

July 2024 Subject Matter Eligibility Examples

The following examples should be used in conjunction with USPTO guidance on subject matter eligibility, which is incorporated into the Manual of Patent Examining Procedure (MPEP) 2106 and discussed in the [2024 Guidance Update on Patent Subject Matter Eligibility, Including on Artificial Intelligence](#). The examples below are hypothetical and only intended to be illustrative of the claim analysis performed using MPEP 2106, and of the particular issues noted below in the Issue Spotting Chart. These examples should be interpreted based on the fact patterns set forth below, as other fact patterns may have different eligibility outcomes. That is, it is not necessary for a claim under examination to mirror an example claim to be subject matter eligible. All claims are analyzed for eligibility in accordance with their broadest reasonable interpretation.

Note that the examples provided here are numbered consecutively beginning with number 47, because 46 examples were previously issued. Appendix 1 to these examples contains a comprehensive index of all 49 of the USPTO's eligibility examples.

The examples are only illustrative of the patent subject matter eligibility analysis. All claims must be ultimately analyzed for compliance with every requirement for patentability. The analyses provided below do not address considerations other than subject matter eligibility under Section 101.

Issue Spotting Chart	Anomaly Detection	Speech Separation	Fibrosis Treatment
Example Number	47	48	49
Claim Type			
Process	•	•	•
Product (Composition of Matter, Manufacture, and/or Machine)	•	•	
Judicial Exception			
Abstract Idea: Mathematical Concept	•	•	•
Abstract Idea: Mental Process	•	•	•
Abstract Idea: Certain Methods of Organizing Human Activity			
Law of Nature			•
Product of Nature			
Multiple exceptions in same claim	•	•	•
No recited exception	•		
Detailed Analysis			
Step 2A Prong One: Generally	•	•	•
Step 2A Prong One: Markedly Different Characteristics analysis			
Step 2A Prong Two: Exception Integrated into a Practical Application	•	•	•
Step 2B: Generally	•	•	•
Step 2B: Claim is eligible because it provides an Inventive Concept			
Considerations Discussed in Step 2A Prong Two and/or Step 2B			
Improvements to Functioning of a Computer or Other Technology	•	•	•
Particular Treatment or Prophylaxis (Prong Two only)			•
Particular Machine			
Particular Transformation			
Other Meaningful Limitations			
Mere Instructions to Apply an Exception	•	•	•
Insignificant Extra-Solution Activity	•	•	•
Field of Use and Technological Environment	•	•	•
Well-Understood, Routine, Conventional (WURC) Activity (Step 2B only)	•	•	•

Example 47. Anomaly Detection

This example illustrates the application of the eligibility analysis to claims that recite limitations specific to artificial intelligence, particularly the use of an artificial neural network to identify or detect anomalies. Claim 1 is eligible because it falls within a statutory category and does not recite any judicial exceptions. Claim 2 is ineligible because it recites a judicial exception (abstract idea), and the claim as a whole does not integrate the exception into a practical application (and is thus directed to an abstract idea), and the claim does not provide significantly more than the exception (does not provide an inventive concept). Claim 3 is eligible because it recites a judicial exception (abstract idea), but the claim as a whole integrates the judicial exception into a practical application by improving network security.

BACKGROUND

The present invention seeks to use an artificial neural network (ANN) to identify or detect anomalies. The use of specially trained ANNs to detect anomalies realizes a number of improvements over traditional methods of detecting anomalies, including more accurate detection of anomalies. The application further provides methods for training an ANN that lead to faster training times and a more accurate model for detecting anomalies.

ANNs are a type of machine learning model used to perform a wide variety of complex tasks, including image recognition, speech recognition, pattern recognition, and detection of anomalies. An ANN is a biologically inspired algorithm that learns from training data. An ANN can be realized through software, hardware, or a combination of software and hardware. The structure of an exemplary ANN has a series of layers, each comprising one or more neurons arranged in one or more neuron arrays. In an exemplary embodiment, a neuron may comprise a register, a microprocessor, and at least one input. Each neuron produces an output, or activation, based on an activation function that uses the outputs of the previous layer and a set of weights as inputs. Each neuron in a neuron array may be connected to another neuron via a synaptic circuit. A synaptic circuit may include a memory for storing a synaptic weight. An exemplary ANN may be a Deep Neural Network having an input layer, an output layer, and a plurality of fully connected hidden layers. ANNs are particularly useful in anomaly detection because they can effectively extract features in linear and nonlinear relationships. In some embodiments, an ANN may be implemented by an application-specific integrated circuit (ASIC). ASICs may be specially customized for a specific artificial intelligence application and provide superior computing capabilities and reduced electricity consumption compared to traditional CPUs.

In some embodiments, training data is generated by receiving continuous data at a computer and using the computer to discretize the continuous data. In some embodiments, the continuous data may be received remotely over a network. The continuous data may be historical data, which the neural network can use to learn patterns to identify or detect potential anomalies. Continuous data is data that is measured and can have any number of possible values. Machine learning models may benefit from being trained with discrete data rather than continuous data. Discrete data can be counted and has a limited number of values. Any type of discretization method may be used to convert continuous data to discrete data, including binning, clustering, and numerical discretization. The ANN is then trained using any known training techniques to generate a

trained neural network which can be used to detect anomalies. The trained ANN monitors incoming data sets to detect anomalies. If the trained ANN detects one or more anomalies, it additionally analyzes the detected anomalies to generate anomaly data which can be output to a user and/or used to re-train the ANN. For example, the anomaly data may explain the type of anomaly or a cause of the anomaly.

A conventional backpropagation algorithm and a conventional gradient descent algorithm may be used to train the neural network. Gradient descent is an optimization algorithm used to minimize differentiable real-valued multivariate functions. Gradient descent begins by initializing the values of parameters and then applying a gradient descent calculation, which uses mathematical calculations to iteratively adjust the values so they minimize a loss function to optimize the ANN. Backpropagation is the mathematical process of calculating the derivatives and gradient descent is the process of adjusting model parameters using the calculated derivatives to minimize the loss function. Backpropagation is a mathematical calculation for supervised learning of ANNs using gradient descent. Given an ANN and an error function, backpropagation is used to calculate the gradient of the error function with respect to the neural network's weights.

Anomaly detection is an important task that impacts any industry that benefits from identifying abnormal data that deviates from expected data or from a general pattern. For example, an intrusion detection system may use the disclosed anomaly detection method to improve detection of malicious network packets. A difficulty in anomaly detection is that a system must define the boundary between ordinary and anomalous data and accurately classify data as ordinary or anomalous. The line between ordinary and anomalous data may be difficult to determine with cases approaching a boundary and based on an application-specific domain. For example, small variations may trigger an identification of an anomaly in network security or medicine while relatively larger deviations may be considered normal in less sensitive applications. Furthermore, malicious actors may attempt to make anomalies appear like ordinary activity. This application provides solutions for using a trained ANN to quickly and accurately identify anomalies as compared to anomaly detection performed using traditional methods.

In some embodiments, the ANN may detect anomalies in a network where the anomalies indicate potential network intrusions or malicious attacks. If the ANN detects one or more anomalies in network traffic, the ANN can additionally determine whether the detected anomaly is associated with a malicious packet. If the detected anomaly is associated with a malicious packet, the ANN may cause a network device to drop the malicious packet and block future traffic from the sender of the malicious packet. By automatically detecting network intrusions or other malicious attacks, the present invention enhances network security by allowing for automatic, proactive remediation of network attacks. In some embodiments, the system may use various detection techniques for detecting potentially malicious network packets and the source of the potentially malicious network packet and can alert a network administrator to potential problems. The system may detect the source of a potentially malicious network packet through a tracing operation or the use of a software tool. The disclosed system detects network intrusions and takes remedial actions, including automatically dropping suspicious packets and blocking traffic from suspicious source addresses without the need for alerting a network administrator. Unlike conventional network remediation solutions, the disclosed method and system are able to identify malicious network packets and take remediation actions, including dropping suspicious

packets and blocking traffic from suspicious source addresses in real time. The disclosed system realizes an improvement in network security by avoiding the delay involved in waiting on a network administrator to react to a network intrusion by automatically dropping suspicious packets and blocking traffic from suspicious source addresses based on anomalies identified by the ANN in real time.

CLAIMS

[Claim 1] An application specific integrated circuit (ASIC) for an artificial neural network (ANN), the ASIC comprising:

a plurality of neurons organized in an array, wherein each neuron comprises a register, a microprocessor, and at least one input; and

a plurality of synaptic circuits, each synaptic circuit including a memory for storing a synaptic weight, wherein each neuron is connected to at least one other neuron via one of the plurality of synaptic circuits.

[Claim 2] A method of using an artificial neural network (ANN) comprising:

- (a) receiving, at a computer, continuous training data;
- (b) discretizing, by the computer, the continuous training data to generate input data;
- (c) training, by the computer, the ANN based on the input data and a selected training algorithm to generate a trained ANN, wherein the selected training algorithm includes a backpropagation algorithm and a gradient descent algorithm;
- (d) detecting one or more anomalies in a data set using the trained ANN;
- (e) analyzing the one or more detected anomalies using the trained ANN to generate anomaly data; and
- (f) outputting the anomaly data from the trained ANN.

[Claim 3] A method of using an artificial neural network (ANN) to detect malicious network packets comprising:

- (a) training, by a computer, the ANN based on input data and a selected training algorithm to generate a trained ANN, wherein the selected training algorithm includes a backpropagation algorithm and a gradient descent algorithm;
- (b) detecting one or more anomalies in network traffic using the trained ANN;
- (c) determining at least one detected anomaly is associated with one or more malicious network packets;

- (d) detecting a source address associated with the one or more malicious network packets in real time;
- (e) dropping the one or more malicious network packets in real time; and
- (f) blocking future traffic from the source address.

ANALYSIS

Claim 1 is eligible.

Claim Interpretation: Under the broadest reasonable interpretation, the terms of the claim are presumed to have their plain meaning consistent with the specification as it would be interpreted by one of ordinary skill in the art. See Manual of Patent Examining Procedure (MPEP) 2111. The claim recites an application specific integrated circuit (ASIC) used for an artificial neural network (ANN). While the background explains that “[a]n ANN can be realized through software, hardware, or a combination of software and hardware,” the broadest reasonable interpretation of the claimed ANN requires hardware because the claimed ASIC is a physical circuit.

Step 1: This part of the eligibility analysis evaluates whether the claim falls within any statutory category. See MPEP 2106.03. The claim recites an ASIC that implements an ANN. The claim is directed to a physical circuit, which is a machine and/or manufacture, and falls within one of the statutory categories of invention. (*Step 1: YES*).

