

**IN THE UNITED STATES DISTRICT COURT FOR
THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

MANY WORLDS 2T INNOVATIONS LLC, Plaintiff, v. OPENAI OPCO, LLC, Defendant.	CIVIL ACTION NO. 2:26-cv-00774 Jury Trial Demanded
--	--

ORIGINAL COMPLAINT FOR PATENT INFRINGEMENT

This is an action for patent infringement where Plaintiff Many Worlds 2T Innovations LLC (“Many Worlds”) makes the following allegations against OpenAI OpCo, LLC (“OpenAI” or “Defendant”).

BACKGROUND

1. OpenAI’s artificial intelligence tools directly implement the technological solutions claimed in the patents asserted in this action, defined below as the Many Worlds Asserted Patents. OpenAI’s system uses dual neural networks—one processing user behavioral data, preferences, and customizations and another processing content characteristics—that generate weighted embeddings subsequently combined to produce content for a user. This architecture mirrors the patented fuzzy network approach that analyzes both user behavior and content attributes through weighted relationship structures. The Many Worlds inventors recognized that prior recommendation engines suffered from critical technical deficiencies: they treated recommendation generation as a black box providing no transparency, and they ignored content

characteristics by relying solely on collaborative filtering of user behavior patterns. The Many Worlds Asserted Patents solved these problems through specific computer-implemented technical architectures including fuzzy network data structures with weighted affinity relationships, contextual neighborhood algorithms for graph traversal, dual-analysis pipelines combining behavioral and content-based inference, vector-based search using behavioral embeddings, automated natural language generation for explanations, and contrast-identification algorithms enabling serendipitous discovery.

PARTIES

2. Plaintiff Many Worlds 2T Innovations LLC is a Texas limited liability company whose principal place of business is 4802 Parkview Place, Austin, TX 78731.

3. Defendant OpenAI OpCo, LLC is a limited liability company organized and existing under the laws of the State of Delaware. OpenAI registered as a foreign entity in the State of Texas effective September 26, 2025. OpenAI may be served with process through its registered agent in Texas, CT Corporation System, 1999 Bryan Street, Suite 900, Dallas, Texas 75201.

JURISDICTION AND VENUE

4. This is an action for infringement of United States patents arising under 35 U.S.C. §§ 271, 281, 283, 284, and 285. This Court has subject matter jurisdiction of the action under 28 U.S.C. § 1331 and § 1338(a).

5. This Court has personal jurisdiction over OpenAI consistent with the Texas long-arm statute and the Due Process Clause. OpenAI is registered to do business in Texas and maintains a registered agent for service of process in Texas. On information and belief, OpenAI has purposefully directed activities at Texas and at residents of this District, including by offering, selling, and delivering subscriptions to the Accused Instrumentalities (defined below) to residents

of Texas and of this District and deriving substantial revenue from those sales; by contracting for and using dedicated compute capacity at the artificial-intelligence data-center campus in Denton County in this District described below (the “Denton Campus”); and by committing acts of patent infringement in Texas and in this District as alleged herein. Many Worlds’ claims arise out of and relate to OpenAI’s contacts with this forum.

6. Venue is proper in this judicial district under 28 U.S.C. § 1400(b) because OpenAI has committed acts of infringement in this District and has a regular and established place of business in this District.

7. OpenAI has a regular and established place of business in this District at the artificial-intelligence data-center campus located at 8171 Jim Christal Road, Denton, Denton County, Texas. Denton County is within the Eastern District of Texas. See 28 U.S.C. § 124(c)(3). Because OpenAI’s place of business and acts of infringement are located in Denton County, venue is proper in this District.

8. The Denton Campus is a fixed, physical place in this District—a permanent, purpose-built facility, not a virtual location. It is an approximately 394-megawatt artificial-intelligence compute campus at which, on information and belief, more than 16,000 graphics processing units were operational and running OpenAI workloads by the end of December 2025. The Denton Campus operates continuously on a steady, uniform, orderly, and methodical basis, and is therefore a regular and established place of business.

9. On information and belief, the Denton Campus is a place of OpenAI’s own business. OpenAI holds the exclusive use of dedicated compute capacity at the Denton Campus under a Master Services Agreement and related reserved-instance order forms, supported by an aggregate commitment of approximately \$22.4 billion and a dedicated financing vehicle established to acquire the hardware deployed there to deliver services to OpenAI. OpenAI’s contractual rights include the

exclusive use of the reserved capacity, a restriction barring suspension of OpenAI's access to the capacity, and the right to approve subcontractors performing core aspects of the services at the Denton Campus. Through these arrangements, OpenAI has established and ratified the Denton Campus, and holds it out, as a place from which OpenAI conducts its business.

10. On information and belief, the business of OpenAI is conducted at the Denton Campus through OpenAI and its agents. The personnel who install, configure, provision, monitor, maintain, and operate the graphics processing units and supporting infrastructure dedicated to OpenAI at the Denton Campus do so for OpenAI's benefit and subject to OpenAI's direction and control, including through OpenAI's contractual right to approve the subcontractors performing core aspects of the services delivered there and OpenAI's contractual right to prevent suspension of its access to the reserved capacity. On information and belief, OpenAI's own personnel further direct, provision, schedule, and monitor OpenAI workloads executing on the dedicated capacity at the Denton Campus. OpenAI has accepted and continues to accept the benefits of that work with knowledge of it, and has thereby established and ratified the Denton Campus as a regular and established place of OpenAI's own business.

11. On information and belief, OpenAI maintains a regular physical presence at the Denton Campus through its employees and agents. On information and belief, OpenAI's own personnel attend the Denton Campus to inspect, commission, accept, benchmark, and audit the compute capacity reserved for OpenAI, and to direct the provisioning and operation of OpenAI workloads executing on that capacity. On information and belief, the personnel who install, configure, provision, monitor, maintain, and operate that capacity act as OpenAI's agents in doing so: they perform that work for OpenAI's benefit, OpenAI holds the right to direct and control the manner in which the core aspects of that work are performed—including through its contractual right to approve the subcontractors who perform those core aspects and its contractual right to

prevent suspension of its access to the reserved capacity—and OpenAI has consented to and accepted the benefit of that work with knowledge of it. *See In re Cray Inc.*, 871 F.3d 1355 (Fed. Cir. 2017); *In re Google LLC*, 949 F.3d 1338 (Fed. Cir. 2020); *Andra Group, LP v. Victoria's Secret Stores, L.L.C.*, 6 F.4th 1283 (Fed. Cir. 2021).

12. Furthermore, on information and belief, OpenAI employs individuals who reside or work in this District—including at least two Infrastructure Engineers who work at OpenAI's Denton Campus.

13. OpenAI has committed acts of infringement in this District. On information and belief, OpenAI uses the Accused Instrumentalities within this District by operating them on the dedicated compute capacity at the Denton Campus, where the accused systems are put into service as a whole and where OpenAI exercises control over, and obtains the beneficial use of, those systems. OpenAI's infringing use therefore occurs, at least in part, within this District.

14. On information and belief, OpenAI has also committed acts of infringement in this District by offering to sell and selling subscriptions to the Accused Instrumentalities to residents of this District, and end users located in this District use the Accused Instrumentalities, which use OpenAI actively induces as alleged herein.

THE MANY WORLDS PATENTS

15. U.S. Patent No. 8,676,742 (the '742 Patent) was filed on October 7, 2011 and issued on March 18, 2014.

16. U.S. Patent No. 10,699,202 (the '202 Patent) was filed on August 27, 2016 and issued on June 30, 2020.

17. U.S. Patent No. 12,307,388 (the '388 Patent) was filed on June 26, 2024 and issued on May 20, 2025.

18. U.S. Patent No. 12,299,603 (the '603 Patent) was filed on June 26, 2024 and issued on May 13, 2025.

19. U.S. Patent No. 8,843,433 (the '433 Patent) was filed on July 23, 2012 and issued on September 23, 2014.

20. Collectively, the '742 Patent, the '202 Patent, the '388 Patent, the '603 Patent, and the '433 Patent are the "Many Worlds Asserted Patents."

21. The Many Worlds Asserted Patents claim priority to: (i) U.S. Provisional Application No. 61/469,052, filed March 29, 2011; (ii) U.S. Provisional Application No. 61/496,025, filed June 12, 2011; and (iii) U.S. Provisional Application No. 61/513,920, filed August 1, 2011.

22. Each asserted claim of each of the Many Worlds Asserted Patents is entitled to the priority date of the provisional applications identified above. The disclosures of those provisional applications describe and support the subject matter of each asserted claim, and each of the Many Worlds Asserted Patents issued from a continuous chain of co-pending applications extending back to those provisional applications. The state of the art relevant to each asserted claim is therefore the state of the art as of March 29, 2011.

23. Many Worlds owns all substantial rights to the Many Worlds Asserted Patents, including the right to sue and recover damages for all infringement thereof including past damages.


