www.sciencemag.org/content/358/6370/1530/suppl/DC1

Supplementary Materials for

What can machine learning do? Workforce implications

Erik Brynjolfsson* and Tom Mitchell

*Corresponding author. Email: erikb@ mit.edu

Published 22 December 2017, *Science* **358**, 1530 (2017)
DOI: 10.1126/science.aap8062

This PDF file includes:

This Appendix includes two components, 1) a short taxonomy of machine learning technology and 2) a 21 question rubric for assessing the suitability of tasks for machine learning, particularly supervised learning.

1. Machine Learning Technology

In this supplemental section we provide additional detail about the variety of machine learning algorithms referred to in the main article.

To understand the types of problems that are suitable for machine learning (SML), it is useful to understand the types of learning problems that cover the vast majority of commercial applications of machine learning. Algorithms for machine learning fall into three broad categories:

1. Supervised learning

This is the most heavily explored type of machine learning and the most common type found in commercial applications, such as those described in Table 1. The supervised learning task is to learn some general function f(x)=y from a set of training examples of input-output (x,y) pairs of the function. For example, we might want to learn a function that takes as input an email message *x* and outputs a binary label *y* to classify whether the email is spam or not, training it by providing a large collection of emails with their correct labels. Thousands of important problems fit this template including classifying which emails are spam and diagnosing which medical patients have which disease. Important prediction problems also fit into this category (Agrawal, Gans and Goldfarb, 2017), for instance, predicting which customers will purchase which products, predicting tomorrow's stock price from today's, or predicting which credit card transactions are likely to be fraudulent. Theoretical results have been developed that give rough guidance on how the accuracy of the learned function improves with increasing numbers of training examples, and how it decreases with the complexity of the function to be learned. (Jordan and Mitchell, 2015)

Table 1 : Examples of Supervised Learning tasks

Input X	**Output Y**	**Application**
Voice recording	Text transcript	Speech recognition
Historical market data	Future market data	Trading bots
Photograph	Caption	Image tagging
Drug chemical properties	Treatment efficacy	Pharma R&D
Store transaction details	Is the transaction fraudulent?	Fraud detection
Recipe ingredients	Customer reviews	Food recommendations

Purchase histories	Future purchase behavior	Customer Retention
Car locations and speed	Traffic flow	Traffic lights
Faces	Names	Face recognition

2. Reinforcement learning.

Like supervised learning, reinforcement learning requires learning a specific target function, but the type of training experience available is very different, as is the learning algorithm. Reinforcement learning tasks typically apply to some agent interacting with its world via a set of actions, where the function to be learned takes the current observed world state as input, and provides the recommended next action as output (i.e., the learned function is f(state)=action). As the agent takes a sequence of actions that move it from state to state it receives a training signal consisting of numerical rewards or penalties following some of these actions. The goal is to learn to choose actions that lead over time to the greatest expected total discounted future reward. This is inherently a more difficult learning problem than supervised learning because the training signal is more indirect and results from multiple choices by the agent. Nevertheless, it is important because it covers significant problems such as learning how to maximize customer satisfaction over multiple interactions, learning how to structure a website to maximize profit resulting from multiple visits, and learning how to minimize energy consumption in a data center. One notable recent use of reinforcement learning was in AlphaGo and AlphaGo Zero (Silver et al., 2016; 2017), the AI systems that defeated the top human Go players. AlphaGo Zero trained itself using reinforcement learning and millions of games played against itself, where the training reward signal occurred only at the end of the game to indicate a win or loss.

3. Unsupervised learning.

Unsupervised learning problems require no training labels, and involve approaches such as clustering the data, or finding its principle components to help the user explore the data without a specific, pre-specified target function to be learned. For example, it can be used to cluster customer records into groups of customers with similar demographics and/or behavior, to help understand which distinct types of customers appear in the data, or to discover different types of objects in images or videos.[1] While this type of learning is used frequently to explore data interactively, it has led to fewer routine commercial applications.

2. Rubric for to determine Suitability for Machine Learning (SML)

Below is the list of task properties we propose, for determining whether a task is suitable for machine learning (SML). Note that although task properties are described here as Boolean properties (e.g., "Task feedback is immediate."), we score them on a scale of 1 through 5, and calculate the overall SML score for a task by combining these scores.