Step 2A, Prong One: This part of the eligibility analysis evaluates whether the claim recites a judicial exception. As explained in MPEP 2106.04, subsection II, a claim “recites” a judicial exception when the judicial exception is “set forth” or “described” in the claim. There is no judicial exception recited in the claim. The claim recites a plurality of neurons, which are hardware components comprising a register and a microprocessor, and a plurality of synaptic circuits which together form an ANN. The claim does not recite any abstract ideas, such as a mathematical concept, mental process, or a method of organizing human activity, such as a fundamental economic concept or managing interactions between people. See MPEP 2106.04(a)(2). While ANNs may be trained using mathematics, there is no mathematical concept recited in the claim. Because the claim does not recite a judicial exception (*Step 2A, Prong One: NO*), it cannot be directed to one (*Step 2A: NO*). The claim is eligible.

Claim 2 is ineligible.

Claim Interpretation: Under the broadest reasonable interpretation, the terms of the claim are presumed to have their plain meaning consistent with the specification as it would be interpreted by one of ordinary skill in the art. See MPEP 2111.

Steps (a) and (b) recite receiving and discretizing continuous training data to generate input data. The term “continuous data” is recognized as having its plain meaning of any data that is measured and can take on any number of possible values. The plain meaning of discrete data, as

supported by the third paragraph of the background, is data that can be counted, has a limited number of values, and is more suitable for use as training data.

The claim does not put any limits on how the continuous data is received, but the background supports the plain meaning of “receiving” as encompassing receiving the data remotely over a network. The claim also does not limit the plain meaning of “discretizing,” which, as explained in the background, includes any known discretization method, including binning and clustering, as well as numerical discretization, such as rounding continuous data values or performing other basic mathematical calculations that can be performed mentally (see the third paragraph of the background). Step (c) recites training an ANN using a selected algorithm. The training algorithm is a backpropagation algorithm and a gradient descent algorithm. When given their broadest reasonable interpretation in light of the background, the backpropagation algorithm and gradient descent algorithm are mathematical calculations. The plain meaning of these terms are optimization algorithms, which compute neural network parameters using a series of mathematical calculations. The fourth paragraph of the background supports the plain meaning by stating the “gradient descent begins by initializing the values of parameters and then applying a gradient descent calculation, which uses mathematical calculations to iteratively adjust the values so they minimize a loss function.” The background also states that “backpropagation is a mathematical calculation for supervised learning of ANNs using gradient descent.”

Steps (a), (b), and (c) are all recited as being performed by a computer. The recited computer is recited at a high level of generality, i.e., as a generic computer performing generic computer functions.

Step (d) recites detecting one or more anomalies in a data set using the trained ANN. The claim does not provide any details about how the trained ANN operates or how the detection is made, and the plain meaning of “detecting” encompasses mental observations or evaluations, e.g., a computer programmer’s mental identification of an anomaly in a data set.

Step (e) recites analyzing the one or more detected anomalies using the trained artificial neural network to generate anomaly data. The step of analyzing includes both determining that an anomaly has been detected and may further include suggesting a type or cause of the anomaly. The plain meaning of “analyzing” encompasses evaluating information, which in this claim is limited to evaluating detected anomalies to generate anomaly data by the trained ANN. The claim does not limit how the analysis (evaluation) is performed, and there is nothing about a detected anomaly itself that would limit how it can be analyzed. As explained in the background, “the anomaly data may explain the type of anomaly or a cause of the anomaly.” The claim does not include any additional details that explain the analysis of detected anomalies.

Regarding step (f), the step of outputting the anomaly data merely requires a generic output using the trained ANN. The claim does not impose any limits on how the data is output or require any particular components that are used to output the anomaly data.

Based on the plain meaning of the words in the claim, the broadest reasonable interpretation of claim 2 is a method that receives continuous training data at a computer, uses the computer to discretize the continuous training data to generate input data, trains the ANN using the input data

and a selected backpropagation algorithm and gradient descent algorithm, detects and analyzes anomalies in a data set using the trained ANN, and outputs anomaly data from the trained ANN. The claimed discretizing, detecting, and analyzing steps encompass mental choices or evaluations, and the claimed discretizing and training using a backpropagation algorithm and gradient descent algorithm encompasses performing mathematical calculations.

Step 1: This part of the eligibility analysis evaluates whether the claim falls within any statutory category. See MPEP 2106.03. The claim recites at least one step or act, including receiving continuous training data. Thus, the claim is to a process, which is one of the statutory categories of invention. (*Step 1: YES*).

Step 2A, Prong One: This part of the eligibility analysis evaluates whether the claim recites a judicial exception. As explained in MPEP 2106.04, subsection II, a claim “recites” a judicial exception when the judicial exception is “set forth” or “described” in the claim.

As discussed above, the broadest reasonable interpretation of steps (b), (d), and (e) is that those steps fall within the mental process groupings of abstract ideas because they cover concepts performed in the human mind, including observation, evaluation, judgment, and opinion. See MPEP 2106.04(a)(2), subsection III.

Specifically, step (b) recites discretizing continuous training data to generate input data by processes including rounding, binning, or clustering continuous data, which may be practically performed in the human mind using observation, evaluation, judgment, and opinion. For example, the claimed discretizing of continuous data encompasses observing continuous data and performing an evaluation, such as rounding the continuous data. Step (d) recites detecting one or more anomalies in a data set using the trained ANN. Under its broadest reasonable interpretation when read in light of the specification, the “detecting” encompasses mental observations or evaluations that are practically performed in the human mind. For example, the claimed detecting of anomalies in a data set encompasses observing data in a data set and performing an evaluation by comparing anomalous and non-anomalous data. Step (e) recites analyzing the one or more detected anomalies using the trained ANN to generate anomaly data. Step (e) encompasses performing evaluation, judgment, and opinion to make a determination about detected anomalies. Under its broadest reasonable interpretation when read in light of the specification, the “analyzing” encompasses mental processes practically performed in the human mind by observation, evaluation, judgment, and opinion. See MPEP 2106.04(a)(2), subsection III.

As discussed above, the broadest reasonable interpretation of discretizing in step (b) also encompasses mathematical concepts (e.g., rounding data values) that can be performed mentally. Step (c) requires specific mathematical calculations (a backpropagation algorithm and a gradient descent algorithm) to perform the training of the ANN and therefore encompasses mathematical concepts.

“Unless it is clear that a claim recites distinct exceptions, such as a law of nature and an abstract idea, care should be taken not to parse the claim into multiple exceptions, particularly in claims involving abstract ideas.” MPEP 2106.04, subsection II.B. However, if possible, the examiner

should consider the limitations together as a single abstract idea rather than as a plurality of separate abstract ideas to be analyzed individually. “For example, in a claim that includes a series of steps that recite mental steps as well as a mathematical calculation, an examiner should identify the claim as reciting both a mental process and a mathematical concept for Step 2A, Prong One to make the analysis clear on the record.” MPEP 2106.04, subsection II.B. Under such circumstances, however, the Supreme Court has treated such claims in the same manner as claims reciting a single judicial exception. *Id.* (discussing *Bilski v. Kappos*, 561 U.S. 593 (2010)). Here, steps (b), (d), and (e) fall within the mental process grouping of abstract ideas, and steps (b) and (c) fall within the mathematical concepts grouping of abstract ideas. Limitations (b)-(e) are considered together as a single abstract idea for further analysis. (*Step 2A, Prong One: YES*).

Step 2A, Prong Two: This part of the eligibility analysis evaluates whether the claim as a whole integrates the recited judicial exception into a practical application of the exception or whether the claim is “directed to” the judicial exception. This evaluation is performed by (1) identifying whether there are any additional elements recited in the claim beyond the judicial exception, and (2) evaluating those additional elements individually and in combination to determine whether the claim as a whole integrates the exception into a practical application. See MPEP 2106.04(d). The claim recites the additional elements of “(a) receiving, at a computer, continuous training data,” “using the trained ANN” in limitations (d) and (e), and “(f) outputting the anomaly data from the trained ANN.” The claim also recites that steps (b) and (c) are performed by a computer.

The limitations “(a) receiving, at a computer, continuous training data” and “(f) outputting the anomaly data from the trained ANN” are mere data gathering and output recited at a high level of generality, and thus are insignificant extra-solution activity. See MPEP 2106.05(g) (“whether the limitation is significant”). In addition, all uses of the recited judicial exceptions require such data gathering and output, and, as such, these limitations do not impose any meaningful limits on the claim. These limitations amount to necessary data gathering and outputting. See MPEP 2106.05.

Further, limitations (a), (b), and (c) are recited as being performed by a computer. The computer is recited at a high level of generality. In limitation (a), the computer is used as a tool to perform the generic computer function of receiving data. See MPEP 2106.05(f). In limitations (b) and (c), the computer is used to perform an abstract idea, as discussed above in Step 2A, Prong One, such that it amounts to no more than mere instructions to apply the exception using a generic computer. See MPEP 2106.05(f).

The limitations in (d) and (e) reciting “using the trained ANN” provide nothing more than mere instructions to implement an abstract idea on a generic computer. See MPEP 2106.05(f). MPEP 2106.05(f) provides the following considerations for determining whether a claim simply recites a judicial exception with the words “apply it” (or an equivalent), such as mere instructions to implement an abstract idea on a computer: (1) whether the claim recites only the idea of a solution or outcome i.e., the claim fails to recite details of how a solution to a problem is accomplished; (2) whether the claim invokes computers or other machinery merely as a tool to

perform an existing process; and (3) the particularity or generality of the application of the judicial exception.

The judicial exception of “detecting one or more anomalies in a data set using the trained ANN” and “analyzing the one or more detected anomalies using the trained ANN to generate anomaly data” is performed “using the trained ANN.” The trained ANN is used to generally apply the abstract idea without placing any limits on how the trained ANN functions. Rather, these limitations only recite the outcome of “detecting one or more anomalies” and “analyzing the one or more detected anomalies” and do not include any details about how the “detecting” and “analyzing” are accomplished. See MPEP 2106.05(f).

The recitation of “using a trained ANN” in limitations (d) and (e) also merely indicates a field of use or technological environment in which the judicial exception is performed. Although the additional element “using a trained ANN” limits the identified judicial exceptions “detecting one or more anomalies in a data set using the trained ANN” and “analyzing the one or more detected anomalies using the trained ANN to generate anomaly data,” this type of limitation merely confines the use of the abstract idea to a particular technological environment (neural networks) and thus fails to add an inventive concept to the claims. See MPEP 2106.05(h).

Even when viewed in combination, these additional elements do not integrate the recited judicial exception into a practical application (*Step 2A, Prong Two: NO*), and the claim is directed to the judicial exception. (*Step 2A: YES*).

Step 2B: This part of the eligibility analysis evaluates whether the claim as a whole amounts to significantly more than the recited exception i.e., whether any additional element, or combination of additional elements, adds an inventive concept to the claim. See MPEP 2106.05.

As explained with respect to Step 2A, Prong Two, there are four additional elements. The additional element of “using the trained ANN” in limitations (d) and (e) are at best mere instructions to “apply” the abstract ideas, which cannot provide an inventive concept. See MPEP 2106.05(f).

