AIPLA

American Intellectual Property Law Association

October 16, 2024

The Honorable Kathy K. Vidal
Under Secretary of Commerce for Intellectual Property and
Director of the United States Patent and Trademark Office
600 Dulany Street
Alexandria, VA 22314
Via Federal Rulemaking Portal at https://www.regulations.gov

RE: **Comments on 2024 Guidance Update on Patent Subject Matter Eligibility, Including on Artificial Intelligence (89 Fed. Reg. 137, Docket No.: PTO-P-2024-0026, July 17, 2024; FR Doc No: 2024-15377)**

Dear Director Vidal:

The American Intellectual Property Law Association is pleased to offer its comments to the United States Patent and Trademark Office ("USPTO" or "Office") in response to the Request for Comments ("the RFC") Regarding the 2024 Guidance Update on Patent Subject Matter Eligibility, Including on Artificial Intelligence ("the Guidance").

Founded in 1897, the American Intellectual Property Law Association ("AIPLA") is a national bar association of approximately 7,000 members including professionals engaged in private or corporate practice, in government service, and in the academic community. AIPLA members represent a wide and diverse spectrum of individuals, companies, and institutions involved directly or indirectly in the practice of patent, trademark, copyright, trade secret, and unfair competition law, as well as other fields of law affecting intellectual property. Our members represent both owners and users of intellectual property. Our mission includes helping establish and maintain fair and effective laws and policies that stimulate and reward invention while balancing the public's interest in healthy competition, reasonable costs, and basic fairness.

Introductory Comments

As noted in Executive Order 14110,[1] AI offers the potential to solve some of society's most difficult challenges. Promoting innovation that will unlock AI's potential will, as

---

[1] Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, *The White House* (Oct. 30, 2023), https://www.whitehouse.gov/briefing-room/presidential-

explained in the order, require investment in development and research and "tackling novel intellectual property questions and other problems to protect inventors and creators." For these reasons, the order directs the Director of the Office to clarify issues related to AI patentable subject matter by issuing additional guidance to USPTO patent examiners and applicants that addresses considerations at the intersection of AI and IP, including, as the USPTO Director deems necessary, updated guidance on patent eligibility to address innovation in AI.

AIPLA applauds this goal as patents must offer reasonably clear and predictable protection to encourage investment in any innovative technology, including AI. Unfortunately, the current state of the law on patent eligible subject matter makes the scope of patent protection for AI inventions unclear and thus uncertain. This is one reason why AIPLA has supported and continues to support legislative reform intended to clarify the law to ensure appropriate protection for critical emerging technologies, such as AI, to incentivize investment in them.

Until legislative reform occurs, the USPTO must, of course, apply current laws. While the paucity of judicial decisions applying judicially created exceptions in the context of inventions involving AI presents challenges, there is an opportunity for the USPTO to establish guidance that brings at least some consistency, and thus predictability, when examining claims for AI inventions. AIPLA is, therefore, pleased by USPTO's efforts to tackle difficult questions by updating its patent subject matter eligibility guidance to address AI inventions. The 2024 Guidance Update on Patent Subject Matter Eligibility, specifically with respect to Artificial Intelligence ("Guidance") addresses the vital importance of clarity in the application of 35 U.S.C. §101 to AI inventions.

The Guidance acknowledges that application of the *Alice*/*Mayo* test for analyzing subject matter eligibility "can be challenging for AI inventions," and communicates dedication to remedying unfair hindrances to patentability for AI inventions.[2] We recognize the challenges the Office faces in developing guidance that aligns with judicial decisions while remaining flexible for technological advancements and commend the efforts in providing such guidance. We believe that certain key issues remain, however, that need addressing to ensure fairness, accuracy, and clarity in the examination process, especially with respect to AI technologies. Our primary observations and recommendations are briefly summarized as follows:

---

actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/.

[2] U.S. Patent & Trademark Office, *2024 Guidance Update on Patent Subject Matter Eligibility, Including on Artificial Intelligence*, 89 FR 58128, 58134 (July 17, 2024) ("III. Update on Certain Areas of the USPTO's Patent Subject Matter Eligibility Guidance Applicable to AI Inventions").

- Clear and consistent application of the *Alice*/*Mayo* test for analyzing the subject matter eligibility of AI inventions is necessary.

  o Step 2A, Prong One: There is a need for more detailed guidance on determining whether AI-related claims recite an abstract idea.

  o Step 2A Prong Two: There is a need for clarification on the proper evaluation of what a claim is "directed to," particularly relating to analyzing the claim as a whole.

  o Step 2B: There is a need for guidance for analyzing AI claims to determine whether they amount to "significantly more" than a judicial exception.

- Subject Mattery Eligibility Examples 47-49 should be clearly tailored to the subject matter eligibility of AI itself and should directly address the known challenges that are most likely to arise during the examination process for AI inventions. While AIPLA appreciates the Office's efforts in providing examples specific to AI inventions, AIPLA is concerned that the new examples as proposed do not reflect modern AI practices, technological realities, and typical claiming approaches.

