Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
accountability	1) relates to an allocated responsibility. The responsibility can be based on regulation or agreement or through assignment as part of delegation; 2) For systems, a property that ensures that actions of an entity can be traced uniquely to the entity; 3) In a governance context, the obligation of an individual or organization to account for its activities, for completion of a deliverable or task, accept the responsibility for those activities, deliverables or tasks, and to disclose the results in a transparent manner.	ISO/IEC_TS_5723:2022(en)	"accountable" (adjective vs. noun): answerable for actions, decisions, and performance	ISO/IEC_TS_5723:2022(en)								
accuracy	Closeness of computations or estimates to the exact or true values that the statistics were intended to measure.	OECD	A qualitative assessment of correctness or freedom from error.	FDA_Glossary	The measure of an instrument's capability to approach a true or absolute value. It is a function of precision and bias.	FDA_Glossary	The accuracy of a machine learning system is measured as the percentage of correct predictions or classifications made by the model over a specific data set. It is typically estimated using a test or "hold out" sample, other than the one(s) used to construct the model. Its complement, the error rate, is the proportion of incorrect predictions on the same data.	Raynor	measure of closeness of results of observations, computations, or estimates to the true values or the values accepted as being true	ISO/IEC_TS_5723:2022(en)		
actionable recourse	The ability of a person to change the decision of the model through actionable input variables.	Ustun,_Berk	affected users of the system have the ability to challenge the decision they receive	Varshney,_Kush	Readily available independent mechanisms by which each individuals' complaints and disputes are investigated and expeditiously resolved at no cost to the individual.	Voight,_Paul					recourse, counterfactual explanation, appeal and override	
active learning	A proposed method for modifying machine learning algorithms by allowing them to specify test regions to improve their accuracy. At any point, the algorithm can choose a new point x, observe the output and incorporate the new (x, y) pair into its training base. It has been applied to neural networks, prediction functions, and clustering functions.	Raynor	Active learning (also called "query learning," or sometimes "optimal experimental design" in the statistics literature) is a subfield of machine learning and, more generally, artificial intelligence. The key hypothesis is that, if the learning algorithm is allowed to choose the data from which it learns—to be "curious," if you will—it will perform better with less training.	settles_active_2009	the process of learning through activities and/or discussion in class, as opposed to passively listening to an expert. It emphasizes higher-order thinking and often involves group work.	Freeman_et_al_2014						
active learning agent	[a machine learning algorithm that can] decide what actions to take [with regards to its training data, in contrast to a passive learning agent, which is limited to a fixed policy].	Russell_and_Norvig									passive learning agent	
activity	Work that an organization performs using business processes; can be singular or compound.	IEEE_Guide_IPA	Set of cohesive tasks of a process.	CSRC								
adaptive dynamic programming	An adaptive dynamic programming (or ADP) agent takes advantage of the constraints among the utilities of states by learning the transition model that connects them and solving the corresponding Markov decision process using dynamic programming.	Russell_and_Norvig	A means of learning a model and a reward function from observations that then uses value or policy iteration to obtain the utilities or an optimal policy; makes optimal use of the local constraints on utilities of states imposed through the neighborhood structure of the environment.	Russell_and_Norvig								
adaptive learning	Updating predictive models online during their operation to react to concept drifts	Gama,_Joao										
adversarial action	actions characterised by mala fide (malicious) intent and/or bad faith.	FBPML_Wiki										
adversarial example	Machine learning input sample formed by applying a small but intentionally worst-case perturbation ... to a clean example, such that the perturbed input causes a learned model to output an incorrect answer.	NISTIR_8269_Draft	Samples generated from real samples with carefully designed imperceptible perturbations	Zhang,_Yonggang							adversarial perturbation	
adversarial machine learning	A practice concerned with the design of ML algorithms that can resist security challenges, the study of the capabilities of attackers, and the understanding of attack consequences.	Reznik,_Leon	The field to study vulnerabilities of machine learning approaches in adversarial settings and to develop techniques to make learning robust to adversarial manipulation.	Vorobeychik								
adversary	The agent who conducts or intends to conduct detrimental activities, perhaps by creating an adversarial example.	NISTIR_8269_Draft	Individual, group, organization, or government that conducts or has the intent to conduct detrimental activities.	CSRC								
adverse action notice	A notification of i) a refusal to grant credit in substantially the amount or on substantially the terms requested in an application unless the creditor makes a counteroffer (to grant credit in a different amount or on other terms) and the applicant uses or expressly accepts the credit offered; ii) A termination of an account or an unfavorable change in the terms of an account that does not affect all or substantially all of a class of the creditor's accounts or iii) A refusal to increase the amount of credit available to an applicant who has made an application for an increase.	ECOA										
adverse impact ratio	A substantially different rate of selection [such as in hiring] which works to the disadvantage of members of a race, sex, or ethnic group.	Cadient_EEOC	privileged and unprivileged groups receiving different outcomes irrespective of the decision maker's intent and irrespective of the decision-making procedure. Quantified as the ratio: disparate impact ratio = $P(\hat{y}(X) = \text{fav} \mid\mid Z = \text{unpr}) / P(\hat{y}(X) = \text{fav} \mid\mid Z = \text{priv})$ where $P(\hat{y}(X) = \text{fav})$ is the favorable label, $(Z = \text{priv})$ is the privileged group, and $(Z = \text{unpr})$ is the unprivileged group.	Varshney,_Kush	Determining what constitutes disparate impact at a statistical level is also far from straightforward. Historically, statisticians and regulators have used a variety of methods to detect its occurrence under existing legal standards. Statisticians have, for example, used a group fairness metric called the "80 percent rule" (it's also known as the "adverse impact ratio") as one central indicator of disparate impact. Originating in the employment context in the 1970s, the ratio consists of dividing the proportion of the selected group in the disadvantaged class by the proportion of selected members of the advantaged group. A ratio below 80% is generally considered to be evidence of discrimination. Other metrics, such as standardized mean difference or marginal effects analysis, have been used to detect unfair outcomes in AI as well.	HBR_Andrew_Burt_how_to_ensure					disparate impact ratio, relative risk ratio	
agile	a development approach that delivers software in increments by following the principles of the Manifesto for Agile Software Development.	Gartner	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.	NSCAI								
AI principles	[An overarching concept, value, belief, or norm that guides AI development, testing, and deployment across the AI lifecycle. The OECD] identifies five complementary values-based principles for the responsible stewardship of trustworthy AI and calls on AI actors to promote and implement them: inclusive growth, sustainable development and well-being; human-centred values and fairness; transparency and explainability; robustness, security and safety; and accountability.	OECD_CAI_recommendation									Are these definitions of what an AI principle is or a list of definitions?	
algorithm	A set of step-by-step instructions. Computer algorithms can be simple (if it's 3 p.m., send a reminder) or complex (identify pedestrians).	Hutson,_Matthew	A set of computational rules to be followed to solve a mathematical problem. More recently, the term has been adopted to refer to a process to be followed, often by a computer.	Comptroller_Office	Formulae given to a computer in order for it to complete a task (i.e. a set of rules for a computer).	Reznik,_Leon	precise rules for transforming specified inputs into specified outputs in a finite number of steps	knuth_art_1981	algorithms are step-by-step procedures for solving problems. For concreteness, we can think of them simply as being computed programs, written in some precise computer languages	garey_computers_1979		
algorithmic aversion	biased assessment of an algorithm which manifests in negative behaviours and attitudes towards the algorithm compared to a human agent.	Ekaterina_et_al_2020										
algorithm-in-the-loop	[a] framework [that] centers human decision making, providing a more precise lens for studying the social impacts of algorithmic decision making aids; . . . processes that employ algorithmic aids to enhance human decision making. In contrast to the human-in-the-loop paradigm, which privileges algorithms as the central focus and uses people to improve algorithmic performance, the algorithm-in-the-loop perspective privileges people as the central focus and uses algorithms to improve human decision making. . . . [It] emphasizes developing systems for integration into sociotechnical contexts rather than for isolated decision making. In terms of evaluation, it emphasizes the human's decisions—rather than the algorithm's decisions—as the primary outcome of interest.	Ben_Green_Yiling_Chen										
alignment	ensur[ing] that powerful AI is properly aligned with human values. ... The challenge of alignment has two parts. The first part is *technical* and focuses on how to formally encode values or principles in artificial agents so that they reliably do what they ought to do. ... The second part of the value alignment question is *normative*. It asks what values or principles, if any, we ought to encode in artificial agents.	Gabriel_2020										
amplification	[an act of amplifying, which is] to make larger or greater (as in amount, importance, or intensity).	Merriam-Webster_amplify	This criterion, disparity amplification, deals with the disparity in positive classification rates, which is a widely accepted measure of discriminatory effect in both law and computer science. It stipulates that a disparity in the output of the model is justified by a commensurate disparity in the construct, thereby allowing accurate models even when the base rates are different for different protected groups, as equalized odds, predictive parity, and calibration do.	yeom_avoiding_2021	Let [construct space] Y' and [prediction space] $\hat{Y}$ be categorical. Then, a model exhibits disparity amplification if $d_{tv}(\hat{Y} \mid Z=0, \hat{Y} \mid Z=1) > d_{tv}(Y' \mid Z=0, Y' \mid Z=1)$. dtv is the total variation distance defined as follows. Let $Y0$ and $Y1$ becategorical random variables with finite supports Y0 and Y1. Then,the total variation distance between $Y0$ and $Y1$ is$d_{tv}(Y0,Y1) = 12\Sigma_{y \in Y0 \cup Y1} \quad Pr[Y0=y] - Pr[Y1=y]$ .In the special case where $Y0, Y1 \in \{0, 1\}$, the total variation distancecan also be expressed as $\lvert Pr[Y0=1] - Pr[Y1=1] \rvert$.	yeom_avoiding_2021						
analytics	Analytics is the application of scientific & mathematical methods to the study & analysis of problems involving complex systems. There are three distinct types of analytics: * Descriptive Analytics gives insight into past events, using historical data. * Predictive Analytics provides insight on what will happen in the future. * Prescriptive Analytics helps with decision making by providing actionable advice.	informs_analytics_2022										
annotation	Further documentation accompanying a requirement.	IEEE_Soft_Vocab	[the act of] mak[ing] or furnish[ing] critical or explanatory notes or comment	Merriam-Webster_annotate								
anomaly	Anything observed in the documentation or operation of a system that deviates from expectations based on previously verified system, software, or hardware products or reference documents.	IEEE_Soft_Vocab	Condition that deviates from expectations, based on requirements specifications, design documents, user documents, or standards, or from someone's perceptions or experiences.	SP800-160								
anonymization	The process in which individually identifiable data is altered in such a way that it no longer can be related back to a given individual. Among many techniques, there are three primary ways that data is anonymized. Suppression is the most basic version of anonymization and it simply removes some identifying values from data to reduce its identifiability. Generalization takes specific identifying values and makes them broader, such as changing a specific age (18) to an age range (18-24). Noise addition takes identifying values from a given data set and switches them with identifying values from another individual in that data set. Note that all of these processes will not guarantee that data is no longer identifiable and have to be performed in such a way that does not harm the usability of the data.	IAPP_Privacy_Glossary	process that removes the association between the identifying dataset and the data subject	CSRC								
anthropomorphism	the attribution of distinctively human-like feelings, mental states, and behavioral characteristics to inanimate objects, animals, and in general to natural phenomena and supernatural entities	Anthropomorphism_in_AI_2020	a particular human-like interpretation of existing physical features and behaviors that goes beyond what is directly observable	Anthropomorphism_in_AI_2020								
application	A software program hosted by an information system.	SP800-37	A hardware/software system implemented to satisfy a particular set of requirements.	CSRC	software or a program that is specific tothe solution of an application problem	aime_measurement_2022 citing ISO/IEC TR 24030						
application programming interface (API)	a software contract between the application and client, expressed as a collection of methods or functions. . . it defines the available functions you can execute; . . . the intermediary interface between the client and the application.	Hands-On_Smart_Contract_Dev										

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
artificial general intelligence (AGI)	Algorithms that perform a wide variety of tasks and switch simultaneously from one activity to another in the manner that humans do.	Brookings_Institution	a machine that's as intelligent as a human and equally capable of solving the broad range of problems that require learning and reasoning	Wallace,_Brian	Human-like intelligence, which can be applied widely as opposed to narrow AI, which can only be applied to one particular problem or task. Also called "strong" AI as opposed to "weak" AI.	AI_Ethics_Mark_Coeckelbergh					strong AI	
artificial intelligence (AI)	Interdisciplinary field, usually regarded as a branch of computer science, dealing with models and systems for the performance of functions generally associated with human intelligence, such as reasoning and learning.	Reznik,_Leon	the field concerned with developing techniques to allow computers to act in a manner that seems like an intelligent organism, such as a human would. The aims vary from the weak end, where a program seems "a little smarter" than one would expect, to the strong end, where the attempt is to develop a fully conscious, intelligent, computer-based entity. The lower end is continually disappearing into the general computing background, as the software and hardware evolves.	Raynor	the study of ideas to bring into being machines that respond to stimulation consistent with traditional responses from humans, given the human capacity for contemplation, judgment and intention. Each such machine should engage in critical appraisal and selection of differing opinions within itself. Produced by human skill and labor, these machines should conduct themselves in agreement with life, spirit and sensitivity, though in reality, they are imitations.	Shubendhu_and_Vijay	a field of study that is adept at applying intelligence to vast amounts of data and deriving meaningful results	Wallace,_Brian	The application of computational tools to address tasks traditionally requiring human analysis.	Comptroller_Office	machine learning; data science	
artificial intelligence learning	The ingestion of a corpus, application of semantic mapping, and relevant ontology of structured and/or unstructured data that yields inference and correlation leading to the creation of useful conclusive or predictive capabilities in a given knowledge domain. Strong AI learning also includes the capability of creating unique hypotheses, attributing data relevance, processing data relationships, and updating its own lines of inquiry to further the usefulness of its purpose.	IEEE_Guide_IPA										
artificial narrow intelligence (ANI)	[an AI system that] is designed to accomplish a specific problem-solving or reasoning task.	OECD_Artificial_Intelligence_in_Society	Artificial Narrow Intelligence, also known as weak or applied intelligence, represents most of the current artificial intelligent systems which usually focus on a specific task. Narrow AIs are mostly much better than humans at the task they were made for: for example, look at face recognition, chess computers, calculus, and translation. The definition of artificial narrow intelligence is in contrast to that of strong AI or artificial general intelligence, which aims at providing a system with consciousness or the ability to solve any problems. Virtual assistants and AlphaGo are examples of artificial narrow intelligence systems.	AI_in_Medical_Imaging_glossary							weak intelligence; applied intelligence	
artificial neural networks	A computing system, made up of a number of simple, highly interconnected processing elements, which processes information by its dynamic state response to external inputs.	Reznik,_Leon	A good definition of ANN, is given by Haykin [1] describing ANN as a massively parallel combination of simple processing unit which can acquire knowledge from environment through a learning process and store the knowledge in its connections.	guresen_definition_2011	Definition 1. A directed graph is called an Artificial Neural Network (ANN) if it has x at least one start node (or Start Element; SE), x at least one end node (or End Element; EE), x at least one Processing Element (PE), x all the nodes used must be Processing Elements (PEs), except start nodes and end nodes, x a state variable ni associated with each node i, x a real valued weight wki associated with each link (ki) from node k to node i, x a real valued bias bi associated with each node i, x at least two of the multiple PEs connected in parallel, x a learning algorithm that helps to model the desired output for given input. x a flow on each link (ki) from node k to node i, that carries exactly the same flow which equals to nk caused by the output of node k , x each start node is connected to at least one end node, and each end node is connected to at least one start node, x no parallel edges (each link (ki) from node k to node i is unique).							
assessment	Action of applying specific documented criteria to a specific software module, package or product for the purpose of determining acceptance or release of the software module, package or product.	IEEE_Soft_Vocab	the action or an instance of making a judgment about something : the act of assessing something : APPRAISAL	Merriam-Webster_assessment								
asset	Item, thing, or entity that has potential or actual value to an organization. Item that has been designed for use in multiple contexts.	IEEE_Soft_Vocab										
attack	Action targeting a learning system to cause malfunction.	NISTIR_8269_Draft	Any kind of malicious activity that attempts to collect, disrupt, deny, degrade, or destroy information system resources or the information itself.	CSRC								
attribute	Property associated with a a set of real or abstract things that is some characteristic of interest.	IEEE_Soft_Vocab	A quantity describing an instance. An attribute has a domain defined by the attribute type, which denotes the values that can be taken by an attribute.	Kohavi,_Ron	property or characteristic of an object that can be distinguished quantitatively or qualitatively by human or automated means	aime_measurement_2022, citing ISO/IEC TR 24029-1						
audit	Systematic, independent, documented process for obtaining records, statements of fact, or other relevant information and assessing them objectively, to determine the extent to which specified requirements are fulfilled.	IEEE_Soft_Vocab	To conduct an independent review and examination of system records and activities in order to test the adequacy and effectiveness of data security and data integrity procedures, to ensure compliance with established policy and operational procedures, and to recommend any necessary changes.	FDA_Glossary	Independent examination of a software product, software process, or set of software processes to assess compliance with specifications, standards, contractual agreements, or other criteria	NASA_Soft_Standards	Independent review conducted to compare the various aspects of the laboratory's performance with a standard for that performance. Also defined as a systematic, independent and documented process for obtaining audit evidence and evaluating it objectively to determine the extent to which audit criteria are fulfilled.	UNODC_Glossary_QA_GLP				
audit log	A chronological record of system activities, including records of system accesses and operations performed in a given period.	SP800-37										
authenticity	property that an entity is what it claims to be	ISO/IEC_TS_5723:2022(en)										
automation	Independent machine-managed choreography of the operation of one or more digital systems.	IEEE_Guide_IPA	conversion of processes or equipment to automatic operation, or the results of the conversion	IEEE_Soft_Vocab	The system functions with no/little human operator involvement; however, the system performance is limited to the specific actions it has been designed to do. Typically these are well-defined tasks that have predetermined responses (i.e., simple rule-based responses).	DOD_TEVV						
automation bias	over-relying on the outputs of AI systems	David_Leslie_Morgan_Briggs	It refers to a well-documented human propensity to automatically defer to automated systems, despite warning signals or contradictory information from other sources. In other words, human actors are found to uncritically abdicate their decision making to automation.	alon-barkat_human_2023								
autonomic	A monitor-analyze-plan-execute (MAPE) computer system capable of sensing environments, interpreting policy, accessing knowledge (data --- information --- knowledge), making decisions, and initiating dynamically assembled routines of choreographed activity to both complete a process and update the set of environmental variables that enables the autonomic system to self-manage its own operation and the processes it oversees. An autonomic system is identified by eight characteristics: a) Knows the resources to which it has access, what its capabilities and limitations are, and how and why it is connected to other systems. b) Is able to configure and reconfigure itself depending on the changing computing environment. c) Is able to optimize its performance to ensure the most efficient computing process. d) Is able to work around encountered problems either by repairing itself or routing functions away from the trouble. e) Is able to detect, identify, and protect itself against various types of attacks to maintain overall system security and integrity. f) Is able to adapt to its environment as it changes by interacting with neighboring systems and establishing communication protocols. g) Relies on open standards and requires access to proprietary environments to achieve full performance. h) Is able to anticipate the demand on its resources transparently to users.	IEEE_Guide_IPA										
autonomous vehicle	[an] automobile, bus, tractor, combine, boat, forklift, etc. . . . capable of sensing its environment and moving safely with little or no human input.	Introduction_to_Information_Systems										
autonomy	The system has a set of intelligence-based capabilities that allows it to respond to situations that were not pre-programmed or anticipated (i.e., decision-based responses) prior to system deployment. Autonomous systems have a degree of self-government and self-directed behavior (with the human's proxy for decisions).	DOD_TEVV	1. a state of independence and self-determination in an individual, a group, or a society. According to some theories, an inordinate focus on self-determination and achievement represents a risk factor for the development of major depressive disorder. 2. in self-determination theory more specifically, the experience of acting from choice, rather than feeling pressured to act. This form of autonomy is considered a fundamental psychological need that predicts well-being.	APA_autonomy								
availability	Ensuring timely and reliable access to and use of information.	SP800-37	The property that data or information is accessible and usable upon demand by an authorized person.	NIST_SP_800	property of being accessible and usable on demand by an authorized entity	ISO/IEC_TS_5723:2022(en)						
back test	the quantitative evaluation of a model's performance both from a statistical and trading perspective	The_Science_of_Algorithmic_Trading_and_Portfolio_Management										
backpropagation	The way many neural nets learn. They find the difference between their output and the desired output, then adjust the calculations in reverse order of execution.	Hutson,_Matthew	A classical method for error propagation when training Artificial Neural Networks (ANNs). For standard backpropagation, the parameters of each node are changed accoring to the local error gradient. The method can be very slow to converge although it can be improved through the use of methods that slow the error propagation and by batch processing. Many alternate methods such as the conjugate gradient and Levenberg-Marquardt algorithms are more effective and reliable.	Raynor								
bad actor	individuals or entities who are responsible for cyber incidents against enterprises, governments, and users.	Mark_Ciampa_2021	someone with objectives of studying and using cyber security techniques and tools for personal or private gain through malicious or threat activity.	Thomas_Edgar							black hat, threat actor	
bagging	Bagging predictors is a method for generating multiple versions of a predictor and using these to get an aggregated predictor.	Breiman_Leo	In this approach we generate B different bootstrapped training data sets. We then train our method onthe bth bootstrapped training set in order to get ˆ f*b (x), and finally averageall the predictions, to obtainˆ fbag(x) =1B Bb=1ˆ f*b(x). This is called bagging.	hastie_introduction_2014								
back-testing	A form of outcomes analysis that involves the comparison of actual outcomes with modeled forecasts during a development sample time period (in-sample back-testing) and during a sample period not used in model development (out-of-time back-testing), and at an observation frequency that matches the forecast horizon or performance window of the model.	Comptroller_Office										
batched automation	Process automation execution of intentionally segregated work processes that are able to be processed irrespective of their contextual placement within a service.	IEEE_Guide_IPA										
benchmark	Standard against which results can be measured or assessed; Procedure, problem, or test that can be used to compare systems or components to each other or to a standard.	IEEE_Soft_Vocab	An alternative prediction or approach used to compare a model's inputs and outputs to estimates from alternative internal or external data or models.	Comptroller_Office	The term benchmarking is used in machine learning (ML) to refer to the evaluationand comparison of ML methods regarding their ability to learn patterns in 'benchmark'datasets that have been applied as 'standards'. Benchmarking could be thought of simplyas a sanity check to confirm that a new method successfully runs as expected and canreliably find simple patterns that existing methods are known to identify.	olson_pmlb_2017						

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
bias	A systematic error. In the context of fairness, we are concerned with unwanted bias that places privileged groups at systematic advantage and unprivileged groups at systematic disadvantage.	AI_Fairness_360	(computational bias) An effect which deprives a statistical result of representativeness by systematically distorting it, as distinct from a random error which may distort on any one occasion but balances out on the average.	OECD	(legal/ethics) Discrimination against or in favor of particular individuals or groups. In the context of ethics and politics, the question arises whether a particular bias is unjust or unfair.	AI_Ethics_Mark_Coeckelbergh	(systemic bias) systematic difference in treatment of certain objects, people or groups in comparison to others	measurement_iso22989_2022	(mathematical) A point estimator \theta_hat is said to be an unbiased estimator fo \theta if E(\theta_hat) = \theta for every possible value of \theta. If \theta_hat is not unbiased, the difference E(\theta_hat) - \theta is called the bias of \theta	devore_probability_2004		
bias mitigation algorithm	A procedure for reducing unwanted bias in training data or models.	AI_Fairness_360										
bias testing	As it relates to disparate impact, courts and regulators have utilized or considered as acceptable various statistical tests to evaluate evidence of disparate impact. Traditional methods of statistical bias testing look at differences in predictions across protected classes, such as race or sex. In particular, courts have looked to statistical significance testing to assess whether the challenged practice likely caused the disparity and was not the result of chance or a nondiscriminatory factor.	SP1270										
big data	Extremely large data sets that are statistically analyzed to gain detailed insights. The data can involve billions of records and require substantial computer-processing power. Datasets are sometimes linked together to see how patterns in one domain affect other areas. Data can be structured into fixed fields or unstructured as free-flowing information. The analysis of big datasets, often using AI, can reveal patterns, trends, or underlying relationships that were not previously apparent to researchers.	Brookings_Institution	consists of extensive datasets primarily in the characteristics of volume, variety, velocity, and/or variability□ that require a scalable architecture for efficient storage, manipulation, and analysis	NIST_1500								
binning	a technique of lumping small ranges of values together into categories, or "bins," for the purpose of reducing the variability (removing some of the fine structure) in a data set.	Pyle,_Dorian._Data_Preparation_as_a_Process										
biometric data	personal data resulting from specific technical processing relating to the physical, physiological or behavioural characteristics of a natural person, which allow or confirm the unique identification of that natural person, such as facial images or dactyloscopic data;	GDPR	an individual's physiological, biological, or behavioral characteristics, including information pertaining to an individual's deoxyribonucleic acid (DNA), that is used or is intended to be used singly or in combination with each other or with other identifying data, to establish individual identity. Biometric information includes, but is not limited to, imagery of the iris, retina, fingerprint, face, hand, palm, vein patterns, and voice recordings, from which an identifier template, such as a faceprint, a minutiae template, or a voiceprint, can be extracted, and keystroke patterns or rhythms, gait patterns or rhythms, and sleep, health, or exercise data that contain identifying information.	CCPA	A measurable physical characteristic or personal behavioral trait used to recognize the identity, or verify the claimed identity, of an applicant. Facial images, fingerprints, and iris scan samples are all examples of biometrics.	SP800-12				personal data; processing	
boosting	Boosting works by sequentially applying a classificationalgorithm to reweighted versions of the training data and then takinga weighted majority vote of the sequence of classifiers thus produced.	friedman_additive_2000	A machine learning technique that iteratively combines a set of simple and not very accurate classifiers (referred to as "weak" classifiers) into a classifier with high accuracy (a "strong" classifier) by upweighting the examples that the model is currently misclassifying	aime_measurement_2022, citing Machine Learning Glossary by Google								
breach	The loss of control, compromise, unauthorized disclosure, unauthorized acquisition, or any similar occurrence where: a person other than an authorized user accesses or potentially accesses personally identifiable information; or an authorized user accesses personally identifiable information for another than authorized purpose.	CSRC										
broad artificial intelligence (broad AI)	Complex, computational, cognitive automation system capable of providing descriptive, predictive, prescriptive, and limited deductive analytics with relevance and accuracy exceeding human expertise in a broad, logically related set of knowledge domains.	IEEE_Guide_1PA	a sophisticated and adaptive system, which successfully performs any cognitive task by virtue of its sensory perception, previous experience, and learned skills.	Hochreiter,_Sepp								
built-in test	Equipment or software embedded in the operational components or systems, as opposed to external support units, which perform a test or sequence of tests to verify mechanical or electrical continuity of hardware, or the proper automatic sequencing, data processing, and readout of hardware or software systems.	SP1011										
bug-bounty	Reward given to independent security researchers, penetrations testers, and white hat hackers for discovering exploitable software vulnerabilities and sharing this knowledge with the operator of a particular bug-bounty program (BBP).	Kuehn,_Andreas										
business process	A defined set of business activities that represent the steps or tasks required to achieve a business objective, including the flow and use of information, participants, and human or digital resources.	IEEE_Guide_1PA										
business process management	Discipline involving any combination of modeling, automation, execution, control, measurement and optimization of business activity flows, in support of enterprise goals, spanning systems, employees, customers, and partners within and beyond the enterprise boundaries.	IEEE_Guide_1PA										
business rule	Definition, constraint, dependency, or decision criteria that determine the method of execution of a task or tasks, or influences the order of execution of a task or tasks. Business rules assert control, or influence the behavior, of a business process within computing systems.	IEEE_Guide_1PA										
calibration	A comparison between a device under test and an established standard, such as UTC(NIST). When the calibration is finished, it should be possible to state the estimated time offset and/or frequency offset of the device under test with respect to the standard, as well as the measurement uncertainty.	CSRC	operation that, under specified conditions, in a first step, establishes a relation between the quantity values with measurement uncertainties provided by measurement standards and corresponding indications with associated measurement uncertainties and, in a second step, uses this information to establish a relation for obtaining a measurement result from an indication	aime_measurement_2022, citing ISO/IEC Guide 99	Set of operations that establish, under specified conditions, the relationship between values indicated by a measuring instrument or measuring system, or values represented by a material measure, and the corresponding known values of a measurand.	UNODC_Glossary_QA_GLP						
capability	measure of capacity and the ability of an entity, person or organization to achieve its objectives	ISO/IEC_TS_5723:2022(en)										
case	Single entry, single exit multiple way branch that defines a control expression, specifies the processing to be performed for each value of the control expression, and returns control in all instances to the statement immediately following the overall construct.	IEEE_Soft_Vocab										
causal inference	an intellectual discipline that considers the assumptions, study designs, and estimation strategies that allow researchers to draw causal conclusions based on data. The term 'causal conclusion' used here refers to a conclusion regarding the effect of a causal variable (often referred to as the 'treatment' under a broad conception of the word) on some outcome(s) of interest.	Jennifer._Hill										
causative	acting as the cause of something	cambridge_causative_2023										
chatbot	Conversational agent that dialogues with its user (for example: empathic robots available to patients, or automated conversation services in customer relations).	COE_AI_Glossary	A chatbot is a computer program which responds like an intelligent entity when conversed with. The conversation may be through text or voice. Any chatbot program understands one or more human languages by Natural Language Processing	Khanna,_Aniruth								
choreography	An ordered sequence of system-to-system message exchanges between two or more participants. In choreography, there is no central controller, responsible entity, or observer of the process.	IEEE_Guide_1PA										
classification	When the output is one of a finite set of values (such as sunny, cloudy or rainy), the learning problem is called classification, and is called Boolean or binary classification if there are only two values.	AIMA	task of assigning collected data to target categories or classes.	aime_measurement_2022, citing ISO/IEC TR 24030								
classifier	A model that predicts categorical labels from features.	AI_Fairness_360										
clustering	Detecting potentially useful clusters of input examples.	AIMA	The basic problem of clustering may be stated as follows: Given a set of data points, partition them into a set of groups which are as similar as possible.	aggarwal_clustering_2013	the tendency for items to be consistently grouped together in the course of recall. This grouping typically occurs for related items. It is readily apparent in memory tasks in which items from the same category, such as nonhuman animals, are recalled together.	APA_clustering						
cognitive automation	The identification, assessment, and application of available machine learning algorithms for the purpose of leveraging domain knowledge and reasoning to further automate the machine learning already present in a manner that may be thought of as cognitive. With cognitive automation, the system performs corrective actions driven by knowledge of the underlying analytics tool itself, iterates its own automation approaches and algorithms for more expansive or more thorough analysis, and is thereby able to fulfill its purpose. The automation of the cognitive process refines itself and dynamically generates novel hypotheses that it can likewise assess against its existing corpus and other information resources.	IEEE_Guide_1PA										
cognitive computing	Complex computational systems designed to — Sense (perceive the world and collect data); — Comprehend (analyze and understand the information collected); - Act (make informed decisions and provide guidance based on this analysis in an independent way); and — Adapt (adapt capabilities based on experience) in ways comparable to the human brain.	IEEE_Guide_1PA										
column	In the context of relational databases, a column is a set of data values, all of a single type, in a table.	techopedia_column_2022										
COMPAS controversy	A canonical example [of algorithmic bias] comes from a tool used by courts in the United States to make pretrial detention and release decisions. The software, Correctional Offender Management Profiling for Alternative Sanctions (COMPAS), measures the risk of a person to recommit another crime. Judges use COMPAS to decide whether to release an offender or to keep him or her in prison. An investigation into the software found a bias against African-Americans: COMPAS is more likely to have higher false positive rates for African-American offenders than Caucasian offenders in falsely predicting them to be at a higher risk of recommitting a crime or recidivism	Mehrabi,_Ninareh										
computer vision	The digital process of perceiving and learning visual tasks in order to interpret and understand the world through cameras and sensors.	NSCAI	An image understanding task that automatically builds a description not only of the image itself, but of the three dimensional scene that it depicts.	NISIR_82-2582								

