

A Privacy Officer's Guide to Artificial Intelligence and Machine Learning

Introduction

Advanced algorithms, machine learning, and Artificial Intelligence are appearing across digital and technology sectors from healthcare and financial institutions, to on-line advertising, education profiling, and social media platforms. Understanding these new AI and related processes presents new challenges for privacy officers and others responsible for data governance in companies from retailers to cloud service providers. Even in advance of concrete legal or regulatory compliance concerns, AI raises ethical and practical challenges as companies strive to maximize its potential benefits while preventing new or foreseeable harms.

While technologies change and evolve, privacy concerns remain similarly focused on personal awareness and empowerment, as well as concerns about exploitation or discrimination of individual consumers or users. It is critical that privacy concerns are addressed now – early in the cycle of applying AI and ML programs into new and existing industries, platforms, and applications – to ensure that individuals are treated with respect and dignity, and retain the discretion and granular controls necessary over their own information.

This guide is designed to describe the privacy challenges associated with the implementation of AI to new and existing products and services by commercial organizations, and to outline recommendations for privacy officials across industries and contexts.

Most applications of Machine Learning and AI require exceptionally large data sets, sometimes called “training data”, in order to power the “learning” potential of these systems. Big Data has been recognized for years for its potential to be searchable for patterns and correlations previously unknowable. However, it is only using recent increases in computing power and algorithmic applications that we can improve on traditional analytics by creating systems that can process information in unimaginably large quantities, and adjust their analysis over time with limited human involvement.

Definitions

Artificial Intelligence:

While there is no one-size-fits-all definition, AI is generally understood to be a field of computer science focused on designing systems using algorithmic techniques, somewhat inspired by knowledge of the human brain, capable of performing tasks that, if performed by a human, would be said to require intelligence. Systems that solve problems commonly associated with human intelligence, such as cognitive learning, problem solving, and pattern recognition. For a system to be said to exhibit artificial intelligence, it should be able to perceive, adapt, and implement changes to its own processes, and then operate using those processes for further data analysis.

Machine Learning:

A collection of algorithms that can learn from and make predictions based on recorded data, identify patterns and functions with defined levels of uncertainty, extract statistical connections and structures from data, and classify data based on outcomes.

Deep Learning:

A subset of Machine Learning. Systems set up to solve real world problems using neural networks to mimic human decision-making. The machine achieves this by training itself on large amounts of data. This process involves layering algorithms and simulating the complexity of connections among inputs and outputs across layers to simulate the human neural system. The algorithms are not limited to explainable relationships but rely on non-linear relationships to interact based on a series of factors from the data provided.

DEEP LEARNING – EXAMPLES

From IBM DeveloperWorks: A Beginners Guide

“Deep learning is a relatively new set of methods that's changing machine learning in fundamental ways. Deep learning isn't an algorithm, per se, but rather a family of algorithms that implement deep networks with unsupervised learning. These networks are so deep that new methods of computation, such as GPUs, are required to build them (in addition to clusters of compute nodes).

“Deep learning algorithms have also been applied to facial recognition, identifying tuberculosis with 96 percent accuracy, self-driving vehicles, and many other complex problems.

“However, despite the results of applying deep learning algorithms, problems exist that we have yet to solve. A recent application of deep learning to skin cancer detection found that the algorithm was more accurate than a board-certified dermatologist. But, where dermatologists could enumerate the factors that led to their diagnosis, there's no way to identify which factors a deep learning program used in its classification. This is called deep learning's black box problem.

“Another application, called Deep Patient, was able to successfully predict disease given a patient's medical records. The application proved to be considerably better at forecasting disease than physicians—even for schizophrenia, which is notoriously difficult to predict. So, even though the models work well, no one can reach into the massive neural networks to identify why.”

