and policy
DPC	Domestic Policy Council

DRL	Bureau of Democracy, Human Rights, and Labor
E	
EB	Bureau of Economic and Business Affairs
ECRA	Export Control Reform Act of 2018
EDT	emerging and disruptive technology
E.O.	Executive Order
EOP	Executive Office of the President
ERI	Electronics Resurgence Initiative
ESA	European Space Agency
ETC	Emerging Technology Coalition
ETTAC	Emerging Technology Technical Advisory Committee
EU	European Union
EUV	extreme ultraviolet
EXIM	Export-Import Bank of the United States
F	
FAIR	Facebook AI Research
FAR	Federal Acquisition Regulation
FARA	Foreign Agents Registration Act
FBI	Federal Bureau of Investigation
FCC	Federal Communications Commission
FCEN	Financial Crimes Enforcement Network
FDA	U.S. Food and Drug Administration
FFRDC	Federally Funded Research and Development Center
FIRRMA	Foreign Investment Risk Review Modernization Act of 2018

FISMA	Federal Information Security Modernization Act
FPGA	field-programmable gate array
FSI	Foreign Service Institute
FTQC	fault-tolerant quantum computer
FWCI	field-weighted citation impact
FY	fiscal year
FYDP	Future Years Defense Plan
G	
G20	Group of 20
GAN	generative adversarial network
GAO	U.S. Government Accountability Office
GDP	gross domestic product
GEC	Global Engagement Center at Department of State
GGE	Group of Governmental Experts
GIST	Global Innovation through Science and Technology
GPAI	Global Partnership on Artificial Intelligence
GPS	Global Positioning System
GPT-3	Generative Pre-trained Transformer 3
GPU	graphics processing unit
GSA	U.S. General Services Administration
H	
HPC	high-performance computing
HHMI	Howard Hughes Medical Institute
HHS	Health and Human Services

HR	human resources
HQE	highly qualified expert
HSI	human-system interactions
HUMINT	human intelligence
I	
I&W	indication(s) and warning(s)
IARPA	Intelligence Advanced Research Projects Activity
IC	U.S. Intelligence Community
IC ITE	Intelligence Community Information Technology Environment
ICRC	International Committee of the Red Cross
ICT	information and communications technology
IDDI	International Digital Democracy Initiative
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronic Engineers
IER	International Entrepreneur Rule
IFBHR	Internet Freedom and Business & Human Rights Section
IFI	international financial institution
IHL	International Humanitarian Law
IMINT	imagery intelligence
INL	Bureau of International Narcotics and Law Enforcement Affairs
IoT	internet of things
IP	intellectual property
IPEC	U.S. Intellectual Property Enforcement Coordinator
IPA	Intergovernmental Personnel Act
IPHE	International Partnership for Hydrogen and Fuel Cells in the Economy

ISAC	Information Sharing and Analysis Center
ISIS	Islamic State of Iraq and Syria
ISN	Bureau of International Security and Nonproliferation
ISO	International Organization for Standardization
ISR	intelligence, surveillance, and reconnaissance
ISTS	International Science and Technology Strategy
IT	information technology
ITF-CCAD	International Task Force to Counter and Compete Against Disinformation
IT SRMC	IT Modernization Senior Risk Management Council
ITU	International Telecommunication Union
ITU-T	International Telecommunication Union–Telecommunication Standardization Sector
IUSSTF	Indo-U.S. Science and Technology Forum

J

JAIC	Joint Artificial Intelligence Center
JCF	Joint Common Foundation
JCIDS	Joint Capabilities Integration and Development System
JIATF	Joint Interagency Task Force
JROC	Joint Requirements Oversight Council
JSP	Joint Strategic Plan
JWAC	Joint Warfare Analysis Center

K

K-12	kindergarten to 12th grade
------	----------------------------

L

LAWS	lethal autonomous weapon systems
LKIE	learning, knowledge, and information exchange
LOAC	Law of Armed Conflict

M

M&A	mergers and acquisitions
M&S	modeling and simulation
MAIEI	Montreal AI Ethics Institute
MAIRI	Multilateral AI Research Institute
MASINT	Measurement and signature intelligence
MCC	Millennium Challenge Corporation
MDA	Milestone Decision Authorities
MDAP	Major Defense Acquisition Program
MediFOR	Media Forensics
MEMT	Multi-Engine Machine Translation
Mila	Montreal Institute for Learning Algorithms
MIT	Massachusetts Institute of Technology
MITE	Malign Information Threat Executive
ML	machine learning

N

NAIRR	National Artificial Intelligence Research Resource
NASA	National Aeronautics and Space Administration
NATO	North Atlantic Treaty Organization
NBA	National Basketball Association

NCEAI	National Council for Expanding American Innovation
NCPS	National Cybersecurity Protection System
NCSC	National Counterintelligence and Security Center
NCTC	National Counterterrorism Center
NDAA	National Defense Authorization Act
NDEA	National Defense Education Act
NDS	National Defense Strategy
NEA	Nuclear Energy Agency
NEC	National Economic Council
NIFA	National Institute of Food and Agriculture
NIH	National Institutes of Health
NIJ	National Institute of Justice
NIS	National Intelligence Strategy
NISQ	noisy intermediate-scale quantum
NIST	National Institute of Standards and Technology
NITRD	Networking and Information Technology Research and Development
NLP	natural language processing
NLU	natural language understanding
NOAA	National Oceanic and Atmospheric Administration
NQCO	National Quantum Coordination Office
NQI	National Quantum Initiative
NRDC	National Reserve Digital Corps
NSF	National Science Foundation
NSTC	National Science and Technology Council
NSA	National Security Agency

NSC	National Security Council
NSF	National Science Foundation
NSIB	National Security Innovation Base
NSIN	National Security Innovation Network
NSS	National Security Strategy
NTF	National Technology Foundation
NTS	National Technology Strategy
NTIA	National Telecommunications and Information Administration
NVLAP	National Voluntary Laboratory Accreditation Program
O	
ODNI	Office of the Director of National Intelligence
OECD	Organisation for Economic Co-operation and Development
OES	Bureau of Oceans and International Environmental and Scientific Affairs
OISE	Office of International Science and Engineering
OMB	Office of Management and Budget
ONR	Office of Naval Research
OPM	U.S. Office of Personnel Management
ORSA	operational research and systems analysis
OSD	Office of the Secretary of Defense
OSINT	open-source intelligence
OSTP	Office of Science and Technology Policy
OTA	Other Transaction Authority
OUSD (A&S)	Office of the Under Secretary of Defense for Acquisition and Sustainment
OUSD (I&S)	Office of the Under Secretary of Defense for Intelligence and Security
OUSD (R&E)	Office of the Under Secretary of Defense for Research and Engineering

OZ	Opportunity Zone
P	
PaaS	platforms as a service
PAL	Permissive Action Link
PCAST	President's Council of Advisors on Science and Technology
P/CLR	Privacy, Civil Liberties, and Civil Rights
P/CRCL	Privacy, Civil Rights, and Civil Liberties
PCLOB	Privacy and Civil Liberties Oversight Board
PCT	Patent Cooperation Treaty
PDDNI	Principal Deputy Director of National Security
PE	program element
PED	processing, exploitation, and dissemination
PGNN	physics-guided neural network
PhD	doctoral graduate
PIA	Privacy Impact Assessment
PII	personally identifiable information
PLA	People's Liberation Army
PM	Bureau of Political-Military Affairs
PM	program manager
PoR	Program of Record
PPBE	Planning, Programming, Budget, and Execution
PPML	privacy-preserving machine learning

Q

QED-C	Quantum Economic Development Consortium
-------	---

QIS	Quantum Information Science
-----	-----------------------------

QPU	quantum processing unit
-----	-------------------------

R

R&D	research and development
-----	--------------------------

RAI	responsible AI
-----	----------------

RDT&E	research, development, test, and evaluation
-------	---

REN-ISAC	Research and Education Networks Information and Sharing Analysis Center
----------	---

RL	reinforcement learning
----	------------------------

RMF	Risk Management Framework
-----	---------------------------

RPA	robotic process automation
-----	----------------------------

S

S&E	science and engineering
-----	-------------------------

SBIR	Small Business Innovation Research Program
------	--

S/CCI	Office of the Coordinator for Cyber Issues
-------	--

SDK	software development kit
-----	--------------------------

SDO	standards developing organization
-----	-----------------------------------

SemaFor	semantic forensics
---------	--------------------

SEP	"standard essential" patents
-----	------------------------------

SFS	scholarship for service
-----	-------------------------

SGE	Special Government Employee
-----	-----------------------------

SIAC	Strategic Intelligence Analysis Cell
------	--------------------------------------

SIGINT	signals intelligence
--------	----------------------

SMART	Science, Mathematics, and Research for Transformation
SME	semiconductor manufacturing equipment
SMIC	Semiconductor Manufacturing International Corporation
SORN	System of Records Notice
SSD	Strategic Security Dialogue
S&T	science and technology
STAS	Office of the Science and Technology Adviser to the Secretary of State
State/Q	Under Secretary of State for Science, Research and Technology
STEM	science, technology, engineering, and mathematics
STTR	Small Business Technology Transfer Program
SWaP	size, weight, and power

T

TCC	Technology Competitiveness Council
T&E	test(ing) and evaluation
TET	Technology Engagement Team at Department of State
TEVV	test(ing) and evaluation, verification and validation
TRC	Technology Research Center
TSMC	Taiwan Semiconductor Manufacturing Corporation
TTCP	Technical Cooperation Program

U

UARC	University Affiliated Research Center
U.K.	United Kingdom
UN	United Nations
U.S.	United States

U.S.C.	United States Code
USAF	U.S. Air Force
USAID	U.S. Agency for International Development
USASOC	U.S. Army Special Operations Command
USCIS	U.S. Citizenship and Immigration Services
USD(A&S)	Under Secretary of Defense for Acquisition and Sustainment
USD(R&E)	Under Secretary of Defense for Research and Engineering
USDA	U.S. Department of Agriculture
USDSA	U.S. Digital Service Academy
USERRA	Uniformed Services Employment and Reemployment Rights Act
USISTEF	United States–India Science & Technology Endowment Fund
USPTO	U.S. Patent and Trademark Office
USTDA	U.S. Trade and Development Agency
USTR	Office of the U.S. Trade Representative

V

VA U.S. Department of Veterans Affairs

VCJCS Vice Chairman of the Joint Chiefs of Staff

W

WH White House

WIPO World Intellectual Property Organization

WLIF Warfighting Lab Incentive Fund

Numbers

3D three-dimensional

3SIIF Three Seas Initiative Investment Fund

5G fifth-generation standard for broadband cellular networks

Appendix C: Key Considerations for the Responsible Development and Fielding of Artificial Intelligence (Abridged)

Prefatory Note:

The paradigm and recommended practices described here stem from the Commission's line of effort dedicated to Ethics and Responsible Artificial Intelligence (AI). The Commission has recommended that heads of departments and agencies critical to national security (at a minimum, the Department of Defense, Intelligence Community, Department of Homeland Security, Federal Bureau of Investigation, Department of Energy, Department of State, and Department of Health and Human Services) should implement the Key Considerations as a paradigm for the responsible development and fielding of AI systems. This includes developing processes and programs aimed at adopting the paradigm's recommended practices, monitoring their implementation, and continually refining them as best practices evolve.

This approach would set the foundation for an intentional, government-wide, coordinated effort to incorporate recommended practices into current processes for AI development and fielding. However, our overarching aim is to allow agencies to continue to have the flexibility to craft policies and processes according to their specific needs. The Commission is mindful of the required flexibility that an agency needs when conducting the risk assessment and management of an AI system, as these tasks will largely depend on the context of the AI system.

This recommendation, along with a set of recommended considerations and practices, was made originally in July 2020. Here we present a revised and updated version as part of the Commission's Final Report. Many of the points made here are also reflected in Chapter 7 of the report.

The content herein is an abridged version of the content included in the extended version, which will be featured on NSCAI's website in March 2021 at www.nsc.ai.gov. In the more comprehensive document, we provide additional details and references for technical implementers.

Introduction

The Commission acknowledges the efforts undertaken to date to establish ethics guidelines for AI systems.¹ While some national security agencies have adopted,² or are in the process of adopting, AI principles,³ other agencies have not provided such guidance. In cases where principles are offered, it can be difficult to translate the high-level concepts into concrete actions. In addition, agencies would benefit from the establishment of greater consistency in policies to further the responsible development and fielding of AI technologies across government.

This Commission has identified five broad categories of challenges and made recommendations for both responsibly developing and fielding AI systems. These recommendations include immediate actions and future work the U.S. government should undertake to help establish best practices to overcome these challenges. Collectively, they form a paradigm for aligning AI system development and AI system behavior to goals and values. The first section, *Aligning Systems and Uses with American Values and the Rule of Law*, provides guidance specific to implementing systems that abide by American values, most of which are shared by democratic nations. The section also covers aligning the run-time behavior of systems to the related, more technical encodings of objectives, utilities, and trade-offs. The four following sections (on *Engineering Practices*, *System Performance*, *Human-AI Interaction*, and *Accountability & Governance*) serve in support of core American values and further outline practices needed to develop and field AI systems that are understandable, reliable, robust, and trustworthy.

Recommended practices span multiple phases of the AI lifecycle and establish a baseline for the responsible development and fielding of AI technologies. The Commission uses “development” to refer to “designing, building, and testing during development and prior to deployment” and “fielding” to refer to “deployment, monitoring, and sustainment.”

The Commission recommends that heads of departments and agencies implement the Key Considerations as a paradigm for the responsible development and fielding of AI systems. This includes developing policies and processes to adopt the paradigm's recommended practices, monitor their implementation, and continually refine them as best practices evolve. These recommended practices should apply both to systems that are developed by departments and agencies as well as to those that are acquired. Systems acquired (whether commercial off-the-shelf systems or through contractors) should be subjected to the same rigorous standards and recommended practices in the acquisitions and acceptance processes. As such, the government organization overseeing the bidding

process should require that vendors articulate how their practices align with the Key Considerations' recommended practices in their proposals, submissions, and bids.

In each of the five sections that follow, we first provide a conceptual overview of the scope and importance of the topic. We then illustrate examples of a current challenge relevant to national security departments that underscores the need to adopt recommended practices in this area. Then, we provide a list of recommended practices that agencies should adopt, acknowledging research, industry tools, and exemplary models within government that could support agencies in the adoption of recommended practices. Finally, in areas where best practices do not exist or are especially challenging to implement, we note the need for future work as a priority; this includes, for example, R&D and standards development. We also identify potential areas in which collaboration with allies and partners would be beneficial for interoperability and trust and note that the Key Considerations can inform potential future efforts to discuss military uses of AI with strategic competitors.

I. Aligning Systems and Uses with American Values and the Rule of Law

(1) Overview

Our values guide our decisions and our assessment of their outcomes. Our values shape our policies, our sensitivities, and how we balance trade-offs among competing interests. America's values, and commitment to upholding them, are reflected in the U.S. Constitution and U.S. laws, regulations, policies, and processes.

One of the seven principles we set forth in the Commission's Interim Report (November 2019) is the following:

The American way of AI must reflect American values—including having the rule of law at its core. For federal law enforcement agencies conducting national security investigations in the United States, that means using AI in ways that are consistent with constitutional principles of due process, individual privacy, equal protection, and non-discrimination. For American diplomacy, that means standing firm against uses of AI by authoritarian governments to repress individual freedom or violate the human rights of their citizens. And for the U.S. military, that means finding ways for AI to enhance its ability to uphold the laws of war and ensuring that current frameworks adequately cover AI.

Values established in the U.S. Constitution, and further operationalized in legislation, include freedoms of speech and assembly as well as the rights to due process, inclusion, fairness, non-discrimination (including equal protection), and privacy (including protection from unwarranted government interference in one's private affairs). These values are codified in the U.S. Constitution and the U.S. Code.⁴ International treaties that the United States has ratified also demonstrate our values by affirming our commitments to human rights and human dignity.⁵ Within America's national security departments, our commitment to

protecting and upholding privacy and civil liberties is further embedded in the policies and programs of the Intelligence Community (IC),⁶ the Department of Homeland Security,⁷ the Department of Defense (DoD),⁸ and oversight entities (e.g., the Privacy and Civil Liberties Oversight Board).⁹ In the military context, core values such as distinction and proportionality are embodied in the nation's commitment to, and the DoD's policies to uphold, the Uniform Code of Military Justice and the Law of Armed Conflict (LOAC).¹⁰

Other values are reflected in treaties, rules, and policies, such as the Convention Against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment¹¹; the DoD's Rules of Engagement¹²; and the DoD's directive concerning autonomy in weapon systems.¹³ While not an exhaustive list of U.S. values, the paradigm of considerations and recommended practices for AI that we introduce resonates with these values, as they have been acknowledged as critical by the U.S. government and national security departments and agencies. Further, many of these values are common to America's like-minded partners, who share a commitment to democracy, human dignity, and human rights.

Our values demand that the development and fielding of AI respect these foundational values and that they enable human empowerment as well as accountability. They require that the operation of AI systems and components be compliant with our laws and international legal commitments and with our departmental policies. In short, American values must inform the way we develop and field AI systems and the way our AI systems behave in the world.

(2) Examples of Current Challenges

Machine learning (ML) techniques can assist DoD agencies with large-scale data analyses to support and enhance decision-making about personnel. As an example, the Proposed New Disability Construct (PNDC) seeks to leverage data analyses to identify service members on the verge of ineligibility for deployment due to concerns with their readiness. Other potential analyses, including factors that lead to success or failure in promotion, can support personnel evaluations. Caution and proven practices are needed, however, to avoid pitfalls in fairness and inclusiveness, several of which have been highlighted in high-profile challenges in areas like criminal justice, recruiting and hiring, and face recognition.¹⁴ Attention should be paid to challenges with decision support systems like PNDC to avoid harmful disparate impact.¹⁵ Likewise, factors weighed in performance evaluations and promotions must be carefully considered to avoid inadvertently reinforcing existing biases through ML-assisted decisions.¹⁶

(3) Recommendations for Adoption

A. *Developing uses and building systems that behave in accordance with American values and the rule of law.* To implement core American values, it is important to:

1. *Employ technologies and operational policies that align with privacy preservation, fairness, inclusion, human rights, and the law of armed conflict (LOAC).* Technologies and policies throughout the AI lifecycle should support achieving these goals. They should ensure that AI uses and systems are consistent with these values and mitigate the risk that AI system uses/outcomes will violate these values.

- An explicit analysis of outcomes that would violate these values should be performed. Policy should prohibit disallowed outcomes that would violate the values above. During system development, analysis of system-specific disallowed outcomes should be performed.¹⁷ As the technology advances, applications evolve, and our understanding of the implications of use grows, these policies should periodically be refreshed.

B. *Representing objectives and trade-offs.* Another important practice for aligning AI systems with values is to consider values as (1) embodied in choices about engineering trade-offs and (2) explicitly represented in the goals and utility functions of an AI system.¹⁸ Recommended practices for representing objectives and trade-offs include the following:

1. *Consider and document value considerations in AI systems by specifying how trade-offs with accuracy are handled.* This includes documenting the choices made when selecting operating thresholds that have implications for performance, such as the ratio of true positive and false positive rates or the precision (how many selected items are relevant?) versus recall (how many relevant items are selected?). For example, consider a system designed to recommend if a person entering the U.S. should be pulled aside for more detailed inspection and interview. Precision refers to how many of the people selected for additional processing are valid security concerns; recall refers to how many valid security concerns are flagged for added processing. The trade-off is between allowing a valid security concern to slip past review and detaining persons who are not a security concern. Setting thresholds to increase precision (i.e., reduce the number of persons detained needlessly) will drive down recall (i.e., detain fewer valid security concerns).

2. *Consider and document value considerations in AI systems that rely on representations of objective or utility functions,* especially when assigning weighting that captures the importance of different goals for the system. As an illustration of multiple goals and value weights, consider shopping for a new car. A buyer may identify factors that are important in the decision, such as gas mileage, safety, reliability, and performance. These clearly interact in some cases—for example, gas mileage and performance are likely in tension, and safety is likely correlated partly with vehicle size, which is likely in

tension with gas mileage. When reviewing a set of new cars, the best pick for a buyer will depend on the priorities placed on these factors.

3. *Conduct documentation, reviews, and set limits that reflect disallowed outcomes (through constraints on allowed performance) to ensure compliance with values.*

(4) Recommendations for Future Action

Future R&D. R&D is needed to advance capabilities for preserving and ensuring that developed or acquired AI systems will act in accordance with American values and the rule of law. For instance, the Commission notes the need for R&D to assure that the personal privacy of individuals is protected in the acquisition and use of data for AI system development. This includes advancing ethical practices with the use of personal data, including disclosure and consent about data collection and use models (including uses of data to build base models that are later retrained and fine-tuned for specific tasks), the use of anonymity techniques and privacy-preserving technologies, and uses of related technologies such as multiparty computation (to allow collaboration on the pooling of data from multiple organizations without sharing data sets). Additionally, we need to understand the compatibility of data usage policies and privacy-preserving approaches with regulatory approaches such as the European Union’s General Data Protection Regulation (GDPR).

II. Engineering Practices

(1) Overview

The government and its partners (including vendors), should adopt recommended practices for creating and maintaining trustworthy and robust AI systems that are *auditable* (able to be interrogated and yield information at each stage of the AI lifecycle to determine compliance with policy, standards, or regulations¹⁹); *traceable* (to understand the technology, development processes, and operational methods applicable to AI capabilities, for example with transparent and auditable methodologies, data sources, and design procedure and documentation²⁰); *interpretable* (to understand the value and accuracy of system output²¹); and *reliable* (to perform in the intended manner within the intended domain of use²²). There are no broadly directed best practices or standards to guide organizations in the building of AI systems that are consistent with designated AI principles, but potential approaches, minimal standards, and engineering proven practices are available.²³

Additionally, several properties of the engineering methods and models used in ML (e.g., data-centric methods) are associated with weaknesses that make the systems brittle and exploitable in specific ways—and vulnerable to failure modalities not seen in traditional software systems. Such failures can rise inadvertently or as the intended results of malicious attacks and manipulation.²⁴ Recent frameworks integrate adversarial attacks²⁵ and unintended faults throughout the lifecycle²⁶ into a single taxonomy that describes both intentional and unintentional failure modes.²⁷

Intentional failures are the result of malicious actors explicitly attacking some aspect of AI system behavior. Taxonomies (e.g., from NIST) on malicious attacks explain the rapidly developing Adversarial Machine Learning (AML) landscape. Attacks span ML training and testing, and each has associated defenses.²⁸ Categories of intentional failures introduced by adversaries include *training data poisoning* attacks (contaminating training data), *model inversion* (recovering training data used in the model through careful queries), and *ML supply chain attacks* (compromising the ML model as it is being downloaded for use).²⁹ National security uses of AI will be the subject of sustained adversarial efforts; AI developed for this community must remain current with a rapidly developing understanding of the nature of vulnerabilities to attacks as these attacks grow in sophistication. Technical and process advances that contribute to reducing vulnerability and to detecting and alerting about attacks must also be monitored routinely.

Unintentional failures can be introduced at any point in the AI development and deployment lifecycle. In addition to faults that can be inadvertently introduced into any software development effort, distinct additional failure modes can be introduced for ML systems.

Examples of unintentional AI failure modes include *reward hacking* (when AI systems learn to achieve a programmed goal in a way that contradicts the programmer's intent) and *distributional shifts* (when a system is tested in one kind of environment but is unable to adapt to changes in other kinds of environments).³⁰ Another area of failure is the inadequate specification of objectives (as described in Section 1 above on *Representing Objectives and Trade-offs*), leading to unexpected and costly behaviors and outcomes.³¹ As AI systems that are separately developed and tested are composed and interact with other AI systems (within one's own services, forces, and agencies, and between U.S. systems and those of allies, adversaries, and potential adversaries), additional unintentional failures can occur.³²

(2) Examples of Current Challenges

To make high-stakes decisions, and often in safety-critical contexts, the DoD and IC must be able to depend on the integrity and security of the data used to train some kinds of ML systems. The challenges of doing so have been echoed by the leadership of the DoD and the IC,³³ including concerns with detecting adversarial attacks such as data poisoning.

(3) Recommendations for Adoption

Critical engineering practices needed to operationalize AI principles (such as “traceable” and “reliable”³⁴) are described in the non-exhaustive list below. These practices span development and fielding of AI systems.

1. *Refine design and development requirements, informed by the concept of operations and risk assessment*, including characterization of failure modes and associated impacts. Conduct systems analysis of operations and identify mission success

metrics and potential functions that can be performed by AI technology. Incorporate early analyses of use cases and scenario development, assess general feasibility and compliance with disallowed outcomes expressed in policy. Critically assess reproducibility (how readily research results can be replicated by a third party) and technical maturity. This includes broad stakeholder engagement and hazard analysis with multidisciplinary experts who ask key questions about potential disparate impacts and document the process undertaken to ensure fairness and the lack of unwanted bias in the ML application.³⁵ The feasibility of meeting these requirements may trigger a review of whether and where it is appropriate to use AI in the system being proposed.

- *Risk assessment.* Trade-offs and risks, including a system's potential societal impact, should be discussed with a diverse, interdisciplinary group. This includes an analysis of the system's potential societal impact and of the impacts of the system's failure modes. Risk-assessment questions should be asked about critical areas relevant to the national security context, including privacy and civil liberties, LOAC, human rights,³⁶ system security, and the risks of a new technology being leaked, stolen, or weaponized.³⁷

2. *Produce documentation of the AI lifecycle.* Whether building and fielding an AI system or "infusing AI" into a preexisting system, require documentation in certain areas.³⁸ These include the data used in ML technologies and the origin of the data³⁹; algorithm(s) used to build models, model characteristics, and intended uses of the AI capabilities; connections between and dependencies within systems, and associated potential complications; the selected testing methodologies, performance indicators, and results for models used in the AI component; and required maintenance (including re-testing requirements) and technical refresh (including for when a system is used in a different scenario/setting or if the AI system is capable of online learning or adaptation).

3. *Leverage infrastructure to support traceability, including auditability and forensics.* Invest resources and build capabilities that support the traceability of AI systems. Traceability captures key information about the system's development and deployment process for relevant personnel to adequately understand the technology.⁴⁰ Audits should support analyses of specific actions and characterizations of longer-term performance and assure that performance on tests of the system and on real-world workloads meet requirements.

4. *For security and robustness, address intentional and unintentional failures.*

- *Adversarial attacks and use of robust ML methods.* Expand notions of adversarial attacks to include various ML attacks⁴¹ (as described above) and seek latest technologies that demonstrate the ability to detect and notify operators of attacks and also tolerate attacks (i.e., to enable systems to withstand or to degrade gracefully when targeted by a deliberate attack).⁴²

- *Follow and incorporate advances in intentional and unintentional ML failures.* Given the rapid evolution of the field of study of intentional and unintentional ML failures, national security organizations must follow and adapt to the latest knowledge about failures and proven practices for system monitoring, failure detection, engineering, and protections during operation. Related efforts and R&D focus on developing and deploying robust AI methods.⁴³
- *Adopt a DevSecOps lifecycle for AI systems focused on potential failure modes.* This includes developing and regularly refining threat models to capture and characterize various attacks, establish a matrixed focus for developing and refining threat models, and ensuring DevSecOps addresses ML development, fielding, and when ML systems are under attack.⁴⁴
- *Limit consequences of system failure through system architecture.* Build an overall system architecture that monitors component performance and handles errors when anomalies are detected; build AI components to be self-protecting and self-checking; and include aggressive stress testing under conditions of intended use.

5. *Conduct red teaming* for both intentional and unintentional failure modalities. Bring together multiple perspectives to rigorously challenge AI systems, exploring the risks, limitations, and vulnerabilities in the context in which they'll be deployed (i.e., red teaming).