I. More Guidance for Determining When a Claim Involving AI Recites an Abstract Idea Would Be Helpful

Section III of the Guidance addresses Step 2A, Prong One of the *Alice*/*Mayo* test under the first subheading. AIPLA believes there is insufficient guidance for examiners concerning the evidence needed to support a determination that a claim recites an abstract idea under Step 2A, Prong One. The Guidance provides no assistance in striking a balance between too little detail, which <u>could</u> lead to a mental process, and too much detail, which leads to a mathematical concept characterization. As a result, Examples 39 and 47 conflict. In Example 39, broadly claimed "training" is not abstract. In Example 47, more narrow claimed training including applying back propagation and gradient descent algorithms inexplicably becomes an abstract mathematical concept. The examples are in conflict because those skilled in the art would consider both descriptions of training to have the same meaning. At a minimum, the Guidance should explain why removing from Example 47 the limitation "applying back propagation and gradient descent algorithms" to match Example 39 makes it no longer abstract under Prong One. In particular, the Guidance should evaluate whether omitting the description of the

mathematical concept leads to a statutory subject matter (SSM) claim, requiring no further SSM assessment, or if it leads to a mental process, necessitating clearer guidance.

The existing guidance for a Step 2A, Prong One analysis is muddled by ambiguity, particularly relating to the identification of mental processes and mathematical concepts. This lack of clarity enables an excessive amount of flexibility during examination such that it is difficult for applicants to predict whether claim wording will be identified as an abstract idea or will overcome a §101 rejection once the examiner has identified an abstract idea. AIPLA urges the Office to offer more detailed guidance for evaluating AI-related claims at Step 2A, Prong One. The Guidance should help examiners better assess whether an AI invention should truly be considered abstract and help applicants claim their inventions without triggering Prong One. This will help avoid improper rejections, overcome proper rejections, and ensure that AI-related claims are evaluated in a manner that accurately reflects their technical contribution.

II. Proper Analysis of What a Claim is "Directed To" at Step 2A, Prong Two of the *Alice*/*Mayo* Test

Step 2A, Prong Two is intended to "distinguish claims that are 'directed to' the recited judicial exception from claims that are not 'directed to' the recited judicial exception."[3] To ensure that Prong Two functions as intended, AIPLA recommends that the Step 2A, Prong Two analysis should be more robust for computer implemented inventions including, but not limited to, AI inventions, emphasizing the significant need for proper consideration of all claim elements. AIPLA believes there is a need for additional guidance and examiner training emphasizing the importance of evaluating whether claim elements as a whole contribute to the technical improvement or practical application of the AI invention, rather than in many cases separating additional elements and subsequently summarily discarding these elements as insignificant.

III. More Robust Guidance is Necessary for Step 2B of the *Alice*/*Mayo* Test

The Guidance addresses Step 2A of the *Alice*/*Mayo* test at length but fails to meaningfully address Step 2B. AIPLA remains concerned that insufficient guidance exists for analyzing the eligibility of computer implemented inventions including, but not limited to, AI inventions at Step 2B. The Guidance fails to adequately address how to assess whether AI-

---

[3] *See* MPEP 2106.04(II)(2).

related claim elements, when viewed as a whole, amount to "significantly more" than a judicial exception. Given that AI technology typically involves highly technical implementations and specific improvements to computational processes, a more tailored Step 2B analysis for AI inventions is necessary.

To this end, AIPLA recommends that the Office develop clearer guidelines for identifying inventive concepts in AI-related claims, especially where the technological improvement lies in the structure, function, or efficiency of the AI invention itself, rather than in peripheral aspects of the invention.

IV. Concerns about Examples 47-49

While we commend the Office's attempt to provide more concrete examples on subject matter eligibility analysis in the field of AI, AIPLA has concerns regarding the practical applicability of supplemental examples ("Examples 47-49"), addressed in Section V of the Guidance. Accordingly, a more detailed analysis of Examples 47-49 is provided in the attached Appendix; the key points are summarized below.

A. <u>The Proposed AI Examples Provide Limited Value in Practice</u>

AIPLA is concerned that Examples 47-49 are not representative of AI Inventions and claims applicants are filing with the USPTO. They are based on unlikely or technically inaccurate or misleading fact patterns, rely on outdated technologies, and offer claim examples that would likely not be pursued by applicants. Consequently, they are likely to provide only limited assistance to examiners and applicants.

An AI invention broadly covers the software and/or the hardware on which it runs. Generally, the elements of AI inventions could relate to the algorithm, the training data, the infrastructure supporting the training database, the hardware running the algorithm, the hardware that supports the algorithm accessing the training data, the input hardware (e.g., sensors) that collect new data for the training dataset, and the output results of implementing the algorithm.

Examples 47-49 do not include claims directed to innovations in AI itself. That is, none of the eligible claims provided in Examples 47-49 derive their eligibility from the AI itself. For example, claim 1 of Example 47 was deemed eligible because it recites a unique arrangement of ASIC components that would not be found in AI inventions. Example 49 indicates that the only eligible subject matter is due to a novel compound, which could have been protected with a claim directed to a composition, without reciting any limitations relating to AI. Other eligible

claims derive their eligibility from external factors, such as packet processing for network security, speech signal modification, or creating new compounds for eye drops.

Furthermore, none of the examples derive eligibility based on an improvement of the AI model architecture, such as a novel approach to model training, or an improvement to the AI itself as a product that can run more efficiently on a resource-constrained device.

B. <u>Examples 47-49 Rely on Problematic Claim Interpretations</u>

AIPLA is concerned that the Broadest Reasonable Interpretation ("BRI") relied upon for the analysis of representative claims within Examples 47-49 fails to comply with the practices outlined within the MPEP[4] in part because many of them conflate eligibility analysis with interpretation. A claim must be first interpreted and then reviewed according to the prescribed eligibility analysis. The approach taken in the examples using claim interpretation consequently opens the door for mischaracterization of AI-related claim elements as mental steps or mathematical operation. While claims might recite limitations that amount to no more than routine and conventional applications of a judicial exception, they plainly are not claiming a mental process or mathematical algorithm, as such.