Neural Networks:

Networks of overlapping algorithms designed to be similar to the human nervous system. These use the stages of learning in ML to give AI-based programs the ability to solve complex problems by breaking them down into levels of data. A neural network generally consists of three categories; an input layer, one or more middle or hidden layers, and an output layer. There are many combinations for how data travels between the layers. Some systems flow through directly

from input to output; others form loops or send the data back and forth within or between layers in the network before the output is produced. The first level of the network may only evaluate a limited number of pixels in an image, for example, and then once the initial analysis is done, the network will pass those findings to the next level, which will evaluate more pixels, create or evaluate metadata, and then pass it forward again.

Backpropagation: “the true power of neural networks is their multilayer variant. Training single-layers is straightforward, but the resulting network is not very powerful. The question became, How can we train networks that have multiple layers? This is where backpropagation came in. Backpropagation is an algorithm for training neural networks that have many layers. It works in two phases. The first phase is the propagation of inputs through a neural network to the final layer (called feedforward). In the second phase, the algorithm computes an error, and then backpropagates this error (adjusting the weights) from the final layer to the first.”¹

Supervised Learning

Supervised learning programs start with a data set containing training examples that have already been assigned correct labels by a programmer/[human]. For example, when learning to classify handwriting, a supervised learning algorithm takes thousands of pictures of handwritten symbols with associated labels containing the correct letter or number each image represents. The algorithm will then learn the relationship between the images and the variations that may exist for it still to represent that figure, and apply that learned relationship to classify completely new, unlabelled images.

Clustering: the algorithm organizes a set of feature vectors into clusters based on one or more attributes of the data. “One of the simplest algorithms that you can implement in a small amount of code is called k-means. In this algorithm, k indicates the number of clusters in which you can assign samples. You can initialize a cluster with a random feature vector, and then add all other samples to their closest cluster (given that each sample represents a feature vector and a Euclidean distance used to identify “distance”). As you add samples to a cluster, its centroid—that is, the center of the cluster—is recalculated. The algorithm then checks the samples again to ensure that they exist in the closest cluster and ends when no samples change cluster membership. Although k-means is relatively efficient, you must specify k in advance. Depending on the data, other approaches might be more efficient, such as hierarchical or distribution-based clustering.”²

Decision Trees: Closely related to clustering is the decision tree. A decision tree is a predictive model about observations that lead to some conclusion. Conclusions are represented as leaves on the tree, while nodes are decision points where an observation diverges. Decision trees are built from decision tree learning algorithms, where the data set is split into subsets based on attribute value tests (through a process called recursive partitioning).³

¹ <https://www.ibm.com/developerworks/library/cc-beginner-guide-machine-learning-ai-cognitive/index.html>

² <https://www.ibm.com/developerworks/library/cc-beginner-guide-machine-learning-ai-cognitive/index.html>

³ Id.

Unsupervised Learning

Unlike supervised learning, programs using unsupervised learning use data sets that have no existing labels for the data being analysed. The system uses alternative methods to classify and understand the data, and learn from it to make future predictions on new data. One common process is called “clustering” where the data is sorted into groups by similarity. The goal of clustering is to create groups of data points such that points in different clusters are dissimilar while points within a cluster are similar.

Reinforcement Learning

In the absence of existing training data, the agent learns from experience. It collects the training examples (“this action was good, that action was bad”) through trial-and-error as it attempts its task, with the goal of maximizing long-term reward. This type of learning concentrates on how an AI 'agent' should behave in order to get the most out of its work. The machine picks an action or a sequence of actions, and gets a reward. “This is used when teaching machines to play and win games but needs a large number of trials to learn even simple tasks.”⁴

Discussion

Almost all current applications of AI have been in “Specific” or “Narrow” AI – that is, learning processes that perform well on specialized tasks. “General” AI, systems that can perform well across a variety of cognitive domains, remains mostly notional at this time. Current use cases of Specific AI include image recognition, natural language (speech) processing, real-time language translation, predictive analytics of personal behaviour, and individual profiling.

AI systems may be implemented in various contexts. In some cases, the system is built by the developer to achieve a certain level of accuracy, then put into operation, or provided to the industrial user, with no further learning expected. This type of model is static or fixed. Results should still be continually evaluated using test data at periodic intervals.

