IP Protection of AI

As companies develop innovative artificial intelligence ("AI") technologies, they face difficult questions on how to protect them. U.S. courts have historically been skeptical of patent protection for software and computer-implemented technologies. That skepticism can extend to AI advances embodied in software and functional algorithms (which is what most AI innovations consist of) rather than specific hardware.

Recent decisions suggest there may be opportunities to protect AI software and algorithms through patents—including for generative technologies. But doing so can be tricky. AI companies must watch for pitfalls that have undermined attempts to patent similar technology. Trade secret protection may offer a better blend of risks and rewards than patents for some AI innovations—particularly where they face uncertain prospects for patent eligibility.

I. Background: Patenting Software

The U.S. patent system allows patenting "any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof," subject to the other requirements of the Patent Act. 35 U.S.C. § 101. This can include innovations in software where those innovations constitute a "process, machine, manufacture, or composition of matter" under § 101.

The U.S. patent system does *not* allow patents on "abstract ideas"—for example, laws of nature, fundamental economic practices, mathematical formulas, methods of organizing human activity, or mental processes. This has made some efforts to patent software and computer-implemented technology challenging. Where (1) algorithms in software simply replicate human mental processes or business practice, (2) a standard computer (*e.g.*, a personal computer or laptop) is used to automate computing a mathematical formula or performing a conventional business practice historically performed by humans, or (3) a patent claim recites the functionality to be performed by software and/or computers, without describing a specific implementation, courts typically conclude the invention is not patentable because it represents an "abstract idea" under § 101.

Challenges to software patents typically focus on whether the patent is eligible for patent protection. The Supreme Court set forth a two-step test to determine patent eligibility in *Alice Corporation v. CLS Bank*, 573 U.S. 208 (2014). First, courts consider whether the patent claim is directed towards an abstract idea—for example, attempting to patent a mathematical formula or mental process. Second, if the patent claim is directed to an abstract idea, courts then "search for an inventive concept, or some element or combination of elements sufficient to ensure that the claim in practice amounts to significantly more than a patent on an ineligible concept." *DDR Holdings, LLC v. Hotels.com, L.P.*, 773 F.3d 1245, 1255 (Fed. Cir. 2014). To survive at step two of this analysis a claim cannot merely recite the application of an abstract idea using conventional and well-understood techniques (*e.g.*, implementing the abstract idea on a standard computer using routine data processing functions); the claim must provide a specific technical solution to a particular technical problem in order to recite an "inventive concept."

Application of these steps are illustrated by two of the leading cases in support of patenting software, *Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327, 1335 (Fed. Cir. 2016) and *McRO, Inc. v. Bandai Namco Games Am. Inc.*, 837 F.3d 1299, 1314 (Fed. Cir. 2016). In *Enfish*, the Federal Circuit considered the patentability of a method of storing data in a table. The representative claim recited:

> A data storage and retrieval system for a computer memory, comprising:
>
> means for configuring said memory according to a logical table, said logical table including:

> a plurality of logical rows, each said logical row including an object identification number (OID) to identify each said logical row, each said logical row corresponding to a record of information;
>
> a plurality of logical columns intersecting said plurality of logical rows to define a plurality of logical cells, each said logical column including an OID to identify each said logical column; and
>
> means for indexing data stored in said table.

Enfish, 822 F.3d at 1336. The court explained that this claim represented a technique for a "self-referential" data model, which unlike conventional approaches could (1) store multiple data types in a single table, and (2) could contain, as a column value, a reference to a row. *Id.* at 1332. It also noted that the patent specification described how the claimed method improved computer processing itself by allowing for faster searching of data and more effective storage of images and unstructured text. The defendant argued that the "the invention's ability to run on a general-purpose computer doom[ed] the claims." *Id.* at 1338. But the Federal Circuit explained that unlike patent-ineligible claims that "recited use of an abstract mathematical formula on any general purpose computer," the claims at issue here were directed to "a specific improvement to computer functionality." *Id.* at 1338. The court concluded that claims such as this "directed to a specific implementation of a solution to a problem in the software arts . . . are not directed to an abstract idea." *Id.* at 1339.

