EDITOR: Joshua J. Yi, josh@joshuayipatentlaw.com

DEPARTMENT: MICRO LAW

A Review of the *University of Tennessee Research Foundation v. Anthropic PBC*

Joshua J. Yi, *The Law Office of Joshua J. Yi, PLLC, Austin, TX, 78750, USA*

On 20 July 2026, the University of Tennessee Research Foundation ("UTRF") filed a lawsuit in the District of Delaware alleging patent infringement by Anthropic PBC. UTRF alleged that Anthropic infringed U.S. Patent Nos. 10,019,470 and 10,095,718, describing them as the result of "groundbreaking" research (see paragraph 1 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

The complaint describes that UTRF is the "non-profit organization responsible for protecting, licensing, and commercializing intellectual property developed by The University of Tennessee's faculty and researchers" (see paragraph 2 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). UTRF described itself as "among the oldest university research foundations in the United States" and that its mission is to "promote, support, and carry out the research mission of The University of Tennessee" (see paragraph 6 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). UTRF asserted that it owns the patents-in-suit, "including the right to recover for past infringement" (see paragraph 7 of Univ. of *Tenn. Rsch. Found. v. Anthropic PBC*[1]).

The complaint describes that Anthropic is a Delaware corporation (more why this is important below) with a principal place of business in San Francisco (see paragraph 8 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). UTRF describes that "Anthropic is organized as a Public Benefit Corporation; however, it is a for-profit corporation" (see paragraph 9 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). UTRF further describes that "in its short history, [Anthropic] has been no friend to the developers and owners of intellectual property rights" (see paragraph 9 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). More specifically, UTRF described that Judge William Alsup of the Northern District of California (where Anthropic has its principal place of business) concluded that

> Before buying books for its central library, Anthropic downloaded over seven million pirated copies of books, paid nothing, and kept these pirated copies in its library even after deciding it would not use them to train its AI (at all or ever again). Authors argue Anthropic should have paid for these pirated library copies. This order agrees. (See paragraph 9 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC.*[1])

UTRF described that "Anthropic has agreed to pay more than $1.5 billion to resolve the copyright claims at issue in that case" (see paragraph 9 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). UTRF alleged that "Anthropic's cavalier approach to others' intellectual property rights in the development of its products extends beyond the use of copyrighted material[,]" but also includes infringing UTRF's patent rights (see paragraph 10 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

UTRF argues that venue—which is the district where the lawsuit is filed—is proper because Anthropic is a Delaware corporation (see paragraph 4 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). Venue for patent cases is governed by §1400(b), which recites "Any civil action for patent infringement may be brought in the judicial district where the defendant resides, or where the defendant has committed acts of infringement and has a regular and established place of business." The judicial district where a defendant is incorporated in is considered to be the "judicial district where the defendant resides." Because Delaware is the judicial district where many corporations and over two-thirds of Fortune 500 companies are incorporated in, Delaware is a logical place for a plaintiff to file a patent infringement lawsuit against many companies.

UTRF describes that the problem that the patents were trying to solve was that while "[m]any AI systems utilize artificial neural networks[,]" "conventional neural network systems often relied upon fixed

0272-1732 © 2026 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies.
Digital Object Identifier 10.1109/MM.2026.3735782
Date of current version 22 October 2026.

architectures, manually designed network structures, and limited mechanisms for understanding, monitoring, or visualizing how neural networks arrived at particular outputs or decisions" and that there were "challenges in designing artificial neural networks capable of adapting to changing conditions, efficiently identifying useful neural substructures, reusing learned functionality, and improving performance without excessive computational cost or manual intervention" (see paragraphs 15, 17, and 18 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

UTRF described that the Asserted Patents "describe, among other things, techniques for constructing neuroscience-inspired artificial neural networks, dynamically adapting neural network architectures, identifying and reusing useful neural substructures, visualizing neural pathways and network activity, monitoring neural network operation, implementing biologically inspired learning mechanisms, and improving the performance and operation of artificial intelligence systems" and that the "patented technologies addressed technical problems associated with the construction, operation, training, monitoring, explainability, and improvement of artificial neural networks years before artificial intelligence systems became widely deployed in commercial products and services" (see paragraphs 25 and 26 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