- To mitigate intentional failure modes, assume an offensive posture and use methods to make systems more resistant to adversarial attacks, work with adversarial testing tools, and deploy teams dedicated to trying to break systems and make them violate rules for appropriate behavior.⁴⁵
- To mitigate unintentional failure modes, test ML systems per a thorough list of realistic conditions they are expected to operate in. When selecting third-party components, consider the impact that a security vulnerability in them could have on the security of the larger system into which they are integrated. Have an accurate inventory of third-party components and a plan to respond when new vulnerabilities are discovered.
- Organizations should consider establishing broader enterprise-wide communities of AI red teaming capabilities that could be applied to multiple AI developments (e.g., at a DoD service or IC element level, or higher).

(4) Recommendations for Future Action

- *Documentation strategy.* As noted in our First Quarter Recommendations, a common documentation strategy is needed to ensure sufficient documentation by all national security departments and agencies.⁴⁶ In the meantime, agencies should pilot documentation approaches across the AI lifecycle to help inform such a strategy.
- *Standards.* To improve traceability, future work is needed by standard-setting bodies, alongside national security departments/agencies and the broader AI community, to develop audit trail requirements per mission needs for high-stakes AI systems including safety-critical applications (e.g., weapon system controls).
- *Future R&D.* R&D is needed to advance capabilities for cultivating more robust methods that can overcome adverse conditions; to advance approaches that enable assessment of types and levels of vulnerability and immunity; and to tolerate attacks. R&D is also needed to advance capabilities to support risk assessment, including standards, methods, and metrics for evaluating degrees of auditability, traceability, interpretability, explainability, and reliability. For interpretability in particular, R&D is also needed to improve our understanding of the efficacy of interpretability tools and possible interfaces.

III. System Performance

(1) Overview

Fielding AI systems in a responsible manner includes establishing confidence that the technology will perform as intended. An AI system's performance must be assessed,⁴⁷ including assessing its capabilities and blind spots with data representative of real-world scenarios or with simulations of realistic contexts,⁴⁸ and its reliability, robustness (i.e., resilience in real-world settings, including withstanding adversarial attacks on AI components), and security during development and deployment.⁴⁹ System performance must also measure compliance with requirements derived from values such as fairness.

Testing protocols and requirements are essential for measuring and reporting on system performance. (Here, "testing" broadly refers to what the DoD calls "Test and Evaluation, Verification and Validation" [TEVV]. This testing includes both what DoD refers to as Developmental Test and Evaluation and Operational Test and Evaluation.) AI systems present new challenges to established testing protocols and requirements as they increase in complexity, particularly for operational testing. However, existing methods like high-fidelity performance traces and means for sensing shifts (e.g., changes in the statistical distribution of data in operation versus model training) allow for the continuous monitoring of an AI system's performance.

When evaluating system performance, it is especially important to take into account holistic, end-to-end system behavior—the consequence of the interactions and relationships among system elements rather than the independent behavior of individual elements. While system engineering and national security communities have focused on system of systems engineering for years, specific attention must be paid to undesired interactions and emergent performance in AI systems. Multiple relatively independent AI systems can be viewed as distinct agents interacting in the environment of the system of systems, and some of these agents will be humans in and on the loop. Industry has encountered and documented problems in building “systems of systems” out of multiple AI systems.⁵⁰ A related problem is encountered when the performance of one model in a pipeline changes, degrading the overall pipeline behavior.⁵¹ As America’s AI-intensive systems may increasingly be composed and/or interoperable with allied AI-intensive systems, these become important topics for coordination with allies.

(2) Examples of Current Challenges

Unexpected interactions and errors commonly occur in integrated simulations and exercises, illustrating the challenges of predicting and managing behaviors of systems composed of multiple components. Intermittent failures can transpire after composing different systems; these failures are not necessarily the result of any one component having errors, but rather are due to the interactions of the composed systems.⁵²

(3) Recommendations for Adoption

Critical practices to ensure optimal system performance are described in the following non-exhaustive list:

A. Model training and model testing procedures should cover key aspects of performance and appropriate performance metrics.

1. Use regularly updated standards for testing and reporting of system performance. Standards for metrics and reporting are needed to adequately:
 - a. Achieve consistency across testing and test reporting for critical areas.
 - b. Test for blindspots.⁵³
 - c. Test for fairness. When testing for fairness, conduct sustained fairness assessments throughout development and deployment and document deliberations made on the appropriate fairness metrics to use. Agencies should conduct outcome and impact analysis to detect when subtle assumptions in the system show up as unexpected and undesired outcomes in the operational environment.⁵⁴
 - d. Articulate system performance. Clearly document system performance and communicate to the end user the meaning/significance of such performance metrics.

2. *Consider and document the representativeness of the data and model for the specific context at hand.* When using classification and prediction technologies, explicitly consider and document challenges with representativeness of data used in analyses and the fairness/accuracy of inferences and recommendations made with systems leveraging that data when applied in different populations/contexts.

3. *Evaluate an AI system's performance relative to current benchmarks* where possible. Such benchmarks should assist in determining if a proposed AI system's performance meets or exceeds current best performance.

4. *Evaluate aggregate performance of human-machine teams.* Consider that the current benchmark might be the current best performance of a human operator or the composed performance of the human-machine team. Where humans and machines interact, it is important to measure the aggregate performance of the team rather than the AI system alone.⁵⁵

5. *Provide sustained attention to reliability and robustness.* Employ tools and techniques to carefully bound assumptions of robustness of the AI component in the larger system architecture. Provide sustained attention to characterizing the actual performance (for normal and boundary conditions) throughout development and deployment.⁵⁶ For systems of particularly high potential consequences of failure, considerable architecture and design work will have been put into making the overall system fail-safe.

6. *For systems of systems, test machine-machine/multi-agent interaction.* Individual AI systems will be combined in various ways in an enterprise to accomplish broader missions beyond the scope of any single system, which can introduce its own problems.⁵⁷ As a priority during testing, challenge (or "stress test") interfaces and usage patterns with boundary conditions and assumptions about the operational environment and use.

B. Maintenance and deployment

Given the dynamic nature of AI systems, best practices for maintenance are also critically important. Recommended practices include:

1. *Specify maintenance requirements* for datasets as well as for systems, given that their performance can degrade over time.⁵⁸

2. *Continuously monitor and evaluate AI system performance*, including the use of high-fidelity traces to determine continuously if a system is going outside of acceptable parameters.⁵⁹

3. *Conduct iterative model testing and validation.* Training and testing that provide characteristics on capabilities might not transfer or generalize to specific settings of usage; thus, testing and validation may need to be done recurrently, and at strategic intervention points, but especially for new deployments and classes of tasks.⁶⁰

4. *Monitor and mitigate emergent behavior.* There will be instances when systems are composed in ways not anticipated by the developers, thus requiring monitoring the actual performance of the composed system and its components.

(4) *Recommendations for Future Action*

- *Future R&D.* R&D is needed to advance capabilities for TEVV of AI systems to better understand how to conduct persistent and iterative TEVV and build checks and balances into an AI system. Improved methods are needed to explore, predict, and control individual AI system behavior so that when AI systems are composed into systems of systems, their interaction does not lead to unexpected negative outcomes.
- *Metrics.* Progress on a common understanding of TEVV concepts and requirements is critical for progress in widely used metrics for performance. Significant work is needed to establish what appropriate metrics should be used to assess system performance across attributes for responsible AI according to applications/context profiles. (Such attributes, for example, include fairness, interpretability, reliability, and robustness.) Future work is needed to develop: (1) definitions, taxonomy, and metrics needed to enable agencies to better assess AI performance and vulnerabilities; and (2) metrics and benchmarks to assess reliability and intelligibility of produced model explanations. In the near term, guidance is needed on: (1) standards for testing intentional and unintentional failure modes; (2) exemplar data sets for benchmarking and evaluation, including robustness testing and red teaming; and (3) defining characteristics of AI data quality and training environment fidelity (to support adequate performance and governance).⁶¹
- *International collaboration and cooperation.* Collaboration is needed to align on how to test and verify AI system reliability and performance, including along shared values (such as fairness and privacy). Such collaboration will be critical among allies and partners for interoperability and trust. Additionally, these efforts could potentially include dialogues between the U.S. and strategic competitors on establishing common standards of AI safety and reliability testing to reduce the chances of inadvertent escalation.

IV. Human-AI Interaction & Teaming

(1) *Overview*

Responsible AI development and fielding requires striking the right balance of leveraging human and AI reasoning, recommendation, and decision-making processes. Ultimately,

all AI systems will have some degree of human-AI interaction as they all will be developed to support humans. And some systems will serve as more than just support tools and will adopt roles of teammates that actively collaborate with humans.

(2) Examples of Current Challenges

There is an opportunity to develop AI systems to complement and augment human understanding, decision-making, and capabilities. Decisions about developing and fielding AI systems for specific domains or scenarios should consider the relative strengths of AI capabilities and human intellect across the expected range of tasks, considering AI system maturity or capability and how people and machines might coordinate.

Designs and methods for human-AI interaction can be employed to enhance human-AI teaming.⁶² Methods in support of effective human-AI interaction can help AI systems understand when and how to engage humans for assistance, when AI systems should take initiative to assist human operators, and, more generally, how to support the creation of effective human-AI teams. In engaging with end users, it may be important for AI systems to infer and share with end users well-calibrated levels of confidence about their inferences, to provide human operators with an ability to weigh the importance of machine output or pause to consider details behind a recommendation more carefully. Methods, representations, and machinery can be employed to provide insight about AI inferences, including the use of interpretable machine learning.⁶³

Research directions include developing and fielding machinery aimed at reasoning about human strengths and weaknesses, such as recognizing and responding to the potential for costly human biases of judgment and decision-making in specific settings.⁶⁴ Other work centers on mechanisms to consider the ideal mix of initiatives, including when and how to rely on human expertise versus on AI inferences.⁶⁵ As part of effective teaming, AI systems can be endowed with the ability to detect the focus of attention, workload, and sensitivity to interruption of human operators and consider these inferences in decisions about when and how to engage with operators.⁶⁶ Directions of effort include developing mechanisms for identifying the most relevant information or inferences to provide end users with different skill levels in different settings.⁶⁷ Consideration must be given to the prospect of introducing bias, including potential biases that may arise because of the configuration and sequencing of rendered data. For example, IC research⁶⁸ shows that confirmation bias can be triggered by the order in which information is displayed, and this order can consequently impact or sway intel analyst decisions. Careful design and study can help to identify and mitigate such bias.

(3) Recommendations for Adoption

Critical practices to ensure optimal human-AI interaction are described in the non-exhaustive list below. These recommended practices span the entire AI lifecycle.

A. Identification of functions of humans in design, engineering, and fielding of AI.

1. Given AI and human capabilities and complementarities, as well as requirements for accountability and human judgment, define the tasks of humans and the goals and mission of the human-machine team across the AI lifecycle. This entails noting needs for feedback loops, including opportunities for oversight.

2. Define functions and responsibilities of humans during system operation and assign them to specific individuals. Functions and responsibilities will vary for each domain and project and should be periodically revisited.

B. Explicit support of human-AI interaction and collaboration.

1. *Extend human-AI design methodologies and guidelines.* AI systems designs should take into account the defined tasks of humans in human-AI collaborations in different scenarios; ensure that the mix of human-machine actions in the aggregate is consistent with the intended behavior and accounts for the ways that human and machine behavior can co-evolve⁶⁹; and also avoid automation bias (that places unjustified confidence in the results of the computation) and unjustified reliance on humans in the loop as fail-safe mechanisms. Practices should allow for auditing of the human-AI pair and designs should be transparent to allow for an understanding of how the AI is working day-to-day, supported by an audit trail if things go wrong. Based on context and mission need, designs should ensure usability of AI systems by AI experts, domain experts, and novices, as appropriate.

2. *Employ algorithms and functions in support of interpretability and explanation.* Algorithms and functions that provide individuals with task-relevant knowledge and understanding should take into account that key factors in an AI system's inferences and actions can be understood differently by various audiences (e.g., real-time operators, engineers and data scientists, and oversight officials). Interpretability and explainability exists in degrees. In this regard, interpretability intersects with traceability, audit, and documentation practices.

3. *Design systems to provide cues to human operator(s) about the level of confidence the system has in its results or behaviors.* AI system designs should appropriately convey uncertainty and error bounding. For instance, a user interface should convey system self-assessment of confidence alerts when the operational environment is significantly different from the environment the system was trained for and indicate internal inconsistencies that call for caution.

4. *Refine policies for machine-human initiative and handoff.* Policies, and aspects of human-computer interaction, system interface, and operational design, should define when and how information or tasks should be passed from a machine to a human operator and vice versa.

5. *Leverage traceability to assist with system development and understanding.* Traceability processes must capture details about human-AI interaction to retroactively understand where challenges occurred, and why, in order to improve systems and their use for redress. Infrastructure and instrumentation⁷⁰ can also help assess humans, systems, and environments to gauge the impact of AI at all levels of system maturity and to measure the effectiveness and performance for hybrid human-AI systems in a mission context.

6. *Conduct training.* Train and educate individuals responsible for AI development and fielding, including human operators, decision-makers, and procurement officers.⁷¹

(4) Recommendations for Future Action

- *Future R&D.* R&D is needed to advance capabilities of AI technologies to perceive and understand the meaning of human communication, including spoken speech, written text, and gestures. This research should account for varying languages and cultures, with special attention to diversity given that AI often performs worse in cases impacting gender and racial minorities. It is also needed to improve human-machine teaming, including disciplines and technologies centered on decision sciences, control theory, psychology, economics (human aspects and incentives), and human factors engineering. R&D for human-machine teaming should also focus on helping systems understand human blind spots and biases and optimizing factors such as human attention, human workload, ideal mixing of human and machine initiative, and passing control between the human and machine. R&D also is needed to optimize the ability of humans and AI to work together to undertake complex, evolving tasks in a variety of environments, as well as for diverse groupings of machines to cooperate with each other, with broader systems, and with human counterparts to achieve shared objectives.

- *Training.* Ongoing work is needed to train the workforce that will interact with, collaborate with, and be supported by AI systems. In its First Quarter Recommendations, the Commission provided recommendations for such training. Operators should receive training on the specifics of the system and application, the fundamentals of AI and data science, and refresher trainings (e.g., when systems are deployed in new settings and unfamiliar scenarios, and when predictive models are revised with new data, as performance may shift with updates and introduce behaviors unfamiliar to operators).

V. Accountability and Governance

(1) Overview

National security departments and agencies must specify who will be held accountable for both specific system outcomes and general system maintenance and auditing, in what way, and for what purpose. Government must address the difficulties in preserving human accountability, including for end users, developers, testers, and the organizations employing AI systems. End users and those affected by the actions of an AI system should have the opportunity to appeal an AI system's determinations. Accountability and appellate processes must exist for AI decisions, inferences, recommendations, and actions.

(2) Examples of Current Challenges

If a contentious outcome occurs, overseeing entities need the technological capacity to understand what in the AI system caused this. For example, if a soldier uses an AI-enabled weapon and the result violates international law of war standards, an investigating body or military tribunal should be able to re-create what happened through audit trails and other documentation. Without policies requiring such technology and the enforcement of those policies, proper accountability would be elusive, if not impossible. Moreover, auditing trails and documentation will prove critical as courts begin to grapple with whether AI system determinations reach the requisite standards to be admitted as evidence. Building the traceability infrastructure to permit auditing (as described in *Engineering Practices*) will increase the costs of building AI systems and take significant work—a necessary investment given our commitment to accountability, discoverability, and legal compliance.

(3) Recommendations for Adoption

Critical accountability and governance practices are identified in the non-exhaustive list below.

1. *Appoint full-time responsible AI leads* to join senior leadership. Every department and agency critical to national security and each branch of the armed services, at a minimum, should have a dedicated, full-time responsible AI lead who is part of the senior leadership team. Such leads should oversee the implementation of the Key Considerations recommended practices alongside the department or agency's respective AI principles.

2. *Identify responsible actors.* Determine and document the people accountable for a specific AI system or any given part of the system and the processes involved. This includes identifying who is responsible for the development or procurement; operation (including the system's inferences, recommendations, and actions during usage), and maintenance of an AI system, as well as the authorization of a system and enforcement of policies for use. Determine and document the mechanism/structure for holding such actors accountable and to whom it should be disclosed for proper oversight.

3. *Require technology to strengthen accountability processes and goals.* Document the chains of custody and command involved in developing and fielding AI systems to know who was responsible at which point in time. Improving traceability and auditability capabilities will allow agencies to better track a system's performance and outcomes.⁷² Policy should establish requirements about information that should be captured about the development process and about system performance and behavior in operation.

4. *Adopt policies to strengthen accountability and governance.* Identify or, if lacking, establish policies that allow individuals to raise concerns about irresponsible AI development/fielding (e.g., via an ombudsman). This requires ensuring a governance structure is in place to address grievances and harms if systems fail, which supports feedback loops and oversight to ensure that systems operate as they should.

Agencies should institute specific oversight and enforcement practices, including auditing and reporting requirements; a mechanism that would allow thorough review of the most sensitive/high-risk AI systems to ensure auditability and compliance with responsible use and fielding requirements; an appealable process for those found at fault for developing or using AI irresponsibly; and grievance processes for those affected by the actions of AI systems. Agencies should leverage best practices from academia and industry for conducting internal audits and assessments,⁷³ while also acknowledging the benefits offered by external audits.⁷⁴

5. *Support external oversight.* Remain responsive and facilitate oversight through documentation processes and other policy decisions.⁷⁵ For instance, supporting traceability and specifically documentation to audit trails will allow for external oversight.⁷⁶ Self-assessment alone might prove to be inadequate in all scenarios.⁷⁷ Congress can provide a key oversight function throughout the AI lifecycle, asking critical questions of agency leaders and those responsible for AI systems.

(4) Recommendations for Future Action

Currently no external oversight mechanism exists specific to AI in national security. Notwithstanding the important work of Inspectors General in conducting internal oversight, open questions remain as to how to complement current practices and structures.

Appendix C - Endnotes

¹ Examples of efforts to establish ethics guidelines are found within the U.S. government, industry, and internationally. See, e.g., *Draft Memorandum for the Heads of Executive Departments and Agencies: Guidance for Regulation of Artificial Intelligence Applications*, Office of Management and Budget (Jan. 1, 2019), <https://www.whitehouse.gov/wp-content/uploads/2020/01/Draft-OMB-Memo-on-Regulation-of-AI-1-7-19.pdf>; Jessica Fjeld & Adam Nagy, *Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI*, Berkman Klein Center (Jan. 15, 2020), <https://cyber.harvard.edu/publication/2020/principled-ai>; *OECD Principles on AI*, OECD (last visited June 17, 2020), <https://www.oecd.org/going-digital/ai/principles/>; *Ethics Guidelines for Trustworthy AI*, European Commission at 26-31 (April 8, 2019), <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>; *Assessment List for Trustworthy Artificial Intelligence (ALTAI) for Self-assessment*, European Commission (July 17, 2020), <https://ec.europa.eu/digital-single-market/en/news/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment>.

² C. Todd Lopez, *DOD Adopts 5 Principles of Artificial Intelligence Ethics*, U.S. Department of Defense (Feb. 5, 2020), <https://www.defense.gov/Explore/News/Article/Article/2094085/dod-adopts-5principles-of-artificial-intelligence-ethics/> [hereinafter Lopez, DoD Adopts 5 Principles].

³ See Ben Huebner, *Presentation: AI Principles*, Intelligence and National Security Alliance 2020 Spring Symposium: Building an AI-Powered IC (March 4, 2020), <https://www.insaonline.org/2020-spring-symposium-building-an-ai-powered-ic-event-recap/>.

⁴ See, e.g., U.S. Const. amendments I, IV, V, and XIV; Americans with Disabilities Act of 1990, 42 U.S.C. § 12101 et seq.; Title VII of the Consumer Credit Protection Act, 15 U.S.C. §§ 1691-1691f; Title VII of the Civil Rights Act of 1964, 42 U.S.C. § 2000e et seq.

⁵ International Covenant on Civil and Political Rights, UN General Assembly, United Nations, Treaty Series, vol. 999, at 171 (Dec. 16, 1966), <https://www.refworld.org/docid/3ae6b3aa0.html>. As noted in the Commission's *Interim Report*, America and its like-minded partners share a commitment to democracy, human dignity, and human rights. *Interim Report*, NSCAI (Nov. 2019), <https://www.nscai.gov/previous-reports/>. Many, but not all nations, share commitments to these values. Even when values are shared, however, they can be culturally relative, for instance, across nations, owing to interpretative nuances.

⁶ See, e.g., Daniel Coats, *Intelligence Community Directive 107*, Office of the Director of National Intelligence (Feb. 28, 2018), <https://fas.org/irp/dni/icd/icd-107.pdf> (on protecting civil liberties and privacy); *IC Framework for Protecting Civil Liberties and Privacy and Enhancing Transparency Section 702*, Intel.gov (Jan. 2020), https://www.intelligence.gov/index.php/ic-on-the-record/guide-to-posted-documents#SECTION_702-OVERVIEW (on privacy and civil liberties implication assessments and oversight); *Principles of Professional Ethics for the Intelligence Community*, Office of the Director of National Intelligence (last accessed June 17, 2020), <https://www.dni.gov/index.php/who-we-are/organizations/clpt/clpt-related-menus/clpt-related-links/ic-principles-of-professional-ethics> (on diversity and inclusion).

⁷ See, e.g., *Privacy Office*, U.S. Department of Homeland Security (last accessed June 3, 2020), <https://www.dhs.gov/privacy-office#>; *CRCL Compliance Branch*, U.S. Department of Homeland Security (last accessed May 15, 2020), <https://www.dhs.gov/compliance-branch>.

⁸ See Samuel Jenkins & Alexander Joel, *Balancing Privacy and Security: The Role of Privacy and Civil Liberties in the Information Sharing Environment*, IAPP Conference 2010 (2010), <https://dpcl.dod.defense.gov/Portals/49/Documents/Civil/IAPP.pdf>.

⁹ See *Projects*, U.S. Privacy and Civil Liberties Oversight Board (last visited June 17, 2020), <https://www.pclob.gov/Projects>.

¹⁰ See *Department of Defense Law of War Manual*, U.S. Department of Defense (Dec. 2016), <https://dod.defense.gov/Portals/1/Documents/pubs/DoD%20Law%20of%20War%20Manual%20-%20June%202015%20Updated%20Dec%202016.pdf?ver=2016-12-13-172036-190> [hereinafter DoD Law of War Manual]; see also *AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense: Supporting Document*, DoD Defense Innovation Board (Oct. 31, 2019), https://media.defense.gov/2019/Oct/31/2002204459/-1/-1/0/DIB_AI_Principles_supporting_document.pdf ("More than 10,000 military and civilian lawyers within DoD advise on legal compliance with regard to the entire range of DoD activities, including the Law of War. Military lawyers train DoD personnel on Law of War requirements, for example, by providing additional Law of War instruction prior to a deployment of forces abroad. Lawyers for a Component DoD organization advise on the

issuance of plans, policies, regulations, and procedures to ensure consistency with Law of War requirements. Lawyers review the acquisition or procurement of weapons. Lawyers help administer programs to report alleged violations of the Law of War through the chain of command and also advise on investigations into alleged incidents and on accountability actions, such as commanders' decisions to take action under the Uniform Code of Military Justice. Lawyers also advise commanders on Law of War issues during military operations.”).

¹¹ Convention against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment, United Nations General Assembly (Dec. 10, 1984), <https://www.ohchr.org/en/professionalinterest/pages/cat.aspx>.

¹² See DoD Law of War Manual at 26 (“Rules of Engagement reflect legal, policy, and operational considerations, and are consistent with the international law obligations of the United States, including the law of war.”).

¹³ See *Department of Defense Directive 3000.09 on Autonomy in Weapon Systems*, U.S. Department of Defense (Nov. 21, 2012), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/300009p.pdf> (“Autonomous and semi-autonomous weapon systems shall be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force.”).

¹⁴ See, e.g., *Report on Algorithmic Risk Assessment Tools in the U.S. Criminal Justice System*, Partnership on AI, <https://www.partnershiponai.org/report-on-machine-learning-in-risk-assessment-tools-in-the-u-s-criminal-justice-system/>; Jeffrey Dastin, *Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women*, Reuters (Oct. 10, 2018), <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazonscraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G> [*hereinafter* Dastin, Amazon Scraps Secret AI Recruiting Tool]; Andi Peng et al., *What You See Is What You Get? The Impact of Representation Criteria on Human Bias in Hiring*, Proceedings of the 7th AAAI Conference on Human Computation and Crowdsourcing (Oct. 2019), <https://arxiv.org/pdf/1909.03567.pdf>; Patrick Grother, et al., *Face Recognition Vendor Test (FRVT) Part Three: Demographic Effects*, National Institute of Standards and Technology (Dec. 2019), <https://doi.org/10.6028/NIST.IR.8280>.

¹⁵ PNDC provides predictive analytics to improve military readiness; enable earlier identification of service members with potential unfitting, disabling, or career-ending conditions; and offer opportunities for early medical intervention or referral into disability processing. To do so, PNDC provides recommendations at multiple points in the journey of the non-deployable service member through the Military Health System to make “better decisions” that improve medical outcomes and delivery of health services. This is very similar to the OPTUM decision support system that recommended which patients should get additional intervention to reduce costs. Analysis showed millions of U.S. patients were processed by the system, with substantial disparate impact on Black patients compared to white patients. Shaping development from the start to reflect bias issues (which can be subtle) would have produced a more equitable system and avoided scrutiny and suspension of system use when findings were disclosed. Heidi Ledford, *Millions of Black People Affected by Racial Bias in Health Care Algorithms*, Nature (Oct. 26, 2019), <https://www.nature.com/articles/d41586-019-03228-6>.

¹⁶ See e.g., Dastin, Amazon Scraps Secret AI Recruiting Tool.

¹⁷ This combined approach of stable policy-level disallowed outcomes and system-specific disallowed outcomes is consistent with DoD practices for system safety, for example. See *Department of Defense Standard Practice: System Safety*, U.S. Department of Defense (May 11, 2012), <https://www.dau.edu/cop/armysoh/DAU%20Sponsored%20Documents/MIL-STD-882E.pdf>. Depending on the context, mitigating harm per values and disallowed outcomes might entail the use of fail-safe technologies. See Eric Horvitz, *Reflections on Safety and Artificial Intelligence, Exploratory Technical Workshop on Safety and Control for AI* (June 27, 2016), http://erichorvitz.com/OSTP-CMU_AI_Safety_framing_talk.pdf. See also Dorsa Sadigh & Ashish Kapoor, *Safe Control Under Uncertainty with Probabilistic Signal Temporal Logic*, Proceedings of Robotics: Science and Systems XII (2016), <https://www.microsoft.com/en-us/research/wp-content/uploads/2016/11/RSS2016.pdf>.

¹⁸ Mohsen Bayati, et al., *Data-Driven Decisions for Reducing Readmissions for Heart Failure: General Methodology and Case Study*, PLOS One Medicine (Oct. 8, 2014), <https://doi.org/10.1371/journal.pone.0109264>; Eric Horvitz & Adam Seiver, *Time-Critical Action: Representations and Application*, Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence (Aug. 1997), <https://arxiv.org/pdf/1302.1548.pdf>.

Appendix C - Endnotes

¹⁹ See Inioluwa Deborah Raji, et al., *Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing*, ACM FAT (Jan. 3, 2020), <https://arxiv.org/abs/2001.00973> [hereinafter Raji, Closing the AI Accountability Gap].

²⁰ See Lopez, DoD Adopts 5 Principles.

²¹ *Model Interpretability in Azure Machine Learning*, Microsoft (Nov. 16, 2020), <https://docs.microsoft.com/en-us/azure/machine-learning/how-to-machine-learning-interpretability>.

²² Lopez, DoD Adopts 5 Principles.

