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 another[7] 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 preductive 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 designs[42] 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/advmlthreatmatrix/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 Model*s, 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 Definition*s, 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://tonyroyphilosophy.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.