24. Many Worlds and any prior owner of the Many Worlds Asserted Patents have neither made nor sold unmarked articles that practice any claim of the Many Worlds Asserted Patents. Many Worlds is therefore entitled to collect pre-filing damages for the full period allowed by law for infringement of the Many Worlds Asserted Patents.

MANY WORLDS' TECHNOLOGICAL INNOVATIONS

25. At the time of the Many Worlds Asserted Patents' inventions, recommendation engines were largely limited to observing a user's past behavior and then attempting to predict what the user may like in the future based on the user's past choices. Further, a feature that prior recommendation engines lacked was the ability to explain why a user received a recommendation. This gap deprived the user of valuable information that may give context to a specific recommendation from a recommendation engine and thus improve the user experience.

26. Prior to the Many Worlds Asserted Patents, recommendation engines predominantly used collaborative filtering: analyzing which users consumed which content, grouping users with similar consumption patterns, and recommending to each user what similar users consumed. This approach had critical technical limitations that degraded recommendation quality. First, pure collaborative filtering is content-agnostic. It treats all content items as opaque identifiers without analyzing what the content actually contains. The system might know that users who liked item #47 also liked item #93, but it has no understanding whether these are short stories or long novels, jazz or rock music, or comedy or drama podcasts. Second, this content-blind approach creates cold-start problems: new content with few user interactions cannot be effectively recommended even if its attributes would appeal to many users. Third, the black-box nature provides no transparency, as users receive recommendations with no explanation of why the system believes they will enjoy the content, reducing trust and engagement. These were not merely business problems but specific technical deficiencies in how computer systems processed and utilized recommendation data.

27. Figure 3 of the '742 Patent presents an example fuzzy network that includes content objects, topic objects, and links between the objects.



28. The '742 Patent discloses that “[t]opic objects 710t are encapsulations that contain meta-information 36c and relationships to other objects (not shown), but do not contain an embedded pointer to reference associated information. The topic object 710t thus essentially operates as a ‘label’ to a class of information.”¹ With respect to content objects, the '742 Patent states:


¹ '742 patent, 6:64-7:2.

Content objects 710c, as shown in FIG. 4C, are encapsulations that optionally contain meta-information 36c and relationships to other objects 710 (not shown). Additionally, content objects 710c may include either an embedded pointer to information or the information 714 itself (hereinafter, “information 714”).²

29. The fuzzy network’s weighted affinity relationships constitute a specific technological tool that improves computer system performance. Unlike conventional binary database relationships that store only whether items are connected, the fuzzy network maintains numerical weights (typically normalized between 0 and 1) indicating relationship strength, enabling probabilistic recommendations that traditional binary structures cannot produce. These weighted affinities are computed and updated automatically from observed user behavior across large user populations, which requires machine processing of behavioral data to generate affinity scores, storage of multidimensional weight matrices, and retrieval of those weights during recommendation generation.

² ’742 patent, 7:6-12.

Figure 5C



30. The claims of the '742 Patent describe how a system using fuzzy networks provides the user a better recommendation than prior systems. Claim 1 of the '742 Patent reads:

1. A computer-implemented method, comprising:

invoking a contextual scope function executed on a processor-based computing device, wherein the contextual scope function establishes a contextual neighborhood within a computer-implemented fuzzy network-based structure;

receiving a recommendation, wherein the recommendation is generated by a computer-implemented discovery function that generates the recommendation in accordance with the contextual neighborhood and an inference from a plurality of usage behaviors; and

receiving the recommendation, wherein the recommendation is generated in accordance with the contextual neighborhood, wherein the contextual neighborhood is based on a selected object.

The '742 Patent also claims the system that carries out that architecture and further requires that the breadth of the contextual neighborhood be set by a user. Claim 8 of the '742 Patent reads:

8. A computer-implemented system, comprising:

a contextual scope function executed on a processor-based computing device that establishes a contextual neighborhood within a computer-implemented fuzzy network-based structure;

a contextual neighborhood scoping function that establishes the scope of the contextual neighborhood at a level set by a user; and

a computer-implemented discovery function that generates a recommendation in accordance with the contextual neighborhood and an inference from a plurality of usage behaviors.

31. Claim 8 requires that the scope of the contextual neighborhood be established “at a level set by a user.” The breadth of the neighborhood is therefore not fixed by the system designer; it is a parameter exposed to and set by the user, and the recommendation that the discovery function generates depends on the neighborhood the user has bounded.

32. The claims of the '742 Patent require specific technological components and operations that improve computer-implemented recommendation systems. Claim 1 recites three interconnected technical requirements: (1) invoking a contextual scope function executed on a processor-based computing device, wherein the contextual scope function establishes a contextual neighborhood within a fuzzy network-based structure; (2) receiving a recommendation generated by a computer-implemented discovery function that operates on the contextual neighborhood together with an inference from a plurality of usage behaviors; and (3) basing the contextual

neighborhood on a selected object. Claim 8 recites the corresponding system and adds the requirement that the scope of the contextual neighborhood be established at a level set by a user, so that the breadth of the neighborhood from which the recommendation is drawn is a parameter under the user's control. These claims specify concrete computer operations manipulating a specific data structure through specific functions to improve the relevance of computer-generated recommendations.

33. As another example of the technological innovations claimed by the Many Worlds Asserted Patents, the '202 Patent addresses the long-standing and pervasive problem of a user not having context for a recommendation. The prior art could leave a user baffled as to why a recommendation was made. The Many Worlds Asserted Patents addressed this deficiency directly. Claim 1 of the '202 Patent reads:

1. A computer-implemented method, comprising:

accessing automatically a first plurality of values, wherein each of the first plurality of values is determined by automatically analyzing text that is associated with a content object, wherein each of the plurality of values is based upon an automatically inferred degree of relevancy between an element of the text and the content object;

accessing automatically a second plurality of values, wherein each of the second plurality of values corresponds to one or more topics of a plurality of topics, and wherein a magnitude of each of the second plurality of values is based upon an automatically determined inference from one or more user behaviors; and

generating automatically a communication for delivery to a user of a computer-implemented system, wherein the communication comprises a plurality of words, wherein the plurality of words are selected in accordance with the first plurality of values, the second plurality of values, and one or more syntactical rules.

34. The '202 Patent's claim 1 specifies technical operations for generating explanations: (1) automatically analyzing text associated with content objects through natural language processing algorithms to extract semantic elements and compute relevancy scores; (2)

automatically accessing user-topic values with magnitudes determined through behavioral inference algorithms; and (3) automatically generating communications by computationally selecting words based on both value sets plus syntactical rules—requiring algorithmic natural language generation that constructs grammatically correct explanatory phrases matching high-relevance content elements to high-affinity user topics. This automated explanation generation solves the prior art’s “black box” problem through specific technical means: coupling content analysis outputs with user profile outputs through a natural language generation engine that applies syntactic constraints. The claim does not merely require “explaining” recommendations (a result) but specifies the technical pipeline by which explanations are computationally constructed from analyzed data.

35. As another example of the technological innovations claimed by the Many Worlds Asserted Patents, the ’388 Patent addresses the long-standing and pervasive problem of a user not having the ability to control the serendipity of a recommendation. In prior art systems, the user had little or no control over the amount of serendipity of a result they could receive. The Many Worlds Asserted Patents addressed this deficiency directly. Claim 1 of the ’388 Patent reads:

1. A communication tuning method performed by a computer-implemented system comprising:

receiving information from a user;

providing a serendipity tuning control to the user;

receiving from the serendipity tuning control, responsive to the user’s interaction with the serendipity tuning control, a serendipity tuning setting;

applying the serendipity tuning setting to a probabilistic selection algorithm;

selecting, responsive to the received information, a first syntactical element from a plurality of syntactical elements in accordance with the probabilistic selection algorithm;

selecting a second syntactical element from the plurality of syntactical elements in accordance with the first syntactical element and the probabilistic selection algorithm;

generating a communication comprising the first syntactical element and the second syntactical element, wherein the communication comprises a theme that is derived from a vector representation of a content object; and

delivering the communication to the user.

36. The '388 Patent's claim 1 solves the problem of too-predictable responses from a system through specific technical operations on affinity vector data structures: (1) receiving input from a user; (2) providing the user with a serendipity control and receiving a serendipity tuning setting from that control; (3) generating a first syntactical element in accordance with the probabilistic selection algorithm; (4) generating a second syntactical element based on the first syntactical element and the probabilistic selection algorithm; and (5) generating and delivering a communication to the user based on the first syntactical element and the second syntactical element. This probabilistic communication generation system requires specific vector operations: probabilistic generation of multiple syntactical elements and a vector-represented content object. A human cannot use a serendipity control to generate syntactical elements based on a probabilistic selection algorithm—this necessarily requires computer implementation using vector processing capabilities.