AIPLA recommends that Examples 47-49 be revised to correct flaws in the BRI of the presented claims, as well as providing clarification that AI processes should not unilaterally be mischaracterized as mental processes or mathematical concepts. By taking this step, the Office will help ensure that AI-related processes are not unfairly excluded from eligibility due to this mischaracterization.

C. <u>Inconsistencies in the Application of the *Alice*/*Mayo* Framework</u>

For reasons detailed in the Appendix, each of Examples 47-49 contradicts other analyses within the previously released eligibility examples, or the MPEP. Of particular note are the contradictions between Example 47 and Example 39, and the rationale provided for the analysis that fails to align with judicial precedent, e.g., *SRI International, Inc. v. Cisco Systems, Inc*.

Examples 47-49, if written correctly, would provide an opportunity to make up for the lack of guidance within the MPEP or controlling jurisprudence relating to the identification of abstract ideas within AI-related claims. For example, it is important to clearly distinguish between AI operations that are based on mathematical concepts, such as the training of AI models or certain input data processing operations, and ineligible claims that present a risk of monopolizing a mathematical concept. The examples would also allow the Office to mitigate

---

[4] *See* MPEP 2111.

the risk of §101 analysis shortcuts during examination of AI inventions. Efforts to decrease the frequency of improper §101 rejections due to poor claim interpretation practices and failure to consider key factors of the *Alice/Mayo* test can streamline the examination process to the benefit of applicants and examiners alike.
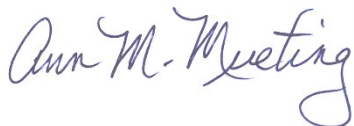
AIPLA recommends that the Office consider reviewing Examples 47-49 for consistency and cohesiveness, and it is with hope that the in-depth review provided within the Appendix can assist in the revision process.

Conclusion

In conclusion, while AIPLA again commends the Office in its rapid efforts to respond to EO 14110, AIPLA believes that improvements to the Guidance are essential. These include providing more specific guidance at each prong of the *Alice/Mayo* test to address the unique challenges computer implemented and, in particular, AI technologies present, modernizing the examples to better reflect current AI practices, and correcting inconsistencies in the application of the eligibility framework under §101.

AIPLA appreciates the Office's continued efforts to refine its subject matter eligibility guidance, particularly in the rapidly evolving field of artificial intelligence. AIPLA believes that the issues raised in this letter, if addressed, will lead to more effective guidance that benefits both patent applicants and examiners. AIPLA welcomes the opportunity to collaborate further and provide additional feedback as needed.

Thank you for considering our comments.

Ann M. Mueting
President
American Intellectual Property Law Association

APPENDIX: DETAILED COMMENTS ON EXAMPLES 47-49

AIPLA COMMENTS IN RESPONSE TO THE 2024 UPDATED GUIDANCE ON SUBJECT MATTER ELIGIBILITY

DETAILED COMMENTS ON EXAMPLE 47

Example 47 relates to a variety of embodiments that include an ASIC used for an artificial neural network ("ANN"), the application of an ANN to anomaly detection and identification in a data set, and detection of malicious network packets. While AIPLA appreciates the efforts of the Office to provide practical guidance for various structures and methods related to an ANN, the analysis provided with this example is undermined by improper claim interpretations, as well as being technologically incorrect and inconsistent with previous implementations of the *Alice/Mayo* framework.

AIPLA's concerns with Example 47 are as follows:

1)   The fact pattern and claim language are technologically flawed, limiting the applicability of the analysis to claims that may be presented to the Office.

2)   The Guidance creates a seemingly inescapable trap between different judicial exception categories with regard to drafting and interpretation, at least for ANN-related claims.

3)   The Guidance impermissibly divorces verbs and corresponding subjects, rather than looking at the claim limitations as a whole, thereby providing a template to reject essentially all computer process claims.

4)   The Guidance intertwines the involvement of a judicial exception in the claims as a whole with the recitation of the judicial exception.

5)   The Guidance incorrectly equates the limited preemption inherent in the claims with the broad preemption of the technology to which the claims are directed.

AIPLA recommends addressing these issues, as discussed in this section, to improve the effectiveness of Example 47 in assisting examiners and the public in evaluating patent claims.

1.   *Overview of Key §101 Considerations and Aspects of AI Technology Relevant to Example 47*

In prior correspondence, AIPLA suggested that AI inventions can include improvements to the architecture of AI engines, model training, acquisition of training data, inference and specific outputs from AI engines, as well as hardware on which AI systems run.[5] Guidance on the application of subject matter eligibility ("SME") analysis of claims to AI-

---

[5] AIPLA Comments to Request for Comments on Patenting Artificial Intelligence Inventions, 84 Fed. Reg. 44889 (August 27, 2019) [Docket No. PTO–C–2019–0029] , at 2 (November 8, 2019).

related inventions should provide sensible examples of eligible advances in each of these AI categories.

AI architectures are based on mathematical concepts. These mathematical concepts may be recited directly or inferentially in claims to AI-related inventions. Evaluation of "abstract ideas" like the mathematical concepts within §101 eligibility analysis has become untethered from the original objective articulated in *Diehr* and reiterated in *Mayo, citing Diehr*. The objective of the §101 judicial exception in *Diehr* was to prevent monopolization of a natural law, embodied in a mathematical formula. Unlike the objectives for §§102, 103 and 112, the objective of identifying abstract ideas has become unclear and counterintuitive due to the multiple meanings of "abstract" beyond just mathematical concepts.[6] Due to the lack of statutory guidance and conflicting judicial decisions, the Office has an opportunity to provide needed leadership in the area of articulating a §101 objective applicable to AI, which can anchor the analysis of whether claims are based on or directed to mathematical concepts, as there is sparse precedent.

