Before the

U.S. Patent and Trademark Office

In the Matter of

Request for Comments
on Patenting Artificial Intelligence Inventions

Docket Number PTO-C-2019-0029

Comments of Engine Advocacy & The Electronic Frontier Foundation

Tyler D. Robbins
Certified Law Student
Phil Malone
Counsel for Commenters

Juelsgaard Intellectual Property
& Innovation Clinic
Stanford Law School
(650) 724-1900
pmalone@stanford.edu

November 8, 2019

Contents

About the Commenters

Engine Advocacy – Engine Advocacy is a non-profit policy, research, and advocacy organization that supports high-growth, high-tech startups. Engine works with federal, state, and local government; international advocacy organizations; and a community of growth-oriented technology startups nationwide to support the development of technology entrepreneurship. Engine conducts research, organizes events, and spearheads campaigns to educate elected officials, entrepreneurs, and the general public on issues vital to fostering technological innovation, including improving patent quality. Engine works with many patent owners and innovators, including artificial intelligence startups. Engine has seen the detrimental impact improperly designed patents can have on innovation and appreciates the Patent Office's attention to this important subject area.

The Electronic Frontier Foundation – EFF is a non-profit civil liberties organization that has worked for more than 25 years to protect consumer interests, innovation, and free expression in the digital world. EFF and its more than 30,000 dues-paying members care deeply about ensuring that intellectual property law in this country serves the goal set forth in the Constitution: promoting the progress of science and technological innovation. To ensure the voices of consumers, end users, and developers are heard, EFF has often provided comments on behalf of the public's interest in the patent system to the USPTO, including on patent-eligibility requirements and their impact on innovation in the software industry.

Introduction

Artificial intelligence ("AI") is a quintessential disruptive technology. It has already significantly affected aspects of our everyday lives, from healthcare[1] to entertainment.[2] And it is difficult to imagine an industry or sector AI will not touch in the future. While many AI technologies are already ubiquitous, we are still in the early stages of an AI revolution, with myriad new advanced techniques and commercial applications on the horizon.

But, despite its transformative tendencies, AI does not need to disrupt the U.S. patent system. Patents have adapted to accommodate revolutionary technologies in the past, such as computer software and genetic engineering. While our patent policies should account for the value of emerging AI technologies – and we commend the Patent Office for seeking public input – the U.S. patent system does not now need substantial changes to accommodate AI.

Existing statutes, regulation, guidance, and case law map well onto the types of AI inventions commonly produced today and on the immediate horizon. The following comments focus on how the current frameworks for subject matter eligibility under § 101 and enabling disclosure under § 112 should apply to AI inventions for promoting progress and establishing high-quality patents in the field.

AI technologies perform tasks that conventionally require human intelligence, such as learning, reasoning, and perception. Usually, these technologies are implemented as computer software or hardware. AI is a broad discipline, including technologies such as expert systems, fuzzy logic, and

---

[1] Brian Kalis, Matt Collier & Richard Fu, *10 Promising AI Applications in Health Care*, Harvard Business Review (May 10, 2018), https://hbr.org/2018/05/10-promising-ai-applications-in-health-care.

[2] *Meson: Workflow Orchestration for Netflix Recommendations*, Netflix Technology Blog (May 31, 2016), https://medium.com/netflix-techblog/meson-workflow-orchestration-for-netflix-recommendations-fc932625c1d9.

robotics. But, regardless of the specific technology, most AI innovations today involve machine learning methods.

Machine learning methods solve problems without being explicitly programed. These methods have roots in statistical modeling and largely use statistical methods. Generally, both statistics and machine learning develop mathematical models from analyzing the inputs and outputs of a process. However, whereas traditional statistics tries to define a model of the process itself, machine learning methods try to predict the outputs of that process without trying to model or understand how it works.[3] By treating the process as unknown while trying to functionally approximate it, machine learning methods can learn to perform incredibly complicated and not well-understood tasks, such as object detection, that would be difficult, if not impossible, to explicitly program.