37. As another example of the technological innovations claimed by the Many Worlds Asserted Patents, the '603 Patent addresses the problem of search systems that retrieve by matching words rather than by comparing meaning. Claim 1 of the '603 Patent reads:

1. A computer-implemented search method comprising:

generating by applying one or more computer-implemented neural networks to each of a plurality of content objects at least one distinct affinity value that is associated with each of the plurality of content objects, wherein

each of the at least one distinct affinity values represents a degree of relationship with the associated content object;

generating a plurality of affinity vectors within a multi-dimensional affinity space, wherein each of the plurality of affinity vectors corresponds to one of the plurality of content objects and comprises the at least one distinct affinity value associated with the one of the plurality of content objects;

receiving information from a user's input into a search interface;

generating by applying a computer-implemented neural network one or more affinity values that each represent a degree of relationship associated with the received information;

generating an affinity vector comprising the one or more affinity values associated with the received information;

comparing using a mathematical-based vector comparison algorithm the affinity vector associated with the received information and each of the plurality of affinity vectors;

selecting one or more of the plurality of affinity vectors in accordance with the comparing of the affinity vector associated with the received information and the each of the plurality of affinity vectors; and

providing to the user one or more of the plurality of content objects that each correspond to one of the selected one or more of the plurality of affinity vectors.

38. Claim 1 of the '603 Patent specifies an ordered sequence of machine operations: a neural network is applied to each stored content object to generate affinity values; those values are assembled into per-object affinity vectors within a multi-dimensional affinity space; the user's search input is run through a neural network to generate its own affinity vector; the query vector is compared against the stored vectors using a mathematical vector-comparison algorithm; vectors are selected on the basis of that comparison; and the corresponding content objects are returned to the user. The claim recites how the search is performed, not merely that content is retrieved.

39. As a further example of the technological innovations claimed by the Many Worlds Asserted Patents, the '433 Patent addresses the problem that prior recommenders drew on either

what a user had done or what content was about, but not both within a single recommendation function. Claim 1 of the '433 Patent reads:

1. A computer-implemented method, comprising:
 - identifying a computer-implemented content object to serve as a contextual reference;
 - accessing a first plurality of values, wherein each of the first plurality of values is based on an inference of a first user's interest level with respect to a topic, wherein each of the interest level inferences is derived from a plurality of behaviors;
 - accessing a second plurality of values, wherein each of the second plurality of values is based on an inferred degree of relevance to a topic of one or more phrases that are contained in the content object; and
 - generating a recommendation comprising one or more objects for delivery to the first user, wherein the recommendation is generated by a recommendation function executed on a processor-based device in accordance with the first plurality of values and the second plurality of values.

40. Claim 1 of the '433 Patent recites a specific two-input architecture: a first set of values inferring the user's interest in a topic from a plurality of the user's behaviors; a second set of values inferring how relevant the phrases within a content object are to a topic; and a single recommendation function, executed on a processor-based device, that generates the recommendation in accordance with both sets. A system that uses one signal or the other does not satisfy the claim. The fusion of the two within one function is the claimed advance.

41. The Many Worlds Asserted Patents do not preempt the fields in which they operate. Each claim is confined to the particular architecture it recites, and numerous alternative approaches to recommendation, search, and communication generation remained available then and remain available now without practicing any claim: collaborative filtering that does not analyze content characteristics; keyword, Boolean, and inverted-index search that does not encode stored content and queries into a shared affinity space; retrieval scopes fixed by the system designer rather than set

by a user; deterministic or greedy text generation that exposes no user-facing tuning control over the selection of successive elements; and template-driven output that does not select words under content-derived relevancy values, behavior-derived topic values, and syntactical rules together. The claims therefore leave the underlying concepts available to others by other means.

42. The specifications of the Many Worlds Asserted Patents identify the specific technical shortcomings of the systems that preceded them and describe how the claimed architectures overcome those shortcomings. Many Worlds alleges as fact that, as of the priority dates of the Many Worlds Asserted Patents, the conventional approaches in the field were these: hierarchical and relational data structures in which an object was joined to others by single-parent or binary links; collaborative filtering that grouped users by overlapping consumption history and treated content items as opaque identifiers without analyzing their contents; keyword, Boolean, and inverted-index matching for search; template-driven or hand-authored text output; and retrieval scopes fixed by the system designer rather than exposed to the user. The claimed architectures depart from each of those conventions in the specific respects recited in the claims.

43. The Many Worlds Asserted Patents do not claim the application of generic or previously established techniques to a new data environment. The claims are not confined to any field of use, industry, or category of subject matter, and they do not recite a desired outcome to be produced by unspecified means. They instead delineate the computational structure by which the outcome is produced: the graded relationship structure and the user-set boundary of the neighborhood computed within it; the assembly of per-object values into vectors within a shared multi-dimensional affinity space and the mathematical comparison carried out in that space; the tuning setting applied to the probabilistic selection of each successive syntactical element; and the three separately computed constraint sets under which the words of a communication are chosen.

The specifications describe each of those structures and explain how they differ from and improve upon what came before.

44. The claimed inventions are not the automation of tasks that people previously performed by hand, and their advantage over the prior art is not merely that a computer performs a known task faster. There was no antecedent human practice of maintaining multiple simultaneous graded relationships between the same pair of objects and computing a bounded neighborhood within that structure at a breadth set by a user, of encoding stored content and a user query into a shared multi-dimensional affinity space and retrieving content by mathematical comparison of the resulting vectors, or of governing the probabilistic selection of each successive element of a generated communication through a control exposed to the user. Each of these operations exists only in, and because of, the computer architectures the claims recite.

45. Each allegation in this section concerning the state of the art, the conventionality of the claimed elements and of their ordered combinations, and the technical improvements the claimed architectures achieve is made as an allegation of fact, is supported by the specifications of the Many Worlds Asserted Patents, and is to be taken as true and construed in Many Worlds' favor at the pleading stage. Whether a claim element or an ordered combination of elements was well-understood, routine, and conventional is a question of fact that cannot be resolved against Many Worlds on the pleadings. *Berkheimer v. HP Inc.*, 881 F.3d 1360 (Fed. Cir. 2018); *Aatrix Software, Inc. v. Green Shades Software, Inc.*, 882 F.3d 1121 (Fed. Cir. 2018); *Cooperative Entertainment, Inc. v. Kollektive Technology, Inc.*, 50 F.4th 127 (Fed. Cir. 2022).

THE ACCUSED INSTRUMENTALITIES

46. The Accused Instrumentalities are OpenAI's artificial intelligence products and services that implement the retrieval, personalization, recommendation, and communication-

generation functionality accused herein, including but not limited to: ChatGPT (including all versions and tiers—Free, Plus, Pro, Team, Enterprise, and Edu); ChatGPT Memory (including Saved Memories and Reference Chat History); ChatGPT Pulse; ChatGPT Shopping Research; ChatGPT Deep Research; ChatGPT Projects; ChatGPT Search; Company Knowledge in ChatGPT and its connectors; OpenAI Vector Stores and the Vector Store Search endpoint; the File Search tool in the Responses API; Assistants API File Search (through its discontinuation on August 26, 2026); the OpenAI Chat Completions API and Completions API, including the temperature, top_p, and logit_bias parameters; OpenAI Structured Outputs and the GPT-5 custom-tool grammar parameter; and the OpenAI embeddings API and the text-embedding-3 model family (collectively, the “Accused Instrumentalities”). The Accused Instrumentalities are accessible at chatgpt.com and via OpenAI’s developer platform at platform.openai.com. The Accused Instrumentalities execute on OpenAI’s processor-based cloud infrastructure, including OpenAI’s dedicated compute capacity at the Denton Campus.

47. ChatGPT with Memory is a computer-implemented conversational AI service running on OpenAI’s processor-based cloud infrastructure. Memory builds persistent user interest profiles from conversational behaviors (recency, frequency, context matching, and explicit user statements). Saved Memories are stored, topic-keyed entries the system picks up during interactions; Reference Chat History is a system-managed index over the user’s prior conversations used to recall information across sessions. On each user request, the Memory feature and Reference Chat History feature execute server-side processing pipelines: the system embeds the current user message, queries the user’s saved-memories store and Reference Chat History index for entries whose stored embeddings have high similarity to the query, selects a top-K subset of memory entries and prior conversation chunks, and injects them into the model context for response generation.

48. ChatGPT Pulse is OpenAI's proactive research feature that, each night, synthesizes information from the user's memory, chat history, and direct feedback (including "curate" requests and thumbs-up/down feedback on cards) to learn what is most relevant to the user, then delivers personalized research cards. Pulse draws on multiple per-user behavioral signals (memories, chat history, card interactions, curate feedback, connected apps) and, for ChatGPT Pro subscribers, is integrated with Shopping Research so it can proactively surface shopping guides based on previous conversations.

49. ChatGPT Shopping Research is OpenAI's personalized shopping agent (launched November 24, 2025), powered by a version of GPT-5 mini trained with reinforcement learning specifically for shopping tasks. When a user expresses shopping intent, Shopping Research reads trusted sites, cites reliable sources, and synthesizes information across many sources to produce product research, then delivers a personalized buyer's guide. It taps into ChatGPT's memory and prior conversations to recall the user's historical preferences and interests and contextualize future searches, prioritizing items based on user preferences without requiring explicit prompts.

