

Statistical Methods and Artificial Intelligent (AI) with Application to Medicine

Xiao-Hua Andrew Zhou, Ph.D

Chair Professor and Chair, Department of Biostatistics, Peking University

Oct 31, 2019

Definition of Intelligence

- Intelligence can be generally described as the ability to perceive information, and retain it as knowledge to be applied towards adaptive behaviors within an environment or context.
- AI: Any device that perceives its environment and takes actions that maximize its chance of success at some goal.
- While there are many different definitions of intelligence, they all essentially involve learning, understanding, and the application of the knowledge learned to achieve one or more goals.

Functions of AI

- AI is about duplicating what the (human) brain DOES
 - Turing Test
- AI is about duplicating what the human brain SHOULD DO
 - Rationality

AI: From Data to Knowledge

- Artificial Intelligence is the study of ideas which enable computers to do the things that make people seem intelligent
- If we could understand the functioning in health and in disease of the human body in sufficient depth to model the detailed disease processes which disturb health, then, at least In principle, we could perform diagnosis by fitting our model to the actually observable characteristics of the patient at hand.
- Further, we could try out possible therapies on the model to select the optimum one to use on the patient.
- Unfortunately, although biomedical research strives for such a depth of understanding, it has not been achieved in virtually any area of medical practice.

Artificial Intelligent (AI) in Medicine

AI can support both the creation and the use of medical knowledge:

- Human cognition is a complex set of phenomena, and AI systems can relate to it in two very different ways.
- Strong AI are interested in creating computer systems whose behaviour is at some level indistinguishable from humans.
- Success in strong AI would result in computer minds that might reside in autonomous physical beings like robots, or perhaps live in 'virtual' worlds like the information space created by something like the Internet.

Artificial Intelligent (AI) in Medicine

AI can support both the creation and the use of medical knowledge:

- Weak AI: An alternative approach to strong AI is to look at human cognition and decide how it can be supported in complex or difficult situations. For example, a fighter pilot may need the help of intelligent systems to assist in flying an aircraft that is too complex for a human to operate on their own.
- These 'weak' AI systems are not intended to have an independent existence, but are a form of 'cognitive prosthesis' that supports a human in a variety of tasks.

Examples with Medical knowledge Systems

- Agents for information retrieval. Software 'agents' can be sent to search for and retrieve information, for example on the Internet, that is considered relevant to a particular problem. The agent contains knowledge about its user's preferences and needs, and may also need to have medical knowledge to be able to assess the importance and utility of what it finds.
- Image recognition and interpretation.
- Many medical images can now be automatically interpreted, from plane X-rays through to more complex images like angiograms, CT and MRI scans. This is of value in mass-screenings, for example, when the system can flag potentially abnormal images for detailed human attention.

Expert laboratory information systems

- One of the most successful areas in which expert systems are applied is in the clinical laboratory.
- Practitioners may be unaware that while the printed report they receive from a laboratory was checked by a pathologist, the whole report may now have been generated by a computer system that has automatically interpreted the test results.

Expert laboratory information systems

Examples:

- The PUFF system for automatic interpretation of pulmonary function tests has been sold in its commercial form to hundreds of sites world-wide (Snow et al., 1988). PUFF went into production at Pacific Presbyterian Medical Centre in San Francisco in 1977, making it one of the very earliest medical expert systems in use.
- GermWatcher checks for hospital-acquired (nosocomial) infections, which represent a significant cause of prolonged inpatient days and additional hospital charges (Kahn et al., 1993). Microbiology culture data from the hospital's laboratory system are monitored by GermWatcher, using a rule-base containing a combination of national criteria and local hospital infection control policy.

Expert laboratory information systems

- A more general example of this type of system is PEIRS (Pathology Expert Interpretative Reporting System) (Edwards et al., 1993).
- During its period of operation, PEIRS interpreted about 80-100 reports a day with a diagnostic accuracy of about 95%.
- It accounted for about 20% of all the reports generated by the hospital's Chemical Pathology Department.

Expert laboratory information systems

- Laboratory expert systems usually do not intrude into clinical practice. Rather, they are embedded within the process of care, and with the exception of laboratory staff, clinicians working with patients do not need to interact with them.
- For the ordering clinician, the system prints a report with a diagnostic hypothesis for consideration, but does not remove responsibility for information gathering, examination, assessment and treatment.
- For the pathologist, the system cuts down the workload of generating reports, without removing the need to check and correct reports.

