PROOF

Current Trends in
Neurology

Mini-Review

The future is here: How machine learning will impact neurology

Jay Hegdé[1,*] and Eugene Bart[2]

[1]Departments of Neuroscience and Regenerative Medicine, Ophthalmology, Culver Vision Discovery Institute, Medical College of Georgia, Augusta University, Augusta, GA; [2]Palo Alto Research Center (PARC), Palo Alto, CA, USA.

ABSTRACT

Machine learning (ML) is a rapidly developing branch of computer science that allows computers to learn complex tasks and behaviors. This enables computers to perform many real-world tasks that have hitherto resisted automation. ML is no longer a matter of futuristic fiction or mere laboratory curiosity. Recent advances have made it possible to tackle significant real-world problems, and ML is already making a significant impact in almost every aspect of modern life: from smart phones to high finance, from fuel injection systems in cars to cars that drive themselves, from computers that speak to robots that help build other robots, and from law enforcement to warfare. In today's world, any smart machine uses some form of ML. Thus, every walk of human life that was once the exclusive domain of extensively trained, highly skilled human professionals is either already being impacted or stands to be impacted by ML. Not surprisingly, ML already has a substantial presence in medicine. In addition to being used for more mundane tasks such as bookkeeping and dispensation, it is also being used to digest vast amounts of cancer research data, to help customize cancer therapies for individual patients, analyze radiological images and other clinical data, and to supervise medical education and training. Applications of ML in neurology are likely to fall under two broad categories: (i) assisting neurological patients with their daily lives (e.g. helping compensate for sensory or motor deficits), or (ii) assisting neurologists in various tasks (e.g. analyzing of neuroimaging data to help make evidence-based decisions customized for each individual patient). Given the unnervingly immense potential of ML, its impact on medicine in general and neurology in particular is likely to increase in the future.

KEYWORDS: artificial intelligence (AI), big data, deep learning, machine learning, neural networks, robotics.

1. Introduction

Since the advent of modern digital computers in the 1940s, computing technology has continuously gained ascendancy in industrialized countries around the world. Today, it is all but impossible to think of a facet of modern society completely free of the influence of computers. In addition to the more conventional computers we use at work and at home, computers of one ilk or another help operate our cars, phones, banks, and grocery stores. Supercomputers, often huge enough to fill whole floors of large buildings, and dedicated data centers that sometimes use as much power as moderately-sized towns, process enormous amounts of data for nation-states, corporations, and other institutions. The list is endless.

Until very recently, however, computers had very obvious strengths and limitations. They were notably good at number crunching, and at rule-based, regimented tasks. Yet they were flummoxed by

*Corresponding author: jay@hegde.us

the real world and could not handle its complex, chancy, messy nature. Thus, computers could gamely handle the fuel injection in your car, but they could not drive the car. Corporations could trust them to handle the tedious tasks of ordering and balancing the books, but they were poor at predicting which products were likely to fare better in the market. Some of the more recent, much ballyhooed supercomputers were good at finding a needle in some haystacks, but not others. For example, computers could screen millions of fingerprints for matches, but were hopeless in finding patterns of phone calls or identifying specific statements within phone calls to help identify potential terrorists. This seemed to be the sole domain of humans. Thus, until recently, computers and humans seemed to nicely complement each other's strengths and weaknesses. Even when a computer, IBM's Watson, beat the reigning world chess champion, it did not seem like such a big leap due to chess being a rule-based game in which all possible moves by either player can, in principle, be calculated in advance. Thus, it seemed like just another number-crunching exercise, but on steroids.

The situation has changed rather dramatically in the last decade or so. Instead of being applied to technically rarefied examples, ML has finally matured enough to tackle some significant real-world problems. Self-driving cars have become fairly common on public roads in many metropolitan areas. Virtual assistants, such as Siri, Alexa, and Cortana etc., have begun to be used by real people to assist with real-world tasks. In February 2011, a newer incarnation of Watson appeared on the popular American television game show *Jeopardy*! and soundly beat two human champions, Ken Jennings and Brad Rutter (Fig. 1). Clearly, this was no mere number crunching exercise. Being a *Jeopardy*! contestant is a uniquely human endeavor. Specifically, the entire match is conducted in conversational English, so the computer had to muster a uniquely human facility: natural language understanding [1]. It had to come up with questions for answers such as, "Wanted for general evil-ness; last seen at the tower of Barad-dur; it's a giant eye, folks. Kinda hard to miss" in the category 'Literary Character AFB', and "This 2-word phrase means the power to take private property for public use;

Fig. 1. IBM's question-answering computer Watson, with fellow contestants Ken Jennings and Brad Rutter on the television game show *Jeopardy*!. A graphical avatar of Watson was shown at the podium, and the actual computer. While successes in chess games and on game shows have garnered Watson much publicity, the impact of machine learning on other fields such as medicine, and by companies and institutions other than IBM, has received far less publicity. "*Jeopardy*!" courtesy Sony Pictures Television.

It's ok as long as there is just competition" as a Daily Double in the category 'Legal "E"s'. Over the two-day match, Watson responded correctly 66 times, including in the case of both the aforementioned cues, and incorrectly nine times.

Watson's shortcomings and quirks during the performance were even more revealing. It could neither understand nor produce human speech. Indeed, in a telling showbiz compromise, Watson received the answers electronically as the human contestants heard it from the host and responded in writing. Watson, unlike its human competitors, never wagered in round numbers. Its final *Jeopardy*! wagers for the two matches were $947 and $17,973, respectively. It had evident trouble with some categories, especially those with short clues containing only a few words. It did not even buzz in for two whole categories, 'Actors who direct' and 'One buck or less'. In a delicious bit of irony, it failed to come up with a single correct response in the category 'Also on Your Computer Keys'. In the final *Jeopardy*! round on the first day, the category was 'U.S. Cities', and the clue was "Its largest airport is named for a World War II Hero; its second largest for a World War II Battle". The correct answer was Chicago, but Watson famously came up with "What is Toronto?????". What these shortcomings reveal is that the human "common sense" that we take for granted is still elusive for Watson despite its otherwise impressive performance.

There are other computing systems that are specialized to perform one or another of a highly skilled, real-world task that was once thought to be the exclusive domain of trained human experts. Most of us are familiar with Siri, Cortana, or Alexa, intelligent personal assistant apps that can understand and produce human speech. Self-driving cars and face recognition technology are two other well-known computing applications. What all of these computing systems have in common, and what makes them successful in complex, real-world tasks is a new style of computing called ML [2-4].

Several factors enabled this recent progress. First, ML algorithms matured and have become better understood. Second, enormous computational resources have become readily available. Finally, a combination of widely available sensors (such as cameras) and transmission technology, which allow for capturing, collecting and storing large amounts of training data necessary for ML algorithms, is now available to improve ML applications.

In this review, we will briefly review the state-of-the-art of ML, and how it is already starting to make an impact, and is likely to make a more substantial future impact, on medicine in general and neurology in particular. Instead of attempting an exhaustive review of all ML applications in neurology, we will highlight a few developments and trends that are crucial to understanding this trend.

2. What is machine learning?

Traditionally, computers have dealt with situations where both the task and the exact procedure for performing the task are precisely specified. A familiar example is the addition of numbers. For any two numbers, the correct result of adding them up is precisely mathematically defined. In addition, a formal procedure (algorithm) used for performing addition is well-known (recall the elementary school algorithm for addition with carry). Each step of this algorithm is also precisely defined and self-contained (i.e., requires no additional knowledge). It is possible to prove that executing the algorithm will in fact always produce the correct answer. Together, these properties make it easy and efficient to explicitly program the computer to perform the task in question.