50. The Accused Instrumentalities implement dual-analysis architectures that process both user behavioral data and content characteristics through separate computational pipelines. On information and belief, the Accused Instrumentalities use dual neural networks—one processing user behavioral data, preferences, and customizations and another processing content characteristics—that generate weighted embeddings subsequently combined to produce content for a user. On information and belief, OpenAI's Vector Stores implement a plurality of nodes—file chunks, each carrying an identifier, text, metadata, and an embedding vector—joined by graded affinities, namely the cosine similarities between those embedding vectors. On information and belief, OpenAI's retrieval substrate returns an explicit numerical relevance score for each retrieved chunk, reflecting an inferred degree of relevance of the phrases contained in that content object.

COUNT 1 – INFRINGEMENT OF U.S. PATENT NO. 8,676,742

51. Many Worlds repeats and incorporates by reference each preceding paragraph as if fully set forth herein, and further states:

52. On March 18, 2014, the '742 Patent was duly and legally issued by the United States Patent and Trademark Office for an invention entitled "Contextual Scope-Based Discovery Systems."

53. Many Worlds is the owner of the '742 Patent, with all substantial rights in and to that patent, including the sole and exclusive right to prosecute this action and enforce the '742 Patent against infringers, and to collect damages for all relevant times.

54. OpenAI's Accused Instrumentalities, including ChatGPT, OpenAI Vector Stores, the File Search tool in the Responses API, ChatGPT Memory, ChatGPT Shopping Research, ChatGPT Deep Research, ChatGPT Projects, and ChatGPT Search, infringe at least claims 1 and 8 of the '742 Patent. For claim 8, the Accused Instrumentalities provide: a contextual scope function executed on a processor-based computing device that establishes a contextual neighborhood within a computer-implemented fuzzy network-based structure (the vector-store retrieval index over embedded content, in which a search returns the subset of objects lying within a specified degree of relatedness of a base object, each carrying its own numerical relatedness score); a contextual neighborhood scoping function that establishes the scope of the contextual neighborhood at a level set by a user (OpenAI's documented file-search ranking controls, including the `score_threshold` relevance threshold on a documented 0-to-1 scale, the maximum-results setting on a documented 1-to-50 scale, the re-ranking stage that can be enabled or disabled, and the hybrid-search weighting between semantic and keyword matching, by which the user configuring the search sets the breadth of the returned neighborhood); and a computer-

implemented discovery function that generates a recommendation in accordance with the contextual neighborhood and an inference from a plurality of usage behaviors (on information and belief, the recommendation functionality of ChatGPT Memory and ChatGPT Shopping Research, which draws on a plurality of the user's behaviors over time).

55. For claim 1, on information and belief, the Accused Instrumentalities perform each step of the claimed method: OpenAI invokes a contextual scope function executed on a processor-based computing device that establishes a contextual neighborhood within the computer-implemented fuzzy network-based structure described above; a recommendation is received that is generated by a computer-implemented discovery function in accordance with that contextual neighborhood and an inference from a plurality of usage behaviors; and the contextual neighborhood is based on a selected object, namely the base object—the user's query, the selected file or product, or the memory entry—from which the degree-of-relatedness comparison is measured.

56. In the alternative, to the extent any step of claim 1 of the '742 Patent is performed by a user or a developer rather than by OpenAI, OpenAI is liable for that performance because OpenAI directs and controls it. OpenAI conditions a user's and a developer's receipt of the benefit of the Accused Instrumentalities—the retrieval and recommendation output itself—upon performance of the accused steps, and OpenAI establishes the manner and the timing of that performance through its published API specifications, its required request parameters and permitted parameter ranges, its documented default values, and the server-side pipeline that executes the accused operations. On information and belief, no user or developer can obtain the accused output except by performing the accused steps in the manner and at the times OpenAI prescribes. The

performance of those steps is therefore attributable to OpenAI, which is liable for direct infringement of the claimed method under 35 U.S.C. § 271(a).

57. OpenAI directly infringes at least claims 1 and 8 of the '742 Patent through the embedding-based retrieval and recommendation functionality described above, which establishes bounded contextual neighborhoods within a graded relationship structure and generates recommendations from those neighborhoods together with inferences drawn from user behaviors.

58. OpenAI has directly infringed and continues to directly infringe at least claims 1 and 8 of the '742 Patent, in violation of 35 U.S.C. § 271(a), by making, using, importing, offering for sale, and/or selling the Accused Instrumentalities without authority in the United States and will continue to do so unless enjoined by this Court.

59. OpenAI has committed these acts of infringement without license or authorization.

60. OpenAI has induced and continues to induce infringement of at least claim 1 of the '742 Patent under 35 U.S.C. § 271(b). OpenAI has had knowledge of the '742 Patent at least as of the filing of this Complaint. With knowledge of the '742 Patent, OpenAI has actively induced end users to directly infringe at least claim 1 of the '742 Patent by: (i) providing the Accused Instrumentalities to users with instructions on how to invoke the contextual retrieval and discovery features that practice the claimed method; (ii) marketing and promoting those retrieval and recommendation capabilities to encourage users to engage with them; (iii) providing user interfaces and documented configuration controls designed to facilitate user interaction with the accused contextual scope functionality; (iv) providing help documentation, tutorials, and support materials instructing users how to discover content through the accused systems; and (v) encouraging users through notifications, emails, and in-app prompts to engage with the recommendations the accused systems generate. OpenAI took these actions with specific intent to

induce infringement, or at minimum with willful blindness to the infringing nature of its contextual discovery system.

61. The claims of the '742 Patent are patent-eligible under 35 U.S.C. § 101. They are not directed to an abstract idea but to a specific computer architecture: a fuzzy network-based structure in which objects are joined by graded rather than binary relationships; a contextual scope function that establishes a bounded contextual neighborhood within that structure; a scoping function that sets the breadth of that neighborhood at a level set by a user; and a discovery function that generates a recommendation from the resulting neighborhood together with an inference drawn from a plurality of usage behaviors. A claimed data structure that departs from conventional structures and thereby improves how a computer stores and retrieves information is an improvement in computer capability, not an abstract idea. *Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327 (Fed. Cir. 2016). The claims recite the mechanism by which the result is achieved, not the result itself. *McRO, Inc. v. Bandai Namco Games Am., Inc.*, 837 F.3d 1299 (Fed. Cir. 2016).

62. Even if the claims of the '742 Patent were treated as directed to an abstract idea, they recite an inventive concept in their ordered combination. Many Worlds alleges as fact that, as of the priority date of the '742 Patent, a structure permitting multiple simultaneous graded relationships between the same pair of objects, coupled with a scope function that computed a bounded neighborhood by traversing those graded relationships at a breadth set by the user, and a discovery function that generated recommendations from that neighborhood together with behavioral inference, was not well-understood, routine, or conventional. Whether a claim element or an ordered combination of elements was well-understood, routine, and conventional is a question of fact. *Berkheimer v. HP Inc.*, 881 F.3d 1360 (Fed. Cir. 2018); *Aatrix Software, Inc. v. Green Shades Software, Inc.*, 882 F.3d 1121 (Fed. Cir. 2018).

63. The specification of the '742 Patent describes the technical deficiencies of the structures that preceded it and explains how the claimed fuzzy network-based structure, the contextual scope function computed within it, and the user-set neighborhood boundary overcome those deficiencies. The claims are confined to that architecture and do not reach the concept of recommending content by other means. Claims directed to a specific means or method that solves a problem in an existing technological process are not directed to an abstract idea. *Koninklijke KPN N.V. v. Gemalto M2M GmbH*, 942 F.3d 1143 (Fed. Cir. 2019); *CardioNet, LLC v. InfoBionic, Inc.*, 955 F.3d 1358 (Fed. Cir. 2020). That the claimed structure runs on standard computing equipment does not make it ineligible. *Thales Visionix Inc. v. United States*, 850 F.3d 1343 (Fed. Cir. 2017). Where the complaint and the written description together raise a plausible factual question whether the claimed combination was inventive, dismissal is improper. *Cooperative Entertainment, Inc. v. Kollektive Technology, Inc.*, 50 F.4th 127 (Fed. Cir. 2022); *Cellspin Soft, Inc. v. Fitbit, Inc.*, 927 F.3d 1306 (Fed. Cir. 2019).

64. By engaging in the conduct described herein, OpenAI has caused injury to Many Worlds, and Many Worlds has been damaged and continues to be damaged as a result thereof and OpenAI is thus liable to Many Worlds for infringement of the '742 Patent, pursuant to 35 U.S.C. § 271.