The basic development process of an AI system that uses machine learning is:

1. Defining the problem to solve.
2. Gathering and preparing data for training the AI system.
3. Selecting and building the machine learning model(s) for use in the AI system.
4. A loop of training, testing, and tuning the AI system until it is either ready for deployment or some aspect of it needs to be redesigned.

See the Appendix for diagrams of this development process, a general machine learning model, and the basic model training process.[4] Common

[3] Leo Breiman, *Statistical Modeling: The Two Cultures*, 16 Statistical Science 199 (2001), http://www2.math.uu.se/~thulin/mm/breiman.pdf.

[4] To create the diagrams, the commenters relied on two sources: Victor Roman, *How to Develop a Machine Learning Model from Scratch*, Towards Data Science (Dec. 23, 2016),

examples of machine learning methods include neural networks, support vector machines, and decision trees.

The request for comments refers to two types of AI inventions: Inventions produced by AI itself and inventions that use AI to solve tasks. AI-produced inventions may pose many difficult and interesting questions for the patent system, such as whether an AI can legally be the inventor, which entity – if any – should own the patent, and how an AI-inventor affects obviousness questions and the definition of an ordinarily skilled artisan.

From the perspective of promoting innovation, however, the second question – how the patent system should handle inventions that use AI – is currently more pertinent. Inventors file an increasing number of AI patent applications each year,[5] with claims we consider to fall into three general categories: (1) Methods for developing AI systems to perform specific tasks; (2) technical improvements to the development process; and (3) the AI system created at the end of the process. Depending on the scope of the claims and detail of the specification, inventions in each category may encounter issues with subject matter eligibility or enabling disclosure in a patent application.

These comments will address three points responsive to questions from the Patent Office's request for comments. First, regarding subject matter eligibility and application of the *Alice/Mayo* test, the Patent Office should prevent the issuance of patents which generically claim the performance of abstract ideas using generic, conventional AI technology, as such overbroad claiming would preempt downstream innovation (Question 5). Second, AI invention can be an unpredictable art, so satisfying the enablement requirement should require detailed descriptions of both the model

---

https://towardsdatascience.com/machine-learning-general-process-8f1b510bd8af; Ayush Pant, *Workflow of a Machine Learning Project,* Towards Data Science (Jan. 10, 2019), https://towardsdatascience.com/workflow-of-a-machine-learning-project-ec1dba419b94.

[5] *WIPO Technology Trends 2019: Artificial Intelligence*, World Intellectual Property Organization (2019), https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf.

and the input (Question 7). Finally, any changes to the patent system for AI should promote innovation and disclosure (Question 11).

5. Are there any patent eligibility considerations unique to AI inventions?

The current framework for assessing patent eligible subject matter should apply to patents on AI

AI inventions do not require unique patent eligibility law or policy considerations. Recently decided cases, especially since *Alice*,[6] have helped to reduce the proliferation of weak, overbroad patents by refining subject matter eligibility standards. The Patent Office should be careful these standards do not regress when considering AI inventions. The *Alice/Mayo* test[7] should apply to AI inventions just as well as any other invention. While an AI invention embodied in software is more likely to be directed towards an abstract idea than an AI invention embodied in specific hardware, resolving this issue should require no special analytical steps outside applying the *Alice/Mayo* test as usual.

AI in Specific Hardware – Inventions where AI is implemented in specific hardware, rather than a generic computer, should not encounter novel eligibility issues under the *Alice/Mayo* test. For example, claims involving AI microprocessors or machines that use trained AI to complete a specific task may not be directed towards abstract ideas. Such claims would not recite any mathematical formula, method of organizing human activity, or mental processes. And, even if they did, such inventions often implement AI as part of a greater system in which the hardware provides a

[6] Alice Corp. Pty. Ltd. v. CLS Bank Int'l, 573 U.S. 208 (2014).

[7] *See Manual of Patent Examining Procedure, Section 2106*, U.S. Patent & Trademark Office (last revised Jan. 2018) (citing *Alice*, 573 U.S. at 216; Mayo Collaborative Servs. v. Prometheus Labs., Inc., 566 U.S. 66, 71 (2012)).

technological solution to a technological problem. We expect that for many hardware AI inventions, there would be few cases that progress past the first step of the *Alice/Mayo* test.