The meaning of "abstract idea" applied in SME analysis is at odds with the meaning of "abstract" as understood in other fields. For instance, in Examples 39 and 47, "training" a neural network is *not* abstract for SME purposes but reciting explicitly that "training" an artificial neural network includes "backpropagation and gradient descent" moves training into the "abstract" mathematical concept bucket. In plain English, saying explicitly that training a neural network includes backpropagation and gradient descent[7] makes the training *more* concrete and *less* abstract. However, the Examples teach that the *linguistically more concrete* limitation is *legally more abstract* for SME. This does nothing to serve the objective of preventing monopolization of a natural law. In the examples, saying backpropagation and gradient descent explicitly -- instead of merely implying such technologies -- counterintuitively crosses the line from a non-abstract limitation based on math to reciting an abstract mathematical concept.

---

[6] Sachs (Mar. 28, 2014) *Alice at Court: Stepping Through the Looking Glass of the Merits Briefs in Alice v. CLS Bank – Part I*, IPWatchdog at < https://ipwatchdog.com/2014/03/28/merits-briefs-in-alice-v-cls/id=48742/ > ("What's an Abstract Idea?").

[7] Training an artificial neural network with training examples is understood to be based on backpropagation and gradient descent, because that is by far the dominant training technology. A few experimental alternatives have been proposed, but a person skilled in the art would expect training an artificial neural network with examples to be based on backpropagation and gradient descent, in the absence of any statement to the contrary. *See, e.g., Backpropagation*, in Wikipedia at <https://en.wikipedia.org/wiki/Backpropagation>

Similar issues arise in the analysis of other examples, which contradict the outcome for analogous technology in *Ex parte Hannun*.[8]

2. *Evaluation of Mathematical Concepts and Mental Processes as Abstract Ideas*

Examples that tie recited mathematical concepts to a clearly stated objective would be much easier for applicants and examiners to apply to real claims, e.g., the subject matter eligibility of feature engineering of inputs to neural networks. Tying the mathematical concepts exception to an objective would make it easier in interviews to discuss, for instance, the circumstances in which feature engineering of inputs to convolutional neural networks or other AIs is sufficient to render a claim eligible under §101, allowing examination to proceed to §§102, 103, and 112 analysis.

The analysis of "mental processes" in AI claims also requires clarification.[9] For example, claims that are directed towards the training of a neural network or the preparation of training data, which often involve millions of training examples, should not be categorized as a mental process when the required steps necessarily require a computer (e.g., internal data type conversion from real to integer). This is especially true when the data volume is too large to reasonably be processed by human editing. If the required data preparation is obvious, this should be addressed under §103, not §101. Similarly, processing inputs through thousands or millions of nodes in a convolutional neural network ("CNN") cannot reasonably be carried out in the human mind, even using a pen and paper.

3. *Claim Interpretations within Example 47 fail to adequately address anomaly detection as a fundamental and practical application of machine learning classification.*

Interpretation of Example 47 Claim 1 ("Claim 1") should go beyond the preamble alone to confirm that the claim elements follow the direction of the preamble to hardware. AIPLA agrees that it is reasonable to conclude that the hardware components in the claims do not recite an abstract idea. However, the claim could benefit from revision for technical accuracy. The

[8] In *Ex parte Hannun*, the spectrographic segment subject matter is determined eligible, whereas the spectrographic segment in Example 48 is determined to be abstract due to the referenced short-term Fourier transform. *Ex parte Hannun*, No. 2020-003120, 2021 WL 979302 (P.T.A.B. Mar. 22, 2021).) [Informative, applying 2019 revised guidance].
[9] Sachs, *supra*, *Alice at Court Part I* (distinguishing between preemption of fundamental truths such as mathematical and natural laws and disembodied concepts including mental processes).

requirements of Claim 1, including the allocation of a microprocessor to each neuron in the ANN and separately allocating a coefficient memory to each neuron, do not reflect a realistic hardware design.

Drafting errors in the Background impact the description of the processing network, of converting continuous data to discrete data, and of data used for malicious packet detection. An early draft of Example 47, summarized on the USPTO's web site, appears to have allocated a register to each neuron, among the thousands or millions of neurons in a neural network. The published Claim 1 allocates a whole microprocessor to each neuron, which is unrealistic.

Of the three provided claims within Example 47, Claim 2 best demonstrates why patentability of artificial intelligence and neural network technologies should be judged primarily under 35 U.S.C. §103 instead of 35 U.S.C. §101. Claim 2 is directed to anomaly detection, and is directed at a high level to a fundamental capability of neural networks that would be recognized as obvious to those skilled in the art but for the recitation of a two-stage ANN, the purpose of which is not described in the Background. A Google Scholar search for anomaly detection and neural networks, limited to articles more than 25 years old, returned 16,200 hits, so Claim 2 is likely obvious or anticipated, unless the two-stage ANN is inventive.

The claims further suffer from confusion relating to the conversion of "continuous data" to discrete data, which one skilled in the art would interpret as converting real numbers to integers. Example 47 references "binning and clustering" as known discretization methods. Clustering would never be used for this purpose and nothing in the Background enables use of clustering for data-type conversion. As a matter of computer architecture and the sheer volume of data used for training, data type conversion should not be described in a way that can be interpreted as a mental process. For malicious packet detection, there is little chance of using "continuous" data that has "any number of possible values" during malicious packet detection, nor does the Background section sufficiently address this issue. It would be helpful for Example 47 either to eliminate (or to point out) these drafting errors for educational purposes.