In contrast, many tasks humans face involve significantly more ambiguity and vagueness. Consider, for example, the task of driving a car. The first source of ambiguity stems from the absence of a precise definition of what constitutes "good driving". Indeed, it is possible for two reasonable people to disagree on an evaluation of someone's driving or on "best practices" recommendations. The second source of ambiguity stems from the fact that even when there is an agreement on what defines "good driving", there is no formal procedure for accomplishing it. Instead, vague instructions are used by human drivers, such as "adjust speed to conditions" or "exercise caution". These instructions generally require a significant amount of additional knowledge to interpret correctly. That knowledge is often called "common sense" and is usually completely left out of human conversations. It proved extremely difficult to endow computers with such common sense, as illustrated above by Watson's shortcomings. In addition, even in the presence of common sense knowledge, the previously stated instructions are themselves ambiguous (for example, it is easy for two reasonable people to disagree as to whether a particular speed is appropriate for the current conditions).

ML studies algorithms that attempt to solve such vague, ambiguously specified tasks. Due to a crisp, well-defined solution strategy being unavailable for such tasks, it has turned out to be extremely difficult to explicitly program computers to perform such tasks. Instead, ML algorithms rely on observing a few "training examples" of the task in question, and learning to generalize from these examples in a meaningful manner based on feedback about their performance. A familiar example of ML is where a virtual assistant (VA) like Siri or Cortana learns what we sound like by making us speak a few times. ML reflects the similarity of the human learning process by performing a task and learning from examples (although the underlying mechanisms are often different).

3. Types of ML tasks

ML tasks, as well as approaches to solving them, are usually subdivided into reinforcement learning, supervised learning, and unsupervised learning (for more detailed reviews, see [2-9]). These subdivisions are useful for understanding and analyzing the field and are described below. However, it should be noted that these subdivisions are neither precise nor mutually exclusive. In addition, solving practical real-world tasks often requires crossing these subdivisions.

3.1. Reinforcement learning

Reinforcement learning (RL) is the broadest and most general category. RL can assist in navigating and overcoming everyday problems. RL problems consist of the following components:

- An "outside world", whose state is generally not known precisely.
- An "agent" that can
  - Perform observations on the world to obtain information (usually only partial) about its current state, and
  - Perform actions that may change both the agent's and the world's state.

In the above "car driving" example, these components would correspond to the following:

- The "agent" would correspond to the driver.
- The "outside world" would include the car being driven, as well as other cars on the road, the road itself, and any relevant surroundings (for example, air clarity, position of the sun, cloud cover, potential obstacles on the pavement, and so on).
- The "observations" the agent can perform would include looking at the car's dashboard, listening to the engine noise, and noticing vestibular sensations, as well as visually observing the world outside the car windows and listening to road sounds.
- The actions would correspond to any control inputs for the car, such as pressing pedals, turning the steering wheel, operating signals or other accessories, etc.

Other examples of tasks well-suited for the RL framework include other daily tasks such as washing hands, cooking a meal or loading the dishwasher, as well as highly skilled and specialized tasks such as performing a surgery. Potential applications to neurology therefore range from assisting patients with their daily life to assisting surgeons with complex procedures.

A wide range of approaches can be used to approach RL tasks. One of the earliest approaches involved providing "rewards" to the agent after a favorable state was reached (for example, when the desired destination was reached in a car or when a meal was successfully cooked). RL reflects the similarity of this methodology to classical operant conditioning. However, multiple other methods have been used as well, such as "imitation learning" (when the agent observes and tries to mimic the behavior of another agent, such as a skilled human).

3.2. Supervised learning

Supervised learning (SL) describes a class of simplified decision tasks that involve producing a label for a given (query) data point. A typical example is looking at a photograph of an animal and determining whether that animal is a zebra or a donkey. Depending on the label type, SL tasks can be further divided into several categories:

- When the label only has two possible values (such as "0" and "1" or "zebra" and "donkey"), the task is called "binary classification".

- When more than two values are possible (for example, if any of 100 types of animals need to be identified in a photograph), the task is called “multi-class classification”.
- When the label is a continuous value rather than one of a discrete set of possibilities, the task is called “regression”. An example of a regression task is estimating the distance to an object in the photograph.
- When the label is a complex object (such as a graph), the task is called “structured learning” or “structured prediction”.

Technically, all of these tasks could be considered special cases of the RL paradigm described above. For example, to cast the “zebra vs. donkey” task in RL framework, we could consider the photograph to be “the world”, any information extracted from the photograph to be the agent’s observations, and the decision to classify the photograph as “zebra” or “donkey” to be the agent’s actions. Therefore, in principle, the techniques of RL could be applied to solve SL problems as well. However, in practice, more specialized techniques are usually used.

Many applications of SL could be envisioned for neurology. For example, when doctors assess a patient at a high level, they perform the binary classification task of labeling the patient as “healthy” or “requiring follow-up”. Also, when making a specific diagnosis, this could be classified as a multi-class classification problem, whereas when writing a radiology report, it would be classified as a structured prediction problem.

3.3. Unsupervised learning

In unsupervised learning (UL), the task is to somehow “describe” a collection of data in the absence of any labels. A familiar example of an UL task is clustering, i.e. grouping together data points that are “similar”. Another relevant example is “anomaly detection”. In this task, the goal is to infer a normalcy model from the examples and identify those examples that deviate most from “normal”.

Applications to neurology include organizing large amounts of data, such as clustering a long-term EEG recording to reveal several major modes of brain activity or automated monitoring of a video feed from an elderly patient’s home to detect a change in behavior that could indicate an episode of a disease.

4. Approaches to solving ML tasks

Although a thorough review of ML techniques is beyond the scope of this paper, we will give a few examples that illustrate some common themes in ML.

4.1. Supervised learning

For simplicity, we will start with SL methods. Consider the problem of observing an image and classifying (i.e., recognizing) the object in the image as either a “zebra” or a “donkey” (see Fig. 2A).

In the classical SL pipeline, the raw input data is first mapped into a feature space by a process called “feature extraction”. These features are properties of the input images that are easily expressed in numerical form and help distinguish between categories. For the example task, the presence of stripes seems like a reasonable way to distinguish between the categories. Therefore, we might calculate two features from each image, “number of horizontal stripes” and “number of vertical stripes”. Each image can now be described by these two properties and thus plotted on a 2-dimensional scatter plot (Fig. 2B). The location of each point is determined by the features, and the symbol used for each point represents the category (with the circles representing donkeys, and the dashes representing zebras).

As can be seen, it becomes quite easy to draw a line on the scatter plot, such that all circles (corresponding to donkeys) are on one side of the line, and all dashes (corresponding to zebras) are on the other. This separating line (called “separating hyperplane” in higher dimensions) can be used to classify new images into the two categories. For example, the image represented by the cross with a question mark in Fig. 2B is likely of a donkey because it is on the “circle” side of the separating line.

To summarize, SL is subdivided into two phases. In the first phase (actual “learning”), a large labeled dataset is used, called a “training set”. Features are extracted from the data points (usually more than two and often, in hundreds or thousands). A separating hyperplane is identified and recorded. Often, this process is very computationally intensive and requires specialized server hardware. This separating hyperplane is the output of the learning phase. In general, the output of the learning phase is called a “classifier” because it can be used to classify new data points into categories.