65. As a direct and proximate result of OpenAI's infringement of the '742 Patent, Many Worlds has suffered monetary damages and is entitled to a monetary judgment in an amount adequate to compensate Many Worlds for OpenAI's past infringement pursuant to 35 U.S.C. § 284, but in no event less than a reasonable royalty, together with interest and costs.

66. In addition, the infringing acts and practices of OpenAI have caused, are causing, and, unless such acts or practices are enjoined by the Court, will continue to cause immediate and

irreparable harm and damage to Many Worlds for which there is no adequate remedy at law, and for which Many Worlds is entitled to injunctive relief pursuant to 35 U.S.C. § 283. As such, Many Worlds is entitled to compensation for any continuing and/or future infringement up until the date that OpenAI is finally and permanently enjoined from further infringement.

67. OpenAI has had actual knowledge of the '742 Patent at least as of the date when it was notified of the filing of this action. By the time of trial, OpenAI will have known and intended (since receiving such notice) that its continued actions would infringe and actively induce the infringement of one or more claims of the '742 Patent. OpenAI's continued infringement after receiving such notice has been and continues to be willful, deliberate, and in reckless disregard of Many Worlds' rights in the '742 Patent. Many Worlds reserves the right to allege pre-suit willfulness should discovery show that OpenAI had knowledge of the '742 Patent before this action was filed.

68. Many Worlds has been damaged as a result of the infringing conduct by OpenAI alleged above. Thus, OpenAI is liable to Many Worlds in an amount that adequately compensates it for such infringements, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

69. Many Worlds is entitled to collect pre-filing damages for the full period allowed by law for infringement of the '742 Patent.

COUNT 2 – INFRINGEMENT OF U.S. PATENT NO. 10,699,202

70. Many Worlds repeats and incorporates by reference each preceding paragraph as if fully set forth herein, and further states:

71. On June 30, 2020, the '202 Patent was duly and legally issued by the United States Patent and Trademark Office for an invention entitled "Inferential-based Communications Method and System."

72. Many Worlds is the owner of the '202 Patent, with all substantial rights in and to that patent, including the sole and exclusive right to prosecute this action and enforce the '202 Patent against infringers, and to collect damages for all relevant times.

73. OpenAI's Accused Instrumentalities, including ChatGPT, the Responses API carrying the hosted File Search tool against an OpenAI Vector Store, OpenAI Structured Outputs, the GPT-5 custom-tool grammar parameter, ChatGPT Memory, ChatGPT Pulse, and ChatGPT Shopping Research, infringe at least claims 1, 8, 15, and 20 of the '202 Patent. For claim 1, OpenAI automatically accesses a first plurality of values determined by automatically analyzing text that is associated with a content object, wherein each of the plurality of values is based upon an automatically inferred degree of relevancy between an element of the text and the content object (on information and belief, the per-chunk relevance scores that OpenAI's file-search documentation describes as a score returned with each retrieved chunk, together with a documented score threshold running from 0.0 to 1.0 that determines which chunks reach the generation step); automatically accesses a second plurality of values, wherein each of the second plurality of values corresponds to one or more topics and wherein a magnitude of each is based upon an automatically determined inference from one or more user behaviors (on information and belief, the per-user, topic-organized store of preferences, interests, and goals that ChatGPT Memory builds automatically from the user's own chat history); and automatically generates a communication for delivery to a user, wherein the communication comprises a plurality of words selected in accordance with the first plurality of values, the second plurality of values, and one or more syntactical rules (OpenAI's Structured

Outputs feature converts a supplied JSON Schema into a context-free grammar, and GPT-5 custom tools accept a grammar directly, which is then enforced token by token during sampling).

74. OpenAI directly infringes at least claims 1, 8, 15, and 20 of the '202 Patent through its server-side text-generation platform, which analyzes content text to produce relevancy values, maintains behavior-derived topic values for the user, and generates communications whose words are selected under those values together with enforced syntactical rules.

75. OpenAI has directly infringed and continues to directly infringe at least claims 1, 8, 15, and 20 of the '202 Patent, in violation of 35 U.S.C. § 271(a), by making, using, importing, offering for sale, and/or selling the Accused Instrumentalities without authority in the United States and will continue to do so unless enjoined by this Court.

76. OpenAI has committed these acts of infringement without license or authorization.

77. OpenAI has induced and continues to induce infringement of at least claim 1 of the '202 Patent under 35 U.S.C. § 271(b). OpenAI has had knowledge of the '202 Patent at least as of the filing of this Complaint. With knowledge of the '202 Patent, OpenAI has actively induced end users and developers to directly infringe at least claim 1 of the '202 Patent by: (i) providing the Accused Instrumentalities with documentation instructing users how to perform the accused retrieval and constrained-generation operations; (ii) marketing and promoting those capabilities to encourage their use; (iii) providing APIs, parameters, and user interfaces designed to facilitate the accused operations; (iv) publishing API references, guides, and cookbooks instructing developers how to configure file search, score thresholds, and grammar-constrained output; and (v) encouraging use through notifications, release notes, and in-product prompts. OpenAI took these actions with specific intent to induce infringement, or at minimum with willful blindness to the infringing nature of its inferential-based communications system.

78. The claims of the '202 Patent are patent-eligible under 35 U.S.C. § 101. They are not directed to an abstract idea but to a specific machine architecture for generating a communication under three simultaneous and separately computed constraints: values derived by automatically analyzing the text of a content object and scoring the relevancy of its elements; values organized by topic whose magnitudes are derived by inference from observed user behaviors; and syntactical rules that govern how the words of the outgoing communication may be assembled. Claims reciting a particular set of computational rules for producing an output are not abstract merely because the output could loosely be described as a message. *McRO*, 837 F.3d 1299. Claims that build a profile by inferring characteristics from observed behavior, rather than by matching against a static list, are eligible on that basis. *Finjan, Inc. v. Blue Coat Sys., Inc.*, 879 F.3d 1299 (Fed. Cir. 2018).

79. Even if the claims of the '202 Patent were treated as directed to an abstract idea, they recite an inventive concept in their ordered combination. An inventive concept may reside in a non-conventional and non-generic arrangement of otherwise known components. *BASCOM Global Internet Servs., Inc. v. AT&T Mobility LLC*, 827 F.3d 1341 (Fed. Cir. 2016); *Amdocs (Israel) Ltd. v. Openet Telecom, Inc.*, 841 F.3d 1288 (Fed. Cir. 2016). Many Worlds alleges as fact that, as of the priority date of the '202 Patent, generating a communication whose words were selected under the simultaneous constraint of content-derived relevancy values, behavior-derived topic values, and syntactical rules was not well-understood, routine, or conventional; conventional systems employed one or the other of those inputs, and did not condition word selection on all three. That question is one of fact. *Berkheimer*, 881 F.3d 1360; *Aatrix*, 882 F.3d 1121.

80. The '202 Patent issued on June 30, 2020, following examination under the USPTO's post-Alice subject-matter-eligibility framework. The claims do not preempt the generation of explanatory or personalized messages: systems that assemble output from fixed templates, that

condition output on content analysis alone, or that condition output on a behavioral profile alone all remain outside the claims. A claimed improvement in the manner in which a computer generates and presents information to a user is a technological improvement, not an abstract idea. *Core Wireless Licensing S.A.R.L. v. LG Electronics, Inc.*, 880 F.3d 1356 (Fed. Cir. 2018); *TecSec, Inc. v. Adobe Inc.*, 978 F.3d 1278 (Fed. Cir. 2020). Whether the claimed ordered combination was conventional is a factual question that cannot be resolved against Many Worlds on the pleadings. *Cooperative Entertainment, Inc. v. Kollektive Technology, Inc.*, 50 F.4th 127 (Fed. Cir. 2022).

81. By engaging in the conduct described herein, OpenAI has caused injury to Many Worlds, and Many Worlds has been damaged and continues to be damaged as a result thereof and OpenAI is thus liable to Many Worlds for infringement of the '202 Patent, pursuant to 35 U.S.C. § 271.

82. As a direct and proximate result of OpenAI's infringement of the '202 Patent, Many Worlds has suffered monetary damages and is entitled to a monetary judgment in an amount adequate to compensate Many Worlds for OpenAI's past infringement pursuant to 35 U.S.C. § 284, but in no event less than a reasonable royalty, together with interest and costs.

83. In addition, the infringing acts and practices of OpenAI have caused, are causing, and, unless such acts or practices are enjoined by the Court, will continue to cause immediate and irreparable harm and damage to Many Worlds for which there is no adequate remedy at law, and for which Many Worlds is entitled to injunctive relief pursuant to 35 U.S.C. § 283. As such, Many Worlds is entitled to compensation for any continuing and/or future infringement up until the date that OpenAI is finally and permanently enjoined from further infringement.

84. OpenAI has had actual knowledge of the '202 Patent at least as of the date when it was notified of the filing of this action. By the time of trial, OpenAI will have known and intended (since receiving such notice) that its continued actions would infringe and actively induce the

infringement of one or more claims of the '202 Patent. OpenAI's continued infringement after receiving such notice has been and continues to be willful, deliberate, and in reckless disregard of Many Worlds' rights in the '202 Patent. Many Worlds reserves the right to allege pre-suit willfulness should discovery show that OpenAI had knowledge of the '202 Patent before this action was filed.