In other online or dynamic applications, continuous learning is intended, using the real world data generated during actual operation. For these programs to learn, they need experience from the information obtained under real world conditions. Ideally, this system will continue to evolve and enhance the accuracy of their analysis and function. For example, a person’s spending pattern may change based on moving to a new home, getting a new job, or getting married or having a child. The machine would see the progressive steps of this change in behavior and adapt to the new pattern for its predictive function.

Regardless of the methods used for machine learning, the result will be a “model”, which can then be fed with new data to produce the desired analysis or decision. This might be a label, a translation, a sentence recommendation, a degree of probability on consumer behavior, and so on.

⁴ <http://www.wired.co.uk/article/machine-learning-ai-explained>

Of particular relevance to consumer privacy applications of ML techniques: 1) the functions of learning software that make decisions regarding or impacting individual people, and 2) the aspects of data use that may – whether intentionally or unintentionally – imbed human bias or existing systemic discrimination into the program’s learning cycles.

One of the reasons for these concerns is the difficulty of having sufficient transparency in ML or AI systems because of the complexity of the algorithms and relationships. Particularly in systems utilizing neural networks, there is often not a single fixed algorithm that can be isolated, or easily explained. In fact, mathematicians can demonstrate that the more complex the system, the more accurate the results are, but the harder is it to explain how the data input results in the system’s decision or output. In simpler systems, it is easier to demonstrate or explain how certain factors impact the output. Unfortunately, the simpler the system, the less accurate or less useful the analysis provided.

Importance of the Data Sets

The design aspects of the underlying system cannot be overemphasized. Decisions made when defining categories, identifying fields, and establishing relationships impact both the efficiency of the model, and what outputs will be generated. When the outcome is identifying patterns over time, some variations or range of accuracy may be acceptable. However, when the systems are generating recommendations that impact individuals – such as with credit evaluations or sentencing guidelines – any inherent bias in the weighting of various factors or choosing which ones to use have real-world results that must be fair and defensible.

The impact of certain sensitive data such as race or gender may affect outcomes in unforeseen or undesirable ways. While the goal of machine learning programs is frequently to create more objective evaluation or analysis and fairer outcomes, if the training data is inherently biased because of its source (that is, if the human-run process from which it was derived included bias in the collection or correlation of the data included in the data set) then those biases will carry over into the machine learning system. It takes expert understanding and evaluation to prevent or discern this, including programmers who are able to use more algorithm to “audit” or evaluate the testing outcomes and ensure that such distortions do not occur.⁵ This is important regardless of the quality or bias inherent in the original training data. A system may also “learn” to be biased based on information input after it is in operation.⁶ (example – sentencing guidelines racially biased?)

The quantity and nature of the data included is critical as well. Too much data – too many unrelated factors – and the desired results risk being obscured. But having too little data per record, or more likely,

⁵ ([From Norway paper](#)) One of the major privacy challenges of AI systems is bias. For example, a research study found substantial disparities in the accuracy of three commercial face recognition systems conducting automated facial analysis. The study found that the commercial systems’ training data were overwhelmingly composed of lighter-skinned subjects. The study showed that facial recognition of darker-skinned females had error rates up to over 30 %, compared to the error rate for lighter-skinned males of maximum 0.8%.

<http://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf>

⁶ MS Tae Twitter bot example

too few records overall, risks the outcomes being non-representative of the total populations to which it will apply.

(example – too many white/young/male records creating outcomes not applicable to women or minorities; facial recognition? Others?)

There are important choices to be considered in whether to use “real” data in the training sets or simulated data, as well. The more personally identifiable information required for the particular application, the more sensitive it may be to use data from real individuals in developing the process. Express consent may be appropriate in instances where PII is being used in research or development significantly unrelated to its original purpose when collected.

(example – genetic data? Others?)