²³ Jessica Cussins Newman, *Decision Points in AI Governance: Three Case Studies Explore Efforts to Operationalize AI Principles* (May 5, 2020), Berkeley Center for Long-Term Cybersecurity, <https://cltc.berkeley.edu/ai-decision-points/>; Raji, *Closing the AI Accountability Gap*; Miles Brundage, et al., *Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims* (April 20, 2020), <https://arxiv.org/abs/2004.07213> [hereinafter Brundage, *Toward Trustworthy AI Development*]; Saleema Amershi, et al., *Software Engineering for Machine Learning: A Case Study*, Microsoft (March 2019), https://www.microsoft.com/en-us/research/uploads/prod/2019/03/amershi-icse-2019-Software_Engineering_for_Machine_Learning.pdf.

²⁴ Dario Amodei, et al., *Concrete Problems in AI Safety*, arXiv (July 25, 2016), <https://arxiv.org/abs/1606.06565>.

²⁵ Guofu Li, et al., *Security Matters: A Survey on Adversarial Machine Learning*, arXiv (Oct. 23, 2018), <https://arxiv.org/abs/1810.07339>; Elham Tabassi et al., *NISTIR 8269: A Taxonomy and Terminology of Adversarial Machine Learning (Draft)*, National Institute of Standards and Technology (Oct. 2019), <https://csrc.nist.gov/publications/detail/nistir/8269/draft>.

²⁶ José Faria, *Non-Determinism and Failure Modes in Machine Learning*, 2017 IEEE 28th International Symposium on Software Reliability Engineering Workshops (Oct. 2017), <https://ieeexplore.ieee.org/document/8109300>.

²⁷ Ram Shankar Siva Kumar, et al. *Failure Modes in Machine Learning* (Nov. 11, 2019), <https://docs.microsoft.com/en-us/security/engineering/failure-modes-in-machine-learning> [hereinafter Kumar, *Failure Modes in Machine Learning*].

²⁸ See Elham Tabassi et al., *NISTIR 8269: A Taxonomy and Terminology of Adversarial Machine Learning (Draft)*, National Institute of Standards and Technology (Oct. 2019), <https://csrc.nist.gov/publications/detail/nistir/8269/draft>. See also Kumar, *Failure Modes in Machine Learning*.

²⁹ For 11 categories of attack, and associated overviews, see the Intentionally-Motivated Failures Summary in Kumar, *Failure Modes in Machine Learning*.

³⁰ For more on reward hacking, see Jack Clark, et al., *Faulty Reward Functions in the Wild* (Dec. 21, 2016), <https://openai.com/blog/faulty-reward-functions/>. For more on distributional shifts, see Colin Smith, et al., *Hazard Contribution Modes of Machine Learning Components*, AAAI-20 Workshop on Artificial Intelligence Safety (SafeAI 2020) (Feb. 7, 2020), <https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20200001851.pdf> (Unexpected performance represents emergent runtime output, behavior, or effects at the system level, e.g., through unanticipated feature interaction ... that was also not previously observed during model validation.).

³¹ Thomas Dietterich & Eric Horvitz, *Rise of Concerns About AI: Reflections and Directions*, Communications of the ACM at 38-40 (Oct. 2015), http://erichorvitz.com/CACM_Oct_2015-VP.pdf. See also Ziad Obermeyer et al., *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, Science (Oct. 25, 2019), <https://science.sciencemag.org/content/366/6464/447>.

³² Kumar, *Failure Modes in Machine Learning*.

³³ For concerns about generative adversarial networks (GANS) voiced by Gen. Shanahan, JAIC, see Don Rassler, *A View from the CT Foxhole: Lieutenant General John N.T. "Jack" Shanahan, Director, Joint Artificial Intelligence Center, Department of Defense*, Combating Terrorism Center at West Point (Dec. 2019), <https://ctc.usma.edu/view-ct-foxhole-lieutenant-general-john-n-t-jack-shanahan-director-joint-artificial-intelligence-center-department-defense/>. Concerns about GANS, information authenticity, and reliable and understandable systems were voiced by Dean Souleles, IC. See *Afternoon Keynote*, Intelligence and National Security Alliance 2020 Spring Symposium: Building an AI-Powered IC (March 4, 2020), <https://www.insonline.org/2020-spring-symposium-building-an-ai-powered-ic-event-recap/>.

³⁴ See Lopez, DOD Adopts 5 Principles.

³⁵ There is no single definition of fairness. System developers and organizations fielding applications must work with stakeholders to define fairness and provide transparency via disclosure of assumed definitions of fairness. Definitions or assumptions about fairness and metrics for identifying fair inferences and allocations should be explicitly documented. This should be accompanied by a discussion of alternate definitions and rationales for the current choice. These elements should be documented internally as ML components and larger systems are developed. This is especially important as establishing alignment on the metrics to use for assessing fairness encounters an added challenge when different cultural and policy norms are involved when collaborating on development and use with allies.

³⁶ For more on the importance of human rights impact assessments of AI systems, see *Report of the Special Rapporteur to the General Assembly on AI and Its Impact on Freedom of Opinion and Expression*, UN Human Rights Office of the High Commissioner (Aug. 29, 2018), <https://www.ohchr.org/EN/Issues/FreedomOpinion/Pages/ReportGA73.aspx>. For an example of a human rights risk assessment for AI in categories such as nondiscrimination and equality, political participation, privacy, and freedom of expression, see Mark Latonero, *Governing Artificial Intelligence: Upholding Human Rights & Dignity*, Data Society (Oct. 2018), https://datasociety.net/wp-content/uploads/2018/10/DataSociety_Governing_Artificial_Intelligence_Upholding_Human_Rights.pdf.

³⁷ For exemplary risk assessment questions that IARPA has used, see Richard Danzig, *Technology Roulette: Managing Loss of Control as Many Militaries Pursue Technological Superiority*, Center for a New American Security at 22 (June 28, 2018), <https://s3.amazonaws.com/files.cnas.org/documents/CNASReport-Technology-Roulette-DoSproof2v2.pdf?mtime=20180628072101>.

³⁸ Documentation recommendations build off of a legacy of robust documentation requirements. See *Department of Defense Standard Practice: Documentation of Verification, Validation, and Accreditation (VV&A) For Models and Simulations*, Department of Defense (Jan. 28, 2008), <https://acqnotes.com/Attachments/MIL-STD-3022%20Documentation%20of%20VV&A%20for%20Modeling%20%20Simulation%2028%20Jan%2008.pdf>.

³⁹ For an industry example, see Timnit Gebru, et al., *Datasheets for Datasets*, Microsoft (March 2018), <https://www.microsoft.com/en-us/research/publication/datasheets-for-datasets/>. For more on data, model, and system documentation, see *Annotation and Benchmarking on Understanding and Transparency of Machine Learning Lifecycles (ABOUT ML)*, an evolving body of work from the Partnership on AI about documentation practices at <https://www.partnershiponai.org/about-ml/>. Documenting caveats of re-use for both data sets and models is critical to avoid "off-label" use harms, as one senior official notes. David Thornton, *Intelligence Community Laying Foundation for AI Data Analysis*, Federal News Network (Nov. 1, 2019), <https://federalnewsnetwork.com/allnews/2019/11/intelligence-community-laying-the-foundation-for-ai-data-analysis/>.

⁴⁰ Jonathan Mace, et al., *Pivot Tracing: Dynamic Causal Monitoring for Distributed Systems*, Communications of the ACM, Vol. 63 No. 3, at 94-102 (March 2020), <https://m-cacm.acm.org/magazines/2020/3/243034-pivot-tracing/fulltext> [hereinafter Mace, Pivot Tracing].

⁴¹ Aleksander Madry, et al., *Towards Deep Learning Models Resistant to Adversarial Attacks*, MIT (Sept. 4, 2019), <https://arxiv.org/abs/1706.06083> [hereinafter Madry, Towards Deep Learning Models Resistant to Adversarial Attacks].

⁴² See e.g., *id.*; Thomas Dietterich, *Steps Toward Robust Artificial Intelligence*, Association for the Advancement of Artificial Intelligence (Fall 2017), <https://www.aaai.org/ojs/index.php/aimagazine/article/view/2756/2644>; Eric Horvitz, *Reflections on Safety and Artificial Intelligence* (June 27, 2016), http://erichorvitz.com/OSTP-CMU_AI_Safety_framing_talk.pdf.

Appendix C - Endnotes

⁴³ On adversarial attacks on ML, see Kevin Eykholt, et al., *Robust Physical-World Attacks on Deep Learning Visual Classification*, IEEE Conference on Computer Vision and Pattern Recognition at 1625-1634 (June 18-23, 2018), <https://ieeexplore.ieee.org/document/8578273>. On directions with robustness, see Madry, *Towards Deep Learning Models Resistant to Adversarial Attacks*. For a more exhaustive list of sources see *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

⁴⁴ Ram Shankar Siva Kumar, et al., *Adversarial Machine Learning—Industry Perspectives*, 2020 IEEE Symposium on Security and Privacy (SP) Deep Learning and Security Workshop (May 21, 2020), <https://arxiv.org/abs/2002.05646>.

⁴⁵ Dou Goodman, et al., *Advbox: A Toolbox to Generate Adversarial Examples That Fool Neural Networks* (Aug. 26, 2020), <https://arxiv.org/abs/2001.05574>.

⁴⁶ See *First Quarter Recommendations*, NSCAI (March 2020), <https://www.nscai.gov/previous-reports/>. Ongoing efforts to share best practices for documentation among government agencies through GSA's AI Community of Practice further indicate the ongoing need and desire for common guidance.

⁴⁷ Ben Shneiderman, *Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy*, International Journal of Human-Computer Interaction 2020 at 495-504 (March 23, 2020), <https://doi.org/10.1080/10447318.2020.1741118> [hereinafter Shneiderman, *Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy*].

⁴⁸ However, test protocols must acknowledge that test sets may not be fully representative of real-world usage.

⁴⁹ Brundage, *Toward Trustworthy AI Development*; Ece Kamar, et al., *Combining Human and Machine Intelligence in Large-Scale Crowdsourcing*, Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (June 2012), <https://dl.acm.org/doi/10.5555/2343576.2343643> [hereinafter Kamar, *Combining Human and Machine Intelligence in Large-Scale Crowdsourcing*].

⁵⁰ One example is “Hidden Feedback Loops,” where systems that learn from external-world behavior may also shape the behavior they are monitoring. See D. Sculley, et al., *Machine Learning: The High Interest Credit Card of Technical Debt*, Google (2014), <https://research.google/pubs/pub43146/>.

⁵¹ Megha Srivastava, et al., *An Empirical Analysis of Backward Compatibility in Machine Learning Systems*, KDD'20 (Aug. 11, 2020), <https://arxiv.org/abs/2008.04572> [hereinafter Srivastava, *An Empirical Analysis of Backward Compatibility in Machine Learning Systems*].

⁵² David Sculley, et al., *Hidden Technical Debt in Machine Learning Systems*, Proceedings of the 28th International Conference on Neural Information Processing Systems (Dec. 2015), <https://dl.acm.org/doi/10.5555/2969442.2969519>.

⁵³ Ramya Ramakrishnan, et al., *Blind Spot Detection for Safe Sim-to-Real Transfer*, Journal of Artificial Intelligence Research 67 at 191-234 (Feb. 4, 2020), <https://www.jair.org/index.php/jair/article/view/11436>.

⁵⁴ See Microsoft's AI Fairness checklist as an example of an industry tool to support fairness assessments; Michael A. Madaio, et al., *Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI*, CHI 2020 (April 25-30, 2020), <http://www.jennvw.com/papers/checklists.pdf> [hereinafter Madaio, *Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI*].

⁵⁵ Kamar, *Combining Human and Machine Intelligence in Large-Scale Crowdsourcing*.

⁵⁶ See Shneiderman, *Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy*.

⁵⁷ Cynthia Dwork, et al., *Individual Fairness in Pipelines*, arXiv (April 12, 2020), <https://arxiv.org/abs/2004.05167>; Srivastava, *An Empirical Analysis of Backward Compatibility in Machine Learning Systems*.

⁵⁸ *Artificial Intelligence (AI) Playbook for the U.S. Federal Government*, Artificial Intelligence Working Group, ACT-IAC Emerging Technology Community of Interest (Jan. 22, 2020), <https://www.actiac.org/act-iac-white-paper-artificial-intelligence-playbook>.

⁵⁹ Ori Cohen, *Monitor! Stop Being A Blind Data-Scientist*, Towards Data Science (Oct. 8, 2019), <https://towardsdatascience.com/monitor-stop-being-a-blind-data-scientist-ac915286075f>; Mace, Pivot Tracing at 94-102.

⁶⁰ Eric Breck, et al., *The ML Test Score: A Rubric for ML Production Readiness and Technical Debt Reduction*, 2017 IEEE International Conference on Big Data (Dec. 11-14, 2017), <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8258038&tag=1>.

⁶¹ The 2021 NDAA expansion of the National Institute of Standards & Technology (NIST) mission authorizes the standards body to provide such guidance: “National Institute of Standards and Technology Activities (Title LIII, Sec. 5301)—expands NIST mission to include advancing collaborative frameworks, standards, guidelines for AI, supporting the development of a risk-mitigation framework for AI systems, and supporting the development of technical standards and guidelines to promote trustworthy AI systems.” Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

⁶² Saleema Amershi, et al., *Guidelines for Human-AI Interaction*, CHI '19: Proceedings of the CHI Conference on Human Factors in Computing Systems (May 2019), <https://dl.acm.org/doi/10.1145/3290605.3300233>.

⁶³ Rich Caruana, et al., *Intelligible Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-day Readmission*, Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Aug. 10-13, 2015), <https://www.semanticscholar.org/paper/Intelligible-Models-for-HealthCare%3APredicting-Risk-Caruana-Lou/cb030975a3dbcdf52a01cbd1c140711332313e13>.

⁶⁴ Eric Horvitz, *Reflections on Challenges and Promises of Mixed-Initiative Interaction*, AI Magazine (Summer 2007), http://erichorvitz.com/mixed_initiative_reflections.pdf.

⁶⁵ Eric Horvitz, *Principles of Mixed-Initiative User Interfaces*, Proceedings of CHI '99 ACM SIGCHI Conference on Human Factors in Computing Systems (May 1999), <https://dl.acm.org/doi/10.1145/302979.303030>; Kamar, Combining Human and Machine Intelligence in Large-Scale Crowdsourcing.

⁶⁶ Eric Horvitz, et al., *Models of Attention in Computing and Communications: From Principles to Applications*, Communications of the ACM at 52-59 (March 2003), <https://cacm.acm.org/magazines/2003/3/6879-models-of-attention-in-computingand-communication/fulltext>.

⁶⁷ Eric Horvitz & Matthew Barry, *Display of Information for Time-Critical Decision Making*, Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence (Aug. 1995), <https://arxiv.org/pdf/1302.4959.pdf>.

⁶⁸ There has been considerable research in the IC on the challenges of confirmation bias for analysts. Some experiments demonstrated a strong effect that the sequence in which information is presented alone can shape analyst interpretations and hypotheses. Brant Cheikes, et al., *Confirmation Bias in Complex Analyses*, MITRE (Oct. 2004), https://www.mitre.org/sites/default/files/pdf/04_0985.pdf. This highlights the care that is required when designing the human-machine teaming when complex, critical, and potentially ambiguous information is presented to analysts and decision-makers.

⁶⁹ Shneiderman, Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy at 495-504. An example of co-evolution of machine and human behavior is in ML spam filters. As human spammers determine what characteristics are getting email flagged as spam, they change how they generate spam, which requires the spam-detection models to evolve in a constant “arms race.”

⁷⁰ Infrastructure includes tools (hardware and software) in the test environment that support monitoring system performance (such as the timing of exchanges among systems or the ability to generate test data). Instrumentation refers to the presence of monitoring and additional interfaces to provide insight into a specific system under test.

⁷¹ Jamie Berryhill, et al., *Hello, World: Artificial Intelligence and Its Use in the Public Sector*, OECD Working Papers on Public Governance (Nov. 21, 2019), <https://doi.org/10.1787/726fd39d-en>.

Appendix C - Endnotes

⁷² See Raji, Closing the AI Accountability Gap.

⁷³ See *Id.* (“In this paper, we present internal algorithmic audits as a mechanism to check that the engineering processes involved in AI system creation and deployment meet declared ethical expectations and standards, such as organizational AI principles”); see also Madaio, Co-Designing Checklists to Understand Organizational Challenges and Opportunities Around Fairness in AI.

⁷⁴ For more on the benefits of external audits, see Brundage, Toward Trustworthy AI Development. For an agency example, see Aaron Boyd, *CBP Is Upgrading to a New Facial Recognition Algorithm in March*, Nextgov.com (Feb. 7, 2020), <https://www.nextgov.com/emerging-tech/2020/02/cbp-upgrading-new-facialrecognition-algorithm-march/162959/> (highlighting a NIST algorithmic assessment on behalf of U.S. Customs and Border Protection).

⁷⁵ Maranke Wieringa, *What to Account for When Accounting for Algorithms*, Proceedings of the 2020 ACM FAT Conference (Jan. 2020), <https://doi.org/10.1145/3351095.3372833>.

⁷⁶ Raji, Closing the AI Accountability Gap.

⁷⁷ Brundage, Toward Trustworthy AI Development.

Technical Glossary to the Key Considerations Appendix

This glossary provides a working set of definitions specific to the NSCAI Key Considerations. The Commission acknowledges that the definitions of the terms below may diverge from other scholarly or government definitions and were developed to be accessible to a broad audience.

AI Component: A software object that uses AI, meant to interact with other components, encapsulating certain functionality or a set of functionalities. An AI component has a clearly defined interface and conforms to a prescribed behavior common to all components within an architecture.¹

AI Lifecycle: The steps for managing the lifespan of an AI system: 1) Specify the system's objective. 2) Build model. 3) Test the AI system. 4) Deploy and maintain the AI system. 5) Engage in a feedback loop with continuous training and updates.²

AI System: A system designed or adapted to interact with an anticipated operational environment to achieve one or more intended purposes while complying with applicable constraints and that uses AI to provide a substantial part of its capabilities.³

Artificial Intelligence (AI): The ability of a computer system to solve problems and to perform tasks that have traditionally required human intelligence to solve.

Auditability: A characteristic of an AI system in which its software and documentation can be interrogated and yield information at each stage of the AI lifecycle to determine compliance with policy, standards, or regulations.

DevSecOps: Enhanced engineering practices that improve the lead time and frequency of delivery outcomes, promoting a more cohesive collaboration between development, security, and operations teams as they work toward continuous integration and delivery.

Differential Privacy: A criterion for a strong, mathematical definition of privacy in the context of statistical and ML analysis used to enable the collection, analysis, and sharing of a broad range of statistical estimates, such as averages, contingency tables, and synthetic data, based on personal data while protecting the privacy of the individuals in the data.⁴

False Negative: An example in which the predictive model mistakenly classifies an item as in the negative class. For example, a false negative describes the situation in which a junk-email model specifies that a particular email message is not spam (the negative class), when the email message actually is spam, leading to frustration of the junk message appearing in an end user's inbox.⁵ In a higher-stakes example, a false negative captures the case in which a medical diagnostic model misses identifying a disease that is present in a patient.

False Positive: An example in which the model mistakenly classifies an item as in the positive class. For example, the model inferred that a particular email message was spam (the positive class), but that email message was actually not spam, leading to delays in an end user reading a potentially important message.⁶ In a higher-stakes situation, a false positive describes the situation in which a disease is diagnosed as present when the disease is not present, potentially leading to unnecessary and costly treatments.

High-Fidelity Performance Traces: A commonly used technique useful in debugging and performance analysis. Concretely, trace recording implies detection and storage of relevant events during run-time, for later off-line analysis. High fidelity traces refers to the amount of fine-grained detail captured in the traces.⁷

Human Factors Engineering: The discipline that takes into account human strengths and limitations in the design of interactive systems that involve people, tools and technology, and work environments to ensure safety, effectiveness, and ease of use.⁸

Human in the Loop: The term describes a system architecture in which active human judgment and engagement are part of the operation of a system, and a human is an integral part of the system behavior. An example is the human operator of a remotely piloted vehicle or a decision support system that makes recommendations for a human to decide on.

Human on the Loop: This term describes a system architecture in which a human has a supervisory role in the operation of the system but is not an integral part of the system behavior. An example is an operator monitoring a fleet of warehouse robots—they operate autonomously but can be shut down if the operator determines something is wrong.

Machine Learning (ML): The study or the application of computer algorithms that improve automatically through experience.⁹ Machine learning algorithms build a model based on training data in order to perform a specific task, like aiding in prediction or decision-making processes, without necessarily being explicitly programmed to do so.

Model Testing: Testing assesses the performance of a trained model against new, previously unseen inputs, to demonstrate that the model generalizes to produce accurate results beyond just the training data.¹⁰

Model Training: Training a model simply means learning (determining) good values for all of the internal parameters that determine the model's performance. In supervised learning, for example, a machine learning model is trained by examining many labeled examples and attempting to find a model that minimizes the discrepancies between the real (labelled) values and the values produced by the model.¹¹

Technical Glossary to the Key Considerations Appendix

Multi-Party Federated Learning: A machine learning architecture in which many clients (e.g., mobile devices or whole organizations) collaboratively train a model under the orchestration of a central server (e.g., a service provider) while keeping the training data decentralized. It can mitigate many of the systemic privacy risks and costs resulting from traditional, centralized machine learning and data science approaches.¹² However, it does introduce new attack vectors that must be addressed.¹³

Precision: A metric for classification models. Precision identifies the frequency with which a model was correct when classifying the positive class. It answers the question “How many selected positive items are true positive?” For example, the percentage of messages flagged as spam that are spam.¹⁴

Privacy-Preserving AI: Techniques for protecting the privacy of people associated with the training data from adversarial attacks. These techniques include federated learning and differential privacy.¹⁵

Recall: A metric for classification models. Recall identifies the frequency with which a model correctly classifies the true positive items. It answers the question “How many true positive items were correctly classified?” For example, the percentage of spam messages that were flagged as spam.¹⁶

Reliable AI: An AI system that performs in its intended manner within the intended domain of use.

Robust AI: An AI system that is resilient in real-world settings, such as an object-recognition application that is robust to significant changes in lighting. The phrase also refers to resilience when it comes to adversarial attacks on AI components.

Run-Time Behavior: The behavior of a program while it is executing (i.e., running on one or more processors).

Trustworthy AI: Trustworthy AI has three components: (1) it should be lawful, ensuring compliance with all applicable laws and regulations; (2) it should be ethical, demonstrating respect for, and ensuring adherence to, ethical principles and values; and (3) it should be robust, both from a technical and social perspective, because, even with good intentions, AI systems can cause unintentional harm.¹⁷

Technical Glossary to the Key Considerations Appendix - Endnotes

¹ See NIST, *NISTIR 7298 Rev. 3, Glossary of Key Information Security Terms* (July 2019), <https://csrc.nist.gov/glossary/term/component>.

² Note that for data-driven AI systems step 2 is expanded and replaced with 2.a) Acquire data to meet the objective, and 2.b) Train the AI system on the data; and these two steps are usually repeated, with data acquisition and training continuing until desired performance objectives are attained. For further discussion on the ML lifecycle, see Saleema Amershi, et al., *Software Engineering for Machine Learning: A Case Study*, IEEE Computer Society (May 2019), <https://www.microsoft.com/en-us/research/publication/software-engineering-for-machine-learning-a-case-study/>.

³ See Hilary Sillitto, et al., *Systems Engineering and System Definitions*, International Council on Systems Engineering, (Jan. 8, 2019), https://www.incose.org/docs/default-source/default-document-library/final_se-definition.pdf.

⁴ Kobbi Nissim, et al., *Differential Privacy: A Primer for a Non-technical Audience*, Working Group of the Privacy Tools for Sharing Research Data Project, Harvard University, (Feb. 14, 2018), https://privacytools.seas.harvard.edu/files/privacytools/files/pedagogical-document-dp_new.pdf.

⁵ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.

⁶ *Id.*

⁷ See Johan Kraft, et al., *Trace Recording for Embedded Systems: Lessons Learned from Five Industrial Projects*, Runtime Verification at 315-329, https://link.springer.com/chapter/10.1007%2F978-3-642-16612-9_24.

⁸ See *Human Factors Engineering*, U.S. Department of Health and Human Services: Agency for Healthcare Research and Quality (Sept. 2019), <https://psnet.ahrq.gov/primer/human-factors-engineering>.

⁹ Thomas M. Mitchell, *Machine Learning*, McGraw-Hill (1997).

¹⁰ See Rob Ashmore, et al., *Assuring the Machine Learning Lifecycle: Desiderata, Methods, and Challenges*, arXiv at 4 (May 2019), <https://arxiv.org/abs/1905.04223>.

¹¹ See *Descending into ML: Training and Loss*, Google (last accessed Feb. 15, 2021), <https://developers.google.com/machine-learning/crash-course/descending-into-ml/training-and-loss>.

¹² Peter Kairouz, et al., *Advances and Open Problems in Federated Learning*, arXiv (Dec. 10, 2019), <https://arxiv.org/pdf/1912.04977.pdf>.

¹³ See Vale Tolpegin, et al., *Data Poisoning Attacks Against Federated Learning Systems*, ArXiv (Aug. 11, 2020), <https://arxiv.org/abs/2007.08432>; Arjun Nitin Bhagoji, et al., *Analyzing Federated Learning Through an Adversarial Lens*, arXiv (Nov. 25, 2019), <https://arxiv.org/abs/1811.12470>.

¹⁴ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.

¹⁵ For a discussion on how privacy-preserving machine learning works, see Roxanne Heston & Helon Toner, *Have Your Data and Use It Too: A Federal Initiative for Protecting Privacy While Advancing AI*, Day One Project (Jan. 23, 2020), <https://www.dayoneproject.org/post/have-your-data-and-use-it-too-a-federal-initiative-for-protecting-privacy-while-advancing-ai>; see also Georgios Kaissis, et al., *Secure, Privacy-Preserving and Federated Machine Learning in Medical Imaging*, Nature Machine Intelligence at 305-311 (June 8, 2020), <https://doi.org/10.1038/s42256-020-0186-1>.

¹⁶ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.

¹⁷ See *Ethics Guidelines for Trustworthy AI*, European Commission: High-Level Expert Group on Artificial Intelligence at 5 (April 8, 2019), <https://ec.europa.eu/futurium/en/ai-alliance-consultation>.

Appendix D:

Draft Legislative Language

The following legislative text represents the Commission staff's best efforts to capture the Commission's final recommendations in legislative form. The Commission defers to the House and Senate members, staff, and legislative counsels as to appropriate drafting.

CHAPTER 1: EMERGING THREATS IN THE AI ERA

Blueprint for Action

Combatting Malign Information Operations Enabled by AI

Recommendation: A National Strategy for the Global Information Domain.

Congress should direct the Executive Branch to transmit a National Strategy for the Global Information Domain that categorizes the global information domain as an arena of competition vital to the national security of the United States.

SEC. ____.—NATIONAL STRATEGY FOR THE GLOBAL INFORMATION DOMAIN.—

(a) IN GENERAL.—Not later than 270 days after the date of the enactment of this Act, the President shall transmit to Congress a National Strategy for the Global Information Domain that addresses the global information domain as an arena of competition vital to the national security of the United States.

(b) ISSUES ADDRESSED.—The National Strategy for the Global Information Domain required by subsection (a) shall, at a minimum:

(1) Prioritize the global information domain as an arena for international competition;

(2) Detail how adversarial state and non-state actors are attempting to define and control the global information domain in order to shape global opinion and achieve strategic advantage;

(3) Account for the critical role of artificial intelligence-enabled malign information in the efforts of adversarial state and non-state actors to achieve these goals;

(4) Identify and prioritize actions to defend, counter, and compete against malign information operations as a national security threat;

(5) As necessary, update critical infrastructure designations and require relevant departments and agencies to update sector-specific plans to reflect emerging technologies; and

(6) Establish organizational structures for U.S. national security agencies to counter and compete against the threat.

