Project Name:	AI Anomaly detection method	**Start Date:**	2024-01-02
Project Number:	2401	**Completion Date:**	2024-12-23

Project Details:

Scientific or Technological Objectives:

Measurement	Current Performance	Objective	Has results?
New ASIC chip design (synapses/micron)	10	20	Yes
ANN Methodology (x)	(not set)	(not set)	No
Malicious attack detection (%)	97	99.9	Yes

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] 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) for detecting one or more anomalies in a data set using the trained ANN;

[Claim 3] A method of using an artificial neural network (ANN) to detect malicious network packets dropping the one or more malicious network packets in real time; and blocking future traffic from the source address.

Field of Science/Technology:

Micro-electronics (2.02.03)

Project Details:

Intended Results: Develop new processes, Develop new materials, devices, or products, Improve existing processes
Work locations: Research Facility
Key Employees:
Evidence types:

Scientific or Technological Advancement:

Uncertainty #1: Circuit design HARDWARE

[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.

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

The most significant underlying key variables are:

materials, layouts, connection methods

Technology or Knowledge Base Level:

Activity #1-1: Circuit design Eligible for patent (Fiscal Year 2024)

Methods of experimentation:

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. 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.

[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.

Results:

- New ASIC chip design: 17 synapses/micron (70% of goal)

Conclusion:

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.

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.

The claim is eligible,

AUTHORS NOTE: THIS OUTLINES THE RELATIVE EASE IN DEFINING TECHNOLOGICAL ADVANCEMENT FOR HARDWARE ITEMS SUCH AS CIRUITS VS. AI PROCESSES WHERE THE "HUMAN" USER INPUT CAN NOT BE CLEARY DEFINED.

Significant variables addressed: connection methods, layouts, materials

Uncertainty #2: ANN design SOFTWARE / AI

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.

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.

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.

The most significant underlying key variables are:

detection techniques, variation sensitivity, boundary definition - ordinary vs anomalous

Technology or Knowledge Base Level:

Activity #2-1: General anomaly detection - Ineligible (Fiscal Year 2024)

Methods of experimentation:

[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.

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.

Results:

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.

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."

Conclusion:

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).

AUTHORS NOTE: THIS OUTLINES SOME POTENTIAL DOCUMENTATION ISSUES AND OPPORTUNITIES FOR SRED CLAIMS AS WELL AS PATENTS.

Activity #2-2: Packet method for security - Eligible (Fiscal Year 2024)

Methods of experimentation:

[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.

Results:

- Malicious attack detection: 99.5 % (86% of goal)

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

"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.

Conclusion:

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.

AUTHORS NOTE: THE PATENT RULES PROVIDE ELIGIBILITY AS A RESULT OF IMPROVING A "PRODUCT" (FUNCTIONING OF A COMPUTER) . IN THE CASE OF HARDWARE / SOFTWARE THIS MAY PROVIDE SOME LESSONS FOR SR&ED RELATED CLAIMS - I.E. FOCUS ON IMPROVED PRODUCTS VS PROCESSES?

Significant variables addressed: boundary definition - ordinary vs anomalous, detection techniques, variation sensitivity