AI in Software – Most AI inventions are a system of algorithms implemented in computer software that run on generic computer hardware. In the context of subject matter eligibility, these inventions should be considered a subcategory of software inventions and subjected to standard software patent analysis. Valid AI patent claims, for example, could include methods for training an AI to perform a task or improvements to the training process itself.

AI software inventions should not be any more patentable than typical software. If anything, AI patent claims are more likely to be directed towards abstract ideas under *Alice* and subsequent case law than other software inventions. This is primarily for two reasons:

First, unlike typical software that is programed with explicit constructions, AI software often uses self-learning algorithms to achieve tasks. Essentially, these algorithms – many of which are now conventional, widely available "off-the-shelf" technology – optimize mathematical models for approximating an opaque phenomenon from its inputs and outputs. Claiming such a process, without some additional limitations, can be the sort of data manipulation and generation the Federal Circuit considered patent-ineligible in *Digitech*[8] and *SAP America*.[9] Additionally, a claim just for the trained model itself should be ineligible because it is a mathematical representation of the relationship between the input and output data.

[8] Digitech Image Techs., LLC v. Electronics for Imaging, Inc., 758 F.3d 1344, 1351 (Fed. Cir. 2014) ("Without additional limitations, a process that employs mathematical algorithms to manipulate existing information to generate additional information is not patent eligible.").

[9] SAP Am., Inc. v. InvestPic, LLC, 898 F.3d 1161, 1163 (Fed. Cir. 2018) (holding "mathematical calculations based on selected information and the presentation of the results of those calculations" not patent eligible).

Second, AI inventions, by definition, try to perform tasks traditionally requiring human intelligence. Broad claims for AI inventions may thus be directed towards concepts analogous to human mental work, from conscious processes, like making predictions, to implicit processes, like perception. AI technologies are powerful information processing systems that can perform a wide variety of tasks. The simple idea of using an AI system to do a typical human mental task, without more, is not the type of inventive contribution warranting patent protection. As *Alice* established with computers, merely using AI to carry out an abstract idea, such as mental steps, should not be enough to transform an abstract idea into patent eligible subject matter.

Overall, software patents should be the ceiling for AI software patents. AI inventions can be subjected to the same analysis and case law as other software. While, by nature, AI software tends to be a little more abstract, AI inventions pose no subject matter eligibility issues outside the domain of the current frameworks of the U.S. patent system.

7. How can patent applications for AI inventions best comply with the enablement requirement, particularly given the degree of unpredictability of certain AI systems?

AI inventions are an unpredictable art that require extra detail for an enabling disclosure

Patents on AI systems must properly disclose both the model and the input to meet the enablement requirement. AI systems thrive in an uncertain world. They identify imperceptible patterns and make unexplainable predictions from noisy, complex, and incomplete information. Much like their environment, AI systems themselves are volatile. Their self-learning algorithms require minimal human intervention and often use random processes to explore possible solutions. It can be difficult to tell whether a

system learned a generalized solution to a given problem or a solution that only works with its training data. There is no guarantee running the same algorithm on the same dataset will even reach the same result. These characteristics make AI inventions a highly unpredictable art.

In the patent context, this unpredictability can pose major problems for § 112's enablement requirement, especially for patents that claim either an AI system for doing a specific task or a method for training such a system. In addition to being an inherently random and unpredictable process, designing and training AI systems involves significant discretion and creativity. Even a person having ordinary skill in the art may need to engage in a lengthy trial-and-error process to recreate the invention without appropriately detailed disclosures. The patent system already has a framework in place to prevent such undue experimentation.[10] As with other unpredictable arts, AI inventions require a more specific description for an enabling disclosure than more predictable arts.

At minimum, patents relating to AI systems performing a task must disclose in detail both the machine learning models used and how to train them. AI's strength is that it can learn to perform tasks without being explicitly programmed. Popular machine learning methods, such as deep learning, effectively create black boxes that make predictions about the world without explaining how they do it. Thus, an AI system's usefulness and functionality happens outside an inventor's direct control and often understanding. Essentially, an inventor's role in the process is setting up the proper environment for an AI system to learn its task. Enabling patents must teach the public how to set up this environment.