CHAPTER 2: FOUNDATIONS OF FUTURE DEFENSE

Blueprint for Action

Recommendation: Drive Change through Top-Down Leadership.

In the National Defense Authorization Act (NDAA) for Fiscal Year (FY) 2022, establish a Steering Committee on Emerging Technology and National Security Threats and designate that it be tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.

SEC. ____.—ROLE OF INTELLIGENCE COMMUNITY IN STEERING COMMITTEE ON EMERGING TECHNOLOGY.—

Section 236 of the William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, is amended—

(1) in subsection (b), by—

(A) redesignating paragraph (8) as paragraph (9); and

(B) inserting the following new paragraph before redesignated paragraph (9):

“(8) One or more representatives of the Intelligence Community, to include the Principal Deputy Director of National Intelligence.”

(2) by redesignating paragraph (c) as paragraph (d); and inserting the following new paragraph before redesignated paragraph (d):

“(c) LEADERSHIP.—The Steering Committee shall be chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.”

The Steering Committee on Emerging Technology recommendation is also featured in Chapters 3 and 5.

Recommendation: Build the Technical Backbone.

Prioritize funding for the Department’s digital ecosystem and associated activities. The Armed Services Committees should use the FY 2022 NDAA to direct the Department of

Defense to develop a resourcing plan for the digital ecosystem that establishes, sustains, and incentivizes use of its various components as enterprise-wide, enduring resources. The Committees should also authorize the obligation of funds to begin work on the ecosystem.

SEC. ____.—RESOURCING PLAN FOR DIGITAL ECOSYSTEM.—

(a) IN GENERAL.—Within one year after the date of the enactment of this Act, the Secretary of Defense shall develop a plan for the development of a modern digital ecosystem that embraces state of the art tools and modern processes to enable development, testing, fielding, and continuous update of artificial intelligence-powered applications at speed and scale from headquarters to the tactical edge.

(b) CONTENTS OF PLAN.—At a minimum, the plan required by subsection (a) shall include—

(1) an open architecture and an evolving reference design and guidance for needed technical investments in the proposed ecosystem that address issues including common interfaces, authentication, applications, platforms, software, hardware, and data infrastructure; and

(2) a governance structure, together with associated policies and guidance, to drive the implementation of the reference throughout the Department on a federated basis.

Recommendation: Train and Educate Warfighters.

Component 1: Integrate Digital Skill Sets and Computational Thinking into Military Junior Leader Education.

Require the military services to integrate digital skills and computational thinking into pre-commissioning and entry-level training.

SEC. ____.—INTEGRATING DIGITAL SKILL SETS AND COMPUTATIONAL THINKING INTO MILITARY JUNIOR LEADER EDUCATION.—

Not later than 270 days after the date of the enactment of this Act, the Chief of Staff of the Army, the Chief of Naval Operations, the Chief of Staff of the Air Force, and the Commandant of the Marine Corps (collectively, the Service Chiefs) shall expand the curriculum for military junior leader education to incorporate appropriate training material related to problem definition and curation, a conceptual understanding of the artificial intelligence lifecycle, data collection and management, probabilistic reasoning and data visualization, and data-informed decision-making. Whenever possible, the new training and education should include the use of existing artificial intelligence-enabled systems and tools.

Component 2: Integrate Emerging and Disruptive Technologies into Service-level Professional Military Education.

Require the military services to integrate emerging and disruptive technologies into service-level Professional Military Education.

SEC. ____.—INTEGRATION OF MATERIAL ON EMERGING TECHNOLOGIES INTO PROFESSIONAL MILITARY EDUCATION.—Not later than one year after the date of the enactment of this Act, the Secretary of Defense, in consultation with the Joint Chiefs of Staff, shall ensure that the curriculum for professional military education is revised in each of the military services to incorporate periodic courses on militarily significant emerging technologies that increasingly build the knowledge base, vocabulary, and skills necessary to intelligently analyze and utilize emerging technologies in the tactical, operational, and strategic levels of warfighting and warfighting support.

SEC. ____.—SHORT COURSE ON EMERGING TECHNOLOGIES FOR SENIOR CIVILIAN AND MILITARY LEADERS.—

(a) IN GENERAL.—Not later than one year after the date of the enactment of this Act, the Secretary of Defense shall establish a short course on emerging technologies for general and flag officers and senior executive-level civilian leaders. The short course shall be taught on an iterative, two-year cycle and shall address the most recent, most relevant technologies and how these technologies may be applied to military and business outcomes in the Department of Defense.

(b) THROUGHPUT OBJECTIVES.—In assessing participation in the short course authorized by subsection (a), the Secretary of Defense shall ensure that:

(1) In the first year that the course is offered, no fewer than twenty percent of general flag officers and senior executive-level civilian leaders are certified as having passed the short course required by subsection (a); and

(2) In each subsequent year, an additional ten percent of general flag officers and senior executive-level civilian leaders are certified as having passed such course, until such time as eighty percent of such officers and leaders are so certified.

Component 3: Create Emerging and Disruptive Technology Coded Billets in the Department of Defense.

Require the Department of Defense to create emerging and disruptive technology critical billets that must be filled by emerging technology certified leaders.

SEC. ____.—EMERGING TECHNOLOGY-CODED BILLETS WITHIN THE DEPARTMENT OF DEFENSE.—

(a) IN GENERAL.—Not later than one year after the date of the enactment of this Act, the Secretary of Defense shall ensure that the military services—

(1) code appropriate billets to be filled by emerging technology-qualified officers; and

(2) develop a process for officers to become emerging technology-qualified.

(b) APPROPRIATE POSITIONS.—Emerging technology-coded positions may include, as appropriate—

(1) positions responsible for assisting with acquisition of emerging technologies;

(2) positions responsible for helping integrate technology into field units;

(3) positions responsible for developing organizational and operational concepts;

(4) positions responsible for developing training and education plans; and

(5) leadership positions at the operational and tactical levels within the military services.

(c) QUALIFICATION PROCESS.—The process for qualifying officers for emerging technology-coded billets shall be modeled on a streamlined version of the joint qualification process and may include credit for serving in emerging technology focused fellowships, emerging technology focused talent exchanges, emerging technology focused positions within government, and educational courses focused on emerging technologies.

Recommendation: Accelerate Adoption of Existing Digital Technologies.

Component 3: Expand Use of Specialized Acquisition Pathways and Contracting Approaches.

Authorize the use of a rapid contracting mechanism for the software acquisition pathway.

SEC. ____.—RAPID CONTRACTING MECHANISM FOR SOFTWARE ACQUISITION.—

(a) IN GENERAL.—Not later than 270 days after the date of the enactment of this Act, the Secretary of Defense shall establish an agile contracting mechanism to support the software acquisition pathway developed pursuant to section 800 of the National Defense Authorization Act for Fiscal Year 2020 and embedded in Department of Defense Directives 5000.02 and 5000.87.

(b) CHARACTERISTICS.—The agile contracting mechanism established pursuant to subsection (a) shall authorize processes pursuant to which—

(1) a contract is awarded on the basis of statements of qualifications and past performance data submitted by contractors, supplemented by discussions with two or more contractors determined to be the most highly-qualified, without regard to price;

(2) the contract identifies the contractor team to be engaged for the work, and substitutions shall not be made during the base contract period without the advance written consent of the contracting officer;

(3) the contractor reviews existing software in consultation with the user community and utilizes user feedback to define and prioritize software requirements, and to design and implement new software and software upgrades, as appropriate;

(4) an independent, non-advocate cost estimate is developed in parallel with engineering of the software, leveraging agile cost estimation best practices rather than counting source lines of code; and

(5) value-based performance metrics are established and can be automatically generated by users to address issues such as deployment rate and speed of delivery, response rate such as the speed of recovery from outages and cybersecurity vulnerabilities, and assessment and estimation of the size and complexity of software development effort.

Component 4: Modernize the Budget and Oversight Processes for Digital Technologies.
Update title 10, Section 181 to designate USD(R&E) Co-Chair and Chief Science Advisor to the JROC.

SEC. ____.—ENHANCED ROLE OF UNDER SECRETARY OF DEFENSE FOR RESEARCH AND ENGINEERING ON THE JOINT REQUIREMENTS OVERSIGHT COUNCIL.—Section 181 of title 10, United States Code, is amended—

(1) in subsection (b), by—

(A) inserting “the Secretary of Defense and” before “the Chairman of the Joint Chiefs of Staff”;

(B) redesignating paragraphs (2) through (6) as paragraphs (3) through (7);

(C) inserting a new paragraph (2), as follows:

“(2) leveraging awareness of global technology trends, threats, and adversary capabilities to address gaps in joint military capabilities and validate technical feasibility of requirements developed by the military services;” and

(D) in redesignated paragraphs (4)(B) and (5) by inserting “the Secretary of Defense and” before “the Chairman of the Joint Chiefs of Staff”;

(2) in subsection (c), by—

(A) striking “Chairman of the Joint Chiefs of Staff for making recommendations about” in paragraph (1)(A) and inserting “Council for”;

(B) redesignating subparagraphs (B) through (E) of paragraph (1) as subparagraphs (C) through (F);

(C) adding a new paragraph (1)(B), as follows:

“(B) The Under Secretary of Defense for Research and Engineering, who is the co-Chair of the Council and is the Chief Science Advisor to the Council.”;

(D) by striking in paragraph (2) “(B), (C), (D), and (E)” and inserting “(C), (D), (E), and (F)”;

(E) by amending paragraph (3) to read as follows:

“(3) In making any recommendation to the Secretary and the Chairman of the Joint Chiefs of Staff pursuant to this section, the Co-Chairs of the Council shall provide any dissenting view of members of the Council with respect to such recommendation.”; and

(3) in subsection (d), by—

(A) striking subparagraph (1)(D); and

(B) redesignating subparagraphs (E) through (H) of paragraph (1) as paragraphs (D) through (G).

Direct the Secretary of Defense to establish the dedicated AI fund.

SEC. ____.—ARTIFICIAL INTELLIGENCE DEVELOPMENT AND PROTOTYPING FUND.—

(a) IN GENERAL.—The Secretary of Defense shall establish a fund to be known as the “Artificial Intelligence Development and Prototyping Fund” to support operational prototyping and speed the transition of artificial intelligence-enabled applications into both service-specific and joint mission capabilities with priority on joint mission capabilities for Combatant Commanders. The Fund shall be managed by the Under Secretary of Defense for Research and Engineering, in consultation with the Joint Artificial Intelligence Center, the Joint Staff, and the military services.

(b) TRANSFER AUTHORITY.—Amounts available in the Fund may be transferred to a military department for the purpose of carrying out a development or prototyping program selected by the Under Secretary of Defense for Research and Engineering for the purposes described in paragraph (1). Any amount so transferred shall be credited to the account to which it is transferred. The transfer authority provided in this subsection is in addition to any other transfer authority available to the Department of Defense.

(c) CONGRESSIONAL NOTICE.—The Under Secretary of Defense for Research and Development shall notify the congressional defense committees of all transfers under paragraph (2). Each notification shall specify the amount transferred, the purpose of the transfer, and the total projected cost and estimated cost to complete the acquisition program to which the funds were transferred.

CHAPTER 3: AI AND WARFARE

Blueprint for Action

Recommendation: Establish AI-readiness performance goals.

Require the Secretary of Defense to establish performance objectives and accompanying metrics for AI and digital readiness and provide an update to Congress no later than 120 days after approving these goals.

SEC. ____.—ARTIFICIAL INTELLIGENCE READINESS GOALS.—

(a) IN GENERAL.—Not later than one year after the date of the enactment of this Act, the Secretary of Defense shall review the potential applications of artificial intelligence and digital technology to Department of Defense platforms, processes and operations, and establish performance objectives and accompanying metrics for the incorporation of artificial intelligence and digital readiness into such platforms, processes and operations.

(b) SKILLS GAPS.—As a part of the review required by subsection (a), the Secretary shall direct the military departments and defense components to—

(1) conduct a comprehensive review of skill gaps in the fields of software development, software engineering, knowledge management, data science, and artificial intelligence;

(2) assess the number and qualifications of civilian personnel needed for both management and specialist tracks in such fields;

(3) assess the number of military personnel (officer and enlisted) needed for both management and specialist tracks in such fields; and

(4) establish recruiting, training, and talent management goals to achieve and maintain staffing levels needed to fill identified gaps and meet the Department's needs for skilled personnel.

(c) REPORT TO CONGRESS.—Not later than 120 days after the completion of the review required by subsection (a), the Secretary shall report to Congress on the findings of the review and any action taken or proposed to be taken by the Secretary to address such findings.

Recommendation: Promote AI interoperability and the adoption of critical emerging technologies among allies and partners.

Component 6: Modify authorities and processes in order to improve DoD's ability to conduct international capability development.

SEC. ____.—ENHANCED AUTHORITY TO ENTER INTO COOPERATIVE RESEARCH AND DEVELOPMENT AGREEMENTS WITH INTERNATIONAL PARTNERS.—

(a) AUTHORITY OF SECRETARY OF DEFENSE.—Section 2350a of title 10, United States Code, is amended—

(1) In subsection (a), by—

(A) Adding a new subparagraph (F) at the end of paragraph (2), as follows:

“(F) Any business, academic or research institution, or other non-governmental entity organized pursuant to the laws of a country referred to in subparagraphs (C), (D) and (E), subject to the consent of the country involved.”;

(B) Amending paragraph (3) by striking “a country referred to in subparagraph (E) of paragraph (2),” and inserting “a country referred to in subparagraph (E) of paragraph (2) or a non-governmental entity referred to in subparagraph (F) of such paragraph,”; and

(C) Adding a new paragraph (4), as follows:

“(4) The Secretary may delegate the authority to enter memoranda of understanding pursuant to this section to the secretary of a military department, the Director of the Joint Artificial Intelligence Center, and the Director of the Defense Advanced Research Projects Agency, subject to such terms and conditions as may be necessary to ensure that any agreements entered are consistent with the foreign policy and defense policy of the United States.”; and

(2) In paragraph (1) of subsection (b), by striking “will improve, through the application of emerging technology,” and inserting “is likely to improve, through the application or enhancement of emerging technology,”;

(3) In subsection (c), by adding at the end the following new sentence:
 “If a foreign partner is expected to contribute significantly to the development of a new or novel capability, full consideration shall be given to non-monetary contributions, including the value of research and development capabilities and the strategic partnerships.”

(b) AUTHORITY OF THE PRESIDENT.—Section 2767 of title 22, United States Code, is amended—

(1) in subsection (c), by adding at the end the following new sentence:
 “If a foreign partner is expected to contribute significantly to the development of a new or novel capability, full consideration shall be given to non-monetary contributions, including the value of research and development capabilities and the strategic partnerships.”

(2) in subsection (f), by inserting before the semicolon in subparagraph (4) the following: “(and a description of any non-monetary contributions made by such participants)”; and

(3) in subsection (j), by—

(A) amending the title to read as follows: “Cooperative project agreements with friendly foreign countries not members of NATO and with non-governmental organizations in NATO and friendly non-NATO countries”; and

(B) amending paragraph (2) to read as follows:

“(2) The President may enter into a cooperative project agreement with any business, academic or research institution, or other non-governmental entity organized pursuant to the laws of NATO member or a friendly foreign country that is not a member of NATO, subject to the consent of the country involved.”

CHAPTER 5: AI AND THE FUTURE OF NATIONAL INTELLIGENCE

Blueprint for Action

Recommendation: Empower the IC’s science and technology leadership.

Designate the Director of S&T within ODNI as the IC CTO and grant that position additional authorities for establishing policies on, and supervising, IC research and engineering, technology development, technology transition, appropriate prototyping activities, experimentation, and developmental testing activities.

Grant the Director of National Intelligence sufficient budgetary authorities to enforce technical standards across the IC, including the ability to fence or otherwise withhold funding for programs that are not compliant with established common standards and policies.

SEC. ____.—CHIEF TECHNOLOGY OFFICER FOR THE INTELLIGENCE COMMUNITY.—

Section 3030 of title 50, United States Code, is amended—

(1) in subsection (a), by striking “who shall be appointed by the Director of National Intelligence” and inserting “who shall be appointed by the Director of National Intelligence and shall serve as the Chief Technology Officer for the Intelligence Community.”; and

(2) in subsection (c), by—

(A) redesignating paragraphs (2) through (5) as paragraphs (4) through (7); and

(B) inserting new paragraphs (2) and (3), as follows:

“(2) establish policies for the intelligence community on research and engineering, technology development, technology transition, prototyping activities, experimentation, and developmental testing, and oversee the implementation of such policies;

“(3) establish common technical standards and policies necessary to rapidly scale artificial intelligence-enabled applications across the intelligence community;”.

Suggested Report Language: The Chief Technology Officer for the Intelligence Community shall collect information on each Intelligence Community element’s compliance with applicable standards and policies for artificial intelligence research and development, and shall provide such information to the Director of National Intelligence. The Intelligence Committees encourage the Director of National Intelligence to closely review the compliance information and place a temporary hold on an Intelligence Community element that fails to execute artificial intelligence research and development funds in accordance with the applicable standards and policies.

Establish a fund that would allow the DNI to identify and invest in AI applications with outsized potential that may not have an identified source of agency or program funding as they near the end of their S&T life cycle.

SEC. ____.—ARTIFICIAL INTELLIGENCE CRITICAL APPLICATIONS FUND FOR THE INTELLIGENCE COMMUNITY.—

(a) IN GENERAL.—The Director of National Intelligence shall establish a fund

to be known as the “Artificial Intelligence Critical Applications Fund” to support agile development and fielding of artificial intelligence-enabled applications with exceptional potential for the intelligence community. The Fund shall be managed by the Director of Science and Technology, in consultation with the National Intelligence Science and Technology Committee established pursuant to section 3030 of title 50, United States Code.

(b) TRANSFER AUTHORITY.—Amounts available in the Fund may be transferred to any element of the intelligence community for the purpose of carrying out a development or fielding program selected by the Director of Science and Technology for the purposes described in subsection (a). Any amount so transferred shall be credited to the account to which it is transferred. The transfer authority provided in this subsection is in addition to any other transfer authority available to the Director of National Intelligence and the intelligence community.

(c) CONGRESSIONAL NOTICE.—The Director of National Intelligence shall notify the congressional intelligence committees and the congressional appropriations committees of all transfers under paragraph (2). Each notification shall specify the amount transferred, the purpose of the transfer, and the total projected cost and estimated cost to complete the acquisition program to which the funds were transferred.

Establish a 10-year, \$1,000,000,000 Program of Record to provide long-term, predictable funding for technologies identified in the technology annex to the National Intelligence Strategy.

SEC. ____.—ARTIFICIAL INTELLIGENCE TECHNOLOGY ROADMAP AND FUNDING PLAN FOR THE INTELLIGENCE COMMUNITY.—

(a) IN GENERAL.—Not later than 180 days after the date of the enactment of this Act, the Director of National Intelligence, in consultation with the Secretary of Defense, shall develop a technology annex to the National Intelligence Strategy and a ten-year plan to provide long-term, predictable funding of up to one billion dollars to implement the steps identified in such annex.

(b) CONTENTS OF TECHNOLOGY ANNEX.—The technology annex required by subsection (a) shall provide a technology roadmap for the adoption of artificial intelligence-enabled applications to solve operational intelligence requirements, including:

(1) A description of challenges faced in the intelligence community’s efforts to analyze the global environment and monitor technological advancements, adversarial capability development, and emerging threats;

(2) Identification of technical capabilities, including artificial intelligence capabilities, needed to enable steps to address each challenge;

(3) A prioritized, time-phased plan for developing or acquiring such technical capabilities, that takes into account research and development timelines, a strategy for public private partnerships, and a strategy for connecting researchers to end users for early prototyping, experimentation, and iteration;

(4) Any additional or revised acquisition policies and workforce training requirements that may be needed to enable intelligence community personnel to identify, procure, integrate, and operate the technologies identified in the annex;

(5) Identification of infrastructure requirements for developing and deploying technical capabilities, including:

(A) data, compute, storage, and network needs;

(B) a resourced and prioritized plan for establishing such infrastructure; and

(C) an analysis of the testing, evaluation, verification, and validation requirements to support prototyping and experimentation and a resourced plan to implement them, including standards, testbeds, and red-teams for testing artificial intelligence systems against digital “denial & deception” attacks.

(6) Consideration of human factor elements associated with priority technical capabilities, including innovative human-centric approaches to user interface, human-machine teaming, and workflow integration;

(7) Consideration of interoperability with allies and partners, including areas for sharing of data, tools, and intelligence products; and

(8) Flexibility to adapt and iterate annex implementation at the speed of technological advancement.

Recommendation: Improve coordination between the IC and DoD.

Revise the National Defense Authorization Act for Fiscal Year 2021 (FY 2021 NDAA) provision authorizing a Steering Committee on Emerging Technology by designating it to be tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.

See Chapter 2 recommendation “Drive Change through Top-Down Leadership” for proposed legislative text.

Recommendation: Aggressively pursue security clearance reform for clearances at the Top Secret level and above, and enforce security clearance reciprocity among members of the IC.

Congress should require the DNI to develop an implementation plan for security clearance reform for clearances at the Top Secret and above level including detailed timelines and metrics.

Congress should require the DNI and the directors of the major intelligence services to regularly report on progress to the oversight committees.

SEC. ____.—IMPLEMENTATION PLAN FOR SECURITY CLEARANCE REFORM.—

(a) PLAN REQUIRED.—Not later than 180 days after the date of the enactment of this Act, the Director of National Intelligence shall develop an implementation plan for security clearance reform for clearances at the Top Secret level and above. The implementation plan shall include, at a minimum:

(1) detailed implementation metrics and timelines;

(2) steps to be taken to collaborate with the private sector and academia to develop data-informed behavioral approaches to understanding risk factors and security clearance adjudication; and

(3) steps to be taken to reform identity management and ensure seamless security clearance reciprocity across the intelligence community (including any enforcement mechanisms that may be needed to ensure such reciprocity).

(b) REPORTS REQUIRED.—Not later than 270 days after the date of the enactment of this Act and annually for five years thereafter, the Director of National Intelligence shall report to the congressional intelligence committees on the implementation of the plan required by subsection (a) and the progress that has been made toward security clearance reform.

CHAPTER 6: TECHNICAL TALENT IN GOVERNMENT

Blueprint for Action

Recommendation: Create a National Reserve Digital Corps.

NATIONAL RESERVE DIGITAL CORPS ACT OF 2021

SECTION. 1.—SHORT TITLE.—This Act may be cited as the “National Reserve Digital Corps Act of 2021”.

SEC. 2.—ESTABLISHMENT OF NATIONAL RESERVE DIGITAL CORPS.—

(a) IN GENERAL.—Subpart I of part III of title 5, United States Code, is amended by inserting after chapter 102 the following new chapter:

CHAPTER 103—NATIONAL RESERVE DIGITAL CORPS

SEC. 10301. Establishment.

SEC. 10302. Definitions.

SEC. 10303. Organization.

SEC. 10304. Work on Behalf of Federal Agencies.

SEC. 10305. Digital Corps Scholarship Program.

SEC. 10306. Duration of Pilot Program.

SEC. 10307. Authorization of Appropriation.

SEC. 10301. ESTABLISHMENT.—For the purposes of attracting, recruiting, and training a corps of world-class digital talent to serve the national interest and enable the Federal Government to become a digitally proficient enterprise, there is established within the Office of Management and Budget a pilot program for a civilian National Reserve Digital Corps, whose members shall serve as special government employees, working not fewer than 30 days per year as short-term advisors, instructors, or developers in the Federal Government.

SEC. 10302. DEFINITIONS.—

(a) DIRECTOR.—The term “Director” means the Director of the Office of Management and Budget.

(b) NODE.—The term “node” means a group of persons or team organized under the direction of a node leader to provide digital service to one or more Federal agencies pursuant to an agreement between the Office of Management Budget and each other Federal agency.

(c) NODE LEADER.—The term “node leader” means a full time government employee, as defined by section 2105 of title 5, United States Code, selected under this Act to lead one or more nodes, who reports to the Director or the Director’s designee.

(d) NODE MEMBER.—The term “node member” means a special government employee, as defined by section 202 of title 18, United States Code, selected under this Act to work at least 38 days per fiscal year and report to a node leader in furtherance of the mission of a specified node.

SEC. 10303. ORGANIZATION.—

(a) NODES AND NODE LEADERS.—The National Reserve Digital Corps shall be organized into nodes, each of which shall be under the supervision of a node leader .

(b) ADMINISTRATIVE SUPPORT.—The National Reserve Digital Corps shall receive funding and administrative support from the Office of Management and Budget,

which shall be responsible for selecting node leaders, establishing standards, ensuring that nodes meet government client requirements, maintaining security clearances, establishing access to an agile development environment and tools, and facilitating appropriate technical exchange meetings.

(c) HIRING AUTHORITY.—

(1) Direct Hiring Authority of Node Members.—The Director of the Office of Management and Budget, on the recommendation of a node leader, may appoint, without regard to the provisions of subchapter I of chapter 33 (other than sections 3303 and 3328 of such chapter), a qualified candidate to a position in the competitive service in the Office of Management and Budget to serve as a node member. This provision shall not preclude the Director from hiring additional employees, including full time government employees, as defined by section 2105 of title 5, United States Code.

(2) Term and Temporary Appointments of Node Members.—The Director of the Office of Management and Budget, on the recommendation of a node leader, may make a noncompetitive temporary appointment or term appointment for a period of not more than 18 months, of a qualified candidate to serve as a node member in a position in the competitive service for which a critical hiring need exists, as determined under section 3304 of title 5, United States Code, without regard to sections 3327 and 3330 of such title.

SEC. 10304. WORK ON BEHALF OF FEDERAL AGENCIES.—

(a) PURPOSE.—Each node shall undertake projects to assist Federal agencies by providing digital education and training, performing data triage and providing acquisition assistance, helping guide digital projects and frame technical solutions, helping build bridges between public needs and private sector capabilities, and related tasks.

(b) AUTHORITIES.—Projects may be undertaken—

(1) on behalf of a Federal agency—

(A) by direct agreement between the Office of Management and Budget and the Federal agency; or

(B) at the direction of the Office of Management and Budget at the request of the Federal agency; or

(2) to address a digital service need encompassing more than one Federal agency—

(A) at the direction of the Office of Management and Budget; or

(B) on the initiative of a node leader.

SEC. 10305. DIGITAL CORPS SCHOLARSHIP PROGRAM.—

(a) IN GENERAL.—The Director shall establish a National Reserve Digital Corps scholarship program to provide full scholarships to competitively selected students who commit to study specific disciplines related to national security digital technology .

(b) SERVICE OBLIGATION.—Each student, prior to commencing the Digital Corps Scholarship Program, shall sign an agreement with respect to the student's commitment to the United States. The agreement shall provide that the student agree to the following:

(1) a commitment to serve as an intern in a Federal agency for at least six weeks during each of the summers before their junior and senior years; and

(2) a commitment to serve in the National Reserve Digital Corps for six years after graduation.

(c) PROGRAM ELEMENTS.—In establishing the program, the Director shall determine the following—

(1) Eligibility standards for program participation;

(2) Criteria for establishing the dollar amount of a scholarship, including tuition, room and board;

(3) Repayment requirements for students who fail to complete their service obligation;

(4) An approach to ensuring that qualified graduates of the program are promptly hired and assigned to node leaders; and

(5) Resources required for the implementation of the program.

(d) CONTINUING EDUCATION.—The Director shall establish a training and continuing education program to fund educational opportunities for members of the National Digital Reserve Corps, including conferences, seminars, degree and certificate granting programs, and other training opportunities that are expected to increase the digital competencies of the participants.

(e) IMPLEMENTATION.—

(1) Not later than six months after the date of the enactment of this Act, the Director shall establish the administrative support function and issue guidance for the National Reserve Digital Corps, which shall include the identification of points of contact for node leaders at Federal agencies.