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
concept drift	Use of a system outside the planned domain of application, and a common cause of performance gaps between laboratory settings and the real world.	SP1270	an online supervised learning scenario when the relation between the input data and the target variable changes over time.	Gama,_Joao	Systems that classify or predict a concept (e.g., credit ratings or computer intrusion monitors) over time can suffer performance loss when the concept they are tracking changes. This is referred to as concept drift. This can either be a natural process that occurs without a reference to the system, or an active process, where others are reacting to the system (e.g., virus detection).	Raynor						
confidentiality	Data confidentiality is a property of data, usually resulting from legislative measures, which prevents it from unauthorized disclosure.	OECD	Preserving authorized restrictions on information access and disclosure, including means for protecting personal privacy and proprietary information.	CSRC	The property that data or information is not made available or disclosed to unauthorized persons or processes.	NIST_SP_800	A property that information is not disclosed to users, processes, or devices unless they have been authorized to access the information.	CISA				
confusion matrix	A matrix showing the predicted and actual classifications. A confusion matrix is of size LxL, where L is the number of different label values	Kohavi,_Ron										
consent	'Consent' of the data subject means any freely given, specific, informed and unambiguous indication of the data subject's wishes by which he or she, by a statement or by a clear affirmative action, signifies agreement to the processing of personal data relating to him or her.	GDPR	"Consent" means any freely given, specific, informed, and unambiguous indication of the consumer's wishes by which the consumer, or the consumer's legal guardian, a person who has power of attorney, or a person acting as a conservator for the consumer, including by a statement or by a clear affirmative action, signifies agreement to the processing of personal information relating to the consumer for a narrowly defined particular purpose. Acceptance of a general or broad terms of use, or similar document, that contains descriptions of personal information processing along with other, unrelated information, does not constitute consent. Hovering over, muting, pausing, or closing a given piece of content does not constitute consent. Likewise, agreement obtained through use of dark patterns does not constitute consent.	CCPA						personal data	
constituent system	independent system that forms part of a system of systems (SoS) (note: Constituent systems can be part of one or more SoS. Each constituent system is a useful system by itself, having its own development, management, utilization, goals, and resources, but interacts within the SoS to provide the unique capability of the SoS).	ISO/IEC_TS_5723:2022(en)										
constraint	Specification of what may be contained in a data or metadata set in terms of the content or, for data only, in terms of the set of key combinations to which specific attributes (defined by the data structure) may be attached.	OECD	A limitation or implied requirement that constrains the design solution or implementation of the systems engineering process and is not changeable by the enterprise	IEEE_Soft_Vocab								
construct validity	the degree to which the application of constructs to phenomena is warranted with respect to the research goals and questions.	Wieringa,_Roel_J.	Construct validation is involved whenever a test is to be interpreted as a measure of some attribute or quality which is not "operationally defined." The problem faced by the investigator is, "What constructs account for variance in test performance?"	cronbach_construct_1955	Established experimentally to demonstrate that a survey distinguishes between people who do and do not have certain characteristics. It is usually established experimentally.	fink_survey_2010	Establishing construct validity means demonstrating, in a variety of ways, that the measurements obtained from measurement model are both meaningful and useful.	jacobs_measurement_2023				
content harms	the psychological, social, physical, or other harms experienced by someone while they are interacting with content that is algorithmically recommended to them.	Chi,_Gao,_Ma									harms of representation	
content validity	Refers to the extent to which a measure thoroughly and appropriately assesses the skills or characteristics it is intended to measure.	fink_survey_2010	the extent to which a test measures a representative sample of the subject matter or behavior under investigation. For example, if a test is designed to survey arithmetic skills at a third-grade level, content validity indicates how well it represents the range of arithmetic operations possible at that level. Modern approaches to determining content validity involve the use of exploratory factor analysis and other multivariate statistical procedures.	APA_content_validity								
contestability	A contestable statement, claim, legal decision, etc. is one that is possible to argue about or try to have changed because it may be wrong	cambridge_contestable_2023										
context	The context is the circumstances, purpose, and perspective under which an object is defined or used.	OECD	The immediate environment in which a function (or set of functions in a diagram) operates	IEEE_Soft_Vocab	the interrelated conditions in which something exists or occurs.	Merriam-Webster_context						
context control												
contextual learning	A computing system with sufficient knowledge regarding its purpose that it understands the source, relevance, and utility of data and inputs.	IEEE_Guide_1PA										
context-of-use	The Context of Use is the actual conditions under which a given artifact/software product is used, or will be used in a normal day to day working situation.	interaction_context_2023	comprises a combination of users, goals, tasks, resources, and the technical, physical and social, cultural and organizational environments in which a system, product or service is used[...] can include the interactions and interdependencies between the object of interest and other systems, products or services.	ISO_9241-11:2018								
controllability	property of a system that allows a human or another external agent to intervene in the system's functioning; such a system is heteronomous.	ISO/IEC_TS_5723:2022(en)										
control class	(control group) the set of observations in an experiment or prospective study that do not receive the experimental treatment(s). These observations serve (a) as a comparison point to evaluate the magnitude and significance of each experimental treatment, (b) as a reality check to compare the current observations with previous observation history, and (c) as a source of data for establishing the natural experimental error.	nist_statistics_2012										
controller	'Controller' means the natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data; where the purposes and means of such processing are determined by Union or Member State law, the controller or the specific criteria for its nomination may be provided for by Union or Member State law;	GDPR									personal data; processor	
corpus (corpora)	A deliberately assembled collection of knowledge and data (structured and/or unstructured) believed to contain relevant information on a topic or topics to be used by software systems for which useful analysis, prediction, or outcome is being sought.	IEEE_Guide_1PA										
correlation	In its most general sense correlation denoted the interdependence between quantitative or qualitative data. In this sense it would include the association of dichotomised attributes and the contingency of multiply-classified attributes.	OECD	The correlation coefficient of two random variables $_1, and $_2, denoted \rho ($_1,$_2) is: \rho($_1, $_2) = Cov($_1, $_2)/\sqrt{Var($_1)*Var($_2)}	box_statistics_2005								
counterfactual explanation	Statements taking the form: Score p was returned because variables V had values (v1, v2,...) associated with them. If V instead had values (v1', v2',...) score p' would have been returned.	wachter_counterfactual_2018										
counterfactual fairness	Our definition of counterfactual fairness captures the intuition that a decision is fair towards an individual if it the same in (a) the actual world and (b) a counterfactual world where the individual belonged to a different demographic group.	kusner_counterfactual_2017	Given a predictive problem with fairness considerations, where A, X and Y represent the protectedattributes, remaining attributes, and output of interest respectively, let us assume that we are given acausal model (U, V, F), where V = A \cup X. We postulate the following criterion for predictors of Y .Definition 5 (Counterfactual fairness). Predictor ^Y is counterfactually fair if under any context X = x and A = a, P(^Y_{A <- a}(U) = y	X = x, A = a) = P(^Y_{A <- a'}(U) = y | X = x,A = a); (i) for all y and for any value a' attainable by A.	kusner_counterfactual_2017	A fairnessmetric that checks whether a classifier produces the same result for one individualas it does for another individual who is identical to the first, except withrespect to one or more sensitive attributes. Evaluating a classifier for counterfactualfairness is one method for surfacing potential sources of bias in a model	aime_measurement_2022; citing Machine Learning Glosssary by Google					
countermeasure	Actions, devices, procedures, techniques, or other measures that reduce the vulnerability of a system. Synonymous with security controls and safeguards.	SP800-37	Actions, devices, procedures, or techniques that meet or oppose (i.e., counters) a threat, a vulnerability, or an attack by eliminating or preventing it, by minimizing the harm it can cause, or by discovering and reporting it so that corrective action can be taken.	GWUC							safeguard; security control	
criterion validity	compares responses to future performance or to those obtained from other, more well-established surveys. Criterion validity is made up two subcategories: predictive and concurrent. Predictive validity refers to the extent to which a survey measure forecasts future performance. A graduate school entry examination that predicts who will do well in graduate school has predictive validity. Concurrent validity is demonstrated when two assessments agree or a new measure is compared favorably with one that is already considered valid.	fink_survey_2010	an index of how well a test correlates with an established standard of comparison (i.e., a criterion). Criterion validity is divided into three types: predictive validity, concurrent validity, and retrospective validity. For example, if a measure of criminal behavior is valid, then it should be possible to use it to predict whether an individual (a) will be arrested in the future for a criminal violation, (b) is currently breaking the law, and (c) has a previous criminal record.	APA_criterion_validity							criterion-referenced validity; criterion-related validity	
crowdsource	a type of participative online activity in which an individual, an institution, a non-profit organization, or company proposes to a group of individuals of varying knowledge, heterogeneity, and number, via a flexible open call, the voluntary undertaking of a task. The undertaking of the task, of variable complexity and modularity, and in which the crowd should participate bringing their work, money, knowledge and/or experience, always entails mutual benefit. The user will receive the satisfaction of a given type of need, be it economic, social recognition, self-esteem, or the development of individual skills, while the crowdsourcer will obtain and utilize to their advantage what the user has brought to the venture, whose form will depend on the type of activity undertaken.	Enrique										
customer	The beneficiary of the execution of an automated task, process, or service.	IEEE_Guide_1PA										
cybersecurity	Prevention of damage to, protection of, and restoration of computers, electronic communications systems, electronic communications services, wire communication, and electronic communication, including information contained therein, to ensure its availability, integrity, authentication, confidentiality, and nonrepudiation.	SP800-37										
dark pattern	"Dark pattern" means a user interface designed or manipulated with the substantial effect of subverting or impairing user autonomy, decisionmaking, or choice, as further defined by regulation.	CCPA										
data	Characteristics or information, usually numerical, that are collected through observation.	OECD	re-interpretable representation of information ina formalized manner suitable for communication, interpretation or processing	aime_measurement_2022, citing ISO/IEC TR 24029-1								
data analytics	The analysis of data to gather substantive insights. Researchers use statistical techniques to find trends or patterns in the data, which give them a better understanding of a range of different topics. Data analytic approaches are used in many businesses and organizations to track day-to-day activities and improve operational efficiency.	Brookings_Institution	Data analysis is the process of transforming raw data into usable information, often presented in the form of a published analytical article, in order to add value to the statistical output.	OECD	the process of applying graphical, statistical, or quantitative techniques to a set of observations or measurements in order to summarize it or to find general patterns.	APA_data_analysis						
data cleaning	Data Cleaning is the process of identifying, correcting, or removing inaccurate or corrupt data records	Ranschaert,_Erik										
data control	management oversight of information policies for an organization's information; observing and reporting on how processes are working and managing issues.	Egnyte										

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
data dredging	A statistical bias in which testing huge numbers of hypotheses of a dataset may appear to yield statistical significance even when the results are statistically nonsignificant.	SP1270									statistical bias; p-hacking	
data drift	The change in model input data that leads to model performance degradation.	Microsoft_Azure_documentation										
data-driven	Data-driven decision making (DDD) refers to the practice of basing decisions on the analysis of data rather than purely on intuition.	provost_data_2013										
data fabric	A data corpus, after the application of semantic mapping, relevant ontologies, and data seeding sufficient for artificial intelligence (AI) or machine learning algorithms to provide meaningful insight, prediction, and/or prescription.	IEEE_Guide_1PA										
data fusion	A process in which data, generated by multiple sensory sources, is integrated and/or correlated to create information, knowledge, and/or intelligence that may be displayed for user or be actionable to accomplish the tasks.	SP1011	The process of combining data from multiple sources to produce more accurate, consistent, and concise information than that provided by any individual data source.	Munir,_Arslan								
data governance	A set of processes that ensures that data assets are formally managed throughout the enterprise. A data governance model establishes authority and management and decision making parameters related to the data produced or managed by the enterprise.	CSRC	refers to a system, including policies, people, practices, and technologies, necessary to ensure data management within an organization	NIST_1500								
data mining	Techniques that analyze large amounts of information to gain insights, spot trends, or uncover substantive patterns. These approaches are used to help businesses and organizations improve their processes or identify associations that shed light on relevant questions. Data mining often involves more use of algorithms, especially machine learning, than traditional statistics.	Brookings_Institution	Data Mining is the process of data analysis and information extraction from large amounts of datasets with machine learning, statistical approaches. and many others.	Ranschaert,_Erik	computational process that extracts patternsby analysing quantitative data from different perspectives and dimensions, categorizingthem, and summarizing potential relationships and impacts	aime_measurement_2022 citinig ISO/IEC 22989						
data point	a discrete unit of information.	TechTarget_data_point	the information we feed into the machine learning model.	Morris_John_data_point								
data poisoning	Machine learning systems trained on user-provided data are susceptible to data poisoning attacks, whereby malicious users inject false training data with the aim ofcorrupting the learned model	Steinhardt,_Jacob										
data preparation	We define data preparation as the set of preprocessing operations performed in early stages of a data processing pipeline, i.e., data transformations at the structural and syntactical levels	hameed_data_2020										
data proxy	Data that are closely related to and serve in place of data that are either unobservable or immeasurable.	Comptroller_Office										
data quality	degree to which the characteristics of data satisfy stated and implied needs when used under specified conditions	IEEE_Soft_Vocab	The dimensions of the IMF definition of "data quality" are: - integrity; - methodological soundness; - accuracy and reliability; - serviceability; - accessibility. There are a number of prerequisites for quality. These comprise: - legal and institutional environment; - resources; - quality awareness.	OECD								
data science	The field that combines domain expertise, programming skills, and knowledge of mathematics and statistics to extract meaningful insights from data.	Reznik,_Leon	Data-Driven Science, or Data Science, is an interdisciplinary field of employing computing algorithms to extract knowledge or insights from data acquired from different sources.	Ranschaert,_Erik	Methodology for the synthesis of useful knowledge directly from data through a process of discovery or of hypothesis formulation and hypothesis testing	NIST_1500	Interdisciplinary science that uses statistics, algorithms, and other methods to extract meaningful and useful patterns from data sets—sometimes known as "big data." Today, *machine learning* is often used in this field. Next to analysis of data, data science is also concerned with the capturing, preparation, and interpretation of data.	AI_Ethics_Mark_Coeckelbergh			artificial intelligence (AI); machine learning (ML)	
data scientist	A practitioner who has sufficient knowledge in the overlapping regimes of business needs, domain knowledge, analytical skills, and software and systems engineering to manage the end-to-end data processes in the analytics life cycle.	NIST_1500										
data seeding	The intentional introduction of initial state conditions, influencing factors, and outcomes (both successful and unsuccessful) in a data fabric to create sufficient machine learning analysis signals to enable encouragement/discouragement to enrich deterministic relationships between data elements in a given information domain.	IEEE_Guide_1PA										
data wrangling	process by which the data required by an application is identified, extracted, cleaned and integrated, to yield a data set that is suitable for exploration and analysis.	Furche,_Tim										
decision	A conclusion reached after consideration of business rules and relevant data within a given process.	IEEE_Guide_1PA	Types of statements in which a choice between two or more possible outcomes controls which set of actions will result.	IEEE_Soft_Vocab								
decision point	A point within a business process where the process flow can take one of several alternative paths, including recursive.	IEEE_Guide_1PA										
decision subject												
decision tree	Tree-structure resembling a flowchart, where every node represents a test to an attribute, each branch represents the possible outcomes of that test, and the leaves represent the class labels.	Reznik,_Leon	In this chapter, we describe tree-based methods for regression andclassification. These involve stratifying or segmenting the predictor spaceinto a number of simple regions. In order to make a prediction for a givenobservation, we typically use the mean or the mode of the training observationsin the region to which it belongs. Since the set of splitting rules usedto segment the predictor space can be summarized in a tree, these types ofapproaches are known as decision tree methods.	james_statistical_2014								
decision-making	the cognitive process resulting in the selection of a belief or a course of action among several possible alternative options. It could be either rational or irrational. The decision-making process is a reasoning process based on assumptions of values, preferences and beliefs of the decision-maker. Every decision-making process produces a final choice, which may or may not prompt action.	Wikipedia_Decision-making	the cognitive process of choosing between two or more alternatives, ranging from the relatively clear cut (e.g., ordering a meal at a restaurant) to the complex (e.g., selecting a mate). Psychologists have adopted two converging strategies to understand decision making: (a) statistical analysis of multiple decisions involving complex tasks and (b) experimental manipulation of simple decisions, looking at the elements that recur within these decisions.	APA_decision_making	Decision making requires that the decision maker make a choice between two or more alternatives (note that doing nothing can be viewed as making a choice). The alternative selected results in real or imaginary consequences to the decision maker. Judgment is a closely related process in which a person rates or assigns values to attributes of the alternatives considered. A rational decision maker seeks desirable consequences and attempts to avoid undesirable consequences.	Lehto_Nanda_2021	a choice of action – of what to do or not do. Decisions are made to achieve goals, and they are based on beliefs about what actions will achieve the goals.	Baron_Thinking_and_Deciding				
decision support system	a computer program application used to improve a company's decision-making capabilities. It analyzes large amounts of data and presents an organization with the best possible options available[; they] bring together data and knowledge from different areas and sources to provide users with information beyond the usual reports and summaries. This is intended to help people make informed decisions.	TechTarget_decision_support_system	a model- or knowledge-based system intended to support managerial decision making in semistructured or unstructured situations. A DSS is not meant to replace a decision maker, but to extend his/her decision making capabilities. It uses data, provides a clear user interface, and can incorporate the decision maker's own insights.	Burstein_Holsapple								
decommission	the total or partial removal of existing components and their corresponding sub-components from Production and any relevant environment, minimising risks and impacts, ensuring policy compliance, and maximizing the financial benefits (i.e., optimizing the cost reduction).	IG1390M_AIOps_Decommission_v1.0.0										
deductive analytics	Insights, reporting, and information answering the question, "What would likely happen IF...?" Deductive analytics evaluates causes and outcomes of possible future events.	IEEE_Guide_1PA									deductive reasoning	
deep learning	A subset of machine learning that relies on neural networks with many layers of neurons. In so doing, deep learning employs statistics to spot underlying trends or data patterns and applies that knowledge to other layers of analysis. Some have labeled this as a way to "learn by example" and a technique that "perform[s] classification tasks directly from images, text, or sound" and then applies that knowledge independently.[4] Deep learning requires extensive computing power and labeled data, and is used in medical research, automated vehicles, electronics, and manufacturing, among other areas.	Brookings_Institution	Deep learning is a broad family of techniques for machine learning in which hypotheses take the form of complex algebraic circuits with tunable connection strengths. The word "deep" refers to the fact that the circuits are typically organized into many layers, which means that computation paths from inputs to outputs have many steps. Deep learning is currently the most widely used approach for applications such as visual object recognition, machine translation, speech recognition, speech synthesis, and image synthesis; it also plays a significant role in reinforcement learning applications.	Russell_and_Norvig	Machine learning method based on characterization of data learning.	Reznik,_Leon	A form of *machine learning* that uses neural networks with several layers of "neurons": simple interconnected processing units that interact.	AI_Ethics_Mark_Coeckelbergh	[an approach to AI that allows] computers to learn from experience and understand the world in terms of a hierarchy of concepts, with each concept defined through its relation to simpler concepts. By gathering knowledge from experience, this approach avoids the need for human operators to formally specify all the knowledge that the computer needs. The hierarchy of concepts enables the computer to learn complicated concepts by building them out of simpler ones. If we draw a graph showing how these concepts are built on top of each other, the graph is deep, with many layers.	deeplearningbook_intro		
deepfake	Digital images and audio that are artificially altered or manipulated by AI and/or deep learning often to make someone do or say something he or she did not actually do or say. Pictures or videos can be edited to put someone in a compromising position or to have someone make a controversial statement, even though the person did not actually do or say what is shown. Increasingly, it is becoming difficult to distinguish artificially manufactured material from actual videos and images.	Brookings_Institution	A digital picture or video that has been maliciously edited using an algorithm in a way that makes the video appear authentic.	GWUC								
deletion	Of an <X>, the action of destroying an instantiated <X>.	IEEE_Soft_Vocab										
denial-of-service	The prevention of authorized access to resources or the delaying of time-critical operations. (Time-critical may be milliseconds or it maybe hours, depending upon the service provided).	SP800-12	An attack that prevents or impairs the authorized use of information system resources or services.	CISA	when legitimate users are unable to access information systems, devices, or other network resources due to the actions of a malicious cyber threat actor. Services affected may include email, websites, online accounts (e.g., banking), or other services that rely on the affected computer or network. A denial-of-service condition is accomplished by flooding the targeted host or network with traffic until the target cannot respond or simply crashes, preventing access for legitimate users. DoS attacks can cost an organization both time and money while their resources and services are inaccessible.	ST04-015						
denigration	(denigrate) to attack the reputation of	merriam-webster_dictionary_2022	(denigrate) to deny the importance or validity of	merriam-webster_dictionary_2022								
dependability	<of an item> ability to perform as and when required (note 1: includes availability, reliability, recoverability, maintainability, and maintenance support performance, and, in some cases, other characteristics such as durability, safety and security. Note 2: used as a collective term for the time-related quality characteristics of an item).	ISO/IEC_TS_5723:2022(en)										
deployment	Phase of a project in which a system is put into operation and cutover issues are resolved	IEEE_Soft_Vocab										
descriptive analytics	Insights, reporting, and information answering the question, "Why did something happen?" Descriptive analytics determines information useful to understanding the cause(s) of an event(s).	IEEE_Guide_1PA										

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
deterministic	modelling [that] produces consistent outcomes for a given set of inputs, regardless of how many times the model is recalculated. The mathematical characteristics are known in this case. None of them is random, and each problem has just one set of specified values as well as one answer or solution. The unknown components in a deterministic model are external to the model. It deals with the definitive outcomes as opposed to random results and doesn't make allowances for error.	Sourabh_Mehta_deterministic										
deterministic algorithm	An algorithm that, given the same inputs, always produces the same outputs.	CSRC										
developer	A general term that includes developers or manufacturers of systems, system components, or system services; systems integrators; vendors; and product resellers. Development of systems, components, or services can occur internally within organizations or through external entities.	SP800-37	Individual or organization that performs development activities (including requirements analysis, design, testing through acceptance) during the system or software life-cycle process.	IEEE_Soft_Vocab								
diagnostic analytics	Insights, reporting, and information answering the question, "Why did something happen?" Diagnostic analytics determines information useful to understanding the cause(s) of an event(s).	IEEE_Guide_1PA										
diagnostics	Pertaining to the detection and isolation of faults or failures	IEEE_Software_Vocab										
differential privacy	Differential privacy is a method for measuring how much information the output of a computation reveals about an individual. It is based on the randomised injection of "noise". Noise is a random alteration of data in a dataset so that values such as direct or indirect identifiers of individuals are harder to reveal. An important aspect of differential privacy is the concept of "epsilon" or ε, which determines the level of added noise. Epsilon is also known as the "privacy budget" or "privacy parameter".	privacy-enhancing_technologies	For two datasets D and D' that differ in at most one element, a randomized algorithm M guarantees \emph{(ϵ, δ)-differential privacy} for any subset of the output S if M satisfies: \begin{equation} Pr[M(D) \in S] \leq exp(\epsilon)*Pr[M(D') \in S] + \delta \end{equation} Furthermore, when $\delta = 0$ an algorithm M is said to guarantee \emph{ϵ-differential privacy}	gong_differential_2020								
differential validity	Differential validity states that the validities in two applicant populations are unequal, that is, $pi \neq pa$.	hunter_differential_1979										
digital labor	Digital automation of information technology systems and/or business processes that successfully delivers work output previously performed by human labor or new work output that would typically or alternatively have been performed by human labor.	IEEE_Guide_1PA										
digital workforce	The collective suite of automation technologies delivering existing or new work output as applied in a business; the manifestation of digital labor.	IEEE_Guide_1PA										
dimension	The dimension of an object is a topological measure of the size of its covering properties. Roughly speaking, it is the number of coordinates needed to specify a point on the object.	wolfram_math_2022	Distinct components that a multidimensional construct encompasses	IEEE_Soft_Vocab								
dimension reduction	Dimensionality reduction is the process of taking data in a high dimensional space and mapping it into a new space whose dimensionality is much smaller	Shalev-Shwartz,_Shai										
discrimination	Disadvantageous treatment of a person based on belonging to a category rather than on individual merit.	Žliobaitė_Indrė										
disparate impact	Facially neutral practices that might nevertheless have an unjustified adverse impact on members of a protected class.	Lipton,_Zachary	For Predictor Y and Sensitive Impact S. Definition 6.2 Disparate Impact (DI) = P[Y' = 1	S != 1]/P[Y' = 1	S = 1]	friedler_comparative_2019								
disparate treatment	Intentional discrimination, including (i) decisions explicitly based on protected characteristics; and (ii) intentional discrimination via proxy variables (e.g literacy tests for voting eligibility).	Lipton,_Zachary										
distributional robustness	Optimizing the predictive accuracy for a whole class of distributions instead of just a single target distribution.	Meinshausen,_Nicolai										
diversity	Diversity refers to anything that sets one individual apart from another, including the full spectrum of human demographic differences as well as the different ideas, backgrounds, and opinions people bring.	Seth_Boden_2020	Diversity: The term diversity is used to describe individual differences (e.g. life experiences, learning and working styles, personality types) and group/social differences (e.g., race, socio-economic status, class, gender, sexual orientation, country of origin, ability, intellectual traditions and perspectives, as well as cultural, political, religious, and other affiliations) that can be engaged to achieve excellence in teaching, learning, research, scholarship, and administrative and support services.	GWU_diversity_and_inclusion	any dimension that can be used to differentiate groups and people from one another. It means respect for and appreciation of differences. But it's more than this. We all bring with us diverse perspectives, work experiences, life styles and cultures. Here in ODE we know the power of diversity is unleashed when we respect and value differences. Diversity is defined by who we are as individuals. HUD recognizes that its strength comes from the dedication, experience, talents, and perspectives of every employee. Diversity encompasses the range of similarities and differences each individual brings to the workplace, including but not limited to national origin, language, race, color, disability, ethnicity, gender, age, religion, sexual orientation, gender identity, socioeconomic status, veteran status, and family structures. We define workforce diversity as a collection of individual attributes that together help us pursue organizational objectives efficiently and effectively. In simple terms, diversity is the mix.	HUD_diversity_and_inclusion	the practice of including the many communities, identities, races, ethnicities, backgrounds, abilities, cultures, and beliefs of the American people, including underserved communities.	EO_DEIA_2021			inclusion	
documentation	Collection of documents on a given subject; written or pictorial information describing, defining, specifying, reporting, or certifying activities, requirements, procedures, or results.	IEEE_Soft_Vocab										
domain	Distinct scope, within which common characteristics are exhibited, common rules observed, and over which a distribution transparency is preserved.	IEEE_Soft_Vocab	A set of elements, data, resources, and functions that share a commonality in combinations of: (1) roles supported, (2) rules governing their use, and (3) protection needs.	SP800-160	<artificial intelligence> specific field of knowledgeor expertise	aime_measurement_2022, citing ISO/IEC 2382						
domain expertise	Domain expertise implies knowledge and understanding of the essential aspects of a specific field of inquiry.	McCue,_Colleen										
domain shift	Differences between the source and target domain data	Stacke,_Karin									distributional shift	
drinking your own champagne	The practice in which tech workers use their own product consistently to see how well it works and where improvements can be made.	kelley_dogfooding_2022									dogfooding, eating your own dogfood	
dynamic process	The process in which one or more paths are defined and may be utilized based on the conditions present at the time of execution.	IEEE_Guide_1										
eavesdropping	An attack in which an attacker listens passively to the authentication protocol to capture information that can be used in a subsequent active attack to masquerade as the claimant.	Reznik,_Leon	An attack in which an attacker listens passively to the authentication protocol to capture information that can be used in a subsequent active attack to masquerade as the claimant.	CSRC	A form of active wiretapping attack in which the attacker intercepts and selectively modifies communicated data to masquerade as one or more of the entities involved in a communication association.	NIST_CSRC_man-in-the-middle_attack	An attack in which an attacker is positioned between two communicating parties in order to intercept and/or alter data traveling between them. In the context of authentication, the attacker would be positioned between claimant and verifier, between registrant and CSP during enrollment, or between subscriber and CSP during authenticator binding.	NIST_CSRC_man-in-the-middle_attack	An attack where the adversary positions himself in between the user and the system so that he can intercept and alter data traveling between them.	NIST_CSRC_man-in-the-middle_attack	man-in-the-middle; interception attack	
edge case	a problem or situation, especially in computer programming, that only happens at the highest or lowest end of a range of possible values or in extreme situations:	cambridge_dictionary_2022										
effective challenge	The concept of effective challenge is used to improve AI implementation at large financial services organizations in the US. An interpretation of an effective challenge is that, when building AI systems, one of the best ways to guarantee good results is to actively challenge and review each step of the development process. Of course, a culture of effective challenge must apply to everyone developing an AI system, even so-called "rock-star" engineers and data scientists.	Darnell_Coss_Hall										
embedding	An embedding is a representation of a topological object, manifold, graph, field, etc. in a certain space in such a way that its connectivity or algebraic properties are preserved. For example, a field embedding preserves the algebraic structure of plus and times, an embedding of a topological space preserves open sets, and a graph embedding preserves connectivity. One space X is embedded in another space Y when the properties of Y restricted to X are the same as the properties of X.	wolfram_math_2022										
emergent risks											ontological uncertainty	
emulation	The use of a data processing system to imitate another data processing system, so that the imitating system accepts the same data, executes the same programs, and achieves the same results as the imitated system.	IEEE_Soft_Vocab										
end event	An activity, task, or output that describes or defines the conclusion of a process.	IEEE_Guide_1PA										
engineer	n. 3a: a designer or builder of engines; b: a person who is trained in or follows as a profession a branch of engineering; c: a person who carries through an enterprise by skillful or artful contrivance; 4: a person who runs or supervises an engine or an apparatus. v. 1: to lay out, construct, or manage as an engineer.	Merriam-Webster_engineer										
ensemble	a machine learning paradigm where multiple models (often called "weak learners") are trained to solve the same problem and combined to get better results. The main hypothesis is that when weak models are correctly combined we can obtain more accurate and/or robust models.	Joseph_Rocca_Ensemble_methods										
environment	Anything affecting a subject system or affected by a subject system through interactions with it, or anything sharing an interpretation of interactions with a subject system	IEEE_Soft_Vocab										
equality of odds	(Equalized odds). We say that a predictor bY satisfies equalized odds with respect to protected attribute A and outcome Y, if bY and A are independent conditional on Y.	hardt_equality_2016	The probability of a person in the positive class being correctly assigned a positive outcome and the probability of a person in a negative class being incorrectly assigned a positive outcome should both be the same for the protected and unprotected group members. In other words, the protected and unprotected groups should have equal rates for true positives and false positives.	Mehrabi,_Ninareh								
equality of opportunity	(Equal opportunity). We say that a binary predictor bY satisfies equal opportunity with respect to A and Y if $Pr[bY = 1 \| A = 0; Y = 1] = Pr[bY = 1 \| A = 1; Y = 1]$.	hardt_equality_2016	The probability of a person in positive class being assigned to a positive outcome should be equal for both protected and unprotected group members. In other words, the protected and unprotected groups should have equal true positive rates.	Mehrabi,_Ninareh								
error	The difference between the observed value of an index and its "true" value. Errors maybe random or systematic. Random errors are generally referred to as "errors". Systematic errors are called "biases".	OECD	Difference between a computed, observed, or measured value or condition and the true, specified, or theoretically correct value or condition.	IEEE_Soft_Vocab	measured quantity value minus a reference quantity value	aime_measurement_2022, citing ISO/IEC Guide 99						