85. Many Worlds has been damaged as a result of the infringing conduct by OpenAI alleged above. Thus, OpenAI is liable to Many Worlds in an amount that adequately compensates it for such infringements, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

86. Many Worlds is entitled to collect pre-filing damages for the full period allowed by law for infringement of the '202 Patent.

COUNT 3 – INFRINGEMENT OF U.S. PATENT NO. 12,307,388

87. Many Worlds repeats and incorporates by reference each preceding paragraph as if fully set forth herein, and further states:

88. On May 20, 2025, the '388 Patent was duly and legally issued by the United States Patent and Trademark Office for an invention entitled "Probabilistically Tunable Conversational Method and System."

89. Many Worlds is the owner of the '388 Patent, with all substantial rights in and to that patent, including the sole and exclusive right to prosecute this action and enforce the '388 Patent against infringers, and to collect damages for all relevant times.

90. OpenAI's Accused Instrumentalities, including ChatGPT, the OpenAI Chat Completions API, the OpenAI Completions API, the temperature and top_p parameters exposed by those endpoints, and OpenAI File Search and Vector Stores when invoked in the same request, infringe at least claims 1, 8, and 15 of the '388 Patent. For claim 1, the Accused Instrumentalities

receive information from a user; provide a serendipity tuning control to the user (the documented temperature parameter, which OpenAI's API reference states will make the output more random at higher values, and the documented top_p nucleus-sampling parameter); receive from that control a serendipity tuning setting (the value supplied in the request payload before generation begins); and apply that setting to a probabilistic selection algorithm (the token-sampling step, which OpenAI describes as operating on the logits generated by the model prior to sampling). The Accused Instrumentalities then select a first syntactical element from a plurality of syntactical elements in accordance with that algorithm and select a second syntactical element in accordance with the first syntactical element and that algorithm (autoregressive token generation, in which each token is selected from the model's vocabulary conditioned on the tokens already generated). The Accused Instrumentalities generate and deliver a communication comprising the first and second syntactical elements, wherein the communication comprises a theme derived from a vector representation of a content object (on information and belief, in requests in which OpenAI's retrieval tooling supplies content objects retrieved by embedding similarity, the theme of the generated communication is derived from those retrieved objects' vector representations). This Count is not asserted against generation in which no retrieval tool is invoked.

91. OpenAI directly infringes at least claims 1, 8, and 15 of the '388 Patent through its generation stack, which exposes documented sampling controls to the user, applies the supplied setting to the probabilistic token-selection step, generates successive syntactical elements each conditioned on the last, and delivers a communication whose theme is derived from a vector representation of a retrieved content object.

92. OpenAI has directly infringed and continues to directly infringe at least claims 1, 8, and 15 of the '388 Patent, in violation of 35 U.S.C. § 271(a), by making, using, importing,

offering for sale, and/or selling the Accused Instrumentalities without authority in the United States and will continue to do so unless enjoined by this Court.

93. OpenAI has induced and continues to induce infringement of at least claim 1 of the '388 Patent under 35 U.S.C. § 271(b). OpenAI has had knowledge of the '388 Patent at least as of the filing of this Complaint. With knowledge of the '388 Patent, OpenAI has actively induced developers and end users to directly infringe at least claim 1 of the '388 Patent by: (i) publishing an API reference that instructs callers how to set the temperature and top_p parameters and describes their effect on output; (ii) marketing and promoting the tunability of its generation endpoints; (iii) providing the endpoints, parameters, and interfaces that make the accused method possible; (iv) publishing guides and cookbooks demonstrating the accused parameter settings together with retrieval tooling; and (v) encouraging use through release notes and in-product prompts. OpenAI took these actions with specific intent to induce infringement, or at minimum with willful blindness to the infringing nature of its probabilistically tunable conversational system.

94. OpenAI has committed these acts of infringement without license or authorization.

95. The claims of the '388 Patent are patent-eligible under 35 U.S.C. § 101. They are not directed to an abstract idea but to a specific mechanism for controlling machine generation: a control provided to the user that yields a tuning setting; the application of that setting to a probabilistic selection algorithm; the selection of a first syntactical element under that algorithm and of a second syntactical element under that algorithm and in accordance with the first; and the assembly of those elements into a communication whose theme is derived from a vector representation of a content object. The claims specify the control, the algorithm it governs, the

dependency between successive selections, and the source of the communication's theme. That is a specific means, not a desired result. *McRO*, 837 F.3d 1299; *Enfish*, 822 F.3d 1327.

96. Even if the claims of the '388 Patent were treated as directed to an abstract idea, they recite an inventive concept in their ordered combination. Many Worlds alleges as fact that, as of the priority date of the '388 Patent, exposing a user-facing control that tuned the probabilistic selection of successive syntactical elements, and deriving the theme of the resulting communication from a vector representation of a content object, was not well-understood, routine, or conventional. The '388 Patent was filed on June 26, 2024 and examined and allowed under the USPTO's post-Alice subject-matter-eligibility framework. Whether a combination was well-understood, routine, and conventional is a question of fact. *Berkheimer*, 881 F.3d 1360; *Aatrix*, 882 F.3d 1121.

97. The claims of the '388 Patent do not apply a generic or previously established generation technique to a new data environment. They recite the mechanism itself: a control provided to the user, a tuning setting received from that control, the application of that setting to the probabilistic selection algorithm, the selection of each successive syntactical element under that algorithm and in dependence on the element before it, and a communication whose theme is derived from a vector representation of a content object. The claims delineate the steps through which control over the variability of machine-generated output is achieved, and they are not limited to any field of use. Cf. *Recentive Analytics, Inc. v. Fox Corp.*, 134 F.4th 1205 (Fed. Cir. 2025). The specification describes the deficiency of prior systems that gave the user no control over the variability of generated output and describes how the claimed control operates on the selection algorithm. Whether that combination was conventional is a question of fact. *Cooperative Entertainment, Inc. v. Kollektive Technology, Inc.*, 50 F.4th 127 (Fed. Cir. 2022).

98. By engaging in the conduct described herein, OpenAI has caused injury to Many Worlds, and Many Worlds has been damaged and continues to be damaged as a result thereof and OpenAI is thus liable to Many Worlds for infringement of the '388 Patent, pursuant to 35 U.S.C. § 271.

99. As a direct and proximate result of OpenAI's infringement of the '388 Patent, Many Worlds has suffered monetary damages and is entitled to a monetary judgment in an amount adequate to compensate Many Worlds for OpenAI's past infringement pursuant to 35 U.S.C. § 284, but in no event less than a reasonable royalty, together with interest and costs.

100. In addition, the infringing acts and practices of OpenAI have caused, are causing, and, unless such acts or practices are enjoined by the Court, will continue to cause immediate and irreparable harm and damage to Many Worlds for which there is no adequate remedy at law, and for which Many Worlds is entitled to injunctive relief pursuant to 35 U.S.C. § 283. As such, Many Worlds is entitled to compensation for any continuing and/or future infringement up until the date that OpenAI is finally and permanently enjoined from further infringement.

101. OpenAI has had actual knowledge of the '388 Patent at least as of the date when it was notified of the filing of this action. By the time of trial, OpenAI will have known and intended (since receiving such notice) that its continued actions would infringe and actively induce the infringement of one or more claims of the '388 Patent. OpenAI's continued infringement after receiving such notice has been and continues to be willful, deliberate, and in reckless disregard of Many Worlds' rights in the '388 Patent. Many Worlds reserves the right to allege pre-suit willfulness should discovery show that OpenAI had knowledge of the '388 Patent before this action was filed.

102. Many Worlds has been damaged as a result of the infringing conduct by OpenAI alleged above. Thus, OpenAI is liable to Many Worlds in an amount that adequately compensates it for such infringements, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

103. Many Worlds is entitled to collect pre-filing damages for the full period allowed by law for infringement of the '388 Patent.

COUNT 4 – INFRINGEMENT OF U.S. PATENT NO. 12,299,603

104. Many Worlds repeats and incorporates by reference each preceding paragraph as if fully set forth herein, and further states:

105. On May 13, 2025, the '603 Patent was duly and legally issued by the United States Patent and Trademark Office for an invention entitled "Vector-Based Search Method and System."

106. Many Worlds is the owner of the '603 Patent, with all substantial rights in and to that patent, including the sole and exclusive right to prosecute this action and enforce the '603 Patent against infringers, and to collect damages for all relevant times.