(2) Not later than one year after the date of the enactment of this Act, the Director shall appoint not fewer than five node leaders under the National Reserve Digital Corps program and authorize the node leaders to begin recruiting reservists and undertaking projects for Federal agencies.

(3) Beginning two years after the date of the enactment of this Act, the Director shall report annually to Congress on the progress of the National Reserve Digital Corps. The Director's report shall address, at a minimum, the following measures of success:

(A) The number of technologists who participate in the National Reserve Digital Corps annually;

(B) Identification of the Federal agencies that submitted work requests, the nature of the work requests, which work requests were assigned a node, and which work requests were completed or remain in progress;

(C) Evaluations of results of National Reserve Digital Corps projects by Federal agencies; and

(D) Evaluations of results of National Reserve Digital Corps projects by reservists.

SEC. 10306. DURATION OF PILOT PROGRAM.—The pilot program under this Act shall terminate no earlier than six years after its commencement.

SEC. 10307. AUTHORIZATION OF APPROPRIATION.—There is authorized to be appropriated \$16,000,000 to remain available until fiscal year 2023 the initial administrative cost, including for the salaries and expenses scholarship and education benefits, for the National Digital Reserve Corps.

Recommendation: Create Digital Talent Recruiting Offices Aligned with Digital Corps.

SEC. ____.—DIGITAL TALENT RECRUITING OFFICES.—

(a) DIGITAL TALENT RECRUITING FOR THE DEPARTMENT OF DEFENSE.—

(1) Not later than 270 days after the date of the enactment of this Act, the Secretary of Defense shall designate a chief digital recruiting officer within the office of the Under Secretary of Defense for Personnel and Readiness to oversee a digital recruiting office to carry out the responsibilities set forth in paragraph (2).

(2) The chief digital recruiting officer shall be responsible for—

(A) identifying Department of Defense needs for specific types of digital talent;

(B) recruiting technologists, in partnership with the military services and defense components, including by attending conferences and career fairs, and actively recruiting on university campuses and from the private sector;

(C) integrating Federal scholarship for service programs into civilian recruiting;

(D) offering recruitment and referral bonuses; and

(E) partnering with human resource teams in the military services and defense components to use direct-hire authorities to accelerate hiring.

(3) The Secretary of Defense shall ensure that the chief digital recruiting officer is provided with personnel and resources sufficient to maintain an office and to carry out the duties set forth in paragraph (2).

(b) DIGITAL TALENT RECRUITING FOR THE INTELLIGENCE COMMUNITY.—

(1) Not later than 270 days after the date of the enactment of this Act, the Director of National Intelligence shall designate a chief digital recruiting officer to oversee a digital recruiting office to carry out the responsibilities set forth in paragraph (2).

(2) The chief digital recruiting officer shall be responsible for—

(A) identifying intelligence community needs for specific types of digital talent;

(B) recruiting technologists, in partnership with components of the intelligence community, by attending conferences and career fairs, and actively recruiting on college campuses;

(C) integrating Federal scholarship for service programs into intelligence community recruiting;

(D) offering recruitment and referral bonuses; and

(E) partnering with human resource teams in the components of the intelligence community to use direct-hire authorities to accelerate hiring.

(3) The Director of National Intelligence shall ensure that the chief digital recruiting officer is provided with personnel and resources sufficient to maintain an office and to carry out the duties set forth in paragraph (2).

(c) DIGITAL TALENT RECRUITING FOR THE DEPARTMENT OF HOMELAND SECURITY.—

(1) Not later than 270 days after the date of the enactment of this Act, the Secretary of Homeland Security shall designate a chief digital recruiting officer to oversee a digital recruiting office to carry out the responsibilities set forth in paragraph (2).

(2) The chief digital recruiting officer shall be responsible for—

(A) identifying Department of Homeland Security needs for specific types of digital talent;

(B) recruiting technologists, in partnership with components of the Department of Homeland Security, by attending conferences and career fairs, and actively recruiting on college campuses;

(C) integrating Federal scholarship for service programs into civilian recruiting;

(D) offering recruitment and referral bonuses; and

(E) partnering with human resource teams in the components of the Department of Homeland Security to use direct-hire authorities to accelerate hiring.

(3) The Secretary of Homeland Security shall ensure that the chief digital recruiting officer is provided with personnel and resources sufficient to maintain an office and to carry out the duties set forth in paragraph (2).

(d) DIGITAL TALENT RECRUITING FOR THE DEPARTMENT OF ENERGY.—

(1) Not later than 270 days after the date of the enactment of this Act, the Secretary of Energy shall designate a chief digital recruiting officer to oversee a digital recruiting office to carry out the responsibilities set forth in paragraph (2).

(2) The chief digital recruiting officer shall be responsible for—

(A) identifying Department of Energy needs for specific types of digital talent;

(B) recruiting technologists, in partnership with Department of Energy programs, by attending conferences and career fairs, and actively recruiting on college campuses;

(C) integrating Federal scholarship for service programs into civilian recruiting;

(D) offering recruitment and referral bonuses; and

(E) partnering with human resource teams in Department of Energy programs to use direct-hire authorities to accelerate hiring.

(3) The Secretary of Energy shall ensure that the chief digital recruiting officer is provided with personnel and resources sufficient to maintain an office and to carry out the duties set forth in paragraph (2).

Recommendation: Grant exemption from OPM General Schedule Qualification Policies for Specific Billets and Position Descriptions.

SEC. ____.—WAIVER OF QUALIFICATION STANDARDS FOR GENERAL SCHEDULE POSITIONS IN ARTIFICIAL INTELLIGENCE.—

(a) DEPARTMENT OF DEFENSE.—Two-star and above commands and their civilian equivalents are authorized to waive any General Schedule qualification standard established by the Office of Personnel Management in the case of any applicant for a position in artificial intelligence who is determined by a hiring manager, in consultation with subject matter experts, to be the best qualified candidate for the position.

(b) OTHER NATIONAL SECURITY AGENCIES.—The Director of the Office of Personnel Management shall establish a process by which the the Attorney General, the Secretary of Homeland Security, the Secretary of State, the Secretary of Commerce, the Director of National Intelligence, and the head of any element of the Intelligence Community may request an exception to any General Schedule qualification standard in any case in

which the agency head determines that national security needs would best be met by hiring managers making an independent judgment about qualifications and pay grades for a position in artificial intelligence with the advice of subject matter experts. The process shall provide for requests to be made for individual billets, for position descriptions, or for categories of individual billets or position descriptions at the discretion of the agency head.

Recommendation: Expand the CyberCorps: Scholarship for Service.

SEC. ____.—AMENDMENT TO THE FEDERAL CYBER SCHOLARSHIP-FOR- SERVICE PROGRAM.—

(a) AMENDMENTS TO TITLE 15, UNITED STATES CODE.—Section 7442 of title 15, United States Code, is amended—

(1) By amending the title to read: “Federal Cyber and Artificial Intelligence Scholarship-for-Service Program”;

(2) in subsection (a), by striking “industrial control system” and all that follows and inserting in lieu thereof “digital engineers, artificial intelligence practitioners, data engineers, data analysts, data scientists, industrial control system security professionals, security managers, and cybersecurity course instructors to meet the needs of the cybersecurity and artificial intelligence missions for Federal, State, local, tribal, and territorial governments.”;

(3) in subsection (b), by—

(A) striking “and” at the end of paragraph (3);

(B) striking the period at the end of paragraph (4) and inserting in lieu thereof “; and”; and

(C) adding a new paragraph (5), as follows:

“(5) provide an opportunity for scholarship recipients to initiate the security clearance process at least one year before their planned graduation date.”; and

(4) in subsection (c), by striking “3 years” and inserting “4 years”.

(b) SAVINGS PROVISION.—Nothing in this section, or an amendment made by this section, shall affect any agreement, scholarship, loan, or repayment under section 302 of the Cybersecurity Enhancement Act of 2014 (15 U.S.C. 7442), in effect on the day before the date of the enactment of this section.

Recommendation: Create a United States Digital Service Academy.

UNITED STATES DIGITAL SERVICE ACADEMY ACT OF 2021

SECTION. 1.—SHORT TITLE.—This Act may be cited as the “United States Digital Service Academy Act of 2021”.

SEC. 2.—ESTABLISHMENT OF ACADEMY.—

(a) ESTABLISHMENT.—There is established as an independent entity within the Federal Government a United States Digital Service Academy (hereafter referred to as the “ACADEMY”), at a location to be determined, to serve as a federally-funded, accredited, degree-granting university for the instruction of selected individuals in digital technical fields and the preparation of selected individuals for civil service with the Federal Government.

(b) DIGITAL TECHNICAL FIELDS DEFINED.—The term “digital technical fields” includes artificial intelligence, software engineering, electrical science and engineering, computer science, molecular biology, computational biology, biological engineering, cybersecurity, data science, mathematics, physics, human-computer interaction, robotics, and design and any additional fields specified in regulations by the Board.

SEC. 3.—ORGANIZATION.—

(a) BOARD OF REGENTS.—The business of the Academy shall be conducted by a Board of Regents (hereafter referred to as the “Board”).

(1) COMPOSITION.—The Board shall consist of nine voting members and ex officio members, as set forth in this subsection.

(2) VOTING MEMBERS.—The President shall appoint, by and with the consent of the Senate, nine persons from civilian life who have demonstrated achievement in one or more digital technical fields, higher education administration, or Federal civilian service, to serve as voting members on the Board. Appointment of the first voting members shall be made not later than 180 days after enactment of this Act.

(3) EX OFFICIO MEMBERS.—Ex officio members shall include—

(A) The Secretary of State;

(B) The Secretary of Defense;

(C) The Attorney General;

(D) The Secretary of Commerce;

(E) The Secretary of Energy;

(F) The Secretary of Homeland Security;

(G) The Director of National Intelligence;

(H) The Director of the Office of Personnel Management; and

(I) such other Federal Government officials as determined by the President.

(2) TERM OF VOTING MEMBERS.—The term of office of each voting member of the Board shall be six years, except that initial terms shall be staggered at two year intervals and any member appointed to fill a vacancy occurring before the expiration of a term shall be appointed for the remainder of such term.

(3) PRESIDENT OF THE BOARD.—One of the members (other than an ex officio member) shall be designated by the President as Chairman and shall be the presiding officer of the Board.

(b) KEY POSITIONS.—There shall be at the Academy the following:

(1) A Superintendent;

(2) A Dean of the Academic Board, who is a permanent professor;

(3) A Director of Admissions; and

(4) A Director of Placement.

(c) SUPERINTENDENT.—The Board shall appoint a Superintendent of the Academy, who shall serve for a term of six years. The Superintendent, acting pursuant to the oversight and direction of the Board, shall be responsible for the day-to-day operations of the Academy and the welfare of the students and the staff of the Academy. The Board shall select the first Superintendent of the Academy no later than 60 days after the Board is established.

(d) ADVISORY BOARD.—The Board of Regents and the Superintendent shall be assisted by an Advisory Board, composed of commercial and academic leaders in digital technical fields and higher education. The Advisory Board shall adhere to the requirements of the Federal Advisory Committee Act, Pub.L. 92–463.

(e) INTERAGENCY WORKING GROUP.—

(1) ESTABLISHMENT.—The Office of Personnel Management shall establish and lead an interagency working group to annually assess and report to the Academy the need for civil servants at agencies in digital technical fields for the purposes of informing Academy student field of study and agency placement.

(2) RESPONSIBILITIES.—The interagency working group shall be responsible for—

(A) establishing a range of Academy graduates needed during the ensuing five-year period, by agency and digital technical field; and

(B) undertaking necessary steps to enable each agency identified to hire Academy graduates into full-time positions in the civil service.

(3) COMPOSITION.—The interagency working group shall consist of the following officials or their designees:

(A) The Secretary of State;

(B) The Secretary of Defense;

(C) The Attorney General;

(D) The Secretary of Commerce;

(E) The Secretary of Energy;

(F) The Secretary of Homeland Security;

(G) The Director of National Intelligence;

(H) The Director of the Office of Personnel Management; and

(I) such other Federal Government officials as determined by the Director of the Office of Personnel Management.

SEC. 4.—FACULTY.—

(a) NUMBER OF FACULTY.—The Superintendent of the Academy may employ as many professors, instructors, and lecturers at the Academy as the Superintendent considers necessary to achieve academic excellence.

(b) FACULTY COMPENSATION.—The Superintendent may prescribe the compensation of persons employed under this section. Compensation and benefits for faculty members of the Academy shall be sufficiently competitive to achieve academic excellence, as determined by the Superintendent.

(c) FACULTY EXPECTATIONS.—Faculty members shall—

(1) possess academic expertise and teaching prowess;

(2) exemplify high standards of conduct and performance;

(3) be expected to participate in the full spectrum of academy programs, including providing leadership for the curricular and extracurricular activities of students;

(4) comply with the standards of conduct and performance established by the Superintendent; and

(5) participate actively in the development of the students through the enforcement of standards of behavior and conduct, to be established in the Academy's rules and regulations.

(d) DEPARTMENT TITLES.—The Superintendent may prescribe the titles of each of the departments of instruction and the professors of the Academy.

SEC. 5.—STUDENT QUALIFICATIONS AND REQUIREMENTS FOR ADMISSION.—

(a) ADMISSIONS REQUIREMENTS.—A student wishing to be admitted to the Academy shall fulfill admission requirements to be determined by the Superintendent and approved by the Board of Regents.

(b) HONOR CODE.—A student wishing to be admitted to the Academy shall sign an Honor Code developed by the Superintendent of the Academy and approved by the Board of Regents. A violation of the honor code may constitute a basis for dismissal from the Academy.

SEC. 6.—APPOINTMENT OF STUDENTS.—

(a) NOMINATIONS PROCESS.—Prospective applicants to the Academy for seats described in paragraphs (1) and (2) of subsection (b) shall follow a nomination process established by the Director of Admissions of the Academy that is similar to the process used for admission to the military academies of the United States Armed Forces.

(b) APPOINTMENTS.—

(1) NOMINEES FOR CONGRESSIONAL SEATS.—Each member of the Senate or the House of Representatives may nominate candidates from the State that the member represents for each incoming first-year class of the Academy .

(2) EXECUTIVE BRANCH NOMINEES.—The President may nominate a maximum of 75 candidates to compete for the executive branch seats.

SEC. 7.—ACADEMIC FOCUS OF THE UNITED STATES DIGITAL SERVICE ACADEMY.—

(a) CURRICULUM.—Each Academy student shall follow a structured curriculum according to the program of study approved by the Board of Regents centered on digital technical fields and incorporating additional core curriculum coursework in history, government, English language arts including composition, and ethics.

(b) DEGREES CONFERRED UPON GRADUATION.—Under such conditions as the Board of Regents may prescribe, once the Academy is accredited, the Superintendent of the Academy may confer a baccalaureate of science or baccalaureate of arts degree upon a graduate of the Academy.

(c) MAJORS AND AREAS OF CONCENTRATION.—Under such conditions as the Board of Regents may prescribe, the Superintendent of the Academy may prescribe requirements for majors and concentrations and requirements for declaring a major or concentration during the course of study.

(d) ADDITIONAL DIGITAL SERVICE OF CIVIL SERVICE PROGRAMMING.—Under such conditions as the Board of Regents may prescribe, the Superintendent of the Academy may prescribe requirements for each Academy student to participate in non-curricular programming during Academy terms and during the summer, which may include internships, summer learning programs, and project-based learning activities.

SEC. 8.—CIVIL SERVICE REQUIREMENTS FOLLOWING GRADUATION.—

(a) CIVIL SERVICE AGREEMENT.—Each Academy student, prior to commencing the third year of coursework, shall sign an agreement with respect to the student's length of civil service to the United States. The agreement shall provide that the student agrees to the following:

(1) The student will complete the course of instruction at the Academy, culminating in graduation from the Academy.

(2) Unless the student pursues graduate education under subsection (f), upon graduation from the Academy, the student agrees to serve in the Federal civil service for not less than five years following graduation from the Academy .

(b) FAILURE TO GRADUATE.—

(1) IN GENERAL.—An Academy student who has completed a minimum of four semesters at the Academy but fails to fulfill the Academy's requirements for graduation shall be—

(A) dismissed from the Academy; and

(B) obligated to repay the Academy for the cost of the delinquent student's education in the amount described in paragraph (2).

(2) AMOUNT OF REPAYMENT.—A student who fails to graduate shall have financial responsibility for certain costs relating to each semester that the student was officially enrolled in the Academy as prescribed by the Superintendent.

(c) FAILURE TO ACCEPT OR COMPLETE ASSIGNED CIVIL SERVICE.—

(1) IN GENERAL.—A student who graduates from the Academy but fails to complete the full term of required civil service shall be obligated to repay the Academy for a portion of the cost of the graduate's education as determined by Academy as set forth in this subsection.

(2) AMOUNT OF REPAYMENT.—In the case of a delinquent graduate who fails to complete all years of public service required under subsection (a)(2) (including any additional years required for graduate education under subsection (f)), the delinquent graduate shall be financially responsible for the cost of the delinquent graduate's education (including the costs of any graduate education), except that the amount of financial responsibility under this paragraph shall be reduced by 20 percent for each year of civil service under subsection (a)(2) that the delinquent graduate did complete.

(d) EXCEPTIONS.—The Superintendent may provide for the partial or total waiver or suspension of any civil service or payment obligation by an individual under this section whenever compliance by the individual with the obligation is impossible or deemed to involve extreme hardship to the individual, or if enforcement of such obligation with respect to the individual would be unconscionable.

(e) STUDENT SALARIES AND BENEFITS.—The Academy shall not be responsible for the salaries and benefits of graduates of the Academy while the graduates are fulfilling the civilian service assignment under this section. All salaries and benefits shall be paid by the employer with whom the Academy graduate is placed.

(f) GRADUATE EDUCATIONS.—An Academy student and the Superintendent may modify the agreement under subsection (a) to provide that—

(1) the Academy shall—

(A) subsidize an Academy student's graduate education; and

(B) postpone the public service assignment required under subsection (a)(2).

(2) the student shall—

(A) accept a civil service assignment under subsection (g) upon the student's completion of the graduate program; and

(B) add two additional years to the student's civil service commitment required under the agreement described in subsection (a) for every year of subsidized graduate education.

SEC. 9.—IMPLEMENTATION PLAN.—

(a) Not later than 180 days after the enactment of this Act, the Superintendent, in consultation with the Advisory Board, shall develop a detailed plan to implement the Academy that complies with the requirements of this section. Upon approval by the Board of Regents, the Superintendent shall present the implementation plan to Congress.

(b) CONTENTS OF PLAN.—The implementation plan described in section (a) shall provide, a minimum, the following:

(1) Identification and securement of an appropriate site for initial Academy build-out with room for future expansion, to include a construction plan and temporary site plan, if necessary;

(2) Identification of gaps in the government's current and envisioned digital workforce by the interagency working group under the Office of Personnel Management as established by section (3)(e);

(3) Establishment of student qualifications and requirements for admission;

(4) Establishment of the student appointment and nomination process;

(5) Establishment of student honor and conduct code to include a plan for student noncompletion of requirements and obligations;

(6) Establishment of the student curriculum;

(7) Establishment of a mechanism for students to select fields of study and annually select agencies and career fields within the limits prescribed by

the interagency working group under the Office of Personnel Management as established by section (3)(e);

(8) Establishment of a mechanism for graduates to transition from the Academy to civil service employment by selected individual agencies;

(9) Determination of the initial Academy departments and faculty needs;

(10) Establishment of faculty and staff requirements and compensation;

(11) Determination of non-academic staff required;

(12) Recruitment and hiring of faculty, including tenure-track faculty, adjunct faculty, part-time faculty and visiting faculty, and other staff as needed;

(13) Identification of nonprofit and private sector partners;

(14) Procurement of outside funds and gifts from individuals and corporations for startup, administrative, maintenance, and infrastructure costs;

(15) Establishment of the process to meet statutory and regulatory requirements for establishing the Academy as an academic institution with degree-granting approval and for applying for degree program specific accreditation and ensuring that the Academy obtains, no later than two years after enactment of this Act, status as an accreditation candidate, as defined by a nationally recognized accrediting agency or association as determined by the Secretary of Education in accordance with section 1099b in title 10, United States Code, before commencing academic operations;

(16) A plan commencing the Academy with an initial class of 500 students three years after enactment of this Act;

(17) Procedures for incorporating accreditation assessments to facilitate ongoing improvements to the Academy; and,

(18) Procedures for assessing the size of the Academy and potential expansion of student enrollment.

SEC. 10.—ADMINISTRATIVE MATTERS.—

(a) FULLY-SUBSIDIZED EDUCATION.—Each Academy student’s tuition and room and board shall be fully subsidized provided that the student completes the requirements of the Academy and fulfills the civil service commitment as determined by the implementation plan in section 9.

(b) GIFT AUTHORITY.—The Board of Regents may accept, hold, administer, and spend any gift, devise, or bequest of real property, personal property, or money made on the condition that the gift, devise, or bequest be used for the benefit, or in connection with, the establishment, operation, or maintenance, of the Academy. The Board of Regents may accept a gift of services, which includes activities that benefit the education, morale, welfare, or recreation of students, faculty or staff, for the Academy.

(1) LIMITATIONS AND PROHIBITIONS.—

(A) IN GENERAL.—The Board of Regents may not accept a gift under this subsection if the acceptance of the gift would reflect unfavorably on the ability of any agency of the Federal Government to carry out any responsibility or duty in a fair and objective manner, or would compromise the integrity or appearance of integrity of any program of the Federal Government or any officer or employee of the Federal Government who is involved in any such program.

(B) FOREIGN GIFTS.—The Board of Regents may not accept a gift of services from a foreign government or international organization under this subsection. A gift of real property, personal property, or money from a foreign government or international organization may be accepted under this subsection only if the gift is not designated for a specific individual.

(C) APPLICABLE LAW.—No gift under this section may be accepted with attached conditions inconsistent with applicable law or regulation.

(D) MISSION.—No gift under this section may be accepted with attached conditions inconsistent with the mission of the Academy .

(E) NAMING RIGHTS.—The Board of Regents may issue regulations governing the circumstances under which gifts conditioned on naming rights may be accepted, appropriate naming conventions, and suitable display standards.

(2) TREATMENT OF GIFTS.—

(A) Gifts and bequests of money, and the proceeds of the sale of property, received under subsection shall be deposited in the Treasury in the account of the Academy as no year money and may be expended in connection with the activities of the Academy as determined by the Board of Regents.

(B) The Board of Regents may pay all necessary expenses in connection with the conveyance or transfer of a gift, devise, or bequest accepted under this section.

(C) For the purposes of Federal income, estate, and gift taxes, any property, money, or services accepted under this subsection shall be considered as a gift, devise, or bequest to or for the use of the United States.

(D) The Comptroller General shall make periodic audits of gifts, devises, and bequests accepted under this section at such intervals as the Comptroller General determines to be warranted. The Comptroller General shall submit to Congress a report on the results of each such audit.

SEC. 11.—INITIAL APPROPRIATION.—There are authorized to be appropriated \$40,000,000 to remain available until expended for the Academy's initial administrative cost and salaries and expenses.

Recommendation: Establish Career Fields for Government Civilians in Software Development, Software Engineering, Data Science, Knowledge Management, and Artificial Intelligence.

SEC. ____.—NEW OCCUPATIONAL SERIES FOR DIGITAL CAREER FIELDS.—Not later than 270 days after the date of the enactment of this Act, the Director of the Office of Personnel Management shall exercise its authority under section 5105 of title 5, United States Code, to establish one or more new occupational series and associated policies covering Federal Government positions in the fields of software development, software engineering, data science, and knowledge management.

SEC. ____.—NEW OCCUPATIONAL SERIES FOR ARTIFICIAL INTELLIGENCE.—Not later than 270 days after the date of the enactment of this Act, the Director of the Office of Personnel Management shall exercise its authority under section 5105 of title 5, United States Code, to establish a new occupational series and associated policies covering Federal Government positions in the field of artificial intelligence.

Recommendation: Establish Digital Career Fields for Military Personnel.

SEC. ____.—MILITARY CAREER FIELDS FOR SOFTWARE DEVELOPMENT, DATA SCIENCE, AND ARTIFICIAL INTELLIGENCE.—Section 230 of the National Defense Authorization Act for Fiscal Year 2020 is amended by adding the following new subsection: “(d) Not later than 270 days after the date of the enactment of this subsection, the Chief of Staff of the Army, the Chief of Naval Operations, the Chief of Staff of the Air Force, and

the Commandant of the Marine Corps (collectively, the Service Chiefs) shall each establish new military career fields for software development, data science, and artificial intelligence that are open to commissioned officers, enlisted personnel and, as appropriate, warrant officers. The Service Chiefs shall utilize the authority provided in sections 605 and 649a to 649k of title 10, United States Code, to ensure that military personnel in these career fields who choose to specialize and focus on technical skill sets rather than pursue leadership positions are not required to move outside their specialties or into management positions to continue to promote.

CHAPTER 8: UPHOLDING DEMOCRATIC VALUES: PRIVACY, CIVIL LIBERTIES, AND CIVIL RIGHTS IN USES OF AI FOR NATIONAL SECURITY

Blueprint for Action

Recommendation Set 1: Increase Public Transparency about AI Use through Improved Reporting.

For AI systems that involve U.S. persons, require AI Risk Assessment Reports and AI Impact Assessments to assess the privacy, civil liberties and civil rights implications for each new qualifying AI system or significant system refresh.

SEC. ___—PRIVACY, CIVIL RIGHTS AND CIVIL LIBERTIES RISK AND IMPACT ASSESSMENTS FOR ARTIFICIAL INTELLIGENCE SYSTEMS.—

(a) IN GENERAL.—The head of a covered agency shall conduct risk and impact assessments of the privacy, civil rights, and civil liberties risks and potential implications of any covered artificial intelligence system utilized by the covered agency and take appropriate steps to mitigate risks and adverse impact of any such system on the privacy, civil rights, and civil liberties of U.S. persons.

(b) DEFINITIONS.—For purposes of this section—

(1) COVERED ARTIFICIAL INTELLIGENCE SYSTEM.—A “covered artificial intelligence system” means a qualified artificial intelligence system or a significant artificial intelligence system refresh as determined by the task force established in section [XX] of this Act that is—

(A) designed to collect, process, maintain, or use information on U.S. persons;

(B) may inadvertently process, maintain, or use information on U.S. persons; or

(C) has a direct impact on U.S. persons.

(2) COVERED AGENCY.—A “covered agency” includes—

(A) the Department of Homeland Security;

(B) the Federal Bureau of Investigation; and

(C) each element of the Intelligence Community, as defined in section 3003(4) of title 50, United States Code.

(3) HEAD OF A COVERED AGENCY.—The “head of a covered agency” shall mean the Secretary of Homeland Security, the Director of the Federal Bureau of Investigation and, for the Intelligence Community, the Director of National Intelligence.

(c) REPORTS REQUIRED.—

(1) ARTIFICIAL INTELLIGENCE SYSTEM RISK ASSESSMENT.—Before acquiring or fielding a covered artificial intelligence system, each covered agency shall conduct an Artificial Intelligence System Risk Assessment (“Risk Assessment”). The Risk Assessment shall—

(A) assess the potential implications of the covered artificial intelligence system on freedom of expression, equal protection, privacy, and due process;

(B) account for the environment in which the covered artificial intelligence system will be deployed, including its interactions with other artificial intelligence tools, programs, and systems that collect personally identifiable information; and

(C) include steps to mitigate and track any risks identified in the assessment.

(2) ARTIFICIAL INTELLIGENCE SYSTEM IMPACT ASSESSMENT.—Each covered agency shall conduct an Artificial Intelligence System Impact Assessment (“Impact Assessment”), no less than once per year, to assess the degree to which a covered artificial intelligence system remains compliant with the constraints and metrics established in the Risk Assessment. The Impact Assessment shall be based on outcomes, impacts, and metrics collected during system use, and shall determine if the existing validation processes should be improved.