In *McRO*, the Federal Circuit considered claims directed to a method for automatically animating the lip synchronization and facial expressions of animated characters during speech, so that their lips and faces would match the script. 837 F.3d 1299. Claim 1 recited:

> A method for automatically animating lip synchronization and facial expression of three-dimensional characters comprising:
>
> obtaining a first set of rules that define output morph weight set stream as a function of phoneme sequence and time of said phoneme sequence;
>
> obtaining a timed data file of phonemes having a plurality of sub-sequences;
>
> generating an intermediate stream of output morph weight sets and a plurality of transition parameters between two adjacent morph weight sets by evaluating said plurality of sub-sequences against said first set of rules;
>
> generating a final stream of output morph weight sets at a desired frame rate from said intermediate stream of output morph weight sets and said plurality of transition parameters; and
>
> applying said final stream of output morph weight sets to a sequence of animated characters to produce lip synchronization and facial expression control of said animated characters.

Id. at 1307–08. The patent further explained that the innovative method used predetermined rules to specify the lip synchronization of a 3–D animated face based on a transcript. Defendants argued that these claims simply used a computer to automate a conventional activity, noting that animators routinely did lip-synchronization by hand. But the court explained that there was "no evidence that the process previously used by animators [was] the same as the process required by the claims." *Id.* at 1314. Instead, it was "the incorporation of the claimed rules, not the use of the computer, that improved the existing technological process by allowing the automation" of the task. *Id.* (brackets and internal quotation marks omitted). The court recognized that "processes that automate tasks that humans are capable of performing are patent eligible if properly claimed," citing to autopilot and facial recognition technologies as examples. *Id.* at 1313 & n.12. The McRO claims recited a patent-eligible innovation by using specific rules that improved on prior approaches, which diverged from traditional manual methods and "achieve[d] an improved technological result in conventional industry practice." *Id.* at 1316. This improvement was therefore patent-eligible subject matter under *Alice*.

As described below, these principles may seem clear-cut on first blush—but their application to various forms of AI innovations can be uncertain.

II. Protecting Innovative AI with Patents

Because of patent eligibility concerns, AI innovations face different prospects for patent eligibility depending on whether they are implemented as hardware or as software. It is relatively straightforward to patent new hardware designs and hardware architectures that are specially geared toward AI—for example, a new and inventive processor designed to execute complex AI models. Specialized hardware constitutes a "machine" under § 101 and is clearly eligible to patent; numerous cases confirm that new and innovative computer hardware is patent eligible, particularly when the design provides technological advantages like improving AI applications. *See, e.g.*, *Visual Memory LLC v. NVIDIA Corp.*, 867 F.3d 1253, 1259 (Fed. Cir. 2017) (patent for a memory system with programable characteristics was patent eligible under § 101 because the claims were "directed to an improved computer memory system, not to the abstract idea of categorical data storage").

Protecting AI innovations embodied as software, or described as an algorithm, may be more challenging. Even the most innovative AI application may rely on algorithms and mathematical techniques that courts view as "abstract" and therefore not patentable. Where AI innovations are described as replicating human mental processes, or performing specific functions or tasks that can be done manually but are instead being done using AI, those innovations may face challenges to patentability. Human mental processes are, by definition, processes that can be carried out by humans without any software. This is precisely what courts often call "abstract" and unpatentable. And simply specifying that the process is performed by a "neural network" or "artificial intelligence" is unlikely to move the needle—that is comparable to the type of "do-it-with-a-computer" claims the Supreme Court held to be abstract and ineligible for patent protection in *Alice Corp. v. CLS Bank Int'l*, 573 U.S. 208, 223 (2014).

Courts have not yet grappled with many cases involving AI patent eligibility. But the decisions thus far show that even impressive and valuable AI software can fall victim to these rules. The following are examples of recent decisions which illustrate the challenges companies face when trying to protect AI through patents.

***In re Bd. of Trustees of Leland Stanford Junior Univ.*, 991 F.3d 1245 (Fed. Cir. 2021)**

In *Stanford*, the Federal Circuit considered whether a particular type of AI model—an "inference" model, trained to infer an individual's specific inherited traits (haplotype) based on that individual's DNA sequence (genotype)—was patent eligible. The patent claim described training and using a specific model (a Hidden Markov Model) to accomplish classification and prediction tasks. Claim 1 of the patent recited:

> 1. A computerized method for inferring haplotype phase in a collection of unrelated individuals, comprising:
>
> receiving genotype data describing human genotypes for a plurality of individuals and storing the genotype data on a memory of a computer system;
>
> imputing an initial haplotype phase for each individual in the plurality of individuals based on a statistical model and storing the initial haplotype phase for each individual in the plurality of individuals on a computer system comprising a processor a memory [sic];
>
> building a data structure describing a Hidden Markov Model, where the data structure contains:

> a set of imputed haplotype phases comprising the imputed initial haplotype phases for each individual in the plurality of individuals;
>
> a set of parameters comprising local recombination rates and mutation rates;
>
> wherein any change to the set of imputed haplotype phases contained within the data structure automatically results in re-computation of the set of parameters comprising local recombination rates and mutation rates contained within the data structure;

> repeatedly randomly modifying at least one of the imputed initial haplotype phases in the set of imputed haplotype phases to automatically re-compute a new set of parameters comprising local recombination rates and mutation rates that are stored within the data structure;
>
> automatically replacing an imputed haplotype phase for an individual with a randomly modified haplotype phase within the data structure, when the new set of parameters indicate that the randomly modified haplotype phase is more likely than an existing imputed haplotype phase;
>
> extracting at least one final predicted haplotype phase from the data structure as a phased haplotype for an individual; and
>
> storing the at least one final predicted haplotype phase for the individual on a memory of a computer system.

Id. at 1247–48 (Fed. Cir. 2021). Essentially, this claim described a specific statistical model, implemented on generic hardware using conventional computer operations (e.g., receiving/storing data).

Stanford argued that the claim was patent eligible, and not abstract, because it described a concrete improvement to previous methods for determining the claimed "inference." The Federal Circuit rejected that argument. The court concluded that training and using software in this way was simply "the use of mathematical calculations and statistical modeling," which, without more, was an ineligible abstract idea. *Id.* at 1250. The court explained that merely improving on prior mathematical calculations or statistical models was not enough; when "the claimed advance . . . merely enhances an ineligible concept," such as improving on mathematical calculations that were themselves not patentable, the improvement is also directed at an abstract idea.[1] *Id.* at 1250–51. The court further concluded that there was no inventive concept because the patented software model was merely math stored as data, and using software in this way merely combined "abstract calculations" with "the well-understood, routine, and conventional steps of receiving and storing data in a computer memory." *Id.* at 1252. The court also explained that "more accurate haplotype predictions than previously achievable under the prior art" was not an inventive concept. *Id.*

Under this reasoning, many AI technologies—even undeniably valuable software that can perform complex operations far faster and more accurately than humans—may face barriers to patent protection. Patent claims that appear to simply improve on the underlying mathematical algorithm, rather than specifically improving on computerized implementations of the underlying mathematical concept, are likely to be viewed as unpatentably "abstract" by courts, despite obvious value to users.

[1] The court did make a distinction between (1) Stanford's claim that the technology resulted in more accurate predictions, which it described as an improvement to a mathematical calculation (which further suggested that the patent was directed to an abstract idea), and (2) Stanford's claim that the technology was an improvement to computing efficiency, which the court found to be forfeited because Stanford did not make the argument below. *See id.* at 1251. If Stanford's argument about improved computing efficiency had not been waived, it is possible that the claim may have survived on that basis. This leaves open the possibility that courts facing similar claims in the future may find that technology patent eligible if the patent owner can argue the patent improves *computing* efficiency—not mathematical accuracy.

***Health Discovery Corp. v. Intel Corp.*, 577 F. Supp. 3d 570 (W.D. Tex. 2021)**

In *Health Discovery*, the plaintiff accused Intel of infringing patents directed to improvements for AI software performing "classification" operations—specifically, improved Support-Vector-Machine (SVM) classification models that can analyze large data sets and classify that data into categories. The representative claim stated:

> 1. A computer-implemented method for identifying patterns in data, the method comprising:
>
> (a) inputting into at least one support vector machine of a plurality of support vector machines a training set having known outcomes, the at least one support vector machine comprising a decision function having a plurality of weights, each having a weight value, wherein the training set comprises features corresponding to the data and wherein each feature has a corresponding weight;
>
> (b) optimizing the plurality of weights so that classifier error is minimized;
>
> (c) computing ranking criteria using the optimized plurality of weights;
>
> (d) eliminating at least one feature corresponding to the smallest ranking criterion;
>
> (e) repeating steps (a) through (d) for a plurality of iterations until a subset of features of pre-determined size remains; and
>
> (f) inputting into the at least one support vector machine a live set of data wherein the features within the live set are selected according to the subset of features.