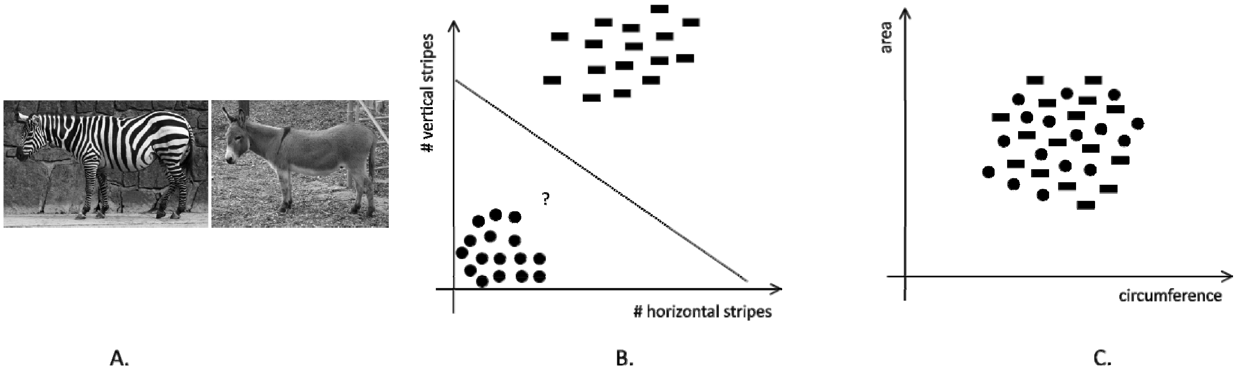

Fig. 2. The problem of classifying objects: An illustrative example. In this and the subsequent figures, the circles denote donkeys, while the dashes denote zebras. See text for details.

In the second phase, what the classifier learned during the first phase is used to assign putative category labels to new, unlabeled data points. To accomplish this, the same set of features is calculated for each new data point. The classifier then determines which side of the separating hyperplane the data point belongs. Compared with learning, this process is usually much less computation-intensive; often, lightweight devices such as cell phones or cameras can run classifiers (although usually not learn them).

The accuracy of the final classifier depends on how well-separated the categories of interest are in the feature space. In Fig. 2B, where the categories are well-separated, the separating line can be expected to produce good results. The reason for this fortunate separation is a good choice of features, guided by the operator's intuition about what properties are likely to help distinguish the categories.

When a poor set of features is used, the accuracy may be expected to drop significantly. This is illustrated in Fig. 2C where we have used "area" and "circumference" as features to describe object shape (rather than presence of stripes to describe their texture). As can be seen, the different symbols (corresponding to different animals) are completely mixed up, making separation impossible.

Figs. 2B and 2C show two extremes (complete separation vs. no separation). Several intermediate configurations often occur in practice as well, illustrated schematically in Fig. 3. In Fig. 3A, the categories are still separable, but not by a straight line, necessitating a different learning algorithm. While such algorithms are available, they come with their own problems and limitations. In Fig. 3B, the categories are only partially separated. This situation is analogous to that arising sometimes in medical diagnosis: some cases (those above the dashed line in Fig. 3B) clearly belong to one category, while many other cases are ambiguous and necessitate further testing. In the context of SL, the analogue of "further tests" would be to calculate additional features and use them to achieve better separation.

5. Importance of choosing the right features

As the figures above illustrate, designing appropriate features is one of the most important tasks in ML. Generally, the operator tries to satisfy two goals: (i) preserve as much information from the original data as possible, and (ii) express it in a manner that makes separation easy (see Figs. 2-4).

Again, two extremes are possible. At one extreme, no actual design takes place and the data is fed to the algorithm as is. For example, in the image classification task above one could simply use the color of each pixel as the feature set (for a 100 x 100 pixel image with 3 color channels per pixel, this would create 30,000 features). While, on the whole, this feature set contains all the information present in the original image, the amount of information in each individual feature is so small that many algorithms would be unable to make use of it (see Fig. 3A).

At the other extreme, a small set of features could be designed that makes the task trivial (see Fig. 2B). While desirable, this is usually difficult to achieve in practice. The reason is that it is often non-obvious

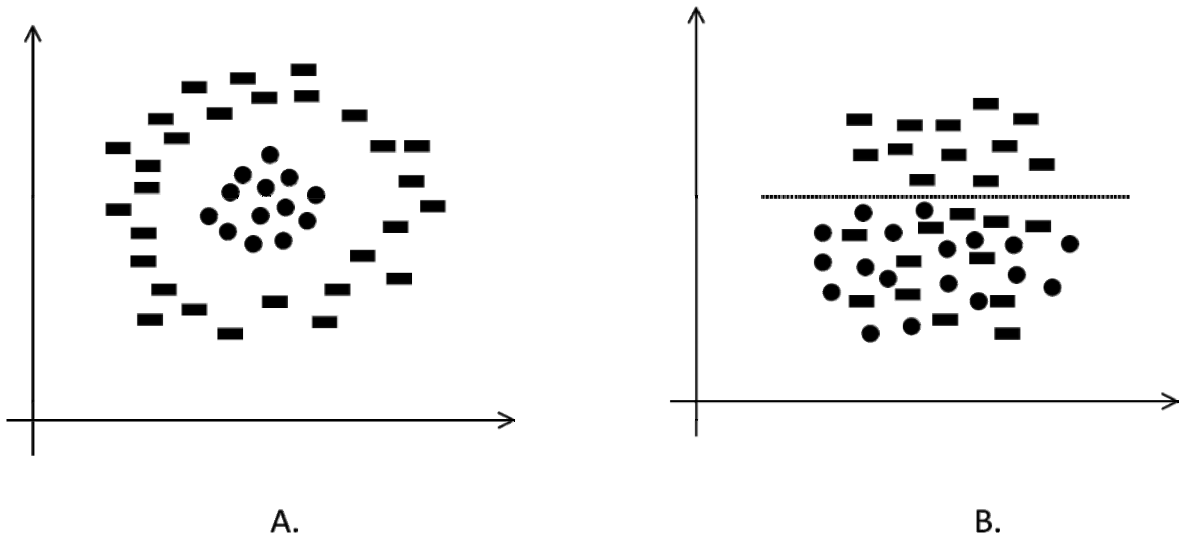

Fig. 3. The choice of features may result in a variety of configurations of data points in the feature space. Some of these configurations may present additional problems or challenges for the learning algorithm. See text for details.

even to human algorithm developers what features are useful for distinguishing two given categories (for example, explaining how you would distinguish between voices of several close friends). This makes detailed feature design time-consuming and expensive.

Another issue that complicates developing SL algorithms is that a large number of features is often necessary (hundreds, thousands, or more). With only two features, it was easy to visualize the task, determine what went wrong (see Figs. 2 and 3), and take the necessary steps to fix the issue. In higher dimensions, visualizing data is extremely difficult, and many of our intuitions (gained from familiarity with 2D and 3D spaces) fail completely in higher dimensions. For example, points in higher-dimensional spaces are generally very sparse, making it harder to generalize (see Fig. 4 for illustration). Thus, this high dimensionality makes diagnosing and solving problems difficult.

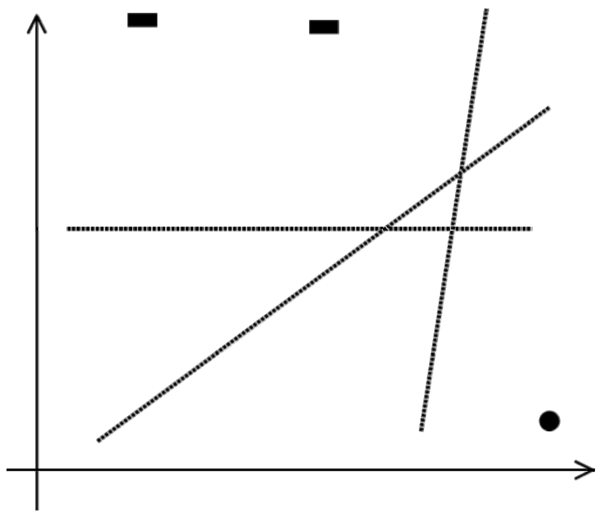

Fig. 4. It is difficult to select the appropriate classifier in a sparsely populated space. While each of the three dashed lines above separates the circles from the dashes, it is difficult to determine which one will generalize best to additional data. See text for details.