(d) NOTICE OF DISCONTINUATION.—Within one year of discontinuing use of any non-public or classified covered artificial intelligence system, a covered agency shall

consider providing notice to the public that the covered artificial intelligence system has been discontinued.

(e) REPORT TO CONGRESS.—The head of each covered agency shall, within 90 days of the date of this Act, submit to Congress a report identifying any additional resources, including staff, needed to carry out the requirements of this section.

This section should be cross-referenced with the recommendation to create a task force to assess the privacy and civil rights and civil liberties implications of AI and emerging technologies, as the definition of a “covered artificial intelligence system” relies on the work of the task force.

Recommendation Set 2: Develop & Test Systems per Goals of Privacy Preservation and Fairness.

Establish third-party testing center(s) to allow independent, third-party testing of national security-related AI systems that could impact U.S. persons.

Require the Department of Justice (DOJ), in consultation with the Privacy and Civil Liberties Oversight Board (PCLOB), to develop binding guidance for the use of third-party testing (e.g., thresholds for high-consequence systems or unprecedented factors) of AI systems.

SEC. ____.—THIRD PARTY TESTING OF ARTIFICIAL INTELLIGENCE SYSTEMS.—

(a) IN GENERAL.—Not later than one year after the date of enactment of this Act, the Director of the National Institute of Standards and Technology shall establish an accreditation program for Third Party Independent Artificial Intelligence Testing Laboratories, as set forth in this section, to conduct independent testing of artificial intelligence systems for covered agencies to assess potential privacy, civil rights, and civil liberties impacts of such systems on U.S. persons.

(b) ARTIFICIAL INTELLIGENCE SYSTEMS REQUIRING TESTING.—The Privacy and Civil Liberties Oversight Board and the Department of Justice shall, in consultation with Privacy and Civil Liberties officers of the covered agencies, propose criteria for when an artificial intelligence system warrants third-party testing for privacy, civil liberties, and civil rights implications for U.S. Persons. Covered agencies shall adopt this criteria, as described in subsection (e).

(c) COVERED AGENCIES.—For the purposes of this section, covered agencies are the elements of the Intelligence Community (as defined in section 3003(4) of title 50, United States Code, and coordinated by the Office of the Director of National Intelligence), the Department of Homeland Security, and the Federal Bureau of Investigation.

(d) ACCREDITATION OF THIRD PARTY ARTIFICIAL INTELLIGENCE TESTING LABORATORIES.—Accreditation of Third Party Artificial Intelligence Testing Laboratories shall be done through the National Institute of Standards and Technology’s National Voluntary Laboratory Accreditation Program (“NVLAP”). In accordance with current NVLAP processes, the National Institute of Standards and Technology shall determine and maintain the authoritative list for approved Third Party Artificial Intelligence Testing Laboratories.

(e) INDEPENDENT TESTING REQUIRED.—Upon the approval of Third Party Artificial Intelligence Testing Laboratories as outlined in subsection (d), a covered agency, prior to procuring or fielding an artificial intelligence system requiring testing, shall institute independent third party testing of the system to assess performance of the system according to attributes listed in section 22A of the National Institute of Standards and Technology Act.

(f) SCOPE OF TESTING.—Each independent Third Party Artificial Intelligence Testing Laboratory accredited pursuant to subsection (d) shall—

(1) utilize metrics relevant to the mission and authorities of the agency that intends to field the artificial intelligence system;

(2) develop approaches to test—

(A) the software product, as installed in a test facility; and

(B) relevant cloud-based services.

(3) establish binding data agreements that enable the agency and other stakeholders to share confidential and proprietary data with the testing entity without fear of inappropriate disclosure; and

(4) collaborate with the covered agency that is seeking testing to reach consensus on appropriate protocols and approaches for handling test data, test results, and analyses.

Recommendation Set 4: Strengthen Oversight and Governance Mechanisms to Address Current and Evolving Concerns.

Strengthen the Privacy and Civil Liberties Oversight Board’s (PCLOB) ability to provide meaningful oversight and advice to the federal government’s use of AI-enabled technologies for counterterrorism purposes.

SEC.____.—OVERSIGHT OF FEDERAL GOVERNMENT USE OF ARTIFICIAL INTELLIGENCE-ENABLED SYSTEMS FOR COUNTERTERRORISM PURPOSES.—

(a) AMENDMENTS TO AUTHORITIES AND RESPONSIBILITIES OF THE PRIVACY AND CIVIL LIBERTIES OVERSIGHT BOARD.—Section 2000ee of title 42, United States Code, is amended—

(1) in paragraph (2) of subsection (d), by—

(A) striking “and” at the end of subparagraph (B);

(B) redesignating subparagraph (C) as subparagraph (D); and

(C) adding a new subparagraph (C), as follows:

“(C) the development and use of artificial intelligence-enabled technologies for counterterrorism purposes; and”;

(2) in subparagraph (1)(A) of subsection (g), by striking the semicolon and adding the following: “and information about artificial intelligence-enabled technologies proposed to be acquired or fielded in the Federal Government (such as documentation of data collection, disclosure and consent processes for artificial intelligence-enabled tools and programs, documentation of models used and supporting training and testing, and any repurposing);”

(b) AMENDMENTS TO AUTHORITIES AND RESPONSIBILITIES OF PRIVACY AND CIVIL LIBERTIES OFFICERS.—Section 2000ee-1 of title 42, United States Code, is amended—

(1) in subsection (a), by—

(A) redesignating paragraphs (3) and (4) as paragraphs (4) and (5); and

(B) inserting a new paragraph (3), as follows:

“(3) provide prior notice to the Privacy and Civil Liberties Oversight Board of the fielding or repurposing of an artificial intelligence-enabled system (including a classified system) that could have an impact on privacy, civil liberties, or civil rights, and provide access to associated impact statements, including System of Record Notices, Privacy Impact Assessments, and Civil Rights and Civil Liberties Impact Assessments;” and

(2) in subsection (d), by striking the semicolon in paragraph (1) and inserting the following: “(including information described in paragraph (a)(3));”.

(c) SELF-ASSESSMENT BY PRIVACY AND CIVIL LIBERTIES OVERSIGHT BOARD.—Not later than 270 days after the date of the enactment of this act, the Privacy and Civil Liberties Oversight Board shall conduct and provide to Congress a self-assessment of any change in resources and organizational structure that may be required to carry out the artificial intelligence-related mission required by this section.

Empower DHS Offices of Privacy and Civil Rights and Civil Liberties.

SEC. ____.—ENHANCED OVERSIGHT OF ARTIFICIAL INTELLIGENCE-ENABLED SYSTEMS AT THE DEPARTMENT OF HOMELAND SECURITY.—

(a) AMENDMENT TO DUTIES AND RESPONSIBILITIES OF CIVIL RIGHTS AND CIVIL LIBERTIES OFFICER.—Section 345 of title 6, United States Code, is amended in paragraph (a)(5), by—

- (1) striking the final “and” in subparagraph (A);
- (2) redesignating subparagraph (B) as subparagraph (C); and
- (3) adding a new subparagraph (B), as follows:

“(B) ensure that the legal and approval processes for the procurement and use of artificial intelligence-enabled systems, including associated data of machine learning systems, provide appropriate consideration to the privacy, civil rights, and civil liberties impacts of such systems; and”.

(b) AMENDMENT TO DUTIES AND RESPONSIBILITIES OF CHIEF PRIVACY OFFICER.—Section 142 of title 6, United States Code, is amended in paragraph (a)(5), by—

- (1) striking the final “and” in subparagraph (A);
- (2) redesignating subparagraph (B) as subparagraph (C); and
- (3) adding a new subparagraph (B), as follows:

“(B) ensure that the legal and approval processes for the procurement and use of artificial intelligence-enabled systems, including associated data of machine learning systems, provide appropriate consideration to the privacy, civil rights, and civil liberties impacts of such systems; and”.

(c) ENHANCED PROCEDURES FOR CONSIDERATION OF PRIVACY AND CIVIL LIBERTIES ISSUES.—Not later than 270 days after the date of the enactment of this Act—

(1) the Secretary of Homeland Security shall revise the legal and approval processes for the procurement and use of artificial intelligence-enabled systems, including associated data of machine learning systems, to ensure that full consideration is given, with the participation of the Department's Chief Privacy Officer and the Officer for Civil Rights and Civil Liberties, to the privacy, civil rights, and civil liberties impacts of such systems; and

(2) the Department's Chief Privacy Officer and the Officer for Civil Rights and Civil Liberties shall report to Congress on any additional staffing or funding resources that may be required to carry out the requirements of this section.

Establish a task force to assess the privacy and civil rights and civil liberties implications of AI and emerging technologies.

SEC. ____.—TASK FORCE ON ORGANIZATIONAL STRUCTURE FOR ARTIFICIAL INTELLIGENCE GOVERNANCE AND OVERSIGHT.—

(a) ESTABLISHMENT.—Not later than 90 days after the date of the enactment of this Act, the President shall appoint a task force to assess the privacy, civil rights, and civil liberties implications of artificial intelligence and emerging technologies. This includes identifying policy and legal gaps and making recommendations to ensure that uses of artificial intelligence and associated data in U.S. government operations comport with freedom of expression, equal protection, privacy, and due process. The task force shall—

(1) assess existing policy and legal gaps for current AI applications and emerging technologies, and make recommendations for—

(A) legislative and regulatory reforms on the development and fielding of AI and emerging technologies; and

(B) institutional changes to ensure sustained assessment and recurring guidance on privacy and civil liberties implications of AI applications and emerging technologies.

(b) MEMBERSHIP OF TASK FORCE.—

(1) The task force shall include—

(A) the Attorney General or his or her designee;

(B) the Director of the Office of Management and Budget or his or her designee;

(C) the Director of the National Institute of Standards and Technology or his or her designee;

(D) the Comptroller General or his or her designee;

(E) the Inspectors General for the following agencies:

(i) the Department of State;

(ii) the Department of the Treasury;

(iii) the Department of Defense;

(iv) the Department of Justice;

(v) the Department of Health and Human Services;

(vii) the Department of Homeland Security;

(viii) the Office of the Director of National Intelligence; and

(ix) the Central Intelligence Agency.

(F) the chief privacy and civil liberties officers of each agency described in subparagraph (E);

(G) the Chair of the Privacy and Civil Liberties Oversight Board;

(H) the Chair of the National Artificial Intelligence Advisory Committee's Subcommittee on Artificial Intelligence and Law Enforcement; and

(I) representatives from civil society, including organizational leaders with expertise in technology, privacy, civil liberties, and civil rights, representatives from industry, and representatives from academia, as appointed by the President.

(2) TASK FORCE CHAIR AND VICE CHAIR.—The President shall designate a Chair and Vice Chair of the task force from among its members.

(c) RESPONSIBILITIES OF TASK FORCE.—The task force established pursuant to subsection (a) shall—

(1) conduct an assessment and make recommendations to Congress and to the President to ensure that the development and fielding of artificial intelligence and other emerging technologies by the Federal Government provides protections for the privacy, civil liberties, and civil rights of U.S. persons as appropriately balanced against critical law enforcement and national security needs;

(2) issue criteria for identifying qualified artificial intelligence systems and significant system refreshes requiring Artificial Intelligence Risk Assessment Reports and Artificial Intelligence Impact Assessments, under section [XX] of this Act;

(3) recommend baseline standards for Federal Government use of biometric identification technologies, including, but not limited to, facial recognition, voiceprint, gait recognition, and keyboard entry technologies;

(4) recommend proposals to address any gaps in Federal law or regulation with respect to facial recognition technologies in order to enhance protections of privacy, civil liberties, and civil rights of U.S. persons;

(5) recommend best practices and contractual requirements to strengthen protections for privacy, information security, fairness, non-discrimination, auditability, and accountability in artificial intelligence systems and technologies and associated data procured by the federal government;

(6) consider updates to and reforms of government data privacy and retention requirements to address implications to privacy, civil liberties, and civil rights;

(7) assess ongoing efforts to regulate commercial development and fielding of artificial intelligence and associated data in light of privacy, civil liberties, and civil rights implications, and as appropriate, consider and recommend institutional or organizational changes to facilitate applicable regulation; and

(8) assess the utility of establishing a new organization within the Federal Government to provide ongoing governance for and oversight over the fielding of artificial intelligence technologies by Federal agencies as technological capabilities evolve over time.

(d) ORGANIZATIONAL CONSIDERATIONS.—In conducting the assessment required by subsection (c)(7), the task force shall consider—

(1) the organizational placement, structure, composition, authorities, and resources that a new organization would require to provide ongoing guidance and baseline standards for—

(A) the Federal Government’s development, acquisition, and fielding of artificial intelligence systems to ensure they comport with privacy, civil liberties, and civil rights and civil liberties law, to include guardrails for their use and to disallow outcomes to be incorporated in policy and embedded in system development; and

(B) providing transparency to oversight entities and the public regarding the Federal Government’s use of artificial systems and the performance of those systems.

(2) the existing interagency and intra-agency efforts to address AI oversight;

(3) the need for and scope of national security carve outs, and any limitations or protections that should be built into any such carve outs; and

(4) the research, development, and application of new technologies to mitigate privacy and civil liberties risks inherent in artificial intelligence systems.

(e) REPORTING.—

(1) Not later than 180 days of establishment, the task force shall issue a report to Congress and the President with its legislative and regulatory recommendations. The task force shall provide periodic updates to the President and the Congress.

(2) Within a year of its establishment, the task force shall issue a report to the President and the Congress with its assessment on organizational considerations, to include any recommendations for organizational changes.

CHAPTER 10: THE TALENT COMPETITION

Blueprint for Action

Recommendation: Pass a National Security Immigration Act.

1) Grant Green Cards to All Students Graduating with STEM PhDs from Accredited American Universities.

2) Double the Number of Employment Based Green Cards.

3) Create an Entrepreneur Visa.

4) Create an Emerging and Disruptive Technology Visa.

NATIONAL SECURITY IMMIGRATION ACT OF 2021

SECTION. 1.—SHORT TITLE.—This Act may be cited as the “National Security Immigration Act of 2021.”

SEC. 2.—GREEN CARDS FOR STUDENTS GRADUATING FROM ACCREDITED AMERICAN UNIVERSITIES WITH DOCTORATES IN THE FIELDS OF SCIENCE, TECHNOLOGY, ENGINEERING, AND MATHEMATICS.—Section 1151 of title 8, United States Code, is amended in subsection (b)(1), by adding a new subparagraph (F), as follows:

“(F) Aliens who have been awarded doctoral degrees in the fields of science, technology, engineering, and mathematics by accredited universities in the United States.”

SEC. 3.—INCREASED AUTHORIZATION FOR EMPLOYMENT-BASED IMMIGRATION.—Section 1151 of title 8, United States Code, as amended by section 2, is further amended in subsection (d)(1)(A) by striking “140,000” and inserting “280,000”.

SEC. 4.—ENTREPRENEUR VISAS FOR HIGH PRIORITY SCIENCE AND TECHNOLOGY FIELDS AS DETERMINED BY NATIONAL SCIENCE FOUNDATION.—Section 1153 of title 8, United States Code, is amended in subsection (b)(5)—

(1) By redesignating subparagraphs (C) and (D) as subparagraphs (D) and (E); and

(2) By adding a new subparagraph (C), as follows:

“(C) PRIORITY FOR ENTREPRENEURS IN CERTAIN SCIENCE AND TECHNOLOGY FIELDS.—

“(i) Priority under this section shall be given to qualified immigrants who engage in new commercial enterprises in high priority science and technology fields, including artificial intelligence-enabled technology fields, as determined by the National Science Foundation.

“(ii) A qualified immigrant under this paragraph section shall not be required to meet the capital investment requirement in clause (A)(i) if the qualified immigrant is one of the principal organizers and operators of a new commercial enterprise described in clause (i).”

SEC. 5.—VISA FOR EMERGING AND DISRUPTIVE TECHNOLOGIES.—Section 1151 of title 8, United States Code, as amended by Sections 2 and 3, is further amended in subsection (b)(1), by adding a new clause (G), as follows:

“(G) Aliens who are students, researchers, entrepreneurs, and technologists in critical emerging and disruptive technology fields, as determined by the National Science Foundation.”

SEC. 6.—DETERMINATIONS BY THE NATIONAL SCIENCE FOUNDATION.—Not later than 180 days after the date of the enactment of this Act, and every three years thereafter, the National Science Foundation shall publish a list of—

(1) high priority science and technology fields in which qualified immigrants will be eligible for consideration for entrepreneur visas under section 1153(b)(5)(C) of title 8, United States Code, as amended; and

(2) critical emerging and disruptive technology fields in which qualified immigrants will be eligible for consideration for student, researcher, and entrepreneur visas under section 1151(b)(1)(G) of title 8, United States Code, as amended.

CHAPTER 11: ACCELERATING AI INNOVATION

Blueprint for Action

Recommendation: Scale and Coordinate Federal AI R&D Funding.

Component 1: Establish a National Technology Foundation.

THE NATIONAL TECHNOLOGY FOUNDATION ACT OF 2021

SECTION 1.—SHORT TITLE.—This Act may be cited as the “National Technology Foundation Act of 2021.”

SEC. 2.—ESTABLISHMENT OF NATIONAL TECHNOLOGY FOUNDATION.—There is established in the executive branch of the Government an independent agency to be known as the National Technology Foundation (hereinafter referred to as the “Foundation”). The Foundation shall consist of a National Technology Board (hereinafter referred to as the “Board”) and a Director of the Foundation (hereinafter referred to as the “Director”).

SEC. 3.—NATIONAL TECHNOLOGY BOARD.—

(a) The Board shall consist of twenty-four members to be appointed by the President and of the Director ex officio. In addition to any powers and functions otherwise granted to it by this chapter, the Board shall establish the policies of the Foundation, within the framework of applicable national policies as set forth by the President and the Congress.

(b) The term of office of each member of the Board shall be six years; except that any member appointed to fill a vacancy occurring prior to the expiration of the term for which his predecessor was appointed shall be appointed for the remainder of such term. Any person, other than the Director, who has been a member of the Board for twelve consecutive years shall thereafter be ineligible for appointment during the two-year period following the expiration of such twelfth year.

SEC. 4.—DIRECTOR OF THE FOUNDATION.—The Director shall be appointed by the President, by and with the advice and consent of the Senate. Before any person is appointed as Director, the President shall afford the Board an opportunity to make recommendations to the President with respect to such appointment. The Director shall receive basic pay at the rate provided for level II of the Executive Schedule under Section 5313 of title 5,

United States Code, and shall serve for a term of six years unless sooner removed by the President.

SEC. 5.—DEPUTY DIRECTOR OF THE FOUNDATION.—The Deputy Director (hereinafter referred to as the “Deputy Director”) shall be appointed by the President, by and with the advice and consent of the Senate. Before any person is appointed as a Deputy Director, the President shall afford the Board and the Director an opportunity to make recommendations to the President with respect to such appointment. The Deputy Director shall receive basic pay at the rate provided for level III of the Executive Schedule under section 5314 of title 5, United States Code, and shall perform such duties and exercise such powers as the Director may prescribe. The Deputy Director shall act for, and exercise the powers of, the Director during the absence or disability of the Director, or in the event of a vacancy in the office of Director.

SEC. 6.—GENERAL AUTHORITY OF THE FOUNDATION.—

(a) The Foundation shall have the authority, within the limits of available appropriations, to do all things necessary to carry out the provisions of this chapter, including, but without being limited thereto, to—

(1) distribute other payments for research and development in priority technology areas through grants, cooperative agreements, and contracts awarded to academic and private sector researchers, nonprofits, and consortia through competitive processes without regard to the provisions of sections 3324(a) and (b) of title 31, United States Code;

(2) establish an innovation unit in which independent program managers, brought into the Foundation on the basis of term appointments, fund proposals from both industry and academia to advance solutions to forward-looking research questions in priority technology areas;

(3) organize prize competitions to catalyze research around significant technology challenge problems;

(4) manage national technology resources, infrastructure, and initiatives that are assigned to the Foundation by statute or executive order;

(5) promote the commercialization of new technologies in priority technology areas and the transfer of such technologies to Federal, State and local government entities; and

(6) serve as a focal point for international research and development collaboration and standards-setting dialogues in priority technology areas.

SEC. 7.—PRIORITY TECHNOLOGY AREAS.—

(a) CORE DIRECTORATES.—The Foundation shall be organized into a set of core directorates, each dedicated to advancing fundamental research into a priority technology area.

(b) PRIORITY TECHNOLOGY AREAS.—Priority technology areas shall include—

- (1) artificial intelligence;
- (2) biotechnology;
- (3) quantum computing;
- (4) semiconductors and advanced hardware;
- (5) robotics and autonomy;
- (6) fifth-generation and advanced networking;
- (7) advanced manufacturing;
- (8) energy technology; and
- (9) any other technology area designated by the Congress or the Board.

(c) REVIEW OF KEY TECHNOLOGY FOCUS AREAS AND SUBSEQUENT LISTS.—

(1) ADDING OR DELETING KEY TECHNOLOGY FOCUS AREAS.—Beginning on the date that is four years after the date of enactment of this Act and every four years thereafter, the Director, acting through the Deputy Director shall—

(A) review the list of key technology focus areas, in consultation with the Board; and

(B) as part of that review, may add or delete key technology focus areas if the competitive threats to the United States have shifted and whether the United States or other nations have advanced or fallen behind in a technological area.

(2) LIMIT ON KEY TECHNOLOGY FOCUS AREAS.—Not more than ten key technology focus areas shall be included on the list of key technology focus areas at any time.

(3) UPDATING FOCUS AREAS AND DISTRIBUTION.—Upon the completion of each review under this subsection, the Director shall make the list of key technology focus areas readily available and publish the list in the Federal Register, even if no changes have been made to the prior list.

SEC. 8.—ADMINISTRATIVE MATTERS.—

(a) HIRING AUTHORITY.—

(1) PRIORITY TECHNOLOGY EXPERTS.—The Director shall have the authority to carry out a program of personnel management authority for the Foundation in the same manner, and subject to the same requirements, as the program of personnel management authority authorized for the Director of the Defense Advanced Research Projects Agency under section 1599h(a)(2) of title 10, United States Code, for the Defense Advanced Research Projects Agency.

(2) HIGHLY QUALIFIED EXPERTS.—In addition to the authority provided under subsection (A), the Director shall have the authority to carry out a program of personnel management authority for the Foundation in the same manner, and subject to the same requirements, as the program to attract highly qualified experts carried out by the Secretary of Defense under section 9903 of title 5, United States Code.

(3) ADDITIONAL HIRING AUTHORITY.—To the extent needed to carry out the duties of the Foundation, the Director shall utilize hiring authorities under section 3372 of title 5, United States Code, to staff the Foundation with employees from other Federal agencies, State and local governments, Indian tribes and tribal organizations, institutions of higher education, and other organizations, as described in that section, in the same manner and subject to the same conditions.

(b) EMPLOYMENT AND COMPENSATION OF CERTAIN PERSONNEL.—

(1) PROGRAM MANAGERS.—The employees of the Foundation may include program managers, who shall perform a role similar to program managers employed by the Defense Advanced Research Projects Agency, for the oversight and selection of programs supported by the Foundation.

(2) COMPENSATION OF MEMBERS OF BOARD.—The members of the Board shall be entitled to receive compensation for each day engaged in the business of the Foundation at a rate fixed by the Chairman but not exceeding the maximum rate payable under section 5376 of title 5, United States Code, and shall be allowed travel expenses as authorized by 5703 of title 5, United States Code. For the purposes of determining the payment of compensation under this subsection, the time spent in travel by any member of the Board shall be deemed as time engaged in the business of the Foundation. Members of the Board and

members of special commissions may waive compensation and reimbursement for traveling expenses.

SEC. 9.—INTERNATIONAL COOPERATION.—

(a) INTERNATIONAL AUTHORITY.—The Foundation is authorized to cooperate in any international technology activities consistent with the purposes of this Act and to expend for such international technology activities such sums within the limit of appropriated funds as the Foundation may deem appropriate.

(b) CONTRACTS AND ARRANGEMENTS.—

(1) The authority to enter into contracts or other arrangements with organizations or individuals in foreign countries and with agencies of foreign countries, as provided in section 1870(c) of title 42, United States Code, and the authority to cooperate in international scientific or engineering activities as provided in subsection (a) of this section, shall be exercised only with the approval of the Secretary of State, to the end that such authority shall be exercised in such manner as is consistent with the foreign policy objectives of the United States.

(2) If, in the exercise of the authority referred to in paragraph (1) of this subsection, negotiation with foreign countries or agencies thereof becomes necessary, such negotiation shall be carried on by the Secretary of State in consultation with the Director.

SEC. 10.—SECURITY PROVISIONS.—

(a) RESEARCH RELATED TO NUCLEAR ENERGY.— The Foundation shall not support any research or development activity in the field of nuclear energy, nor shall it exercise any authority pursuant to section 1870(e) of title 42, United States Code, in respect to that field, without first having obtained the concurrence of the Secretary of Energy that such activity will not adversely affect the common defense and security. To the extent that such activity involves restricted data as defined in the Atomic Energy Act of 1954, the provisions of that Act regarding the control of the dissemination of restricted data and the security clearance of those individuals to be given access to restricted data shall be applicable. Nothing in this chapter shall supersede or modify any provision of the Atomic Energy Act of 1954.

(b) RESEARCH RELATION TO NATIONAL DEFENSE.—

(1) In the case of priority technology area research activities under this Act in connection with matters relating to the national defense, the Secretary of Defense shall establish such security requirements and safeguards, including restrictions with respect to access to information and property, as the Secretary of Defense deems necessary.

(2) Any agency of the Government exercising investigatory functions otherwise within its jurisdiction is authorized to make such investigations and reports as may be requested by the Foundation in connection with the enforcement of security requirements and safeguards, including restrictions with respect to access to information and property, established under paragraph (1) of this subsection.

SEC. 11.—REPORTS.—

(a) INITIAL REPORT.—Not later than one year after the date of enactment of this Act, the Director shall transmit to the Committee on Commerce, Science, and Transportation of the Senate and the Committee on Science, Space, and Technology of the House of Representatives a report regarding the establishment of the Foundation. The report shall include an assessment of the priority technology focus areas as defined in this Act and of authorities that conflict with the National Science Foundation.

(b) ANNUAL REPORTS.—

(1) The Board shall submit to the President and the Congress no later than January 15 of each even numbered year, a report on indicators of the state of the priority technology areas in the United States, as defined in this Act.

(2) The Board shall render to the President and the Congress reports on specific, individual policy matters within the authority of the Foundation (or otherwise as requested by the Congress or the President) related to priority technology areas, as the Board, the President, or the Congress determines the need for such reports.

SEC. 12.—AUTHORIZATION OF APPROPRIATIONS.—

(a) INITIAL APPROPRIATION.—To enable the Foundation to carry out its powers and duties, including the establishment of a physical location, there is authorized to be appropriated to the Foundation \$30,000,000 for the first fiscal year following the enactment of this Act. Appropriations made pursuant to the authority provided in this subsection shall remain available for obligation, for expenditure, or for obligation and expenditure until expended for the Foundation's initial administrative costs and salaries and expenses.

(b) ANNUAL APPROPRIATION.—There are authorized to be appropriated for the Foundation, in addition to the appropriation provided in subsection (a) of this section and any other funds made available to the Foundation, a total of \$51,000,000,000 for fiscal years 2022 through 2026, of which—

(A) \$1,000,000,000 is authorized for fiscal year 2022;

(B) \$5,000,000,000 is authorized for fiscal year 2023;

(C) \$10,000,000,000 is authorized for fiscal year 2024;

(D) \$15,000,000,000 is authorized for fiscal year 2025; and

(E) \$20,000,000,000 is authorized for fiscal year 2026.