---

[10] *See Manual of Patent Examining Procedure, Section 2164.03,* U.S. Patent & Trademark Office (last revised Jan. 2018) ("The amount of guidance or direction needed to enable the invention is inversely related to the amount of knowledge in the state of the art as well as the predictability in the art.").

Proper model disclosure includes the model architecture and objective function

Disclosing a model's architecture and objective function is necessary for an ordinarily skilled artisan to recreate the invention without undue experimentation. Even machine learning models of the same type can vary drastically. For example, convolutional neural networks, which are well suited for tasks with visual data, have an arbitrary number and arrangement of layers, each with a set height, width, and depth.[11] Patents must describe any variations from prior art model architectures in detail. Otherwise, a skilled artisan will be forced to conduct many experiments to find the correct architecture from countless possible ones. For the deep neural networks widely used today, this process becomes more difficult the more layers there are in the network. More layers generally help networks perform more complex tasks.[12] However, deeper layers' functionalities are also increasingly unpredictable and difficult to visualize and understand intuitively.[13] AI inventions using deep learning should contain extra guidance on how to put the model in a position to learn successfully without substantial trial and error.

In addition to the model architecture, enabling disclosures must also include the model's objective function. Objective functions mathematically define the task a model is supposed to perform. They measure the level of "success" for any given set of parameters. Any changes to the objective function mean the model technically is being optimized for a different task.

---

[11] Andrej Karpathy, *CS231n Convolutional Neural Networks for Visual Recognition*, (last visited Nov. 8, 2019), https://cs231n.github.io/convolutional-networks.

[12] Jason Brownlee, *How to Configure the Number of Layers and Nodes in a Neural Network,* Machine Learning Mastery (July 27, 2018), https://machinelearningmastery.com/how-to-configure-the-number-of-layers-and-nodes-in-a-neural-network.

[13] Jason Yosinski et al., *Understanding Neural Networks Through Deep Visualization*, Cornell Univ. Computer Science Dept. (June 22, 2015), https://arxiv.org/abs/1506.06579.

Despite being a relatively small part of the entire training process, the objective function needs to be exact to ensure the AI system is learning as intended. There are many well-documented and commonly used objective functions.[14] A model using one of those objective functions would only have to specify which one to enable a skilled artisan to recreate the AI system. However, any custom objective functions should either be expressed mathematically or in pseudo-code in the specification.

Proper input disclosure requires either a detailed description of a training data's structure or disclosing the training data itself

Describing the model, while necessary, is alone not sufficient for enablement. Proper patent disclosures must also include how an inventor trained the model, most significantly the model's input data. The input is so important to AI systems that phrases such as "garbage in, garbage out" and "data is the new oil" became mantras in the community. Managing the input is at least as crucial to an AI system's success and requires as much inventor ingenuity as preparing the algorithm.[15] Training an AI system, no matter how carefully designed it is, could be impossible without an appropriate dataset.

A patent that only discloses what sort of generic data is needed to train an AI system is not enough to meet the enablement requirement. Given the unpredictability of AI systems, a skilled artisan needs more guidance to avoid undue experimentation in recreating the proper input. Any of the

---

[14] Lars Hulstaert, *Understanding Objective Functions in Neural Networks* (Nov. 4, 2017), https://towardsdatascience.com/understanding-objective-functions-in-neural-networks-d217cb068138.

[15] *In Machine Learning, What is Better: More Data or Better Algorithms,* KD Nuggets (June 2015), https://www.kdnuggets.com/2015/06/machine-learning-more-data-better-algorithms.html; Chen Sun et al., *Revisiting Unreasonable Effectiveness of Data in Deep Learning Era,* Cornell Univ. Computer Science Dept. (July 10, 2017), https://arxiv.org/abs/1707.02968.

steps taken to prepare the data and the characteristics of the dataset's structure could prove dispositive in the system's ability to learn. This is especially important for tasks where the input itself is not readily understandable for humans, such as data with thousands of inputs or that is dimensionally reduced. An inventor does not necessarily know which aspects of any input are essential for the self-learning algorithms to optimize for their tasks. To meet the enablement requirement, a patent for an AI invention should disclose in detail what sort of data the model needed to learn.