Further interpretation errors in Example 47 include taking words out of context, particularly "detecting" and "analyzing" in the two-stage ANN of Claim 2 elements (d) and (e). The limitations read, "detecting one or more anomalies in a data set using the trained ANN" and "analyzing ... using the trained ANN"; however, the language "using the trained ANN" is ignored. For instance, it says "the plain meaning of 'detecting' encompasses mental observations or evaluations." A person skilled in the art, having read the claim language and Background,

would never adopt such an interpretation. Taking words out of context leads to an implausible interpretation.

Similar issues are apparent in the interpretation of Claim 3, particularly for Claim 3 element (b). Interpretation of Claim 3 elements (c)-(f) are performed well, and AIPLA suggests adding further explanation as to why the "determining" of Claim 3 element (c) does not recite a mental process.

> **4.   Application of Step 2A of the Alice/Mayo Test to Example 47 Requires Improvement.**

The evaluation of whether Example 47 steps are "abstract" at Step 2A, Prong One is flawed, primarily because it does not explain how to overcome the asserted bases for the "abstract" determination. Claim 2 step (b) includes data-type conversion referred to as "discretization," which should be viewed as being based on a mathematical concept, but not directed to a mathematical concept, because data conversion is not at the core of the claim. Real to integer data-type conversion is not in danger of being monopolized by this claim element.

Next, the "wherein" clause of training step (c) of Claim 2 does not attempt to preempt the use of backward propagation or gradient descent by reciting the algorithmic functions involved. This claim wording involves or is based on math. The wording distinguishes itself from *Gottschalk*[10], in which the claim was directed to a mathematical algorithm for data type conversion that was recited in detailed steps and directed to preempting use of the algorithm by users. AIPLA notes that the primary difference between the training step of Example 47 (abstract) and the training step of Example 39 (not abstract) is the inclusion of a wherein clause that merely states a widely accepted definition of training. The extra words in Example 47, as compared to Example 39, make the claim easier to understand, but do not change how a person of ordinary skill in the art would understand the claim as a whole. The extra words, which are otherwise implied, should not render the claim "abstract."

Overall, anomaly detection is "not the type of human activity that §101 is meant to exclude." *See*, MPEP §2106.04(a)(2)(III), quoting *SRI Int'l, Inc. v. Cisco Systems, Inc.,* 930 F.3d 1295, 1303 (Fed. Cir. 2019) (re detecting suspicious network activity from anomalies). Furthermore, in MPEP §2106.04(d)(1), *SRI Int'l, Inc.,* 930 F.3d at 1303 is cited again as an example of a claim that provides an improvement to a technical field. AIPLA believes that it is

---

[10] *Gottschalk v. Benson*, 409 U.S. 63 (1972).

unnecessarily confusing to draft exemplary ineligible claims that are highly similar to those of *SRI Int'l Inc.*, which is cited within the MPEP and other Office materials as an example of eligible subject matter, unless this juxtaposition is clearly addressed with specific explanations to distinguish between the contradictory outcomes.

At Step 2A Prong Two, the analysis paraphrases claim language and labels an extended paraphrase of the claim as "abstract." This is a legal error, according to the recent PTAB opinion in *Ex parte Annakov*.[11,12]

The application of Step 2A, Prong Two to the claims of Example 47 provides an example of the tendency for claims under examination to be paraphrased in Step 2A, Prong One, which subsequently leads to a failure to properly apply Step 2A, Prong Two to the claim as a whole, as required by *Diehr* and reiterated in MPEP §2106.04(d). During the Prong Two analysis, it is common for examiners to filter out any claim elements that were paraphrased in the Prong One analysis and ignore most of the claim as a result, which can be seen in Example 47. The Example 47 analysis, instead of paraphrasing and filtering, should apply Prongs One and Two as instructed by *Diehr* and MPEP §2106.04(d). The abstract idea should be concisely restated in Prong One analysis rejection to allow an applicant to decide how best to proceed, instead of allowing examiners to broadly label as abstract quotes or paraphrases of substantial parts of claims. Restatement of the identified abstract idea(s) will guide examiners toward applying Prong Two to the claim as a whole.

Analysis of whether the claim as a whole is directed to a practical application under Step 2A Prong Two should entail articulation of the problem addressed and identification of the improvements realized by the claimed solution over the prior solution.[13] Description of the prior anomaly detection technology on which the claim improves is missing from analysis of Example 47 Claim 2.

---

[11] *See also, e.g.,* PatentlyO, "Directed to" vs. "Reciting" vs. "Involving" an Abstract Idea, September 13, 2024. This is an ordinary decision, not yet elevated to informative or precedential.

[12] *Ex parte Annakov*, No. 2018-005666, 2019 WL 5587082 (P.T.A.B. Oct. 25, 2019). (in which the Board held that paraphrasing multiple elements of the claim and collectively labelling them an abstract idea was reversable error "because such a broad-brush approach in identifying purported recitations of abstract ideas is overly inclusive, and runs counter to the Revised Guidance, as incorporated in the MPEP. ... 'Examiners should accordingly be careful to distinguish claims that **recite** an exception (which require further eligibility analysis) and claims that merely **involve** an exception (which are eligible and do not require further eligibility analysis.'" Id., quoting MPEP §2106.04(II)(A)(1) (emphasis in opinion)).