107. OpenAI's Accused Instrumentalities, including ChatGPT, the standalone Vector Store Search endpoint, the File Search tool in the Responses API, Assistants API File Search, and ChatGPT retrieval over user-uploaded files, infringe at least claims 1, 8, and 15 of the '603 Patent. For claim 1, OpenAI applies one or more computer-implemented neural networks to each of a plurality of content objects to generate at least one distinct affinity value associated with each content object (when a file is uploaded to an OpenAI vector store, OpenAI splits it into chunks and embeds each chunk; OpenAI describes the endpoint that produces its embeddings as using neural network models); generates a plurality of affinity vectors within a multi-dimensional affinity space,

each corresponding to one content object (OpenAI describes each embedding as a vector of floating-point (e.g., decimal) numbers.

108. OpenAI directly infringes at least claims 1, 8, and 15 of the '603 Patent through its retrieval stack, which embeds stored content objects into a multi-dimensional vector space using neural network models, embeds the user's query into the same space, compares the query vector against the stored vectors by cosine similarity, selects the highest scoring vectors, and returns the corresponding content.

109. OpenAI has directly infringed and continues to directly infringe at least claims 1, 8, and 15 of the '603 Patent, in violation of 35 U.S.C. § 271(a), by making, using, importing, offering for sale, and/or selling the Accused Instrumentalities without authority in the United States and will continue to do so unless enjoined by this Court.

110. OpenAI has committed these acts of infringement without license or authorization.

111. OpenAI has induced and continues to induce infringement of at least claim 1 of the '603 Patent under 35 U.S.C. § 271(b). OpenAI has had knowledge of the '603 Patent at least as of the filing of this Complaint. With knowledge of the '603 Patent, OpenAI has actively induced developers and end users to directly infringe at least claim 1 of the '603 Patent by: (i) publishing API references and guides instructing callers how to create vector stores, upload and embed files, and run semantic searches against them; (ii) marketing and promoting the semantic-search capabilities of its retrieval products; (iii) providing the endpoints, tools, and interfaces that perform the accused method; (iv) publishing embeddings guides and cookbooks demonstrating query-to-object cosine-similarity retrieval; and (v) encouraging use through release notes and in-product prompts. OpenAI took these actions with specific intent to induce infringement, or at minimum with willful blindness to the infringing nature of its vector-based search system.

112. The claims of the '603 Patent are patent-eligible under 35 U.S.C. § 101. They are not directed to an abstract idea but to a specific retrieval architecture: applying a neural network to each stored content object to generate affinity values; assembling those values into per-object affinity vectors within a multi-dimensional affinity space; applying a neural network to the user's query to generate a query affinity vector; comparing the query vector against the stored vectors with a mathematical vector-comparison algorithm; and selecting and returning content on the basis of that comparison. The claims recite the computational structure by which meaning-based retrieval is accomplished, which is an improvement in how a computer stores and searches information rather than an abstract idea. *Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327 (Fed. Cir. 2016); *Data Engine Techs. LLC v. Google LLC*, 906 F.3d 999 (Fed. Cir. 2018).

113. Even if the claims of the '603 Patent were treated as directed to an abstract idea, they recite an inventive concept in their ordered combination. Many Worlds alleges as fact that, as of the priority date of the '603 Patent, encoding both stored content and a user query into a shared multi-dimensional affinity space by neural network and retrieving content by mathematical comparison of those vectors was not well-understood, routine, or conventional. The '603 Patent was filed on June 26, 2024 and examined and allowed under the USPTO's post-Alice subject-matter-eligibility framework. Whether a combination was well-understood, routine, and conventional is a question of fact. *Berkheimer v. HP Inc.*, 881 F.3d 1360 (Fed. Cir. 2018); *Aatrix Software, Inc. v. Green Shades Software, Inc.*, 882 F.3d 1121 (Fed. Cir. 2018).

114. The claims of the '603 Patent do not claim the application of generic machine learning to a new data environment. They recite the retrieval architecture itself: the generation of at least one distinct affinity value for each of a plurality of stored content objects; the assembly of those values into per-object affinity vectors within a multi-dimensional affinity space; the

generation of a corresponding affinity vector from the information a user enters into a search interface; the comparison of the query vector against each stored vector by a mathematical-based vector-comparison algorithm; the selection of vectors in accordance with that comparison; and the return of the corresponding content objects. The claims delineate the steps through which meaning-based retrieval is accomplished, and they are not confined to any field of use, industry, or category of content searched. Cf. *Recentive Analytics, Inc. v. Fox Corp.*, 134 F.4th 1205 (Fed. Cir. 2025). A claimed improvement in the manner in which a computer stores, indexes, and searches information is a technological improvement. *Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327 (Fed. Cir. 2016); *Data Engine Techs. LLC v. Google LLC*, 906 F.3d 999 (Fed. Cir. 2018); *Koninklijke KPN N.V. v. Gemalto M2M GmbH*, 942 F.3d 1143 (Fed. Cir. 2019).

115. By engaging in the conduct described herein, OpenAI has caused injury to Many Worlds, and Many Worlds has been damaged and continues to be damaged as a result thereof and OpenAI is thus liable to Many Worlds for infringement of the '603 Patent, pursuant to 35 U.S.C. § 271.

116. As a direct and proximate result of OpenAI's infringement of the '603 Patent, Many Worlds has suffered monetary damages and is entitled to a monetary judgment in an amount adequate to compensate Many Worlds for OpenAI's past infringement pursuant to 35 U.S.C. § 284, but in no event less than a reasonable royalty, together with interest and costs.

117. In addition, the infringing acts and practices of OpenAI have caused, are causing, and, unless such acts or practices are enjoined by the Court, will continue to cause immediate and irreparable harm and damage to Many Worlds for which there is no adequate remedy at law, and for which Many Worlds is entitled to injunctive relief pursuant to 35 U.S.C. § 283. As such, Many Worlds is entitled to compensation for any continuing and/or future infringement up until the date that OpenAI is finally and permanently enjoined from further infringement.

118. OpenAI has had actual knowledge of the '603 Patent at least as of the date when it was notified of the filing of this action. By the time of trial, OpenAI will have known and intended (since receiving such notice) that its continued actions would infringe and actively induce the infringement of one or more claims of the '603 Patent. OpenAI's continued infringement after receiving such notice has been and continues to be willful, deliberate, and in reckless disregard of Many Worlds' rights in the '603 Patent. Many Worlds reserves the right to allege pre-suit willfulness should discovery show that OpenAI had knowledge of the '603 Patent before this action was filed.

119. Many Worlds has been damaged as a result of the infringing conduct by OpenAI alleged above. Thus, OpenAI is liable to Many Worlds in an amount that adequately compensates it for such infringements, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

120. Many Worlds is entitled to collect pre-filing damages for the full period allowed by law for infringement of the '603 Patent.

COUNT 5 – INFRINGEMENT OF U.S. PATENT NO. 8,843,433

121. Many Worlds repeats and incorporates by reference each preceding paragraph as if fully set forth herein, and further states:

122. On September 23, 2014, the '433 Patent was duly and legally issued by the United States Patent and Trademark Office for an invention entitled "Integrated Search and Adaptive Discovery System and Method."

123. Many Worlds is the owner of the '433 Patent, with all substantial rights in and to that patent, including the sole and exclusive right to prosecute this action and enforce the '433 Patent against infringers, and to collect damages for all relevant times. The '433 Patent and U.S.

Patent No. 8,600,926, to which its terminal disclaimer refers, were conveyed to Many Worlds by the same recorded instrument and remain commonly owned.

124. OpenAI's Accused Instrumentalities, including ChatGPT, OpenAI File Search and vector-store retrieval, the File Search tool in the Responses API, and ChatGPT Memory, infringe at least claims 1, 8, and 17 of the '433 Patent. For claim 1, OpenAI identifies a computer-implemented content object to serve as a contextual reference (a file added to an OpenAI vector store, which OpenAI chunks, embeds, and indexes, and against which subsequent relevance scoring is computed); accesses a first plurality of values, each based on an inference of the user's interest level with respect to a topic derived from a plurality of behaviors (ChatGPT Memory, which builds a persistent per-user record of interests and preferences automatically from the user's conversations, weighting how recent a detail is and how often the user returns to a topic); accesses a second plurality of values, each based on an inferred degree of relevance to a topic of one or more phrases contained in the content object (the per-result similarity score that OpenAI's retrieval returns for each retrieved chunk of the content object); and, on information and belief, generates a recommendation comprising one or more objects for delivery to the first user, wherein the recommendation is generated by a recommendation function executed on a processor-based device in accordance with the first plurality of values and the second plurality of values.

125. OpenAI directly infringes at least claims 1, 8, and 17 of the '433 Patent through its retrieval-augmented recommendation platform, which represents content objects as embedding vectors and scores them for relevance, infers the user's interests by topic from a plurality of usage behaviors, and generates recommendations for delivery to the user.

126. OpenAI has directly infringed and continues to directly infringe at least claims 1, 8, and 17 of the '433 Patent, in violation of 35 U.S.C. § 271(a), by making, using, importing,

offering for sale, and/or selling the Accused Instrumentalities without authority in the United States and will continue to do so unless enjoined by this Court.