Additional elements (a) and (f) were both found to be insignificant extra-solution activity in Step 2A, Prong Two, because they were determined to be insignificant limitations as necessary data gathering and outputting. However, a conclusion that an additional element is insignificant extra-solution activity in Step 2A, Prong Two should be re-evaluated in Step 2B. See MPEP 2106.05, subsection I.A. At Step 2B, the evaluation of the insignificant extra-solution activity consideration takes into account whether or not the extra-solution activity is well understood, routine, and conventional in the field. See MPEP 2106.05(g).

As discussed in Step 2A, Prong Two above, the recitations of “(a) receiving continuous training data” and “(g) outputting the anomaly data from the trained ANN” are recited at a high level of generality. These elements amount to receiving or transmitting data over a network and are well-understood, routine, conventional activity. See MPEP 2106.05(d), subsection II.

As discussed in Step 2A, Prong Two above, the recitation of a computer to perform limitations (a), (b), and (c) amounts to no more than mere instructions to apply the exception using a generic computer component.

Even when considered in combination, these additional elements represent mere instructions to implement an abstract idea or other exception on a computer and insignificant extra-solution activity, which do not provide an inventive concept. (*Step 2B: NO*).

Claim 3 is eligible.

Claim Interpretation: Under the broadest reasonable interpretation, the terms of the claim are presumed to have their plain meaning consistent with the specification as it would be interpreted by one of ordinary skill in the art. See MPEP 2111.

Step (a) recites “training . . . the ANN based on input data . . . to generate a trained ANN.” The claim recites performing the training using a backpropagation algorithm and a gradient descent algorithm. When given their broadest reasonable interpretation in light of the disclosure, the backpropagation algorithm and gradient descent algorithm are mathematical calculations. The plain meaning of these terms are optimization algorithms which compute neural network parameters using a series of mathematical calculations. The background supports this plain meaning in the fourth paragraph (“gradient descent begins by initializing the values of parameters and then applying a gradient descent calculation, which uses mathematical calculations to iteratively adjust the values so they minimize a loss function.”) Limitation (a) is also recited as being performed by a computer. The recited computer is recited at a high level of generality.

Step (b) recites “detecting one or more anomalies in network traffic using the trained ANN.” The claim does not provide any details about how the trained ANN operates or how the detection is made, and the plain meaning of “detecting” encompasses mental observations or evaluations (e.g., a computer programmer’s mental identification of an anomaly in a data set).

Step (c) recites “determining at least one detected anomaly is associated with one or more malicious network packets.” Under its broadest reasonable interpretation, this step only requires associating a detected anomaly with a malicious network packet. This step does not require the use of any specific process or component for associating a detected anomaly with one or more malicious network packets.

Step (d) further recites “detecting a source address associated with the one or more malicious network packets in real time.” The broadest reasonable interpretation of “detecting a source address associated with one or more malicious network packets in real time” is that the detection is a computer function that is described at a high level. Specifically, the sixth paragraph of the background notes that a computer may perform the detection by a tracing operation or by using a software tool.

Steps (e) and (f) further specify remedial actions that are executed to remediate or prevent network intrusions. The claimed steps of (e), automatically dropping the one or more malicious

network packets, and (f), blocking future traffic from the source address, provide specific computer solutions that use the output from the ANN to provide security solutions to the detected anomalies. As indicated in paragraph six of the background, the system may “automatically” perform dropping of malicious network packets and blocking future traffic without the need for any action by a network administrator. Instead, the ANN may make decisions about whether a network packet is potentially malicious and take action to drop malicious network packets and block future traffic.

Step 1: This part of the eligibility analysis evaluates whether the claim falls within any statutory category. See MPEP 2106.03. The claim recites a series of steps and, therefore, is a process. See MPEP 2106.03. (*Step 1: YES*).

Step 2A, Prong One: This part of the eligibility analysis evaluates whether the claim recites a judicial exception. As explained in MPEP 2106.04, subsection II, a claim “recites” a judicial exception when the judicial exception is “set forth” or “described” in the claim.

Step (a) provides use of specific mathematical calculations (a backpropagation algorithm and a gradient descent algorithm) to perform the training of the ANN and therefore encompasses mathematical concepts.

As discussed above, the broadest reasonable interpretation of steps (b) and (c) is that they fall within the mental process groupings of abstract ideas because they cover concepts performed in the human mind, including observation, evaluation, judgment, and opinion. See MPEP 2106.04(a)(2), subsection III.

Specifically, “detecting one or more anomalies in network traffic” is a mental process because the claimed detection is a process that is practically performed in the human mind by a human observing network traffic data and using “evaluation, judgment, and opinion” to detect whether an anomaly has occurred. “[D]etermining that at least one detected anomaly is associated with one or more malicious network packets” only requires associating a detected anomaly with a malicious network packet as stated above. The claim recites limitations that fall within the mental process grouping of abstract ideas.

“Unless it is clear that a claim recites distinct exceptions, such as a law of nature and an abstract idea, care should be taken not to parse the claim into multiple exceptions, particularly in claims involving abstract ideas.” MPEP 2106.04, subsection II.B (discussing *Bilski v. Kappos*, 561 U.S. 593 (2010)). Here, step (a) recites a mathematical concept, and steps (b) and (c) recite mental processes; therefore, claim 3 recites multiple abstract ideas. As discussed above for claim 2, in this case, it is appropriate to consider the limitations together as a single abstract idea rather than as a plurality of separate abstract ideas to be analyzed individually. Limitations (d)-(f) do not recite mental processes because they cannot be practically performed in the human mind. That is, the human mind is not equipped to detect a source address associated with malicious network packets, drop the malicious network packets in real time, and block future traffic as recited in the claim. See MPEP 2106.04(a)(2), subsection III.A (discussing *SRI Int’l, Inc. v. Cisco Systems, Inc.*, 930 F.3d 1295, 1303 (Fed. Cir. 2019)). As step (a) and steps (b)-(c) fall within different groupings of abstract ideas (i.e., mathematical concepts and mental processes, respectively),

these limitations are considered together as a single abstract idea for further analysis. (*Step 2A, Prong One: YES*).

Step 2A, Prong Two: This part of the eligibility analysis evaluates whether the claim as a whole integrates the recited judicial exception into a practical application of the exception. This evaluation is performed by (1) identifying whether there are any additional elements recited in the claim beyond the judicial exception, and (2) evaluating those additional elements individually and in combination to determine whether the claim as a whole integrates the exception into a practical application. See MPEP 2106.04(d).

One way to determine integration into a practical application is when the claimed invention improves the functioning of a computer or improves another technology or technical field. To evaluate an improvement to a computer or technical field, the specification must set forth an improvement in technology and the claim itself must reflect the disclosed improvement. See MPEP 2106.04(d)(1) and 2106.05(a).

The claim recites the additional elements of “(d) detecting a source address associated with the one or more malicious network packets,” “(e) dropping the one or more malicious network packets,” and “(f) blocking future traffic from the source address.” The claim also recites that limitation (a) is performed by a computer.

In limitation (a), the computer is used to perform an abstract idea, as discussed above in Step 2A, Prong One such that it amounts to no more than mere instructions to apply the exception using a generic computer. See MPEP 2106.05(f). In limitation (b), use of a trained ANN does not integrate the abstract idea of limitation (b) into a practical application for similar reasons as explained above in limitation (d) of Claim 2. Additionally, the recitation of “network traffic” generally links the abstract idea recited in limitation (b) to a particular field of use. See MPEP 2106.05(h).

The consideration of whether the claim as a whole includes an improvement to a computer or to a technological field requires an evaluation of the specification and the claim to ensure that a technical explanation of the asserted improvement is present in the specification, and that the claim reflects the asserted improvement. See MPEP 2106.04(d)(1). According to the background section, existing systems use various detection techniques for detecting potentially malicious network packets and can alert a network administrator to potential problems. The disclosed system detects network intrusions and takes real-time remedial actions, including dropping suspicious packets and blocking traffic from suspicious source addresses. The background section further explains that the disclosed system enhances security by acting in real time to proactively prevent network intrusions.

The claimed invention reflects this improvement in the technical field of network intrusion detection. Steps (d)-(f) provide for improved network security using the information from the detection to enhance security by taking proactive measures to remediate the danger by detecting the source address associated with the potentially malicious packets. Specifically, the claim reflects the improvement in step (d), dropping potentially malicious packets in step (e), and blocking future traffic from the source address in step (f). These steps reflect the improvement

described in the background. Thus, the claim as a whole integrates the judicial exception into a practical application such that the claim is not directed to the judicial exception.

The additional elements in steps (d)-(f), when considered in combination, integrate the abstract idea into a practical application because the claim improves the functioning of a computer or technical field. See MPEP 2106.04(d)(1) and 2106.05(a). The claimed invention reflects this improvement in the technical field of network intrusion detection. Thus, the claim as a whole integrates the judicial exception into a practical application (*Step 2A, Prong Two: YES*), such that the claim is not directed to the judicial exception. (*Step 2A: NO*). The claim is eligible.

Example 48. Speech Separation

This example illustrates the application of the eligibility analysis to claims that recite artificial intelligence-based methods of analyzing speech signals and separating desired speech from extraneous or background speech. The claims, including a deep neural network, are hypothetical and are loosely based on current speech separation literature. Claim 1 is ineligible. It is “directed to” a judicial exception because it recites a judicial exception (an abstract idea), and the claim as a whole does not integrate the exception into a practical application (and thus it is directed to an abstract idea), and the claim does not amount to significantly more than the exception (does not provide an inventive concept). Claim 2 recites the same judicial exception as claim 1 but is eligible because the claim as a whole improves speech-separation technology and thus integrates the exception into a practical application of separating speech and is therefore not “directed to” the judicial exception. Claim 3 is eligible because it recites a judicial exception (an abstract idea), but the claim as a whole integrates the exception into a practical application by improving speech-to-text transcription and is therefore not “directed to” the judicial exception.

BACKGROUND

Over the last several years, there has been a vast increase in the creation and consumption of audiovisual multimedia content. Smartphones, gaming consoles, and head-mounted devices are equipped with microphones and cameras for recording. These devices also incorporate technology that allows a user to control the device with voice commands. When the same devices are used to record an event, the devices may capture voice commands as part of the recording or fail to filter the voices of nearby speakers from the recording. These recordings could be used for simple replay on social media, for personal use, or in applications like automatic captioning or transcription.