Id. at 577. This claim described a technique in which: (1) an SVM classifier would be trained on data to recognize features; (2) the features would be ranked by importance; and (3) the lowest-ranked feature would be removed; and then these three steps would be iterated until only a specified number of features for classification remained. But just as in *Stanford*, the claims did not recite any specific computer or hardware implementation; beyond the bare recitation of a "computer-implemented method" and an unspecified "support vector machine," nothing tied the claim to a computer at all.

Intel challenged the patents as ineligible subject matter under 35 U.S.C. § 101. Plaintiff HDC argued that this method was an improvement on previous feature reduction techniques, and solved a problem in which SVM models would become overwhelmed by an excessive number of features to classify. Intel argued that the patent described mental processes that were abstract and could be performed without a computer, rendering them unpatentable.

The district court recognized that "[i]t is difficult to extract a unified theory" of patent eligibility under § 101, and suggested HDC's patent was close to the line. *Id.* at 583. But ultimately, because the technology at issue was machine learning, the court found the *Stanford* decision described above most instructive, as opposed to cases dealing with similar claims in other areas of technology, where the Federal Circuit had upheld patent eligibility. Based on that determination, the court held that the claims "merely describe[d] improving a mathematical analysis," which represented an unpatentable abstract idea. *Id.* at 584. At step 2 of the *Alice* inquiry, the court considered whether there was some inventive concept sufficient to transform the abstract idea into a patent-eligible application. The court was not persuaded by HDC's assertions that patented method was important to the field and was used in variety of applications, or that the original academic paper describing the method had been cited more than eight thousand times. Because HDC had failed to show that its claim was an "innovation in [a] non-abstract application realm," the patent failed under step 2. *Id.* at 586.

Health Discovery again illustrates the challenges that AI innovations will face for patent protection under § 101. Even where software improves on a machine learning algorithm and enables machine learning implementations that are faster and more efficient than prior techniques, courts viewing those inventions through the lens of *Stanford*'s logic may ultimately view those improvements as improvements to the underlying math itself, rather than an improvement in the functioning of the computer itself, and conclude that the inventions are not patent eligible.

***Recentive Analytics, Inc. v. Fox Corp.*, No. CV 22-1545-GBW, 2023 WL 6122495 (D. Del. Sept. 19, 2023)**

In *Recentive Analytics*, the district court considered patent eligibility for AI models trained to create optimized schedules for live events—for example, using AI to optimize a touring schedule for a performer, or using AI to determine the ratings-maximizing TV schedule for programming that could be shown at various times on different networks. One representative claim stated:

> A computer-implemented method of dynamically generating an event schedule, the method comprising:
>
>> receiving one or more event parameters for series of live events, wherein the one or more event parameters comprise at least one of venue availability, venue locations, proposed ticket prices, performer fees, venue fees, scheduled performances by one or more performers, or any combination thereof;
>>
>> receiving one or more event target features associated with the series of live events, wherein the one or more event target features comprise at least one of event attendance, event profit, event revenue, event expenses, or any combination thereof;
>>
>> providing the one or more event parameters and the one or more event target features to a machine learning (ML) model, wherein the ML model is at least one of a neural network ML model and a support vector ML model;
>>
>> iteratively training the ML model to identify relationships between different event parameters and the one or more event target features using historical data corresponding to one or more previous series of live events, wherein such iterative training improves the accuracy of the ML model;
>>
>> receiving, from a user, one or more user-specific event parameters for a future series of live events to be held in a plurality of geographic regions;
>>
>> receiving, from the user, one or more user-specific event weights representing one or more prioritized event target features associated with the future series of live events;
>>
>> providing the one or more user-specific event parameters and the one or more user-specific event weights to the trained ML model;
>>
>> generating, via the trained ML model, a schedule for the future series of live events that is optimized relative to the one or more prioritized event target features;
>>
>> detecting a real-time change to the one or more user-specific event parameters;
>>
>> providing the real-time change to the trained ML model to improve the accuracy of the trained ML model; and
>>
>> updating, via the trained ML model, the schedule for the future series of live events such that the schedule remains optimized relative to the one or more prioritized event target features in view of the real-time change to the one or more user-specific event parameters.

Id. at *2–3.