1. Information needed to complete the task (inputs) and outputs can be explicitly specified in machine-readable format
1: It is very difficult or impossible to identify particular inputs and outputs (e.g. emotions, ideas, impressions)
3: It is possible to create rankings or partial representability of inputs and outputs
5: It is easy to quantify results on a machine/computer (e.g. calculations, concrete inputs and outputs)

2. Task information is recorded or recordable by computer
1: Input and training outputs in separate, incompatible databases, or not online at all (e.g. no dataset exists at all)
3: Some inputs and training outputs online/can be observed by computer (e.g. sensors can be installed at low cost to generate necessary data - includes places where images, video, or sound data can be recorded but may not be yet).
5: Inputs and training outputs already online in consistent format (e.g. queries can be made to a database to get necessary information)

3. Task does not require a wide range of complex outputs (mental and/or physical)
*1: Task requires a range of different mental **or** physical abilities (e.g. performing a difficult surgery, completing an obstacle course, acting in a film, creating a plan, negotiating)*
3: Tasks requires moderate variety in responding and acting (e.g. walking a dog, supervising a construction site, inspecting a building)
5: Task requires very little variety in thinking or action/movement (e.g. sorting mail, doing taxes, labeling an image)

4. Task feedback (on the success of outputs) is immediate or available through plentiful historical data.
1: Feedback is never received or takes a long very long time (e.g. predictions may take months or years until accuracy can be proven, it takes years often, to determine the success of writing a business plan)
3: Feedback is received but response time is inconsistent/unclear and/or unclear on if it is beneficial to progress (e.g. teaching quality can be measured through standardized tests and observations, which are neither immediate nor long-term)
5: Instant results/feedback when the action is executed (e.g. calculations can be immediately categorized as either correct or incorrect)

5. The task output is error tolerant
1: A mistake could lead to serious harm, injury, or death to those involved, or could lead to lasting negative consequences (e.g. mistake during surgery, mistake at a nuclear facility)
3: A mistake will have negative consequences, but can fixed with some work (e.g. a construction mistake, or human resources slip-up will be noticed and reprimanded, but would not result in termination of employment or injury)
5: A mistake can be easily fixed, and holds few, if any negative consequences (e.g. a slip up in factory work or mail sorting mistakes could go potentially unnoticed)

6. It is not important that outputs are perceived to come from a human
1: Task fundamentally requires human connection (e.g. therapy, teaching, making a speech, delivering bad news like a diagnosis)
3: Task could be done by a non-human, but might cause frustration or inefficiency (e.g. customer service)
5: Task requires little human connection, empathy, or emotional intelligence (e.g. telemarketing, preparing taxes, performing calculations, lifting boxes)

7. Task does not require complex, abstract reasoning
1: Task requires intuition or highly involved reasoning (e.g. writing a plan, administering therapy, coming up with a research proposal/plan, teaching)
3: Task requires some reasoning, but can mostly be broken into well-defined rules (e.g. playing chess, sorting mail)
5: Task is mainly perception and does not require complex reasoning skills (i.e. can be done in less than 1 second)

8. Task is principally concerned with matching information to concepts, labels, predictions, or actions
1: The task does not have clear, consistent inputs or outputs (e.g. open-ended travel vacation)
3: The task potentially has well-defined inputs and outputs, but does not require mapping of the two (e.g. middle management, camp counselor, kindergarten teacher)
5: The task has clear inputs and outputs. The purpose of the task is to determine how the inputs affect the outputs (e.g. translating one language to another, matching an image to words describing the image)

9. Task does not require detailed, wide-ranging conversational interaction with a customer or other person
1: Task requires explaining something deeply to another person, or having a deep conversation that cannot be predicted in advance (e.g. therapy, negotiation)
3: Task involves communicating but about a relatively small, pre-set range of topics (e.g. giving instructions or directions, answering/asking specific/common questions)
5: Task doesn't require any form of communication/conversation with another person (e.g. solving equations, lifting objects, observing)

10. Task is highly routine and repeated frequently
1: Task is not routine, and different approaches must be taken every time (e.g. negotiating contracts, fighting a fire, treating rare and specific issues)
3: Task is routine, but might be done differently each time (e.g. waiting tables, operating a multi-purpose machine, teaching a class)
5: Task is very repetitive, and must be done in the same way each time (e.g. working in an assembly line, delivering things along a route, being a cashier)

11. Task is describable with rules
1: The task has no clear, well-known set of rules on what is and is not effective (e.g. inventing a product, dealing with exceptions, painting a painting, writing a book)
3: The task may be subject to some general rules (e.g. cutting hair, making a chair, overseeing a construction site, driving a car)
5: The task can be fully described by a detailed set of rules (e.g. folding origami, following a recipe, determining loan eligibility via formula)