Typical human listeners can easily perceive separate sources in an acoustic mixture. For example, a typical human being is capable of paying attention to a single conversation in a crowded restaurant even though the listener is surrounded by other audible conversations. When computers interpret spoken language, they receive an audio signal from a microphone that contains all the audio the microphone has picked up. For a computer to “pay attention” to a single conversation or speaker, the relevant speech must be separated from the rest of the audio signal. Traditional computer-based speech separation techniques perform well in distinguishing and separating different classes of audio (e.g., human speech and background noise) but perform poorly in separating audio from sources belonging to the same class (e.g., speech from different speakers). Some existing solutions rely on separating speech based on volume, but a volume-based approach is unreliable in environments in which speakers may vary in how loudly they speak or in their distance from the microphone. Other solutions require training the input device to recognize a particular voice but require the user to explicitly interact with the device to provide the training data. Further, speech separation systems are ill-suited for distinguishing a conversation between individuals of interest as opposed to commands issued by a single user. As a result, using traditional speech separation techniques may cause important information that was captured during the recording to be deleted, or unwanted data not being deleted at all. When

these techniques are used in pre-processing stages of speech-to-text or automatic speech recognition systems, the transcription quality suffers. There exists a need to remove these unwanted utterances from the audio, both for user privacy and for providing a high-quality recording or transcription.

Artificial neural networks (ANNs) offer a promising solution to separate speech signals from different sources. Applicant has filed a patent application that discloses a system which receives as input, a mixed speech signal x , from an audio recording device (e.g., a microphone) recording an event. The system uses a deep neural network (DNN), a type of ANN, that operates to promote separation of the features during clustering. Specifically, the DNN learns high-level feature representations of the signal x by mapping the feature representations to the embedding space. Subsequently, clustering is performed on these feature representations, with each cluster representing a distinct speech source, thereby separating speech signals of different sources s_n , where $n \in \{1, \dots, N\}$, identified in the mixed speech signal. The signals so separated can be used for downstream applications such as transcription, removal of voice commands from a recording, and uploading into a social networking site.

The DNN could be an autoencoder, recurrent neural network, or convolutional neural network. In one embodiment, the DNN is pretrained for speech separation on audio datasets comprising speech segments from multiple speakers. The training data may comprise manually labeled audio that specifies separate speakers and separate conversations, so that the DNN can learn to distinguish both individual voices and speech context. Using these deep learning techniques has the advantage of using a simple clustering to achieve the separated speech signals of the different sources, as explained below.

Whereas a traditional Fourier transform describes the various frequency components of an entire signal, speech has frequency components that change over time. A short-time Fourier transform (STFT) is a mathematical tool for obtaining a representation of a signal when frequency components change over time by performing a sequence of Fourier transforms of smaller “windows” or “frames” of the signal. In one embodiment, a mixed speech signal $x(t)$ is divided into T overlapping frames of equal length. This mixed speech signal is processed in the STFT domain as follows. For each frame, commonly known temporal features like pitch, variance, and zero-crossing rate are extracted and represented by a feature matrix FM_{tj} , where t is the frame index and j is the feature index. In addition to the temporal feature extraction, a spectrogram of each frame is generated by converting the mixed audio signal into time-frequency (TF) domain S_{tf} by using the STFT, where f is the frequency bin index. At the end of this step, each frame is represented by X_t , which corresponds to a spectrogram S_t and a corresponding row of the feature matrix FM_t .

The DNN learns the high-level feature representations of the input mixed speech signal x . Specifically, the DNN converts these feature representations X_t , obtained from the spectrograms S_t and the corresponding feature matrices FM_t , into multi-dimensional embedding vectors V . These embedding vectors V are assigned to the TF bins as a global function of the input signal ($V = f_{\theta}(X)$, where f_{θ} represents the function of the DNN). The DNN assigns embedding vectors V to each TF region such that the Euclidean distances between embedding vectors of TF bins dominated by the same source are minimized and the Euclidean distances between embedding

vectors of TF bins dominated by different sources are maximized. Thus, embedding vectors V for all TF bins, representing the different sources, are calculated.

Next, clustering is performed using a k-means clustering algorithm to separate different speech sources s_n , of the mixed signal. Embedding vectors V are clustered into k distinct groups, with each group representing a distinct speech source from s_n . The clustering algorithm arbitrarily chooses k initial centers C . Then, until the algorithm converges, embedding vectors V are assigned to their closest cluster center and each center is moved to the mean of its currently assigned clustering subset. By the end of this process, the embedding vectors V are partitioned into clusters corresponding to the different constituent sources, s_n .

Binary time-frequency masks are used to separate the signals by using a binary matrix to indicate which portions of a representation should be turned on or off. In audio processing, a binary mask is a matrix of binary values that correspond to sources such that it is multiplied with a spectrogram to include or exclude portions of the audio. The binary time-frequency mask for each speaker is obtained using the cluster assignments by assigning 1 to all the TF bins corresponding to the respective speaker and 0 to the rest of the TF bins. Inverse STFT converts the obtained separated signals into the time domain. The speech signals so separated can be used for multiple different downstream applications as discussed below.

Because the DNN assigns the embedding vectors V as a function of the entire input signal, the embedding vectors V take into account the global properties of the input signal, thereby allowing the k distinct groups to correspond to the N sources identified in the mixed speech signal, thereby providing a superior speech separation. This feature of the invention is an improvement over prior speech separation methods because it allows for blind speech separation (i.e., the system is not required to have prior knowledge of the number of speakers and does not need to be trained on speech from the different constituent sources of the mixed audio signal). Due to this, the DNN could be trained with mixed speech signals comprising a fewer number of speakers and could be used to separate speech signals from a larger number of sources. Also, as this speech separation process uses both temporal and spatial features of the speech signal and derives embedding vectors V based on the global properties of the input signal, it performs well with inter-speaker variability within the same audio class for downstream applications like automatic speech recognition (ASR). Test results have consistently shown that conventional ASR systems that used the disclosed process during pre-processing stages significantly reduced the gap in transcription performance for accented speakers over traditional speech-to-text systems.

One application of this process is the removal of voice commands or background conversations from a recording of an event like a baseball game. A speech signal from an undesired source, say s_s , is identified by conventional means, for example, by obtaining an audio clip corresponding to the undesired source. In one embodiment, a user may request a separated speech signal by providing a sample of the undesired audio signal to be redacted from the mixed audio signal x . The system compares the temporal features of the sample with the source signals s_n , to find the undesired source signal s_s , that is the closest match. After the inverse STFT step, the various speech signals from all the frames T are stitched together, excluding the speech signal from s_s . In a preferred embodiment, an overlap-add method is used to reconstruct the entire clean speech signal. The output is therefore a clean audio signal x' , such that x' includes speech signals from

all the sources s_n , where $n \in \{1, \dots, N\}$, excluding the speech signal from the source s_s . This clean audio signal x' is transmitted for storage in a remote location to be used for downstream applications like uploading into a social networking site.

Another application is real-time speech transcription or transcription of recorded audio. In one exemplary embodiment, a user may request a transcript of a desired source signal s_d , within the mixed speech signal x , during playback of a recorded audio, using a graphical user interface (GUI). After the inverse STFT step, the speech signal from only the desired source s_d is transmitted to a speech-to-text system. The ASR or speech-to-text system extracts spectral features from the desired source s_d using conventional means and generates a sequence of words, which are then converted into text. The system displays the translation result as text on the GUI.

Applicant has described that the systems discussed in the patent application could be implemented by one or more processors coupled with one or more non-transitory computer readable media. The methods described herein can be performed by execution of computer-readable instructions stored on a non-transitory computer-readable storage medium (e.g., random-access memory, flash memory, magnetic/optical storage, etc.) by the processor(s). The GUI is hardware or a combination of both hardware and software. The GUI is coupled to the systems described above and is configured to receive user instructions and outputs transcripts of audio selected by the user.

CLAIMS

[Claim 1] A speech separation method comprising:

- (a) receiving a mixed speech signal x comprising speech from multiple different sources s_n , where $n \in \{1, \dots, N\}$;
- (b) converting the mixed speech signal x into a spectrogram in a time-frequency domain using a short time Fourier transform and obtaining feature representation X , wherein X corresponds to the spectrogram of the mixed speech signal x and temporal features extracted from the mixed speech signal x ; and
- (c) using a deep neural network (DNN) to determine embedding vectors V using the formula $V = f_\theta(X)$, where $f_\theta(X)$ is a global function of the mixed speech signal x .

[Claim 2] The speech separation method of claim 1 further comprising:

- (d) partitioning the embedding vectors V into clusters corresponding to the different sources s_n ;
- (e) applying binary masks to the clusters to create masked clusters;
- (f) synthesizing speech waveforms from the masked clusters, wherein each speech waveform corresponds to a different source s_n ;

- (g) combining the speech waveforms to generate a mixed speech signal x' by stitching together the speech waveforms corresponding to the different sources s_n , excluding the speech waveform from a target source s_s such that the mixed speech signal x' includes speech waveforms from the different sources s_n and excludes the speech waveform from the target source s_s ; and
- (h) transmitting the mixed speech signal x' for storage to a remote location.

[Claim 3] A non-transitory computer-readable storage medium having computer-executable instructions stored thereon, which when executed by one or more processors, cause the one or more processors to perform operations comprising:

- (a) receiving a mixed speech signal x comprising speech from multiple different sources s_n , where $n \in \{1, \dots, N\}$, at a deep neural network (DNN) trained on source separation;
- (b) using the DNN to convert a time-frequency representation of the mixed speech signal x into embeddings in a feature space as a function of the mixed speech signal x ;
- (c) clustering the embeddings using a k-means clustering algorithm;
- (d) applying binary masks to the clusters to obtain masked clusters;
- (e) converting the masked clusters into a time domain to obtain N separated speech signals corresponding to the different sources s_n ; and
- (f) extracting spectral features from a target source s_d of the N separated speech signals and generating a sequence of words from the spectral features to produce a transcript of the speech signal corresponding to the target source s_d .

ANALYSIS

Claim 1 is ineligible.

Claim interpretation: Under the broadest reasonable interpretation, the terms of the claim are presumed to have their plain meaning consistent with the specification as it would be interpreted by one of ordinary skill in the art. See Manual of Patent Examining Procedure (MPEP) 2111.

Regarding step (a), the claim does not put any limits on how the mixed speech signal is received. The broadest reasonable interpretation of a mixed speech signal encompasses audible, spoken speech from different sources. The mixed speech signal could be received with a microphone in a user device or other sensor that converts sound into an electrical signal.

Regarding step (b), the claim specifies that the mixed speech signal is converted into a spectrogram in the time-frequency domain using an STFT. A feature representation X , corresponding to the spectrogram and temporal features extracted from the mixed speech signal

x , is obtained. The claim does not specify how the temporal features and the spectrogram of the mixed speech signal are obtained.

Regarding step (c), the claim specifies that a formula is used to determine embedding vectors based on the result of step (b). The claim further specifies that a DNN is used in the determination, but the claim does not include any details about the DNN or how it operates.

The broadest reasonable interpretation of claim 1 is a method of receiving spoken audio from different sources, deriving a temporal feature representation and a spectrogram of the audio, and using a DNN to calculate embedding vectors based on the temporal feature representation and a spectrogram using a mathematical formula.

Step 1: This part of the eligibility analysis evaluates whether the claim falls within any statutory category. See MPEP 2106.03. The claim recites the steps or acts of receiving a mixed speech signal, converting the mixed speech signal, and using a DNN to determine embedding vectors, and thus is a process (a series of steps or acts). A process is a statutory category of invention. (*Step 1: YES*).

