

Glossary

This glossary is derived from Winnie Soon and Geoff Cox's [Aesthetic Programming: A Handbook of Software Studies](#) (London: Open Humanities Press, 2020), and published with kind permission of the authors, to provide a shared vocabulary for this volume.

Artificial Intelligence (AI)

AI research is focused on developing computational systems that can perform tasks and activities normally considered to require human intelligence. The term 'AI' is often used interchangeably with 'machine learning' and 'deep learning', but there are key distinctions to be made. To explain: 'You can think of deep learning, machine learning and artificial intelligence as a set of Russian dolls nested within each other, beginning with the smallest and working out. Deep learning is a subset of machine learning, and machine learning is a subset of AI, which is an umbrella term for any computer program that does something smart. In other words, all machine learning is AI, but not all AI is machine learning, and so forth.' (<https://pathmind.com/wiki/ai-vs>) For a more critical explanation, see the essay and diagram by Kate Crawford and Vladan Joler, in 'Anatomy of an AI System: The Amazon Echo as an anatomical map of human labor, data and planetary resources' (2018), <https://anatomyof.ai/>. For a decolonial perspective on AI, see Shakir Mohamed, Marie-Therese Png, William Isaac, 'Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence', *Philosophy & Technology*, Springer (July 12, 2020).

Dataset

A dataset (or data set) is a collection of facts. In the case of tabular data, a dataset corresponds to one or more database table, where every column of a table represents a particular variable, and each row refers to a given record of the dataset in question. Datasets are used by developers to test, train and evaluate the performance of their algorithms. The algorithm is said to 'learn' from the examples contained in the dataset and this constitutes the worldview of the algorithm (with colonial implications). See Nicolas Malevé's 'An Introduction to Image Datasets' (2019) for a useful summary of key concepts and concerns, and the extent to which datasets have become significant cultural objects (<https://unthinking.photography/articles/an-introduction-to-image-datasets>).

Deep Learning

Deep learning is the subfield of machine learning that designs and evaluates training algorithms and architectures for modern neural network models. It is part of a broader collection of machine learning methods based on artificial neural networks with representation learning. Learning can be 'supervised', 'semi-supervised' or 'unsupervised'. See John D. Kelleher, *Deep Learning* (Cambridge, MA: MIT Press, 2019).

Generative Adversarial Networks (GANs)

A GAN comprises two conflicting neural nets — a 'Generator' that forges new data, and a 'Discriminator' that distinguishes this fake data created by the Generator from real data. These nets challenge each other with increasingly realistic fakes, both optimising their strategies until their generated data is indistinguishable from the real data. This is an 'unsupervised' method of training that doesn't rely on the tagging of input images by humans, since the machine generates groupings based on its own analysis. The workshop 'Adversarial Hacking in the Age of AI' asked the question whether critical theory could learn from this system (that seems to resonate with dialectical materialism), in which everything is considered to be in a process of transformation through contradiction, and becomes a technical reality. The published outline provides a useful description of what is at stake: 'Adversarial attacks are an instance of how a machine-learning classifier is tricked into perceiving something that is not there, like a 3D-printed model of a turtle that is classified as a rifle. The computer vision embedded in a driverless car can be confused and not recognize street signs. Artists Adam Harvey, Zach Blas & Jemina Wyman, and Heather Dewey-Hagborg have utilized adversarial processes in their projects in order to subvert and critically respond to facial recognition systems. But this is not just about computer vision. Scientists in Bochum, Germany recently studied how psychoacoustic hiding can oppose the detection of automatic speech

recognition systems.’ See <https://2020.transmediale.de/content/adversarial-hacking-in-the-age-of--i-call-for-proposals>. For more on GANS, see Ian J. Goodfellow, et al, ‘Generative Adversarial Networks’ IPS’14: Proceedings of the 27th International Conference on Neural Information Processing Systems – Volume 2 (2014), pp. 2672–2680.

Machine Learning (ML)

Machine learning is broadly defined as a collection of models, statistical methods and operational algorithms that are used to analyse experimental or observational data. The term itself was coined by Arthur Samuel in 1959 during his game development research at IBM, which ultimately aimed to reduce or even eliminate the need for ‘detailed programming effort’, using learning through generalisation in order to achieve pattern recognition. See Arthur L. Samuel, ‘Some Studies in Machine Learning Using the Game of Checkers’, *IBM Journal of research and development* 3 (3) (1959), pp. 210–229. ML involves the development and evaluation of algorithms that enable computers to learn from experience. Generally, the concept of experience is represented as a dataset of historic events, and learning involves identifying and extracting useful patterns from a dataset. ML algorithms take a dataset as input and return a model that encodes the patterns the algorithm extracted (or learned) from the data. But to what extent does this process of generalisation present a problem inasmuch as the overall idea of learning implies new forms of control over what and how something becomes known and how decisions are made? See Adrian Mackenzie, *Machine Learners: Archaeology of a Data Practice* (Cambridge, MA: MIT Press, 2017).

Model

In machine learning, a model is a computer program that encodes the patterns that the machine learning algorithm has extracted from a dataset. There are many different types of machine learning models, but put simply, a model is created (or trained) by running a machine learning algorithm on a dataset. Once the model has been trained, it can then be used to analyse new instances. Named to reflect some of the problems associated with this, ‘All Models’ is a mailing list of critical AI studies hosted by the research group KIM at the Karlsruhe University of Arts and Design. The ‘About’ page prophetically states that ‘All Models are wrong, but some are useful’, to point to some of the limits of statistics and machine learning, and how, for instance, mainstream AI discourse stresses the need for unbiased data and algorithms to ensure fair representation, but overlooks the intrinsic limits of any statistical technique. Herein lies the politics and the way in which traditional forms of power (such as those related to gender, race, and class discrimination) are amplified. In summary, ‘All Models questions all models!’. See ‘About’ (allmodels.ai).