An inventor has at least two options for sufficiently disclosing the data necessary for training an AI system. First, an inventor could describe the training data in the specification. The description should include all data characteristics necessary for the AI to learn the task. For example, for a dataset of facial images, these characteristics might include the quantity and demographics of the people, the lighting conditions and camera angles of the images, or any adjustment filters and sampling applied to the final input. If appropriate, the figures might also include some drawings of mock inputs. Such guidance, in contrast to a terse direction such as "collect a large sample of facial data," would allow a skilled artisan to recreate a similar dataset without first trying to figure out the data's necessary features through guessing and experimenting.

Alternatively, an inventor could opt to disclose the dataset directly. For example, if an AI system were trained using a commonly available dataset, such as ImageNet,[16] citing it would be a sufficient disclosure. An inventor might also self-publish the data. These data disclosure methods echo the patent procedure for biological deposits. Like AI systems, those self-replicating organisms function in invention-relevant ways outside the inventor's direct control. And, as with biological deposits, a dataset description may not convey all the necessary information to enable a skilled artisan to make or use an AI invention. The Patent Office could even create and maintain a

[16] ImageNet Homepage (last visited Nov. 8, 2019), http://image-net.org/index.

data depository. The public could request samples from the depository and inventors would be able to meet the enablement requirement without self-publishing the data.

Overall, the patent system already has the tools necessary to assess whether patent applications comply with the enablement requirement. AI inventions are a particularly unpredictable art, highly dependent on the inventor's design choices and the input data. Patent applications could greatly help enablement analysis by indicating how unpredictable the AI system and training process are. Assuming it met the eligibility requirement, a system implementing well-known models, like AlexNet,[17] on openly available datasets would require little extra disclosure because these systems are so well-documented. However, if a completely custom algorithm was trained under specific circumstances, even a skilled artisan would need to know how exactly to set up the model's environment to recreate the invention.

Because the enabling requirement is a high bar for AI inventions, a patent application with sufficient details for an enabling disclosure may signal that other requirements for patent eligibility have also been met. By requiring more detailed descriptions, it seems likely an AI patent application that meets the enablement requirement can also meet the written description requirement. Moreover, detailed descriptions in the specification must also correlate to precision in the claims. Precise claims are essential to make the public aware of what exactly the invention is, so they know if they infringe the patent. The enablement and written description requirements must be carefully and thoroughly vetted during prosecution to ensure that AI patents serve their notice and disclosure functions.

---

[17] Alex Krizhevsky et al., *ImageNet Classification with Deep Convolutional Neural Networks*, Neural Information Processing Systems Conference (2012), https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf.

11. Are there any other issues pertinent to patenting AI inventions that we should examine?

The AI field already exhibits substantial innovation and disclosure under the current patent framework

The U.S. patent system exists to promote innovation and disclosure. Over the past decade, few fields have experienced a greater explosion in innovation than AI. There are new developments and advancements in AI technology all the time, as investment, research, and interest in the industry increased drastically over the past few years.[18] AI development thrives everywhere, from universities and research institutions to startups and large companies. To remain globally competitive, the U.S. needs to keep investing in AI research and education.[19]

Notably, this innovation and growth in AI all happened under the current patent framework. AI development does not need extra incentive or stimulation. Thus, when considering any changes to how it evaluates AI patent applications, the Patent Office should carefully consider how those changes might impact this "Golden Age" of AI and whether they will hamper innovation. Easing the requirements for patentable AI inventions could have negative effects. In particular, the risk of permitting weak, overbroad patents, like those that plagued the patent system prior to *Alice*, could end up restricting downstream innovation in this field.