6. Deep learning

Many recent success stories of ML involve a family of approaches called “deep learning” (DL). For example, a DL machine recently beat the world champion at the game of Go [10], a highly significant achievement that remained out of reach for the past 50 years.

These DL approaches were inspired by early studies of the brain function. The building blocks of a DL machine are called “neurons”, and they simulate a few basic properties of biological neurons. This includes the ability to integrate inputs from a few simulated “synapses” and, under the right conditions, to simulate "firing" by outputting appropriate signal into outgoing connections. Many (millions or even billions) of these units are combined into an artificial neural network, which is then trained to perform the necessary task. Just like with biological brains, artificial neural networks are extremely flexible and can be applied to a broad range of SL, UL, and RL problems.

In DL, the artificial neural network is arranged in layers. These layers can be thought of as similar to the hierarchy of layers in biological brains that effect "feed-forward" processing (for example, in the visual cortex). Akin to their biological counterparts, the layers of the artificial neural network enable it to learn a hierarchy of increasingly complex features. This allows a deep network to take raw data (for example, individual image pixels) as input and learn to represent it in a highly sophisticated manner automatically, without the need for explicit feature design by a human operator. Thus, to some extent, deep networks are capable of automated feature design. This significantly alleviates some of the feature design problems mentioned above and explains much of the recent success of DL. For a further review of DL, see [8,11].

7. Explainability

As ML is applied to increasingly consequential real-world problems, a significant issue (dubbed "explainability") is making the behavior of these complex machines understandable to their human operators or collaborators. Most viewers of *Jeopardy*! do not have an intuitive understanding of why Watson worked well on some questions and poorly on others, but this lack of understanding made the experience even more entertaining. Yet if similar technology were to be applied, for example, to automated radiogram screening, understanding the failure modes of the system would be critical, as any missed detections would potentially endanger human lives. Thus, it is important for ML systems to perform not only with high accuracy, but also in a manner which their human operators can readily and intuitively grasp and predict. This is a current area of active research in ML.

8. Applications to medical decision-making

In a real sense, clinical problems represent a microcosm of real-world problems. No two patients are exactly alike, nor is the same patient exactly the same at different times. Clinical decision-making at every level involves sifting through a large amount of information of varying degrees of relevance to the case at hand.

One of the obvious ML applications useful to memory is to help cope with the so-called 'big data' problems. For instance, one specific application of IBM Watson at the Memorial Sloan Kettering Cancer Center in New York City is to sift through vast (and constantly enlarging) clinical research literature in lung cancer care to find best treatment options for a given cancer patient, based on the individual patient's clinical profile [12]. Another comparable application of Watson scours the research literature to help identify RNA-binding proteins (RBPs) altered in amyotrophic lateral sclerosis (ALS) [13] (also see [14-17]).

A different set of ML applications involve the processing of medical imaging and image analysis [18]. For obvious reasons, such applications are increasingly prevalent in diagnostic radiology and pathology [19, 20].

9. Some illustrative applications of ML specific to neurology

9.1. Treatment of paralysis: "Mind reading" machines

Fig. 5 schematically illustrates an actual case in which a brain-machine interface (BMI) helps a quadriplegic patient move her arm just by imagining moving the arm. Briefly, a 96-channel microelectrode array was surgically implanted near the 'hand notch' gyral region of the left primary motor cortex of a 24 year-old male patient. The array continuously reads and transmits the neuronal activity data to a relatively small external data processing module running ML implementations (see [21] for details). The module converted the motor planning data collected from the neurons into motor impulses usable by the muscles of the forearm in real terms. The motor impulses were conveyed to the patient's right forearm muscles through a custom built high-resolution sleeve that stimulated the relevant muscles. Using this system, the patient was able to make specific, planned hand movements such as grasping, manipulating, and releasing objects with his hand just by thinking about performing the action.

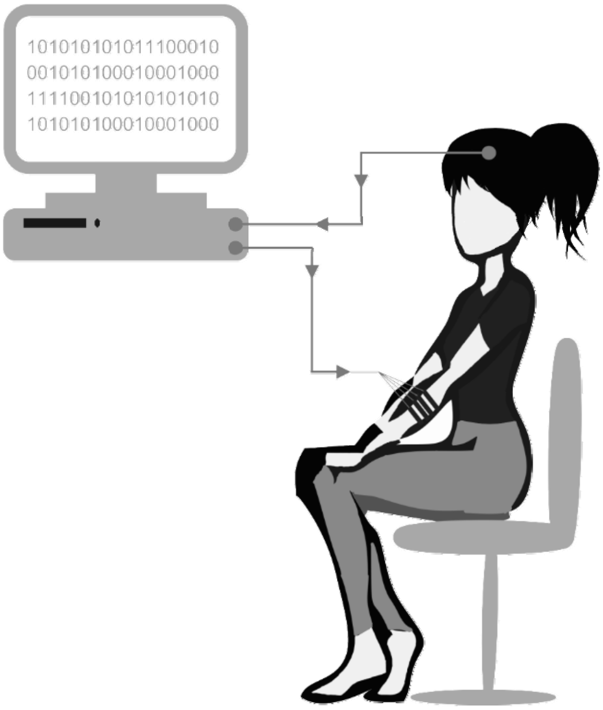

Fig. 5. 'Mind reading' machines. This figure schematically illustrates how a BMI helps an actual quadriplegic patient move his arm just by thinking about moving it. See text for details.

Related advances have been made in other cases using comparable ML setups. For instance, human patients, as well as monkeys, have been able to move robotic arms by similarly planning the corresponding volitional limb movements [22-24]. Recently, a similar system helped a paralyzed man to use a mind-controlled robotic arm to haptically feel objects [25, 26]. In this case, the microelectrode arrays from the motor cortex gathered the motor planning data necessary to control the robotic arm, while a second array of stimulating electrodes delivered the haptic information from the robot's fingers to the patient's somatosensory cortex.

9.2. Applications to evidence-based medicine

In the neurological context, a discrimination map typically refers to a map of the brain created using ML methods (for reviews, see [27, 28]). It is a classification tool that can be consulted to distinguish, or discriminate, between two or more patient populations. The intended use of a discrimination map is illustrated by a recent neuroradiological study by Colij *et al.* [29] (see Fig. 6).

Multiple previous studies using positron emission tomography (PET) or arterial spin labeling (ASL) have found localized cerebral hypoperfusion in patients with Alzheimer's disease (AD) compared to control subjects. Patients with mild cognitive impairment (MCI) also show hypoperfusion, but at levels intermediate between AD patients and control subjects. Obviously, it would be clinically valuable if perfusion levels could be used as a biomarker for the early diagnosis of AD and for prognosis for patients with early AD diagnosis. On the other hand, how can such data on a group biomarker help in making clinical decisions about a given individual patient? Note, incidentally, that this is typical of the quandary faced by all clinicians in evidence-based medicine.

Colij *et al.* [29] used ML methods to compute discrimination maps that can help in multiple binary clinical decisions. For instance, they showed that perfusion levels in the bilateral parietal lobe and hippocampus, indicated by the colored regions in Fig. 6A, can differentiate between patients with AD *vs*. patients with subjective cognitive decline (SCD) with 89% accuracy (area under the receiver operating characteristic curve, 0.93). Even more importantly, the weighting of the various brain voxels in these regions, indicated by the color-coding in Fig. 6, can be used as a tool to make diagnostic decisions about a given individual patient. Suppose the neurologist must decide whether a given individual patient must be diagnosed with AD or SCD. If the measured perfusion levels of the given patient, appropriately weighted by the colored voxels in Fig. 6A, exceed a predetermined criterion value, then the patient can be 'classified' as an AD patient with 89% accuracy. Otherwise, the patient can be placed in the SCD category. Discrimination maps for making related binary decisions are shown in Figs. 6B and 6C.