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
error propagation	the way in which uncertainties in the variables affect the uncertainty in the calculated results.	Dorf_2018									propgation of uncertainty; proprgation of error	
ethics	definition 1a: "a set of moral principles : a theory or system of moral values"; definition 1b: "the principles of conduct governing an individual or a group"; definition 1c: "a consciousness of moral importance"; definition 1d: "a guiding philosophy"; definition 2: "a set of moral issues or aspects (such as rightness)"; definition 3: "the discipline dealing with what is good and bad and with moral duty and obligation"	Merriam-Webster_ethic	n. 1. the branch of philosophy that investigates both the content of moral judgments (i.e., what is right and what is wrong) and their nature (i.e., whether such judgments should be considered objective or subjective). The study of the first type of question is sometimes termed normative ethics and that of the second metaethics. Also called moral philosophy. 2. the principles of morally right conduct accepted by a person or a group or considered appropriate to a specific field. In psychological research, for example, proper ethics requires that participants be treated fairly and without harm and that investigators report results and findings honestly. See code of ethics; professional ethics; research ethics. —ethical adj.	APA_ethics								
ethics by design	An approach to technology ethics and a key component of responsible innovation that aims to integrate ethics in the design and development stage of the technology. Sometimes formulated as "embedding values in design." Similar terms are "value-sensitive design" and "ethically aligned design."	AI_Ethics_Mark_Coeckelbergh										
evaluation	(1) systematic determination of the extent to which an entity meets its specified criteria; (2) action that assesses the value of something	aime_measurement_2022. citing ISO/IEC 24765									Test, Evaluation, Verification and Validation (TEVV)	
evasion	In Evasion Attacks, the adversary solves a constrained optimization problem to find a small input perturbation that causes a large change in the loss function and results in output misclassification.	tabassi_adversarial_2019										
example	definition 1: "one that serves as a pattern to be imitated or not to be imitated"; definition 3: "one that is representative of all of a group or type"; definition 4: "a parallel or closely similar case especially when serving as a precedent or model"; definition 5: "an instance (such as a problem to be solved) serving to illustrate a rule or precept or to act as an exercise in the application of a rule"	Merriam-Webster_example										
exception	An event that occurs during the performance of the process that causes a diversion from the normal flow of the process. Exceptions are generated by an unanticipated event within a process due to an undefined or unknown input, undefined or unexpected outcome, or unforeseen sequencing of a task or event.	IEEE_Guide_1PA										
execute	To carry out a plan, a task command, or another instruction	SP1011	To carry out an instruction, process, or computer program; directing, managing, performing, and accomplishing the project work, providing the deliverables, and providing work performance information.	IEEE_Soft_Vocab								
executive	one that exercises administrative or managerial control	Merriam-Webster_executive										
ex-nomination	Ex-nomination is the harm of eliminating social identity by almost ignoring its existence. This term comes from Barthes where he coined it to describe what the bourgeoisie do to hide their name and identity by not referring to themselves as such to naturalize bourgeois ideology. This can show up in some of the same examples as mentioned above, as ex-nomination can present itself in technology not recognizing a certain class of people with facial recognition technology or by having implicit biases towards certain adjectives to describe certain classes	Blank,_Abagayle_Lee										
experiment	a series of observations conducted under controlled conditions to study a relationship with the purpose of drawing causal inferences about that relationship. An experiment involves the manipulation of an independent variable, the measurement of a dependent variable, and the exposure of various participants to one or more of the conditions being studied. Random selection of participants and their random assignment to conditions also are necessary in experiments.	apa_experiment_2023	A study of a fundamental physical process by the use of one or more computer simulators. Like empirical experiments, input variables (factors) are systematically changed to assess their impact upon simulator outputs (responses). Unlike empirical experiments, the simulator responses are deterministic, and this has implications: Computer experiments can appropriately have their factors with intermediate levels and the scope, especially the number of runs, can be more ambitious. Further, modeling methods based on interpolators (especially kriging) emerge as a viable approach. Good practice is to use Latin hypercubes for computer experiments, and advanced nonparametric modeling methods such as kriging, neural networks, and multivariate adaptive regression splines (MARS) in the data analysis stage. Important applications of computer experimental methods are for determining process optima and for evaluating process tolerances.	nist_statistics_2012								
expert system	A form of AI that attempts to replicate a human's expertise in an area, such as medical diagnosis. It combines a knowledge base with a set of hand-coded rules for applying that knowledge. Machine-learning techniques are increasingly replacing hand coding.	Hutson,_Matthew	Intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution.	Reznik,_Leon	An expert system is an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution.	OECD	Computer system that provides for expertly solving problems in a given field or application area by drawing inferences from a knowledge base developed from human expertise.	IEEE_Soft_Vocab	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. 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.	NSCAI		
expertise	The accumulation of specialized knowledge is often called *expertise*. Passive expertise is a type of knowledge-based specialization that arises from experiences in life and one's position in a society or culture. Formal expertise is the result of a self-selection of a domain of knowledge that is mastered deliberately and for which there are clear benchmarks of success.	Schneider_McGrew_in_Flanagan_McDonough_2018										
explainability	The ability to provide a human interpretable explanation for a machine learning prediction and produce insights about the causes of decisions, potentially to line up with human reasoning.	NISTIR_8269_Draft	Within the context of AI, the extent to which AI decisioning processes and outcomes are reasonably understood.	Comptroller_Office	The ability to explain or be explained. In the context of ethics, it refers to the ability to explain to others why you have done something or why you have made a decision; this is part of what it means to be responsible.	AI_Ethics_Mark_Coeckelbergh	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).	NSCAI			interpretability	
explainable artificial intelligence (XAI)	XAI seeks to make AI more understandable and interpretable, and therefore trustworthy. One of the complaints about artificial intelligence is the lack of transparency in how it operates. Many algorithm developers don't reveal the data that go into applications or how various factors are weighted and analyzed. That leads to a situation where outsiders cannot understand or explain how AI reached the outcome or decision that it did. That lack of explainability can lead people to suspect the worst about AI, and thus not trust AI in general or certain AI applications in particular. XAI seeks to help describe either the overall function of AI or the specific way it reaches decisions.	Brookings_Institution	AI that can explain to humans its actions, decisions, or recommendations, or can provide sufficient information about how it came to its result.	AI_Ethics_Mark_Coeckelbergh								
explainer	Functionality for providing details on or causes for fairness metric results.	AI_Fairness_360										
explanation	Systems deliver accompanying evidence or reason(s) for all outputs.	NISTIR_8269_Draft	The explanation principle obligates AI systems to supply evidence, support, or reasoning for each output.	NISTIR_8312								
exploratory	Exploratory Data Analysis (EDA) is an approach/philosophy for data analysis that employs a variety of techniques (mostly graphical) to 1. maximize insight into a data set; 2. uncover underlying structure; 3. extract important variables; 4. detect outliers and anomalies; 5. test underlying assumptions; 6. develop parsimonious models; and 7. determine optimal factor settings.	nist_statistics_2012										
external validity	A study has external validity to the degree that its results can be extended (generalized) beyond the limited research setting and sample in which they were obtained	bordens_Research_2011	the extent to which the results of research or testing can be generalized beyond the sample that generated them. The more specialized the sample, the less likely will it be that the results are highly generalizable to other individuals, situations, and time periods.	APA_external_validity								
facial recognition (FR)	A technology for identifying specific people based on pictures or videos. It operates by analyzing features such as the structure of the face, the distance between the eyes, and the angles between a person's eyes, nose, and mouth. It is controversial because of worries about privacy invasion, malicious applications, or abuse by government or corporate entities. In addition, there have been well-documented biases by race and gender with many facial recognition algorithms.	Brookings_Institution	Records the spatial geometry of distinguishing features of the face. Different vendors use different methods of facial recognition, however, all focus on measures of key features of the face.	Woodward	Face recognition algorithms, however, have no built-in notion of a particular person. They are not built to identify particular people; instead they include a face detector followed by a feature extraction algorithm that converts one or more images of a person into a vector of values that relate to the identity of the person. The extractor typically consists of a neural network that has been trained on ID-labeled images available to the developer. In operations, they act as generic extractors of identity-related information from photos of persons they have usually never seen before. Recognition proceeds as a differential operator: Algorithms compare two feature vectors and emit a similarity score. This is a vendor-defined numeric value expressing how similar the parent faces are. It is compared to a threshold value to decide whether two samples are from, or represent, the same person or not. Thus, recognition is mediated by persistent identity information stored in a feature vector (or "template").	NISTIR_8280							
fair-washing	promoting the false perception that a machine learning model respects some ethical values	aivodji_fairwashing_2019										
fairness (another entry for "algorithmic fairness"?)	"'cultural assumptions" regarding "the regulation of [human] life effected by stated and unstated rules of interaction;" rules that most interactants see as "generally applicable" and "reasonable." (We have to get the full definition from the book...	Anna Wierzbicka, English: Meaning and Culture (Oxford: Oxford University Press, 2006), 152–54										
fairness metric	A quantification of unwanted bias in training data or models.	AI_Fairness_360	A mathematical definition of "fairness" that is measurable. Some commonly used fairness metrics include: equalized odds predictive parity counterfactual fairness demographic parity Many fairness metrics are mutually exclusive; see incompatibility of fairness metrics.	google_glossary_2023								

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
false negative	An example in which the predictive model mistakenly classifies an item as in the negative class.	NSCAI	an outcome where the model *incorrectly* predicts the *negative* class.	google_dev_classification-true-false-positive-negative	A false negative is denying an applicant who should be approved	Varshney,_Kush	1. An instance in which a security tool intended to detect a particular threat fails to do so. 2. Incorrectly classifying malicious activity as benign.	CSRC_false_negative			Type II error (in statistics)	
false positive	An example in which the model mistakenly classifies an item as in the positive class	NSCAI	an outcome where the model *incorrectly* predicts the *positive* class.	google_dev_classification-true-false-positive-negative	A false positive is approving an applicant who should be denied	Varshney,_Kush	1. An alert that incorrectly indicates that a vulnerability is present. 2. An alert that incorrectly indicates that malicious activity is occurring. 3. An instance in which a security tool incorrectly classifies benign content as malicious. 4. Incorrectly classifying benign activity as malicious. 5. An erroneous acceptance of the hypothesis that a statistically significant event has been observed. This is also referred to as a type I error. This is also referred to as a type 1 error. When "health-testing" the components of a device, it often refers to a declaration that a component has malfunctioned – based on some statistical test(s) – despite the fact that the component was actually working correctly.	CSRC_false_positive			Type I error (in statistics)	
fault tolerance	The ability of a system or component to continue normal operation despite the presence of hardware or software faults	SP1011										
favorable label	A label whose value corresponds to an outcome that provides an advantage to the recipient. The opposite is an unfavorable label.	AI_Fairness_360										
feature	An attribute containing information for predicting the label.	AI_Fairness_360										
feature extraction	a more general method in which one tries to develop a transformation of the input space onto the lowdimensional subspace that preserves most of the relevant information	khalid_feature_2014										
feature importance	how important the feature was for the classification performance of the model; a measure of the individual contribution of the corresponding feature for a particular classifier, regardless of the shape (e.g., linear or nonlinear relationship) or direction of the feature effect	saarela_feature_2021										
feature shift	Unlike joint distribution shift detection, which cannot localize which features caused the shift, we define a new hypothesis test for each feature individually. Naïvely, the simplest test would be to check if the marginal distributions have changed for each feature (as explored by [25]); however, the marginal distribution would be easy for an adversary to simulate (e.g., by looping the sensor values from a previous day). Thus, marginal tests are not sufficient for our purpose. Therefore, we propose to use conditional distribution tests. More formally, our null and alternative hypothesis for the j-th feature is that its full conditional distribution (i.e., its distribution given all other features) has not shifted for all values of the other features.	kulinski_feature_2020										
federated learning	a learning model which addresses the problem of data governance and privacy by training algorithms collaboratively without transferring the data to another location.	Public_Health_and_Informatics_M1E_2021										
feedback loop	describes the process of leveraging the output of an AI system and corresponding end-user actions in order to retrain and improve models over time. The AI-generated output (predictions or recommendations) are compared against the final decision (for example, to perform work or not) and provides feedback to the model, allowing it to learn from its mistakes.	C3_ai_feedback_loop									closed-loop learning	
fitting	Fitting is the process of verifying whether the data item value is in the previously specified interval.	OECD										
firmware	Computer programs and data stored in hardware - typically in read-only memory (ROM) or programmable read-only memory (PROM) - such that the programs and data cannot be dynamically written or modified during execution of the programs.	SP800-37										
Forecasting	Estimate or prediction of conditions and events in the project's future based on information and knowledge available at the time of the forecast. The information is based on the project's past performance and expected future performance, and includes information that could impact the project in the future, such as estimate at completion and estimate to complete.	IEEE_Soft_Vocab	Predicting the future as accurately as possible, given all of the information available, including historical data and knowledge of any future events that might impact the forecasts.	Hyndman,_Rob								
four-fifths rule	a rule of thumb under which [federal enforcement agencies] will generally consider a selection rate for any race, sex, or ethnic group which is less than four-fifths (4/5ths) or eighty percent (80%) of the selection rate for the group with the highest selection rate as a substantially different rate of selection. . . . This "4/5ths" or "80%" rule of thumb is not intended as a legal definition, but is a practical means of keeping the attention of the enforcement agencies on serious discrepancies in rates of hiring, promotion and other selection decisions.	EEOC_Q&A_Employee_Selection										
fraud detection	Monitoring the behavior of populations of users in order to estimate, detect, or avoid undesirable behavior.	Kou,_Yufeng	detecting and recognizing fraudulent activities as they enter systems and report them to a system manager.	Behdad								
fully autonomous	Accomplishes its assigned mission, within a defined scope, without human intervention while adapting to operational and environmental conditions	SP1011										
generative adversarial network (GAN)	Generative Adversarial Networks, or GANs for short, are an approach to generative modeling using deep learning methods, such as convolutional neural networks. Generative modeling is an unsupervised learning task in machine learning that involves automatically discovering and learning the regularities or patterns in input data in such a way that the model can be used to generate or output new examples that plausibly could have been drawn from the original dataset.	Brownlee,_Jason	A pair of jointly trained neural networks that generates realistic new data and improves through competition. One net creates new examples (fake Picassos, say) as the other tries to detect the fakes.	Hutson,_Matthew	Generative adversarial networks (GANs) consist of two competing neural networks—a generator network that tries to create fake outputs (such as pictures), and a discriminator network that tries to determine whether the outputs are real or fake. A major advantage of this structure is that GANs can learn from less data than other deep learning algorithms.	CRS_AI	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.	NSCAI				
generative artificial intelligence	[a kind of artificial intelligence] capable of generating new content such as code, images, music, text, simulations, 3D objects, videos, and so on. It is considered an important part of AI research and development, as it has the potential to revolutionize many industries, including entertainment, art, and design.	Arham_Islam_History_2023	describes algorithms (such as ChatGPT) that can be used to create new content, including audio, code, images, text, simulations, and videos.	McKinsey_generative_AI								
global	[An approach that] [t]ries to understand the model as a whole.	arun_opportunities_2020	A global explanation produces a model that approximates the non-interpretable model.	NISTIR_8312_Full								
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. AI governance may include policies on the nature of AI applications developed and deployed versus those limited or withheld.	NSCAI	A framework of policies, rules, and processes for ensuring direction, management and accountability.	SP1270								
graph	Diagram that represents the variation of a variable in comparison with that of one or more other variables. Diagram or other representation consisting of a finite set of nodes and internode connections called edges or arcs.	IEEE_Soft_Vocab	A graph (sometimes called an undirected graph to distinguish it from a directed graph, or a simple graph to distinguish it from a multigraph) is a pair G = (V, E), where V is a set whose elements are called vertices (singular: vertex), and E is a set of paired vertices, whose elements are called edges (sometimes links or lines).	wikipedia_graph_2023								
graphical 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.	NSCAI										
graphical user interface (GUI)	A GUI is a type of computer human interface on a computer. It solves the blank screen problem that confronted early computer users. These early users sat down in front of a computer and faced a blank screen, with only a prompt. The computer gave the user no indication what the user was to do next. GUIs are an attempt to solve this blank screen problem. At a conceptual level, a computer human interface is a "means by which people and computers communicate with each other"	jansen_graphical_1998										
ground truth	information provided by direct observation as opposed to information provided by inference	Collins_Dictionary_ground_truth	value of the target variable for a particular item of labelledinput data	aime_measurement_2022, citing ISO/IEC 22989	In most accounts of supervised (machine) learning, the ground truth is considered to be the "dependent variable" that is predicted by a collection of features (independent variables).	Muller,_Michael						
group fairness	The goal of groups defined by protected attributes receiving similar treatments or outcomes.	AI_Fairness_360	Treat different groups equally	Mehrabi,_Ninareh								
hacker	Unauthorized user who attempts to or gains access to an information system.	Reznik,_Leon	Technically sophisticated computer enthusiast who uses his or her knowledge and means to gain unauthorized access to protected resources.	IEEE_Soft_Vocab								
hallucination	generated content that is nonsensical or unfaithful to the provided source content[; ...] there are two main types of hallucinations, namely intrinsic hallucination and extrinsic hallucination. [An intrinsic hallucination is a] generated output that contradicts the source content; [an extrinsic hallucination is a] generated output that cannot be verified from the source content (i.e., output can neither be supported nor contradicted by the source).	Survey_of_Hallucination_in_NLG	when a bot confidently says something that is not true.	Liam_Tung_2022_Meta_hallucination								
hardware	Physical equipment used to process, store, or transmit computer programs or data	IEEE_Soft_Vocab										
harm	An undesired outcome [whose] cost exceeds some threshold[; ...] the key points in the definition of safety are that: costs have to be sufficiently high in some human sense for events to be harmful, and that safety involves reducing both the probability of expected harms and the possibility of unexpected harms.	Engineering_safety_in_machine_learning	to damage, injure or hurt.	Black's_Law_Dictionary_harm								
harmful bias	Harmful bias can be either conscious or unconscious. Unconscious, also known as implicit bias, involves associations outside conscious awareness that lead to a negative evaluation of a person on the basis of characteristics such as race, gender, sexual orientation, or physical ability.3,14 Discrimination is behavior; discriminatory actions perpetrated by individuals or institutions refer to inequitable treatment of members of certain social groups that results in social advantages or disadvantages	humphrey_addressing_2020										
harms of allocation	unfairly assigned opportunities or resources due to algorithmic intervention[; ...] when a system [distributes] or withholds certain groups an opportunity or a resource. [They are] immediate, easily quantifiable, discrete, and transactional.	Lim,_Swee_Kiat_harms										

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
harms of bias	a skew that produces a type of harm[; ...] further classifie[d] into harms of allocation and harms of representation.	Lim_Swee_Kiat_harms										
harms of representation	algorithmically filtered depictions that are discriminatory. [They are] long term, difficult to formalize, diffuse, and cultural.	Lim_Swee_Kiat_harms									content harms	
human activity recognition (HAR)	the art of identifying and naming activities using [a]rtificial [i]ntelligence (AI) from the gathered activity raw data by utilizing various sources (so-called devices).	Gupta_et_al_HAR_2022										
human–assisted	The type of human–robot–interaction that that refers to situations during which human interactions are needed at the level of detail of task plans, i.e., during the execution of a task	SPI011										
human capital management	A C–suite business discipline that develops enterprise human capital strategies and ensures the human capital portfolio is effectively managed. Human capital management provides decision support by combining business and workforce intelligence to the development of enterprise human capital strategies: how to leverage people and their ideas effectively to achieve bottom–line business goals such as growing the business, increasing market share, margins, share price, and decreasing SG&A costs, as well as improving business processes, benefiting from technology investments, and increasing productivity.	Merritt,_Linda										
human–computer interaction (HCI)	methods and approaches for designing and architecting user interfaces and the interactions between humans and computer (or information) technology.	Poore_Lawrence_ARLIS_2023-01										
human–cognitive bias	Human–cognitive biases relate to how an individual or group perceives AI system information to make a decision or fill in missing information, or how humans think about purposes and functions of an AI system. Human biases are omnipresent in decision–making processes across the AI lifecycle and system use, including the design, implementation, operation, and maintenance of AI.	NIST_AI_RMF_1.0	Systematic error in judgment and decision–making common to all human beings which can be due to cognitive limitations, motivational factors, and/or adaptations to natural environments.									
human–enabled machine learning	Detection, correlation, and pattern recognition generated through machine–based observation of human operation of software systems capturing successful or unsuccessful operations to enable the creation of a useful predictive analytics capability.	IEEE_Guide_1PA										
human experiment	anything done to an individual to learn how it will affect [that person].	Bassiouni_Baffes_Evrard										
human–in–the–loop	An AI system that requires human interaction.	DOD_Modeling_and_Simulation_Glossary										
human–machine teaming (HMT)	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, computing methods that ideally complement people, methods that optimize goals of teamwork, and designs that enhance human–AI interaction.	NSCAI	methods and approaches for coordinating the functions and actions of (semi) autonomous machine capabilities and human users, which are granted equal weighting.	Poore_Lawrence_ARLIS_2023-01							human–AI teaming	
human–operator–intervention	The need for human interaction in a normally fully autonomous behavior due to some extenuating circumstances.	SPI011										
human subjects	a living individual about whom an investigator (whether professional or student) conducting research: (i) Obtains information or biospecimens through intervention or interaction with the individual, and uses, studies, or analyzes the information or biospecimens; or (ii) Obtains, uses, studies, analyzes, or generates identifiable private information or identifiable biospecimens.	45_CFR_46_2018_Requirements_(2018_Common_Rule)									participant	
human system integration (HSI)	methods and approaches for testing and optimizing all human–related considerations from a "whole–system" or "system–of–systems" level.	Poore_Lawrence_ARLIS_2023-01										
human values	Artificial intelligence systems use data we generate in our daily lives and as such are a mirror of our interests, weaknesses, and differences. Artificial intelligence, like any other technology, is not value–neutral. Understanding the values behind the technology and deciding on how we want our values to be incorporated in AI systems requires that we are also able to decide on how and what we want AI to mean in our societies. It implies deciding on ethical guidelines, governance policies, incentives, and regulations. And it also implies that we are aware of differences in interests and aims behind AI systems developed by others according to other cultures and principles. *See note.	Virginia_Dignum_Responsibility_and_Artificial_Intelligence	The quantity which a function f takes upon application to a given quantity.	wolfram_math_2022	what is important to people in their lives, with a focus on ethics and morality. Work to date has emphasized human well–being, dignity, and justice	Batya_Friedman_VSD_Introduction					
hyperparameters	the parameters that are used to either configure a ML model (e.g., the penalty parameter C in a support vector machine, and the learning rate to train a neural network) or to specify the algorithm used to minimize the loss function (e.g., the activation function and optimizer types in a neural network, and the kernel type in a support vector machine).	On_Hyperparameter_Optimization										
hypothesis testing	A term used generally to refer to testing significance when specific alternatives to the null hypothesis are considered.	OECD										
impact	the force of impression of one thing on another : a significant or major effect	Merriam–Webster_Impact	a powerful effect that something, especially something new, has on a situation or person	cambridge_impact_2023								
impact assessment	a risk management tool that seeks to ensure an organization has sufficiently considered a system's relative benefits and costs before implementation. In the context of AI, an impact assessment helps to answer a simple question: alongside this system's intended use, for whom could it fail?	Bipartisan_Policy_Center_impact_assessments										
impersonation	A malicious individual is able to impersonate a legitimate data subject to the data controller. The adversary forges a valid access request and goes through the identity verification enforced by the data controller. The data controller sends to the adversary the data of a legitimate data subject. Defeating impersonation is the primary objective of any authentication protocol. The result of this attack is a data breach (e.g. bloggers [sic] pretend to be someone they are not in order to wheedle out the information they are seeking obtaining information illegaly which they then sell for a specified price).	Security_Analysis_of_Subject_Access										
in–processing	Techniques that modify the algorithms in order to mitigate bias during model training. Model training processes could incorporate changes to the objective (cost) function or impose a new optimization constraint.	SPI270	Techniques that try to modify and change state–of–the–art learning algorithms to remove discrimination during the model training process.	Mehrabi,_Ninareh								
in–processing algorithm	A bias mitigation algorithm that is applied to a model during its training.	AI_Fairness_360										
incident	a situation in which AI systems caused, or nearly caused, real–world harm.	AI_Incident_Database	the occurrence of a technical event that affects the integrity of a Product and/or Model.	FBPML_Wiki	an alleged harm or near harm event to people, property, or the environment where an AI system is implicated.	AI_Incident_Editors	Adverse event(s) in a computer system or networks caused by a failure of a security mechanism, or an attempted or threatened breach of these mechanisms.	Hasan,_Raza				
incident response	a public official response to an incident ... from an entity (i.e. company, organization, individual) allegedly responsible for developing or deploying the AI or AI system involved in said incident.	AIID_incident_response										
inclusion a cultural and environmental feeling of belonging and sense of uniqueness. It represents the extent to which employees feel valued, respected, encouraged to fully participate, and able to be their authentic selves.		Seth,_Boden_2020	describ[ing] the active, intentional, and ongoing engagement with diversity -- in people, in the curriculum, in the co–curriculum, and in communities (e.g. intellectual, social, cultural, geographic) with which individuals might connect.	GWU_diversity_and_inclusion	a state of being valued, respected and supported. It's about focusing on the needs of every individual and ensuring the right conditions are in place for each person to achieve his or her full potential. Inclusion should be reflected in an organization's culture, practices and relationships that are in place to support a diverse workforce. Inclusion is the process of creating a working culture and environment that recognizes, appreciates, and effectively utilizes the talents, skills, and perspectives of every employee; uses employee skills to achieve the agency's objectives and mission; connects each employee to the organization; and encourages collaboration, flexibility, and fairness. We define inclusion as a set of behaviors (culture) that encourages employees to feel valued for their unique qualities and experience a sense of belonging.	HUD_diversity_and_inclusion	the recognition, appreciation, and use of the talents and skills of employees of all backgrounds.	EO_DEIA_2021		diversity	
independence	Of software quality assurance (SQA), situation in which SQA is free from technical, managerial, and financial influences, intentional or unintentional	IEEE_Soft_Vocab	Two events are independent if the occurrence of one event does not affect the chances of the occurrence of the other event. The mathematical formulation of the independence of events A and B is the probability of the occurrence of both A and B being equal to the product of the probabilities of A and B (i.e., P(A and B) = P(A)P(B))	nist_800_2010	In simple terms, inclusion is getting the mix to work together.							
individual fairness	The goal of similar individuals receiving similar treatments or outcomes.	AI_Fairness_360	Give similar predictions to similar individuals	Mehrabi,_Ninareh	A fairness metricthat checks whether similar individuals are classified similarly	aime_measurement_2022 citing Machine Learning Glossary by Google						
inference	The stage of ML in which a model is applied to a task. For example, a classifier model produces the classification of a test sample.	NISTIR_8269_Draft										
information input component	One of the three components of a model. This component delivers assumptions and data to the model.	Comptroller_Office										
information security	preservation of confidentiality, integrity and availability of information; in addition, other properties, such as authenticity, accountability, non–repudiation, and reliability can also be involved.	ISO/IEC_TS_5723:2022(en)										
input	Data received from an external source	IEEE_Soft_Vocab										
insider attack	Those who are within [an] organisation may have authorised access to vast amounts of sensitive company records that are essential for maintaining competitiveness and market position, and knowledge of information services and procedures that are crucial for daily operations. . . . [and] should an individual choose to act against the organisation, then with their privileged access and their extensive knowledge, they are well positioned to cause serious damage.	IEEE_Caught_in_the_Act										

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
in silico	carrying out some experiment by means of a computer simulation	World_Wide_Words_In_silico									computer simulation testing	
instance	Discrete, bounded thing with an intrinsic, immutable, and unique identity. Individual occurrence of a type	IEEE_Soft_Vocab	A single object of the world from which a model will be learned, or on which a model will be used (e.g., for prediction).	Kohavi,_Ron								
instance weight	A numerical value that multiplies the contribution of a data point in a model.	AI_Fairness_360										
integrity	Degree to which a system, product, or component prevents unauthorized access to, or modification of, computer programs or data.	IEEE_Soft_Vocab	Guarding against improper information modification or destruction, and includes ensuring information non-repudiation and authenticity.	CSRC	The property whereby information, an information system, or a component of a system has not been modified or destroyed in an unauthorized manner.	CISA	<data> property whereby data have not been altered in an unauthorized manner since they were created, transmitted, or stored; <systems> property of accuracy and completeness	ISO/IEC_TS_5723:2022(en)	the quality of moral consistency, honesty, and truthfulness with oneself and others.	APA_integrity		
intelligent process automation	A preconfigured software instance that combines business rules, experience-based context determination logic, and decision criteria to initiate and execute multiple interrelated human and automated processes in a dynamic context. The goal is to complete the execution of a combination of processes, activities, and tasks in one or more unrelated software systems that deliver a result or service with minimal or no human intervention.	IEEE_Guide_1PA										
interaction	Action that takes place with the participation of the environment of the object.	IEEE_Soft_Vocab										
internal validity	The ability of your research design to adequately test your hypotheses	bordens_research_2010	the degree to which a study or experiment is free from flaws in its internal structure and its results can therefore be taken to represent the true nature of the phenomenon. In other words, internal validity pertains to the soundness of results obtained within the controlled conditions of a particular study, specifically with respect to whether one can draw reasonable conclusions about cause-and-effect relationships among variables.	APA_internal_validity								
interoperability	The ability of software or hardware systems or components to operate together successfully with minimal effort by end user	SP1011	Degree to which two or more systems, products or components can exchange information and use the information that has been exchanged.	IEEE_Soft_Vocab	The ability for tools to work together in execution, communication, and data exchange under specific conditions.	NIST_1500						
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.	NSCAI	The ability to explain or to present an ML model's reasoning in understandable terms to a human	aime_measurement_2022, citing Machine Learning Glossary by Google							explainability	
interpretable model	An interpretable machine learning model obeys a domain-specific set of constraints to allow it (or its predictions, or the data) to be more easily understood by humans. These constraints can differ dramatically depending on the domain.	rudin_interpretable_2022										
intervenability	the property that intervention is possible concerning all ongoing or planned privacy relevant data processing[...] the data subjects themselves should be able to intervene with regards to the processing of their own data ... [to ensure] that data subjects have the ability to control how their data is processed and by whom.	Covert_et_al										
kill switch	a form of safety mechanism used to completely shut off a device in case of an emergency situation where it cannot be shut off using the normal process or if immediate shut off is required.	Techopedia_kill_switch										
knowledge	The sum of all information derived from diagnostic, descriptive, predictive, and prescriptive analytics embedded in or available to or from a cognitive computing system.	IEEE_Guide_1PA	<artificial intelligence> abstracted informationabout objects, events, concepts or rules, their relationships and properties, organizedfor goal-oriented systematic use	aime_measurement_2022, citinig ISO/IEC 22989								
label	A value corresponding to an outcome.	AI_Fairness_360	target variable assigned to a sample	aime_measurement_2022, citing ISO/IEC 22989								
label shift	Under label shift, the label distribution p(y) might change but the class-conditional distributions p(x|y) do not. ... We work with the label shift assumption, i.e., ps(x|y) = pt(x|y)	saurabh_label_2020										
large language model (LLM)	a class of language models that use deep-learning algorithms and are trained on extremely large textual datasets that can be multiple terabytes in size. LLMs can be classed into two types: generative or discriminatory. Generative LLMs are models that output text, such as the answer to a question or even writing an essay on a specific topic. They are typically unsupervised or semi-supervised learning models that predict what the response is for a given task. Discriminatory LLMs are supervised learning models that usually focus on classifying text, such as determining whether a text was made by a human or AI.	AI_Assurance_2022									language model	
language model	A language model is an approximative description that captures patterns and regularities present in natural language and is used for making assumptions on previously unseen language fragments.	Gustavii,_Ebba									large language model (LLM)	
learning	A procedure in artificial intelligence by which an artificial intelligence program improves its performance by gaining knowledge.	Dennis_Mercadal	the acquisition of novel information, behaviors, or abilities after practice, observation, or other experiences, as evidenced by change in behavior, knowledge, or brain function. Learning involves consciously or nonconsciously attending to relevant aspects of incoming information, mentally organizing the information into a coherent cognitive representation, and integrating it with relevant existing knowledge activated from long-term memory.	APA_learning								
least privilege	The principle that a security architecture should be designed so that each entity is granted the minimum system resources and authorizations that the entity needs to perform its function.	CSRC	The security objective of granting users only those accesses they need to perform their official duties.	SP-800-12								
lemmatization	the process of grouping together the different inflected forms of a word so they can be analyzed as a single item.	Artasanchez_Josh_AI_with_Python	in natural language processing[, ...] working with words according to their root lexical components	Techopedia_lemmatization	grouping together words with the same root or lemma but with different inflections or derivatives of meaning so they can be analyzed as one item.	Techslang_lemmatization	the grouping together of different forms of the same word.	TechTarget_lemmatization				
linear model	[a supervised learning algorithm that uses] a simple formula to find a best-fit line through a set of data points.	dataiku_ML_and_linear_models	(linear) An operator L^~ is said to be linear if, for every pair of functions f and g and scalar t, L^~(f+g)=L^~f+L^~g and L^~(tf)=tL^~f.	wolfram_mathworld_2022								
local	Mainly focus on explanation of individual data instances. Generates one explanation map g perdata x \in X.	arun_opportunities_2020	A local explanation explains a subset of decisions or is a per-decision explanation.	NISTIR_8312_Full								
localization	Creation of a national or specific regional version of a product.	IEEE_Soft_Vocab										
logistic model	(logistic equation) The continuous version of the logistic model is described by the differential equation dN)/(dt)=rN(K-N))/K, (1) where r is the Malthusian parameter (rate of maximum population growth) and K is the so-called carrying capacity (i.e., the maximum sustainable population). Dividing both sides by K and defining x=N/K then gives the differential equation dx)/(dt)=rx(1-x), (2) which is known as the logistic equation and has solution x(t)=1/(1+(1/(x_0)-1)e^(-rt)). (3) The function x(t) is sometimes known as the sigmoid function.	wolfram_mathworld_2022										
machine learning	A general approach for determining models from data.	AI_Fairness_360	Machine Learning is the study of computer algorithms that improve automatically through experience.	Mitchell,_Tom	Machine learning is based on algorithms that can learn from data without relying on rules-based programming.	Pyle,_and_San_José	The study or the application of computer algorithms that improve automatically through experience. 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	NSCAI	A subcategory of artificial intelligence; a method of designing a sequence of actions to solve a problem that optimizes automatically through experience and with limited or no human intervention.	Comptroller_Office		
machine observation	Machine detection and interpretation of relevant and meaningful events and conditions that impact operation of the computer system itself or other dependent mechanisms or processes essential to the purpose of the system.	IEEE_Guide_1PA										
malicious actor	See bad actor.											
malware	Hardware, firmware, or software that is intentionally included or inserted in a system for a harmful purpose.	Reznik,_Leon	Software that compromises the operation of a system by performing an unauthorized function or process.	CISA							trojan horse	
materiality	Refers to the significance of a matter in relation to a set of financial or performance information. If a matter is material to the set of information, then it is likely to be of significance to a user of that information	OECD										
McNamara fallacy	presum[ing] that (A) quantitative models of reality are always more accurate than other models; (B) the quantitative measurements that can be made most easily must be the most relevant; and (C) factors other than those currently being used in quantitative metrics must either not exist or not have a significant influence on success. This flawed approach to reasoning is also known as the quantitative fallacy.	McNamara_Fallacy									quantitative fallacy	
measurement	(Quantitative) (1) act or process of assigning a number or category to an entity to describe an attribute of that entity; (2) assignment of numbers to objects in a systematic way to represent properties of the object; (3) use of a metric to assign a value (e.g., a number or category) from a scale to an attribute of an entity; (4) set of operations having the object of determining a value of a measure; (5) assignment of values and labels to aspects of software engineering work products, processes, and resources plus the models that are derived from them, whether these models are developed using statistical or other techniques; (6) figure, extent, or amount obtained by measuring	aime_measurement_2022, citing ISO/IEC 24765	(Qualitative) (1) a way of learning about social reality [...][that uses] approaches [...] to explore, describe, or explain social phenomena[;] unpack the meaning people ascribe to activities, situations, events, or [artifacts]; build a depth of understanding about some aspect of social life; build "thick descriptions" (see Clifford Geertz, 1973) of people in naturalistic settings; explore new or underresearched areas; or make micro–macro links (illuminate connections between individuals–groups and institutional and/or cultural contexts). (2) [approaches that] can make visible and unpick the mechanisms which link particular variables, by looking at the explanations, or accounts, provided by those involved.	Leavy_OHQR_Intro	Qualitative measurement engages research methods and techniques to provide information about the nature of phenomenon. Qualitative methods are designed for systematic collection, organization, description and interpretation of non-numeric (textual, verbal or visual) data (Hammarberg et. al, 2016). Qualitative measurement generally answers questions about why, for whom, when, and how something is (or is not) observed, whereas quantitative measurement answers questions about what is observed. Elements assessed using qualitative measurement may include contextual norms or meaning, socio-cultural dynamics, individual or collective beliefs, and complex multi-component interactions or interventions (Busetto et. al, 2020).	Hammarberg_2016,_Busetto_2020	Documentation of assumptions and methods used is a foundational element of qualitative measurement, as the choice of single or combined methods is made based on the phenomenon and its context (Russell & Gregory, 2003). When appropriately paired, qualitative and quantitative measurement can provide corroboration or elaboration, demonstrate use cases, and/or identify conditions for complementarity or contradiction (Brannen, 2005).	Russell_2003,_Brannen_2005			