127. OpenAI has committed these acts of infringement without license or authorization.

128. OpenAI has induced and continues to induce infringement of at least claim 1 of the '433 Patent under 35 U.S.C. § 271(b). OpenAI has had knowledge of the '433 Patent at least as of the filing of this Complaint. With knowledge of the '433 Patent, OpenAI has actively induced end users and developers to directly infringe at least claim 1 of the '433 Patent by: (i) providing the Accused Instrumentalities with documentation instructing users how to create vector stores, upload and index content, and run retrieval against them; (ii) marketing and promoting the personalized retrieval and recommendation capabilities of those products; (iii) providing interfaces, APIs, and configuration controls designed to facilitate the accused operations; (iv) publishing guides and help-center materials instructing users how to use Memory together with retrieval; and (v) encouraging use through notifications, release notes, and in-product prompts. OpenAI took these actions with specific intent to induce infringement, or at minimum with willful blindness to the infringing nature of its integrated search and adaptive discovery system.

129. The claims of the '433 Patent are patent-eligible under 35 U.S.C. § 101. They are not directed to an abstract idea but to a specific two-input recommendation architecture: a first set of values inferring a user's interest level in a topic from a plurality of observed behaviors; a second set of values inferring the degree of relevance to a topic of phrases contained within a content object; and a single recommendation function that generates the recommendation in accordance with both. Claims that build a profile by inferring characteristics from observed behavior, rather than by matching against a static list, are patent-eligible on that basis. *Finjan, Inc. v. Blue Coat Sys., Inc.*, 879 F.3d 1299 (Fed. Cir. 2018). Claims reciting a particular set of computational rules

for producing an output are not abstract merely because the output could loosely be described as a recommendation. *McRO, Inc. v. Bandai Namco Games Am., Inc.*, 837 F.3d 1299 (Fed. Cir. 2016).

130. Even if the claims of the '433 Patent were treated as directed to an abstract idea, they recite an inventive concept in their ordered combination; an inventive concept may reside in a non-conventional and non-generic arrangement of otherwise known components. *BASCOM Global Internet Servs., Inc. v. AT&T Mobility LLC*, 827 F.3d 1341 (Fed. Cir. 2016); *Amdocs (Israel) Ltd. v. Openet Telecom, Inc.*, 841 F.3d 1288 (Fed. Cir. 2016). Many Worlds alleges as fact that, as of the priority date of the '433 Patent, integrating behavior-derived per-topic interest inference with phrase-level in-content relevance inference within a single recommendation function was not well-understood, routine, or conventional; conventional recommenders employed collaborative filtering or content-based filtering, but did not condition a single recommendation function on both. Whether a claim element or an ordered combination was well-understood, routine, and conventional is a question of fact. *Berkheimer v. HP Inc.*, 881 F.3d 1360 (Fed. Cir. 2018); *Aatrix Software, Inc. v. Green Shades Software, Inc.*, 882 F.3d 1121 (Fed. Cir. 2018).

131. By engaging in the conduct described herein, OpenAI has caused injury to Many Worlds, and Many Worlds has been damaged and continues to be damaged as a result thereof and OpenAI is thus liable to Many Worlds for infringement of the '433 Patent, pursuant to 35 U.S.C. § 271.

132. As a direct and proximate result of OpenAI's infringement of the '433 Patent, Many Worlds has suffered monetary damages and is entitled to a monetary judgment in an amount adequate to compensate Many Worlds for OpenAI's past infringement pursuant to 35 U.S.C. § 284, but in no event less than a reasonable royalty, together with interest and costs.

133. In addition, the infringing acts and practices of OpenAI have caused, are causing, and, unless such acts or practices are enjoined by the Court, will continue to cause immediate and irreparable harm and damage to Many Worlds for which there is no adequate remedy at law, and for which Many Worlds is entitled to injunctive relief pursuant to 35 U.S.C. § 283. As such, Many Worlds is entitled to compensation for any continuing and/or future infringement up until the date that OpenAI is finally and permanently enjoined from further infringement.

134. OpenAI has had actual knowledge of the '433 Patent at least as of the date when it was notified of the filing of this action. By the time of trial, OpenAI will have known and intended (since receiving such notice) that its continued actions would infringe and actively induce the infringement of one or more claims of the '433 Patent. OpenAI's continued infringement after receiving such notice has been and continues to be willful, deliberate, and in reckless disregard of Many Worlds' rights in the '433 Patent. Many Worlds reserves the right to allege pre-suit willfulness should discovery show that OpenAI had knowledge of the '433 Patent before this action was filed.

135. Many Worlds has been damaged as a result of the infringing conduct by OpenAI alleged above. Thus, OpenAI is liable to Many Worlds in an amount that adequately compensates it for such infringements, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

136. Many Worlds is entitled to collect pre-filing damages for the full period allowed by law for infringement of the '433 Patent.

**OPENAI SHOULD BE PERMANENTLY ENJOINED FROM INFRINGING THE MANY
WORLDS ASSERTED PATENTS**

137. Although Many Worlds is seeking money damages in this complaint, those money damages are inadequate to fully compensate Many Worlds for OpenAI's infringement of the Many

Worlds Asserted Patents. OpenAI's continued unlicensed use impairs Many Worlds' ability to license the Many Worlds Asserted Patents on exclusive or field-limited terms. Furthermore, an entity—whether practicing or non-practicing—is entitled to use its assets to partner with or exclude whomever it chooses. OpenAI's infringement will continue to cause Many Worlds irreparable harm, and money damages are insufficient to remedy OpenAI's infringement. The balance of hardships and public interest also support an injunction against OpenAI. Many Worlds is therefore entitled to a permanent injunction against OpenAI or, in the alternative, to an ongoing royalty for OpenAI's continued infringement after judgment.

DEMAND FOR JURY TRIAL

138. Plaintiff hereby demands a jury trial for all issues so triable.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff prays for judgment as follows:

A. Declaring that OpenAI has directly infringed and induced infringement of, and continues to directly infringe and induce infringement of, U.S. Patent Nos. 8,676,742; 10,699,202; 12,307,388; 12,299,603; and 8,843,433;

B. Awarding damages to Many Worlds in an amount no less than a reasonable royalty for OpenAI's direct infringement and induced infringement of U.S. Patent Nos. 8,676,742; 10,699,202; 12,307,388; 12,299,603; and 8,843,433, together with prejudgment and post-judgment interest and without limitation by 35 U.S.C. § 287;

C. Declaring that OpenAI's infringement of the Many Worlds Asserted Patents has been willful, and awarding enhanced damages up to three times the amount found or assessed pursuant to 35 U.S.C. § 284;

D. Awarding injunctive relief pursuant to 35 U.S.C. § 283 enjoining OpenAI, its officers, agents, servants, employees, and all persons in active concert or participation with them, from further infringement of any claim of U.S. Patent Nos. 8,676,742; 10,699,202; 12,307,388; 12,299,603; and 8,843,433;

E. In the alternative to a permanent injunction, awarding Many Worlds an ongoing royalty for OpenAI's continued infringement after judgment;

F. Ordering an accounting and awarding supplemental damages for all infringement not accounted for by the verdict, together with prejudgment and post-judgment interest;

G. Awarding Many Worlds its attorneys' fees pursuant to 35 U.S.C. § 285 or as otherwise permitted by law; and

H. Awarding all other costs and relief that the Court deems just and proper.

Date: August 31, 2026

Respectfully submitted,

By: /s/ Kevin E. Cadwell

Kevin E. Cadwell, *Lead Attorney*

kcadwell@cadwellthomas.com

Texas Bar No. 24036304

Lisa M. Thomas

lthomas@cadwellthomas.com

Texas Bar No. 24079455

Bradley S. Bowling

bbowling@cadwellthomas.com

Texas Bar No. 24040555

CADWELL THOMAS LLP

5373 W. Alabama St., Suite 457

Houston, Texas 77056

Phone: (713) 360-1560

Fax: (940) 233-8587

Mark D. Siegmund

Texas Bar No. 24117055

msiegmund@cjsjlaw.com

“Grace” Shuya Yang

Texas Bar No. 24144144

gyang@cjsjlaw.com

CHERRY JOHNSON SIEGMUND JAMES P.C.

7901 Fish Pond Road, 2nd Floor

Waco, Texas 76710

Tel: (254) 732-2242

Fax: (866) 627-3509

Jack Shaw (admitted in this District)

California Bar No. 309382

jshaw@cjsjlaw.com

Ryan Loveless

Texas Bar No. 24036997

rloveless@cjsjlaw.com

William D. Ellerman

Texas Bar No. 24007151

wellerman@cjsjlaw.com

CHERRY JOHNSON SIEGMUND JAMES P.C.

8140 Walnut Hill Lane, Suite 105

Dallas, Texas 75231

Tel: (972) 987-0709

Fax: (866) 627-3509

ATTORNEYS FOR PLAINTIFF MANY WORLDS 2T INNOVATIONS LLC