The Commission acknowledges additional authorities may be required to establish the NTF, including administrative, financial, and educational authorities mirroring those of the National Science Foundation, and that amendments to the NSF's statutory authorities may be required to alleviate duplication of duties. The Commission is ready to work with Congress to address such provisions.

Component 4: Invest in Talent that Will Transform the Field.

Direct and fund establishment of an AI Innovator Award.

Direct and fund establishment of a team-based AI research award.

SEC. ____.—ARTIFICIAL INTELLIGENCE AWARD PROGRAM.—

(a) ARTIFICIAL INTELLIGENCE INNOVATOR AWARD.—

(1) IN GENERAL.—The Director of the National Science Foundation shall partner with a nonprofit organization as described in subsection (c) to establish an Artificial Intelligence Innovator Award program to recognize and support the research of leaders in the field of artificial intelligence.

(2) ARTIFICIAL INTELLIGENCE INNOVATOR AWARD RECIPIENTS.—The Artificial Intelligence Award Selection Committee as described in subsection (d) shall select no fewer than 10 and no more than 20 award recipients each year. Recipients shall be selected for five-year, renewable award terms, based on a proven track record of prior innovation, a proposed general research program, a commitment to spend 75 percent of the recipients' time on research, and the committee's assessment of the potential of the research to generate breakthroughs in the area of artificial intelligence. Award amounts shall be determined by the selection committee with the objective of covering the full salary and benefits of the researcher and the cost of associated support staff and research equipment.

(b) ARTIFICIAL INTELLIGENCE TEAM AWARD.—

(1) IN GENERAL.—The Director of the National Science Foundation shall partner with a nonprofit organization as described in subsection (c) to establish an Artificial Intelligence Team Award program to support interdisciplinary research directed at applying artificial intelligence to solve complex problems or pursuing use-inspired basic research efforts to advance a fundamental understanding of the science of artificial intelligence in a manner that provides a significant benefit to society.

(2) ARTIFICIAL INTELLIGENCE TEAM AWARD RECIPIENTS.—The Artificial Intelligence Innovator Awards Selection Committee as described in paragraph (d) shall select no fewer than five and no more than 10 team recipients each year. Recipients shall be selected for five-year, nonrenewable terms, based on team qualifications, commitment to multi-disciplinary approaches, and innovative research proposals. Award amounts shall be determined by the selection committee with the objective of covering the cost of carrying out the proposed research proposal.

(c) NONPROFIT ORGANIZATION PARTNER.—The National Science Foundation shall partner with a nonprofit organization active in the field of computer science and artificial intelligence that maintains the requisite expertise and connections to the artificial intelligence research community to identify promising talent and invest in innovative ideas and to manage the award programs described in subsections (a) and (b), including to administer the programs and arrange the annual meeting.

(d) ARTIFICIAL INTELLIGENCE AWARD SELECTION COMMITTEE.—Recipients of the Artificial Intelligence Innovator Award and the Artificial Intelligence Team Award shall be selected by a rotating committee of artificial intelligence experts known as the Artificial Intelligence Award Selection Committee. The Committee shall consist of members chosen for their first-hand experience in artificial intelligence research and their familiarity with the frontiers of the field. Committee member selection shall be made by the nonprofit organization partner identified under subsection (c), in consultation with the Director of the National Science Foundation or designee.

(e) ANNUAL MEETING.—The Director of the National Science Foundation shall sponsor an annual meeting of recipients of the Artificial Intelligence Innovator Award and the Artificial Intelligence Team Award, at which the award recipients shall share information on the progress of their work.

(f) OTHER SOURCES OF FUNDING.—Nothing in this section shall be interpreted to preclude a recipient of an Artificial Intelligence Innovator Award or an Artificial Intelligence Team Award from pursuing supplemental government research grant or other research support provided by individuals, nonprofits and corporations, provided that such additional funding does not interfere with the recipient's commitment to the research program or require the assignment of ownership of intellectual property in a manner that would be inconsistent with the provisions of the Bayh-Dole Act, Public Law 96-517.

(g) INDEPENDENT REVIEW.—The Director of the National Science Foundation shall engage an independent entity to conduct a review to assess the successes and failures of the awards program authorized by this section, evaluate the impact of the funding level and award term on the research conducted by participants, and recommend any needed changes to the program (including any expansion or contraction in the number of

awards). The findings of the independent review shall be delivered to Congress not later than seven years after the commencement of the program.

(h) AUTHORIZATION OF APPROPRIATION.—

(1) There is authorized to be appropriated for each of the fiscal years 2022 through 2028 \$125,000,000 for the Artificial Intelligence Innovator Award.

(2) There is authorized to be appropriated for the Artificial Intelligence Team Award—

(A) \$50,000,000 for fiscal year 2022;

(B) \$100,000,000 for fiscal year 2023;

(C) \$150,000,000 for fiscal year 2024;

(D) \$200,000,000 for fiscal year 2025; and

(E) \$250,000,000 for fiscal years 2026 through 2028.

Recommendation: Leverage Both Sides of the Public-Private Partnership.

Component 2: Form a Network of Regional Innovation Clusters Focused on Strategic Emerging Technologies.

SEC. ____.—ESTABLISHMENT OF A NATIONAL NETWORK FOR REGIONAL INNOVATION IN EMERGING TECHNOLOGIES.—

(a) ESTABLISHMENT OF NATIONAL PROGRAM OFFICE.—The Secretary of Commerce shall establish, within the National Institute of Standards and Technology, a National Program Office for Regional Innovation in Emerging Technologies (referred to in this section as the ‘National Program Office’).

(b) DUTIES AND RESPONSIBILITIES.—The National Program Office, in coordination with representatives of Federal agencies with experience in and missions related to emerging technologies, shall—

(1) oversee the planning, development, management, and coordination of a National Network for Regional Innovation in Emerging Technologies (referred to in this section as the “National Network”);

(2) develop, not later than one year after the date of enactment, and update not less frequently than once every three years thereafter, a strategic plan to guide the development of the National Network to include identification of priority emerging technologies critical to national security or national competitiveness;

(3) use a competitive process to designate and provide financial assistance to regional innovation clusters that enable United States leadership in emerging technologies and support regional economic development throughout the United States;

(4) establish within each regional innovation cluster in the National Network a Technology Research Center for the purpose of facilitating collaboration among regional innovation cluster participants;

(5) establish such procedures, processes, and criteria as may be necessary and appropriate to coordinate the activities of the National Network and to maximize participation in and coordination with the National Network by Federal agencies that field or operate systems that incorporate emerging technologies;

(6) establish a clearinghouse of public information related to the activities of the National Network; and

(7) act as a convener of the National Network.

(c) DESIGNATION OF AND FINANCIAL ASSISTANCE IN SUPPORT OF REGIONAL INNOVATION CLUSTERS.—The National Program Office shall use a competitive process to designate and provide financial assistance to regional innovation clusters based on the following criteria:

(1) the equitable distribution of regional innovation clusters throughout the United States, taking into account factors such as proximity to the research and development facilities of Federal agencies, the level of support from state and local governments, the presence of and value proposition for leading firms and research institutions in relevant fields, and the size and education level of the local workforce;

(2) the capacity of regional innovation clusters to support the research, development, and commercialization of specific emerging technologies in areas that are critical to United States national competitiveness; and

(3) the clear potential for future development of regional innovation clusters that are not yet established technology hubs.

(d) TECHNOLOGY RESEARCH CENTERS.—The National Program Office shall establish within each regional innovation cluster in the National Network a Technology Research Center for the purpose of facilitating collaboration between regional innovation cluster participants. The Technology Research Centers shall—

(1) form sustained partnerships with anchor institutions in the region;

(2) host researchers on temporary assignments from Federal agencies, establish talent exchanges with local firms and research institutions, and fund multi-year, post-doctoral fellowships for the commercialization of research;

(3) host program managers from Federal agencies responsible for transitioning basic research into commercially viable technologies, identifying national security use cases and end users within the Federal Government, and initiating new Federal Government contracts to support technology transition;

(4) facilitate low cost access by regional innovation cluster participants to computing resources, curated datasets, testing infrastructure and ranges, and other research and development facilities owned or operated by the Federal government;

(5) establish intellectual property sharing agreements with regional innovation cluster participants to encourage Federal government adoption of commercial technologies; and

(6) when appropriate, provide for the publication of research in the open-source domain to encourage advances in the science and technology community more broadly.

(e) OTHER MATTERS.—

(1) RECOMMENDATIONS.—In developing and updating the strategic plan under subsection (b)(2), the National Program Office shall solicit recommendations and advice from a wide range of stakeholders, including industry, small and medium-sized enterprises, research universities, community colleges, state and local elected officials, and other relevant organizations and institutions on an ongoing basis.

(2) REPORT TO CONGRESS.—Upon completion of the strategic plan required by subsection (b)(2) or an update thereof, the National Program Office shall transmit the strategic plan to the Committee on Commerce, Science, and Transportation of the Senate and the Committee on Science, Space, and Technology of the House of Representatives.

(3) DETAILEES.—Any Federal Government employee may be detailed to the National Program Office without reimbursement. Such detail shall be without interruption or loss of civil service status or privilege.

(f) DEFINITIONS.—

(1) REGIONAL INNOVATION CLUSTER.—The term “regional innovation cluster” means a geographically bounded network of similar, synergistic, or complementary entities that —

(A) are engaged in or with a particular industry sector and its related sectors;

(B) have active channels for business transactions and communication;

(C) share specialized infrastructure, labor markets, and services; and

(D) leverage the region’s unique competitive strengths to stimulate innovation and create jobs.

(2) EMERGING TECHNOLOGIES.—For the purposes of this section the term “emerging technologies” may include such technologies as artificial intelligence, microelectronics, quantum computing, biotechnology, any associated, enabling or successor technologies, or any technologies identified by the National Program Office to be critical to national security or national competitiveness.

(g) AUTHORIZATION OF APPROPRIATION.—There is authorized to be appropriated to the Secretary of Commerce to carry out this section \$5,000,000 for fiscal year 2022.

CHAPTER 14: TECHNOLOGY PROTECTION*Blueprint for Action*

Recommendation: Reform CFIUS for Emerging Technology Competition.

Amend CFIUS’ authorizing legislation to require competitors to disclose investments in “sensitive technologies” to CFIUS.

SEC. ____ . REVIEW OF SENSITIVE TRANSACTIONS INVOLVING COUNTRIES OF SPECIAL CONCERN.

(a) TECHNICAL AMENDMENTS.—Section 721(a) of the Defense Production Act of 1950 (50 USC 4565(a)) is amended by redesignating paragraphs (4), (5), (6), (7), (8), (9), (10), (11), (12), and (13) as paragraphs (5), (6), (7), (9), (10), (11), (12), (13), (15), and (16), respectively.

(b) DEFINITION OF COUNTRY OF SPECIAL CONCERN.—Section 721(a) of the Defense Production Act of 1950 (50 USC 4565(a)) is amended by inserting after paragraph (3) the following:

“(4) COUNTRY OF SPECIAL CONCERN.—The term “country of special concern” means any country that is—

“(A) subject to export restrictions pursuant to section 744.21 of title 15, Code of Federal Regulations;

“(B) determined by the Secretary of State to be a state sponsor of terrorism; or

“(C) determined by the Committee to have a demonstrated or declared strategic goal of acquiring a type of technology or infrastructure that would have an adverse impact on United States leadership in areas related to national security, and is specified in regulations prescribed by the Committee.”

(c) DEFINITION OF SENSITIVE TECHNOLOGY.—Section 721(a) of the Defense Production Act of 1950 (50 USC 4565(a)) is amended by inserting after redesignated paragraph (7) the following:

“(8) SENSITIVE TECHNOLOGY.—The term ‘sensitive technology’ means any technology that is determined by the Committee to be necessary for maintaining or increasing the technological advantage of the United States over countries of special concern with respect to national defense, intelligence, or other areas of national security, or gaining such an advantage over such countries with respect to national defense, intelligence, or other areas of national security in areas where such an advantage may not exist, and is not a critical technology as defined in paragraph (7) of this subsection, and is specified in regulations prescribed by the Committee.

(d) DEFINITION OF SENSITIVE TRANSACTION INVOLVING A COUNTRY OF SPECIAL CONCERN.—Section 721(a) of the Defense Production Act of 1950 (50 USC 4565(a)) is amended by inserting after redesignated paragraph (13) the following:

“(14) SENSITIVE TRANSACTION INVOLVING A COUNTRY OF SPECIAL CONCERN.—The term ‘sensitive transaction involving a country of special concern’ means any investment in an unaffiliated United States business by a foreign person that—

“(A) is—

“(i) a national or a government of, or a foreign entity organized under the laws of, a country of special concern; or

“(ii) a foreign entity—

“(I) over which control is exercised or exercisable by a national or a government of, or by a foreign entity organized under the laws of, a country of special concern; or

“(II) in which the government of a country of special concern has a substantial interest; and

“(B) as a result of the transaction, could achieve—

“(i) influence, other than through voting of shares, on substantive decision making of the United States business regarding the use, development, acquisition, or release of sensitive technologies, as defined in this section; or—

“(ii) access to material nonpublic technical information related to sensitive technologies, as defined in this section, in the possession of the United States business.”

(e) DEFINITION OF COVERED TRANSACTIONS.—Section 721(a) of the Defense Production Act of 1950 (50 USC 4565(a)) is amended—

(1) in redesignated paragraph (5)(B)—

(A) in clause (iv)(I), by striking “or”;

(B) in clause (iv)(II), by striking the period and inserting “; or”; and

(C) by adding at the end the following:

“(III) a sensitive transaction involving a country of special concern.”

(2) by redesignating clause (v) as clause (vi) and inserting after clause (iv) the following:

“(v) Any sensitive transaction involving a country of special concern.”

(f) INFORMATION REQUIRED IN ANNUAL REPORT TO CONGRESS.—Section 721(m)(2) of the Defense Production Act of 1950 (50 USC 4565(m)(2)) is amended by adding at the end the following:

“(L) Identification of each country designated as a country of special concern along with an explanation of the rationale for such designation.

“(M) Identification of each technology designated as a sensitive technology along with an explanation of the rationale for such designation.”

(g) MANDATORY DECLARATIONS.—Section 721(b)(1)(C)(v)(IV)(bb)(AA) of the Defense Production Act of 1950 (50 USC 4565(b)(1)(C)(v)(IV)(bb)(AA)) is amended by inserting before the period “or is a sensitive transaction involving a country of special concern”.

(h) CONFORMING AMENDMENTS.—Title 50, United States Code, is amended—

(1) in section 4817(a)(1)(B) by striking “section 4565(a)(6)(A)” and inserting “section 4565(a)(7)(A)”;

(2) in section 4565(b)(4)(B)(ii) (section 721(b)(4)(B)(ii) of the Defense Production Act of 1950) by striking “subsection (a)(4)(B)(ii)” and inserting “subsection (a)(5)(B)(ii)”;

(3) in section 4565(b)(1)(c)(v)(III)(bb)(AA) (section 721(b)(1)(c)(v)(III)(bb)(AA) of the Defense Production Act of 1950) by striking “subsection (a)(4)(B) (iii)” and inserting “subsection (a)(5)(B)(iii)”;

(4) in section 4565(b)(1)(c)(v)(III)(bb)(BB) (section 721(b)(1)(c)(v)(III)(bb)(BB) of the Defense Production Act of 1950) by striking “subsection (a)(4)(B)(iii)” and inserting “subsection (a)(5)(B)(iii)”;

(5) in section 4565(b)(1)(c)(v)(III)(cc) (section 721(b)(1)(c)(v)(III)(bb)(BB) of the Defense Production Act of 1950) by striking “subsection (a)(4)(B)(iii)(II)” and inserting “subsection (a)(5)(B)(iii)(II)”.

Recommendation: Build Capacity to Protect the Integrity of the U.S. Research Environment. Establish a government-sponsored independent entity focused on research integrity.

SEC. ____.—Establishment of University Affiliated Research Center Focused on Research Integrity.—

(a) AGREEMENT AUTHORIZED.—Not later than 180 days after the date of the

enactment of this Act, the Secretary of Defense, acting through the Under Secretary of Defense for Research and Engineering and in consultation with the Director of the Office of Science and Technology Policy and other appropriate members of the Federal research community, shall enter into an agreement with a college or university to establish a University Affiliated Research Center to act as a center of excellence on research integrity and provide information and advice on research security.

(b) RESEARCH PURPOSES.—The University Affiliated Research Center established pursuant to subsection (a) shall—

(1) Maintain open source materials to serve university vetting of international engagement and risk management, including databases and risk assessment tools;

(2) Provide tailored guidance to research organizations for decision support on matters related to research security and integrity;

(3) Conduct comprehensive studies and regular reports on the state of foreign influence on U.S. research;

(4) Undertake independent investigations on research integrity;

(5) Develop education materials and tools for U.S. universities to build annual training and compliance initiatives; and

(6) Manage dialogue with stakeholder communities and provide a venue for information sharing among research organizations and Federal agencies.

*Recommendation: Counter Foreign Talent Recruitment Programs.
Mandate and resource compliance operations.*

SEC. ____.—Enhanced Review of Risk Posed by Applicants for Federal Grants.—

(a) ENHANCED REVIEW REQUIRED.—Not later than 180 days after the date of the enactment of this Act, the Director of the Office of Management and Budget shall revise section 200.206 of Part 2 of the Code of Federal Regulations to ensure that Federal grant-making agencies maintain compliance operations to guard against malign foreign talent recruitment programs and to prescribe standardized disclosure and accountability measures to support such compliance operations.

(b) DEFINITION.—For the purposes of this section, a “malign foreign talent recruitment program” is an effort directly or indirectly organized, managed, or funded by a foreign government to recruit science and technology professionals or students (regardless of citizenship or national origin) engaged in research funded by a federal agency to share information with or otherwise act on behalf of such foreign government.

Amend the Foreign Agents Registration Act.

SEC. ____.—AMENDMENT TO FOREIGN AGENTS REGISTRATION ACT. —Section 611 of title 22, United States Code, is amended in paragraph (1) of subsection (c) by—

- (1) Striking “and” at the end of clause (iv); and
- (2) Inserting at the end a new clause (v), as follows:

“(v) directly or indirectly organizes, manages, or funds an effort to recruit science and technology professionals or students (regardless of citizenship or national origin) engaged in research funded by a Federal agency to share information with or otherwise act on behalf of a foreign government; and”.

CHAPTER 15: A FAVORABLE INTERNATIONAL TECHNOLOGY ORDER

Blueprint for Action

Recommendation: Develop and Implement a Comprehensive U.S. National Plan to Support International Technology Efforts.

Core Goal #1: Shape International Technical Standards.

Establish a grant program to enable small- and medium-sized U.S. AI companies to participate in international standardization efforts.

SEC. ____.—SUPPORT FOR INDUSTRY PARTICIPATION IN INTERNATIONAL STANDARDS ORGANIZATIONS.—

(a) IN GENERAL.—Not later than 180 days after the date of the enactment of this Act, the Administrator of the Small Business Administration shall establish a program to support participation by small business concerns in meetings and proceedings of international standards organizations in the development of voluntary technical standards.

(b) GRANTS AUTHORIZED.—In carrying out the program authorized by subsection (a), the Administrator shall award competitive, merit-reviewed grants, to small business concerns to cover the reasonable costs, up to a specified ceiling, of participation of employees of such businesses in meetings and proceedings of international standards organizations. Participation may include regularly attending meetings, contributing expertise and research, proposing new work items, volunteering for leadership roles such as convenors and editors, and being early adopters of emerging standards. Recipients of awards under this subsection shall not be required to provide a matching contribution.

(c) AWARD CRITERIA.—The Administrator may provide under this section a grant award to covered entities that:

- (1) demonstrate deep technical expertise in key emerging technologies, including Artificial Intelligence and related technologies;

(2) commit personnel with such expertise to regular participation in international bodies responsible for setting standards for such technologies over the period of the grant; and

(3) agree to participate in efforts to coordinate between the U.S. government and industry to ensure protection of national security interests in the setting of international standards.

(d) EVALUATION.—In issuing awards under this section, the Administrator shall coordinate with the Director of the National Institute of Standards and Technology who shall provide support in the assessment of technical expertise in emerging technologies and standards setting needs.

(e) DEFINITIONS.—In this section:

(1) ADMINISTRATOR.—The term “Administrator” means the Administrator of the Small Business Administration.

(2) COVERED ENTITY.—The term “covered entity” means a small business concern that is incorporated in and maintains a primary place of business in the United States.

(3) SMALL BUSINESS CONCERN.—The term “small business concern” has the same definition as set out in section 632 of title 15, United States Code.

(f) AUTHORIZATION OF APPROPRIATION.—There is authorized to be appropriated for fiscal year 2022 and each fiscal year thereafter \$1,000,000 to carry out the program authorized in this section.

Core Goal #2: Implement a Coordinated U.S. National Policy for the IDDI.

Create an allocated Emerging Technology Fund for foreign operations and related programs of USAID and the Department of State.

SEC. ____.—EMERGING TECHNOLOGY FUND.—

(a) ESTABLISHMENT.—There is established within the Department of State an Emerging Technology Fund (“Fund”) to facilitate holistic planning of digital foreign assistance, digital development projects, emerging technology programs, and other related initiatives of the Department of State and the United States Agency for International Development and to ensure the efficient management, coordination, operation, and utilization of such resources.

(b) FUNDING.—Funds otherwise available for the purposes of subsection (a) may be deposited in such Fund.

(c) AVAILABILITY.—Amounts deposited into the Fund shall remain available until expended.

(d) EXPENDITURES FROM FUND.—Amounts deposited in the Fund shall be available for the purposes of subsection (a).

(e) TRANSFER AUTHORITY.—Amounts available in the Fund may be transferred to any account of the Department of State or the United States Agency for International Development authorized by the Secretary of State for the purposes of carrying out a program described in subsection (a). Any amount so transferred shall be credited to the account to which it is transferred. The transfer authority provided in this subsection is in addition to any other transfer authority available to the Department of State.

Recommendation: Enhance the United States' Position as an International Digital Research Hub.

Component #2: Establish the Multilateral AI Research Institute (MAIRI).

SEC. ____.—MULTILATERAL ARTIFICIAL INTELLIGENCE RESEARCH INSTITUTE.—

(a) ESTABLISHMENT.—Not later than 180 days after the date of the enactment of this Act, the Director of the National Science Foundation (“Director”) shall establish a Multilateral Artificial Intelligence Research Institute (“MAIRI”) that leverages the National Artificial Intelligence Research Institutes as well as contributions from international partners, U.S. Government agencies, and non-governmental partners to facilitate international collaborative research and development initiatives involving artificial intelligence (“AI”). MAIRI shall have both a physical center located in the United States and a virtual presence.

(b) PURPOSE.—The purpose of MAIRI shall be to facilitate collaboration of international artificial intelligence research, foster international artificial intelligence innovation, and develop the next generation global artificial intelligence workforce in a manner that comports with democratic values and helps to preserve free and open societies.

(c) INTERNATIONAL PARTNERS.—As authorized by section 1872 of title 42, United States Code, the Director, in coordination with the Secretary of State, shall seek to develop partnerships with foreign governments that have existing research agreements and collaborative relationships with the United States. The Director of MAIRI shall provide for international partners to collaborate in the governance of MAIRI, contingent upon appropriate contributions of financial support.

(d) OTHER PARTNERS.—To further the goals of MAIRI, the Director shall seek, as necessary, partnerships with other U.S. Federal departments and agencies, and their national laboratories, and non-governmental partners, such as from industry, academia, research institutions, and philanthropies on a project-by-project basis.

(e) FACILITATION.—The Director, in coordination with the Secretary of State, shall facilitate the operations of MAIRI by creating a trusted learning cloud and associated compute capacity to facilitate international collaborative research by enabling access to needed resources, compute, and data for shared innovation, research, and development

(f) RESEARCH AGENDA.—MAIRI shall work with international partners, as well as U.S. Government partners, as needed, to—

(1) develop principles for multilateral artificial intelligence research, which address the importance of research integrity, the need for transparency, the necessity of open data and data sharing, the development of risk-benefit frameworks, and the use of merit-based competition reviews for research proposals; and

(2) develop research priorities that leverage members' capabilities and may include the development of—

(A) shared, secure compute resources, including joint benchmarking projects and data sharing, pooling, and storing initiatives founded on commonly agreed principles that ensure trust, privacy and security;

(B) privacy-preserving artificial intelligence and machine learning technologies, including technologies like federated learning and on-device prediction that enable remote execution, encrypted computation through multi-party computation and homomorphic encryption, and differential privacy; and

(C) smart city technologies, aligned with democratic values, that promote sustainability as well as norms that should guide standards development at bodies like the ITU and technical standards bodies.

(g) SOLICITATION AUTHORIZED.—The Director is authorized to issue one or more solicitations to create a physical facility to support the establishment of MAIRI. Any such solicitation shall provide for the selection of an awardee on a competitive, merit-reviewed basis.

(h) FINANCIAL ASSISTANCE TO ESTABLISH AND SUPPORT MAIRI.—Subject to the availability of funds appropriated for this purpose, the Director, the Secretary of Energy, the Secretary of State, the Secretary of Commerce, and other Federal agency heads may award financial assistance, as determined by an agency head, to establish and support MAIRI and associated research.

(i) AUTHORIZATION OF APPROPRIATION.—There is authorized to be appropriated for fiscal years 2022 through 2027, in such funds as may be required, for the purpose of—

(1) establishing and maintaining a physical center for MAIRI in the United States;

(2) carrying out MAIRI research initiatives in cooperation with the National Science Foundation, the Department of Energy, the Department of State, and other appropriate federal agencies;

(3) creating a trusted learning cloud and associated compute capacity to facilitate international collaborative research;

(4) U.S. researchers' travel and associated expenses to participate in MAIRI workshops, conferences, and similar events; and

(5) the establishment of an endowment fund in cooperation with international partners.

Recommendation: Reorient U.S. Foreign Policy and the Department of State for Great Power Competition in the Digital Age.

Expedite necessary reorganization of the Department of State by passing legislation to create an Under Secretary for Science, Research and Technology (Q).

SEC. ____.—UNDER SECRETARY OF STATE FOR SCIENCE, RESEARCH AND TECHNOLOGY.—

(a) POSITION ESTABLISHED.—Subsection (b) of section 2651a of title 22, United States Code, is amended—

(1) in paragraph (1), by striking “6” and inserting “7”;

(2) by redesignating paragraph (4) as paragraph (5); and

(3) by inserting before redesignated paragraph (5) the following new paragraph:

“(4) UNDER SECRETARY FOR SCIENCE, RESEARCH AND TECHNOLOGY. There shall be in the Department of State, among the Under Secretaries authorized by paragraph (1), an Under Secretary for Science, Research and Technology, who shall have primary responsibility to assist the Secretary and the Deputy Secretary on matters related to international science and technology policy.”

(b) REORGANIZATION REQUIRED.—Not later than 180 days after the date of the enactment of this Act, the Secretary of State shall develop a plan to consolidate the

science and technology policy functions of the Department in a single division under the leadership of the Under Secretary for Science, Research and Technology.

CHAPTER 16: ASSOCIATED TECHNOLOGIES

Blueprint for Action

Recommendation: Foster a Vibrant Domestic Quantum Fabrication Ecosystem.

Enact a package of provisions that incentivizes the domestic design and manufacturing of quantum computers and their constituent materials.

SEC. ____.—TAX CREDIT FOR DOMESTIC DESIGN AND MANUFACTURING OF QUANTUM COMPUTERS AND CONSTITUENT MATERIALS.—

Section 41(d) of title 26, United States Code, is amended by adding at the end a new paragraph (5), as follows—

“(5) SPECIAL RULE FOR DOMESTIC DESIGN AND MANUFACTURING OF QUANTUM COMPUTERS AND CONSTITUENT MATERIALS.—

“(A) With regard to domestic design and manufacturing of qualified quantum computers and constituent materials, the term ‘qualified research’ shall include, in addition to research described in paragraph (1)—

“(i) the development and production of qualified quantum computers and constituent materials in the United States; and

“(ii) the training of United States persons with regard to the development and production of qualified quantum computers and constituent materials.