The plaintiff argued that the AI-based optimization technique described in the patent provided a "better and more optimized event schedule than what a human could achieve without the claimed techniques," *id.* at *6, including because "machine learning can identify patterns and details imperceptible to humans, and thereby optimizes . . . in a different way than the human brain would or could," *id.* at *9. The court rejected this as a basis for patentability. According to the court, "mathematical algorithms, without more, [are] essentially mental processes within the abstract-idea category," and "[b]ecause machine learning is algorithmic in nature," the claims were abstract under § 101. As the court noted, the claims "do not require that the machine learning process be complex," instead referring to standard "regress" and "decision trees" as forms of machine learning. *Id.*

The plaintiff also argued that the patents were distinct from ineligible mental processes because "the claimed processes require too much data and computing power for the human brain to do," and were therefore distinct from mental processes. *Id.* The court rejected this argument as well; "humans can engage in the mathematical techniques to perform machine learning (albeit slowly)." *Id.* Ultimately, the court concluded that merely referencing machine learning generically or noting speed enhancements generally was not enough to render claims patentable; "requir[ing] generic machine learning or operations that a human could not perform as a computer using machine learning" was simply a faster way to carry out mathematical calculations that were themselves unpatentable abstract ideas. *Id.* at *10. At step 2, the court explained that the claim was not an "unconventional technological solution . . . to a technological problem," and instead was merely the implementation of the abstract idea on a computer. *Id.* at *13.

While the court in *Recentive* ultimately concluded the claims were not patent eligible, the court did note that patent claims drawn to improved neural networks could be patent eligible. *Id.* at *6. As the court noted, "novel methods to improve prior art neural networks" could be patent eligible—for example, "an expanded training set using mathematical transformations and the minimization of false positives using a distinctive training method." *Id.* But the patent in *Recentive* did not compare favorably to such a patent on improved neural networks; the *Recentive* patent "applied generic machine learning techniques to a pre-existing process" rather than "improving a prior art machine learning technique." *Id.* As the court noted, the patent itself stated that "any suitable machine learning technique can be used," rather than reciting a specific and novel improvement on machine learning techniques. *Id.*

***Angel Techs. Grp. LLC v. Facebook Inc.*, No. CV 21-8459-CBM, 2022 WL 3093232 (C.D. Cal. June 30, 2022)**

In *Angel Technology Group*, at issue was a patent claiming the method of "tagging" individuals in photos, storing that information in a database, and using "artificial intelligence algorithms" to locate tagged individuals in additional photos. *Angel Techs. Grp. LLC v. Facebook Inc.*, No. CV 21-8459-CBM, 2022 WL 3093232, at *4 (C.D. Cal. June 30, 2022). One representative claim stated:

> In a multi-user computer network, a method for obtaining and displaying information relating to existence of at least one user of a computer network in an image comprising:
>
> > identifying users of said computer network and assigning a unique user identification to a user of said computer network;
> >
> > storing said unique user identifier in an users database wherein said database is accessible by other computers of said computer network;

> obtaining image data from at least one uploading user of said computer network and assigning a unique image identifier to said image data;
>
> storing said unique image identifier in an images database wherein said database is accessible by other computers of said computer network;
>
> obtaining identification data from a first tagging user of said computer network, wherein said identification data comprises said unique image identifier and a pictured user unique identifier of a user of said computer network pictured in said image data;
>
> storing said identification data from said first tagging user in an identifications database accessible by other computers of said network whereby a user identifier may be associated with one or more image identifiers and an image identifier may be associated with one or more users identifiers.

The district court determined that "identifying people in photos and recording that information is a commonplace human concept," and that the use of artificial intelligence algorithms did not make the method patent-eligible because the claim "provide[d] no details on how to implement these algorithms, how they [would] work, or how Plaintiff improved upon it." *Id.* And it explained that "all this claim discloses about artificial intelligence is that it will be used which is the very essence of an abstract idea." *Id.*

***PUREPREDICTIVE, Inc. v. H20.AI, Inc.*, No. 17-CV-03049-WHO, 2017 WL 3721480 (N.D. Cal. Aug. 29, 2017)**

In *Purepredictive*, the claim at issue described a "predictive analytics factory"—essentially, a method for automatically generating application-specific data analytics models without manual input or customization from a user or expert. 2017 WL 3721480 (N.D. Cal. Aug. 29, 2017). The claim described a three-step method: (1) randomly generate different predictive models; (2) evaluate the effectiveness of each model; and (3) apply different subsets of the prediction models for different subsets of the data based on the evaluations. The full representative claim recited:

> A method for a predictive analysis factory, the method comprising:
>
> pseudo-randomly generating a plurality of learned functions based on training data without prior knowledge regarding suitability of the generated learned functions for the training data, the training data received for forming a predictive ensemble customized for the training data;
>
> evaluating the plurality of learned functions using test data to generate evaluation metadata indicating an effectiveness of different learned functions at making predictions based on different subsets of test data; and
>
> forming the predictive ensemble comprising a subset of multiple learned functions from the plurality of learned functions, the subset of multiple learned functions selected and combined based on the evaluation metadata, the predictive ensemble comprising a rule set synthesized from the evaluation metadata to direct different subsets of the workload data through different learned functions of the multiple learned functions based on the evaluation metadata.

The plaintiff argued that the claim was patent eligible because it solved a specific problem in predictive analytics and made improvements to a computer-related technology. The court determined that because predictive analytics involved "the basic mathematical process of . . . regression modeling, or running data through an algorithm," the method claimed by the patent was "directed to the patent-ineligible abstract concept of testing and refining mathematical algorithms," which only "invoke[d] computers as a tool for this process." *Id.* at *5. The court also rejected the plaintiff's assertion that the claim was an improvement to computer-related technology, noting that the plaintiff "fail[ed] to identify any previously existing technology that its claims improve[d] upon." *Id.* At step 2, the court rejected the plaintiff's argument that the claim was "an unconventional solution in its field" and "an improvement to existing technology." *Id.* at *6. The court explained that the claims were not specifically tied to a problem, and even if the method was effective, it was "simply an implementation of the basic concept of predictive analytics on an apparatus, computer program product, or other medium." *Id.* at *7.

USPTO Guidance

The United States Patent and Trademark Office has provided some guidance as to the types of AI and machine learning claims that would be patent eligible.

Example 39 from the USPTO subject matter eligibility guidelines considers the following claim:

> A computer-implemented method of training a neural network for facial detection comprising:
>
> collecting a set of digital facial images from a database;
>
> applying one or more transformations to each digital facial image including mirroring, rotating, smoothing, or contrast reduction to create a modified set of digital facial images;
>
> creating a first training set comprising the collected set of digital facial images, the modified set of digital facial images, and a set of digital non-facial images;
>
> training the neural network in a first stage using the first training set;
>
> creating a second training set for a second stage of training comprising the first training set and digital non-facial images that are incorrectly detected as facial images after the first stage of training;
>
> and training the neural network in a second stage using the second training set.

Manual of Patent Examining Procedure (MPEP) 2106.04(a)(1)(vii). The PTO Guidance explains the claim should survive *Alice* step 1 because although it is "based on mathematical concepts, the mathematical concepts are not recited ***in*** the claims," and it "does not recite a mental process because the steps are not practically performed in the human mind." *Id.* As referenced above, the *Recentive* court analyzed Example 39 and took no issue with its result, acknowledging that claims directed to improving the underlying neural network used for a machine learning application may be patent eligible. Companies may therefore maximize their chances of patenting innovative AI software by hewing more closely to the claim in Example 39.

This approach to securing patent protection for AI was successful in *Ex Parte Hannun*, No. 2018-003323, 2019 WL 7407450 (P.T.A.B. Apr. 1, 2019). There the PTAB reversed the Patent Examiner's rejection of a patent claim that described using deep learning and neural networks for speech recognition. The claim at issue stated:

> A computer-implemented method for transcribing speech comprising:
>
> receiving an input audio from a user;
>
> normalizing the input audio to make a total power of the input audio consistent with a set of training samples used to train a trained neural network model;
>
> generating a jitter set of audio files from the normalized input audio by translating the normalized input audio by one or more time values;

> for each audio file from the jitter set of audio files, which includes the normalized input audio:
>
> generating a set of spectrogram frames for each audio file;
>
> inputting the audio file along with a context of spectrogram frames into a trained neural network;
>
> obtaining predicted character probabilities outputs from the trained neural network; and
>
> decoding a transcription of the input audio using the predicted character probabilities outputs from the trained neural network constrained by a language model that interprets a string of characters from the predicted character probabilities outputs as a word or words.