Problem #1: Health care is a label desert and the advent of one-shot learning

- Modern artificial intelligence is data hungry.
- To make speech recognition on your Android phone accurate, Google trains a deep neural network on roughly 10,000 hours of annotated speech.
- These annotations, called labels, are essential to make techniques like deep learning work.
- Can AI work well in situations where we may have 1000x fewer labels than we are used to?
- There is already some promising work in this direction. First, it was unsupervised learning that sparked interest in deep learning from 2006 to 2009?, namely, pre-training and autoencoders, which can find structure in data that is completely unlabeled.

Problem #1: Health care is a label desert and the advent of one-shot learning

- More recently, hybrid techniques such as semi-supervised sequence learning have established that you can make accurate predictions with less labeled data if you have a lot of unlabeled data.
- The most extreme setting is one-shot learning, where the algorithm learns to recognize a new pattern after being given only one label.
- Underlying many of these techniques is the idea that large amounts of unlabeled data may substitute for labeled data.
- Is that a real advance? Maybe.
- With the proliferation of sensors, unlabeled data is now cheap. Some medical studies are making use of that already.

Problem #2: Deployment and the outside-in principle

- There are numerous reasons why more expert systems are not in routine use (Coiera, 1994).
- Some require the existence of an electronic medical record system to supply their data, and most institutions and practices do not yet have all their working data available electronically.
- Others suffer from poor human interface design and so do not get used even if they are of benefit.

Problem #2: Deployment and the outside-in principle

- Much of the reluctance to use systems simply arose because expert systems did not fit naturally into the process of care, and as a result using them required additional effort from already busy individuals.
- It is also true, but perhaps dangerous, to contribute some of the reluctance to use early systems upon the technophobia or computer illiteracy of healthcare workers.
- If a system is perceived by those using it to be beneficial, then it will be used. If not, independent of its true value, it will probably be rejected.

Challenges in Gaining Knowledge with Big Data

- Inference: the problem of turning big data into knowledge, where knowledge often is expressed in terms of entities that are not present in the data per se but are present in models that one uses to interpret the data.
- Statistical principles are necessary to justify the inferential leap from data to knowledge
- Overlooking this foundation may yield results that are not useful at best, or harmful at worst.
- In any discussion of massive data and inference, it is essential to be aware that it is quite possible to turn data into something resembling knowledge when actually it is not.
- Moreover, it can be quite difficult to know that this has happened.

Application of Statistics in AI with Big Data

- The field of statistics has developed tools that can address many challenges in artificial intelligent
- Data structure discovery (unsupervised learning)
 - Clustering or graphical modeling
 - Multiple hypothesis testing
 - Subgroup identification for best treatment
- Supervised learning
 - Regression model
 - Classification models

Machine Learning

Machine Learning:

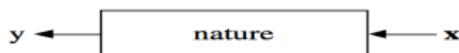
- A collection of methods for automated learning from data
- Many ideas originated in computer science
- Many ideas originated in statistics
- The best machine learning tools have good statistical properties

Statistical Learning

- Statistical learning: tools for understanding data.
- The same as machine learning but term is used to emphasize the statistical aspects

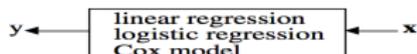
Difference Between Statistical and Machine Learning

Input variables x ; the response variables y :

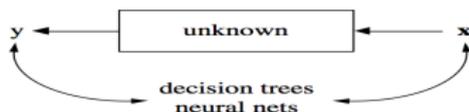


Statistical Approach: Assume a stochastic data model for the inside of the black box.

Response variables = $f(\text{predictor variables, random noise, parameters})$:



The machine learning approach: Find a function $f(x)$ —an algorithm that operates on x to predict the responses y



Breiman (2001).

Contrast between Machine and Statistical Learning

- Some machine learning methods have been used in statistics for decades or longer: For example, Linear regression, logistic regression, and other statistical models. These and many other statistical methods are used to summarize, predict or learn from the data.
- Some non-statistical approaches (e.g., from computer science) bring unique strengths to machine learning:
- Support vector machines estimate only the sign of the classification function rather than the entire distribution.
- Improved robustness, greater ease in addressing non-linearity and high dimensional covariates.
- Some learning tasks not part of traditional statistics

Contrast between Statistical and Machine Learning

Tibshirani

Machine learning	Statistics
network, graphs	model
weights	parameters
learning	fitting
generalization	test set performance
supervised learning	regression/classification
unsupervised learning	density estimation, clustering