Step 2A, Prong One: This part of the eligibility analysis evaluates whether the claim recites a judicial exception. As explained in MPEP 2106.04, subsection II, a claim “recites” a judicial exception when the judicial exception is “set forth” or “described” in the claim. The claim recites the step (b) of “converting the mixed speech signal x into a spectrogram in a time-frequency domain using an STFT, and obtaining feature representation X , wherein X corresponds to the spectrogram of the mixed speech signal x and temporal features extracted from the mixed speech signal x .” In this claim, “converting” the mixed speech signal x into a spectrogram in time-frequency domain involves a mathematical operation using an STFT. Therefore, the claim recites a mathematical operation for converting a signal from one domain to another using a specific transform function. The claim also recites a step (c) that determines “embedding vectors V using the formula $V = f_{\theta}(X)$, where $f_{\theta}(X)$ is a global function of the input signal.” The recited formula is clearly a mathematical formula or equation, and the determination is a mathematical calculation. Thus, the claim recites a mathematical formula or equation as well as a mathematical calculation, both of which fall within the mathematical concepts grouping of abstract ideas. As explained in the MPEP, when a claim recites multiple abstract ideas that fall in the same or different groupings, examiners should consider the limitations together as a single abstract idea, rather than as a plurality of separate abstract ideas to be analyzed individually. See MPEP 2106.04, subsection II.B. As the steps (b) and (c) fall within the same grouping of abstract ideas (i.e., mathematical concepts), these limitations are considered together as a single abstract idea for further analysis. (*Step 2A, Prong One: YES*).

Step 2A, Prong Two: This part of the eligibility analysis evaluates whether the claim as a whole integrates the recited judicial exception into a practical application of the exception. This evaluation is performed by (1) identifying whether there are any additional elements recited in the claim beyond the judicial exception, and (2) evaluating those additional elements individually and in combination to determine whether the claim as a whole integrates the exception into a practical application. See MPEP 2106.04(d).

The claim recites a first additional element of “receiving a mixed speech signal x comprising speech from multiple different sources s_n , where $n \in \{1, \dots, N\}$ ” in step (a). As explained above, step(a) is claimed at a high level of generality and could describe receiving the mixed speech signal with a microphone in a user device or other sound sensor. The element amounts to mere data gathering. It is necessary to acquire the data in order to use the recited judicial exception to perform the calculation (i.e., converting the mixed speech signal into a time-frequency domain representation). The “receiving” element does not impose any other meaningful limits on the claim. Therefore, the additional limitation is insignificant extra-solution activity. See MPEP 2106.05(g).

The method also recites a second additional element in step (c) of “using a deep neural network (DNN) to determine embedding vectors $V \dots$ ” When determining whether a claim simply recites a judicial exception with the words “apply it” (or an equivalent), such as mere instructions to implement an abstract idea on a computer, examiners may consider: (1) whether the claim recites only the idea of a solution or outcome i.e., the claim fails to recite details of how a solution to a problem is accomplished; (2) whether the claim invokes computers or other machinery merely as a tool to perform an existing process; and (3) the particularity or generality of the application of the judicial exception. See MPEP 2106.05(f). Here, there are no details about a particular DNN or how the DNN operates to derive the embedding vectors other than that it is being used to determine the embedding vectors. The DNN is used to generally apply the abstract idea (i.e., perform the mathematical calculation using the recited mathematical equation) without placing any limitation on how the DNN operates to derive the embedding vectors as a function of the input signal. In addition, the limitation recites only the idea of determining embedding vectors using a DNN without details on how this is accomplished. The claim omits any details as to how the DNN solves a technical problem, and instead recites only the idea of a solution or outcome. Also, the claim invokes a generic DNN merely as a tool for making the recited mathematical calculation rather than purporting to improve the technology or a computer. See MPEP 2106.05(f). Therefore, the limitation represents no more than mere instructions to apply the judicial exception on a computer. It can also be viewed as nothing more than an attempt to generally link the use of the judicial exception to the technological environment of computers.

The disclosure identifies a technical problem encountered in the field of speech-separation and provides this invention as a solution to the identified speech-separation problem. The disclosure clearly describes how this invention offers an improvement over existing speech-separation methods by providing a particular speech-separation technique that solves the problem of separating speech from different speech sources belonging to the same class, while not requiring prior knowledge of the number of speakers or speaker-specific training. In particular, the improvement is achieved by determining embedding vectors as a function of the input signal, partitioning those vectors into clusters, and synthesizing a reconstructed mixed speech signal based on these clusters. The claim, however, only requires determining the embedding vectors and therefore does not reflect the improvement discussed in the disclosure. The recited generic DNN merely adds a generic computer component to perform the method and therefore fails to provide an improvement to the technology or technical field. See MPEP 2106.05(a). Even when viewed in combination, these additional elements do not integrate the recited judicial exception into a practical application (*Step 2A, Prong Two: NO*), and the claim is directed to the judicial exception. (*Step 2A: YES*).

Step 2B: This part of the eligibility analysis evaluates whether the claim as a whole amounts to significantly more than the recited exception, i.e., whether any additional element, or combination of additional elements, adds an inventive concept to the claim. See MPEP 2106.05.

At Step 2A, Prong Two, the second additional element in step (c), “using a deep neural network,” was found to represent no more than mere instructions to apply the judicial exception on a computer using generic computer components. The analysis under Step 2A, Prong Two is carried through to Step 2B. Further, the first additional element in step (a) was found to be insignificant extra-solution activity. However, a conclusion that an additional element is insignificant extra-solution activity in Step 2A should be re-evaluated in Step 2B. See MPEP 2106.05, subsection I.A. At Step 2B, the re-evaluation of the insignificant extra-solution activity consideration takes into account whether or not the extra-solution activity is well understood, routine, and conventional in the field. See MPEP 2106.05(g). Here, the step of receiving a mixed speech signal is mere data gathering that is recited at a high level of generality, and as discussed in the disclosure, is well-understood (e.g., the first paragraph of the background explains that smartphones and other devices have long been equipped to receive a mixed speech signal via the microphones integrated into the devices). Therefore, this limitation remains insignificant extra-solution activity even upon reconsideration and does not amount to significantly more.

Even when considered in combination, these additional elements represent mere instructions to apply an exception and insignificant extra-solution activity, and therefore do not provide an inventive concept (*Step 2B: NO*). The claim is not eligible.

Claim 2 is eligible.

Claim interpretation: Under the broadest reasonable interpretation, the terms of the claim are presumed to have their plain meaning consistent with the specification as it would be interpreted by one of ordinary skill in the art. See MPEP 2111. Claim 2 is a dependent claim that depends from and requires all the limitations of claim 1.

Regarding step (d), the claim does not place any limits on how the embedding vectors are partitioned into clusters corresponding to the different sources. The clustering could be performed using a k-means algorithm as indicated in the disclosure or any other algorithm known to a person of ordinary skill in the art.

Step (e) requires applying binary masks to the clusters. The plain meaning of “applying a binary mask” to a person of ordinary skill in the art is a mathematical operation of using a binary matrix to indicate which portions of a representation should be turned on or off. Such masking can be performed in any way known in the art, for example, by performing bitwise operations on two numbers or multiplying a binary matrix with another numerical representation, etc.

Step (f) synthesizes the speech waveforms from the masked clusters, where each waveform corresponds to a different source of the mixed speech signal. The background section explains that this synthesis is performed by converting the masked clusters into separate speech signals in the time domain, corresponding to the different sources in the mixed speech signal, using inverse STFT.

Step (g) combines the separated speech waveforms of step (f) by stitching together the speech waveforms corresponding to the different sources s_n , excluding the speech waveform from the target source s_s such that the resultant mixed speech signal excludes at least one speech signal from a target source and includes speech signals from other sources. The background section indicates that an overlap-add method is used to perform the stitching to reconstruct the entire clean speech signal, but step (g) is broad enough to cover any method known in the art to perform the stitching.

Step (h) specifies that the reconstructed mixed speech signal is transmitted for storage to a remote location. Note that this step does not require that the reconstructed mixed speech signal is actually stored, but that the signal is transmitted with the intended result of being stored at a remote location.

As described above with respect to claim 1, the broadest reasonable interpretation of claim 2 is a method of receiving spoken audio from different sources, deriving a temporal feature representation and a spectrogram of the audio, and using a DNN to calculate embedding vectors based on the temporal feature representation and a spectrogram using a mathematical formula. The embedding vectors are then partitioned into clusters, the clusters are modified using binary masks, and the modified clusters are synthesized into separate speech signals. A new, combined mixed speech signal is created by excluding at least one speech signal from one source and including speech signals from other sources. The combined mixed speech signal is then transmitted.

Step 1: As discussed above with respect to claim 1, the claim recites the steps or acts of receiving a mixed speech signal, converting the mixed speech signal, and using a DNN to determine embedding vectors. Dependent claim 2 adds further steps of partitioning the embedding vectors into clusters, applying binary masks to the clusters and synthesizing the result, combining the resultant signals into a mixed speech signal, and transmitting the mixed speech signal. Claim 2 is therefore a process (a series of steps or acts). A process is a statutory category of invention. (*Step 1: YES*).

Step 2A, Prong One: This part of the eligibility analysis evaluates whether the claim recites a judicial exception. MPEP 2106.04, subsection II explains that a claim “recites” a judicial exception when the judicial exception is “set forth” or “described” in the claim. As discussed above with respect to claim 1, steps (b) and (c) recite mathematical concepts. Step (d) recites “partitioning the embedding vectors V into clusters corresponding to the different sources s_n .” The claim places no limits on how this partitioning is performed. That is, nothing in the claim element precludes the step from practically being performed in the mind. For example, “partitioning . . . into clusters” encompasses a human arbitrarily selecting groups of vectors and mentally assigning them to clusters. The recitation of a DNN in this claim does not negate the mental nature of these limitations because the claim here merely uses the DNN as a tool to perform the otherwise mental process. See MPEP 2106.04(a)(2), subsection III.C. The claim therefore recites a mental process.

The claim also recites step (e) - “applying binary masks to the clusters to create masked clusters.” This step recites a mathematical operation for generating numbers based on binary calculations. Therefore, the claim recites a further mathematical calculation, which falls within the mathematical concept grouping of abstract ideas.

“Unless it is clear that a claim recites distinct exceptions, such as a law of nature and an abstract idea, care should be taken not to parse the claim into multiple exceptions, particularly in claims involving abstract ideas.” MPEP 2106.04, subsection II.B (discussing *Bilski v. Kappos*, 561 U.S. 593 (2010)). Thus, if possible, the examiner should consider the limitations together as a single abstract idea rather than as a plurality of separate abstract ideas to be analyzed individually. As discussed above, steps (b), (c), and (e) of claim 2 recite mathematical concepts and step (d) recites a mental process. An examiner should identify the claim as reciting both a mental process and a mathematical concept for Step 2A, Prong One, and consider the limitations (b)-(e) together as a single abstract idea for further analysis. See MPEP 2106.04, subsection II.B.