Examples of AI in practice (needed here separate from earlier examples? Some categories for further discussion included)

Image recognition

1. *Image (Facial) recognition (define/describe – test cases in commercial use)*

Natural language processing

2. *Personal assistants*

Autonomous machines

3. *Cars, Roombas (in future versions), personal robots, surveillance drones*

Profiling and prediction

4. *Financial, criminal, college acceptance, etc.⁷*

Privacy challenges

The privacy implications of AI and ML systems are challenging under the traditional guidelines such as the Fair Information Practice Principles. Data protection processes such as using data in its encrypted state when possible, or applying de-identification techniques, may prove insufficient or impossible in these new use cases.

Depending on the specific legal regime, the Fairness principle will be affected by the restrictions on what personal data may be considered when making certain types of decisions that affect individuals. Under the GDPR, for example, the model “must not emphasise information relating to racial or ethnic origin,

⁷ https://www.cdcr.ca.gov/rehabilitation/docs/FS_COMPAS_Final_4-15-09.pdf
<https://doc.wi.gov/Pages/AboutDOC/COMPAS.aspx>

political opinion, religion or belief, trade union membership, genetic status, health status or sexual orientation if this would lead to arbitrary discriminatory treatment.”⁸

Because of the iterative and evolving nature of AI systems, all privacy considerations must be continually assessed and re-evaluated throughout the life of the application as actual requirements may change.

Data Minimization

Data minimization must be evaluated differently in light of the extremely large quantities of data required to “teach” these systems. In order to achieve the potential of AI systems, they are designed to handle the large amounts of data required to identify patterns and connections and provide precise outputs and analysis. This is what distinguishes AI from traditional computational systems, and allows for the promise of great advancements, in contexts ranging from medical research to personalized services.

Nevertheless, the Data Minimization principle requires that the data collected is sufficient and limited to what is necessary for the particular purpose. Even in machine learning contexts, including unnecessary data remains a risk, both to the accuracy and efficiency of the system, and to the individuals from whom the data is derived. The opportunities presented by these applications must be balanced against the appropriate protections for individuals’ personal data.

Purpose Limitation; Notice and Consent

Another unique challenge of AI and ML contexts clearly identifying the purpose and use limitations at the original collection of the data, because it is not always possible to predict in advance what the algorithm will “learn” and what applications may become clear during development. This is true even in the context of applications of “narrow” AI, and almost inevitable in the more general use applications being studied and designed.

In addition to selecting the relevant and necessary data, the developer must consider how to achieve scope the notice and consent to frame the overall objective comprehensively enough to be include the variety of ways in which the data may be used, but still provide enough specifics for consent to be appropriately informed. Such considerations and decisions should be carefully documented.

(example/s)

There will almost certainly be applications developed subsequently or adjacent to the primary objective which will require that previously collected data be used. “Function creep,” the tendency of an organization to find new uses for collected data that are unrelated to the purpose for which is was collected, is a problem already in commercial data use, and will only become more tempting with more powerful programs. And while it is likely that many beneficial uses promising economic incentives or personalized products will be identified, If the new uses are significantly different from the original service or use to which the individual consented, new consent will be required.

(example/s; education data?)

⁸ Norway paper

Access and Control

Perhaps most challenging, even as systems grow in the amount of data required, and the applications for which it is used, is the requirement that individuals be aware of who holds their data, and what processes are in place for them to verify its accuracy, request corrections, or be able to have it deleted.

(example/s; and GDPR implications?)

AI systems are even more likely than past systems create new data and metadata about individuals as a result of the use of the system. For dynamic models in particular, new profiles, analyses, and conclusions are constantly being generated, and the connection to particular individuals may be hard to manage and trace. This leads to the next principle and the challenges of transparency.

Transparency

As discussed above, depending on the type of machine learning employed, understanding how a particular decision or outcome was reached might not be possible. This is sometimes referred to as “the black box” problem with algorithms.

A core privacy principle is the right to know what information about oneself is collected, created or retained, and how that information is used. Privacy is fundamentally about an individual’s right to control personal information, which can never occur unless the individual knows when PII is being processed.