[13] *See, e.g., McRO, Inc. v. Bandai Namco Games Am. Inc.*, 837 F.3d 1299, 1316, 120 USPQ2d 1091, 1103 (Fed. Cir. 2016) (methods of automatic lip synchronization and facial expression animation using computer-implemented rules were not directed to an abstract idea).

5.   **The Step 2B analysis of Example 47 does not represent best practices for identifying well-understood, routine, conventional activity.**

At Step 2B, the "apply it"[14] and "well-understood, routine, and conventional" ("WURC") analyses are too broad-brushed to be helpful for training or to satisfy the MPEP requirement for clarity of a rejection. The "apply it" analysis of the two-stage ANN in claim elements (d) and (e) of Claim 2 does not address inputs to, structure of, or outputs from the ANN. The Guidance fails to suggest amendments to (d) and (e) sufficient for SME, which would provide valuable educational benefit for applicants. The WURC analysis fails to cite cases that support the reasoning that prior cases found the claim limitations to be WURC, in accordance with MPEP 2106.07(a)(III). High-level citations to the MPEP without additional rationale fail to provide sufficient clarity for an applicant to learn from the provided examples.

The ANN in Claim 2 elements (d) and (e) has two stages, binary classification followed by multi-category classification. This sequence of operations is unusual, as an ANN could perform the multi-category classification in just one stage. As a result, Example 47 compounds errors in interpretation of these steps, described above, by ignoring the requirement of Step 2B to consider the claim as a whole. In programming, it has long been common practice to discuss the inputs to, process of, and outputs from ("IPO") any computer program, including AI technology like an ANN. The reasoning in Example 47 does not address any part of IPO. This leaves applicants and examiners without any guidance as to how Claim 2 elements (d) and (e) could be amended. In claim 2, the inputs raise §112 issues, both for clarity and written description. Neither the Background nor the claim is clear about the inputs. Consequently, the errors within the formulation of Example 47 are better addressed under §112, not §101.

The process that applies the ANN is more clearly claimed than the inputs, but the process is ignored in the reasoning. It can be determined from the holding in *Dolly, Inc. v. Spalding & Evenflo Cos.*, 16 F.3d 394 (Fed. Cir. 1994) that a single stage ANN would not infringe the claimed two-stage structure. The Background is insufficient in explaining the inventive contribution of the two-stage ANN, but the "apply it" analysis does not even acknowledge that a two-stage structure is positively claimed. On the output side, the reasoning concludes for Claim 2 element (e) that the output "anomaly data" is a type or cause of the anomaly. However, in the following discussion of Claim 2 element (f), the reasoning ignores the conclusion reached two

---

[14] *See* MPEP §2106.04(d)(I) ("Merely reciting the words "apply it" (or an equivalent) with the judicial exception, or merely including instructions to implement an abstract idea on a computer, or merely using a computer as a tool to perform an abstract idea, as discussed in MPEP § 2106.05(f)").

sentences earlier for element (e). The Step 2B analysis should be revised to discuss the ANN's IPO and to explain how the claim or description would be amended to satisfy §101. A successful amendment should be straightforward, and could be structured using the similar, patent eligible claims approved in *SRI Int'l, Inc. v. Cisco Systems, Inc.* as a template.[15] Without a counter example demonstrating a route to eligibility, there is a risk that Example 47 may be misinterpreted to imply that using a trained ANN in a production mode, as opposed to training the ANN, never constitutes eligible subject matter.

The WURC analysis in Example 47 (and in Example 48) is not sufficiently clear without further explanation, which should ideally include citations to specific case law. The court in *Berkheimer v. HP, Inc.,* 881 F.3d 1360, 1368, 125 USPQ2d 1649, 1654 (Fed. Cir. 2018) required a factual basis for a WURC determination, and MPEP §2106.07(a)(III) clearly outlines various procedures that satisfy this requirement. One valid option includes a citation to a previously decided court case that found a substantially similar claim limitation to constitute WURC activity. MPEP §§2106.05(d)(I)(2)(b) and 2106.07(a)(III)(B) call for citation to one or more court decisions to support an assertion that a claim limitation constitutes WURC activity, such as the cases listed in §2106.05(d)(II). This dovetails with 35 U.S.C. §132[16] to allow the applicant to determine how best to proceed, i.e., whether to distinguish the claim limitation in question from the cited court decision or to amend the claim. The analysis within Example 47 provides explanation only by way of citation to MPEP §2106.05(d), subsection (II). As such, the explanation simply refers the reader to a long list of nonspecific examples from the case law but fails to articulate rationale specific to the claim language of Example 47. A rejection under §101 formulated in such a way would fail to satisfy the requirements as outlined within the MPEP, nor does the analysis satisfy the statutory requirements of §101 or the relevant case law.

AIPLA suggests that Example 47 be revised to ensure consistency with Example 39 such that the claim element reciting model training is eligible subject matter. If this issue were addressed, it would be more straightforward to address the evaluation of whether the claim is based on mathematical concepts or directed to mathematical concepts in a clearer way.

6. *Example 47 Summary*

---

[15] *SRI Int'l, Inc. v. Cisco Systems, Inc.*, 930 F.3d 1295, 1304 (Fed. Cir. 2019) (declining to identify the claimed collection and analysis of network data as abstract because "the human mind is not equipped to detect suspicious activity by using network monitors and analyzing network packets as recited by the claims").
[16] *See also* MPEP §2141.