Step (f) synthesizes the speech waveforms from the masked clusters, where each waveform corresponds to a different source of the mixed speech signal. This step requires converting the result of step (e) into separate speech signals in the time domain. Synthesizing speech waveforms from a cluster of numbers is not a process that can be practically performed in the human mind. Further, while the synthesis would involve a mathematical calculation, the claim does not specify any mathematical formula, calculations, or relationships. In addition, step (f) is not a method of organizing human activity because it does not fall within the enumerated sub-groupings of fundamental economic principles or practices, commercial or legal interactions, and managing personal behavior and relationships or interactions between people.

Similarly, step (g) combines the speech waveforms to generate a mixed speech signal by stitching together speech waveforms corresponding to the different sources excluding the speech signal from a target source. Even though the disclosure explains that the stitching could be performed by an overlap-add method, which is a mathematical operation, the claim recites no details of how the stitching is performed. Additionally, while the claim recites variables, variables on their own are not mathematical relationships, formulas, or calculations. Therefore, the combining step is merely based on or involves a mathematical concept but does not recite a mathematical concept. See MPEP 2106.04(a)(2), subsection I. Generating a mixed speech signal such that it includes speech signals from different sources and excludes the speech signal from a target source is not a process that can be practically performed in a human mind. Therefore, step (g) is neither a mathematical concept nor a mental process. In addition, step (g) is not a method of organizing human activity because it does not fall within the enumerated sub-groupings.

As discussed above, limitations (b)-(e) are considered as a single abstract idea, and the claim is considered as reciting limitations which fall within the mathematical concepts grouping of abstract ideas. (*Step 2A, Prong One: YES*).

Step 2A, Prong Two: This part of the eligibility analysis evaluates whether the claim as a whole integrates the recited judicial exception into a practical application of the exception. This evaluation is performed by (1) identifying whether there are any additional elements recited in the claim beyond the judicial exception, and (2) evaluating those additional elements

individually and in combination to determine whether the claim as a whole integrates the exception into a practical application. See MPEP 2106.04(d).

As discussed above with respect to claim 1, step (a) recites a data gathering step of receiving a mixed speech signal, and step (c) recites using a DNN to determine embedding vectors, which is equivalent to the words “apply it.”

Step (h) recites “transmitting the mixed speech signal x' for storage to a remote location.” As discussed above, the mixed speech signal x' is transmitted with the intended result of being stored at a remote location for future use, for example, for audio playback or to be uploaded into a social media website. This limitation is merely a post-solution step of transmitting data output—a nominal addition to the claim that does not meaningfully limit the claim. Therefore, step (h) is an insignificant extra-solution activity. See MPEP 2106.05(g). The remaining additional limitations beyond the abstract idea recited in limitations (b), (c), (d), and (e) are synthesizing speech waveforms from the masked clusters, recited in limitation (f), and generating the mixed speech signal excluding the speech signal from a target source, recited in limitation (g).

Step (f) recites “synthesizing speech waveforms from the masked clusters, wherein each speech waveform corresponds to a different source s_n ,” and step (g) recites “combining the speech waveforms to generate a mixed speech signal x' by stitching together the speech waveforms corresponding to the different sources s_n , excluding the speech waveform from a target source s_s such that the mixed speech signal x' includes speech signals from the different sources s_n , where $n \in \{1, \dots, N\}$, and excludes the speech signal from the target source s_s .” Steps (f) and (g) integrate the abstract idea into a practical application. The disclosure explains that devices that capture audio cannot properly distinguish different speech sources belonging to the same class and that the current available solutions do not adequately address this problem because they require a target user, whose speech is to be recognized, to explicitly interact with the device to provide training data. The disclosure states that this invention offers an improvement over existing speech-separation methods by providing a particular speech-separation technique that solves the problem of separating speech from different speech sources belonging to the same class, while not requiring prior knowledge of the number of speakers or speaker-specific training. The claim reflects the improvement discussed in the disclosure by reciting details of how the DNN aids in the cluster assignments to correspond to the sources identified in the mixed speech signal, which are then synthesized into separate speech waveforms in the time domain and converted into a mixed speech signal, excluding audio from the undesired source. See MPEP 2106.05(a). While on their own steps (b)-(e) recite judicial exceptions, steps (f) and (g) are directed to creating a new speech signal that no longer contains extraneous speech signals from unwanted sources. The claimed invention reflects this technical improvement by including these features. Further, converting clusters into separate speech waveforms and generating a mixed speech signal from the separate speech waveforms are not insignificant extra-solution activity, mere instructions to apply the exception, or mere field of use limitations. Rather, these steps reflect the improvement described in the disclosure. Accordingly, the claim is directed to an improvement to existing computer technology or to the technology of speech separation, and the claim integrates the abstract idea into a practical application. (*Step 2A, Prong Two: YES*). The claim is eligible. (*Step 2A: NO*).

Claim 3 is eligible.

Claim interpretation: Under the broadest reasonable interpretation, the terms of the claim are presumed to have their plain meaning consistent with the specification as it would be interpreted by one of ordinary skill in the art. See MPEP 2111. The preamble specifies that the claim is to a non-transitory computer-readable storage medium containing instructions that when executed by the one or more processors with which it is associated, cause the processor to perform the receiving, generating, and producing steps recited in the claim. The disclosure gives random-access memory, flash memory, magnetic/optical storage, etc. as examples of a non-transitory computer-readable storage medium but the claim does not specify the type of non-transitory computer-readable storage medium.

Regarding step (a), the claim does not place any limits on how the mixed speech signal is received. The broadest reasonable interpretation of a mixed speech signal encompasses audible, spoken speech from different sources. The mixed speech signal could be received with a microphone in a user device or other sensor that converts sound into an electrical signal. The claim requires a DNN that receives this mixed speech signal to be trained for source separation.

Regarding step (b), the claim specifies that the DNN is used to convert a time-frequency representation of the mixed speech signal to embeddings in the feature space but does not provide any structural details about the DNN itself. The specification discusses processing the mixed speech signal in STFT domain to obtain temporal features and spectrograms, which are then used by the DNN to determine embeddings V in the feature space as a function of the input signal. As the claim does not specify how the time-frequency representation of the mixed speech signal is obtained or how it is converted into embeddings as a function of the signal x , this sub-step can be performed as indicated in the specification or by any other method known to a person having ordinary skill in the art.

Step (c) requires clustering the embeddings using a k-means clustering algorithm but does not place any limits on how the algorithm is implemented. The clustering could be performed using the k-means algorithm either as described in the disclosure or in other ways known to a person having ordinary skill in the art.

Step (d) obtains masked clusters by applying binary masks to the clusters. The plain meaning of “applying a binary mask” to a person of ordinary skill in the art is a mathematical operation of using a binary matrix to indicate which portions of a representation should be turned on or off. Such masking can be performed in any way known in the art, for example, by performing bitwise operations on two numbers or multiplying a binary matrix with another numerical representation, etc.

Step (e) requires converting the masked clusters into separate speech signals in the time domain, corresponding to the different sources in the mixed speech signal. The claim does not specify how the conversion is performed.

The final step (f) requires extracting spectral features from only one target source s_d of the N separated signals from the output of step (e) and generating a sequence of words from the

spectral features to produce a transcript of the speech signal corresponding to the target source s_d . The disclosure states that the desired speech signal from the output of the inverse STFT step is transmitted to an ASR, which extracts spectral features from the desired source s_d using conventional means and generates a sequence of words, which are then converted into text. The claim does not specify a particular way in which the extracting and generating is done; therefore, the step of extracting the spectral features and generating a sequence of words can be implemented by ASR systems known in the art to produce the transcript.

Therefore, the broadest reasonable interpretation of claim 3 is a non-transitory computer-readable storage medium that stores instructions, which when executed by a processor, cause the processor to perform the steps of receiving a mixed speech signal constituting audio from different sources by a DNN that calculates embedding vectors from time-frequency representation of the signal. The embeddings are then partitioned into clusters and the clusters are converted into separate speech signals in the time domain. Only one specific separated speech signal of these separated speech signals is converted into text to produce a transcript.

Step 1: This part of the eligibility analysis evaluates whether the claim falls within any statutory category. See MPEP 2106.03. The preamble specifies that the claim is to a non-transitory computer-readable medium, which causes the one or more processors it is associated with to perform a series of steps. The disclosure gives random-access memory, flash memory, and magnetic/optical storage as non-limiting examples of a non-transitory computer-readable storage medium. The broadest reasonable interpretation of the claim covers only statutory embodiments of a computer-readable medium in light of the disclosure and not a transitory signal. A non-transitory computer-readable storage medium falls within the “manufacture” category of invention. (*Step 1: YES*).

Step 2A, Prong One: This part of the eligibility analysis evaluates whether the claim recites a judicial exception. As explained in MPEP 2106.04, subsection II, a claim “recites” a judicial exception when the judicial exception is “set forth” or “described” in the claim. Step (b) requires converting a time-frequency representation of the mixed speech signal into embeddings in a feature space as a function of the mixed speech signal, which is a mathematical equation written in text format. Step (c) requires clustering the embeddings by a k-means clustering algorithm, which is a mathematical calculation. Step (d) obtains masked clusters by applying binary masks to the clusters, which is also a mathematical calculation. Thus, the claim recites mathematical calculations which fall within the mathematical concepts grouping of abstract ideas.

Step (e) requires converting a cluster of points in feature space into speech signals in the time domain, which is not a process that can be practically performed in the human mind. Further, while the conversion may be based on mathematical concepts, the claim does not specify any mathematical formulae, calculations, or relationships.

Finally, step (f) requires extracting spectral features from only one target source s_d of the N separated signals from the output of step (e) and generating a sequence of words from the spectral features to produce a transcript of the speech signal corresponding to the target source s_d . Extracting spectral features from a signal and generating a sequence of words from these extracted features to produce a transcript is not a process that can practically be performed in the

human mind. While this extraction and generation involves mathematical operations, the claim does not specify any mathematical formulae, calculations, or relationships. Neither do steps (e) and (f) fall within the enumerated sub-groupings of a method of organizing human activity. Therefore, steps (e) and (f) do not recite a judicial exception.

As explained in the MPEP, when a claim recites multiple abstract ideas that fall in the same or different groupings, examiners should consider the limitations together as a single abstract idea rather than as a plurality of separate abstract ideas to be analyzed individually. See MPEP 2106.04, subsection II.B. As discussed above, limitations (b)-(d) recite mathematical concepts. As all these steps (b)-(d) fall within the same grouping of abstract ideas (i.e., mathematical concepts), these limitations are considered together as a single abstract idea for further analysis. (*Step 2A, Prong One: YES*).