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
measurement method	generic description of a logical organization of operations used in a measurement	aime_measurement_2022, citing ISO/IEC Guide 99	logical sequence of operations, described generically, usedin quantifying an attribute with respect to a specified scale	aime_measurement_2022, citing ISO/IEC 24765								
measurement model	The initial confirmatory factory analysis (CFA) model that underlies the structural model [that] tests the adequacy (as indexed by model fit) of the specified relations whereby indicators are linked to their underlying construct.	Little_2013	A statistical model that links unobservable theoretical constructs, operationalized as latent variables, and observable properties—i.e., data about the world	jackman_oxford_2008								
measurability	ability to assess an attribute of an entity against a metric (note 1: "measurable" is the adjective form of "measurability")	ISO/IEC_TS_5723:2022(en)										
membership inference	given a machine learning model and a record, determining whether the record was used as part of the model's training dataset or not.											
metadata	Metadata is data that defines and describes other data.	OECD	Data that describe other data.	IEEE_Soft_Vocab	Data employed to annotate other data with descriptive information, possibly including their data descriptions, data about data ownership, access paths, access rights, and data volatility.							
metric	defined measurement method and measurement scale	ISO/IEC_TS_5723:2022(en)	(1) quantitative measure of the degree to which a system, component, or process possesses a given attribute; (2) defined measurement method and the measurement scale; c.f., measure in this section above	aime_measurement_2022, citing ISO/IEC 24765								
minimization	(Part of the ICO framework for auditing AI) AI systems generally require large amounts of data. However, organisations must comply with the minimisation principle under data protection law if using personal data. This means ensuring that any personal data is adequate, relevant and limited to what is necessary for the purposes for which it is processed. [...] The default approach of data scientists in designing and building AI systems will not necessarily take into account any data minimisation constraints. Organisations must therefore have in place risk management practices to ensure that data minimisation requirements, and all relevant minimisation techniques, are fully considered from the design phase, or, if AI systems are bought or operated by third parties, as part of the procurement process due diligence	ICO_data_minimisation	a data controller should limit the collection of personal information to what is directly relevant and necessary to accomplish a specified purpose. They should also retain the data only for as long as is necessary to fulfil that purpose. In other words, data controllers should collect only the personal data they really need, and should keep it only for as long as they need it. The data minimisation principle is expressed in Article 5(1)(c) of the GDPR and Article 4(1)(c) of Regulation (EU) 2018/1725, which provide that personal data must be "adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed".	EDPS_data_minimization								
mitigation	the process of lessening the severity of stereotypical/unjust associations and disparate model performance.	Arjun_Subramonian_bias_mitigation										
mixed methods	In mixed methods, the researcher collects and analyzes both qualitative and quantitative data rigorously in response to research questions and hypotheses; integrates the two forms of data and their results; organizes these procedures into specific research designs that provide the logic and procedures for conducting the study; and frames these procedures within theory and philosophy.	Creswell_Clark_mixed_methods	research in which the inquirer or investigator collects and analyzes data, integrates the findings, and draws inferences using both qualitative and quantitative approaches or methods in a single study or a program of study.	Lisa_M._Given_SAGE								
MLOPS	MLOps (machine learning operations) stands for the collection of techniques and tools for the deployment of ML models in production.	symeonidis_MLOps_2022										
model	A function that takes features as input and predicts labels as output.	AI_Fairness_360	Machine Learning algorithms and data processing designed, developed, trained and implemented to achieve set outputs, inclusive of datasets used for said purposes unless otherwise stated.	FBPML_Wiki	A model is a formalised expression of a theory or the causal situation which is regarded as having generated observed data. In statistical analysis the model is generally expressed in symbols, that is to say in a mathematical form, but diagrammatic models are also found. The word has recently become very popular and possibly somewhat over-worked.	OECD	A conceptual, mathematical, or physical representation of phenomenon observed in a system of ideas, events, or processes. In computationally-based models used in AI, phenomenon are often abstracted for mathematical representation, which means that characteristics that can not be represented mathematically may not be captured in the model.	SPI270	A quantitative method, system, or approach that applies statistical, economic, financial, or mathematical theories, techniques, and assumptions to process input data into quantitative estimates. A model consists of three components: an information input component, which delivers assumptions and data to the model; a processing component, which transforms inputs into estimates; and a reporting component, which translates the estimates into useful business information.	Comptroller_Office		
model assertion	Model assertions are arbitrary functions over a model's input and output that indicate when errors may be occurring	Kang_Daniel										
model card	short documents accompanying trained machine learning models that provide benchmarked evaluation in a variety of conditions, such as across different cultural, demographic, or phenotypic groups (e.g., race, geographic location, sex, Fitzpatrick skin type) and intersectional groups (e.g., age and race, or sex and Fitzpatrick skin type) that are relevant to the intended application domains. [They] also disclose the context in which models are intended to be used, details of the performance evaluation procedures, and other relevant information.	Model_Cards_for_Model_Reporting										
model debugging	Model debugging aims to diagnose a model's failures.	Jain_Saachi										
model decay	Model decay depicts that the performance of the model is degrading over time	Nayak_Pragati										
model editing	An area of research that aims to enable fast, data-efficient updates to a pre-trained base model's behavior for only a small region of the domain, without damaging model performance on other inputs of interest	Mitchell_Eric										
model extraction	Adversaries maliciously exploiting the query interface to steal the model. More precisely, in a model extraction attack, a good approximation of a sensitive or proprietary model held by the server is extracted (i.e. learned) by a dishonest user who interacts with the server only via the query interface.	Chandrasekaran_Varun									model inversion; model stealing	
model governance	Model Governance is the name for the overall internal framework of a firm or organization that controls the processes for Model Development, Model Validation and Model Usage, assign responsibilities and roles etc.	open_risk_2022										
model inventory	in the context of Risk Management, [...] a database/[management information system] developed for the purpose of aggregating quantitative model related information that is in use by a firm or organization.	ORM_model_inventory										
model overlay	Judgmental or qualitative adjustments to model inputs or outputs to compensate for model, data, or other known limitations. A model overlay is a type of override.	Comptroller_Office										
model risk management	model risk management encompasses governance and control mechanisms such as board and senior management oversight, policies and procedures, controls and compliance, and an appropriate incentive and organizational structure	Fed_Reserve										
model suite	A group of models that work together.	Comptroller_Office										
model training	the phase in the data science development lifecycle where practitioners try to fit the best combination of weights and bias to a machine learning algorithm to minimize a loss function over the prediction range	C3.ai_Model_Training	process to determine or to improve the parameters of a machine learning model, based on a machine learning algorithm, by using training data	aime_measurement_2022, citing ISO/IEC 22989								
model validation	the set of processes and activities intended to verify that models are performing as expected.	yields.io_model_validation	the set of principles, practices and organizational arrangements supporting a rigorous (audited) model development and validation cycle.	Open_Risk_Manual_model_validation								
monitoring	Examination of the status of the activities of a supplier and of their results by the acquirer or a third party.	IEEE_Soft_Vocab	Continual checking, supervising, critically observing or determining the status in order to identify change from the performance level required or expected.	SP800-160								
moral agency	The capacity for moral action, reasoning, judgment, and decision making, as opposed to merely having moral consequences.	AI_Ethics_Mark_Coeckelbergh										
moral patiency	The moral standing of an entity in the sense of how that entity should be treated.	AI_Ethics_Mark_Coeckelbergh										
naive Bayes	The naive Bayes classifier is a Bayesian learning method that has been found to be useful in many practical applications. It is called "naive" because it incorporates the simplifying assumption that attribute values are conditionally independent, given the classification of the instance. The naive Bayes classifier applies to learning tasks where each instance x is described by a conjunction of attribute values and where the target function f(x) can take on any value from some finite set V.	Mitchell_Tom										
natural language processing	A computer's attempt to "understand" spoken or written language. It must parse vocabulary, grammar, and intent, and allow for variation in language use. The process often involves machine learning.	Hutson, _Matthew	The ability of a machine to process, analyze, and mimic human language, either spoken or written.	NSCAI								
neural network	Neural networks Also known as artificial neural network, neural net, deep neural net; a computer system inspired by living brains.	Ranschaert, _Erik	A model that, taking inspiration from the brain, is composed of layers (at least one of which is hidden) consisting of simple connected units or neurons followed by nonlinearities	aime_measurement_2022, citing Machine Learnign Glossary by Google	A highly abstracted and simplified model of the human brain used in machine learning. A set of units receives pieces of an input (pixels in a photo, say), performs simple computations on them, and passes them on to the next layer of units. The final layer represents the answer.	Hutson, _Matthew						
nondiscrimination	the practice of treating people, companies, countries, etc. in the same way as others in order to be fair:	Cambridge Dictionary	In the context of machine learning non-discrimination can be defined as follows: (1) people that are similar in terms non-protected characteristics should receive similar predictions, and (2) differences in predictions across groups of people can only be as large as justified by non-protected characteristics.	Žliobaitė_Indre	the elimination of all existing discriminatory conditions, whether purposeful or inadvertent.	FWS_062_FW_1	the practice of treating people, companies, countries, etc. in the same way as others in order to be fair	Cambridge_Dictionary_non-discrimination				
normal flow	The intended flow of a process originating from a start event, continuing through all defined activities, and concluding successfully to its defined end event.	IEEE_Guide_IPA										
normalization	Conceptual procedure in database design that removes redundancy in a complex database by establishing dependencies and relationships between database entities. Normalization reduces storage requirements and avoids database inconsistencies.	OECD	The process of convertingan actual range of values into a standard range of values, typically -1to +1 or 0 to 1	aime_measurement_2022, citing Machine Learnign Glossary by Google								
objective evidence	data supporting the existence or verity of something (note: can be obtained through observation, measurement, test, or other means).	ISO/IEC_TS_5723:2022(en)										
observation	a piece of information received online from users, sensors, or other knowledge sources	poole_mackworth_observation	the careful, close examination of an object, process, or other phenomenon for the purpose of collecting data about it or drawing conclusions.	APA_observation								
off-label use	[when] data published for one task are used for another[. ...] off-label usage could lead to biased, overly optimistic results of machine-learning algorithms. The underlying cause is that public data are processed with hidden processing pipelines that alter the data features.	misuse_of_public_data										
offline learning	implies ... a static dataset that [one] know[s] from the start and the parameters of [one's] machine learning algorithm are adjusted to the whole dataset at once often loading the whole dataset into memory or in batches.	Ben_Auffarth_2021										

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
online learning	fitting [one's] model incrementally as the data flows in (streaming data).	Ben_Auffarth_2021										
ontology	A set of concepts and categories in a subject area or knowledge domain that shows their properties and the relationships among them to enable interoperability among disparate elements and systems and specify interfaces to independent, knowledge-based services for the purpose of enabling certain kinds of automated reasoning.	IEEE_Guide_1PA										
opacity	[to receive] the output of [an] algorithm (the classification decision) [and to not] have any concrete sense of how or why a particular classification has been arrived at from inputs.	Jenna_Burrell	A description of some deep learning systems [that] take an input and provide an output, but the calculations that occur in between are not easy for humans to interpret.	Hutson,_Matthew	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.	NSCAI	When AI system processes, functions, output or behavior are unavailable to all stakeholders - usually an antonym for transparency				black box; unexplainable	
operationalization	Putting AI systems or related concepts into use so they can be measured.											
operator	A role assumed by the person performing remote control or teleoperation, semi-autonomous operations, or other human-in-the-loop types of operations	SP1011	Individual or organization that performs the operations of a system.	IEEE_Soft_Vocab	Individual or organization that performs the operations of a system.	SP800-160						
opt-in	an individual makes an active affirmative indication of choice via a user interface signaling a desire to share their information with third parties.	IAPP_Privacy_Glossary									privacy; consent; opt-out	
opt-out	an individual makes an active affirmative indication of choice via a user interface signaling a desire not to share their information with third parties.	IAPP_Privacy_Glossary									privacy; consent; opt-in	
outcome	something that follows as a result or consequence	merriam_webster_outcome										
outlier	An outlier is a data point that is far from other points.	Russell_and_Norvig	An outlier is a data value that lies in the tail of the statistical distribution of a set of data values.	OECD	Values distant from mostother values. In machine learning, any of the following are outliers:• Weights with high absolute values• Predicted values relatively far away from the actual values• Input data whose values are more than roughly 3 standard deviations fromthe meanOutliers often cause problems in model training. Clipping is one way of managingoutliers	alme_measurement_2022 citing Machine Learning Glossary by Google						
output	Data transmitted to an external destination	IEEE_Soft_Vocab	Process by which an information processing system, or any of its parts, transfers data outside of that system or part	IEEE_Soft_Vocab								
overfitting	In statistics and machine learning, overfitting occurs when a model tries to predict a trend in data that is too noisy. Overfitting is the result of an overly complex model with too many parameters. A model that is overfitted is inaccurate because the trend does not reflect the reality of the data. An overfitted model is a model with a trend line that reflects the errors in the data that it is trained with, instead of accurately predicting unseen data.	Ranschaert,_Erik	Given a hypothesis space H, a hypothesis h element of H is said to overfit the training data if there exists some alternative hypothesis h' element of H, such that h has smaller error than h' over the training examples, but h' has a smaller error than h over the entire distribution of instance.	Mitchell,_Tom								
package	a folder with all the code and metadata needed to train and serve a machine learning model.	about_ML_packages										
parametric	A learning model that summarizes data with a set of parameters of fixed size (independent of the number of training examples)	Russell_and_Norvig										
parent process	A process that may contain one or more sub-processes, activities, and tasks.	IEEE_Guide_1PA										
parity	Bit(s) used to determine whether a block of data has been altered. Rationale: Term has been replaced by the term "parity bit".	NIST_CSRC_parity	the quality or state of being equal or equivalent	Merriam-Webster_parity								
participation	engag[ing] multiple stakeholders in deliberative processes in order to achieve consensus.	Sloane_et_al_2020										
participant	A computer system, data, input, business rule, human intervention, and other contributor to the flow of a process.	IEEE_Guide_1PA	a living individual about whom an investigator (whether professional or student) conducting research: (i) Obtains information or biospecimens through intervention or interaction with the individual, and uses, studies, or analyzes the information or biospecimens; or (ii) Obtains, uses, studies, analyzes, or generates identifiable private information or identifiable biospecimens.	45_CFR_46_2018_Requirements_(2018_Common_Rule)							human subject	
passive learning agent	A passive learning agent has a fixed policy that determines its behavior. An active learning agent gets to decide what actions to take.	Russell_and_Norvig									active learning agent	
personal data	'Personal data' means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.	GDPR	(l) "Personal information" means information that identifies, relates to, describes, is reasonably capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular consumer or household. Personal information includes, but is not limited to, the following if it identifies, relates to, describes, is reasonably capable of being associated with, or could be reasonably linked, directly or indirectly, with a particular consumer or household: (A) Identifiers such as a real name, alias, postal address, unique personal identifier, online identifier, Internet Protocol address, email address, account name, social security number, driver's license number, passport number, or other similar identifiers. (B) Any personal information described in subdivision (e) of Section 1798.80. (C) Characteristics of protected classifications under California or federal law. (D) Commercial information, including records of personal property, products or services purchased, obtained, or considered, or other purchasing or consuming histories or tendencies. (E) Biometric information. (F) Internet or other electronic network activity information, including, but not limited to, browsing history, search history, and information regarding a consumer's interaction with an internet website application, or advertisement. (G) Geolocation data. (H) Audio, electronic, visual, thermal, olfactory, or similar information. (I) Professional or employment-related information. (J) Education information, defined as information that is not publicly available personally identifiable information as defined in the Family Educational Rights and Privacy Act (20 U.S.C. Sec. 1232g; 34 C.F.R. Part 99). (K) Inferences drawn from any of the information identified in this subdivision to create a profile about a consumer reflecting the consumer's preferences, characteristics, psychological trends, predispositions, behavior, attitudes, intelligence, abilities, and aptitudes. (L) Sensitive personal information. (2) "Personal information" does not include publicly available information or lawfully obtained, truthful information that is a matter of public concern. For purposes of this paragraph, "publicly available" means: information that is lawfully made available from federal, state, or local government records, or information that a business has a reasonable basis to believe is lawfully made available to the general public by the consumer or from widely distributed media; or information made available by a person to whom the consumer has disclosed the information if the consumer has not restricted the information to a specific audience. "Publicly available" does not mean biometric information collected by a business about a consumer without the consumer's knowledge. (3) "Personal information" does not include consumer information that is deidentified or aggregate consumer information.	CCPA								
policy	The general principles by which a government is guided in its management of public affairs, or the legislature in its measures. This term, as applied to a law, ordinance, or rule of law, denotes its general purpose or tendency considered as directed to the POLICY	law_policy_2023	A policy defines the learning agent's way of behaving at a given time	sutton_reinforcement_2018								
positionality	Awareness and discussion of ones' social and institutional position with regards to research, particularly of power imbalances, and limitations the researcher may have because of differences in lived experience.	Malik_2021	the researcher's starting points and standpoints before and during inquiry, as well as the conditions shaping the research situation, process, and product.	Charmaz_Henwood							reflexivity	
post-hoc explanation	also known as the reverse engineering approach[:] tries to reconstruct explanations for decisions made by a black-box[:] . . . can be further divided into global and local methods. Global explanations concern understanding the overall logic and behavior of a black-box model, while local explanations try to find correlations between feature values of a record and an outcome. Global methods may use different types of explanators to open various black-boxes for different problems, e.g. decision trees for explaining random forests, prototypes to explain ensemble models, feature importance scores, decision trees for identifying training samples responsible for mispredictions , and minimal feature adjustment for reverting a class label.	Moradi_Samwald	Post-hoc explainability targets models that are not readily interpretable by design by resorting to diverse means to enhance their interpretability, such as text explanations, visual explanations, local explanations, explanations by example, explanations by simplification and feature relevance explanations techniques. Each of these techniques covers one of the most common ways humans explain systems and processes by themselves.	NISTIR_8312_Full	Post-hoc explainability targets models that are not readily inter- pretable by design by resorting to diverse means to enhance their in- terpretability, such as text explanations, visual explanations, local expla- nations, explanations by example, explanations by simplification and feature relevance explanations techniques. Each of these techniques covers one of the most common ways humans explain systems and processes by themselves.	barredo_explainable_2020						
post-processing	Typically performed with the help of a holdout dataset (data not used in the training of the model). Here, the learned model is treated as a black box and its predictions are altered by a function during the post-processing phase. The function is deduced from the performance of the black box model on the holdout dataset.	SP1270	Performed after training by accessing a holdout set that was not involved during the training of the model. If the algorithm can only treat the learned model as a black box without any ability to modify the training data or learning algorithm, then only post-processing can be used in which the labels assigned by the black-box model initially get reassigned based on a function during the post-processing phase.	Mehrabi,_Ninareh								
post-processing algorithm	A bias mitigation algorithm that is applied to predicted labels.	AI_Fairness_360										
practical significance	speaks to the magnitude of the relationship [between two variables] and whether or not that magnitude is important.	Mind_on_Statistics	a conceptual framework for evaluating discrimination cases developed primarily on statistical evidence that is the subject of increasing interest and discussion by some in the equal employment opportunity (EEO) field.	DOI:_Practical_Significance							statistical significance (often paired in contrast to this); substantive significance (synonym)	
pre-processing algorithm	A bias mitigation algorithm that is applied to training data.	AI_Fairness_360										

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
precision	A metric for classification models. Precision identifies the frequency with which a model was correct when classifying the positive class.	NSCAI	closeness of agreement between indications or measuredquantity values obtained by replicate measurements on the same or similarobjects under specified conditions	aime_measurement_2022, citing ISO/IEC Guide 99	A metric for classification models.Precision identifies the frequency with which a model was correct when predictingthe positive class. That is:Precision = True Positive/(True Positive + False Positive)	aime_measurement_2022, citing Machine Learning Glossary by Google	Closeness of agreement between independent test results obtained under prescribed conditions. It is generally dependent on analyte concentration, and this dependence should be determined and documented. The measure of precision is usually expressed in terms of imprecision and computed as a standard deviation of the test results. Higher imprecision is reflected by a larger standard deviation. Independent test results means results obtained in a manner not influenced by any previous results on the same or similar material. Precision covers repeatability and reproducibility [19]. Alternatively, precision is a measure for the reproducibility of measurements within a set, that is, of the scatter or dispersion of a set about its central value. Precision depends only on the distribution of random errors and does not relate to the true value or specified value.	UNODC_Glossary_QA_GLP				
prediction	Forecasting quantitative or qualitative outputs through function approximation, applied on input data or measurements.	NSCAI	primary output of an AI system when provided with input data or information	aime_measurement_2022, citing ISO/IEC 22989								
predictive analysis	The organization of analyses of structured and unstructured data for inference and correlation that provides a useful predictive capability to new circumstances or data.	IEEE_Guide_1PA										
predictive analytics	Insights, reporting, and information answering the question, "What is likely to happen?" Predictive analytics support high confidence foretelling of future event (s).	IEEE_Guide_1PA										
preprocessing	Transforming the data so that the underlying discrimination is mitigated. This method can be used if a modeling pipeline is allowed to modify the training data.	SP1270	Techniques that try to transform the data so the underlying discrimination is removed. If the algorithm is allowed to modify the training data, then pre-processing can be used.	Mehrabi, _Ninareh								
prescriptive analytics	Insights, reporting, and information answering the question, "What should I do about it?" Prescriptive analytics determines information that provides high confidence actions necessary to recover from an event or fulfill a need.	IEEE_Guide_1PA										
privacy	freedom from intrusion into the private life or affairs of an individual	ISO/IEC_TS_5723:2022(en)	freedom from intrusion into the private lifeor affairs of an individual when that intrusion results from undue or illegalgathering and use of data about that individual	aime_measurement_2022, citing ISO/IEC TR 24029-1								
privacy-by-design	Embedding privacy measures and privacy enhancing technologies directly into the design of information technologies and systems.	ENISA									data-protection-by-design (def: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A02016R0679-20160504&qid=1532348683434)	
privacy-enhancing technology	A coherent system of ICT (Information and Communications Technology) measures that protects privacy by eliminating or reducing personal data or by preventing unnecessary and/or undesired processing of personal data, all without losing the functionality of the information system.	PET_Handbook										
privileged protected attribute	A value of a protected attribute indicating a group that has historically been at systematic advantage.	AI_Fairness_360										
procedure	Information item that presents an ordered series of steps to perform a process, activity, or task.	IEEE_Soft_Vocab										
process	A sequence or flow of activities in an organization with the objective of carrying out work, which may include a set of activities, events, tasks, and decisions in a sequenced flow that adhere to finite execution semantics. Process levels will generally follow structure at the capability maturity model integration (CMMI) level.	IEEE_Guide_1PA	Set of interrelated or interacting activities that transforms inputs into outputs	IEEE_Soft_Vocab								
process flow	The defined representation of the overall progression of how a process is intended to be performed, including all exceptions.	IEEE_Guide_1PA										
processing	'Processing' means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction.	GDPR	"Processing" means any operation or set of operations that are performed on personal information or on sets of personal information, whether or not by automated means.	CCPA							personal data; processing	
processing environment	the combination of software and hardware on which the Application runs.	Law_Insider_processing_environment										
processor	'Processor' means a natural or legal person, public authority, agency or other body which processes personal data on behalf of the controller.	GDPR	"Processing" means any operation or set of operations that are performed on personal information or on sets of personal information, whether or not by automated means.	CCPA							personal data; processing; controller	
product manager	a specialized product management professional whose job is to manage the planning, development, launch, and success of products/solutions powered by AI, machine learning, and deep learning technologies.	productmanagerHQ_Josh_Fechter										
product owner	[person who is] focused on providing direction and prioritization for the cross-functional AI team, ensuring everyone remains focused on the overall vision and road map. This role is responsible for unifying individuals with diverse skills and backgrounds toward a common goal.	Forbes_Tracy_Kemp										
product velocity	how fast a product can be delivered to the market	Cost_Management_ch15										
productization	[turning the best performing model] into an actual "data product," ready to be used in live services.	Towards_Productizing										
profiling	'Profiling' means any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements.	GDPR	"Profiling" means any form of automated processing of personal information, as further defined by regulations pursuant to paragraph (16) of subdivision (a) of Section 1798.185, to evaluate certain personal aspects relating to a natural person and in particular to analyze or predict aspects concerning that natural person's performance at work, economic situation, health, personal preferences, interests, reliability, behavior, location, or movements.	CCPA	Measuring the characteristics of expected activity so that changes to it can be more easily identified.	CSRC					personal data; processing	
protected attribute	An attribute that partitions a population into groups whose outcomes should have parity. Examples include race, gender, caste, and religion. Protected attributes are not universal, but are application specific.	AI_Fairness_360										
protected class	[a feature] that may not be used as the basis for decisions [and] could be chosen because of legal mandates or because of organizational values. Some common protected [classes] include race, religion, national origin, gender, marital status, age, and socioeconomic status.	MIT_Protected_Attributes	A group of people with a common characteristic who are legally protected from [...] discrimination on the basis of that characteristic. Protected classes are created by both federal and state law.	Practical_Law_protected_class								
prototype	A prototype is an original model constructed to include all the technical characteristics and performances of the new product.	OECD										
provisioning	The granting of access rights and executional privilege to an agent (human or machine) within an application(s) or system(s).	IEEE_Guide_1PA										
proxy	A variable that can stand in for another, usually not directly observable or measurable, variable.	SP1270										
proxy discrimination	a particularly pernicious subset of disparate impact. Like all forms of disparate impact, it involves a facially neutral practice that disproportionately harms members of a protected class. But a practice producing a disparate impact only amounts to proxy discrimination when a second condition is met. In particular, proxy discrimination requires that the usefulness to the discriminator of a facially neutral practice derives, at least in part, from the very fact that it produces a disparate impact. This condition can be met either when the discriminator intends to disparately impact a protected group or when a legally-prohibited characteristic is predictive of the discriminator's goals in ways that cannot be captured more directly by non-suspect data.	Proxy_Discrimination						A variable V in a causal graph exhibits unresolved discrimination if there exists a directed path from A to V that is not blocked by a resolving variable, and V itself is non-resolving.	Mehrabi, _Ninareh (this definiiton is quite technical, though)		
pseudo-anonymization (pseudonymization)	'Pseudonymisation' means the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person;	GDPR	"Pseudonymize" or "Pseudonymization" means the processing of personal information in a manner that renders the personal information no longer attributable to a specific consumer without the use of additional information, provided that the additional information is kept separately and is subject to technical and organizational measures to ensure that the personal information is not attributed to an identified or identifiable consumer.	CCPA	A data management technique to strip identifiers linking data to an individual.	NSCAI					personal data; processing	
pseudoscience	a system of theories, assumptions, and methods erroneously regarded as scientific	Merriam-Webster_pseudoscience										
quality	The totality of features and characteristics of a product or service that bear on its ability to satisfy stated or implied needs.	OECD	<data> degree to which the characteristics of data satisfy stated and implied needs when used under specified conditions; <system> degree to which a set of inherent characteristics of an object fulfils requirements (an object can be a product, process or service)	ISO/IEC_TS_5723:2022(en)								
racialized	A socio-political process by which groups are ascribed a racial identity, whether or not members of the group self-identify as such	AAAS_AI_and_Bias_2022-09										
ranking	a type of machine learning that sorts data in a relevant order[; often used by companies] to optimize search and recommendations.	DEV_ranking	position, order, or standing within a group : RANK	Merriam-Webster_ranking								

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
reason code	Response functions created by linear regression algorithms are probably the most popular, accountable, and transparent class of machine learning models. These models will be referred to here as linear and monotonic. They are transparent because changing any given input feature (or sometimes a combination or function of an input feature) changes the response function output at a defined rate, in only one direction, and at a magnitude represented by a readily available coefficient. Monotonicity also enables accountability through intuitive, and even automatic, reasoning about predictions. For instance, if a lender rejects a credit card application, they can say exactly why because their probability of default model often assumes that credit scores, account balances, and the length of credit history are linearly and monotonically related to the ability to pay a credit card bill. When these explanations are created automatically and listed in plain English, they are typically called reason codes.	Machine_Learning_Interpretability_with_H2O_Driverless_AI	logical and actionable [response] in the event credit is denied or other adverse action is taken; emerge[s] directly from the model that generates the credit score and serve[s] the purpose of explaining to consumers why the consumer was denied credit.	Comparing_scores_and_reason_codes								
recall	A metric for classification models; identifies the frequency with which a model correctly classifies the true positive items.	NSCAI	A metric for classification modelsthat answers the following question: Out of all the possible positive labels,how many did the model correctly identify? That is: Recall = True Positive/(True Positive + false Negative)	aime_measurement_2022, citing Machine Learning Glossary by Google								
recognition	the automatic discovery of regularities in data through the use of computer algorithms and with the use of these regularities to take actions such as classifying the data into different categories.	Pattern_Recognition_and_Machine_Learning	a sense of awareness and familiarity experienced when one encounters people, events, or objects that have been encountered before or when one comes upon material that has been learned in the past.	APA_recognition	to transfer prior learning or past experience to current consciousness: that is, to retrieve and reproduce information; to remember.	APA_recall						
recommendation system	A software tool and techniques that provide suggestion based on the customer's taste to discover new appropriate thing for them by filtering personalized information based on the user's preferences from a large volume of information	Das_Debashis	A subclass of information filtering system that seek to predict 'rating' or 'preference' that a user would give to an item (such as music, books or movies) or social element (e.g. people or group) they had not yet considered, using a model built from the characteristics of an item (content based approaches) or the user's social environment (collaborative filtering approaches)	Sharma_Lalita								
rectification	An individual's right to have personal data about them corrected or amended by a business or other organization if it is inaccurate.	IAPP_Privacy_Glossary										
red-lining	the creation and maintenance of technology practices that further entrench discriminatory practices against already marginalized groups.	Banking_on_Your_Data_Christopher_Gilliard										
red-team	A group of people authorized and organized to emulate a potential adversary's attack or exploitation capabilities against an enterprise's security posture. The Red Team's objective is to improve enterprise cybersecurity by demonstrating the impacts of successful attacks and by demonstrating what works for the defenders (i.e., the Blue Team) in an operational environment. Also known as Cyber Red Team.	CSRC	a role-playing exercise in which a problem is examined from an adversary's or enemy's perspective.	Annie_Jacobsen_2015	The objective of red-teaming exercises is to "obtain a realistic level of risk and vulnerabilities against your technology, people and physical/facilities" [RedTeam Security Consulting, 2016, p. 1]. [See note.]	Pamela_Gosh_2021						
reference class	A class which is intended to describe structure and behavior of object identifiers. Its instances, called references, are passed by-value and indirectly represent objects by substituting for some primitive reference.	IGI_Global_reference_class										
reflexivity	A form of critical thinking that prompts us to consider the 'whys' and 'hows' of research, critically questioning the utility, ethics, and value of what, whom, and how we study	Jamieson_Govaart_Pownall	in qualitative research, the self-referential quality of a study in which the researcher reflects on the assumptions behind the study and especially the influence of his or her own motives, history, and biases on its conduct.	APA_reflexivity							positionality	
regression	Regression is a process of predicting the value to a yes or no label provided it falls on a continuous spectrum of input values, subcategory of supervised learning.	Ranschaert,_Erik	Refers to the process of developing an empirical (data-driven) model to predict and/or explain one or more attributes in a database or set of data. It is most frequently associated with the simple linear model (y=mx+b) taught in most introductory statistics courses; the same ideas have been extended in many directions, including classification problems. When the emphasis is on hypothesis testing and simple models, the regression output is typically a few parameters that provide a direct linkage from the input variables to the predicted variables (or classification). In other situations the emphasis is on explaining as much of the variability in the output variables as is "reasonable" from the input variables. In this case, there are a number of "advanced" techniques, such as smoothing splines, decision trees, neural nets, and so forth, for which there are many "free" parameters. The meaning of any one of these parameters can be obscure. Many Data Mining techniques are, at their core, variations on well-known regression techniques.	Raynor	the prediction of an exact value using a given set of data	Saleh_Alkhalifa_ML_in_Biotech				
reinforcement learning	A type of machine learning in which the algorithm learns by acting toward an abstract goal, such as "earn a high video game score" or "manage a factory efficiently." During training, each effort is evaluated based on its contribution toward the goal.	Hutson,_Matthew	Algorithms, in which an agent decides what to do to perform the given task to maximize the given function.	Reznik,_Leon	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.	NSCAI						
reliability	Reliability refers to the closeness of the initial estimated value(s) to the subsequent estimated values.	OECD	ability of an item to perform as required, without failure, for a given time interval, under given conditions. Note 1 to <system> definition: The time interval duration can be expressed in units appropriate to the item concerned (e.g. calendar time, operating cycles, distance run, etc.) and the units should always be clearly stated. Note 2 to <system> definition: Given conditions include aspects that affect reliability, such as: mode of operation, stress levels, environmental conditions, and maintenance.	ISO/IEC_TS_5723:2022(en)	property of consistentintended behaviour and results	aime_measurement_2022, citing ISO/IEC 22989						
remediation	The process of treating data by cleaning, organizing, and migrating it to a safe and secure environment for optimized usage is called data remediation. Generally [understood] as a process involving deleting unnecessary or unused data. However, the actual process . . . is very detailed and includes several steps, including replacing, updating, or modifying data along with cleaning it, organizing it, and getting rid of unnecessary data.	CPO_Magazine_Amar_Kanagaraj										
representativeness												
reproducibility	Closeness of the agreement between the results of measurements of the same measurand carried out under changed conditions of measurement.	IEEE_Soft_Vocab										
requirement	something essential to the existence or occurrence of something else : CONDITION	Merriam-Webster_requirement										
residual	Residuals are differences between the one-step-predicted output from the model and the measured output from the validation data set. Thus, residuals represent the portion of the validation data not explained by the model.	MathWorks_Residual										
residual analysis	a powerful [statistical] tool to detect the problems associated with the violation of the ANOVA [analysis of variance] assumptions.	fernandez_residual_1992	Residual analysis consists of two tests: the whiteness test and the independence test. According to the whiteness test criteria, a good model has the residual autocorrelation function inside the confidence interval of the corresponding estimates, indicating that the residuals are uncorrelated. According to the independence test criteria, a good model has residuals uncorrelated with past inputs. Evidence of correlation indicates that the model does not describe how part of the output relates to the corresponding input.	MathWorks_Residual								
resilience	The ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions. Resilience includes the ability to withstand and recover from deliberate attacks, accidents, or naturally occurring threats or incidents. The ability of a system to adapt to and recover from adverse conditions.	NISTIR_8269_Draft	<governance> ability to anticipate and adapt to, resist, or quickly recover from a potentially disruptive event, whether natural or man-made; <system> capability of a system to maintain its functions and structure in the face of internal and external change, and to degrade gracefully when this is necessary	ISO/IEC_TS_5723:2022(en)	ability of a system to recover operational conditionquickly following an incident	aime_measurement_2022, citing ISO/IEC 22989						
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.	NSCAI										
result	The consequential outcome of completing a process.	IEEE_Guide_1PA										
retention limit	refers to the amount of information that is stored long-term, and can be measured in volume (the size of the total collected logs in bytes) and time (the number of months or years that logs are stored for).	Industrial_Network_Security_2011										
risk	When considering the negative impact of a potential event, risk is a function of 1) the negative impact, or magnitude of harm, that would arise if the circumstance or event occurs and 2) the likelihood of occurrence (Adapted from: OMB Circular A-130:2016). Negative impact or harm can be experienced by individuals, groups, communities, organizations, society, the environment, and the planet.	NIST_AI_RMF_1.0	The risk to organizational operations (including mission, functions, image, reputation), organizational assets, individuals, other organizations, and the Nation due to the potential for unauthorized access, use, disclosure, disruption, modification, or destruction of information and/or a system.	Reznik,_Leon	A measure of the extent to which an entity is threatened by a potential circumstance or event, and typically a function of: (i) the adverse impacts that would arise if the circumstance or event occurs; and (ii) the likelihood of occurrence.	SP800-12	An uncertain event or condition that, if it occurs, has a positive or negative effect on a project's objectives	IEEE_Soft_Vocab	effect of uncertainty on objectives	ISO_IEC_38507		
risk control	mechanisms at the design, implementation, and evaluation stages [that can be taken] into consideration when developing responsible AI for organizations that includes security risks (cyber intrusion risks, privacy risks, and open source software risk), economic risks (e.g., job displacement risks), and performance risks (e.g., risk of errors and bias and risk of black box, and risk of explainability).	Toward_an_understanding_of_responsible_artificial_intelligence_practices										
risk tiering	when evaluating the desired level of explainability[, h]igh-risk models are worthy of deeper consideration: models with an elevated level of risk and/or models which are related to decisions about individuals (especially decisions which can be impactful), to privacy, to potential predetermination of individual's choice; models using sensitive variables (even if using them is formally legitimate); models related to decisions in healthcare, human resources, allocation of benefits. All of these are examples that should drive a higher model risk rating, and they should meet more stringent explainability expectations. Existing MRM policies need to be updated to ensure they take into consideration the relevant risks when rating AI/ML models and deciding what the desired level of interpretability emphasis to be in any given case.	pwc_Model_Risk_Management_of_AI_and_ML_Systems										
risk tolerance	Risk tolerance refers to the organization's or AI actor's ... readiness to bear the risk in order to achieve its objectives. Risk tolerance can be influenced by legal or regulatoryrequirements.	NIST_AI_RMF_1.0										