The '470 Patent is titled "Method and Apparatus for Constructing, Using and Reusing Components and Structures of an Artificial Neural Network" (see paragraph 28 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). UTRF described that this patent had been cited several hundred times (which is a very large number of citations) in U.S. and international patents by companies such as Apple, Google, Microsoft, Amazon, Samsung, IBM, Facebook, Sony, Intel, LG, GE, Siemens, AT&T, Cisco, Qualcomm, Micron, HP Enterprise, and Motorola Solutions, as well as several universities (see paragraph 30 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

The '470 Patent describes the invention as follows:

> In accordance with an embodiment of a method and apparatus for constructing, using and reusing components and structures of an artificial neural network, a neuroscience-inspired dynamic artificial neural network (NIDA or a dynamic adaptive neural network array (DANNA) of basic neuron and synapse elements or combinations of substructures thereof may be constructed and then utilized to identify useful components, structure and substructures. A useful substructure of an artificial neural network, by way of example, may be an affective system, a learning algorithm, a computational network, a central pattern generator, an image recognizer and the like. For example, in a classification task such as a character recognizer which generally recognizes all characters, an artificial neural network may have implanted therein a numeric digit recognizer. Also, as artificial neural networks are constructed for tasks such as one of control, detection and classification, some network substructures may prove to be unsuccessful for performing certain sub-tasks and so extracted from a network. Some neural pathways may never be used in a special purpose artificial neural network and so, when isolated or identified as such, they may be saved as a potential useful substructure and, if recognized in a given artificial neural network for the same task, be extracted as unsuccessful. It should be noted, however, that useful substructures may cause either excitation or inhibition of activity elsewhere in a network or in a second network with which the structure or first network can communicate. It is therefore preferable to identify both substructures that induce a response and structures that inhibit a response as possibly useful for use, reuse and construction of artificial networks.

UTRF alleged that "Anthropic designs, makes, uses, sells, and/or offers for sale in the United States products and/or services for constructing a central pattern generator of a neuromorphic network for use in a process control application of an external automated process" (see paragraph 35 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). UTRF further alleged that "Anthropic designs, makes, sells, offers to sell, imports, and/or uses Anthropic's Claude Code and its underlying agentic software architecture" (see paragraph 36 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). UTRF refers to these products as these "Anthropic '470 Product(s)" (see paragraph 36 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

UTRF then provided a description of how some parts of the Anthropic '470 Product(s) operate. For example, UTRF describes that the Anthropic '470 Product(s) contain a "background memory consolidation engine (autoDream.ts), which automatically manages, compresses, and restructures memory states in a recurring, cyclic process to control local terminal environments or remote container sandboxes (acting as the external automated system or external device)" (see paragraph 38 of *Univ. of Tenn. Rsch. Found. v.*

Anthropic PBC[1]). UTRF describes that "[d]reaming is a process that looks for patterns and mistakes across your recent agent sessions and their transcripts and automatically produces organized and up-to-date memory content" (see paragraph 39 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

UTRF then described how Anthropic '470 Products meet the elements of some of the claims, in particular Claim 10. Due to the length of Claim 10—13 limitations and 402 words—this article will not summarize all of UTRF's claim mapping allegations, but will only focus on select claim elements: in particular, Limitations [c], [d], [i], [j], [k], and [l]. With respect to Claim 10, Limitation [c], UTRF alleges that

> Anthropic '470 Products "comprise an addressably configurable circuit element for providing inputs or outputs to components of the external automated system or external device comprising one of a display, a camera, a radio, and a scanner device external to the neuromorphic network, the neuromorphic network receiving input data from and transmitting output data to the external automated system or external device. Specifically, the Anthropic '470 Products contain addressable Model Context Protocol (MCP) servers, tool wrappers, and host adapters (hostAdapter.ts). These modules provide addressable inputs and outputs to external hardware and OS-level devices. (See paragraph 51 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC.*[1])

With respect to Claim 10, Limitation [d], UTRF alleges that

> The Anthropic '470 Products perform the step of addressably configuring one synapse circuit element to connect one addressably configured neuron circuit element to another addressably configured neuron circuit element to form an initial addressably configured array comprising an initial two-dimensional neuromorphic network. Specifically, in the Anthropic '470 Products, synapses (connections between neuron nodes) are implemented as at least the following: parent-child tree node references, JSON-RPC channels, and structural parent-child connections. In the Anthropic '470 Products, layout nodes are connected and bound via insertChild and removeChild, which form an initial, addressable two-dimensional network grid of nodes. The layout engine (engine.ts) in the Anthropic '470 Products sets up and forms these connection nodes. (See paragraph 55 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC.*[1])