Step 2A, Prong Two: This part of the eligibility analysis evaluates whether the claim as a whole integrates the recited judicial exception into a practical application of the exception. This evaluation is performed by (1) identifying whether there are any additional elements recited in the claim beyond the judicial exception, and (2) evaluating those additional elements individually and in combination to determine whether the claim as a whole integrates the exception into a practical application. See MPEP 2106.04(d).

Claim 3 recites as step (a) the additional limitation of “receiving a mixed speech signal x comprising speech from multiple different sources s_n , where $n \in \{1, \dots, N\}$.” As explained in the claim interpretation portion above, the limitation is claimed at a high level of generality and could describe receiving the mixed speech signal with a microphone in a user device or other sound sensor. The limitation amounts to mere data gathering. It is necessary to acquire the data in order to use the recited judicial exception to perform the calculations of steps (b), (c), and (d). The limitation does not impose any other meaningful limits on the claim. Therefore, this additional limitation is insignificant extra-solution activity. See MPEP 2106.05(g).

Another additional limitation beyond the abstract idea recited in step (b) is the use of a DNN trained on source separation. When determining whether a claim simply recites a judicial exception with the words “apply it” (or an equivalent), such as mere instructions to implement an abstract idea on a computer, examiners may consider: (1) whether the claim recites only the idea of a solution or outcome i.e., the claim fails to recite details of how a solution to a problem is accomplished; (2) whether the claim invokes computers or other machinery merely as a tool to perform an existing process; and (3) the particularity or generality of the application of the judicial exception. See MPEP 2106.05(f). Here, the claim recites no details about a particular DNN. The DNN is used to generally apply the abstract idea (i.e., perform the mathematical calculation recited in step (b)) without placing any limitation on how the DNN operates to derive the embedding vectors. In addition, the limitation would cover every mode of implementing the recited abstract idea using a DNN. The claim omits any details as to how the DNN solves a technical problem and instead recites only the idea of a solution or outcome. See MPEP 2106.05(f). Therefore, the limitation represents no more than mere instructions to implement the abstract idea recited in step (b), which is equivalent to adding the words “apply it” to the recited judicial exception. In addition, the claim confines the use of the judicial exception recited in step (b) to the technological environment of a DNN by generally linking the use of the judicial

exception to the recited DNN. Therefore, this general DNN recitation does not integrate the judicial exception into a practical application. See MPEP 2106.05(h). Therefore, it can also be viewed as nothing more than an attempt to generally link the use of the judicial exception to a particular field of use or a technological environment.

The remaining additional limitations are step (e), which converts the masked clusters into N separate speech signals in time domain, and step (f), which extracts spectral features from only one target source s_d of the N separate signals from the output of step (e) and generates a sequence of words from the spectral features to produce a transcript. These additional limitations integrate the abstract idea recited in steps (b), (c), and (d) into a practical application of speech-to-text conversion.

The disclosure explains that devices that capture audio perform poorly in distinguishing conversations between individuals of interest from unwanted utterances due to their inability in distinguishing different speech sources belonging to the same class, thereby resulting in poor quality transcriptions of the recorded speech. The disclosure states that this invention offers an improvement over existing speech-separation methods by providing a particular speech-separation technique that solves the problem of separating speech from different speech sources belonging to the same class, and also performing well with inter-speaker variability within the same audio class for transcriptions. The disclosure states that the invention derives embeddings by the DNN based on the global properties of the input signal, which is an improvement over prior art speech separation methods. In addition, the invention uses both temporal and spatial features of the speech signal; this feature of the invention helps a downstream conventional speech-to-text system reduce the gap in transcription performance for accented speakers over traditional speech-to-text methods.

Here, the claim reflects these technical improvements discussed in the disclosure by reciting details of how the DNN trained on source separation aids in the cluster assignments to correspond to the sources identified in the mixed speech signal, which are then converted into separate speech signals in the time domain to generate a sequence of words from the spectral features, thereby making individual transcription of each separated speech signal possible. See MPEP 2106.05(a).

While on their own, the steps (b), (c), and (d) recite an abstract idea, the ordered combination of the steps of receiving a mixed speech signal, processing the speech signal to produce masked clusters, converting the masked clusters into separate signals in time domain, extracting spectral features from one such converted signal, and generating a sequence of words from the extracted spectral features to produce a transcript reflects the technical improvement discussed in the disclosure. Accordingly, the claim is directed to an improvement to existing speech-to-text technology, and the claim integrates the abstract idea recited in steps (b), (c), and (d) into a practical application of speech-to-text conversion of a speech signal corresponding to one source of the mixed speech signal. Thus, the claim as a whole integrates the judicial exception into a practical application (*Step 2A, Prong Two: YES*), such that the claim is not directed to the judicial exception. (*Step 2A: NO*). The claim is eligible.

Example 49. Fibrosis Treatment

This example illustrates the analysis of method claims reciting an artificial intelligence model that is designed to assist in personalizing medical treatment to the individual characteristics of a particular patient. The recited condition (post-implantation inflammation) is hypothetical, but open angle glaucoma and fibrosis are known medical conditions. The recited anti-fibrotic drugs (drug A and Compound X) are hypothetical fibrosis treatments. Lifestyle adjustments, pharmaceutical eye drops, laser eye surgery, and eye drainage device implants (including microstents) are known glaucoma treatments. The recited computer-implemented machine learning model (ezAI model) is hypothetical; however, genome-wide association studies, polygenic risk scores, and single-nucleotide polymorphisms (SNPs) are known in the genomics field. Claim 1 is ineligible because it recites a judicial exception (an abstract idea) and the claim as a whole does not integrate the exception into a practical application (and is thus directed to an abstract idea) and the claim does not provide significantly more than the exception (does not provide an inventive concept). Even though claim 2 recites the same judicial exception, claim 2 is eligible because the claim as a whole integrates the exception into a practical application and is therefore not “directed to” the judicial exception.

BACKGROUND

Glaucoma is a leading cause of blindness globally. The most common form is open angle glaucoma, in which irreversible vision loss results from cell and optic nerve damage caused largely by poor drainage of aqueous humor from the eye. Depending on condition severity and the timing of the diagnosis, treatment can include lifestyle adjustments, pharmaceutical eye drops, laser eye surgery, or drainage device implants to facilitate healthy drainage. While newer drainage devices, such as microstents, are more comfortable than earlier drainage devices, post-surgery scarring and inflammation due to fibrosis remain issues. Commonly prescribed anti-fibrotic drugs, such as drug A, can reduce scarring but do so non-specifically while causing more inflammation (“post-implantation inflammation” or “PI”) that further damages the eye.

Applicant developed a new anti-fibrotic drug, Compound X, that effectively reduces scarring around a microstent implantation site in glaucoma patients at high risk of PI after microstent implant surgery without the undesirable side effects of drug A. In connection with this invention, the applicant filed a patent application disclosing Compound X and describing how it may be topically administered in eye drop form after microstent implant surgery.

The applicant also discovered that PI is a polygenic condition, meaning that it results from the interaction of multiple genes rather than any single gene. Using standard methodologies, the applicant conducted a large-scale genome wide association (GWA) study on PI in glaucoma patients and identified 37 informative single-nucleotide polymorphisms (SNPs) that have a statistically significant association with PI. The application defines “informative SNPs” as the 37 SNPs identified in the GWA study. These informative SNPs serve as genetic markers for the genes that give rise to PI.

From the GWA study data, the applicant developed a polygenic risk score (PRS) model by classic clumping and thresholding methods to provide a weighted PRS. As is known in the art, a PRS is a single value estimate of an individual's relative risk for a given phenotype or condition based on the individual's genotype. In brief, after collecting a sample from a patient, the sample is sequenced and genotyped to provide a genotype dataset. Alleles in the genotype dataset corresponding to informative SNPs selected from the GWA study are identified, tallied, and weighted (using effect sizes derived from GWA study summary statistics as weights). A PRS is then generated as a sum of the weighted values. Sample collection, sequencing, and genotyping for the purposes of providing a genotype dataset may be carried out by conventional methods known in the art.

The specification discloses a method of identifying glaucoma patients at high risk of PI using this PRS model. An identification of a patient as being at high risk of PI is made by ranking a weighted PRS once the score has been generated. The application defines a "glaucoma patient at high risk of PI" as a glaucoma patient having a weighted PRS in the top quartile of PRS values when ranked against reference PRS values established during PRS model development. The disclosure teaches that determining patient risk using a weighted PRS as disclosed and accordingly customizing treatment lends to better prognosis after implant surgery.

The applicant also discloses a computer-implemented machine learning model (referred to as "the ezAI model" for the purposes of this disclosure) and its clinical applications. Given an input of a patient's genotype dataset, the ezAI model calculates a weighted PRS from informative SNPs in the dataset—using multiplication to weight corresponding alleles in the dataset by their effect sizes and addition to sum the weighted values. Using the same weights and informative SNPs, the ezAI model improves upon the base PRS model by determining a risk score and providing a classification in less time.

CLAIMS

[Claim 1] A post-surgical fibrosis treatment method comprising:

- (a) collecting and genotyping a sample from a glaucoma patient to provide a genotype dataset;
- (b) identifying the glaucoma patient as at high risk of post-implantation inflammation (PI) based on a weighted polygenic risk score that is generated from informative single-nucleotide polymorphisms (SNPs) in the genotype dataset by an ezAI model that uses multiplication to weight corresponding alleles in the dataset by their effect sizes and addition to sum the weighted values to provide the score; and
- (c) administering an appropriate treatment to the glaucoma patient at high risk of PI after microstent implant surgery.

[Claim 2] The method of claim 1, wherein the appropriate treatment is Compound X eye drops.

ANALYSIS

Claim 1 is ineligible.

Claim Interpretation: Under the broadest reasonable interpretation, the terms of the claim are presumed to have their plain meaning consistent with the specification as it would be interpreted by one of ordinary skill in the art. See Manual of Patent Examining Procedure (MPEP) 2111.

Claim 1 does not restrict how the sample is collected or genotyped in step (a). The specification explains that these steps may be carried out by any conventional method known in the art.

In step (b), the terms “glaucoma patient . . . at high risk of PI” and “informative SNPs” are given special definitions in the specification. Based on these special definitions, step (b) requires identifying the glaucoma patient as at high risk of PI where the patient’s weighted PRS is in the top quartile of PRSs when ranked against reference PRS values established during PRS model development. Step (b) also requires that the weighted PRS is generated using an ezAI model from the 37 SNPs identified in the GWA study that are present in the patient’s genotype dataset.

Step (b) indicates that the PRS is generated from informative SNPs in the genotype dataset by “an ezAI model.” As explained in the specification, an ezAI model can be a computer-implemented machine learning model that calculates a weighted PRS from alleles corresponding to informative SNPs present in a given input of a patient’s genotype dataset. See MPEP 2111.01, subsection V. The score is calculated using multiplication to weight the alleles in the dataset by their effect sizes and addition to sum the weighted values according to the plain language of the claim.

Step (c) does not require any specific treatment or prophylaxis because the claim merely recites “administering an appropriate treatment.” For instance, treatment can involve administering Compound X or any common anti-fibrotic treatment, such as drug A, after microstent implant surgery.