Id. at *1. The Examiner had rejected this patent under § 101, finding that the steps of the invention could be performed mentally, and that the claim was "directed to using the predicted character probabilities (mathematical formula) to decode a transcription of the input audio into words or text data." *Id.* at *3. The panel disagreed on both counts. It explained that the steps of normalizing input, creating transformed training data, and training a neural network could not be performed mentally. And it found that although the claim used a mathematical formula, the claim did not ***recite*** a mathematical formula. Instead, the features of the technology were integrated into a speech recognition system, which was a practical application. The panel therefore concluded that patent was directed towards practical improvements in speech recognition and was patentable.

III. Drafting Patents Covering AI Inventions: What to Avoid and What to Emphasize

The recent decisions discussed above provide helpful guidance for companies seeking patent protection for AI. Several practical approaches companies can employ based on these decisions include the following:

> Be as specific as possible about the AI improvement at issue. Avoid stating that the innovation uses any "general machine learning algorithm" or "any suitable machine learning technique." References to generic AI models or techniques, or simple regression or other statistical methods, likely will draw comparisons to cases where courts viewed AI patents as improper attempts to patent math, or to improperly preempt the use of AI generically to perform a particular task or application. Concrete and specific enumerated improvements to prior art machine learning techniques may be more favorably received.
>
> Leverage the specification to describe how the claimed innovation improves the functionality of the computer itself or solves a real-world problem. For example, explain how and why a computer-implemented AI model enables a computer to process data more efficiently than prior computer-implemented machine-learning implementations, or how and why combining the computer-implemented AI model with your hardware solves a well-known problem.
>
> Avoid describing the innovation as analyzing and displaying data, or otherwise using computers to manipulate data. Where the AI process at issue simply replicates steps that humans can and do perform manually (even if much more slowly) to accomplish a similar task, courts are typically hostile to patent eligibility.
>
> Do not make the underlying mathematical concepts the focus of the claims. Attempts to claim specific mathematical techniques or statistical methods have drawn fire from courts considering patent eligibility.

Using these approaches, companies may improve their chances of securing patents on innovations in AI. There are no guarantees—but by avoiding the pitfalls described above, AI innovators will increase their odds for meaningful patent protection.

IV. To Patent or Not To Patent: Trade Secrets as an Alternative

Ultimately, companies are not forced to decide between pursuing patents or foregoing intellectual property protections entirely. Protecting AI innovations as trade secrets can be an important alternative to seeking patents, and in many cases may be the only means of doing so. Trade secret protections in lieu of patents may be particularly appropriate where companies do not want to disclose valuable "crown jewel" techniques in patent applications, or where innovations are valuable but seem unlikely to qualify for patent eligibility.

As discussed above, patent protection is not available for abstract ideas, algorithms or other elements of AI innovation. By comparison, the subject matter eligibility for trade secrets is broad and can encompass many different forms of information. Under the Uniform Trade Secrets Act, trade secrets can be "information, including a formula, pattern, compilation, program, device, method, technique, or process." Under the Defend Trade Secrets Act ("DTSA"), trade secrets can be "all forms and types of financial, business, scientific, technical, economic, or engineering information, including patterns, plans, compilations, program devices, formulas, designs, prototypes, methods, techniques, processes, procedures, programs, or codes, whether tangible or intangible, and whether or how stored, compiled, or memorialized physically, electronically, graphically, photographically, or in writing." Thus, even abstract ideas may be subject to trade secret protection in the United States.

Trade secret protections can be substantial. The DTSA protects any information that "derives independent economic value . . . from not being generally known to, and not being readily ascertainable through proper means by, another person who can obtain economic value from . . . the information." 18 U.SC. § 1839(3)(B). Remedies include preliminary or permanent injunctions to prevent improper use of trade secrets, as well as damages (including enhanced damages and fees) for misappropriation.

Whether it is preferable to protect IP through patents or as trade secrets—where there is a clear choice—is highly dependent on the specific circumstances at issue. The following are some of the issues and factors to consider:

Unlike patents, trade secrets are not subject to any statutory subject-matter restrictions or enablement requirements. Even abstract ideas may be trade secrets. But in litigation trade secrets do need to be identified with particularity, typically at the outset of a case. An asserted trade secret must be identified with enough specificity to put a defendant on notice of the protected information, and "to separate it from matters of general knowledge in the trade or of special knowledge of those persons skilled in the trade." *InteliClear, LLC v. ETC Glob. Holdings, Inc.*, 978 F.3d 653, 658 (9th Cir. 2020). Where it is difficult to articulate a trade secret with particularity or precision, it will also be difficult to assert that trade secret.