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
robotic desktop automation (RDA)	The computer application that makes available to a human operator a suite of predefined activity choreography to complete the execution of processes, activities, transactions, and tasks in one or more unrelated software systems to deliver a result or service in the course of human-initiated or -managed workflow.	IEEE_Guide_1PA										
robotic process automation (RPA)	A preconfigured software instance that uses business rules and predefined activity choreography to complete the autonomous execution of a combination of processes, activities, transactions, and tasks in one or more unrelated software systems to deliver a result or service with human exception management.	IEEE_Guide_1PA	Software to help in the automation of tasks, especially those that are tedious and repetitive.	NSCAI								
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.	NSCAI										
robustness	ability of a system to maintain its level of performance under a variety of circumstances	ISO/IEC_TS_5723:2022(en)	The ability of a machine learning model/algorithm to maintain correct and reliable performance under different conditions (e.g., unseen, noisy, or adversarially manipulated data).	NISTIR_8269_Draft								
root-mean-square deviation (RMSD)	of an estimator of a parameter[; ...] the square-root of the mean squared error (MSE) of the estimator. In symbols, if X is an estimator of the parameter t, then RMSE(X) = (E[(X-t)2])½. The RMSE of an estimator is a measure of the expected error of the estimator. The units of RMSE are the same as the units of the estimator.	Glossary_of_Statistical_Terms	a frequently used measure of the differences between values (sample or population values) predicted by a model or an estimator and the values observed	Wikipedia_RMSD							root-mean-square error (RMSE)	
rounding	the transformation of a number expressed in a particular base to a number with fewer digits.	Nick_Higham_1	In finite precision arithmetic[, the process by which] the result of an elementary arithmetic operation [that] does not generally lie in the underlying number system, F, . . . must be mapped back into F.	Nick_Higham_2								
row	describes a single entity or observation and the columns describe properties about that entity or observation. The more rows you have, the more examples from the problem domain that you have.	Machine_Learning_Mastery_Jason_Brownlee										
safety	property of a system such that it does not, under defined conditions, lead to a state in which human life, health, property, or the environment is endangered; [safety involves reducing both the probability of expected harms and the possibility of unexpected harms].	ISO/IEC_TS_5723:2022(en)	freedom from risk which is not tolerable	aime_measurement_2022, citinig ISO/IEC TR 24029-1								
scalability	The ability to increase or decrease the computational resources required to execute a varying volume of tasks, processes, or services.	IEEE_Guide_1PA										
score	A continuous value output from a classifier. Applying a threshold to a score results in a predicted label.	AI_Fairness_360										
screen out	Screen-out discrimination occurs when "a disability prevents a job applicant or employee from meeting—or lowers their performance on—a selection criterion, and the applicant or employee loses a job opportunity as a result."	EEOC_ADA_AI	I) to remove (someone or something that is not suitable for a particular purpose) from a group that is being examined; 2) to prevent (something harmful) from passing through	Merriam-Webster_screen_out								
security	resistance to intentional, unauthorized act(s) designed to cause harm or damage to a system	ISO/IEC_TS_5723:2022(en)	degree to which a product or system (3.38)protects information (3.20) and data (3.11) so that persons or other productsor systems have the degree of data access appropriate to their types and levelsof authorization	aime_measurement_2022, citing ISO/IEC TR 24029-1								
segmentation	The process of identifying homogeneous subgroups within a data table.	Raynor										
selective adherence	Decision-makers are more likely to follow advice (human or algorithmic-based) that matches stereotypical views of the decision subjects.	alon-barkat_human_2023										
self-aware system	A computing platform imbued with sufficient knowledge and analytic capability to make useful conclusions about its inputs, its own processing, and the use of its output so that it is capable of self- judgment and improvement consistent with its purpose.	IEEE_Guide_1PA										
self-diagnosis	Ability of a system to adequately take measurement information from sensors, validate the data, and communicate the processes and results to other devices	SP1011										
self-healing system	A computing system able to perceive that it is not operating correctly and, without human intervention, make the necessary adjustments to restore itself to normalcy.	IEEE_Guide_1PA										
semantic mapping	A strategic schema or framework of metadata labels applied to all data, data groups, data fields, data types, or data content used to introduce new or raw data into a corpus or data fabric to give machine learning algorithms direction for investigating known or potential relationships between data. A semantic map provides a structure for the introduction of new data, information, or knowledge	IEEE_Guide_1PA										
sensitivity analysis	A "what-if" type of analysis to determine the sensitivity of the outcomes to changes in parameters. If a small change in a parameter results in relatively large changes in the outcomes, the outcomes are said to be sensitive to that parameter.	OECD	Sensitivity analysis varies settings of a model's input parameters and assesses resulting changes in model outputs.	Raynor	Sensitivity analysis varies settings of a model's input parameters and assesses resulting changes in model outputs.	mills_study_2010						
sensitive data	a specific set of "special categories" that must be treated with extra security[, including] information about: [r]acial or ethnic origin; [p]olitical opinions; [r]eligious or philosophical beliefs; [t]rade union membership; [g]enetic data; [d]ata related to a person's sex life or sexual orientation; and [b]iometric data (where processed to uniquely identify someone).	IT_Governance_Blog_Luke_Irwin										
sensory digitization	The conversion of typically analog or human sensory perception (e.g., vision, speech) to a digital format useful for machine-to-human interaction or machine processing of traditionally analog sensory information [e.g., optical character recognition (OCR)].	IEEE_Guide_1PA										
service	A collection of coordinated processes that takes one or more kinds of input, performs a value-added transformation, and creates an output that fulfills the needs of a customer [or shareholder].	IEEE_Guide_1PA										
signal detection theory	a framework for interpreting data from experiments in which accuracy is measured.	Signal_Detection_Theory										
shallow learning	Techniques that separate the process of feature extraction from learning itself.	Reznik_Leon										
situational awareness	Perception of elements in the system and/or environment and a comprehension of their meaning, which could include a projection of the future status of perceived elements and the uncertainty associated with that status.	SP800-160										
snake oil	Something proposed as a solution to a problem, but which is of little real or practical value; speech or action which is superficially attractive or convincing but of no real substance or effectiveness.	OED_snake_oil										
socio-technical system	how humans interact with technology within the broader societal context	NIST SP1270	system that includes a combination of technical and human or natural elements	ISO/IEC_TS_5723:2022(en)								
software testing	Activity in which a system or component is executed under specified conditions, the results are observed or recorded, and an evaluation is made of some aspect of the system or component.	IEEE_Soft_Vocab										
sparsity	refers to a matrix of numbers that includes many zeros or values that will not significantly impact a calculation.	Dave_Salvator_sparsity										
specification	A document that specifies, in a complete, precise, verifiable manner, the requirements, design, behavior, or other characteristics of a system or component and often the procedures for determining whether these provisions have been satisfied.	SP800-37										
stakeholder	Individual or organization having a right, share, claim, or interest in a system or in its possession of characteristics that meet their needs and expectations. An individual, group, or organization who may affect, be affected by, or perceive itself to be affected by a decision, activity, or outcome of a project.	IEEE_Soft_Vocab	any individual, group, or organization that can affect, be affected by, or perceive itself to be affected by a decision or activity	ISO/IEC_TS_5723:2022(en)								
standard deviation	The most widely used measure of dispersion of a frequency distribution introduced by K. Pearson (1893). It is equal to the positive square root of the variance. The standard deviation should not be confused with the root mean square deviation.	OECD										
start event	An activity, task, or input that describes or defines the beginning of a process.	IEEE_Guide_1PA										
statistical bias	A systematic tendency for estimates or measurements to be above or below their true values. Statistical biases arise from systematic as opposed to random error. Statistical bias can occur in the absence of prejudice, partiality, or discriminatory intent.	SP1270										
statistical parity	The independence between the protected attribute and the outcome of the decision rule	Besse,_Philippe										
statistical significance	When the probability of obtaining a statistic of a given size due strictly to random sampling error, or chance, is less than the selected alpha level [or the probability of a type I error]; also represents a rejection of the null hypothesis.	Statistics_in_Plain_English	refers to whether a relationship between two or more variables exists beyond a probability expected by chance	The_SAGE_Encyclopedia_of_Communication_Research_Methods								
statistics	Numerical data relating to an aggregate of individuals; the science of collecting, analysing and interpreting such data	OECD										
stereotype	Definition 3b (figurative): A preconceived and oversimplified idea of the characteristics which typify a person, situation, etc.; an attitude based on such a preconception. Also, a person who appears to conform closely to the idea of a type.	OED_stereotype	a set of cognitive generalizations (e.g., beliefs, expectations) about the qualities and characteristics of the members of a group or social category. Stereotypes, like schemas, simplify and expedite perceptions and judgments, but they are often exaggerated, negative rather than positive, and resistant to revision even when perceivers encounter individuals with qualities that are not congruent with the stereotype.	APA_stereotype	Contemporary social psychology typically defines stereotypes as mental representations of a group and its members, and *stereotyping* as the cognitive activity of treating individual elements in terms of higher level categorial properties	Augoustinos_Walker_1998						
stochastic	The adjective "stochastic" implies the presence of a random variable; e.g. stochastic variation is variation in which at least one of the elements is a variate and a stochastic process is one wherein the system incorporates an element of randomness as opposed to a deterministic system.	OECD										

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
straight-through processing (STP)	The successful execution of a service, process, or transaction performed entirely through traditional application platforms with predefined interfaces (i.e., application programming interfaces [APIs]).	IEEE_Guide_1PA										
strawperson	a fallacious argument which irrelevantly attacks a position that appears similar to, but is actually different from, an opponent's position, and concludes that the opponent's real position has thereby been refuted.	Hughes_Lavery_Critical_Thinking										
stress test	Type of performance efficiency testing conducted to evaluate a test item's behavior under conditions of loading above anticipated or specified capacity requirements, or of resource availability below minimum specified requirements	IEEE_Soft_Vocab										
strong AI	AI that is capable of solving almost all tasks that humans can solve	Shevlin_et_al_2019	AI that is as smart and well-rounded as a human. Some say it's impossible. Current AI is weak, or narrow. It can play chess or drive but not both, and lacks common sense.	Hutson,_Matthew							artificial general intelligence (AGI)	
structured data	Data that has a predefined data model or is organized in a predefined way.	NIST_1500										
sub-process	A subordinate process that can be included within a parent process. It can be present and/or repeated within other parent processes.	IEEE_Guide_1PA										
supervised learning	A type of machine learning in which the algorithm compares its outputs with the correct outputs during training. In unsupervised learning, the algorithm merely looks for patterns in a set of data.	Hutson,_Matthew	Algorithms, which develop a mathematical model from the input data and known desired outputs.	Reznik,_Leon	For a computer to process a set of data whose attributes have been divided into two groups and derive a relationship between the values of one and the values of the other. These two groups are sometimes called predictor and targets, respectively. In statistical terminology, they are called independent and dependent variables. Respectively. The learning is "supervised because the distinction between the predictors and the target variables is chosen by the investigator or some other outside agency.	Raynor	a general subset of machine learning in which data, like its associated labels, is used to train models that can learn or generalize from the data to make predictions, preferably with a high degree of certainty.	Saleh_Alkhalifa_ML_in_Biotech				
support vector machines	A supervised machine learning model for data classification and regression analysis. One of the most used classifiers in machine learning. It optimizes the width of the gap between the points of separate categories in feature space.	Ranschaert,_Erik										
surveillance	an outcome of establishing information infrastructures as the basis for administration, proddduction, marketing, entertainment and law enforcement[, involving] garnering personal data for a variety of purposes in a quest for greater efficiency, convenience or safety. Its ethics and politics are inherently ambiguous, but at the same time surveillance is never neutral.	David_Lyon_2007										
system	combination of interacting elements organized to achieve one or more stated purposes	ISO/IEC_TS_5723:2022(en)										
systemic bias	Systemic biases result from procedures and practices of particular institutions that operate in ways which result in certain social groups being advantaged or favored and others being disadvantaged or devalued. This need not be the result of any conscious prejudice or discrimination but rather of the majority following existing rules or norms.	D. Chandler and R. Munday, A Dictionary of Media and Communication. Oxford University Press, Jan. 2011, publication Title: A Dictionary of Media and Communication										
system of systems	set of systems and system elements that interact to provide a unique capability that none of the constituent systems can accomplish on its own (note: can be necessary to facilitate interaction of the constituent systems in the system of systems)	ISO/IEC_TS_5723:2022(en)										
talent acquisition	the process of finding and acquiring skilled human labor for organizational needs and to meet any labor requirement.	IJMAE										
target	a method for solving a problem that an AI algorithm parses its training data to find. Once an algorithm finds its target function, that function can be used to predict results (predictive analysis). The function can then be used to find output data related to inputs for real problems where, unlike training sets, outputs are not included.	TechTarget_target_function									target variable, target value	
task	The performance of a discrete activity with a defined start, stop, and outcome that cannot be broken down to a finer level of detail.	IEEE_Guide_1PA	Required, recommended, or permissible action, intended to contribute to the achievement of one or more outcomes of a process	IEEE_Soft_Vocab	set of activities undertaken in order to achieve a specific goal	aime_measurement_2022, citing ISO/IEC TR 24030						
taxonomy	Taxonomy refers to classification according to presumed natural relationships among types and their subtypes.	OECD										
tech-washing	the practice of slapping a trendy, new label on legacy solutions.	Forbes_Kayvan_Alikhani										
technical control	Security controls (i.e., safeguards or countermeasures) for an information system that are primarily implemented and executed by the information system through mechanisms contained in the hardware, software, or firmware components of the system.	NIST_SP_800-30_Rev_1										
technochauvinism	The belief that technology is always the solution	M. Broussard, Artificial Unintelligence: How Computers Misunderstand the World. MIT Press, 2018.									techno-solutionism	
techno-solutionism	See technochauvinism.											
test	Technical operation to determine one or more characteristics of or to evaluate the performance of a given product, material, equipment, organism, physical phenomenon, process or service according to a specified *procedure*.	UNODC_Glossary_QA_GLP	any activity aimed at evaluating an attribute or capability of a program or system and deteermining that it meets its required results.	William_Hetzel	(1) activity in which a system or component is executedunder specified conditions, the results are observed or recorded, and an evaluationis made of some aspect of the system or component; (2) to conduct anactivity as in (1); (3) set of one or more test cases and procedures.	aime_measurement_2022, citing ISO/IEC 24765	the process of executing a program with the intent of finding errors.	The_Art_of_Software_Testing			Test, Evaluation, Verification and Validation (TEVV)	
Test and Evaluation, Verification and Validation (TEVV)	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.	NSCAI_Report										
third party	an entity that is involved in some way in an interaction that is primarily between two other entities. [Please see note, especially regarding NIST CSRC terms that we might incorporate into this definition.]	TechTarget_third_party										
threat actor	See bad actor.											
three lines of defense	Most financial institutions follow a three-lines-of-defense model, which separates front line groups, which are generally accountable for business risks (the First Line), from other risk oversight and independent challenge groups (the Second Line) and assurance (the Third Line)	AIRS_Penn										
traceability	Ability to trace the history, application or location of an entity by means of recorded identification. ["Chain of custody" is a related term.] Alternatively, traceability is a property of the result of a measurement or the value of a standard whereby it can be related with a stated uncertainty, to stated references, usually national or international standards, i.e. through an unbroken chain of comparisons. In this context, The standards referred to here are measurement standards rather than written standards.	UNODC_Glossary_QA_GLP	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).	NSCAI								
training data	A dataset from which a model is learned.	AI_Fairness_360	a sample from a population of possible examples, and the statistical similarities of each class extracted, or more precisely the significant differences between classes are found.	Ripley,_Brian	samples for training used to fit a machine learningmodel	aime_measurement_2022, citing ISO/IEC 22989						
transaction	Enactment of a process represented by a set of coordinated activities carried out by multiple systems and/or participants in accordance with defined relationships. This coordination leads to an intentional, consistent, and verifiable result across all participants.	IEEE_Guide_1PA										
transfer learning	A technique in machine learning in which an algorithm learns to perform one task, such as recognizing cars, and builds on that knowledge when learning a different but related task, such as recognizing cats.	Hutson,_Matthew										
transformer	A procedure that modifies a dataset.	AI_Fairness_360										
transparency	<information> open, comprehensive, accessible, clear and understandable presentation of information; <systems> property of a system or process to imply openness and accountability	ISO/IEC_TS_5723:2022(en)	Understanding the working logic of the model.	NISTIR_8269_Draft	<organization> property of an organization that appropriate activities and decisions are communicated to relevant stakeholders (3.5.13) in a comprehensive, accessible and understandable manner Note 1 to entry: Inappropriate communication of activities and decisions can violate security, privacy or confidentiality requirements.	iso_22989_2022	<system> property of a system that appropriate information about the system is made available to relevant stakeholders (3.5.13) Note 1 to entry: Appropriate information for system transparency can include aspects such as features, performance, limitations, components, procedures, measures, design goals, design choices and assumptions, data sources and labelling protocols. Note 2 to entry: Inappropriate disclosure of some aspects of a system can violate security, privacy or confidentiality requirements.	iso_22989_2022				
trojan horse	A computer program that appears to have a useful function, but also has a hidden and potentially malicious function that evades security mechanisms, sometimes by exploiting legitimate authorizations of a system entity that invokes the program.	Reznik,_Leon									malware	
true negative	outcome where the model correctly predicts the negative class.	google_dev_classification-true-false-positive-negative										

Terms	Definition 1	Citation 1 [1]	Definition 2	Citation 2	Definition 3	Citation 3	Definition 4	Citation 4	Definition 5	Citation 5	Related terms and synonyms [2]	Legal definition applicable
true positive	an outcome where the model *correctly* predicts the *positive* class.	google_dev__classification-true-false-positive-negative										
trust	the system status in the mind of human beings based on their perception of and experience with the system; concerns the attitude that a person or technology will help achieve specific goals in a situation characterized by uncertainty and vulnerability.	DOD_TEVV	degree to which a user or other stakeholder has confidence that a product or system will behave as intended	aime_measurement_2022_citing ISO/IEC TR 24029-1	the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability.							
trustworthiness	The degree to which an information system (including the information technology components that are used to build the system) can be expected to preserve the confidentiality, integrity, and availability of the information being processed, stored, or transmitted by the system across the full range of threats and individuals' privacy.	SP800-37	Worthy of being trusted to fulfill whatever critical requirements may be needed for a particular component, subsystem, system, network, application, mission, enterprise, or other entity.	SP800-160	ability to meet stakeholders' expectations in a verifiable way; an attribute that can be applied to services, products, technology, data and information as well as to organizations.	ISO/IEC_TS_5723:2022(en)						
trustworthy AI	Characteristics of trustworthy AI systems include: valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy-enhanced, and fair with harmful bias managed.	NIST_AI_RMF_1.0	AI that can be trusted by humans. Conditions for such trust can refer to (other) ethical principles such as human dignity, respect for human rights, and so on, and/or to social and technical factors that influence whether people will want to use the technology. The use of the term "trust" with regard to technologies is controversial.	AI_Ethics_Mark_Coeckelbergh	Trustworthy AI has three components: (i) it should be lawful, ensuring compliance with all applicable laws and regulations (ii) it should be ethical, demonstrating respect for, and ensure adherence to, ethical principles and values and (3) it should be robust, both from a technical and social perspective, since, even with good intentions, AI systems can cause unintentional harm. Trustworthy AI concerns not only the trustworthiness of the AI system itself but also comprises the trustworthiness of all processes and actors that are part of the system's life cycle.	european_ethics_2019					
type I error	The null hypothesis H0 is rejected, even though it is [true]	berthold_guide_2020	false positive rate	james_statistical_2014								
type II error	The null hypothesis H0 is accepted, even though it is [false]	berthold_guide_2020	true positive rate	james_statistical_2014								
uncertainty	Result of not having accurate or sufficient knowledge of a situation; state, even partial, of deficiency of information related to understanding or knowledge of an event, its consequence, or likelihood	IEEE_Soft_Vocab										
underfitting	Underfitting occurs when a statistical model cannot adequately capture the underlying structure of the data.	Ranschaert,_Erik										
underrepresentation	inadequately represented. (See note.)	Merriam-Webster_underrepresented	when members of discernible groups are not consistently present in representative bodies and among measures of well-being in numbers roughly proportionate to their numbers within the population.	Encyclopedia.com_underrepresentation								
unexplainable	impossibility of providing an explanation for certain decisions made by an intelligent system which is both 100% accurate and comprehensible.	Roman_V._Yampolskiy_Unexplainability									black box; opacity	
unstructured data	Data that does not have a predefined data model or is not organized in a predefined way											
unsupervised learning	Algorithms, which take a set of data consisting only of inputs and then they attempt to cluster the data objects based on the similarities or dissimilarities in them.	Reznik,_Leon	Learning techniques that group instances without a pre-specified dependent attribute.	Kohavi,_Ron	A learning strategy that consists in observing and analyzing different entities and determining that some of their subsets can be grouped into certain classes, without any correctness test being performed on acquired knowledge through feedback from external knowledge sources. Note 1 to entry: Once a concept is formed, it is given a name that may be used in subsequent learning of other concepts.	iso_2382_1997						
usability	extent to which a system product or service can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use (note 1: The "specified" users, goals and context of use refer to the particular combination of users, goals and context of use for which usability is being considered; note 2: used as a qualifier to refer to the design knowledge, competencies, activities and design attributes that contribute to usability, such as usability expertise, usability professional, usability engineering, usability method, usability evaluation, usability heuristic). [See also: ISO/IEC 9241-11 Ergonomic of Human-System Interaction — Part 11: Usability: Definitions and Concepts. ISO, Geneva, Switzerland, 2018, https://www.iso.org/standard/63500.html.]	ISO/IEC_TS_5723:2022(en)										
usability testing	refers to evaluating a product or service by testing it with representative users. Typically, during a test, participants will try to complete typical tasks while observers watch, listen and takes notes. The goal is to identify any usability problems, collect qualitative and quantitative data and determine the participant's satisfaction with the product.	Usabilitygov										
user	individual or group that interacts with a system or benefits from a system during its utilization	IEEE_Soft_Vocab	A person, organization, or other entity which requests access to and uses the resources of a computer system or network.	CSRC								
user-centered design	the practice of the following principles, the active involvement of users for a clear understanding of user and task requirements, iterative design and evaluation, and a multi-disciplinary approach	Vredenburg,_Karel	Approach to system design and development that aims to make interactive systems more usable by focusing on the use of the system; applying human factors, ergonomics and usability knowledge and techniques.	IEEE_Soft_Vocab								
validation	Confirmation by examination and provision of objective evidence that the particular requirements for a specific intended use are fulfilled.	UNODC_Glossary_QA_GLP	Confirmation, through the provision of objective evidence, that the requirements for a specific intended use or application have been fulfilled.	IEEE_Soft_Vocab	provides objective evidence that the capability provided by the system complies with stakeholder performance requirements, achieving its use in its intended operational environment; answers the question, "Is it the right solution to the problem?" [C]onsists of evaluating the operational effectiveness, operational suitability, sustainability, and survivability of the system or system elements under operationally realistic conditions.	DOD_TEVV	A continuous monitoring of the process of compilation and of the results of this process.	OECD			Test and Evaluation, Verification, and Validation (TEVV)	
value sensitive design	a theoretically grounded approach to the design of technology that accounts for human values in a principled and systematic manner throughout the design process.	Friedman_et_al_2017										
variable	A variable is a characteristic of a unit being observed that may assume more than one of a set of values to which a numerical measure or a category from a classification can be assigned.	OECD	Quantity or data item whose value can change	IEEE_Soft_Vocab								
variable importance	Represents the statistical significance of each variable in the data in terms of its affect on the model.	H2O.ai_glossary										
variance	The variance is the mean square deviation of the variable around the average value. It reflects the dispersion of the empirical values around its mean.	OECD	A quantifiable deviation, departure, or divergence away from a known baseline or expected value	IEEE_Soft_Vocab								
verifiable	can be checked for correctness by a person or tool	ISO/IEC_TS_5723:2022(en)	provides evidence that the system or system element performs its intended functions and meets all performance requirements listed in the system performance specification and functional and allocated baselines; answers the question, "Did you build the system correctly?"	DOD_TEVV	the goal of designing AI systems that have strong, ideally provable, assurances of correctness with respect to mathematically specified requirements.	Seshia_et_al_2022					Test and Evaluation, Verification and Validation (TEVV)	
washing	a marketing effort designed to imply that a company's brands and products involve artificial intelligence technologies, even though the connection may be tenuous or non-existent.	TechTarget_Ivy_Wigmore										
word embedding	a popular framework to represent text data as vectors which has been used in many machine learning and natural language processing tasks. . . . A word embedding, trained on word co-occurrence in text corpora, represents each word (or common phrase) w as a d-dimensional word vector w~ 2 Rd. It serves as a dictionary of sorts for computer programs that would like to use word meaning. First, words with similar semantic meanings tend to have vectors that are close together. Second, the vector differences between words in embeddings have been shown to represent relationships between words.	Bolukbasi_et_al_Debiasing_Word_Embeddings										