In the above case, the neurological markers are the structural changes within individual brain regions. Since brain regions are extensively interconnected in the healthy brain, changes within regions can be expected to cause, or be caused, by connectivity across brain regions. There is a vast amount of group data that demonstrate the changes in brain connectivity patterns in various neurological conditions (for reviews, see [30-32]). Once again, tools had been hitherto lacking to help translate such group data to the clinical care of individual patients, in part because of the familiar problem of

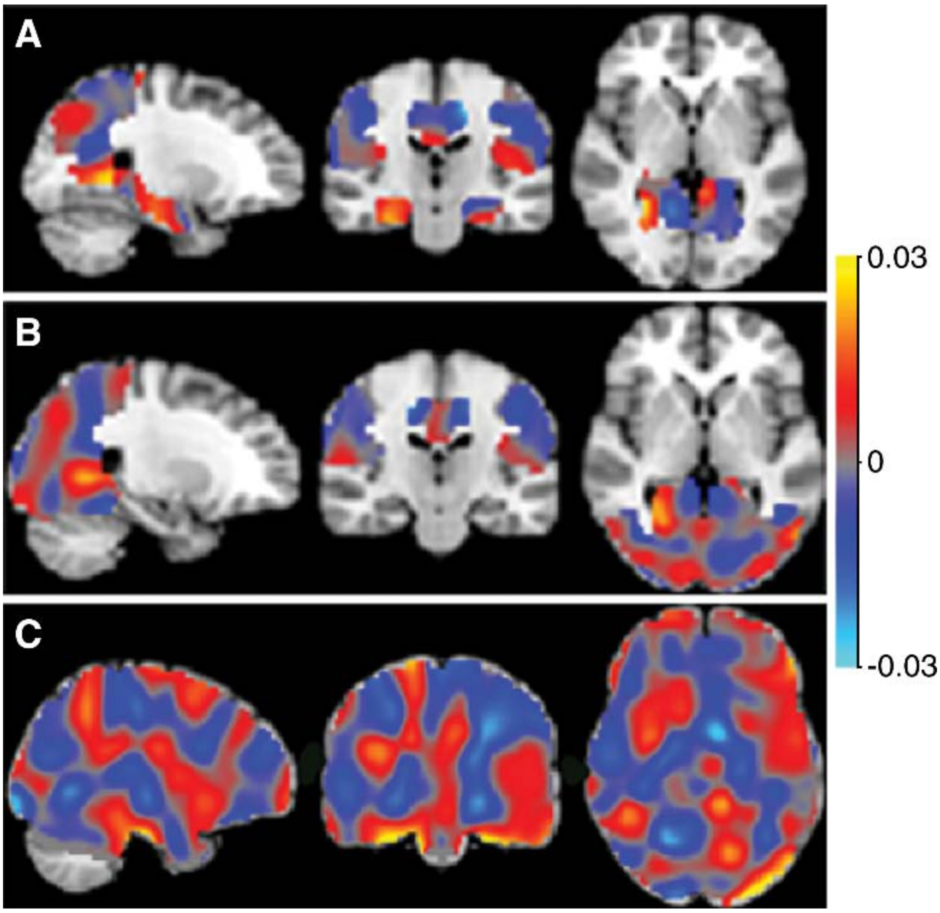

Fig. 6. Discrimination maps. The colored regions represent brain voxels, the perfusion levels of which can be diagnostic in various AD-related conditions. The weights assigned to each relevant voxel is color-coded according to the color bar on right. Each panel denotes regions diagnostic in one pairwise diagnostic decision. (**A**) AD versus SCD. (**B**) AD versus MCI. (**C**) MCI versus SCD. See text for details. Reproduced from Collij, L.E., Heeman, F., Kuijer, J. P., Ossenkoppele, R., Benedictus, M. R., Möller, C., Verfaillie, S. C., Sanz-Arigita, E. J., van Berckel, B. N., van der Flier, W. M., Scheltens, P., Barkhof, F. and Wink, A. M. 2016, Radiology, 281, 865 with permission from Radiological Society of North America.

multivariate data with substantial intra- and inter-patient variability along each variable. ML tools are ideally suited to aid optimal clinical decision-making when the evidence is 'messy' in this fashion. For instance, recent studies have indicated that in the case of Parkinson's Disease (PD), ML methods can use group data on resting state functional connectivity data to make personalized predictions about the progression of motor and non-motor aspects of the disease, and the effectiveness of various interventions such as l-DOPA treatment, exercise therapy, or deep brain stimulation (DBS) (for a recent review, see [33]).

Note that the ML applications in most of these instances are compact enough that they can be run on any standard laptop or tablet [28, 34, 35]. In the aforementioned study [29], the application 'learned' from a few dozen patients from each category which brain regions have the greatest diagnostic value in a given binary diagnostic task, and the optimal weighting of the individual voxels in diagnostically valuable regions. This helps highlight the fact that even such small-scale ML applications, if and when approved by the Food and Drug Administration (FDA), can be of substantial clinical assistance to the neurologist.

9.3. Assistive technologies

The speech-generating device used by the Cambridge physicist Stephen Hawking is perhaps one of the best known smart assistive technologies [36]. Hawking operated this device using a single cheek muscle. Such devices can be operated by the patient using just about any spared motor functionality, including eye movements. This provides a straightforward way of customizing assistive technology for the individual patient. Such smart technologies are versatile tools that can perform functions that neither the patient nor any caregiver can easily perform [37-39].

9.4. Computer-assisted neurological rehabilitation and assisted living

Robots or 'cyber-physical systems' have made their presence felt in assisted living. In addition to fairly simple mini robot vacuums, there are robots that can act like phone caddies for patients in assisted living conditions, whereby they find the patient and bring the phone to him or her when someone calls [40]. Robots can also assist in neurological rehabilitation through therapeutic routines that can be custom-programmed to suit the rehabilitative needs of individual patients [41].

9.5. Closed loop neurostimulation

DBS therapy has been shown to be effective for the treatment of essential tremor, dystonia, and PD [35, 42-45]. Currently approved systems are 'open-loop', in that the stimulating signal is generated in the device and delivered to the brain, and that the device itself receives no information whatsoever from the brain. Recently, 'closed loop' stimulation systems (implantable pulse generator, or IPG) have been developed that can sense the relevant neurophysiological signals and can dynamically adjust the stimulation parameters. ML algorithms carry out the computations that underlie such adaptive brain-computer interfaces (BCIs).

In the case of PD, for instance, the neurophysiological signal monitored by the IPGs are the local field potentials (LFPs) of the basal ganglia measured using surgically implanted, high-impedance electrodes. What the IPGs do with this signal depends on the desired therapeutic goal that the neurostimulation is intended to maintain or restore for the given patient. Given a therapeutic endpoint, ML algorithms continuously decode the 'state' of the relevant neuronal ensemble from the neurophysiological signal, and continually adjust the stimulation as needed 'on the fly' (for reviews, see [46, 47]).

9.6. Visual prostheses

Multiple prostheses are currently approved by the FDA to help restore visual function in visually impaired patients. However, all of these devices are designed to address visual impairment at the level of the eye. No devices addressing purely cortical blindness have been approved to date, self-evidently because, as complex as the eye is, the brain is an infinitely more complex organ, and creating effective prostheses clearly await a synergistic effort from neurological, surgical, and ML experts.