Additionally, the Patent Office should consider how any changes to the current patent framework might impact the high level of disclosure already

---

[18] AI Index Steering Committee, *The AI Index 2018 Annual Report,* Human-Centered AI Initiative, Stanford University (2018), http://cdn.aiindex.org/2018/AI%20Index%202018%20Annual%20Report.pdf

[19] Tom Simonite, *China is Catching Up to the US in AI Research – Fast,* Wired (Mar. 13, 2019), https://www.wired.com/story/china-catching-up-us-in-ai-research.

in the AI development community. AI systems and tools are often distributed for free online under permissive licenses, whether they were developed by for-profit corporations, non-profit reach institutions, or a dedicated hobbyist. For example, TensorFlow,[20] one of the most powerful and widely used libraries for machine learning in the world, is distributed under the free and open source Apache 2.0 license,[21] which includes an explicit patent grant. Patents naturally support the community's open culture by promoting the disclosure of complex AI inventions that might otherwise remain secret in exchange for granting a temporary monopoly on the invention. However, patents without a proper enabling disclosure allow this monopoly without teaching the public how to make or use the invention. Granting these types of patents not only results in patents that fail under the statute, but also undercuts the open ideals of the community, whose work many in the field depend on and profit from.

Finally, the Patent Office should also understand the greater policy contexts of AI inventions when examining issues around AI patents. Opaque AI systems trained with biased datasets are negatively impacting the lives of already marginalized communities.[22] While addressing these issues is not necessarily within the Patent Office's mandate, the patent

---

[20] TensorFlow homepage (last visited Nov. 8, 2019), https://www.tensorflow.org.

[21] *Apache License, Version 2.0,* Apache Software Foundation (last visited Nov. 8, 2019), https://www.apache.org/licenses/LICENSE-2.0.

[22] *See, e.g.*, Julia Angwin et al., *Machine Bias*, ProPublica (May 23, 2016), https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing; Hal Hodson, *Police Mass Face Recognition in the US Will Net Innocent People*, NewScientist (Oct. 20, 2016), https://www.newscientist.com/article/2109887-police-mass-face-recognition-in-the-us-will-net-innocent-people; Natasha Singer, *Amazon Is Pushing Facial Technology That a Study Says Could Be Biased*, N.Y. Times (Jan. 24, 2019), https://www.nytimes.com/2019/01/24/technology/amazon-facial-technology-study.html; Jeffery Dastin, *Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women*, Reuters (Oct. 9, 2018), https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G; Aaron Rieke & Corrine Yu, *Discrimination's Digital Frontier*, Atlantic (Apr. 15, 2019), https://www.theatlantic.com/ideas/archive/2019/04/facebook-targeted-marketing-perpetuates-discrimination/587059.

system itself can nonetheless have an impact. Strong patents with precise claims can better help the public understand how AI systems make decisions that impact their lives. And properly disclosing the data used to train AI systems in order to meet the enablement requirement can give the public the opportunity to check for bias in the dataset or data collection process. With AI inventions, patents have a special opportunity to promote not only the progress of science and useful arts, but also the public's general welfare.