ID	Title of article, chapter, or page	Author(s) and/or Editor(s)	Publication or website (either the main domain or major subdomain)	Volume	Issue	Page(s)	Year	URL	Notes
GDPR	Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)							https://eur-lex.europa.eu/eli/reg/2016/679/oj	
CCPA	California Consumer Privacy Act of 2018							https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5	
AI_Incident_Database	What is an AI incident?	AI Incident Database	AI Incident Database				2022	https://incidentdatabase.ai/research/1-criteria	
Shubendhu_and_Vijay	Applicability of Artificial Intelligence in Different Fields of Life	Shubhendu, Shukla S. and Jaiswal Vijay	International Journal of Scientific Engineering and Research (IJSER)	1	1	28-35	2013	https://www.ijser.in/archives/v1i1/MDExMzA5MTU=.pdf	
Raynor	Glossary of Computer System Software Development Terminology	Raynor, William J., Jr.	The International Dictionary of Artificial Intelligence				1999	https://archive.org/details/internationaldic0000rayn	I access this book using a free Internet Archive account; it can be borrowed in 1-hour stretches This link also can be used: https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.375.8194&rep=rep1&type=pdf
AI_Fairness_360	Glossary	AI Fairness 360	AI Fairness 360					https://aif360.mybluemix.net/resources#glossary	
Mitchell,_Tom	Machine Learning	Mitchell, Tom	Machine Learning				1997	http://www.cs.cmu.edu/~tom/mlbook.html	
Brookings_Institution	The Brookings glossary of AI and emerging technologies	Allen, John R. and Darrell M. West	Brookings Institution				2021	https://www.brookings.edu/blog/techtank/2020/07/13/the-brookings-glossary-of-ai-and-emerging-technologies/	
Brownlee,_Jason	A Gentle Introduction to Generative Adversarial Networks (GANs)	Brownlee, Jason	Machine Learning Mastery				2019	https://machinelearningmastery.com/what-are-generative-adversarial-networks-gans/	
Pyle_and_San_José	An executive's guide to machine learning	Pyle, Dorian and Cristina San José	McKinsey Quarterly				2015	https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/an-executives-guide-to-machine-learning	
Hutson,_Matthew	AI Glossary: Artificial intelligence, in so many words	Hutson, Matthew	Science	357	6346	19	2017	https://www.science.org/doi/10.1126/science.357.6346.19	
FBPML_Wiki	Definitions	Foundation for Best Practices in Machine Learning	FBPML Wiki					https://wiki.fbpml.org/wiki/Definitions	
IAPP_Privacy_Glossary	Glossary of Privacy Terms							https://iapp.org/resources/glossary/	
Reznik,_Leon	Introduction 1.5 Glossary of Basic Terms	Reznik, Leon	Intelligent Security Systems: How Artificial Intelligence, Machine Learning and Data Science Work for and Against Computer Security			xv-xxiv	2022		
IEEE_Guide_IPA	IEEE Guide for Terms and Concepts in Intelligent Process Automation	IEEE Standards Association	IEEE Guide for Terms and Concepts in Intelligent Process Automation						
Russell_and_Norvig		Stuart Russell and Peter Norvig	Artificial Intelligence: A Modern Approach (Fourth Edition)				2021		Definition 1 for "parametric" comes from section 19.7; the page number is missing from the ebook consulted, but the PDF page number is 1269 of 2579.
SP1270	Towards a Standard for Identifying and Managing Bias in Artificial Intelligence	Schwartz, Reva; Apostol Vassilev; Kristen Greene; Lori Perine; Andrew Burt; Patrick Hall	NIST Special Publication 1270					https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf	
SP1011	Autonomy Levels for Unmanned Systems (ALFUS) Framework	Autonomy Levels for Unmanned Systems Working Group Participants	NIST Special Publication 1011				2008	https://www.nist.gov/system/files/documents/el/isd/ks/NISTSP_1011-1-2-0.pdf	
Gartner	Gartner Glossary	Gartner Group						https://www.gartner.com/en/glossary/all-terms	
Varshney,_Kush	Trustworthy Machine Learning	Varshney, Kush R.							
Munir,_Arslan	Artificial Intelligence and Data Fusion at the Edge	Munir, Arslan, Erik Blasch, Jisu Kwon, Joonho Kong, and Alexander Aved	IEEE A&E SYSTEMS MAGAZINE	36	7	62-78	2021	https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=9475883	
Wallace,_Brian	Introduction to Artificial Intelligence for Security Professionals	Wallace, Brian; Sepehr Akhavan-Masouleh; Andrew Davis; Mike Wojnowicz; John H. Brock					2017	http://book.itep.ru/depository/AI/IntroductionToArtificialIntelligenceForSecurityProfessionals_Cylance.pdf	
NSCAI	National Security Commission on Artificial Intelligence: The Final Report	National Security Commission on Artificial Intelligence	National Security Commission on Artificial Intelligence Final Report				2021	https://www.nscai.gov/2021-final-report/	Appendix A: Technical Glossary begins on page 601 of the report (603 of the PDF itself).
OECD	Glossary of Statistical Terms	Organisation for Economic Co-operation and Development					2007	https://ec.europa.eu/eurostat/ramon/coded_files/OECD_glossary_stat_terms.pdf / https://stats.oecd.org/glossary/	
OECD_CAI_recommendation	Recommendation of the Council on Artificial Intelligence	OECD	OECD Legal Instruments				2019	https://legalinstruments.oecd.org/en/instruments/oecd-legal-0449	Adopted on: 21/05/2019
NISTIR_8269_Draft	A Taxonomy and Terminology of Adversarial Machine Learning	Tabassi, Elham;Kevin J. Burns; Michael Hadjimichael; Andres D. Molina-Markham; Julian T. Sexton	Draft NISTIR 8269				2019	https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8269-draft.pdf	
SP800-37	Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy	Joint Task Force Interagency Working Group	NIST Special Publication 800-37 Revision 2		2 (revision 2)		2018	https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-37r2.pdf	
IEEE_Soft_Vocab	Systems and software engineering —Vocabulary		ISO/IEC/IEEE 24765				2017	https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8016712	
Kohavi,_Ron	Glossary of Terms: Special Issue on Applications of Machine Learning and the Knowledge Discovery Process	Kohavi, Ron; Foster Provost	Machine Learning	30		271-274	1998	http://robotics.stanford.edu/~ronnyk/glossary.html	
Mitchell,_Tom	Machine Learning	Mitchell, Tom M.	McGraw-Hill Science/Engineering/Math				1997	https://www.cin.ufpe.br/~cavmj/Machine%20-%20Learning%20-%20Tom%20Mitchell.pdf	
Cyber_Guide	Cyber Security Planning Guide	Federal Communications Commision						https://www.fcc.gov/sites/default/files/cyberplanner.pdf	
CSRC	Information Technology Laboratory Computer Security Resource Center Glossary		NIST					https://csrc.nist.gov/glossary	
AIMA	Artificial Inelligence: A Modern Approach	Russell, Stuart; Peter Norvig	Pearson				2010	https://zoo.cs.yale.edu/classes/cs470/materials/aima2010.pdf	
Breiman_Leo	Bagging Predictors	Breiman, Leo	Machine Learning		24	123-140	1996	https://link.springer.com/content/pdf/10.1007/BF00058655.pdf	
NISTIR_8312	Four Principles of Explainable Artificial 3 Intelligence	Phillips, P. Jonathon; Carina A. Hahn; Peter C. Fontana; David A. Broniatowski; Mark A. Przybocki	Draft NISTIR 8312				2020	https://nvlpubs.nist.gov/nistpubs/ir/2020/NIST.IR.8312-draft.pdf	Full Version Available at https://nvlpubs.nist.gov/nistpubs/ir/2021/NIST.IR.8312.pdf
SP800-12	An Introductin to Information Security	Nieles, Michael; Kelley Dempsey; Victoria Yan Pillitteri	NIST SP 800-12				2017	https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-12r1.pdf	
Steinhardt,_Jacob	Certified Defenses for Data Poisoning Attacks	Steinhardt,_Jacob; Pang Wei Koh; Percy Liang	31st Conference on Neural Information Processing Systems				2017	https://proceedings.neurips.cc/paper/2017/file/9d7311ba459f9e45ed746755a32dcd11-Paper.pdf	
Ranschaert,_Erik	Artificial Intelligence in Medical Imaging: Opportunities, Applications and Risks	Ranschaert, Erik R.; Sergey Morozov; Paul R. Algra	Springer				2019	https://link.springer.com/content/pdf/10.1007/978-3-319-94878-2.pdf	
Blank,_Abagayle_Lee	Computer Vision Machine Learning and Future-Oriented Ethics	Blank, Abagayle Lee	Seattle Pacific University					https://digitalcommons.spu.edu/cgi/viewcontent.cgi?article=1100&context=honorsprojects	
Crawford,_Kate	The Trouble with Bias	Crawford, Kate	Neural Information Processing Systems, Long Beach				2017	https://www.youtube.com/watch?v=fMym_BKWQzk	
COE_AI_Glossary	Artificial Intelligence Glossary		Council of Europe					https://www.coe.int/en/web/artificial-intelligence/glossary	
Kuehn,_Andreas	Analyzing Bug Bounty Programs: An Institutional Perspective on the Economics of Software Vulnerabilities	Kuehn, Andreas; Milton Mueller	2014 TPRC Conference Paper				2014	https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2418812	
Kang,_Daniel	Model Assertions for Monitoring and Improving ML Models	Kang, Daniel; Deepti Raghavan; Peter Baili; Matei Zaharia	3rd MLSys Conference				2020	https://arxiv.org/pdf/2003.01668.pdf	
MathWorks_Residual	What Is Residual Analysis?		MathWorks					https://www.mathworks.com/help/ident/ug/what-is-residual-analysis.html	
Nayak,_Pragati	Concept Drift and Model Decay Detection using Machine Learning Algorithm	Nayak, Pragati Aravind; Pavithra Sriganesh; Rakshitha K.M; Manoj Kumar M.V; Prashanth B S; Sneha H R	12th International Conference on Computing Communication and Networking Technologies (ICCCNT)				2021	https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=9580110	
Hasan,_Raza	Artificial Intelligence Based Model for Incident Response	Hasan, Raza; Salman Mahmood; Akshyadeep Raghav; M. Asim Hasan	2011 International Conference on Information Management, Innovation Management and Industrial Engineering				2011	https://ieeexplore.ieee.org/abstract/document/6114714	
McCue_Colleen	Data Mining and Predictive Analysis: Intelligence Gathering and Crime Analysis	McCue, Colleen	Butterworth-Heinemann				2007	https://www.sciencedirect.com/topics/computer-science/domain-expertise#:~:text=2.1%20Domain%20Expertise.need%20to%20know%20your%20stuff.	
Besse,_Philippe	A Survey of Bias in Machine Learning Through the Prism of Statistical Parity	Besse, Philippe; Eustasio del Barrio; Paula Gordaliza; Jean-Michel Loubes; Laurent Risser	The American Statistician	76	2	188-198	2021	https://www.tandfonline.com/doi/full/10.1080/00031305.2021.1952897	
Muller,_Michael	Designing Ground Truth and the Social Life of Labels	Muller, Michael; Christine T. Wolf; Josh Andres; Michael Desmond; Narendra Nath Joshi; Zahra Ashktorab; Aabhas Sharma; Kristina Brimijoin; Qian Pan; Evelyn Duesterwald; Casey Dugan	Proceedings of the 2021 CHI Conference on Human Factors in Computing System			1-16	2021	https://doi.org/10.1145/3411764.3445402	
Shalev-Shwartz,_Shai	Understanding Machine Learning: From Theory to Algorithms	Shalev-Shwartz, Shai; Shai Ben-David	Cambridge University Press				2014	https://www.cs.huji.ac.il/w~shais/UnderstandingMachineLearning/understanding-machine-learning-theory-algorithms.pdf	
SP800-160	Engineering Trustworthy Secure Systems	Ross, Ron; Mark Winstead; Michael McEvilley	NIST SP 800-160				2022	https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-160v1r1.fpd.pdf	
	Hard choices in artificial intelligence,	Roel Dobbe, Thomas Krendl Gilbert, Yonatan Mintz,	https://www.sciencedirect.com/science/article/pii/S0004370221001065?via%3Dihub						
Khanna,_Anirudh	A Study of Today's A.I. through Chatbots and Rediscovery of Machine Intelligence	Khanna, Anirudh; Bishwajeet, Pandey; Kushagra, Vashishta; Kartik, Kalia; Bhale, Pradeepkumar; Teerath, Das	International Journal of u-and e-Service, Science and Technology	8	7	277-284	2015	http://article.nadiapub.com/IJUNESST/vol8_no7/28.pdf	
NBSIR_82-2582	An Overview of Computer Vision	Gevarter, William B.	NIST 82-2582				1982	https://nvlpubs.nist.gov/nistpubs/Legacy/IR/nbsir82-2582.pdf	
Kou,_Yufeng	Survey of fraud detection techniques	Kou, Yufeng; Chang-Tien, Lu; Sirirat, Sirwongwattana; Yo-Ping, Huang	IEEE International Conference on Networking, Sensing and Control	2		749-754	2004	https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=1297040	
Merritt,_Linda	Human capital management: More than HR with a new name.	Merritt, Linda	*People and Strategy*	30	2	14-16	2007	https://www.proquest.com/docview/224596559?pq-origsite=gscholar&fromopenview=true	
CRS_AI	Artificial Intelligence: Background, Selected Issues, and Policy Considerations	Harris, Laurie A.	Congressional Research Service				2021	https://crsreports.congress.gov/product/pdf/R/R46795	
IJMAE	A Review of the Role of Marketing in Recruitment and Talent Acquisition.	Alashmawy, Ahmad; Rashad, Yazdanifard	International Journal of Management, Accounting and Economics	6	7		2019	https://www.academia.edu/41261685/A_Review_of_the_Role_of_Marketing_in_Recruitment_and_Talent_Acquisition?from=cover_page	
Das,_Debashis	A survey on recommendation system.	Das, Debashis; Laxman, Sahoo; Sujoy, Datta.	International Journal of Computer Applications	160	7		2017	https://www.researchgate.net/profile/Debashis-Das-17/publication/313787463_A_Survey_on_Recommendation_System/links/5d7de0474585155f1e4de908/A-Survey-on-Recommendation-System.pdf	
Hyndman,_Rob	Forecasting: principles and practice	Hyndman, Rob J.; George, Athanasopoulos				14	2018	https://books.google.com/books?hl=en&lr=&id=_bBhDwAAQBAJ&oi=fnd&pg=PA7&dq=Forecasting&ots=Tij_xmXIMJ&sig=c8LjAcmbLDC5QeH_xQno2I_gTr0#v=onepage&q=Forecasting&f=false	
Ustun,_Berk	Actionable Recourse in Linear Classification	Ustun, Berk; Spangher, Alexander; Liu, Yang	Association of Computing Machinery						
Voight,_Paul	The EU General Data Protection Regulation: A Practical Guide	Voight,_Paul; von_dem_Bussche,_Axel	Springer						
ECOA	12 CFR Part 1002 - Equal Credit Opportunity Act (Regulation B)	Consumer Financial Protection Bureau							
Gill,_Navdeep	A Responsible Machine Learning Workflow with Focus on Interpretable Models, Post-hoc Explanation, and Discrimination Testing	Gill, Navdeep; Hall, Patrick; Montgomery, Kim; Schmidt, Nicholas	MDPI						
ISO_IEC_38507	Information Technology - Governance of IT - Governance implications of the use of artificial intelligence by organizations	ISO/IEC	ISO		1		2022	https://www.iso.org/obp/ui/#iso:std:iso-iec:38507:ed-1:v1:en	
Darnell_Coss_Hall	The Future of Analytics	Dan Darnell, Rafael Coss, Patrick Hall	O'Reilly Media Inc.			Ch. 4	2020	https://www.oreilly.com/library/view/the-future-of/9781492091769/ch04.html	
Fed_Reserve	Guidance on Model Risk Management	Patrick M. Parkinson	The Federal Reserve				2011	https://www.federalreserve.gov/supervisionreg/srletters/sr1107a1.pdf	
Jennifer,_Hill	Causal Inference: Overview	Jennifer Hill, Elizabeth A. Stuart	International Encyclopedia of the Social & Behavioral Sciences (Second Edition)			255-260	2015	https://www.sciencedirect.com/science/article/pii/B9780080970868420957	
AIRS_Penn	Artificial Intelligence Risk & Governance	AIRS	The Warton School, University of Pennsylvania				2022	https://ai.wharton.upenn.edu/artificial-intelligence-risk-governance/	
ENISA	Privacy by Design		European Union Agency for Cybersecurity					https://www.enisa.europa.eu/topics/data-protection/privacy-by-design	
Furche,_Tim	Data wrangling for big data: Challenges and opportunities.	Furche, Tim; George, Gottlob; Leonid, Libkin; Giorgio, Orsi; Norman, Paton	Advances in Database Technology—EDBT 2016: Proceedings of the 19th International Conference on Extending Database Technology			473-478	2016	https://www.research.manchester.ac.uk/portal/files/50447231/paper_94_1_.pdf	
Enrique	Towards an integrated crowdsourcing definition.	Estellés-Arolas, Enrique; Fernando, González-Ladrón-de-Guevara	Journal of Information science	38	2	189-200	2012	https://journals.sagepub.com/doi/full/10.1177/0165551512437638	
Behdad	Nature-inspired techniques in the context of fraud detection.	Behdad, Mohammad; Luigi, Barone; Mohammed, Bennamoun; Tim, French.	EEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)	42	6	1273-1290	2012	https://ieeexplore.ieee.org/abstract/document/6392447?casa_token=-mPJh2Do05MAAAAA:c5Eyg4___i64XlBowloZbwwwKNe0KnqAQvjg5romygO6ymRKDl70np3DiDL0mipMXmdmeboDj1hGwa9I	
Woodward	Biometrics: A look at facial recognition.	Woodward Jr, John D., Christopher Horn, Julius Gatune, and Aryn Thomas	RAND CORP SANTA MONICA CA				2003	https://apps.dtic.mil/sti/citations/ADA414520	
Sharma,_Lalita	A survey of recommendation system: Research challenges.	Sharma, Lalita; Anju, Gera	International Journal of Engineering Trends and Technology (IJETT)	4	5	1989-1992	2013	https://d1wqtxts1xzle7.cloudfront.net/38584474/IJETT-V4I5P132_1_-with-cover-page-v2.pdf?Expires=1656295756&Signature=QAnv0QzAEvSK2TSOMFygnk1AGUQrGWtz7PD7~bal6kX9SbiWAwd18feH3nB0kH4EQDA0rEPGrm3V9E38s4eY2C5i52NHA-jY0h2zBYAnLETt5PB1cAMjSZf4qaqNWZCPvQtOpaxTkhV6YB1MLDTmaFCAjKcqQfqm6WPQoax5VLh0hIbF-hlc8p5wXiW7fE9z~OeYpmTJld4doW94bG3OrDUY75EN9cptH-IHlldQRs~jF1POAkJUO5PL91PKVJatYAjeWQn7eNchB2TszrfnDbs6pjCyPXIIvi1WdnQzkKVN0F6N-CA~YtZ6yGKIxTjDuxdcJh0AtfD2cjxLDvNFCA___&Key-Pair-Id=APKAJLOHF5GGSLRBV4ZA	
Žliobaitė_Indrė	A survey on measuring indirect discrimination in machine learning	Žliobaitė, Indrė	CoRR				2018	https://arxiv.org/abs/1511.00148	
Vorobeychik	Adversarial machine learning	Vorobeychik, Yevgeniy; Murat, Kantarcioglu	Synthesis Lectures on Artificial Intelligence and Machine Learning	12	3		2018	https://www.morganclaypool.com/doi/abs/10.2200/S00861ED1V01Y201806AIM039?casa_token=bcr3UzYRz6AAAAAA:M-sh4ANwQuRFYcXk18O4x_x6zu7Qq3P5ZC12MrCgTckeNFm9sXOCAkAFvHtMce7t1A3lpUt4MFA	
Zhang,_Yonggang	Principal component adversarial example	Zhang,_Yonggang; Xinmei, Tian; Ya, Li; Xinchao, Wang; Dacheng, Tao	IEEE Transactions on Image Processing 29			4804-4815	2020	https://ieeexplore.ieee.org/abstract/document/9018372?casa_token=MFIULJJsFBEAAAAA:GVCQa1Ccivt9ETPlYlBmLvK74iIDNzODP-6mo2B4o4nkZG69ORI-HliXK0iXb3GG-Q8csl2vxtBorQ	
NIST_SP_800	Sp 800-66 rev. 1. an introductory resource guide for implementing the health insurance portability and accountability act (hipaa) security rule	Scholl, Matthew A.; Kevin M. Stine; Joan Hash; Pauline Bowen; L. Arnold Johnson; Carla Dancy Smith; Daniel I. Steinberg	National Institute of Standards & Technology				2008	https://dl.acm.org/doi/pdf/10.5555/2206281	
Jain_Saachi	Missingness Bias in Model Debugging	Jain, Saachi; Hadi Salman; Eric Wong; Pengchuan Zhang; Vibhav Vineet; Sai Vemprala; Aleksander Madry	ICLR				2022	https://arxiv.org/abs/2204.08945	
Mehrabi,_Ninareh	A Survey on Bias and Fairness in Machine Learning	Mehrabi, Ninareh ; Fred Morstatter; Nripsuta Saxena; Kristina Lerman; Aram Galstyan	ACM Computing Surveys	54	6	1-35	2022	https://dl.acm.org/doi/abs/10.1145/3457607	
Lipton,_Zachary	Does mitigating ML's impact disparity require treatment disparity?	Lipton, Zachary C.; Alexandra Chouldechova; Julian McAuley	32nd Conference on Neural Information Processing Systems				2018	https://doi.org/10.48550/arXiv.1711.07076	
Gustavii,_Ebba	A Swedish Grammar for Word Prediction	Gustavii, Ebba; Eva Pettersson	Uppsala University - Department of Linguistics				2003	https://www.researchgate.net/publication/2838153_A_Swedish_Grammar_for_Word_Prediction/link/00b4951a5f2645ca02000000/download	
Comptroller_Office	Comptroller's Handbook: Model Risk Management, Version 1.0	Office of the Comptroller of the Currency	Comptroller's Handbook: Model Risk Management, Version 1.0				2021	https://www.occ.gov/publications-and-resources/publications/comptrollers-handbook/files/model-risk-management/index-model-risk-management.html	
CISA	A Glossary of Common Cybersecurity Words and Phrases	National Initiative for Cybersecurity Careers and Studies					2022	https://niccs.cisa.gov/cybersecurity-career-resources/glossary#D	
GWUC	Cyber Glossary	The George Washington University	Natural Security Archive - The Cyber Vault Project				2018	https://nsarchive.gwu.edu/cyber-glossary-b	
SP-800-12	Sp 800-12. an introduction to computer security: The nist handbook.	Guttman, Barbara; Edward A. Roback					1995	https://dl.acm.org/doi/pdf/10.5555/2206203	
Thomas_Edgar	Research Methods for Cyber Security	Thomas W. Edgar; David O. Manz				367-392	2017	https://www.sciencedirect.com/science/article/pii/B9780128053492000157	
Vredenburg,_Karel	A survey of user-centered design practice	Vredenburg, Karel; Ji-Ye, Ma; Paul W., Smith; Tom, Carey	*Proceedings of the SIGCHI conference on Human factors in computing systems*			471-478	2002	https://dl.acm.org/doi/abs/10.1145/503376.503460	
ST04-015	Understanding Denial-of-Service Attacks		CISA					https://www.cisa.gov/uscert/ncas/tips/ST04-015	
Chandrasekaran,_Varun		Chandrasekaran,_Varun; Kamalika, Chaudhuri; Irene, Giacomelli; Somesh, Jha; Songbai, Yan	*29th USENIX Security Symposium (USENIX Security 20)*			1309-1326	2020	https://www.usenix.org/conference/usenixsecurity20/presentation/chandrasekaran	
FDA_Glossary	Glossary of Computer System Software Development Terminology		FDA				1995	https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/inspection-guides/glossary-computer-system-software-development-terminology-895#_top	
Meinshausen,_Nicolai	Causality from a distributional robustness point of view	Meinshausen, Nicolai	*2018 IEEE Data Science Workshop (DSW)*			6-10	2018	https://ieeexplore.ieee.org/abstract/document/8439889?casa_token=8tga5hcskjQAAAAA:PlpvmJtJNGKSNXiBdghGmy670MxN9IPXc3ekZRVEXwOcLoJkh0sxDuWseVzr0EowGRCf8WR-eYLWeU4	

ID	Title of article, chapter, or page	Author(s) and/or Editor(s)	Publication or website (either the main domain or major subdomain)	Volume	Issue	Page(s)	Year	URL	Notes
Stacke,_Karin	Measuring domain shift for deep learning in histopathology.	Stacke, Karin; Gabriel, Eilertsen; Jonas, Unger; Claes, Lundström	*IEEE journal of biomedical and health informatics*	25	2	325-336	2020	https://ieeexplore.ieee.org/abstract/document/9234592?casa_token=RIQcYxte8IIAAAAA:4t6lwNomEw95f3c1ir73BRReG7OzKecUzpVQS_Bk5zIEWA5R75uG-66g9irlblzDDVwu7ut4jAo2i_8	
NASA_Soft_Standards	Software Assurance and Software Safety Standard		*NASA Standard*				2020	https://standards.nasa.gov/standard/nasa/nasa-std-87398	
PET_Handbook	Handbook of privacy and privacy-enhancing technologies.	Van Blarkom, G. W.; John J. Borking; JG Eddy Olk	Privacy Incorporated Software Agent (PISA) Consortium, The Hague 198				2003	https://andrewpatrick.ca/pisa/handbook/Handbook_Privacy_and_PET_final.pdf	
NIST_1500	NIST Big Data Interoperability Framework	Wo L. Chang; Nancy Grady	NIST	1			2019	https://www.nist.gov/publications/nist-big-data-interoperability-framework-volume-1-definitions?pub_id=918927	
Microsoft_Azure_documentation_Detect_data_drift	Detect data drift (preview) on datasets	Microsoft	Azure documentation				2022	https://docs.microsoft.com/en-us/azure/machine-learning/how-to-monitor-datasets?tabs=python	
Egnyte	Egnyte	Egnyte	Data Control: Definition and Benefits				2022	https://www.egnyte.com/guides/governance/data-control	
IG1190M_AIOps_Decommission_v1.0.0	IG1190M AIOps Decommission v1.0.0	AI Operations	IG1190M AIOps Decommission v1.0.0				2022	https://www.tmforum.org/resources/reference/ig1190m-aiops-decommission-v1-0-0/	
peak.ai	AI decision making: the future of business intelligence	Peak.AI	AI decision making: the future of business intelligence				2022	https://peak.ai/hub/blog/ai-decision-making-the-future-of-business-intelligence/	
Hochreiter,_Sepp	Toward a Broad AI	Hochreiter, Sepp	Communications of the ACM	65	4	56-57	2022	https://cacm.acm.org/magazines/2022/4/259402-toward-a-broad-ai/fulltext	
Mitchell,_Eric	Memory-Based Model Editing at Scale.	Mitchell, Eric; Charles, Lin; Antoine, Bosselut; Christopher D., Manning; Chelsea, Finn	International Conference on Machine Learning, PMLR			15817-15831	2022	https://proceedings.mlr.press/v162/mitchell22a.html	
ISO/IEC_TS_5723:2022(en)	ISO/IEC TS 5723:2022(en) Trustworthiness — Vocabulary	ISO/IEC	ISO/IEC				2022	https://www.iso.org/obp/ui/#iso:std:iso-iec:ts:5723:ed-1:v1:en	
H20.ai_glossary	H20.ai Glossary	H20.ai	H20.ai				2022	https://docs.h2o.ai/h2o/latest-stable/h2o-docs/glossary.html	
TechTarget_Ivy_Wigmore	AI Washing	Ivy Wigmore	TechTarget				2017	https://www.techtarget.com/searchenterpriseai/definition/AI-washing	
Roman_V._Yampolskiy_Unexplainability	Unexplainability and Incomprehensibility of Artificial Intelligence	Roman V. Yampokskiy	PhilArchive				2019	https://philarchive.org/archive/YAMUAI	PhilArchive is the website, not the publication.
Forbes_Kayvan_Alikhani	Remember 'Cloud Washing'? It's Happening In RegTech	Kayvan Alikhani	Forbes				2019	https://www.forbes.com/sites/forbestechcouncil/2019/10/14/remember-cloud-washing-its-happening-in-regtech/?sh=2a52bd3a796c	
Ripley,_Brian	Introduction and Examples	Brian D. Ripley	*Pattern Recognition and Neural Networks*			6	1996	https://archive.org/details/patternrecogniti0000ripl/page/6/mode/2up?q=training&view=theater	
C3.ai_Model_Training	Model Training	C3.ai	C3.ai Glossary					https://c3.ai/glossary/data-science/model-training/	
Pyle,_Dorian_Data_Preparation_as_a_Process	Data Preparation as a Process	Dorian Pyle	*Data Preparation for Data Mining*			89-124	1999	https://www.google.com/books/edition/Data_Preparation_for_Data_Mining/hhdVr9F-JfAC?hl=en&gbpv=1&dq=Binning%20is%20a%20technique&pg=PA110&printsec=frontcover	"Binning" definition on page 110
Wieringa,_Roel_J.	Conceptual Frameworks	Roel J. Wieringa	*Design Science Methodology for Information Systems and Software Engineering*			73-92	2014	https://www.google.com/books/edition/Design_Science_Methodology_for_Informati/xLKLBQAAQBAJ?hl=en&gbpv=1&dq=%22Construct+validity+is+defined+by+Shadish+et+al+24+p+506+as+the+degree+to+which+inferences+from+phenomena+to+constructs+are+warranted%22&pg=PA87&printsec=frontcover	"Construct validity" definition on page 87
York,_Dan_Internet_Society	What Is the Splinternet? And Why You Should Be Paying Attention	Dan York	Internet Society				2022	https://www.internetsociety.org/blog/2022/03/what-is-the-splinternet-and-why-you-should-be-paying-attention/	March 23, 2022
Splinternets	Key Definitions	European Parliament, Directorate-General for Parliamentary Research Services, Perarnaud, C., Rossi, J., Musiani, F., et al.	*'Splinternets': Addressing the Renewed Debate on Internet Fragmentation*				2022	https://op.europa.eu/en/publication-detail/-/publication/5a5bfaed-0d52-11ed-b11c-01aa75ed71a1/language-en	"Splinternet" definition on page 4
Regularization_for_Deep_Learning	Regularization for Deep Learning	Ian Goodfellow, Yoshua Bengio, and Aaron Courville	*Deep Learning*			221-266	2016	https://www.google.com/books/edition/Deep_Learning/omivDQAAQBAJ?hl=en&gbpv=1&dq=%22Sparsity+in+this+context+refers+to+the+fact+that+some+parameters+have+an+optimal+value+of+zero%22&pg=PA229&printsec=frontcover	"Sparsity" definition on page 229
Statistics_in_Plain_English	Statistical Significance and Effect Size	Timothy C. Urdan	*Statistics in Plain English*			43-56	2001	https://www.google.com/books/edition/Statistics_in_Plain_English/MwCK52n_wBEC?hl=en&gbpv=1&dq=%22Statistical+significance+When+the+probability+of+obtaining+a+statistic+of+a+given+size+due+strictly+to%22&pg=PA55&printsec=frontcover	"Statistical significance" definition on page 55
The_SAGE_Encyclopedia_of_Communication_Research_Methods	Relationships Between Variables	Mike Allen, ed.	*The SAGE Encyclopedia of Communication Research Methods, Volume 1*				2017	https://www.google.com/books/edition/The_SAGE_Encyclopedia_of_Communication_R/4GFCDgAAQBAJ?hl=en&gbpv=1&dq=%22Statistical+significance+refers+to+whether+a+relationship+between+two+or+more+variables+exists+beyond+a+probability+expected+by+chance%22&pg=PA1413&printsec=frontcover	"Statistical significance" definition on page 1413
The_Science_of_Algorithmic_Trading_and_Portfolio_Management	Evaluation	Robert Kissell	*The Science of Algorithmic Trading and Portfolio Management*				2013	https://www.google.com/books/edition/The_Science_of_Algorithmic_Trading_and_P/FKPND2zz9OoC?hl=en&gbpv=1&dq=%22Back+testing+is+the+quantitative+evaluation+of+a+model%E2%80%99s+performance+both+from+a+statistical+and+trading+perspective%22&pg=PA446&printsec=frontcover	"Back test" definition on page 446 (as "back-testing")
Introduction_to_Information_Systems	Artificial Intelligence	R. Kelly Rainer, Reiner R. Kelly, and Brad Prince	*Introduction to Information Systems: International Adaptation*			410-450	2022	https://www.google.com/books/edition/Introduction_to_Information_Systems/Y75VEAAAQBAJ?hl=en&gbpv=1&dq=%22An+autonomous+vehicle+automobile+bus+tractor+combine+boat+forklift+etc+is+a+vehicle+capable+of+sensing+its+environment+and+moving+safely+with+little+or+no+human+input%22&pg=PA437&printsec=frontcover	"Autonomous vehicle" definition on page 437
OED_snake_oil	Snake oil	Oxford English Dictionary	Oxford English Dictionary				2022	https://www-oed-com.proxygw.wrlc.org/view/Entry/95490133?redirectedFrom=%22snake+oil%22	June 2022 is the citation date
Nick_Higham_1	What Is Rounding?	Nick Higham	Nick Higham				2020	https://nhigham.com/2020/04/28/what-is-rounding/	04/28/2020
Nick_Higham_2	What Is Stochastic Rounding?	Nick Higham	Nick Higham				2020	https://nhigham.com/2020/07/07/what-is-stochastic-rounding/	07/07/2020
pwc_Model_Risk_Management_of_AI_and_ML_Systems	Model Risk Management of AI and Machine Learning Systems	PricewaterhouseCoopers	Model Risk Management of AI and Machine Learning Systems				2020	https://www.pwc.co.uk/data-analytics/documents/model-risk-management-of-ai-machine-learning-systems.pdf	"risk tiering" definition on page 11
Toward_an_understanding_of_responsible_artificial_intelligence_practices	Toward an understanding of responsible artificial intelligence practices	Yichuan Wang, Mengran Xiong, and Hossein G. T. Olya	In: Bui, T.X., (ed.) Proceedings of the 53rd Hawaii International Conference on System Sciences. Hawaii International Conference on System Sciences (HICSS 2020), 20202-01-07 - 2020-01-10, Maui, Hawaii, USA. Hawaii International Conference on System Sciences (HICSS)				2020	https://eprints.whiterose.ac.uk/162719/8/Toward%20an%20Understanding%20of%20Responsible%20Artificial%20Intelligence%20Practices.pdf	"risk control" definition on page 4967
Comparing_scores_and_reason_codes	Comparing scores and reason codes in credit scoring systems: NeuroDecision vs. unconstrained neural networks	Equifax	Comparing scores and reason codes in credit scoring systems: NeuroDecision vs. unconstrained neural networks				2020	https://assets.equifax.com/marketing/US/assets/comparing_scores_whitepaper.pdf	"reason code" definition on page 3
Machine_Learning_Interpretability_with_H20_Driverless_AI	Machine Learning Interpretability with H20 Driverless AI	Patrick Hall, Navdeep Gill, Megan Kurka, and Wen Phan; edited by Angela Bartz	Machine Learning Interpretability with H20 Driverless AI				2022	https://docs.h2o.ai/driverless-ai/latest-stable/docs/booklets/MLIBooklet.pdf	"reason code" definition on page 6
Pattern_Recognition_and_Machine_Learning	Introduction	Christopher M. Bishop	*Pattern Recognition and Machine Learning*			1-66	2006		"recognition" definition on page 1
DOD_TEVV	Technology Investment Strategy 2015-2018	United States Department of Defense's Test and Evaluation, Verification and Validation (TEVV) Working Group	Technology Investment Strategy 2015-2018				2015	https://defenseinnovationmarketplace.dtic.mil/wp-content/uploads/2018/02/OSD_ATEVV_STRAT_DIST_A_SIGNED.pdf	"trust" definition on page 15; "automation" and "autonomy" definitions on page 2; "validation" and "verification" definitions on page 15
Fundamentals_of_Information_Systems_Security	Auditing, Testing, and Monitoring	David Kim and Michael G. Solomon	Fundamentals of Information Systems Security			216-250	2016	https://www.google.com/books/edition/Fundamentals_of_Information_Systems_Secu/Yb4eDQAAQBAJ?hl=en&gbpv=1&dq=%22Audit+logs+Defined+events+that+provide+additional+input+to+audit+activities%22&pg=PA233&printsec=frontcover	"audit log" definition on page 233
AI_Ethics_Mark_Coeckelbergh	Glossary	Mark Coeckelbergh	*AI Ethics*			203-206	2020	https://www.google.com/books/edition/AI_Ethics/Gs_XDwAAQBAJ?hl=en&gbpv=1&dq=%22Trustworthy+AI+AI+that+can+be+trusted+by+humans%22&pg=PA206&printsec=frontcover	"trustworthy AI" definition on page 206; "bias" definition on page 203; "data science" definition on page 203; "deep learning" definition on page 203; "ethics by design" definition on page 203; "explainability" definition on page 204; "explainable artificial intelligence (XAI)" (as "Explainable AI") definition on page 204; "artificial general intelligence (AGI)" definition on page 204 (as "General AI"); "moral agency" definition on page 204; "moral patiency" definition on page 204
Hands-On_Smart_Contract_Dev	Fundamental Concepts of Blockchain	Matt Zand, Xun (Brian) Wu, and Mark Anthony Morris	*Hands-On Smart Contract Development with Hyperledger Fabric V2*						"application programming interface (API)" definition on page 17
TechTarget_target_function	target function	TechTarget	TechTarget				2018	https://www.techtarget.com/whatis/definition/target-function	We could potentially find a more accurate definition for "target"; this is specifically "target function" (as of this inclusion, the page was last updated June 2018). The author is just listed as TechTarget Contributor.
IGI_Global_reference_class	What is Reference Class	IGI Global	IGI Global					https://www.igi-global.com/dictionary/reference-class/35564	This entry cites the Encyclopedia of Information Science and Technology, Second Edition (8 Volumes) as (possibly?) where it is deriving this definition.
CPO_Magazine_Amar_Kanagaraj	Data Remediation and Its Role in Data Security and Privacy	CPO Magazine	CPO Magazine				2022	https://www.cpomagazine.com/data-protection/data-remediation-and-its-role-in-data-security-and-privacy/	February 1, 2022; the definition of "remediation"
DEV_ranking	Machine learning (ML) applications: ranking	DEV Community	DEV Community				2022	https://dev.to/mage_ai/machine-learning-ml-applications-ranking-238d	Definition for "ranking" taken from here; last updated March 17, 2022 on date accessed.
productmanagerHQ_Josh_Fechter	What Does an AI Product Manager Do?	Product Manager HQ	Product Manager HQ					https://productmanagerhq.com/ai-product-manager/	Definition for "product manager" taken from here; last updated on December 22, 2021 on date accessed.
Proxy_Discrimination	Proxy Discrimination in the Age of Artificial Intelligence and Big Data	Anya E. R. Prince; Daniel Schwarcz	Iowa Law Review	105	3		2020	https://ilr.law.uiowa.edu/print/volume-105-issue-3/proxy-discrimination-in-the-age-of-artificial-intelligence-and-big-data	105 Iowa L. Rev. 1257 (2020)
Forbes_Tracy_Kemp	Four Skills Every Successful AI Product Owner Should Possess	Tracy Kemp	Forbes				2021	https://www.forbes.com/sites/forbestechcouncil/2021/08/24/four-skills-every-successful-ai-product-owner-should-possess/?sh=65dbfe423d3d	Definition for "product owner"; August 24, 2021
Merriam-Webster_underrepresented	underrepresented	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/underrepresented	See the note for definition 1 of this term. The entry is in the nominal form (underrepresentation), but the (circular) Merriam-Webster definition is in the adjectival form. I am using the Merriam-Webster definition because the OED as of yet does not have an entry for this term. The OED has quotations from the late 1800s for "under-represent" and "under-representation," but no discussion of the term beyond those citations.
HBR_Andrew_Burt_how_to_ensure	How to Ensure Your AI Doesn't Discriminate	Andrew Burt	Harvard Business Review					https://hbr.org/2020/08/how-to-ensure-your-ai-doesnt-discriminate	August 28, 2020
Cadient_EEOC	Understanding and Avoiding Adverse Impact in Employment Practices	Michael Baysinger; Kristin Worrell	Cadient					https://cadienttalent.com/resources/understanding-and-avoiding-adverse-impact	No date provided. This citation is being used because it cites verbatim an EEOC definition for "adverse impact ratio," but that EEOC source does not seem readily locatable. As a result, I have highlighted in light yellow definition 1 for "adverse impact ratio" because we will need to work to find this original EEOC source. This Cadient report is fine for now, but it should not stand as the final citation for this definition; the original EEOC one should.
Ben_Green_Yiling_Chen	Algorithm-in-the-Loop Decision Making	Ben Green; Yiling Chen	The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)					https://ojs.aaai.org/index.php/AAAI/article/view/7115	Added citation Ben_Green_Yiling_Chen; added the following note: Definition 1 for "algorithm-in-the-loop" is taken from the wording from the abstract, with everything after the first ellipsis being taken from paragraphs 2, and everything after the second ellipsis being taken from paragraph 3.
IT_Governance_Blog_Luke_Irwin	Personal data vs Sensitive Data: What's the Difference?	Luke Irwin	IT Governance Blog					https://www.itgovernance.co.uk/blog/the-gdpr-do-you-know-the-difference-between-personal-data-and-sensitive-data	
OED_stereotype	Stereotype	Oxford English Dictionary	Oxford English Dictionary					https://www-oed-com.proxygw.wrlc.org/view/Entry/189956?rskey=ImZ8YE&result=1#eid	Definition 3b.
Machine_Learning_Mastery_Jason_Brownlee	Machine Learning Terminology from Statistics and Computer Science	Jason Brownlee	Machine Learning Mastery				2016	https://machinelearningmastery.com/data-terminology-in-machine-learning/#:~:text=Row%3A%20A%20row%20describes%20a,problem%20domain%20that%20you%20have.	Originally published March 9, 2016, but updated on August 8, 2019.
Banking_on_Your_Data_Christopher_Gilliard	Prepared Testimony and Statement for the Record of Christopher Gilliard, PhD. Hearing on "Banking on Your Data: the Role of Big Data in Financial Services" before the House Financial Services Committee Task Force on Financial Technology	Christopher Gilliard	Prepared Testimony and Statement for the Record of Christopher Gilliard, PhD. Hearing on "Banking on Your Data: the Role of Big Data in Financial Services" before the House Financial Services Committee Task Force on Financial Technology					https://www.congress.gov/116/meeting/house/110251/witnesses/HHRG-116-BA00-Wstate-GillardC-20191121.pdf	See also https://www.congress.gov/event/116th-congress/house-event/110251 for details about the hearing.
Merriam-Webster_pseudoscience	pseudoscience	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/pseudoscience	
Cost_Management_ch15	Lean Accounting and Productivity Measurement	Don R. Hansen; Maryanne M. Mowen; Dan L. Heitger	*Cost Management*				2021	https://www.google.com/books/edition/Cost_Management/HhQcEAAAQBAJ?hl=en&gbpv=1&dq=%22Velocity+has+to+do+with+how+fast+a+product+can+be+delivered+to+the+market+and+quality+is+concerned+with+providing+a+nondefective+product+with+the+desired+features+to+customers%22&pg=PA783&printsec=frontcover	Definition 1 for "product velocity" appears on page 783 in the "Big Data" explanation box
Towards_Productizing	Towards Productizing AI/ML Models: An Industry Perspective from Data Scientists	Filippo Lanubile, Fabio Calefato, Luigi Quaranta, Maddalena Amoruso, Fabio Fumarola, and Michele Filannino	arXiv					https://arxiv.org/abs/2103.10548	Definition 1 for "productization" appears in the first paragraph of the paper as "productizing"; I have changed around the grammar a bit in the definition within the brackets.
The_Art_of_Software_Testing	The Psychology and Economics of Program Testing	Glenford J. Myers	*The Art of Software Testing*				1979	https://archive.org/details/artofsoftwaretes0000myer/page/4/mode/2up	Definition 1 for "test" on page 5.
William_Hetzel	An Introduction	William C. Hetzel	*The Complete Guide to Software Testing*, 2nd edition				1988	https://archive.org/details/completeguidetos0000hetz/page/6/mode/2up?view=theater	Definition 2 for "test" on page 6.
On_Hyperparameter_Optimization	On Hyperparameter Optimization of Machine Learning Algorithms: Theory and Practice	Li Yang and Abdallah Shami	arXiv					https://arxiv.org/pdf/2007.15745.pdf	
Security_Analysis_of_Subject_Access	Security Analysis of Subject Access Request Procedures: How to Authenticate Data Subjects Safely When They Request for Their Data	Coline Boniface, Imane Fouad, Nataliia Bielova, Cédric Lauradoux, and Cristiana Santos	*Privacy Technologies and Policy (7th Annual Privacy Forum, APF 2019, Rome, Italy, June 13–14, 2019, Proceedings)*				2019	https://www.google.com/books/edition/Privacy_Technologies_and_Policy/SW2cDwAAQBAJ?hl=en&gbpv=1&dq=%22i+Impersonation+data+breach+A+malicious+individual+is+able+to+impersonate+a+legitimate+data+subject+to+the+data+controller%22&pg=PA186&printsec=frontcover	Definition 1 for "impersonation" on page 186
IEEE_Caught_in_the_Act	Caught in the Act of an Insider Attack: Detection and Assessment of Insider Threat	Philip A. Legg, Oliver Buckley, Michael Goldsmith, and Sadie Creese	*2015 IEEE International Symposium on Technologies for Homeland Security (HST*				2015	https://ieeexplore-ieee-org.proxygw.wrlc.org/stamp/stamp.jsp?tp=&arnumber=7446229	Definition 1 for "insider attack" taken from the abstract of this article.
Moradi_Samwald	Post-hoc explanation of black-box classifiers using confident itemsets	Milad Moradi; Matthias Samwald	*Expert Systems with Applications*	165			2021	https://reader.elsevier.com/reader/sd/pii/S0957417420307302?token=858E09A321B5727ECF889090AFB0B943768A26AFE15FB920B20176BD22F82449CE42841047FD99C51A01659C2ECE9695&originRegion=us-east-1&originCreation=20220918190335#b0035	Definition 1 for "post-hoc explanation" appears at the bottom of page 2 and continues into page 3. (In a very strict sense, the first sentence should be sufficient.) The paragraph is reproduced in full in the note for the term's definition.
Mind_on_Statistics	Chapter 4	Jessica M. Utts and Robert F. Heckard	*Mind on Statistics (6th Edition)*				2021	https://www.google.com/books/edition/Mind_on_Statistics/npQMEAAAQBAJ?hl=en&gbpv=1&dq=%22Practical+Versus+Statistical+Significance+Statistical+significance+does+not+necessarily+mean+that+the+relationship+between+the+two+variables+has+practical+significance%22&pg=PA138&printsec=frontcover	Definition 1 for "practical significance" is on page 138.
Jenna_Burrell	How the machine 'thinks': Understanding opacity in machine learning algorithms	Jenna Burrell	*Big Data & Society*			1-12	2016	https://journals.sagepub.com/doi/pdf/10.1177/2053951715622512	
Hughes_Lavery_Critical_Thinking	Glossary	William Hughes and Jonathan Lavery	*Critical Thinking - Concise Edition*				2015	https://www.google.com/books/edition/Critical_Thinking_Concise_Edition/k0idCgAAQBAJ?hl=en&gbpv=1&dq=%222+2+2+straw+man+a+fallacious+argument+which+irrelevantly+attacks+a+position+that+appears+similar+to+but+is+actually+different+from+an+opponent%E2%80%99s+position+and+concludes+that+the+opponent%E2%80%99s+real+position+has+thereby+been+refuted%22&pg=PA271&printsec=frontcover	Definition 1 on page 271
NIST_SP_800-30_Rev_1	NIST Special Publication 800-30 Revision 1: Guide for Conducting Risk Assessments	NIST	NIST Special Publication 800-30 Revision 1: Guide for Conducting Risk Assessments				2012	https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf	Definition 1 for "technical control" appears on page Page B-13 of Appendix B (63 of 95 of the PDF itself).
TechTarget_third_party	third party	TechTarget	TechTarget Glossary				2014	https://www.techtarget.com/whatis/definition/third-party	
Law_Insider_processing_environment	Processing Environment	Insider Law						https://www.lawinsider.com/dictionary/processing-environment	
wachter_counterfactual_2018	"Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GDPR."	Wachter, S., B. D. M. Mittelstadt, and C. Russell.	Harvard Journal of Law and Technology	31	2		2018	https://ora.ox.ac.uk/objects/uuid:86dfcdac-10b5-4314-bbd1-08e6e78b9094	