Currently, U.S. patent law provides protection for the inventions of natural persons only. This means that, at present, discoveries and inventions generated by artificial intelligence are not eligible for patent protection in the United States or in most other jurisdictions in the world.[2] There are, however, no

[2] *See, e.g.,* USPTO decision of 22 April 2020 on Application No. 16/524,350, https://www.uspto.gov/sites/default/files/documents/16524350_22apr2020.pdf;
EPO Board of Appeal decision of 21 December 2021 on EP 18 275 163 and EP 18 275 174, https://www.epo.org/law-practice/case-law-appeals/communications/2021/20211221.html;

restrictions on AI-based inventions under trade secret law. As long as AI-based discoveries are subject to the standard requirements for trade secret protection (*i.e.*, are not generally known or readily ascertainable, have value, and are subject to reasonable efforts to maintain secrecy), they can qualify for protection under U.S. trade secret laws.

Patent protection generally lasts no longer than 20 years. Trade secret protection can last much longer and potentially be unlimited. As long as information is secret and is valuable because it is secret, it can be protectable as a trade secret. But while trade secret protection can be unlimited, it can also be fleeting. Trade secret protection is not guaranteed to last for any period of time. Once a trade secret is generally known in the field (by independent discovery or otherwise), it is no longer protectable.

Unlike patents, trade secret protection may be forfeited if the information is disclosed or "reasonable efforts" were not used to maintain secrecy. Thus, it is important that any disclosure of the trade secrets (including to customers and partners/vendors) is made under appropriate NDAs or other measures to assure confidentiality.

Trade secret protection can have a vast geographic scope. Patents generally provide little protection for foreign conduct. But trade secret claims can be asserted against foreign entities for foreign conduct, as long as a single act in furtherance of the misappropriation occurred in U.S. *See* 18 U.S.C. § 1897(2).

An accused infringer can use knowledge gained from the patent to design around the patent and build non-infringing equivalent technology. Redesign can be much harder for trade secret disputes. Any knowledge gained from misappropriated trade secrets is arguably fruit of the poisonous tree. An accused misappropriator may only be able to "design around" a trade secret by establishing an independent "clean room" development effort or through well-documented reverse engineering.

Unlike patents, trade secret protection cannot protect a later company from developing the same innovations independently and then using those same innovations against companies that first developed those same innovations.

There are many other differences between patent and trade secret protection, both in the United States and abroad. Companies should carefully consider the advantages and disadvantages of each option when selecting the appropriate strategy to protect their AI innovations. Often, the most effective strategy will be to protect some innovations through patents and others as trade secrets. Choosing which innovations will benefit from which form of protection, and carefully crafting that protection, is the key.

In addition to evaluating each innovation to assess whether it should be protected though patents or as trade secrets, companies should consider reviewing existing and pending patent claims for any eligibility issues under *Alice* to strengthen them for potential litigation, and should consider an independent review of their trade secret policies, NDA policies, and employee onboarding-exit process to enhance their trade secret protection.

***

Commissioner of Patents v Thaler, [2022] FCAFC 62 (13 April 2022, Federal Court of Australia), https://www.austlii.edu.au/cgi-bin/viewdoc/au/cases/cth/FCAFC/2022/62.html;
UK IPO patent decision BL O/741/19 of 4 December 2019, https://www.ipo.gov.uk/p-challenge-decision-results/p-challenge-decision-results-bl?BL_Number=O/741/19 (appeal dismissed by UK Court of Appeal (Civil), [2021] EWCA Civ 1374, https://www.bailii.org/ew/cases/EWCA/Civ/2021/1374.pdf).

If you have any questions about the issues addressed in this memorandum, or if you would like a copy of any of the materials mentioned in it, please do not hesitate to contact:

John Quinn
Email: johnquinn@quinnemanuel.com
Phone: 213-443-3200

Sean Pak
Email: seanpak@quinnemanuel.com
Phone: 415-875-6320

Patrick Curran
Email: patrickcurran@quinnemanuel.com
Phone: 617-712-7103

Ron Hagiz
Email: ronhagiz@quinnemanuel.com
Phone: 212-849-7453

Ryan Sullivan
Email: ryansullivan@quinnemanuel.com
Phone: 617-712-7164

To view more memoranda, please visit www.quinnemanuel.com/the-firm/publications/.

To update information or unsubscribe, please email updates@quinnemanuel.com.