Example 47 addresses an important application of AI technology within the field of cybersecurity. Given the parallels between Example 47 and the claims at issue within SRI Int'l, the technological premise of Example 47 offers a promising foundation for demonstrating the eligibility of AI subject matter while building upon existing guidance within the MPEP and case law. However, certain flaws within Example 47 must be resolved in order to achieve this educational potential.

The fact pattern and claim language within Example 47 are technologically flawed, limiting the applicability of the analysis to claims that may be presented to the Office. The Guidance creates a seemingly inescapable trap between different judicial exception categories with regard to drafting and interpretation, especially for claiming neural networks and other AI-related inventions. Example 47 compounds this problem by impermissibly divorcing verbs and their corresponding subjects within the claim language, rather than evaluating the claim as a whole. Encouragement of these problematic practices consequently provides a template to reject most, if not all, computer-implemented method claims. The analysis of Example 47 fails to sufficiently distinguish between involvement of a judicial exception in the claim as a whole from the recitation of a judicial exception. Example 47 incorrectly equates limited preemption inherent in the recited claim language with the broad preemption of the technology to which the claims are directed, thereby opening loopholes in the subject matter eligibility analysis.

DETAILED COMMENTS ON EXAMPLE 48

Example 48 relates to the application of AI to speech signal processing, as an extension of deep clustering. AIPLA acknowledges the Office's efforts to provide practical guidance for this field of current interest. However, as detailed in the review below, AIPLA believes that the analysis provided with this example is undermined by improper claim interpretations and inconsistent application of the SME analysis.

AIPLA's concerns with Example 48 are as follows:

1)   The fact pattern and claim language included within Example 48 contain misleading or technically flawed information, limiting the real-world applicability of the provided analysis.

2)   Although Example 48 extensively addresses characterization of abstract ideas and, more specifically, emphasizes recitation of mathematical concepts, the Step 2A Prong One analyses within Example 48 fail to resolve existing issues of clarity and consistency relating to the identification of abstract ideas within claims.

3) Example 48 implicitly dismisses various legitimate practical applications and technological improvements by not providing claims that derive eligibility from AP input features, processing, or outputs.

7. *Example 48 can benefit from more realistic claims and reference to the background technology upon which they are based.*

The Background section of Example 48 describes a deep neural network ("DNN") such as an autoencoder that is trained to map spectrogram frames for segments of speech into an embedding space, from which the segments can be categorized by speaker. This is recognizable to those skilled in the art as an improvement on deep clustering technology.[17] The speech separation process described includes performing short-term Fourier transforms (STFT) to create spectrogram frames,[18] converting feature representations into embedding vectors, mapping the embedding vectors using a global function, clustering the embedding vectors, and leveraging of binary time-frequency masks.

The claim wording is artificial and leads to an interpretation that makes legally abstract a claim element that is grammatically more concrete than wording within claims that have been previously Board-approved. In Claim 1, for example, formulas are introduced for apparently no reason other than to create impediments to patentability. If the Guidance is to succeed in its goal of fostering AI innovation, it is important to provide credible, authentic claim language that is capable of leading applicants to successfully claim their AI inventions. Claim elements in Example 48 are considered abstract, seemingly based on the use of windowed Fourier transform (i.e., STFT), which has been a tool of signal processing for more than fifty years.[19] Fourier transformation is a standard component for transforming an input from the time domain to the time-frequency domain for processing. The informative decision in *Ex parte Hannun* held that claiming generation of a spectrogram frame, without mentioning use of the STFT, is not abstract. It is unhelpful for Example 48 to explicitly reference the STFT to make the claim

---

[17] Ochieng, "Deep neural network techniques for monaural speech enhancement and separation: state of the art analysis," in Artificial Intelligence Review (2023) 56:S3651-S3703, doi: https://doi.org/10.1007/s10462-023-10612-2. This article appears to be the basis for Example 48, as it includes the same unusual global function.

[18] *Compare*, *Ex parte Hannun*, Appeal 2018-003323, Claim 11 at 2-3 (PTAB designated informative Dec. 11, 2019)("generating a set of spectrogram frames for each audio file" did not recite math, was not abstract).

[19] J. Allen, "Short term spectral analysis, synthesis, and modification by discrete Fourier transform," in *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 25, no. 3, pp. 235-238, June 1977, doi: 10.1109/TASSP.1977.1162950.

abstract, without any comment on how to avoid abstraction or why more concrete claim wording should be considered legally abstract.

In Claim 1, the introduction of a global function defined as "$V = f_\theta(X)$," is more a generic placeholder than a meaningful formula. This generic formula that is unlikely to come up in real-world claims may have been included in the claim merely for the purpose of providing an example of a claim that recites a mathematical formula. Hence, it is difficult to extract meaningful guidance from the evaluation and analysis based on the recited mathematical formula given its generic and arbitrary form.

The USPTO has the opportunity to assert leadership regarding the application of §101 to mathematical formulas and algorithms that do not effectively limit claims, but rather make the claims less abstract to a reader without attempting to monopolize a mathematical algorithm. The courts have not addressed such a detailed issue (*see*, *e.g., Mayo*, *Alice*), and the Office should provide additional guidance for claims that are based on mathematical concepts to refine the courts' analyses.[20] AIPLA does not believe any court has held that added claim wording making explicit what is already implied by claim limitations, essentially a redundancy, should make a claim legally abstract in Prong One analysis.

The "Claim Interpretation" for Claim 1 disregards positively recited limitations of the claim. Element (a) of Claim 1 recites "receiving a mixed speech signal *x* comprising speech from multiple different sources…." The interpretation of element (a) asserts that the "claim does not put any limits on how the mixed speech signal is received." One of ordinary skill in the art would be familiar with the inputs to windowed Fourier transforms. A microphone would typically be used to collect such an audio signal, before the claimed method receives the signal. The omission of technical details that would be understood by those skilled in the art from the example analysis inappropriately pushes the examiner toward considering a concrete technical step to be abstract.