ID	Title of article, chapter, or page	Author(s) and/or Editor(s)	Publication or website (either the main domain or major subdomain)	Volume	Issue	Page(s)	Year	URL	Notes
mills_study_2010	Proposed Internet Congestion Control Mechanisms.	Mills, Kevin L, James J Filliben, Dong Yeon Cho, Edward Schwartz, and Daniel Genin.	NIST SP 500-282.				2010	https://doi.org/10.6028/NIST.SP.500-282	This definitions cites reference [94] in that paper, which is: [94] Law, A. and Kelton, W. D. (2000) Simulation Modeling and Analysis, 3rd edition,McGraw Hill,
friedman_additive_2000	Additive logistic regression: a statistical view of boosting (With discussion and a rejoinder by the authors)	Friedman, Jerome, Trevor Hastie, and Robert Tibshirani.	The Annals of Statistics	28	2	337-407	2000	https://doi.org/10.1214/aos/1016218223	
kusner_counterfactual_2017	Counterfactual Fairness	Kusner, Matt J, Joshua Loftus, Chris Russell, and Ricardo Silva.	Advances in Neural Information Processing Systems (NIPS)				2017	https://proceedings.neurips.cc/paper/2017/hash/a486cd07e4ac3d270571622f4f316ec5-Abstract.html	Abstract is intuitive Definition 1. Definition 2 is the formal definition provided in Section 3
cronbach_construct_1955	Construct Validity in Psychological Tests.	Cronbach, Lee J., and Paul E. Meehl.	Psychological Bulletin	52		281-302	1955	https://doi.org/10.1037/h0040957	Cronbach, Lee J., and Paul E. Meehl. "Construct Validity in Psychological Tests." Psychological Bulletin 52 (1955): 281–302. https://doi.org/10.1037/h0040957.
fink_survey_2010	Survey Research Methods	Fink, A.	International Encyclopedia of Education (Third Edition)			152-160	2010	https://doi.org/10.1016/B978-0-08-044894-7.00296-7	Glossary has definitions that are used
bordens_research_2010	Research Design and Methods: A Process Approach	Kenneth S. Bordens, Bruce B. Abbott	Book (Eighth Edition)				2011	https://www.amazon.com/Research-Design-Methods-Process-Approach/dp/B008BLHYQ8	This is a book and may not be an ideal source
nist_statistics_2012	NIST/SEMATECH e-Handbook of Statistical Methods							https://doi.org/10.18434/M32189	This is an online site. Definition for Experimental Design is in Section 5.1.1
symeonidis_MLOps_2022	MLOps - Definitions, Tools and Challenges	Symeonidis, Georgios, Evangelos Nerantzis, Apostolos Kazakis, and George A. Papakostas	In 2022 IEEE 12th Annual Computing and Communication Workshop and Conference (CCWC)			453-460	2022	https://ieeexplore.ieee.org/document/9720902	This cites the book "Begining MLOps with ML Flow": https://link.springer.com/book/10.1007/978-1-4842-6549-9
hardt_equality_2016	Equality of Opportunity in Supervised Learning	Hardt, Moritz, Eric Price, and Nati and Srebro	Advances in Neural Information Processing Systems (NIPS)			3315-3323	2016	http://papers.nips.cc/paper/6374-equality-of-opportunity-in-supervised-learning.pdf	
EEOC_Q&A_Employee_Selection	Questions and Answers to Clarify and Provide a Common Interpretation of the Uniform Guidelines on Employee Selection Procedures	Equal Employment Opportunity Commission	Questions and Answers to Clarify and Provide a Common Interpretation of the Uniform Guidelines on Employee Selection Procedures					https://www.eeoc.gov/laws/guidance/questions-and-answers-clarify-and-provide-common-interpretation-uniform-guidelines	Definition for "four-fifths rule" appears in section II for the answer to Q. 11 "What is a substantially different rate of selection?"
Engineering_safety_in_machine_learning	Engineering safety in machine learning	Kush R. Varshney	*Information Theory and Applications Workshop* (ITA), 2016				2016	https://ieeexplore-ieee-org.proxygw.wrlc.org/document/7888195	
DOD_Modeling_and_Simulation_Glossary	DoD Modeling and Simulation (M&S) Glossary	United States Department of Defense	DoD Modeling and Simulation (M&S) Glossary				1998	https://web.archive.org/web/20070710104756/http://www.dtic.mil/whs/directives/corres/pdf/500059m.pdf	Definition for "human-in-the-loop" appears on page 124 (P2.8.15)
Model_Cards_for_Model_Reporting	Model Cards for Model Reporting	Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, Timnit Gebru	arXiv					https://arxiv.org/abs/1810.03993	[Submitted on 5 Oct 2018 (v1), last revised 14 Jan 2019 (this version, v2)]
David_Leslie_Morgan_Briggs	*Explaining Decisions Made with AI: A Workbook (Use Case 1: AI-Assisted Recruitment Tool)*	David Leslie; Morgan Briggs	*Explaining Decisions Made with AI: A Workbook (Use Case 1: AI-Assisted Recruitment Tool)*				2021	https://arxiv.org/ftp/arxiv/papers/2104/2104.03906.pdf	Definition for "automation bias" is on page 18
deeplearningbook_intro	Introduction	Ian Goodfellow, Yoshua Bengio; Aaron Courville	*Deep Learning*				2016	https://www.deeplearningbook.org/contents/intro.html	The definition for "deep learning" is taken from the bottom of page 1 and the very top of page 2, slightly restructured to fit the conventions of a definition.
privacy-enhancing_technologies	Chapter 5: Privacy-enhancing technologies (PETs)	UK Information Commissioner's Office	DRAFT Anonymisation, pseudonymisation and privacy enhancing technologies guidance				2022	https://ico.org.uk/media/about-the-ico/consultations/4021464/chapter-5-anonymisation-pets.pdf	The definition for "differential privacy" appears on page 30. This document, as accessed on October 27, 2022, was last updated on September 7, 2022.
Joseph_Rocca_Ensemble_methods	Ensemble methods: bagging, boosting and stacking	Joseph Rocca	Towards Data Science				2019	https://towardsdatascience.com/ensemble-methods-bagging-boosting-and-stacking-c9214a10a205	Date of publication: April 22, 2019
google_dev_classification-true-false-positive-negative	Classification: True vs. False and Positive vs. Negative	Google	Google Machine Learning Education Foundational Courses					https://developers.google.com/machine-learning/crash-course/classification/true-false-positive-negative	
Public_Health_and_Informatics_MIE_2021	A Preliminary Scoping Study of Federated Learning for the Internet of Medical Things	Arshad Farhad; Sandra I. Woolley; Peter Andras	*Public Health and Informatics: Proceedings of MIE 2021*			504-505	2021	https://www.google.com/books/edition/Public_Health_and_Informatics/81A2EAAAQBAJ?hl=en&gbpv=1&dq=%22Federated+learning+1+2+is+a+learning+model+which+addresses+the+problem+of+data+governance+and+privacy+by+training+algorithms+collaboratively+without+transferring+the+data+to+another+location%22&pg=PA504&printsec=frontcover	Definition for "federated learning" appears on page 504
Black's_Law_Dictionary_harm	harm	The Law Dictionary / *Black's Law Dictionary Second Edition*	The Law Dictionary / *Black's Law Dictionary Second Edition*					https://thelawdictionary.org/harm/	
dataiku_ML_and_linear_models	Machine Learning and Linear Models: How They Work (In Plain English)	Katie Gross	Dataiku				2020	https://blog.dataiku.com/top-machine-learning-algorithms-how-they-work-in-plain-english-1	Published February 20, 2020
ORM_model_inventory	Model Inventory	Open Risk Manual						https://www.openriskmanual.org/wiki/Model_Inventory	
yields.io_model_validation	What Is Model Validation?	Eimee V	Yields.io				2020	https://www.yields.io/blog/what-is-model-validation/	Date of publication is February 3, 2020
Open_Risk_Manual_model_validation	Model Validation	Open Risk Manual	Open Risk Manual					https://www.openriskmanual.org/wiki/Model_Validation	
misuse_of_public_data	Implicit data crimes: Machine learning bias arising from misuse of public data	Efrat Shimron; Jonathan I. Tamir; Ke Wang; Michael Lustig	PNAS	119	13		2022	https://www.pnas.org/doi/full/10.1073/pnas.2117203119	Published March 21, 2022
MIT_Protected_Attributes	Module 3: Pedagogical Framework for Addressing Ethical Challenges - Protected Attritbutes and "Fairness Through Unawareness"	MIT Open Courseware	Exploring Fairness in Machine Learning for International Development				2020	https://ocw.mit.edu/courses/res-ec-001-exploring-fairness-in-machine-learning-for-international-development-spring-2020/pages/module-three-framework/protected-attributes/	
Practical_Law_protected_class	Protected Class	Thomson Reuters Practical Law	Thomson Reuters Practical Law					https://content.next.westlaw.com/practical-law/document/Ibb0a38daef0511e28578f7ccc38dcbee/Protected-Class?viewType=FullText&transitionType=Default&contextData=(sc.Default)&firstPage=true0law.	
Dave_Salvator_sparsity	How Sparsity Adds Umph to AI Inference	Dave Salvator	NVIDIA Blogs				2020	https://blogs.nvidia.com/blog/2020/05/14/sparsity-ai-inference/	Published May 14, 2020
saurabh_label_2020	A Unified View of Label Shift Estimation	Garg, Saurabh, Yifan Wu, Sivaraman Balakrishnan, and Zachary Lipton	Advances in Neural Information Processing Systems				2020	https://proceedings.neurips.cc/paper/2020/hash/219e052492f4008818b8adb6366c7ed6-Abstract.html	Not Necessarily the original source of label shift but provides the definition
provost_data_2013	Data Science and its Relationship to Big Data and Data-Driven Decision Making	Provost, Foster and Tom Fawcett	Big Data	1	1	51-59	2013	https://www.liebertpub.com/doi/full/10.1089/big.2013.1508	
hameed_data_2020	Data Preparation: A Survey of Commercial Tools	Hameed, Mazhar and Felix Naumann	ACM SIGMOD Record	49	3	18-29	2020	https://doi.org/10.1145/3444831.3444835	
Merriam-Webster_ranking	ranking	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/ranking	Definition for "ranking" taken from definition 2
hunter_differential_1979	Differential validity of employment tests by race: A comprehensive review and analysis	Hunter, John E. and Frank L. Schmidt and Ronda Hunter	Psychological Bulletin	86		721-735	1979	https://psycnet.apa.org/record/1979-30107-001	
kelley_dogfooding_2022	'Dogfooding'	Kelley, Lora	The New York Times				2022	https://www.nytimes.com/2022/11/14/business/dogfooding.html	
wolfram_math_2022	Wolfram MathWorld: The Web's Most Extensive Mathematics Resource						2022	https://mathworld.wolfram.com/	
aivodji_fairwashing_2019	Fairwashing: the risk of rationalization	Aivodji, Ulrich, Hiromi Arai, Olivier Fortineau, Sébastien Gambs, Satoshi Hara, and Alain Tapp	Proceedings of the 36th International Conference on Machine Learning				2019	https://proceedings.mlr.press/v97/aivodji19a.html	
khalid_feature_2014	A survey of feature selection and feature extraction techniques in machine learning	Khalid, Samina, Tehmina Khalil, and Shamila Nasreen.	2014 Science and Information Conference			372-378	2014	https://ieeexplore.ieee.org/abstract/document/6918213	
tabassi_adversarial_2019	A Taxonomy and Terminology of Adversarial Machine Learning	Tabassi, Elham, Kevin Burns, Michael Hadjimichael, Andres Molina-Markham, and Julian Sexton.	NIST Internal or Interagency Report (NISTIR) 8269 (Draft)				2019	https://doi.org/10.6028/NIST.IR.8269-draft	
measurement_iso22989_2022	ISO/IEC 22989:2022 Information technology — Artificial intelligence — Artificial intelligence concepts and terminology		ISO/IEC 22989:2022				2022	https://www.iso.org/standard/74296.html	Reference taken from AIME "Notes on Measurement" Document
aime_measruement_2022	Notes on Measurement	NIST AIME Team	Unpublished Manuscript				2022	<None>	Cites ISO Standards and NIST IRs; see each term citation for which source this document cites
saarela_feature_2021	Comparison of feature importance measures as explanations for classification models	Mirka Saarela and Susanne Jauhiainen	SN Applied Sciences	3			2021	https://link.springer.com/article/10.1007/s42452-021-04148-9	
poole_mackworth_observation	5.3.1 Background Knowledge and Observations	David Poole and Alan Mackworth	*Artificial Intelligence: Foundations of Computational Agents*				2010	https://artint.info/html/ArtInt_112.html	
kathleen_walch_operationalization	Operationalizing AI	Kathleen Walch	Forbes				2020	https://www.forbes.com/sites/cognitiveworld/2020/01/26/operationalizing-ai/?sh=42ea4b2733df	Originally published on January 26, 2020
about_ML_packages	About ML Packages	UiPath	UiPath AI Center Guide				2021	https://docs.uipath.com/ai-fabric/v0/docs/about-ml-packages	I am guessing on the year; as of this date (11/23/2022), the page says, "Updated 12 months ago"
TechTarget_data_point	data point	Katie Terrell Hanna and Ivy Wigmore	TechTarget				2022	https://www.techtarget.com/whatis/definition/data-point	See note for "point" regarding the full definition here. Last updated July 2022.
Morris_John_data_point	What is a data point in a machine learning model?	John Morris	VProexpert				2022	https://www.vproexpert.com/what-is-a-data-point-in-a-machine-learning-model/	See note for "point" regarding the full definition here. Originally published July 30, 2022.
Artasanchez_Joshi_AI_with_Python	Natural Language Processing	Alberto Artasanchez and Prateek Joshi	*Artificial Intelligence with Python: Your Complete Guide to Building Intelligent Apps Using Python 3.x, 2nd Edition*			351-378	2020	https://www.google.com/books/edition/Artificial_Intelligence_with_Python/P0fODwAAQBAJ?hl=en&gbpv=1&dq=%22Lemmatization+is+the+process+of+grouping+together+the+different+inflected+forms+of+a+word+so+they+can+be+analyzed+as+a+single+item%22&pg=PA356&printsec=frontcover	Definition for "lemmatization" appears on page 356.
Techopedia_lemmatization	Lemmatization	Techopedia	Techopedia					https://www.techopedia.com/definition/33256/lemmatization	
Techslang_lemmatization	What is Lemmatization?	Techslang	Techslang					https://www.techslang.com/definition/what-is-lemmatization/	
TechTarget_lemmatization	Lemmatization	TechTarget contributor	TechTarget				2018	https://www.techtarget.com/searchenterpriseai/definition/lemmatization	Page was last updated February 2018.
Lim_Swee_Kiat_harms	Understanding Bias Part I	Lim Sweet Kiat	Machines Gone Wrong				2019	https://machinesgonewrong.com/bias_i/#two-types-of-harms	See notes column for harms of allocation, harms of bias, and harms of representation
ICO_data_minimisation	Data minimisation and privacy-preserving techniques in AI systems	Information Commissioner's Office	ICO AI Blog				2022	https://ico.org.uk/about-the-ico/media-centre/ai-blog-data-minimisation-and-privacy-preserving-techniques-in-ai-systems/	Published on August 21, 2022.
EDPS_data_minimization	Data Minimization	European Data Protection Supervisor	Data Protection Glossary					https://edps.europa.eu/data-protection/data-protection/glossary/d_en	
Arjun_Subramonian_bias_mitigation	An Introduction to Fairness and Bias Mitigation with AllenNLP	Arjun Subramonian	AI2Blog				2021	https://blog.allenai.org/an-introduction-to-fairness-and-bias-mitigation-with-allennlp-d1b478d44d4c	Published on January 28, 2021
Chi,_Gao,_Ma	Tik Tok Unwrapped	Nicole Chi, Keming Gao, Joanne Ma	Tik Tok Unwrapped				2022	https://www.ischool.berkeley.edu/sites/default/files/sproject_attachments/tiktok_unwrapped_zine_final.pdf	
TechTarget_decision_support_system	decision support system	TechTarget	TechTarget				2021	https://www.techtarget.com/searchcio/definition/decision-support-system	Published May 2021
Burstein_Holsapple	Spreadsheet-Based Decision Support Systems	Frada Burstein and Clyde W. Holsapple	*Handbook on Decision Support Systems 1: Basic Themes*				2008	https://www.google.com/books/edition/Handbook_on_Decision_Support_Systems_1/q_3sRkRKZQwC?hl=en&gbpv=0	Definition for "decision support system" appears on pp. 278-279
Sourabh_Mehta_deterministic	Deterministic vs Stochastic Machine Learning	Sourabh Mehta	Analytics India Magazine				2022	https://analyticsindiamag.com/deterministic-vs-stochastic-machine-learning/	Published May 10, 2022.
jansen_graphical_1998	The Graphical User Interface.	Jansen, Bernard J.	ACM SIGCHI Bulletin	30	2	22-26	1998	https://doi.org/10.1145/279044.279051	
rudin_interpretable_2022	Interpretable Machine Learning: Fundamental Principles and 10 Grand Challenges.	Rudin, Cynthia, Chaofan Chen, Zhi Chen, Haiyang Huang, Lesia Semenova, and Chudi Zhong.	Statistics Surveys	16		1-85	2022	https://doi.org/10.1214/21-SS133	
NISTIR_8312_Full	Four Principles of Explainable Artificial Intelligence	Phillips, P. Jonathon; Carina A. Hahn; Peter C. Fontana; David A. Broniatowski; Mark A. Przybocki	NISTIR 8312				2021	https://nvlpubs.nist.gov/nistpubs/ir/2021/NIST.IR.8312.pdf	
arun_opportunities_2020	Opportunities and Challenges in Explainable Artificial Intelligence (XAI): A Survey.	Das, Arun, and Paul Rad.	Arxiv				2020	https://doi.org/10.48550/arXiv.2006.11371	
open_risk_2022	Open Risk Manual: Midel Governancde		Open Risk Manual				2022	https://www.openriskmanual.org/wiki/Model_Governance	
merriam_webster_outcome	outcome	Merriam-Webster	*Merriam-Webster Dictionary*				2022	https://www.merriam-webster.com/dictionary/outcome	
sutton_reinforcement_2018	Reinforcement Learning: An Introduction	Sutton, Richard, and Andrew Barto	Book, Published by MIT Press				2018		Second Edition
Covert_et_al	Towards a Triad for Data Privacy	Quentin Covert, Mary Francis, Dustin Steinhagen, and Kevin Streff	Proceedings of the 53rd Hawaii International Conference on System Sciences				2020	https://scholarspace.manoa.hawaii.edu/server/api/core/bitstreams/5486a250-cc3c-4227-a752-7d08378afbdf/content	Definition for "intervenability" appears on page 4381
Cami_Rosso	The Human Bias in the AI Machine: How artificial intelligence is subject to cognitive bias	Cami Rosso	Psychology Today				2018	https://www.psychologytoday.com/us/blog/the-future-brain/201802/the-human-bias-in-the-ai-machine	Published February 6, 2018
NIST_CSRC_man-in-the-middle_attack	man-in-the-middle attack (MitM)	NIST Computer Security Resource Center	NIST Computer Security Resource Center					https://csrc.nist.gov/glossary/term/man_in_the_middle_attack	Definitions 1, 2, and 3 of "man-in-the-middle" are taken from this page. However, see the explanatory note for the original sourcing (which is also available on the CSRC page).
cambridge_dictionary_2022	Cambridge Dictionary						2022	https://dictionary.cambridge.org/us/	
kulinski_feature_2020	Feature Shift Detection: Localizing Which Features Have Shifted via Conditional Distribution Tests	Kulinski, Sean M., Saurabh Bagchi, and David I. Inouye	NIPS				2020	https://proceedings.neurips.cc/paper/2020/file/e2d52448d36918c575fa79d88647ba66-Paper.pdf	
settles_active_2009	Active Learning Literature Survey	Burr Settles	Technical Report, University of Wisconsim-Madison, Department of Computer Sciences				2009	https://minds.wisconsin.edu/handle/1793/60660	
informs_analytics_2022	Operations Research & Analytics	INFORMS					2022	https://www.informs.org/Explore/Operations-Research-Analytics	INFORMS is a highly-respected professional society for Operations Researchers
guresen_definition_2011	Definition of Artificial Neural Networks with Comparison to Other Networks	Guresen, Erkam, and Gulgun Kayakutlu	Procedia Computer Science, World Conference on Information Technology	3			2011	https://doi.org/10.1016/j.procs.2010.12.071	
olson_pmlb_2017	PMLB: A Large Benchmark Suite for Machine Learning Evaluation and Comparison.	Olson, Randal S., William La Cava, Patryk Orzechowski, Ryan J. Urbanowicz, and Jason H. Moore	BioData Mining	10	1		2017	https://doi.org/10.1186/s13040-017-0154-4	
devore_probability_2004	Probability and Statistics for Engineering and the Sciences	Jay L. Devore	Book, Published by Thompson, Brooks/Cole, Sixth Edition				2004		
aggarwal_clustering_2013	Data Clustering: Algorithms and Applications	Aggarwal, Charu C., and Chandan K. Reddy	Book, Published by Chapman & Hall/CRC, First Edition				2013		
techopedia_column_2022	Database Column	Techopedia					2022	https://www.techopedia.com/definition/8/database-column#:~:text=In%20the%20context%20of%20relational,documents%20or%20even%20video%20clips.	
box_statistics_2005	Statistics for Experimenters: Design, Innovation, and Discovery	George E. P. Box, J. Stuart Hunter, William G. Hunter	Book, Published by Wiley, Second Edition				2005		
james_statistical_2014	An Introduction to Statistical Learning: With Applications in R	James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani	Book, Published by Springer				2014	http://www-bcf.usc.edu/~gareth/ISL/	
friedler_comparative_2019	A Comparative Study of Fairness-Enhancing Interventions in Machine Learning	Friedler, Sorelle A., Carlos Scheidegger, Suresh Venkatasubramanian, Sonam Choudhary, Evan P. Hamilton, and Derek Roth	In Proceedings of the Conference on Fairness, Accountability, and Transparency				2019	https://doi.org/10.1145/3287560.3287589	
gong_differential_2020	A Survey on Differentially Private Machine Learning [Review Article].	Gong, Maoguo, Yu Xie, Ke Pan, Kaiyuan Feng, and A.K. Qin	IEEE Computational Intelligence Magazine	15	2		2020	https://doi.org/10.1109/MCI.2020.2976185	
Bolukbasi_et_al_Debiasing_Word_Embeddings	Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings	Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai	30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain				2016	https://papers.nips.cc/paper/2016/file/a486cd07e4ac3d270571622f4f316ec5-Paper.pdf	The definition for "word embedding" is taken from a sentence in the abstract and a sentence later on page 1 (the second paragraph of the introduction).
google_glossary_2023	Machine Learning Glossary	Google					2023	https://developers.google.com/machine-learning/glossary	
wikipedia_graph_2023	Graph (discrete mathematics)	Wikipedia	Wikipedia				2023	https://en.wikipedia.org/wiki/Graph_(discrete_mathematics)#Definitions	I realize Wikipedia is not the best source but having a G=(V,E) definition for a graph I think is a foundational definition
Mark_Ciampa_2021	Utilizing Threat Data and Intelligence	Mark Ciampa	*CompTIA CySA+ Guide to Cybersecurity Analyst* (CS0-002)				2021	https://www.google.com/books/edition/CompTIA_CySA+_Guide_to_Cybersecurity_Ana/NwpIEAAAQBAJ?hl=en&gbpv=1&dq=%22In+cybersecurity+a+threat+actor+is+a+term+used+to+describe+individuals+or+entities+who+are+responsible+for+cyber+incidents+against+enterprises+governments+and+users%22&pg=PA29&printsec=frontcover	Definition for "bad actor" is taken from page 29 of this book; Ciampa uses the term "threat actor" and offers "attacker" and "malicious actor" as synonyms.
David_Lyon_2007	Security, Suspicion, Social Sorting	David Lyon	*Surveillance Studies: An Overview*				2007	https://www.google.com/books/edition/Surveillance_Studies/_dTHJgh3-f0C?hl=en&gbpv=1&bsq=surveillance%20is	Definition 1 for "surveillance" is taken from page 96 of this book.
Hartley_and_Zisserman_2003	Projective Geometry and Transformations of 2D	Richard Hartley; Andrew Zisserman	*Multiple View Geometry in Computer Vision*				2003	https://www.google.com/books/edition/Multiple_View_Geometry_in_Computer_Visio/si3R3Pfa98QC?hl=en&gbpv=1&dq=%22Projection+along+rays+through+a+common+point+the+centre+of+pro+jection+defines+a+mapping+from+one+plane+to+another%22&pg=PA34&printsec=frontcover	Definition 1 for "projection" appears on page 34 of this book.
Sloane_et_al_2020	Participation is not a Design Fix for Machine Learning	Mona Sloane, Emanuel Moss, Olaitan Awomolo, Laura Forlano	*Proceedings of the 37th International Conference on Machine Learning*, PMLR				2020	https://arxiv.org/ftp/arxiv/papers/2007/2007.02423.pdf	See note for definition 1 of "participation." See this source for a robust discussion on "participation-washing" and an explanation of different forms of participation (specifically "participation as work," "participation as consultation," and "participation as justice").