In some legal jurisdictions, the inability to explain how a decision was reached could prevent the use of this system. In others, there is a requirement that regardless of the role of the AI system, there must also be a human in the chain who evaluate the individual impact of each decision, and is the final arbiter.⁹ Even where there are not definitive legal requirements, commercial applications might be well-served by considering whether such oversight is an appropriate protection that should be implemented.

Machine learning systems make it fundamentally difficult to understand how a particular outcome was reached. Even if the general process can be outlined, the role of different features, what weights are assigned to particular aspects, how those weights might change based on other inputs over time, and other complex adjustments inherent in the program may be beyond even the designer’s ability to clearly tease out and identify.

In many cases involving more inscrutable forms of machine learning, such as neural networks and deep learning, data use and evaluation changes throughout the system – the presence of other data or the same data in a new context might result in different answers to the same question at different times. This makes it more difficult to determine how a particular result was achieved, and in particular makes it challenging to recreate the parameters of a past analysis when the algorithm has adapted and changed to new inputs in the time since.

(example/s)

Sensitive Information

⁹ See for example GDPR Article 22 on Automated decision-making

A risk particular to machine learning systems is the potential identification of sensitive information being derived. Even if individual data inputs were not sensitive in themselves, the ability of AI's analytical power to combine small bits of data from comprehensive sources means it can identify patterns that predict details or draw conclusions about an individual seemingly unrelated to the original information provided. This might result in determinations about health, religious or political viewpoints, sexual orientation, or other aspects of a person that merit special protection and handling.

(example from Norwegian paper: "For example, one study combined "likes" on Facebook with information from a simple survey and predicted male users' sexual orientation with an accuracy of 88 percent. Moreover, they predicted ethnicity with 95 per cent accuracy, and whether a user was Christian or Muslim with 82 per cent accuracy.")

Risk of re-identification

Similarly, there are other risks in AI systems based on the sheer volume of data included. Traditional de-identification techniques for personal data may prove insufficient, as compilation from several sources magnifies the risk that individuals will be identifiable from data sets which were intended to be anonymous. This renders these protections less effective, or even useless, as a method to prevent privacy risks associated with individual profiling. This is a particular risk when data is collected for AI systems from multiple sources, with no clear oversight of what different categories are thus being connected. In addition, creating a profile that retains some level of obscurity becomes moot if other data sets exist with sufficient overlap to identify particular individuals without significant effort.

Security

AI holds the promise of potentially providing great assistance in assessing and meeting security threats – from predicting and thwarting hackers attacking IT systems, to evaluating physical security threats at particular facilities, or across borders.

However, AI also allows for the development of tools to automate threats and inform physical attacks in ways that are more complex, more cost-effective, and likely more successful than when human intervention was the standard. Criminals use machine learning to more broadly distribute and directly target spam. On-line bots can more successfully solicit personal data, or install malware or other computer viruses. And new threats are emerging from AI beyond the traditional boundary management of on-line or physical security programs. The ability to influence people's behaviour by targeted distribution of false information is a growing concern.

Also, like any data-based system, AI systems are vulnerable to security breaches. Exploiting them might involve seeking out the large amounts of personal data involved, or might mean influencing the factors or weighting to force particular outcomes, whether targeted to individual victims, or to benefit specific attackers. There have been cases of "reverse-engineering" where outside parties replicate the machine

learning algorithm or model based on outputs or queries from the original system. This is a threat to proprietary information, among other possible damages, and almost certainly puts user data at risk.¹⁰

New Technology and Protections

Technology Solutions to AI Privacy Challenges

There are many ways in which new tools and protections can be incorporated into AI and machine learning systems to address some of the challenges identified above. Some of those being developed or explored are described here.

In response to the data minimization concerns, there may be ways other than de-identification or traditional data protection techniques to limit access to individual records or personal data, while still allowing immense data sets to be put to beneficial uses.