“(B) In this paragraph, the term ‘qualified quantum computers and constituent materials’ means—

“(i) any computers have been identified by the Secretary, in consultation with the Secretary of Commerce, as quantum computers; and

“(ii) any components or constituent parts of such computers that have been identified by the Secretary, in consultation with the Secretary of Commerce, as critical to the operation of such computers.”

General Note: Should Congress establish a National Technology Foundation pursuant to the Commission’s Chapter 11 recommendation, Congress should also review conflicting National Science Foundation authorities and delegating appropriate authorities to the NTF.

Appendix E: Funding Recommendation Table

Funding Recommendation Table

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 1 – Emerging Threats in the AI Era	1	Create a Foreign Malign Influence Response Joint Interagency Task Force (JIATF).	Office of the Director of National Intelligence	\$30 million	-
	2	Increase DARPA funding for media authentication, disinformation detection, attribution, and disruption.	Department of Defense: USD(R&E) - DARPA	\$60 million to \$80 million	-
	3	Fund a machine speed AI-enabled cyber defense acceleration study.	Department of Homeland Security	\$10 million	-
	4	Increase DARPA funding for AI-enabled cyber defense research.	Department of Defense: USD(R&E) - DARPA	\$20 million	-
	5	Increase National Institute of Standards and Technology AI testbed funding.	National Institute of Standards and Technology	\$25 million	-
	6	Provide funding for a SolarWinds threat review.	Cyberspace Solarium Commission	\$6.5 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 2 – Foundations of Future Defense	1	Establish a dedicated AI Fund.	Department of Defense: USD(R&E)	\$200 million	-
	2	Increase investments in AI R&D.	Department of Defense	\$8 billion	-
	3	Establish a fund to to accelerate procurement and integration of commercial AI solutions for business applications.	Department of Defense: Joint Artificial Intelligence Center	\$100 million	-
	4	Provide funding to build enterprise data sets.	Department of Defense: Office of the Chief Data Officer	\$125 million	-
	5	Provide funding for technology scouting tools, data, and a technology fellows program.	Department of Defense, USD(R&E)	\$10 million	-
Chapter 3 – AI and Warfare	1	Develop innovative operational concepts that integrate new warfighting capabilities with emerging technologies.	Department of Defense: USD(R&E)	\$5 million	-
	2	Incentivize experimentation with AI-enabled applications through the Warfighting Lab Incentive Fund (WLIF).	Department of Defense: USD(R&E)	\$10 million	-
	3	Encourage a culture of "Thinking Red."	Department of Defense: Joint Warfighting Analysis Center	\$2.5 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 3 – AI and Warfare	4	Direct the military services, in coordination with the Under Secretary of Defense (for Acquisition and Sustainment), the Joint Staff, and the Defense Logistics Agency, and enabled by enterprise services and expertise at the JAIC, to prioritize integration of AI into logistics and sustainment systems wherever possible.	Department of Defense: Office of the Deputy Secretary of Defense	\$100 million	-
	5	Define a joint warfighting network architecture by the end of 2021.	“Department of Defense: Office of the Chief Information Officer”	\$5 million	-
Chapter 5 – AI and the Future of National Intelligence	1	Work with the intelligence community to establish a 10-year, \$1 billion, Program of Record to provide long-term, predictable funding for technologies identified in the technology annex to the National Intelligence Strategy.	Office of the Director of National Intelligence	\$1 billion annually for FYs 2022-2032	-
Chapter 6 – Technical Talent in Government	1	Congress should create a National Reserve Digital Corps.	Office of Management and Budget	\$16 million	-
	2	Congress should establish a STEM Corps.	Department of Defense	\$5 million for FY 2022 & \$5 million for FY 2023	-
	3	Congress should create a United States Digital Service Academy.	New Entity	\$40 million initial appropriation	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail
Chapter 7 – Establishing Justified Confidence in AI Systems	1 Appoint responsible AI leads and supporting staff in each agency critical to national security.	Department of Defense; Office of the Director of National Intelligence; Department of Homeland Security; Federal Bureau of Investigation; Department of State; Department of Energy; & Department of Health and Human Services	\$21.5 million	This funding supports one responsible AI lead and two supporting staff. Additionally, the funding includes responsible AI leads for each of the armed services in the Department of Defense and each of the agencies of the Intelligence Community.
Chapter 8 – Upholding Democratic Values	1 Congress should establish third-party testing center(s) to allow independent, third-party testing of national security-related AI systems that could impact U.S. persons.	National Institute of Standards and Technology	\$1.2 million	-
Chapter 9 – A Strategy for Competition and Cooperation	1 Create a Technology Competitiveness Council.	The White House: Executive Office of the President	\$2 million	-
Chapter 10 – The Talent Competition	1 Congress should pass a new National Defense Education Act.	Department of Education; National Science Foundation	One time appropriation of \$8.2 billion	-
Chapter 11 – Accelerating AI Innovation	1 Establish a National Technology Foundation.	New Entity	\$30 million initial appropriation for start-up expenses; \$1 billion for FY 2022; \$5 billion for FY 2023; \$10 billion for FY 2024; \$15 billion for FY 2025; & \$20 billion for FY 2026	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 11 – Accelerating AI Innovation	1	Increase federal funding of Non-Defense AI R&D at compounding levels.	Multiple agencies, including: the NSCAI proposed National Technology Foundation; National Science Foundation; Department of Energy; National Institute of Standards and Technology; National Institutes of Health; & National Aeronautical and Space Administration	\$2 billion for FY 2022; \$4 billion for FY 2023; \$8 billion for FY 2024; \$16 billion for FY 2025; & \$32 billion for FY 2026	-
	2	Expand the Network of AI Research Institutes.	National Science Foundation	\$200 million for FY 2022; \$200 million for FY 2023; & \$200 million for FY 2024	-
	3	Establish an AI Innovator Award.	National Science Foundation	\$125 million	-
	4	Establish a team-based AI Award.	National Science Foundation	\$50 million for FY 2022; \$100 million for FY 2023; \$150 million for FY 2024; \$200 million for FY 2025; & \$250 million annually for FYs 2026-2028	-
	5	Implement the NAIRR Roadmap.	National Science Foundation	\$30 million	-
	6	Fund an AI Data Program.	Department of Energy	\$25 million	-
	7	Sponsor an Open Knowledge Network.	National Science Foundation	\$25 million	-
	8	Form a network of Regional Innovation Clusters.	National Institute of Standards and Technology	\$200 million for FYs 2022-2026	Funding recommended at \$20 million per Regional Innovation Cluster

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail
Chapter 13 – Microelectronics	1 Increase federal grants for microelectronics manufacturing.	Department of Commerce	\$15 billion total	\$3 billion per project on average
	2 Increase funding for DARPA’s Electronics Resurgence Initiative (ERI).	Department of Defense: USD(R&E) - DARPA	\$400 million for FY 2022 & \$5 billion total for FYs 2022-2026	These funding levels should ramp up on an annual basis as absorptive capacity increases
	3 Increase funding for National Science Foundation semiconductor research.	National Science Foundation	\$300 million for FY 2022 & \$2.5 billion total for FYs 2022-2026	These funding levels should ramp up on an annual basis as absorptive capacity increases
	4 Increase funding for Department of Energy semiconductor research.	Department of Energy	\$400 million for FY 2022 & \$4.5 billion total for FYs 2022-2026	These funding levels should ramp up on an annual basis as absorptive capacity increases
	5 Establish the Advanced Packaging National Manufacturing Program.	National Institute of Standards and Technology	\$1 billion for FY 2022 & \$5 billion total for FYs 2022-2026	-
	6 Establish the National Semiconductor Technology Center.	Department of Commerce in collaboration with the Department of Defense and Department of Energy	\$100 million FY 2022 & \$2 billion total for FYs 2022-2026	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail
Chapter 15 – A Favorable International AI Order	1 Provide funding for U.S. International Development Finance Corporation to execute development financing for technology infrastructure projects.	U.S. International Development Finance Corporation	\$1 billion	-
	2 Provide funding to support U.S. International Development Finance Corporation development financing initiatives.	Department of State; U.S. Agency for International Development	\$200 million	-
	3 Provide funding for U.S. Agency for International Development Digital Strategy.	U.S. Agency for International Development: Bureau of Democracy, Development, and Innovation	\$200 million	-
	4 Provide funding for an Interagency AI Standards team to support National Institute of Standards and Technology AI Standards Coordinator and fund travel and other administrative needs.	National Institute of Standards and Technology; Department of Defense; Department of State; Office of the Director of National Intelligence; Department of Energy; Department of Homeland Security; U.S. Agency for International Development	\$3.3 million	Funding includes five full-time employee (FTE) from National Institute of Standards and Technology and one FTE from each of the following departments and agencies: Department of Defense, Department of State, Office of the Director of National Intelligence, Department of Energy, Department of Homeland Security, and U.S. Agency for International Development.
	5 Provide funding to support grants for small- and medium-sized businesses to participate in international data and technical standards efforts.	Small Business Administration	\$1 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail
Chapter 15 – A Favorable International AI Order	6 Funding for administrative costs associated with establishing an U.S. Center of Expertise relationship with GPAI/OECD.	National Science Foundation	\$1 million	-
	7 Funding for the Multilateral AI Research Initiative (MAIRI), including establishing and maintaining physical center; supporting research initiatives; created a trusted learning cloud resource; and supporting U.S. researchers' travel and involvement in workshops, conferences, and events.	National Science Foundation; Department of State; Department of Energy	\$12.15 million annually for FYs 2022-2027	\$10 million to National Science Foundation/ Department of State/ Department of Energy for research and personnel; \$2M to National Science Foundation for infrastructure; \$150,000 to National Science Foundation for administrative costs.
	8 Provide funding for trusted learning cloud to facilitate collaborative R&D with allies and partners (envisioned as a component of MAIRI).	National Science Foundation; Department of State	\$11.3 million	Funding includes underlying infrastructure, data storage and sharing capacity, grants for researchers, foreign assistance grants.
	9 Provide funding to support grants for scholars and researchers to participate in international data and technical standards efforts.	Department of State	\$5 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail
Chapter 15 – A Favorable International AI Order	10 Provide funding for immediate augmentation and training of U.S. diplomatic corps for efforts related to AI and emerging technology (funding does not include future funding needs which we recommend be determined by a focused planning effort to be undertaken by Department of State).	Department of State	\$8 million	\$550,000 - STAS; \$550,000 - Office of Communication and Information Policy; \$400,000 - Office of Science and Technology Cooperation; \$3.8 million - Regional Technology Officers (12 locations); \$1.25 million - Office of the Special Representative to Silicon Valley; \$450,000 - FSI training.
	11 Provide funding for the Bureau of Cyberspace Security and Emerging Technologies.	Department of State	\$20 million	-
	12 Provide funding for public diplomacy and engagement activities on AI innovation and democratic values.	Department of State	\$5.5 million	-
	13 Provide funding for AI exchange programs to support U.S. values and fund participation by developing countries in multilateral AI activities.	Department of State	\$8.5 million	-
	14 Provide funding for efforts to promote U.S. innovation and values and support American Spaces, Tech Camps, Maker Spaces, Speakers Program, and other initiatives.	Department of State	\$3 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 15 – A Favorable International AI Order	15	Provide funding for tracking and analysis of public opinion to measure impact of engagement efforts and guide strategic planning.	Department of State	\$1 million	-
	16	Provide funding for U.S. Science Envoys and Embassy Science Fellows programs.	Department of State	\$1 million	-
	17	Provide funding to support U.S. leadership in AI through Emerging Technology Coalition and internal programs.	Department of State: Office of the Under Secretary for Economic Growth, Energy, and the Environment (E)	\$5.5 million	Funding includes ETC support, creation of an advisory committee on emerging technology, private sector engagement, multilateral R&D efforts, tech-oriented diplomatic efforts, innovation enhancements.
	18	“Funding to support promotion of human rights and fundamental freedoms in AI context through civil society initiatives, promoting AI and emerging tech to counter censorship, and supporting research and awareness campaigns”	Department of State: Office of the Under Secretary for Civilian Security, Democracy, and Human Rights (J): Bureau of Democracy, Human Rights, and Labor (DRL)	\$1.5 million	-
	19	Provide funding to support use of AI for national security/military applications through cooperation with allies and partners, to include joint exercises, grants, fellowships, and other activities.	“Department of State: Office of the Under Secretary of State for Arms Control and International Security (T)”	\$3 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 15 – A Favorable International AI Order	20	Provide funds to support building technical capacity in emerging democracies and market economies to counter malign influence.	Department of State	\$3 million	-
	21	Provide funds to support research grants on malign influence in AI ecosystems.	Department of State	\$2 million	-
	22	Provide funds to support public diplomacy initiatives on international AI standards and tracking and reporting of impact on public engagement.	Department of State	\$2 million	-
	23	Provide funds to support US Global Innovation through Science and Technology (GIST) Initiative.	Department of State	\$1 million	-
	24	Provide additional funding to support foreign assistance activities around emerging tech and digital infrastructure, to include planning, assessments, and provision of assistance. Funds would support targeted, digital programs in several areas, including rule of law (INL), democracy and human rights (DRL), security cooperation (AVC/PM/ISN), and technical assistance (EB, STAS, others).	Department of State	\$230 million	-

*Unless otherwise noted funding is annual beginning in Fiscal Year 2022.

**All funding figures should be considered initial estimates for consideration by Congress and the Executive Branch.

Appendix F:

Commissioner Bios



Dr. Eric Schmidt, Chair

Dr. Eric Schmidt is an accomplished technologist, entrepreneur, and philanthropist. He joined Google in 2001 and helped grow the company from a Silicon Valley startup to a global leader in technology alongside founders Sergey Brin and Larry Page. Schmidt served as Google's Chief Executive Officer and Chairman from 2001 to 2011, as well as Executive Chairman and Technical Advisor. Under his leadership, Google dramatically scaled its infrastructure and diversified its product offerings while maintaining a strong culture of innovation.

In 2017, he co-founded Schmidt Futures, a philanthropic initiative that bets early on exceptional people making the world better. Schmidt is the host of "Reimagine with Eric Schmidt," a podcast series of conversations with leaders to explore how society can build a brighter future after the COVID-19 pandemic.



The Honorable Robert Work, Vice Chair

Robert Work was the 32nd Deputy Secretary of Defense, serving alongside three Secretaries of Defense from May 2014 to July 2017. In 2001, he retired as a Colonel in the United States Marine Corps after spending 27 years on active duty. He subsequently served as Senior Fellow and Vice President and Director of Studies at the Center for Strategic and Budgetary Assessments. In January 2009, he was asked to join the Obama administration as the 31st Under Secretary of the Navy, and was confirmed in that role in May 2009. Work stepped

down as the Under Secretary in March 2013 to become the Chief Executive Officer for the Center for a New American Security (CNAS). He remained in that position until he assumed the role of Deputy Secretary of Defense in May 2014. He currently is the President and Owner of TeamWork, LLC, which specializes in defense strategy and policy, programming and budgeting, military-technical competitions, revolutions in war, and the future of war.



Safra Catz

Safra A. Catz has served as chief executive officer of Oracle Corporation since 2014 and a member of the company’s board of directors since 2001. She joined Oracle in 1999 and held various positions within the company, including President and Chief Financial Officer, prior to being named CEO. Catz currently serves as a director of The Walt Disney Company and previously served as a director of HSBC Holdings plc.



Dr. Steve Chien

Dr. Steve Chien is a Technical Fellow, Senior Research Scientist, and the Technical Group Supervisor of the Artificial Intelligence Group at the Jet Propulsion Laboratory, California Institute of Technology. Chien has led the deployment of AI software to a wide range of missions. He is currently supporting the development of onboard and ground automated scheduling for the Mars 2020 rover mission, as well as scheduling technologies for the ECOsystem Spaceborne Thermal Radiometer Experiment on Space Station (ECOSTRESS) and Orbiting Carbon


Observatory 3 (OCO-3). Chien has received numerous awards for these efforts, to include Lew Allen Award for Excellence, JPLs highest award recognizing outstanding technical achievements by JPL personnel in the early years of their careers. He has been recognized four times in the NASA Software of the Year competition and has received four NASA medals for his work in AI for space. In 2011, Chien was awarded the inaugural American Institute of Aeronautics and Astronautics Intelligent Systems Award for his contributions to spacecraft autonomy.



The Honorable Mignon Clyburn

Mignon L. Clyburn served as Commissioner on the Federal Communications Commission (FCC) from 2009 to 2018, and acting chair from May to November of 2013. During her nearly nine years at the FCC, Mignon was committed to closing persistent digital and opportunities divides that continue to challenge rural, Native, and low wealth communities. Previously, Clyburn served for 11 years on the South Carolina Public Service Commission. Prior to that, she was the publisher and general manager of the Coastal Times,

a family-founded, Charleston-based weekly newspaper focusing on issues affecting the African American community. Clyburn is currently the principal of MLC Strategies, LLC.



Christopher Darby

Christopher Darby has served as President and CEO of In-Q-Tel since September 2006. He is also a member of its Board of Trustees. Prior to joining In-Q-Tel, Darby was a Vice President and General Manager at Intel, where he oversaw the Middleware Products Division. He joined Intel in August 2005 with the acquisition of Sarvega, a venture-backed supplier of XML networking and security products, where he served as President and CEO. Prior to Sarvega, Darby was the Chairman and CEO of @stake, an Internet security consulting firm ultimately acquired by

Symantec. Before that, Darby served as President and CEO of Interpath Communications, which was later acquired by US Internetworking. Earlier in his career, he held several executive positions at Digital Equipment Corporation (now Hewlett-Packard) and Northern Telecom (now Nortel Networks). Chris began his career at Bell Northern Research.



Dr. Kenneth Ford

Dr. Kenneth Ford is Founder and CEO of the Institute for Human & Machine Cognition. His research interests include AI, human-centered computing, and human performance and resilience. Ford is a Fellow of the Association for the Advancement of AI (AAAI), and a charter Fellow of the National Academy of Inventors. He has received many awards and honors including the Doctor Honoris Causas from the University of Bordeaux in 2005, the 2008 Robert Englemore Award for his work in AI, and the AAAI Distinguished Service Award in 2015. In 2015, he was elected as Fellow of

the American Association for the Advancement of Science and in 2017 was inducted into the Florida Inventors Hall of Fame. Ford has served on the National Science Board, the Air Force Science Advisory Board, and the Defense Science Board. In 2008, he was named as Chairman of the NASA Advisory Council – a capacity in which he served through 2011. In 2010, Ken was awarded NASA's Distinguished Public Service Medal – the highest honor the agency confers.



Dr. José-Marie Griffiths

Dr. José-Marie Griffiths is president of Dakota State University in Madison, South Dakota. Griffiths has spent her career in research, teaching, public service, corporate leadership, economic development, and higher education administration. She has served in presidential appointments to the National Science Board, the U.S. President's Information Technology Advisory Committee, and the U.S. National Commission on Libraries and Information. Griffiths has led projects for more than 28 U.S. federal agencies such as the National Science Foundation

and NASA, and more than 20 major corporations including, AT&T Bell Laboratories and IBM, in more than 35 countries. She also has worked with seven major international organizations, including NATO and the United Nations. She has received over 20 significant awards in science, technology, teaching, and the advancement of women in these fields.



Dr. Eric Horvitz

Dr. Eric Horvitz is a technical fellow at Microsoft, where he serves as the company's first Chief Scientific Officer. Horvitz provides cross-company leadership and perspectives on advances and trends on scientific matters, and on issues and opportunities arising at the intersection of technology, people, and society. He is recognized for his research on challenges and opportunities with the uses of AI technologies amidst the complexities of the open world. Horvitz is the recipient of the Feigenbaum Prize and the Allen Newell Prize for contributions to AI.



Andrew Jassy

Andy Jassy is the founder and CEO of Amazon Web Services (AWS), the world's most comprehensive and broadly adopted cloud platform. Jassy launched AWS in 2006 and has managed an inventive and nimble team that has delivered more than 165 services for compute, storage, networking, databases, analytics, mobile, Internet of Things, Artificial Intelligence, security, hybrid, and enterprise applications. Prior to founding AWS, Jassy held

several leadership positions across Amazon. Shortly after joining the company in 1997, he authored the business plan for Amazon's Music business and served as its Director of Product Management and General Manager. Jassy also started the Amazon Customer Relationship Management team, led marketing for Amazon, and was Technical Advisor (shadow) to Amazon Founder and CEO Jeff Bezos.



Gilman Louie

Gilman Louie is Co-Founder and Partner of Alsop Louie Partners, an early-stage technology venture capital firm founded in 2006. From 1999 until 2006, Louie was the first CEO of In-Q-Tel. Prior to In-Q-Tel, Louie built a career as a pioneer in the interactive entertainment industry, during which he founded and ran a publicly traded company called Spectrum HoloByte, and served as Chief Creative Officer of Hasbro Interactive. He serves as a member of the Board of Directors for the Markle Foundation, Maxar Technologies, Niantic, Lookingglass Cyber Solutions,

Aurora Insights and various other private companies and non-profit foundations. He is also Chairman of the Board of the Federation of American Scientists. Louie has served as a member of the Technical Advisory Group of the United States Senate Select Committee on Intelligence, and as a Commissioner of the National Commission for Review of Research and Development Programs of the United States Intelligence Community. He has received dozens of awards for his achievements, including from the NGA, CIA, and DNI, and in 2002 was named as one of fifty scientific visionaries by Scientific American.



Dr. William Mark

Dr. William Mark leads SRI International's Information and Computing Sciences division, creating new technology in machine learning, virtual personal assistance, trusted systems, and speech and vision analytics. The group also commercializes technology, licensing to corporations and creating spinoff companies such as Siri, Kasisto, CurieAI, and LatentAI. Prior to joining SRI International, Mark headed research groups at National Semiconductor, Lockheed Martin, and the University of Southern California Information Sciences Institute.



Dr. Jason Matheny

Dr. Jason Matheny is the founding director of Georgetown University's Center for Security and Emerging Technology (CSET). Previously he served as Assistant Director of National Intelligence, and Director of IARPA, responsible for the development of breakthrough technologies for the U.S. intelligence community. Before IARPA, he worked at Oxford University, the World Bank, the Johns Hopkins University Applied Physics Laboratory, the Center for Biosecurity, and Princeton University, and was the co-founder of two biotechnology companies.



The Honorable Katharina McFarland

Katharina McFarland serves as Chairman of the Board of Army Research and Development at the National Academies of Science, and as a Director on the Boards of SAIC, Exyn Technologies, and the Procurement Round Table. With more than 30 years of government service, McFarland is widely recognized as a leading subject matter expert on government procurement. She also serves as an advisor to Raytheon Missile Systems Division Senior Advisory Board, Cypress International Senior Strategy Group, Transunion Corporation Advisory

Board, and Sehlke, Inc. Senior Advisory Board. From 2012 to 2017, McFarland served as the Assistant Secretary of Defense for Acquisition and as acting Assistant Secretary of the Army (Acquisition, Logistics & Technology) from 2016-2017. She was President of the Defense Acquisition University from 2010 to 2012, and the Director of Acquisition at the Missile Defense Agency from 2006 to 2010. She has received an Honorary Doctoral of Engineering from the University of Cranfield in the United Kingdom, the Presidential Meritorious Executive Rank Award, the Secretary of Defense Medal for Meritorious Civilian Service Award, the Department of the Navy Civilian Tester of the Year Award, and the Navy and United States Marine Corps Commendation Medal for Meritorious Civilian Service.



Dr. Andrew Moore

Dr. Andrew W. Moore is a distinguished computer scientist with expertise in machine learning and robotics. He became the head of Google Cloud Artificial Intelligence division in January 2019. Moore previously worked at Google from 2006 to 2014 and was the founding director of Google's Pittsburgh engineering office in 2006. He then spent a four-year hiatus at Carnegie Mellon University as the dean of the School of Computer Science. Moore's research interests encompass the field of "big data" — applying statistical methods and mathematical

formulas to massive quantities of information, ranging from web searches to astronomy to medical records, in order to identify patterns and extract meaning from that information. His past research has included improving the ability of robots and other automated systems to sense the world around them and respond appropriately.

Appendix G: Commission Staff and Contributors

EXECUTIVE STAFF

Yll Bajraktari,
Executive Director

Michael J. Lueptow,
General Counsel

Michael L. Gable,
Chief of Staff

Tara M. Rigler,
Director of Strategy,
Communications & Engagements

Angela A. Ponmakha,
Director of Operations,
Designated Federal Officer

LEGISLATIVE AFFAIRS

Brandon McKee

Jenilee Keefe Singer

SENIOR ADVISORS

Dr. Seth Center

Robert Nelson

RESEARCH AND ANALYSIS

Courtney Barno
Dr. Ryan Carpenter
Matthew Cordova
Caroline Danauy
Raina Davis
Tess deBlanc-Knowles
Rama Elluru
Michael Garris

Matthew Gentzel
Charles Howell
LTC Michael Jackson, USA
Rebekah Kennel
Jeffrey Kojac
David Kumashiro
CAPT Lance Lantier, USN
Christie Lawrence

Paul Lekas
Dr. Margaret Lentz
Quinn Lorenz
Justin Lynch
Col Paul “P.J.” Maykish, USAF
Kevin McGinnis
Christopher McGuire

Paul Rhodes
Dr. Christopher Rice
Joe Wang
Parker Wild
Jessica Young
Olivia Zetter

OPERATIONS AND LEGAL TEAM

Chelsea Holt
Sarah Johnson
Brent Myles

Jennifer Sheehan
Angela Stacks
Jamie Tomberlin

INTERNS

Richard Altieri
Madeline Blanchard
Sabrina Broderick
Shaantam Chawla
Devin Davidson
Nickie Deahl
Hudson Dizon
Dylan Halpern
Courtney Lange
Alexander Mann
Nikhil Marda
M. Marin Ruelas Mendoza

Ariana Orne
Sultan Seraj
Katie Stolarczyk
Jaide Tarwid
Christopher Tonelli
Claire Trotter
Samuel Trotter
Aristotle Vainikos
Jackson Valen
Zoe Weinberg
Kate Yeager

The Commission would like to thank Members of Congress, Congressional Staff, government personnel, industry professionals, academia, members of the public, and all others who participated in, advised on, or commented on our work. The unified effort of everyone involved made this document possible. Additionally, the Commission thanks all of our outside contributors whose hard work resulted in this product.

OUTSIDE CONTRIBUTORS

John Bansemer
Susanna Blume
Dr. Anne Bowser
Scott Britt
Kristy Colbert
Mark Cohen
Dr. David Danks
Jeffrey Ding
Dr. Kathleen Fisher
Amanda Foley
Dr. Kristin Gilkes
Dr. Bryce Goodman
Gregory Grant
Erin Hahn
Orin Hoffman
Dajonte Holsey
Dr. Michael Horowitz
Dr. Andrew Imbrie
Taylor Lineberger
Dr. Albana Shehaj
Dr. Paul Scharre
Dr. William Scherlis
Raj Shah
John “Jack” Shanahan
Dr. Bernadette Johnson

Elsa Kania
Dr. Christopher Kirchhoff
Zachary Kuehn
Thomas Kalil
Peter Levine
Frank Long
Michael “Brendan” McCord
Michael McNerney
Tariq Mehmood
Paul Michel
Dr. Nadia Schadlow Murphy
Adam Mossoff
Geoffrey Odlum
Scott Padgett
Jared C. Ponmakha
Douglas Rand
Dr. Heather Roff
Craig Smith
Michael Soos
Dean Souleles
Dr. Barbara Stephenson
Francoise von Trapp
German Wegbrait
Darren Wright
Dr. Amy Zegart

A special thanks to Lirijon Kadriu for designing the Commission’s logo.



NATIONAL
SECURITY
COMMISSION
ON ARTIFICIAL
INTELLIGENCE