Appendix

Basic Machine Learning Development

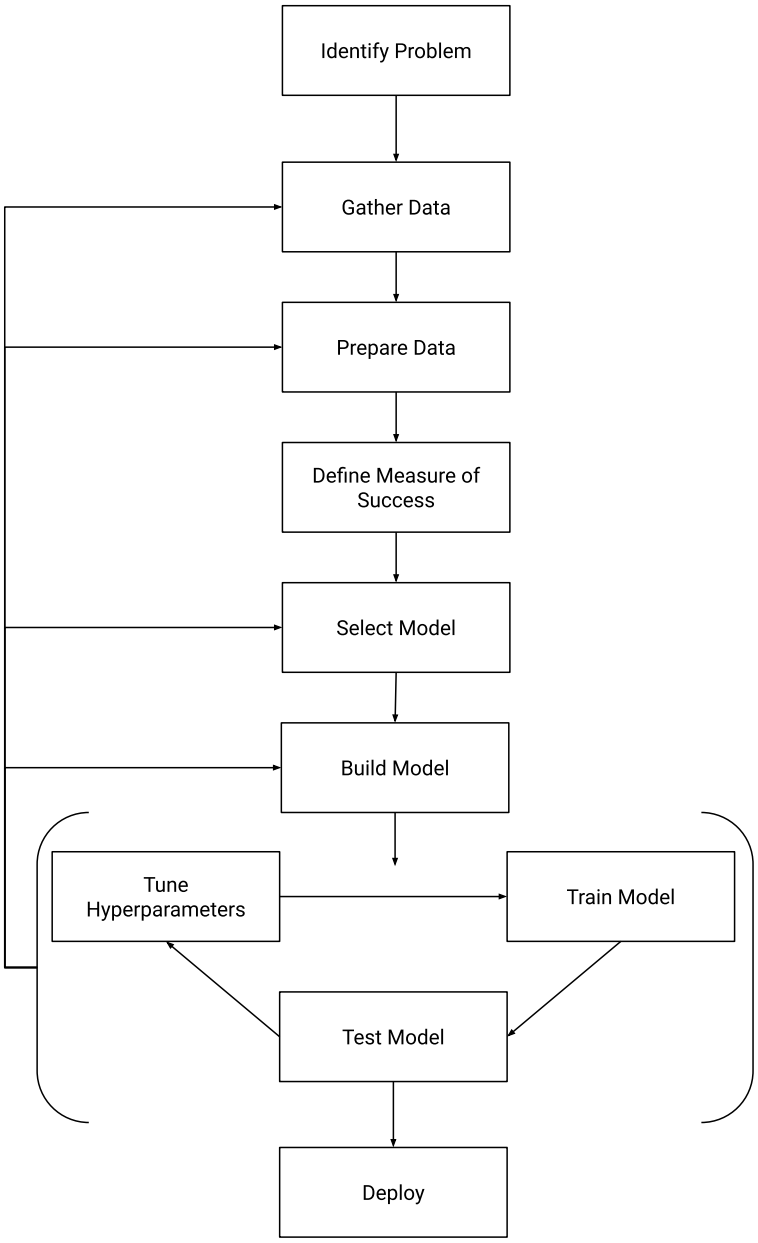

Generalized Machine Learning Model

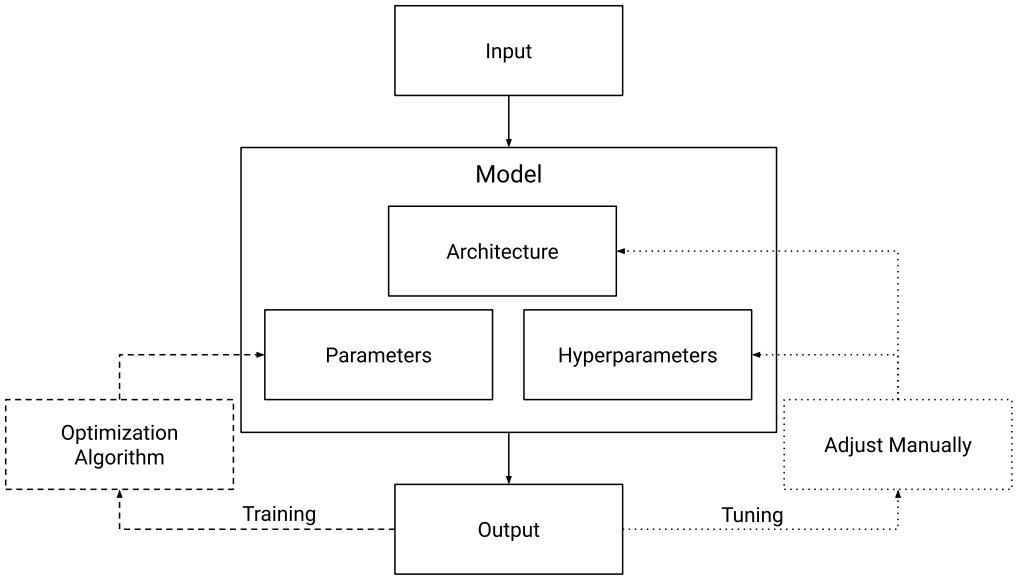

Basic Model Training Process