Differential Privacy

If a database contains PII on individual users, information can be retrieved from the database in a way that ensures the response will contain deliberately-generated “noise.” This means that general information about groups of records can be retrieved for the necessary analysis or processing, but precise details about specific individuals would not be available. In order for the noise to effectively prevent re-identification of individual records, the database must be sufficiently large and diverse such that the removal of a specific individual record will not cause a notably different result to the same query. Thus, the overriding trends or characteristics of the dataset remain consistent and may still be used as needed for the machine learning functions. Differential privacy can therefore add a degree of privacy protection for large datasets in a way that is distinct from a data minimization approach (or something).

Homomorphic encryption:

As briefly referenced above, homomorphic encryption permits the manipulation of data while it remains encrypted. This allows for a broad range of functions to be performed, or data manipulation to still occur, without revealing the clear text content of the data. Each party can maintain the confidentiality of individual records, even while sharing for research or other purposes. This process is still in development, and there will likely always be some limits on the possible uses of encrypted data, but to the extent that it is applicable, this tool will allow much greater flexibility, particularly in situations where the data held by each party is particularly sensitive, and the objectives they hope to accomplish by sharing it particularly significant. It is generally less computationally efficient, and some level of industry standardization may ultimately be required, but many companies and data scientists are exploring these possibilities.

Transfer Learning

¹⁰ <https://www.wired.com/2016/09/how-to-steal-an-ai/>

This is a practice where models that have already been developed using extensive data sets can be used or easily adapted to similar or related purposes without having to start from zero. Basing programming on existing models enables the use of machine learning platforms to perform new functions. The adapted model can achieve a final state with much less data, such that less data needs to be collected-- and more quickly. Similar to open source programming code, there are shared locations where these models reside, for access by new users as needed.

Limited Access

There are tools under development that would allow researchers and others to have some amount of access to datasets that include individual records, without having direct access to all the individual data. Researchers might only be able to access the metadata of the underlying data set, receiving back aggregated results to queries rather than a breakout list of individual records that fit the search criteria. This can be particularly useful for data sets that include small groups, or individual records that are unique enough for easy identification. This is particularly applicable for data required for machine learning processes.¹¹

Explainability

This is a term for the idea that all automated decision-making should be able to be explained. Whether for legal compliance, or simple fairness, individuals have a right to understand what data related to them was used, how it impacted the outcomes, and what the process overall considered or accomplished. This will not be possible in all systems, but where it can be achieved, it will contribute to the understanding and trust in AI-decision-making programs more generally. There are two examples of XAI in development now:

- The Defense Advanced Research Projects Agency (DARPA) is making a significant investment of time and resources into a search for broadly applicable XAI. One current research project is with Oregon State University, to create an AI that can explain its decisions in such a way that most people can understand it. DARPA is committed to achieving breakthroughs in this area.
- LIME is another approach to XAI.¹² It stands for Local Interpretable Model-Agnostic Explanations. It is the development of techniques to explain the predictions of a machine learning classifier, and evaluate its usefulness in various tasks related to trust. This involves changing (or “perturbing”) the input data to see how the model changes. We can then generate an explanation that approximates the underlying model, at a level that can be understood by someone trying to understand the system.

¹¹ RAIRD23 – Norwegian program under development

¹² <https://www.oreilly.com/learning/introduction-to-local-interpretable-model-agnostic-explanations-lime>

Recommendations (desired? For variety of audiences? Or just educational focus at this point...)

Conclusion

(to be written)

AI and Machine Learning - Resources for further reading:

1. Webpage/articles - From 5-60 minutes to read through. Almost all will cover the same basic concepts (what is AI, machine learning, deep learning, strong/weak AI, etc)
 - a. “The CIO’s Guide to Artificial Intelligence”
<https://www.gartner.com/smarterwithgartner/the-cios-guide-to-artificial-intelligence/>
 - b. “The Beginner’s Guide to Understanding Artificial Intelligence” <https://uxdesign.cc/the-beginners-guide-to-understanding-artificial-intelligence-3b93e844591>
 - c. “The AI Revolution: The Road to Superintelligence”
<https://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html>
 - d. IBM- “A Beginner’s Guide to Artificial Intelligence, Machine Learning, and Cognitive Computing” <https://www.ibm.com/developerworks/library/cc-beginner-guide-machine-learning-ai-cognitive/index.html>
 - e. Quartz Guide to AI: <https://qz.com/1046350/the-quartz-guide-to-artificial-intelligence-what-is-it-why-is-it-important-and-should-we-be-afraid/>
 - f. <https://medium.com/machine-learning-for-humans/why-machine-learning-matters-6164faf1df12>
2. EU
 - a. Big data, artificial intelligence, machine learning and data protection. <https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>
 - b. ROYAL SOCIETY: MACHINE LEARNING: THE POWER AND PROMISE OF COMPUTERS THAT LEARN BY EXAMPLE. <https://royalsociety.org/~media/policy/projects/machine-learning/publications/machine-learning-report.pdf>
3. Provision of additional resources:
 - a. Medium- “The Best Machine Learning Resources”: <https://medium.com/machine-learning-for-humans/how-to-learn-machine-learning-24d53bb64aa1>
 - b. Practical explanation/overview for individuals wanting to get experience in or enter the AI field: <http://blog.hackerearth.com/artificial-intelligence-101-how-to-get-started>
4. Interactive Explanation and Courses
 - a. Udacity has a free 4-month “Intro to Artificial Intelligence” course to AI (I believe they are an outgrowth of Stanford’s free computer classes)
<https://www.udacity.com/course/intro-to-artificial-intelligence--cs271>
 - b. Salesforce interactive explanation “Artificial Intelligence Basics”
https://trailhead.salesforce.com/en/modules/ai_basics (2 15-minute basic introductions)
5. PDF-length
 - a. “AI for Dummies” http://gunkelweb.com/coms493/texts/AI_Dummies.pdf
 - b. **Tutorials Point**
 - i. Webpage:
https://www.tutorialspoint.com/artificial_intelligence/artificial_intelligence_neural_networks.htm
 - ii. **PDF**: Shorter than a book, but more in depth than an article. In terms of content, an overview for individuals who are interested in the development side of AI with good descriptions of multiple areas/components inside AI, complete with terminology, examples of application in the real world, and tracking how the machine receives input and produces its output.

https://www.tutorialspoint.com/artificial_intelligence/artificial_intelligence_tutorial.pdf

6. Other:

- a. <https://medium.com/towards-data-science/artificial-intelligence-understanding-the-hype-daee0df04695>
- b. <https://aws.amazon.com/amazon-ai/what-is-ai/>
- c. <https://futurism.com/this-is-what-a-true-artificial-intelligence-really-is/>
- d. <https://code.facebook.com/posts/384869298519962/artificial-intelligence-revealed/>
- e. https://www.bernardmarr.com/default.asp?contentID=958&utm_content=buffer581e4&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer
- f. <http://www.wired.co.uk/article/machine-learning-ai-explained>
- g. <https://www.intel.com/content/www/us/en/analytics/ai-luminary-reza-zadeh-video.html>
- h. <http://www.campaignlive.co.uk/article/introduction-algorithms-machine-learning-ai/1431094>
- i. <https://medium.com/@matthewbiggins/ai-and-the-future-of-ethics-e4286567e742>
- j. <https://medium.com/machine-learning-for-humans/why-machine-learning-matters-6164faf1df12>
 - i. Part 1: Why Machine Learning Matters
 - ii. [Part 2.1: Supervised Learning](#)
 - iii. [Part 2.2: Supervised Learning II](#)
 - iv. [Part 2.3: Supervised Learning III](#)
 - v. [Part 3: Unsupervised Learning](#)
 - vi. [Part 4: Neural Networks & Deep Learning](#)
 - vii. [Part 5: Reinforcement Learning](#)
 - viii. [Appendix: The Best Machine Learning Resources](#)
- k. <https://medium.com/machine-learning-for-humans/how-to-learn-machine-learning-24d53bb64aa1>
- l. <https://bravenewcoin.com/news/how-artificial-intelligence-is-affected-by-bias-in-data/>