SUPPLEMENTAL MATERIAL

A Primer on Artificial Intelligence

Artificial intelligence (AI) is defined as the study of intelligent agents which can perceive the environment and act as humans do, typically for narrowly defined tasks.[1] This is distinct from 'general artificial intelligence' popularized by the entertainment industry and lay press. Machine learning (ML) involves the scientific study of statistical and mathematical models and algorithms which can progressively learn from data, to achieve desired performance on a specific task. ML can be used in tasks which need automation or standardization, and is well suited to scenarios where humans find it difficult to manually develop a set of 'rules' for these desired tasks. ML can be broadly categorized into supervised learning, unsupervised learning, semi-supervised learning and reinforcement learning.[2]

Supervised learning problems involve tasks where one is provided with a set of input data (predictors) and targets, and the objective is to find appropriate functions to map predictors to the targets for which "known-labels" are available. Such problems can be categorized into problems of regression or classification. When target variables are continuous real values, the supervised learning task(s) is known as a regression problem; when the target variables are categorical variables, the tasks are known as classification problems. Common supervised learning algorithms include well-described statistical and mathematical tools such as linear regression,[3] logistic regression,[4] decision tree,[5] random forest (RF),[6] support vector machine (SVM),[7] k-nearest neighbors (k-NN), as well as non-traditional analyses which are well suited to complex data, in particular artificial neural networks (ANN).[8] Unsupervised learning pertains to tasks of knowledge discovery, where there exist only predictors, and target variables are absent. Examples of knowledge discovery tasks are: (1) identifying the underlying distribution of data, (2) discovering natural grouping/clustering within data, (3) simplifying data by reducing the numbers of distinct data axes (dimensionality reduction). Some commonly used unsupervised learning algorithms are: k-means clustering and hierarchical clustering. Semi-supervised learning tasks typically have similar objectives to supervised learning, but seek methods to utilize unlabeled data since large, labeled datasets are often difficult to obtain. In situations where assigning target labels is very expensive, active learning is used. Active learning focuses on methods that best suggest which unlabeled data should be labelled next, to attain the desired supervised learning task while minimizing the need for labeling. Reinforcement learning is concerned with how intelligent agents learn and perform actions in a dynamic environment by maximizing metrics of cumulative rewards.

Typically, an ML pipeline consists of the following steps: (1) data acquisition, (2) data preprocessing, (3) feature extraction, (4) feature selection, and (5) a supervised/unsupervised/ reinforcement learning process. Several ML experts have stated that optimal feature representation ('feature engineering') is a defining pre-requisite for high ML performance.

Deep learning (DL) is a sub-field of ML, which allows very complex data sets to be matched to useful output labels such as image recognition, ECG diagnosis or other 'ground truths'. DL has been enabled by advances in architecture of computational ANNs which provide several levels of abstraction ('layers' of the ANN) to encode complex data relationships. Deep neural network (DNN), Convolutional neural network (CNN), Recurrent Neural Net (RNN), Deep Belief Network (DBN), and Generative Adversarial N*etwork* (*GAN)* are some notable DL architectures.[9] Broadly speaking, these computational learning architectures were inspired by the organization of biological networks in mammalian visual cortex observed by Hubel and Wiesel in the 1950s for

which they, with Sperry, received the 1981 Nobel Prize in Physiology or Medicine.[10] CNNs are particularly good for analyzing images. RNNs are a class of ANNs which work especially well in analyzing sequential data such as time series or natural text. DBN is a probabilistic generative graphical model composed of multiple layers of latent variables.

References

1. Russell SJ, Norvig P. *Artificial intelligence: A modern approach*. Malaysia; Pearson Education Limited; 2016.
2. Seber GA, Lee AJ. *Linear regression analysis*. John Wiley & Sons; 2012.
3. Hosmer Jr DW, Lemeshow S, Sturdivant RX. *Applied logistic regression*. John Wiley & Sons; 2013.
4. Quinlan JR. Induction of decision trees. *Machine learning*. 1986;1:81-106
5. Breiman L. Random forests. *Machine learning*. 2001;45:5-32
6. Boser BE, Guyon IM, Vapnik VN. A training algorithm for optimal margin classifiers. *Proceedings of the fifth annual workshop on Computational learning theory*. 1992:144-152
7. Goodfellow I, Bengio Y, Courville A, Bengio Y. *Deep learning*. MIT press Cambridge; 2016.
8. Krittanawong C, Johnson KW, Rosenson RS, Wang Z, Aydar M, Baber U, Min JK, Tang WHW, Halperin JL, Narayan SM. Deep learning for cardiovascular medicine: A practical primer. *Eur Heart J*. 2019;40:2058-2073
9. Shaw LJ, Blankstein R, Jacobs JE, Leipsic JA, Kwong RY, Taqueti VR, Beanlands RSB, Mieres JH, Flamm SD, Gerber TC, Spertus J, Di Carli MF; on behalf of the American Heart Association Cardiovascular Imaging and Intervention Subcommittee of the Council on Clinical Cardiology; Council on Cardiovascular Radiology and Intervention; Council on Cardiovascular and Stroke Nursing; Council on Cardiovascular Disease in the Young; and Council on Quality of Care and Outcomes Research. *Circ Cardiovasc Imaging.* 2017;10:e000017. DOI: 10.1161/HCI.0000000000000017
10. Wilson R, Clippingdale S. Self-similar neural networks based on a kohonen learning rule. *Neural networks: the official journal of the International Neural Network Society*. 1996;9:747-763