With respect to Claim 10, Limitation [i], UTRF alleges that

> The Anthropic '470 Products comprise a structure of components of the constructed central pattern generator. Specifically, the structure of the central pattern generator is defined by the programmatic orchestration of the task scheduler, background daemon configurations, and concurrent locking mechanisms working in tandem to maintain safety and consistency in background processing. This is confirmed by the structure of the source code of the Anthropic '470 Products, which shows components such as autoDream, cronScheduler, and consolidationLock within a single directory. (See paragraph 62 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC.*[1])

With respect to Claim 10, Limitation [j], UTRF alleges that

> The Anthropic '470 Products comprise a neuron of the neuron and synapse pair of addressably configurable circuit elements having a programmable threshold and a programmable refractory period. Specifically, the Anthropic '470 Products implement "programmable thresholds" and "programmable refractory periods" inside their memory consolidation and polling components. These parameters act as temporal and volume gating rules. The memory consolidation trigger (autoDream.ts) applies programmable thresholds (e.g., requiring 24 elapsed hours and 5 accumulated sessions) before permitting task execution. Furthermore, a refractory period (cooldown scan interval and file lock expiry) is enforced to throttle and prevent rapid overlapping executions, safeguarding system execution against feedback loops. Further, the Anthropic '470 Products contain a file lock utility (consolidationLock.ts) that implements a PIDbacked refractory period to prevent concurrent execution conflicts, enforcing a strict timeout parameter before the lock can be recycled. (See paragraph 63 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC.*[1])

With respect to Claim 10, Limitation [k], UTRF alleges that

> The Anthropic '470 Products comprise a synapse of the neuron and synapse pair of

> addressably configurable circuit elements having a programmable delay or distance parameter value. Specifically, programmable delay parameter values are implemented as jittered backoff intervals, network connection retries, and tool timeouts. In command scope analysis, a programmable "distance" parameter maps to scope-level tracking constraints. This resolves shell variable propagation across sequential commands (varScope) but isolates and copies scope (strictly limiting the propagation distance) across nested pipelines or subshells. Further, the network connection interface in the Anthropic '470 Products (withRetry.ts) implements connection retries as an asynchronous generator that uses custom retry logic. (See paragraph 64 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*.[1])

With respect to Claim 10, Limitation [l], UTRF alleges that

> Responsive to a neuron that is addressably configurable, the Anthropic '470 Products perform the step of generating a sequence of discrete events, the addressably configurable synapse permitting at least two discrete events to propagate together through the synapse. Specifically, the Anthropic '470 Products' synapse (implemented as the communication, stream, or parsing channel) permits multiple discrete events (such as stdout/stderr streams, parallel tool uses, or piped subcommands) to propagate and resolve together in parallel or as concatenated payloads. Further, the syntax utility in the Anthropic '470 Products (ast.ts) processes concatenation nodes to merge and pass multiple distinct arguments together. (See paragraph 65 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*.[1])

The '718 Patent is titled "Method and Apparatus for Constructing a Dynamic Adaptive Neural Network Array (DANNA)" (see paragraph 31 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). UTRF described that this patent also has been cited several hundred times in U.S. and international patents by many of the same companies listed above (see paragraph 33 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

The '718 Patent describes how the invention improves on prior art as follows:

> In accordance with an embodiment of a method and apparatus for constructing a neuroscience-inspired artificial neural network in the form of a dynamic adaptive neural network array (DANNA) of basic neuron and synapse elements that may be constructed of field programmable gate arrays that overcomes the deficiencies of the prior art, the embodiment exhibits five characteristics which differentiate over the prior art. Firstly, it is desirable that a neural network structure evolve over time. Evolution over time means that the learning algorithms of the prior art may not evolve in accordance with data (events) received at input neurons compared with the evolution achieved in accordance with the present invention and the present learning algorithms disclosed herein. Secondly, it is desirable that neural networks may be embedded into a geometric space. This characteristic suggests that the present invention seeks confinement, for example, to a geometric space in a similar manner that it is believed that the human brain and its sensory inputs and outputs are confined to a geometric space. Thirdly, it is desirable that neural networks compose dynamic elements and operate on a continuous time scale. In some embodiments, a discrete time scale may be preferred, for example, for digital or discrete time unit measurement. By dynamic elements is intended the opposite of static elements such as dynamic neurons and dynamic synapses. Also, the concept of continuous time scale means an intention to differentiate from a discrete time scale or one capable of one input at a time, such as one per clock period when the clock period may be increased or decreased in length. Clearly, it is desirable if an artificial neural network is able to receive two inputs at the same time, that is, all inputs are continuously occurring and the network is continuously learning and making decisions with the expectation that the artificial neural network will adapt to its environment. Fourthly, it is desirable if useful substructures in neural network structures can be recognized and reused. That is, for example, the present invention is capable of identifying in a visualization those structures that are acting in response to a given input or continuous series of inputs. Then, theoretically, these same structures that have been identified may be reused to respond to a similar set of inputs. Fifthly, it is desirable if special-purpose emotion-related substructures and neurotransmitters can be incorporated into artificial neural