Based on the plain meaning of the words in the claim, the broadest reasonable interpretation of claim 1 is: a method of collecting and genotyping a sample to provide a genotype dataset, identifying the glaucoma patient by phenotype (at high risk of PI) based on a weighted PRS generated from informative SNPs in the genotype dataset using an ezAI model that calculates the score using multiplication to weight alleles corresponding to the informative SNPs by their effect sizes and addition to sum the weighted values, and administering a treatment to the patient at high risk of PI after microstent implant surgery.

Step 1: This part of the eligibility analysis evaluates whether the claim falls within any statutory category. MPEP 2106.03. The claim recites at least one step or act of determining the patient’s risk for a condition. Thus, the claim is a process, which is a statutory category of invention. (*Step 1: YES*).

Step 2A, Prong One: This part of the eligibility analysis evaluates whether the claim recites a judicial exception. As explained in MPEP 2106.04, subsection II, a claim “recites” a judicial exception when the judicial exception is “set forth” or “described” in the claim.

Limitation (b) in the claim recites “identifying the glaucoma patient as at high risk of PI based on a weighted PRS.” Under its broadest reasonable interpretation consistent with the specification, the plain and ordinary meaning of this limitation requires an evaluation of patient risk by determining whether a patient’s weighted PRS falls within the top quartile of scores when ranked against reference PRS values established during PRS model development. As this step requires comparing a patient’s score against known top-quartile scores, the limitation falls into the “mental process” grouping of abstract ideas because the evaluation can be practically performed in the human mind. This limitation further recites a law of nature because it describes the naturally occurring relationship between a patient’s genotype (a particular combination of genes giving rise to PI) and phenotype (risk for PI).

Limitation (b) also recites “a weighted PRS that is generated from informative SNPs in the genotype dataset by an ezAI model that uses multiplication to weight corresponding alleles in the dataset by their effect sizes and addition to sum the weighted values to provide the score.” As discussed above, the broadest reasonable interpretation of this limitation requires a mathematical calculation. Namely, an arithmetic calculation (multiplication to weight alleles by their effect sizes and addition to sum the weighted values) is required to generate a weighted risk score. Limitation (b) hence recites a “mathematical calculation” and so falls into the “mathematical concepts” grouping of abstract ideas. For instance, given a genotype dataset, a physician would identify, tally, and weight alleles corresponding to informative SNPs present in the dataset by their effect sizes, then sum the resulting values to generate a risk score. See MPEP 2106.04(a)(2), subsection I.C. This limitation also falls into the “mental process” grouping because it requires an assessment of the genotype dataset in order to identify informative SNPs in evaluating patient risk. Moreover, the recited mathematical calculation is simple enough that it can be practically performed in the human mind. Even if most humans would use a physical aid, like a pen and paper or a calculator, to make such calculations, the use of a physical aid would not negate the mental nature of this limitation. See MPEP 2106.04(a)(2), subsection III.B.

As there are no bright lines between the types of judicial exceptions, and many of the concepts identified by the courts as exceptions can fall under several exceptions, MPEP 2106.04, subsection I instructs examiners to “identify . . . the claimed concept (the specific claim limitation(s) that the examiner believes may recite an exception) [that] aligns with at least one judicial exception.” While limitation (b) can be categorized under several exceptions (a mathematical concept-type abstract idea, a mental process-type abstract idea, and a law of nature), it is adequate for an examiner to identify the limitation as falling under at least one judicial exception and to base further analysis on that identification. The remainder of this discussion is premised on the recited exception as an abstract idea. See MPEP 2106.04, subsection II.B. (*Step 2A, Prong One: YES*).

Step 2A, Prong Two: This part of the eligibility analysis evaluates whether the claim as a whole integrates the recited judicial exception into a practical application of the exception. This evaluation is performed by (1) identifying whether there are any additional elements recited in the claim beyond the judicial exception, and (2) evaluating those additional elements individually and in combination to determine whether the claim as a whole integrates the exception into a practical application. MPEP 2106.04(d).

The claim recites the additional elements of “collecting” and of “genotyping a sample from a glaucoma patient to provide a genotype dataset” in limitation (a). The additional element of “collecting” a sample is an insignificant extra-solution activity that amounts to mere data gathering incidental to limitation (b). The additional element of “genotyping” also represents mere data gathering (a set of genotypes from the collected sample) because all uses of the judicial exception require genotyping the collected sample (the mental process in limitation (b) uses the genotype dataset to determine a patient’s risk). See MPEP 2106.05(g). As such, genotyping the sample is also an insignificant extra-solution activity.

The claim also recites the additional element of “administering an appropriate treatment to the glaucoma patient at high risk of PI” in limitation (c). Although this limitation indicates that a treatment is to be administered, it does not provide any information as to how the patient is to be treated or what the treatment is, but instead covers any possible treatment that a medical professional decides to administer to the patient. As such, there are no meaningful constraints on the administering step such that the particular treatment or prophylaxis consideration would apply because it is not limited to any particular manner or type of treatment. See MPEP 2106.04(d)(2). Moreover, like the claims in *Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 566 U.S. 66, 78 (2012), the claim merely provides the identification made in limitation (b) to the relevant audience (such as a physician or another medical professional) and at most adds a suggestion to take that identification into account when treating patients. Limitation (c) may thus be understood as no more than an attempt to generally link the judicial exception to a field of use. See MPEP 2106.05(h). Therefore, limitation (c) fails to meaningfully limit the claim because it does not require any particular application of the abstract idea and therefore amounts only to a generic instruction to “apply” the exception or to a mere indication of the field of use or technological environment in which the abstract idea is performed.

While the disclosure states that “the ezAI model improves upon the base PRS model by determining a risk score and providing a classification in less time,” there is no improvement to the functioning of a computer nor to any other technology. At best, the claimed combination amounts to an improvement to the abstract idea of determining patient risk rather than to any technology. See MPEP 2106.05(a). Thus, even when considering the elements in combination, the claim as a whole does not integrate the recited exception into a practical application. (*Step 2A, Prong Two: NO*). Thus, claim 1 is directed to a judicial exception. (*Step 2A: YES*).

Step 2B: This part of the eligibility analysis evaluates whether the claim as a whole amounts to significantly more than the recited exception, *i.e.*, whether any additional element, or combination of additional elements, adds an inventive concept to the claim. MPEP 2106.05. Additional elements in limitation (a) were considered insignificant extra-solution activities (mere data gathering) in Step 2A, Prong Two. The additional element of “an appropriate treatment” in limitation (c) does not require any particular application of the patient risk determination and is at most an instruction to “apply” the abstract idea.

These additional elements should be re-evaluated in Step 2B, in which the extra-solution activity consideration takes into account whether or not an extra-solution activity is well-known. The data gathering activities in limitation (a) are recited at a high level of generality and have been recognized by the courts as being routine laboratory techniques. See *Genetic Techs. v. Merial*

LLC, 818 F.3d 1369, 1377 (Fed. Cir. 2016) (analyzing DNA to provide sequence information or to detect allelic variants is conventional in the art); MPEP 2106.05(d), subsection II. The specification otherwise only describes carrying out sample collection and genotyping by conventional methods. See paragraph 4 of the Background.

Consequently, for the reasons discussed above, the additional elements individually or in combination with the judicial exception do not provide an inventive concept; so, the claim as a whole does not amount to significantly more than a generic instruction to “apply” the judicial exception. (*Step 2B: NO*). The claim is not eligible.

Claim 2 is eligible.

Claim Interpretation: Under the broadest reasonable interpretation, the terms of the claim are presumed to have their plain meaning consistent with the specification as it would be interpreted by one of ordinary skill in the art. See MPEP 2111. Claim 2 depends from claim 1 and adds a “wherein” clause specifying that the appropriate treatment is Compound X eye drops. As a dependent claim, claim 2 is construed as incorporating by reference all the limitations of the claim (claim 1) that it refers back to and further limits. See MPEP 608.01(n), subsection III. It is important to remember during claim interpretation that no limitations can be disregarded and the mere fact that a limitation appears in a wherein clause does not automatically mean that it is not given weight. Here, when the clause in limitation (c) is considered in view of the specification, it is clear that the wherein clause has patentable weight. Namely, the claim requires that the appropriate treatment be Compound X eye drops, and does not make the appropriate treatment optional or simply express the result of a process step. The claim does not require any particular dosage or frequency of administration.

Step 1: This part of the eligibility analysis evaluates whether the claim falls within any statutory category. MPEP 2106.03. The claim recites at least one step or act of determining the patient’s disease risk. Thus, the claim is a process, which is a statutory category of invention. (*Step 1: YES*).

Step 2A, Prong One: This part of the eligibility analysis evaluates whether the claim recites a judicial exception. As explained in MPEP 2106.04, subsection II, a claim “recites” a judicial exception when the judicial exception is “set forth” or “described” in the claim. Because claim 2 depends from claim 1, thereby incorporating by reference all the limitations of claim 1, it recites an abstract idea in limitation (b) for the reasons discussed above. As a result, it is necessary to continue to analysis under Step 2A, Prong Two. (*Step 2A, Prong One: YES*).

Step 2A, Prong Two: This part of the eligibility analysis evaluates whether the claim as a whole integrates the recited judicial exception into a practical application of the exception. This evaluation is performed by (1) identifying whether there are any additional elements recited in the claim beyond the judicial exception, and (2) evaluating those additional elements individually and in combination to determine whether the claim as a whole integrates the exception into a practical application. MPEP 2106.04(d).

In addition to the abstract idea in claim 1, from which claim 2 depends, claim 2 recites the additional element of “the appropriate treatment is Compound X eye drops.” Step 2A specifically

excludes consideration of whether the additional elements represent well-understood, routine, conventional activity. Instead, analysis of well-understood, routine, conventional activity is done in Step 2B. Therefore, the assessment below does not evaluate the conventionality of additional elements.

Under the broadest reasonable interpretation of the claim, additional element (c) and the wherein clause encompass the administration of a “particular treatment” when considered in the context of the claim as a whole. Specifically, the additional elements have more than a nominal relationship to the judicial exception because they use the abstract idea of determining patient risk of PI in a manner that meaningfully limits it. That is, the abstract idea is used to identify the patient as belonging to a specific patient population (glaucoma patients at high risk of PI), and the patient is then administered a treatment (Compound X eye drops instead of any common anti-fibrotic treatment, such as drug A, after microstent implant surgery) that is particular to that specific patient population (glaucoma patients at high risk of PI). Relying on the determination of patient risk to administer Compound X eye drops to glaucoma patients at high risk of PI after microstent implant surgery is therefore a particular treatment for a medical condition such that the claim as a whole integrates the judicial exception into a practical application. See MPEP 2106.04(d)(2). (*Step 2A, Prong Two: YES*). Therefore, the claim is not directed to a judicial exception. (*Step 2A: NO*). The claim is eligible.