Some of the visual prostheses are small integrated circuit chips that are surgically implanted in the retina (for reviews, see [48-52]). These integrated circuits (ICs) have built-in sensors that sense the light, processors to convert the light signal to electrical signals, and electrodes to convey the signal to the downstream neurons. The processors mimic the retinal processing through fairly simple ML algorithms implemented by the IC [52]. New models also feature output verification and wireless technology for closed-loop systems [50, 53].

In some cases, e.g., when the fibers of the optic nerve are not sufficiently healthy, it is advisable to bypass the oculothalamocortical pathway altogether, and to convey the visual signal directly to the brain. Such devices feature an external video camera. The signals from the camera are converted to brain-compatible electrical signals using smart (think ML-based) processing devices, and are conveyed to the brain through a small (typically $<10\ mm^2$) microelectrode array (see Fig. 7A; also see [54]), typically implanted in the foveal or parafoveal representation of the primary visual cortex. Current iterations of such devices can restore enough of visual function to allow the patient to carry out many everyday visual tasks, such as sorting high-contrast socks.

9.7. Visual prostheses that piggyback on other sensory modalities

It is important to note that visual 'prostheses' do not necessarily entail stimulating electrodes surgically implanted in the brain. One can do an end-run

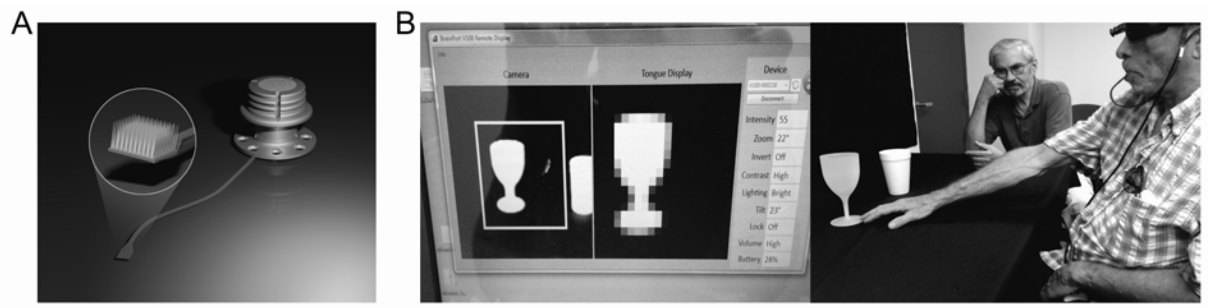

Fig. 7. Some visual prostheses. (**A**) Implantable stimulating microelectrode arrays that convey the neurally encoded visual signal to cortical neuronal ensembles, such as the primary visual cortex or somatosensory cortex. (**B**) Seeing through one's tongue with the BrainPort rehabilitative device. The visually impaired patient wears a pair of sunglasses fitted with a video camera. The video signal from the camera is processed in a small, hand-held device and is converted to electronic signals conveyed to the tongue through a removable transducer worn on the tongue by the patient like a lollipop. After some practice, the patient can perceive the visual world with enough detail to do most everyday tasks. This modus operandi of this device is sometimes described as 'tasking the light' [55]. However, this is a bit of misnomer, because the electrical signals in questions are taken up by the tactile sensors, not the taste buds, in the tongue. Picture in *Panel A*, courtesy of UPMC/Pitt Health Sciences, used with permission [26]. Picture in *Panel B*, courtesy of Wicab; used with permission [66].

around the invasive surgical procedures and, in fact, bypass the early visual processing pathways altogether by using other sensory systems to convey the visual signal to the brain. One such FDA-approved device, BrainPort, gathers the visual signal through a small, wearable video camera which processes the signal using a small, portable (usually hand-held), smart processor, and conveys the visual signal through a small, removable, electrical transducer worn on the tongue like a lollipop that delivers the visual signal as corresponding visual impulses to the tongue [55] (also see Fig. 7B). Use of this device takes a certain sensory learning, or 'getting used to' on the part of the visually impaired patient to perceive the visual world in this fashion. Yet patients can learn this easily and well enough to do most everyday tasks, and even some arduous tasks like rock climbing. The primary reason why this rerouting works is because, with the possible exception of olfaction, the cortical anatomical substrates of the various senses are fundamentally similar. Therefore, such cross-modal therapeutic solutions represent a potentially huge growth area for ML-based rehabilitative solutions.

10. Future prospects

To understand the limitations and future prospects of ML, it is instructive to examine a fundamental theorem colloquially known as the "No Free Lunch" theorem [56, 57]. The theorem considers what happens when the set of problems to be solved is not *a priori* circumscribed. The theorem's conclusion is that under these conditions, any learning algorithm is as good as any other learning algorithm (including random guessing). In other words, if someone took the most modern and sophisticated learning algorithm (such as DL), tested it on all conceivable learning problems, and counted the number of successes, that number would be no higher than the success count of the algorithm that does not learn anything and outputs a random guess each time.

The theorem initially appears to contradict not only the recent successes of ML, but also the ability of biological organisms to learn and adapt. To reconcile the two, recall that in real life, both machine and biological learning systems are not dealing with unconstrained, arbitrary problems. For example, all practical visual classification problems share the common property called smoothness: nearby pixels in the image are likely to come from the same object and have similar color, depth, and surface normal. Existing visual recognition algorithms all depend on this property to achieve better-than-chance success rates. Consequently, those algorithms will perform at worse-than-chance levels on other visual recognition problems (for example, those that do not have smoothness), but that is not an issue since these problems do not seem to appear in practical real-world tasks.

Several conclusions can be made from the preceding discussion. First, despite its strong mathematical underpinnings, ML is an empirical discipline. Successes occur when an algorithm can be tuned to a family of related real-world problems; since fundamentally no algorithm is better than any other, such successes have to be evaluated empirically. Second, algorithms are unlikely to be universal; each new family of problems will require a carefully selected and specifically tuned algorithm for its solution.

In addition to the fundamental limitations of ML described above, some more pragmatic, empirical, and potentially temporary limitations can be outlined. Perhaps the most obvious is that computers, even endowed with ML technology, still lack even basic human "common sense". One example is that a computer trained on the binary classification task of distinguishing between zebras and donkeys will force either a "zebra" or a "donkey" label on every image, even if nothing in the image resembles an animal in any way. A human trained on a similar task would naturally label such images as "neither", but a computer would have to be trained explicitly for that expression. Another example is that a human trained on a task invariably learns much more than strictly necessary to accomplish the task; their learning is considerably "deep". As a result, humans will generalize much more readily to new conditions or new tasks. In computer algorithms, even those considered DL, the learning is still not nearly as deep. Continuing the animals example, a computer that learned to distinguish zebras from donkeys (see Fig. 2) would be at a total loss if it had to determine the approximate age or gender of the animals. In contrast, humans would adapt much faster, even if they were not warned about these additional tasks. Several additional examples showing the lack of common sense have been outlined in the Introduction.

There seem to be no fundamental limitations that prevent an artificial system from acquiring and using "common sense" similar to humans. However, accomplishing this in practice so far has proved difficult. Until this problem is overcome, the lack of common sense will continue to put certain limitations on the applications of ML. First, the tasks that computers solve will be somewhat limited in scope. One might have imagined that a virtual assistant that is able to order pizza from Domino's would also be able to order one from Pizza Hut or to book a ride on Uber. In reality, each of these tasks still have to be programmed separately, and it is likely to be a while before a virtual assistant can truly display some general proficiency with ordering things online. Second, adoption is likely to be fastest for extremely structured scenarios (for example, the earliest adopters of robotics were factories that could exercise extremely precise control over placement of machinery, spare parts, and personnel). Outside of these controlled environments, one is much more likely to see ML used in tasks that "assist" a human operator, such that errors are not critical. A familiar example is automatic face detection on most modern cameras and cell phones: it is quite convenient when it works well, but any errors are easy to circumvent. Critical tasks in uncontrolled environments (such as performing a surgery, given that every human's anatomy is slightly different) are unlikely to be performed completely autonomously for some time. Third, while we have come to expect perfect performance from traditional automation, computers that utilize ML are likely to have "quirks" or errors in their behavior. As their human users, we will have to learn not to expect perfection, but rather anticipate and work around these quirks (much like when working with fellow human beings). One of the goals of research on explainability (outlined above) is to make these quirks easier to understand and predict for human beings.