ID	Title of article, chapter, or page	Author(s) and/or Editor(s)	Publication or website (either the main domain or major subdomain)	Volume	Issue	Page(s)	Year	URL	Notes
NIST_CSRC_parity	parity	NIST Computer Security Resource Center	NIST Computer Security Resource Center					https://csrc.nist.gov/glossary/term/parity	
Dennis_Mercadal	L	Dennis Mercadal	*Dictionary of Artificial Intelligence*				1990	https://archive.org/details/dictionaryofarti0000merc/page/162/mode/2up?view=theater	See page 163 for definition of "learning"; many other specific kinds of "learning" are defined here as well.
Ekaterina_et_al_2020	Why Are We Averse towards Algorithms? A Comprehensive Literature Review on Algorithmic Aversion	Ekaterina Jussupow, Izak Benbasat, and Armin Heinzl	Proceedings of the 28th European Conference on Information Systems (ECIS), An Online AIS Conference				2020	https://aisel.aisnet.org/ecis2020_rp/168/	
Gabriel_2020	Artificial Intelligence, Values, and Alignment	Iason Gabriel	Minds and Machines	30			2020	https://link.springer.com/article/10.1007/s11023-020-09539-2	
nist_800_2010	A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications.	Rukhin, Andrew, Juan Soto, James Nechvatal, Miles Smid, Elaine Barker, Stefan Leigh, Mark Levenson, et al.	NIST SP-800-22ra				2010	https://doi.org/10.6028/NIST.SP.800-22r1a	
iso_2382_1997	Information technology — Vocabulary — Part 31: Artificial intelligence — Machine learning	ISO	ISO/IEC 2382-31				1997		
iso_22989_2022	Information technology — Artificial intelligence — Artificial intelligence concepts and terminology	ISO	ISO/IEC 22989				2022	https://www.iso.org/standard/74296.html	
european_ethics_2019	Ethics Guidelines for Trustworthy AI	High-Level Expert Group on Artificial Intelligence	Draft Report				2019	https://ec.europa.eu/futurium/en/ai-alliance-consultation.1.html	
Friedman_et_al_2017	A Survey of Value Sensitive Design Methods	Batya Friedman, David G. Hendry, and Alan Borning	*Foundations and Trends® in Human–Computer Interaction*				2017	https://www.nowpublishers.com/article/Details/HCI-015	
Bipartisan_Policy_Center_impact_assessments	Explainer: Impact Assessments for Artificial Intelligence	Sean Long, Jeremy Pesner, and Tom Romanoff	Bipartisan Policy Center				2022	https://bipartisanpolicy.org/blog/impact-assessments-for-ai/	Published November 9, 2022
Kimiz_Dalkir_2011	The Value of Knowledge Management	Kimiz Dalkir	*Knowledge Management in Theory and Practice*				2011	https://www.google.com/books/edition/Knowledge_Management_in_Theory_and_Pract/_MrxCwAAQBAJ?hl=en&gbpv=1&dq=%22Qualitative+measures+provide+more+context+and+details+about+the+value+e+g+perceptions+which+are+often+difficult+to+measure+quantitatively%22&pg=PA343&printsec=frontcover	Definition 1 for "qualitative measurement" is on page 343 ("Qualitative measures can serve to...").
Cost_Management_ch2	Basic Cost Management Concepts	Don R. Hansen; Maryanne M. Mowen; Dan L. Heitger	*Cost Management*				2021	https://www.google.com/books/edition/Cost_Management/HhQcEAAAQBAJ?hl=en&gbpv=1&dq=%22Qualitative+measurement+implies+the+use+of+data+expressed+in+categories+such+as+customer+reviews+of+new+model+jet+skis%22&pg=PA38&printsec=frontcover	Definition 2 for "qualitative measurement" is on page 343 ("Qualitative measures can serve to...").
Virginia_Dignum_Responsibility_and_Artificial_Intelligence	Responsibility and Artificial Intelligence	Virginia Dignum	*The Oxford Handbook of Ethics of AI*			215-232	2020	https://www.google.com/books/edition/The_Oxford_Handbook_of_Ethics_of_AI/8PQTEAAAQBAJ?hl=en&gbpv=1&dq=%22Understanding+the+values+behind+the+technology+and+deciding+on+how+we+want+our+values+to+be+incorporated+in+AI+systems+requires+that+we+are+also+able+to+decide+on+how+and+what+we+want+AI+to+mean+in+our+societies%22&pg=PA221&printsec=frontcover	Definition 1 for "values" is taken verbatim from page 221; see the note at the end of the term's row.
yeom_avoiding_2021	Avoiding Disparity Amplification under Different Worldviews	Yeom, Samuel, and Michael Carl Tschantz	FAccT 2021: In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency			273-283	2021	https://doi.org/10.1145/3442188.3445892	
Merriam-Webster_context	context	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/context	Taken from definition 2
jacobs_measurement_2023	Measurement and Fairness	Jacobs, Abigail Z., and Hanna Wallach	FAccT 2021: In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency				2021	https://doi.org/10.1145/3442188.3445901	
Merriam-Webster_impact	impact	Merriam-Webster	Merriam-Webster					https://www.merriam-webster.com/dictionary/impact	Taken from definition 2
Lisa_M._Given_SAGE	M: Mixed Methods Research	Lisa M. Given	*The SAGE Encyclopedia of Qualitative Research Methods*			491-538	2008	https://www.google.com/books/edition/The_SAGE_Encyclopedia_of_Qualitative_Res/byh1AwAAQBAJ?hl=en&gbpv=1&dq=Mixed+methods+is+defined+as+research+in+which+the+inquirer+or+investigator+collects+and+analyzes+data,+integrates+the+findings,+and+draws+inferences+using+both+qualitative+and+quantitative+approaches+or+methods+in+a+single+study+or+a+program+of+study.&pg=PT584&printsec=frontcover	Definition for "mixed methods" appears on page 526.
NIST_AI_RMF_1.0	NIST AI RMF 1.0	NIST	NIST AI RMF 1.0				2023	https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf	Definition 5 for "risk" comes from p. 3 of NIST AI RMF 1.0.
45_CFR_46_2018_Requirements_(2018_Common_Rule)	2018 Requirements (2018 Common Rule)	United States Department of Health and Human Services (HHS)	45 CFR 46				2018	https://www.hhs.gov/ohrp/regulations-and-policy/regulations/45-cfr-46/revised-common-rule-regulatory-text/index.html	
cambridge_causative_2023	Cambridge Dictionary: causative	Cambridge	Cambridge Dictionary				2023	https://dictionary.cambridge.org/us/dictionary/english/causative	
lyons_contestability_2021	Conceptualising Contestability: Perspectives on Contesting Algorithmic Decisions	Lyons, Henrietta, Eduardo Velloso, and Tim Miller	Proceedings of the ACM on Human-Computer Interaction				2021	https://doi.org/10.1145/3449180	
cambridge_contestable_2023	Contestable		Cambridge Dictionary				2023	https://dictionary.cambridge.org/us/dictionary/english/contestable	
Saleh_Alkhalifa_ML_in_Biotech	Supervised Machine Learning	Saleh Alkhalifa	*Machine Learning in Biotechnology and Life Sciences: Build Machine Learning Models Using Python and Deploy Them on the Cloud*			168-233	2022	https://www.google.com/books/edition/Machine_Learning_in_Biotechnology_and_Li/KUVWEAAAQBAJ?hl=en&gbpv=1&dq=%22We+can+define+supervised+learning+as+a+general+subset+of+machine+learning+in+which+data+like+its+associated+labels+is+used+to+train+models+that+can+learn+or+generalize+from+the+data+to+make+predictions+preferably+with+a+high+degree+of+certainty%22&pg=PA168&printsec=frontcover	Definition 4 for "supervised learning" appears on page 168; definition 2 for "regression" also appears on page 168
Schneider_McGrew_in_Flanagan_McDonough_2018	The Cattell-Horn-Carroll Theory of Cognitive Abilities	W. Joel Schneider and Kevin S. McGrew; edited by Dawn P. Flanagan and Erin M. McDonough	*Contemporary Intellectual Assessment: Theories, Tests, and Issues*			73-163	2018	https://www.google.com/books/edition/Contemporary_Intellectual_Assessment/JA1mDwAAQBAJ?hl=en&gbpv=1&dq=%22Formal+expertise+is+the+result+of+a+selfselection+of+a+domain+of+knowledge+that+is+mastered+deliberately+and+for+which+there+are+clear+benchmarks+of+success+Fisher+Keil+2016%22&pg=PA117&printsec=frontcover	Definition 1 for "expertise" appears on page 117.
Little_2013	The Measurement Model	Todd D. Little	*Longitudinal Structural Equation Modeling*			71-105	2013	https://www.google.com/books/edition/Longitudinal_Structural_Equation_Modelin/gzeCu3FjZf4C?hl=en&gbpv=1&dq=%22Measurement+model%22&pg=PA103&printsec=frontcover	Definition 1 for "measurement model" appears on page 103.
Merriam-Webster_executive	executive	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/executive	Definition 1 for "executive" is taken from Merriam-Webster's definition of "executive" 2 of 2 noun 3.
Dorf_2018	Measurement and Instrumentation	Richard C. Dorf	*The Engineering Handbook*			151-160	2018	https://www.google.com/books/edition/The_Engineering_Handbook/l_TLBQAAQBAJ?hl=en&gbpv=1&dq=%22The+propagation+of+uncertainty+is+defined+as+the+way+in+which+uncertainties+in+the+variables+affect+the+uncertainty+in+the+calculated+results%22&pg=SA99-PA711&printsec=frontcover	Definition 1 for "error propagation" is taken from page 152; technically, the textbook offers the term as "propagation of uncertainty," which is synonymous with "error propagation."
Merriam-Webster_example	example	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/example	
Merriam-Webster_ethic	ethic	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/ethic	
Merriam-Webster_anthropomorphism	anthropomorphism	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/anthropomorphism	
berthold_guide_2020	Guide to Intelligent Data Science: How to Intelligently Make Use of Real Data	Berthold, Michael R., Christian Borgelt, Frank Höppner, Frank Klawonn, and Rosaria Silipo	Springer International Publishing				2020	https://doi.org/10.1007/978-3-030-45574-3	
alon-barkat_human_2023	Human–AI Interactions in Public Sector Decision Making: 'Automation Bias' and 'Selective Adherence' to Algorithmic Advice.	Alon-Barkat, Saar, and Madalina Busuioc	Journal of Public Administration Research and Theory	33	1	153–169	2023	https://doi.org/10.1093/jopart/muac007	
humphrey_addressing_2020	Addressing Harmful Bias and Eliminating Discrimination in Health Professions Learning Environments: An Urgent Challenge.	Humphrey, Holly J., Dana Levinson, Marc A. Nivet, and Stephen C. Schoenbaum	Academic Medicine	95	12S		2020	https://doi.org/10.1097/ACM.0000000000003679	
knuth_art_1981	The Art of Computer Programming, Volume 2: Seminumerical Algorithms	Donald Knuth	Addison-Wesley	2			1981		
garey_computers_1979	Computers and Intractability: A Guide to the Theory of NP-Completeness	Michael Garey and David Johnson	W. H. Freeman				1979		
cambridge_impact_2023	Impact		Cambridge Dictionary				2023	https://dictionary.cambridge.org/us/dictionary/english/impact	
law_policy_2023	Policy		The Law Dictionary				2023	https://thelawdictionary.org/policy/#:~:text=Definition%20%26%20Citations%3A,as%20directed%20to%20the%20POLICY	
fernandez_residual_1992	Residual Analysis and Data Transformations: Important Tools in Statistical Analysis	Fernandez, George C. J.	HortScience	27	4	297–300	1992	https://journals.ashs.org/hortsci/view/journals/hortsci/27/4/article-p297.xml	
Leavy_OHQR_Intro	Introduction	Patricia Leavy	*The Oxford Handbook of Qualitative Research*			1-13	2014	https://academic.oup.com/edited-volume/38166/chapter-abstract/332997092?redirectedFrom=fulltext	Definition of "qualitative research" is taken from page 2.
Barbour_2014	The scope and contribution of qualitative research	Rosaline S. Barbour	*Introducing Qualitative Research: A Student's Guide, Second Edition*			11-27	2008; 2014	https://archive.org/details/introducingquali0000barb_j9h1/page/12/mode/2up?view=theater	Definition of "qualitative research" is taken from page 13.
Merriam-Webster_engineer	engineer	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/engineer	Definition takes from both the noun and verb definitions.
AI_Incident_Editors	Editor's Guide	AI Incident Database	AI Incident Database				2023	https://incidentdatabase.ai/editors-guide/	
interaction_context_2023	Context of Use	Interaction Design: The Glossary of Human Computer Interaction	Interaction Design Foundation				2023	https://www.interaction-design.org/literature/book/the-glossary-of-human-computer-interaction/context-of-use	
AAAS_AI_and_Bias_2022-09	Artificial Intelligence and Bias – An Evaluation	M. Karanicolas and M. Knodel	*Artificial Intelligence and the Courts: Materials for Judges*				2022	https://doi.org/10.1126/aaas.adf0788	Definition for "racialized" taken from page 11.
Seth_Boden_2020	Start Here: A Primer on Diversity and Inclusion (Part 1 of 2)	Seth Boden	Harvard Business Publishing				2020	https://www.harvardbusiness.org/start-here-a-primer-on-diversity-and-inclusion-part-1-of-2/	July 23, 2020
GWU_diversity_and_inclusion	Diversity and Inclusion Defined	George Washington University	George Washington University Office for Diversity, Equity and Community Engagement					https://diversity.gwu.edu/diversity-and-inclusion-defined	
HUD_diversity_and_inclusion	Diversity and Inclusion Definitions	U.S. Department of Housing and Urban Development	U.S. Department of Housing and Urban Development					https://www.hud.gov/program_offices/administration/admabout/diversity_inclusion/definitions	
Jamieson_Govaart_Pownall	Reflexivity in quantitative research: A rationale and beginner's guide	Michelle K. Jamieson, Gisela H. Govaart, and Madeleine Pownall	*Social and Personality Psychology Compass*		e12735		2023	https://doi.org/10.1111/spc3.12735	Definition for "reflexivity" appears on page 2. Jamieson, Govaart, and Pownall ("Reflexivity in quantitative research: A rationale and beginner's guide") are paraphrasing Carla Willig's ideas, but in their citation of Willig's book (*Introducing Qualitative Research in Psychology*), they're not providing a page number where they are paraphrasing how Willig defines reflexivity. So they must be paraphrasing the general sense throughout the book, maybe. I guess that gives us the green light to just be fine with citing Jamieson, Govaart, and Pownall as the "original" source for that definition.
Industrial_Network_Security_2011	Monitoring Enclaves	Eric D. Knapp and Joel Langill	*Industrial Network Security: Securing Critical Infrastructure Networks for Smart Grid, SCADA, and Other Industrial Control Systems*				2011	https://www.google.com/books/edition/Industrial_Network_Security/PlOtqouTwaUC?hl=en&gbpv=1&dq=%22Data+retention+refers+to+the+amount+of+information+that+is+stored+long-term,+and+can+be+measured+in+volume+(the+size+of+the+total+collected+logs+in+bytes)+and+time+(the+number+of+months+or+years+that+logs+are+stored+for).%22&pg=PA243&printsec=frontcover	Definition for "retention limit" appears on page 243; appears as "data retention" in the book, but works for "retention limit."
ChatGPT	ChatGPT	OpenAI						https://chat.openai.com/chat	We can use this to cite any moments that we use a ChatGPT-generated definition in lieu of a more authoritative source.
Merriam-Webster_parity	Parity		Merriam-Webster Dictionary				2023	https://www.merriam-webster.com/dictionary/parity	
barredo_explainable_2020	Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI	Barredo Arrieta, Alejandro, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador Garcia, et al.	Information Fusion	58		82–115	2020	https://doi.org/10.1016/j.inffus.2019.12.012	
jackman_oxford_2008	Measurement	Simon Jackman	The Oxford Handbook of Political Methodology				2008		
Hammarberg_2016_Busetto_2020	1. Qualitative research methods: when to use them and how to judge them; 2. How to use and assess qualitative research methods	1. K. Hammarberg, M. Kirkman, and S. de Lacey; 2. Loraine Busetto, Wolfgang Wick, and Christoph Gumbinger	1. *Human Reproduction*; 2. *Neurological Research and Practice*		1. 31	1. 3; 2. 14	1. 2016; 2. 2020	1. https://academic.oup.com/humrep/article/31/3/498/2384737; 2. https://neurolrespract.biomedcentral.com/articles/10.1186/s42466-020-00059-z	The definition combines two sources into one. I (Danny) have organized this citation row using 1. and 2. to mark, respectively, the Hammarberg (et al) and Busetto (et al) pieces. **Citation 1 (rough)**: K. Hammarberg, M. Kirkman, S. de Lacey, Qualitative research methods: when to use them and how to judge them, Human Reproduction, Volume 31, Issue 3, March 2016, Pages 498–501, https://doi.org/10.1093/humrep/dev334. **Citation 2 (rough)**: Busetto, L., Wick, W. & Gumbinger, C. How to use and assess qualitative research methods. Neurol. Res. Pract. 2, 14 (2020). https://doi.org/10.1186/s42466-020-00059-z
Russell_2003_Brannen_2005	1. Evaluation of qualitative research studies; 2. Mixing Methods: The Entry of Qualitative and Quantitative Approaches into the Research Process	1. Cynthia K. Russell and David M. Gregory; 2. Julia Brannen	1. *Evidence-Based Nursing*; 2. *International Journal of Social Research Methodology*	1. 6; 2. 8	2. 8	1. 36-40; 2. 173-184	1. 2003; 2. 2007	1. http://dx.doi.org/10.1136/ebn.6.2.36; 2. https://www.tandfonline.com/doi/full/10.1080/13645570500154642	The definition combines two sources into one. I (Danny) have organized this ctiation row using 1. and 2. to mark these two sources: 1. Russell, C. K., & Gregory, D. M. (2003). Evaluation of qualitative research studies. Evidence Based Nursing, 6(2), 36–40 2. Brannen J. Mixing methods: the entry of qualitative and quantitative approaches into the research process. Int J Soc Res Methodol 2005;8:173–184.
Pamela_Goh_2021	Humans as the Weakest Link in Maintaining Cybersecurity: Building Cyber Resilience in Humans	Pamela Goh. Edited by Majeed Khader, Loo Seng Neo, and Whistine Xiau Ting Chai	*Introduction to Cyber Forensic Psychology: Understanding the Mind of the Cyber Deviant Perpetrators*			287-305	2021	https://www.worldscientific.com/doi/abs/10.1142/9789811232411_0014	Definition 3 for "red-team" appears on page 298.
Annie_Jacobsen_2015	Total Information Awareness	Annie Jacobsen	*The Pentagon's Brain: An Uncensored History of DARPA, America's Top-Secret Military Research Agency*			336-352	2015	https://archive.org/details/pentagonsbrainun0000jaco_c8o3/page/342/mode/2up?q=%22a+role-playing+exercise+in+which+a+problem+is+examined+from+an+adversary%E2%80%99s+or+enemy%E2%80%99s+perspective.%22	Definition 2 for "red-team" appears on page 342. (In the large print edition of this book, it's on page 501. But we should cite the regular edition.)
Ben_Auffarth_2021	Online Learning for Time-Series	Ben Auffarth	*Machine Learning for Time-Series with Python: Forecast, Predict, and Detect Anomalies with State-of-the-art Machine Learning Methods*			209-259	2021	https://www.google.com/books/edition/Machine_Learning_for_Time_Series_with_Py/a7tLEAAAQBAJ?hl=en&gbpv=1&dq=%22On+the+other+hand+offline+learning+the+more+commonly+known+approach+implies+that+you+have+a+static+dataset+that+you+know+from+the+start+and+the+parameters+of+your+machine+learning+algorithm+are+adjusted+to+the+whole+dataset+at+once+often+loading+the+whole+dataset+into+memory+or+in+batches%22&pg=PA210&printsec=frontcover	Definitions 1 for "offline learning" and "online learning" both appear on page 210.
FWS_062_FW_1	062 FW 1, Affirmative Employment Program and Plans	Office for Human Resources of the U.S. Fish & Wildlife Service	U.S. Fish & Wildlife Service				1996	https://www.fws.gov/policy/062fw1.html	
Merriam-Webster_assessment	assessment	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/assessment	
Anthropomorphism_in_AI_2020	anthropomorphism	Arleen Salles, Kathinka Evers, and Michele Farisco	AJOB *Neuroscience*		2	88-95	2020	https://doi.org/10.1080/21507740.2020.1740350	Definition 1 for "anthropomorphism" appears on page 90; definition 2 begins on page 89 and ends on page 90.
OECD_Artificial_Intelligence_in_Society	The technical landscape	OECD	*Artificial Intelligence in Society*			19-34	2019	https://www.google.com/books/edition/Artificial_Intelligence_in_Society/eRmdDwAAQBAJ?hl=en&gbpv=1&dq=%22Artificial+narrow+intelligence+ANI+or+applied+AI+is+designed+to+accomplish+a+specific+problem+solving+or+reasoning+task%22&pg=PA22&printsec=frontcover	The definition for "artificial narrow intelligence (ANI)" appears on page 22.
AI_in_Medical_Imaging_glossary	Glossary	Erik R. Ranschaert, Sergey Morozov, and Paul R. Algra, eds.	*Artificial Intelligence in Medical Imaging: Opportunities, Applications and Risks*			349-364	2019	https://www.google.com/books/edition/Artificial_Intelligence_in_Medical_Imagi/ss6FDwAAQBAJ?hl=en&gbpv=1&dq=%22The+definition+of+artificial+narrow+intelligence+is+in+contrast+to+that+of+strong+AI+or+artificial+general+intelligence+which+aims+at+providing+a+system+with+consciousness+or+the+ability+to+solve+any+problems%22&pg=PA350&printsec=frontcover	The definition for "artificial narrow intelligence (ANI)" appears on page 22.
DOL_Practical_Significance	Practical Significance in EEO Analysis Frequently Asked Questions	U.S. Department of Labor Office of Federal Contract Compliance Programs	U.S. Department of Labor Office of Federal Contract Compliance Programs				2021	https://www.dol.gov/agencies/ofccp/faqs/practical-significance	Page last updated January 15, 2021.
Cambridge_Dictionary_non-discrimination	non-discrimination	Cambridge Dictionary	*Cambridge Dictionary*					https://dictionary.cambridge.org/us/dictionary/english/non-discrimination	
Signal_Detection_Theory	Signal Detection Theory	N.A. Macmillan	*International Encyclopedia of the Social & Behavioral Sciences*			14075-14078	2001	https://doi.org/10.1016/B0-08-043076-7/00677-X	
Techopedia_kill_switch	Kill Switch	Techopedia	Techopedia				2019	https://www.techopedia.com/definition/4001/kill-switch	Page last updated August 8, 2019
Batya_Friedman_VSD_Introduction	Introduction	Batya Friedman and David G. Hendry	*Value Sensitive Design: Shaping Technology with Moral Imagination*			1-17	2019	https://doi-org.proxy.library.georgetown.edu/10.7551/mitpress/7585.003.0002	Definition for values comes from page 4.
C3.ai_feedback_loop	What Is a Feedback Loop?	C3.ai	C3.ai					https://c3.ai/glossary/features/feedback-loop/	
Collins_Dictionary_ground_truth	ground truth	Collins Dictionary	*Collins Dictionary*					https://www.collinsdictionary.com/us/dictionary/english/ground-truth	

ID	Title of article, chapter, or page	Author(s) and/or Editor(s)	Publication or website (either the main domain or major subdomain)	Volume	Issue	Page(s)	Year	URL	Notes
Wikipedia_Decision-making	Decision-making	Wikipedia	Wikipedia					https://en.wikipedia.org/wiki/Decision-making	Check out *The New Science of Management Decision* (1977) to try to locate the original source of the definition.
Shevlin_et_al_2019	The limits of machine intelligence	Henry Shevlin, Karina Vold, Matthew Crosby, and Marta Halina	*EMBO Reports*	20	10	e49177	2019	https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6776890/	
EEOC_ADA_AI	The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees	U.S. Equal Employment Opportunity Commission	U.S. Equal Employment Opportunity Commission				2022	https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence	Definition for "screen out" taken from this page.
Merriam-Webster_screen_out	screen out	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/screen%20out	
apa_experiment_2023	experiment	American Psychological Association (APA)	*APA Dictionary of Psychology*				2023	https://dictionary.apa.org/experiment	
APA_DoP_laboratory_research	laboratory research	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/laboratory-research	
UNODC_Glossary_QA_GLP	Glossary of Terms for Quality Assurance and Good Laboratory Practices	Laboratory and Scientific Section of the United Nations Office on Drugs and Crime	Glossary of Terms for Quality Assurance and Good Laboratory Practices				2009	https://www.unodc.org/documents/scientific/ST_NAR_26_E.pdf	
World_Wide_Words_In_silico	In silico	World Wide Words	World Wide Words					http://www.worldwidewords.org/weirdwords/ww-ins1.htm	
Bassiouni_Baffes_Evrard	An Appraisal of Human Experimentation in International Law and Practice: The Need for International Regulation of Human Experimentation	M. Cheriff Bassiouni, Thomas G. Baffes, and John T. Evrard	*Journal of Criminal Law and Criminology*	72	4	1597-1666	1981	https://scholarlycommons.law.northwestern.edu/cgi/viewcontent.cgi?article=6276&context=jclc	
Merriam-Webster_amplify	amplify	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/amplify	
Gupta_et_al_HAR_2022	Human activity recognition in artificial intelligence framework: a narrative review	Neha Gupta, Suneet K. Gupta, Rajesh K. Pathak, Vanita Jain, Parisa Rashidi, and Jasjit S. Suri	*Artificial Intelligence Review*	55		4755–4808	2022	https://link.springer.com/article/10.1007/s10462-021-10116-x	
Merriam-Webster_annotate	annotate	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/annotate	
Freeman_et_al_2014	Active learning increases student performance in science, engineering, and mathematics	Scott Freeman, Sarah L. Eddy, Miles McDonough, Michelle K. Smith, Nnadozie Okoroafor, Hannah Jordt, and Mary Pat Wenderoth	PNAS	111	23	8410-8415	2014	https://www.pnas.org/doi/full/10.1073/pnas.1319030111	
AI_Assurance_2022	Assuring AI methods for economic policymaking	Anderson Monken, William Ampeh, Flora Haberkorn, Uma Krishnaswamy, and Feras A. Batarseh	*AI Assurance: Towards Trustworthy, Explainable, Safe, and Ethical AI*			371-428	2022	https://www.google.com/books/edition/AI_Assurance/dch6EAAAQBAJ?hl=en&gbpv=1&dq=%22Large+language+models+LLMs+are+a+class+of+language+models+that+use+deep+learning+algorithms+and+are+trained+on+extremely+large+textual+datasets+that+can+be+multiple+terabytes+in+size%22&pg=PA376&printsec=frontcover	The definition for "large language model (LLM)" appears on page 376. This book was edited by Feras A. Batarseh and Laura Freeman.
Poore_Lawrence_ARLIS_2023-01	AI Engineering: An Academic Research Roadmap	Joshua Poore and Craig Lawrence	Applied Research Laboratory for Intelligence and Security (ARLIS)				2023		We have this PDF, but I do not think it is readily available online. See also https://www.arlis.umd.edu/sites/default/files/2022-03/No_AI_In_Teams_FinalReport%20(1).pdf.
Survey_of_Hallucination_in_NLG	Survey of Hallucination in Natural Language Generation	Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung	ACM *Computing Surveys*	55	12	1-38	2023	https://dl.acm.org/doi/10.1145/3571730	Definition for "hallucination" is taken from page marked 248:3.
APA_clustering	clustering	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/clustering	
APA_content_validity	content validity	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/content-validity	
ISO_9241-11:2018	ISO 9241-11:2018(en) Ergonomics of human-system interaction — Part 11: Usability: Definitions and concepts	ISO	ISO Online Browsing Platform				2018	https://www.iso.org/obp/ui/#iso:std:iso:9241:-11:ed-2:v1:en	
APA_criterion_validity	criterion validity	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/criterion-validity	
APA_data_analysis	data analysis	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/data-analysis	This term in the APA dictionary is "data analysis"; the term in our glossary is "data analytics." The terms are not necessarily synonymous, but we are including the "data analysis" definition in with "data analytics."
APA_decision_making	decision making	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/decision-making	
Lehto_Nanda_2021	Decision-Making Models, Decision Support, and Problem Solving	Mark R. Lehto and Gaurav Nanda	*Handbook of Human Factors and Ergonomics*, Fifth Edition			159-202	2021	https://www.wiley.com/en-us/Handbook+of+Human+Factors+and+Ergonomics%2C+5th+Edition-p-9781119636083	The editors of this book are Gavriel Salvendy and Waldemar Karwowski
Baron_Thinking_and_Deciding		Jonathan Baron	*Thinking and Deciding*				2008		Page 6; also cited here: https://www.redalyc.org/journal/5538/553866738009/html/
EO_DEIA_2021	Executive Order on Diversity, Equity, Inclusion, and Accessibility in the Federal Workforce	Joseph R. Biden Jr.	The White House				2021	https://www.whitehouse.gov/briefing-room/presidential-actions/2021/06/25/executive-order-on-diversity-equity-inclusion-and-accessibility-in-the-federal-workforce/	This executive order was published on June 25, 2021.
APA_ethics	ethics	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/ethics	
APA_external_validity	external validity	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/external-validity	
CSRC_false_negative	False Negative	NIST CSRC	Information Technology Laboratory Computer Security Resource Center Glossary					https://csrc.nist.gov/glossary/term/false_negative	
CSRC_false_positive	False Positive	NIST CSRC	Information Technology Laboratory Computer Security Resource Center Glossary					https://csrc.nist.gov/glossary/term/false_positive	
Wilke_Mata_2012	Cognitive Bias	A. Wilke and R. Mata	*Encyclopedia of Human Behavior*	1		531-535	2012	https://s3.amazonaws.com/arena-attachments/557491/b16d97da35ed37a0a022e806cc931a0d.pdf	Edited by V.S. Ramachandran
AIID_incident_response	Defining an "AI Incident Response"	Sean McGregor	Artificial Intelligence Incident Database					https://incidentdatabase.ai/research/5-response/	
APA_integrity	integrity	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/integrity	
APA_internal_validity	internal validity	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/internal-validity	
APA_learning	learning	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/learning	
McNamara_Fallacy	The McNamara Fallacy	Jonathan Cook	The McNamara Fallacy				2023	https://mcnamarafallacy.com/	
Creswell_Clark_mixed_methods		John W. Creswell and Vicki L. Plano Clark	*Designing and Conducting Mixed Methods Research*, Third Edition				2017		I don't know which chapter and page this citation comes from; all I know is that it's somewhere in this book.
APA_observation	observation	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/observation	
Glossary_of_Statistical_Terms	Glossary of Statistical Terms	Philip B. Stark	SticiGui				2019	https://www.stat.berkeley.edu/~stark/SticiGui/Text/gloss.htm	
Wikipedia_RMSD	Root-mean-square-deviation	Wikipedia	Wikipedia					https://en.wikipedia.org/wiki/Root-mean-square_deviation	
APA_recognition	recognition	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/recognition	
APA_recall	recall	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/recall	
APA_stereotype	stereotype	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/stereotype	
Augoustinos_Walker_1998	The Construction of Stereotypes within Social Psychology: From Social Cognition to Ideology	Martha Augoustinos and Iain Walker	*Theory & Psychology*	8	5	629-652	1998	https://doi.org/10.1177/0959354398085003	Definition 3 for "stereotype" is taken from page 631.
APA_autonomy	autonomy	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/autonomy	
Charmaz_Henwood	Grounded Theory Methods for Qualitative Psychology	Kathy Charmaz and Karen Henwood	*The SAGE Handbook of Qualitative Research in Psychology*			238-255	2017	https://doi.org/10.4135/9781526405555	
Malik_2021	Critical Technical Awakenings	Maya Malik and Momin M. Malik	*Journal of Social Computing*	2	4	365-384	2021	https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=9698152	
APA_reflexivity	reflexivity	American Psychological Association (APA)	*APA Dictionary of Psychology*					https://dictionary.apa.org/reflexivity	
Lee_See_2004	Trust in Automation: Designing for Appropriate Reliance	John D. Lee and Katrina A. See	*Human Factors: The Journal of the Human Factors and Ergonomics Society*	46	1	50-80	2004	https://doi.org/10.1518/hfes.46.1.50_30392	
Mayer_Davis_Schoorman_1995	An Integrative Model of Organizational Trust	Roger C. Mayer, James H. Davis, and F. David Schoorman	*The Academy of Management Review*	20	3	709-734	1995	https://doi.org/10.2307/258792	
NISTIR_8280	NISTIR 8280. Face Recognition Vendor Test (FRVT). Part 3: Demographic Effects	Patrick Grother, Mei Ngan, and Kayee Hanaoka	NIST				2019	https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf	
Usabilitygov	Usability Testing	Usability.gov	Usabillity.gov					https://www.usability.gov/how-to-and-tools/methods/usability-testing.html	
Encyclopedia.com_underrepresentation	Underrepresentation	Encyclopedia.com	Encyclopedia.com					https://www.encyclopedia.com/social-sciences/applied-and-social-sciences-magazines/underrepresentation	
Arham_Islam_History_2023	A History of Generative AI: From GAN to GPT-4	Arham Islam	MarkTechPost				2023	https://www.marktechpost.com/2023/03/21/a-history-of-generative-ai-from-gan-to-gpt-4/	Published March 21, 2023
McKinsey_generative_AI	What is generative AI?	McKinsey & Company	McKinsey & Company				2023	https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-generative-ai	Published January 19, 2023
Seshia_et_al_2022	Toward Verified Artificial Intelligence	Sanjit A. Seshia, Dorsa Sadigh, and S. Shankar Sastry	*Communications of the ACM*	65	7	46-55	2022	https://cacm.acm.org/magazines/2022/7/262079-toward-verified-artificial-intelligence/fulltext	
Liam_Tung_2022_Meta_hallucination	Meta warns its new chatbot may forget that it's a bot	Liam Tung	ZDNet				2022	https://www.zdnet.com/article/meta-warns-its-new-chatbot-may-not-tell-you-the-truth/	Published August 8, 2022
Merriam-Webster_requirement	requirement	Merriam-Webster	*Merriam-Webster Dictionary*					https://www.merriam-webster.com/dictionary/requirement	

[1] Add citation to citations sheet and only list ID in these columns

[2] Make sure the spelling matches another term (value in A column)