12. There is no need to explain decisions during task execution
1: Decisions highly impact the lives of others and require justification (e.g. persuasion, long term planning, law-making, courtroom decisions)
3: There is some need to explain decisions, particularly when people ask questions (e.g. doctors performing checkups, operating machinery)
5: There is no need to explain decisions. The task is only concerned with having the correct output, and does not depend on the process through which the output is determined. (e.g. correctly predicting the weather, performing the correct calculation, optimizing allocation of resources, determining quickest route)

13. Task can be converted to answering multiple choice questions, ranking alternative options, predicting a number, or grouping similar objects
1: Task output does not have to do with any of these options (e.g. lifting objects, collecting things, making things)
3: Task output may have something to do with categorizing or identifying things but the rules or criteria are not clear in advance and must be discovered (e.g. Supervising others, making sure that a team stays on track, recommending a plan)
5: Task is focused on applying a rule or pattern, particularly related to sorting or grouping things (e.g. grading food quality, recognizing faces, diagnosing common conditions, sorting mail)

14. Long term planning is not required to successfully complete the task
1: The task is concerned with planning around a timeline of months or years (e.g. supervising research projects, constructing complex buildings, entrepreneurship, crafting long term cancer treatment plans)
3: The task is concerned with a timeline in the range of weeks or days or an indeterminate amount of time (e.g. managing others' workloads, teaching students a specific set of lessons)
5: The task involves an immediate response and isn't concerned with a future impact (e.g. answer questions in a call center, lifting objects, performing calculations)

15. The task requires working with text data or might require working with text in the future:
1: Task does not include working with any text (e.g. making a hamburger, operating machinery)
3: Task may include some light reading (e.g. reading labels, occasionally reading directions)
5: Task includes heavy text processing, reading, or writing (e.g. reading documents, writing a letter)

16. The task requires working with image/video data or might require working with image/video data in the future:

1: Task does not require looking at images or videos, or otherwise using your eyes (e.g. having a phone conversation)
3: Task may occasionally require utilizing images and video (e.g. greeting customers, knowing whether you're alone before entering a password)
5: Task requires analyzing images and videos (e.g. finding defects in products, looking at surveillance footage, classifying objects in pictures, facial recognition)

17. The task requires working with speech data or might require working with speech data in the future:
1: Task does not require listening to or communicating with speech (e.g. independent tasks such as lifting objects, repetitive assembly work)
3: Task may require occasional listening, talking, or communicating (e.g. construction work, being a cashier, financial analyst)
5: Task requires heavy speech processing, or communicating with speech (e.g. telemarketing, translating between languages, giving a speech, having a conversation)

18. The task requires working with other types of data (other than text, image/video, and speech):
1: Task does not require working with digital data in any form (e.g. making handmade art)
3: Task requires working with some types of data at a low frequency (e.g. performing restocking tasks at a grocery store, testing machines for maintenance)
5: Task requires constant interaction with machine or sensor data. (e.g. monitoring temperature/weather, analyzing pricing data, monitoring machine performance)

19. Many components of the task can be completed in a second or less
1: Task takes a long time to complete (e.g. making a plan, writing a book)
3: Task cannot be done instantly, but also does not involve much long-term planning (e.g. performing a surgery, delivering food)
5: Task can be done instantly, or can be broken up into smaller choices that can be done instantly (e.g. Identifying a picture)

20. Each instance, completion, or execution of the task is similar to the other instances in how it is done and these actions can be measured
1: Task primarily involves rare or unique situations that cannot be summarized easily with machine-readable data (e.g. making strategic decisions for a company)
3: Data can be collected but the data output structure is highly varied (e.g. performing different types of surgery will generate different kinds of feedback)
5: Data are already available or can be easily collected (e.g. customer service transcripts, text translation, image classification, stock price movements)

21. Actions in the task must be completed in a very specific order, and practicing the task to get better is easy
1: Task involves many rare-occurring or unique situations that make the task hard to practice (e.g. therapy, negotiating)
3: Task involves some components for which a reward function can be defined (e.g. shipping/receiving in a warehouse, designing mechanical components)
5: Sequences can be repeated and tested over and over again, and there are "right" moves that can be used to generate rewards (e.g. video games, industrial process optimization, investment decision-making)

References

A. K. Agrawal, J. S. Gans, A. Goldfarb, What to Expect From Artificial Intelligence. *MIT Sloan Manag. Rev.* **58**, 23 (2017).

M. I. Jordan, T. M. Mitchell, Machine learning: Trends, perspectives, and prospects. *Science* **349**, 255–260 (2015). doi:10.1126/science.aaa8415 Medline

D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, D. Hassabis, Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016). doi:10.1038/nature16961 Medline

D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel, D. Hassabis, Mastering the game of Go without human knowledge. *Nature* **550**, 354–359 (2017). doi:10.1038/nature24270 Medline