Biological systems provide an "existence proof" that successful learning is possible for a wide variety of problems. In principle, the field of ML can make some headway, as it has so many times in the past, by studying and mimicking these biological systems. Unfortunately, significant gaps exist in our understanding of biological learning systems as well [58-60].

11. What factors are likely to facilitate or hinder application of machine learning to neurology?

Several factors facilitate the development and adoption of ML algorithms for a given domain. One factor is the availability of good features (or a method to learn them). For example, detecting the presence and severity of tremor in PD is likely to be relatively easy given that it is a simple periodic motion readily amenable to Fourier transform analysis. In contrast, detecting presence and severity of coprolalia (involuntary swearing) in Tourette syndrome is likely to be much harder because it is difficult to characterize precisely the contexts in which swearing is socially acceptable (see [61]). A second factor to consider is the ability to limit the scope of the problem or to control the domain. For example, access control based on facial recognition is a relatively "easy" problem. It is very precisely defined (either the face belongs to one of a few authorized users or it does not), and the users are cooperative (the user wishing to gain access will voluntarily pose in frontal view, remove sunglasses, or perform other necessary steps to reduce variability and simplify the task for the computer). In contrast, screening for "suspicious behavior" at an airport checkpoint is much more complicated. The passengers are unlikely to pose for cameras, any relevant facial expressions (such as fear or anger) will be fleeting, and those who have malicious intent will try to mask their expressions, behaviors, and intent. In addition, the task is extremely open-ended as the goal is not to detect any specific facial expression, but rather to evaluate the passenger's mental state and its appropriateness for the current context. Availability of good-quality, large datasets is an important factor for consideration. Facial recognition and recognition of handwritten digits are some of the most studied problems in ML because large datasets were collected and became widely shared several decades ago. In contrast, automated interpretation of brain X-ray images is likely to face many more hurdles. Collecting radiograms is much harder than collecting photographs or handwriting samples because X-ray images are cumbersome to use and X-rays pose non-negligible risks for healthy subjects. In addition, sharing such medical images has a much higher regulatory burden. Also, low baseline human performance should be considered. For example, humans have difficulties when dealing with very large amounts of certain types of data. In those cases, even imperfect performance from a ML algorithm can prove extremely helpful. In medical fields such as cancer therapy, ML is helping in making advances that would not have been possible without ML, as discussed above.

12. Regulatory hurdles

ML applications face, appropriately enough, steep regulatory hurdles worldwide. In the United States, relatively few ML applications have been approved for clinical use by the FDA. The FDA classifies devices that carry out clinical interpretation as Class III devices and, fittingly, requires them to meet extremely high standards. Since it is very difficult and time-consuming to meet these standards, this alone is enough to ensure that ML tools will not be replacing the neurologist anytime soon, if ever. Applications based on ML are much more likely to appear in the clinic as Class II devices, or devices that help take a measurement. It is much easier and faster for developers to develop, test, and validate Class II devices, and to get them approved by the FDA.

13. Can ML ever replace the neurologist?

ML is unlikely to replace the highly trained human professional in safety-critical fields such as medicine, piloting passenger vehicles, aircraft, etc., for the foreseeable future. Even if it were ever technically possible for a smart machine to match or exceed a skilled human professional, the market is likely to be exceedingly resistant to such a change. This is because such professionals provide more than just highly skilled services: they provide a sense of security and beneficence that we simply cannot get from machines. Many of us would not entrust our lives to machines flying an airplane by themselves or robots autonomously performing a heart transplant surgery. Human preference, after all, is the main reason why even some of the mundane predictions by futurists like Arthur Clarke and Ray Kurzweil about life in 2010 [62-65] have not come true.

ACKNOWLEDGEMENTS

We thank Ms. Jennevieve Sevilla and Ms. Fallon Branch for their help in preparing the manuscript. We gratefully acknowledge past and ongoing grant support from the U.S. Army (grants # W911NF-11-1-0105 and W911NF-15-1-0311), National Institutes of Health (grant # R21 NS086356), and National Science Foundation (grant # IOS-1147097).

CONFLICT OF INTEREST STATEMENT

The authors have no conflicts of interest, nor any proprietary or commercial interests, whatsoever to declare.

REFERENCES

1. Ferrucci, D., Levas, A., Bagchi, S., Gondek, D. and Mueller, E. T. 2013, Artificial Intelligence, 199, 93.
2. Summa, M. G. 2012, Statistical learning and data science. Chapman & Hall/CRC computer science & data analysis series. Boca Raton, CRC Press.
3. Bishop, C. M. 2006, Pattern recognition and machine learning. Information science and statistics. New York, Springer.
4. Alpaydin, E. 2010, Introduction to machine learning, Adaptive computation and machine learning (Ed. 2), Cambridge, Mass., MIT Press.
5. Cleophas, T. J. Z. and Aeilko H. 2013, Machine learning in medicine. New York, Springer.
6. Hastie, T. and Robert-Friedman, J. H. 2009, The elements of statistical learning: data mining, inference, and prediction. Springer series in statistics (Ed. 2), New York, NY, Springer.
7. Theodoridis, S. and Konstantinos. 2006, Pattern recognition (Ed. 3), San Diego, CA, Academic Press.
8. Goodfellow, I. and Yoshua Courville, A. 2016, Deep learning. Adaptive computation and machine learning. Cambridge, Massachusetts, The MIT Press.
9. Govindaraju, V. and Rao, C. R. 2013, Machine learning : theory and applications. Handbook of statistics, Amsterdam, Boston, Elsevier/North Holland.
10. Gibney, E. 2016, Nature, 529, 445.
11. LeCun, Y., Bengio, Y. and Hinton, G. 2015, Nature, 521, 436.
12. Upbin, B. 2013, IBM’s Watson Gets Its First Piece Of Business In Healthcare, iForbes.
13. Bakkar, N., Kovalik, T., Lorenzini, I., Spangler, S., Lacoste, A., Sponaugle, K., Ferrante, P., Argentinis, E., Sattler, R. and Bowser, R. 2017, Acta. Neuropathol.,135, 227.
14. Ahmed, M. N., Toor, A. S., O'Neil, K. and Friedland, D. 2017, IEEE Pulse, 8, 4.
15. Hoyt, R. E., Snider, D., Thompson, C. and Mantravadi, S. 2016, JMIR. Public Health Surveill., 2, e157.
16. Chen, Y., Elenee Argentinis, J. D. and Weber, G. 2016, Clin. Ther., 38, 688.
17. Shader, R. I. 2016, Clin. Ther., 38, 1.
18. Wernick, M. N., Yang, Y., Brankov, J. G., Yourganov, G. and Strother, S. C. 2010, IEEE Signal Processing Magazine, 27, 25.
19. Wang, S. and Summers, R. M. 2012, Med. Image Anal., 16, 933.
20. Djuric, U., Zadeh, G., Aldape, K. and Diamandis, P. 2017, NPJ Precision Oncology, 1, 22.
21. Bouton, C. E., Shaikhouni, A., Annetta, N. V., Bockbrader, M. A., Friedenberg, D. A., Nielson, D. M., Sharma, G., Sederberg, P. B., Glenn, B. C., Mysiw, W. J., Morgan, A. G., Deogaonkar, M. and Rezai, A. R. 2016, Nature, 533, 247.
22. Schwartz, A. B. 2004, Annu. Rev. Neurosci., 27, 487.
23. Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S. and Schwartz, A. B. 2008, Nature, 453, 1098.
24. Ajiboye, A. B., Willett, F. R., Young, D. R., Memberg, W. D., Murphy, B. A., Miller, J. P., Walter, B. L., Sweet, J. A., Hoyen, H. A., Keith, M. W., Peckham, P. H., Simeral, J. D., Donoghue, J. P., Hochberg, L. R. and Kirsch, R. F. 2017, Lancet, 389, 1821.
25. Lubin, J. A., Strebe, J. K. and Kuo, J. S. 2017, Neurosurgery, 80, N29.
26. Flesher, S. N., Collinger, J. L., Foldes, S. T., Weiss, J. M., Downey, J. E., Tyler-Kabara, E. C., Bensmaia, S. J., Schwartz, A. B., Boninger, M. L. and Gaunt, R. A. 2016, Sci. Transl. Med., 8, 361ra141.
27. Mahmoudi, A., Takerkart, S., Regragui, F., Boussaoud, D. and Brovelli, A. 2012, Comput. Math Methods Med., 2012, 961257.
28. Schrouff, J., Rosa, M. J., Rondina, J. M., Marquand, A. F., Chu, C., Ashburner, J., Phillips, C., Richiardi, J. and Mourão-Miranda, J. 2013, Neuroinformatics, 11, 319.
29. Collij, L. E., Heeman, F., Kuijer, J. P., Ossenkoppele, R., Benedictus, M. R., Möller, C., Verfaillie, S. C., Sanz-Arigita, E. J., van Berckel, B. N., van der Flier, W. M., Scheltens, P., Barkhof, F. and Wink, A. M. 2016, Radiology, 281, 865.
30. Lubrini, G., MartínMontes, A., Diez-Ascaso, O. and Diez-Tejedor, E. 2018, Neurologia, 33, 187.
31. Gillebert, C. R. and Mantini, D. 2013, Neuroscientist, 19, 509.