Neural Network

A neural network is a machine learning model that is implemented as a network of simple information processing units called neurons. It is possible to create a variety of different types of neural networks by modifying the connections between the neurons in the network. Examples of popular types of neural networks include feedforward, convolutional and recurrent networks. A recurrent neural network, for instance, has a single layer of hidden neurons, the output of which is fed back into this layer with the next input. This feedback (or recurrence) within the network gives the network a memory that enables it to process each input within the context of what it has previously processed. Recurrent neural networks are ideally suited to processing sequential or time-series data. A good example of this is Helen Pritchard and Winnie Soon’s ‘Recurrent Queer Imaginaries’ (2019–20), in which a machine learner ‘Motto Assistant’ continuously reworks queer and feminist manifestos to activate alternative imaginaries.

Reinforcement Learning

Reinforcement Learning is based on interaction with the environment, mapping an analysis of a situation into actions, typically used in robot control and game playing. The learner (or agent) does not have any previous data to base itself on in order to determine or predict which action to take, but rather learns by trial and error. This type of learning finds the optimum possible behaviour or path to take in a

specific environment, mapping state-action pairs to achieve the best result. As in behavioural psychology, reinforcement is used to suggest future actions, like a pet or child getting a treat for doing what it was told. Unlike supervised learning that relies on input training data, the characteristics of reinforcement learning are that the programme understands the environment as a whole, and is able to learn from its experience by evaluating the effectiveness of each action taken: 'trial-and-error search' and 'delayed reward' are based on sequential decisions, computation, repeated attempts and feedback on the success of actions. See Richard S. Sutton, 'Introduction: The Challenge of Reinforcement Learning', in Richard S. Sutton, ed. *Reinforcement Learning. The Springer International Series in Engineering and Computer Science (Knowledge Representation, Learning and Expert Systems)* 173 (Springer, 1992), pp. 5–32. Having mentioned environment, it is important to emphasise that there are worrying environmental costs associated with machine learning. See, for instance, Karen Hao, 'Training a single AI model can emit as much carbon as five cars in their lifetimes', *MITk Technology Review* (June 6, 2019), <https://www.technologyreview.com/s/613630/training-a-single-ai-model-can-mit-as-much-carbon-as-five-cars-in-their-lifetimes/>.

Supervised Learning

Supervised learning is based on a training dataset with input/output pairs as expected answers. A classic example would be spam emails in which an algorithm learns from the sample of emails that are labelled as 'spam' or 'not spam'. The goal of this type of learning is to map the input data to output labels. For example, with new email as the input, what would the predicted output result be? Can it be classified as spam and then moved to the spam mailbox? In mathematical terms, this is expressed as $Y=f(X)$, and the goal is to predict the output variable Y from the new input data (X). But this prediction process relies on classification techniques, for example binary classification (such as 'yes/no', 'spam/not spam', 'male/female') and multi-classification (such as different object labels in visual recognition), which is based on the process of data labelling. This is where inconsistencies arise. Data is categorised in a discrete manner, and there are many elements that might lead to a 'normative' prediction, especially problematic when it comes to complex subjects such as gender, race, and identity, because these operate beyond binary, discrete classification. See Joy Buolamwini, 'Response: Racial and Gender Bias in Amazon Recognition — Commercial AI System for Analyzing Faces'. *Medium* (2019), <https://medium.com/@Joy.Buolamwini/response-racial-and-gender-bias-in-amazon-rekognition-commercial-ai-system-for-analyzing-faces-a289222eeced>.

Unsupervised Learning

Unlike supervised learning, unsupervised learning does not contain a set of labelled data. One of the common tasks with unsupervised learning is 'clustering' (algorithms such as K-means and Hierarchical Clustering). The goal of this technique is to find similarities, providing insights into underlying patterns and relationships of different groups in a dataset using exploratory and cluster analysis. Items in the same group or cluster share similar attributes and metrics. The idea behind clustering is to identify groups of data in a dataset, segregating groups with similar characteristics. It is commonly used in the business and marketing sectors to understand customer preferences so that personalisation and data marketing can be provided by grouping customers based on their purchasing behaviour with regard to certain types of goods. Artists Joana Chicau and Jonathan Reus have developed a performative project 'Anatomies of Intelligence' (<https://anatomiesofintelligence.github.io/>), which uses an unsupervised learning model to develop an understanding of anatomical knowledge and computational learning. In a workshop setting, they suggest that participants think of two features for examining a small image dataset (around fifteen images) — such as 'cuteness' and 'curliness' — and each of the images is rated and sorted according to these features. Each image can then be described by the set of feature values. As a result, several clusters are formed, providing a new perspective on the relations between images in terms of their similarities and differences. It's a simple exercise, but can be scaled up, systematised and automated, for example by

deciding on the number of clusters and calculating the distribution of/distance between data points. This also helps reinforce how algorithms designed to recognise patterns, known as neural networks, operate, being loosely based on a model of the human brain and how it learns to differentiate certain objects from other objects.

Winnie Soon is an artist-researcher and Associate Professor at Aarhus University, Denmark. **Geoff Cox** is Associate Professor and co-Director of the Centre for the Study of the Networked Image at London South Bank University, UK.