> networks. As suggested above, emotions such as fear or anger have been artificially simulated in the prior art individually but not collectively as to the collection of many emotion-related substructures.

For the '718 Patent, UTRF refers to the Anthropic '470 Product(s) as "Anthropic '718 Product(s)" (see paragraph 78 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

UTRF provides a description of how some parts of the Anthropic '718 Product(s) operate. For example, UTRF describes that the Anthropic '718 Product(s) "construct a multi-agent system representing a software-based neuromorphic network," "which represents an addressably reconfigurable multi-dimensional array of weights (interconnected circuit elements) in a high-dim ensional vector/token embedding space" (see paragraph 81 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]). UTRF describes that "Anthropic has described the Anthropic '718 Products as comprising a neural network that contains neurons (features), weights, and circuits (pathways connecting features)" and where "Anthropic '718 Products represent concepts via linear combinations of the network's underlying biological/artificial neurons" (see paragraphs 83 and 84 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

UTRF then described how Anthropic '718 Products meets the elements of some of the claims, in particular Claim 35. With respect to Claim 35, Limitation [a], UTRF alleges that the Anthropic '718 Products "perform the method using the special purpose computer processor" and that they "comprise an artificial neural network configuration structure," where "the special purpose computer being responsive to input data from the external process for addressably configuring the multi-dimensional array" (see paragraphs 81 and 82 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

With respect to Claim 35, Limitation [b], UTRF alleges that the "Anthropic '718 Products contain input neuron elements that are endpoints of the network's weights" (see paragraph 93 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

With respect to Claim 35, Limitation [c], UTRF alleges that the "the output generated by the Anthropic '718 Products (tokens, text block completions, and tool use JSON schemas) represents output neurons" and that "[t]he Anthropic '718 Products contain an output neuron element as the terminal point of the network's weight array" (see paragraphs 94 and 95 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*[1]).

With respect to Claim 35, Limitation [d], UTRF alleges that

> The Anthropic '718 Products perform, under a special purpose program control, one to multiple circuit elements of the addressably configurable multi-dimensional array addressably configured as a neuron or a synapse forming a pathway between the input neuron receiving input data from the external process and the output neuron outputting data to the external process. Specifically, in coordinator mode, the Anthropic '718 Products configure pathways using multiple agents and tool execution loops. The Anthropic '718 Products contain a coordinator that spawns workers via AgentTool (neurons) that communicate via XML <task-notification> (synapses). Furthermore, toolExecution.ts in the Anthropic '718 Products implement self-correcting feedback loops (if a tool execution fails or returns an error, it is caught in a catch block and fed back as a new input to reconfigure the LLM's attention pathway). (See paragraph 97 of *Univ. of Tenn. Rsch. Found. v. Anthropic PBC*.[1])

Unlike other complaints that clearly tie features of the accused products to the claim language, this complaint does not make those ties as explicitly clear.

With respect to the parties' patent litigation history, UTRF has filed approximately a dozen patent infringement lawsuits in the past. In May 2017, UTRF filed nine patent infringement lawsuits against SAP America, Oracle, HP Enterprise, Amazon, Teradata, IBM, Cloudera, and Citrix. All of these lawsuits were voluntarily dismissed in October 2017, which means that the parties in those cases reached a settlement agreement.

Anthropic has been sued by another patent holder, who represented himself without a lawyer. That case was ultimately dismissed for failure to prosecute.

REFERENCE

1. Univ. of Tenn. Rsch. Found. v. Anthropic PBC, No. 1:26-cv-00887, ECF No. 1, Jul. 20, 2026, (D. Del.).

JOSHUA J. YI is a solo practitioner at The Law Office of Joshua J. Yi, PLLC, Austin, TX, 78750, USA. Contact him at josh@joshuayipatentlaw.com.