32. Xiao, H., Yang, Y., Xi, J. H. and Chen, Z. Q. 2015, Neural Regen. Res., 10, 2062.
33. Cerasa, A., Novellino, F. and Quattrone, A. 2016, Curr. Neurol. Neurosci. Rep., 16, 91.
34. Salvatore, C., Battista, P. and Castiglioni, I. 2016, Curr. Alzheimer Res., 13, 509.
35. Lee, D. 2013, Neuron, 78, 233.
36. Harmon, K. 2012, Scientific American.
37. Teipel, S., Babiloni, C., Hoey, J., Kaye, J., Kirste, T. and Burmeister, O. K. 2016, Alzheimers Dement., 12, 695.
38. Casadio, M., Ranganathan, R. and Mussa-Ivaldi, F. A. 2012, J. Mot. Behav., 44, 419.
39. Cooper, R. A., Dicianno, B. E., Brewer, B., LoPresti, E., Ding, D., Simpson, R., Grindle, G. and Wang, H. 2008, Med. Eng. Phys., 30, 1387.
40. Horowitz, B. T. 2010, Scientific American.
41. Volpe, B. T., Krebs, H. I. and Hogan, N. 2001, Curr. Opin. Neurol., 14, 745.
42. Pahwa, R., Lyons, K. E., Wilkinson, S. B., Simpson, R. K. Jr., Ondo, W. G., Tarsy, D., Norregaard, T., Hubble, J. P., Smith, D. A., Hauser, R. A. and Jankovic, J. 2006, J. Neurosurg., 104, 506.
43. Tramontana, M. G., Molinari, A. L., Konrad, P. E., Davis, T. L., Wylie, S. A., Neimat, J. S., May, A. T., Phibbs, F. T., Hedera, P., Gill, C. E., Salomon, R. M., Wang, L., Song, Y. and Charles, D. 2015, J. Parkinsons Dis., 5, 151.
44. Weaver, F. M., Follett, K., Stern, M., Hur, K., Harris, C., Marks, W. J. Jr., Rothlind, J., Sagher, O., Reda, D., Moy, C. S., Pahwa, R., Burchiel, K., Hogarth, P., Lai, E. C., Duda, J. E., Holloway, K., Samii, A., Horn, S., Bronstein, J., Stoner, G., Heemskerk, J. and Huang, G. D. 2009, JAMA., 301, 63.
45. Vidailhet, M., Vercueil, L., Houeto, J. L., Krystkowiak, P., Lagrange, C., Yelnik, J., Bardinet, E., Benabid, A. L., Navarro, S., Dormont, D., Grand, S., Blond, S., Ardouin, C., Pillon, B., Dujardin, K., Hahn-Barma, V., Agid, Y., Destee, A. and Pollak, P. 2007, Lancet. Neurol., 6, 223.
46. Blankertz, B., Acqualagna, L., Dahne, S., Haufe, S., Schultze-Kraft, M., Sturm, I., Uscumlic, M., Wenzel, M. A., Curio, G. and Muller, K. R. 2016, Front Neurosci., 10, 530.
47. Hebb, A. O., Zhang, J. J., Mahoor, M. H., Tsiokos, C., Matlack, C., Chizeck, H. J. and Pouratian, N. 2014, Neurosurg. Clin. N. Am., 25, 187.
48. Dowling, J. 2009, Eye (Lond.), 23, 1999.
49. Luo, Y. H. and da Cruz, L. 2014, Br. Med. Bull., 109, 31.
50. Ng, D. C., Bai, S., Yang, J., Tran, N. and Skafidas, E. 2009, J. Neural Eng., 6, 065004.
51. Picaud, S. and Sahel, J. A. 2014, C. R. Biol., 337, 214.
52. Weiland, J. D., Cho, A. K. and Humayun, M. S. 2011, Ophthalmology, 118, 2227.
53. Chen, K., Lo, Y. K., Yang, Z., Weiland, J. D., Humayun, M. S. and Liu, W. 2013, IEEE Trans. Biomed. Circuits Syst., 7, 326.
54. Fernandez, E., Greger, B., House, P. A., Aranda, I., Botella, C., Albisua, J., Soto-Sanchez, C., Alfaro, A. and Normann, R. A. 2014, Front Neuroeng., 7, 24.
55. Kendrick, M. 2009, Sci. Am., 301, 22.
56. Wolpert, D. H. and Macready, W. G. 1997, IEEE Transactions on Evolutionary Computation, 1, 67.
57. Wolpert, D. 1996, Neural Computation, 8, 1341.
58. Rolls, E. T. 2012, Front Comput. Neurosci., 6, 35.
59. VanRullen, R. 2017, Front Psychol., 8, 142.
60. Kumaran, D., Hassabis, D. and McClelland, J. L. 2016, Trends in Cognitive Sciences, 20, 512.
61. Huys, Q. J., Maia, T. V. and Frank, M. J. 2016, Nat. Neurosci., 19, 404.
62. Clarke, A. C. 1984, 2010, Odyssey Two, Del Rey Science fiction (Ed. 1), New York, Ballantine Books.
63. Clarke, A. C. and Hyams, P. 1985, The Odyssey file (Ed. 1), New York, Ballantine Books.
64. Kurzweil, R. 1990, The age of intelligent machines. MIT Press Cambridge, Mass.
65. Kurzweil, R. 2005, The singularity is near : when humans transcend biology. New York, Viking.
66. Wicab, Inc. 2017.