The interpretation of Claim 1 element (b) points out a critical error in the Background by way of not specifying what "temporal features" are being obtained and processed to supplement the spectrogram frame produced by the STFT. This is significant, because processing of temporal features in addition to spectrogram frames is a primary improvement in Example 48 over prior deep clustering. This issue is a §112 application drafting error, not

---

[20] *See*, Rader, *Rader's Ruminations – Patent Eligibility III: Seven Times the Federal Circuit Has Struck Out* (Mar. 31, 2024) at https://ipwatchdog.com/2024/03/31/raders-ruminations-patent-eligibility-iii-seven-times-federal-circuit-struck/id=174751/.

something that should make the subject matter of claim ineligible for examination and patentability. Similarly to the earlier discussion of §103 confounding problems for Example 47, the example claims should be free of §112 issues that muddy the clarity of the subject matter eligibility analysis and dilute the value as it specifically relates to §101.

In practice, applicants frequently encounter roadblocks during examination of AI-related claims in which the examiner fails to address any of the AI-specific language while interpreting the claim. It would be beneficial if guidance were provided to address and mitigate these interpretations. Example 48, as currently presented however, may exacerbate the problem instead.

8. ***At Step 2A Prongs One and Two, Example 48 emphasizes the recitation of mathematical concepts, but the analysis fails to offer clarity or consistency when identifying abstract ideas within the provided claims.***

Mathematical operations inevitably are involved in AI-related inventions. The MPEP and case law instruct that the mere involvement of such operations does not automatically render a claim "directed to" an abstract idea. The Step 2A Prong One analyses for the respective claims within Example 48 do not adequately distinguish between claims that merely involve or are based on a mathematical concept and claims that recite a mathematical concept.

AIPLA believes that the formula recited within Claim 1, as currently written in its generic form rather than a true closed-form equation, is a better example of a §112 issue than an example of an abstract idea under §101.[21] The recited global function $V = f_\theta(X)$ is such a broad genus that it includes a nearly infinite volume of species. Hence, there is limited benefit in assessing whether the claim is directed to a mathematical concept due to the extent of indefiniteness hinders the ability to discern a meaningful BRI. The technology described in the Background section would not typically be claimed in this way by a skilled practitioner. AIPLA believes that there is clear value present in the technological focus of Example 48. As such, there is considerable potential benefit to the Office providing an example of a Prong One abstract idea analysis on a revised claim that is not confounded by §112 issues.

---

[21] *See* MPEP §2173.04 ("For example, a genus claim that covers multiple species is broad, but is not indefinite because of its breadth, which is otherwise clear. But a genus claim that could be interpreted in such a way that it is not clear which species are covered would be indefinite (e.g., because there is more than one reasonable interpretation of what species are included in the claim).").

If Claim 1 were subject to real-world examination, Claim 1 would be rejected under §112(b) and the examiner would make their best attempt at interpreting the claim for the purposes of articulating any other applicable rejections as early in the prosecution process as is feasible.[22] When the claim limitation reciting the generic global function is interpreted in light of the Background disclosure, the limitation is directed to a feature embedding technique that optimizes scedasticity of the data spread to improve separability of data clusters. If the claim limitation were presented in a way that was more authentically evocative of the technology itself, we think that step (c) of Claim 1 would instead be merely based on mathematical concepts rather than reciting a mathematical concept. This is analogous to the training example provided in MPEP 2106.04(a)(1), which recites the limitation of "applying one or more transformations to the digital images." It is readily apparent to the PHOSITA that applying transformations to data, such as digital images, could also be stated in terms of using a function having the form of $y = f_i(X)$ to execute the transformation, but there is no logical reason as to why one would.

The Step 2A Prong Two analyses fail to adequately consider how the DNN operates in the context of the claims. Example 48 dismisses the DNN as a "generic tool" for making mathematical calculations without acknowledging the specialized nature of the network's training and the way the network is specifically applied to solve the speech separation problem. The DNN is a highly specialized system tailored to perform speech separation, which should qualify as a technical application of AI rather than being characterized as a generic computer component.

Example 48 presents a ripe opportunity for the Office to provide clarification in response to existing ambiguity and flexibility in applying Step 2A Prong Two analysis to computer-implemented inventions. Applicants often find themselves in a "trap" in which AI-related claim limitations are at high risk of being arbitrarily dismissed as merely using a generic computer as a tool without consideration as to the highly specialized hardware that is necessary and integral to the operation of many AI inventions. Example 48 currently lacks meaningful discussion distinguishing between improvements to a technical field (i.e., speech separation) or a particular machine (i.e., the specialized circuitry related to Example 48) integral to the claim. Improvement to the claims in Example 48 would enable the Office to provide additional guidance on evaluating general-purpose vs. special-purpose computers, from which applicants and examiners alike would benefit considerably.

---

[22] *See* MPEP 706, 707.08(g).

Further, the analyses do not appear to provide any insight as to how the Prong Two practical application can be within the field of AI itself, rather than an unrelated field of use. Per MPEP 2106.04(a)(2)(I), claims are directed to mathematical concepts only when the mathematical concept is claimed in the abstract without meaningful integration to, e.g., a technological context. However, if AI itself is being considered an abstract idea in and of itself such that an advancement to the operation or efficiency of AI technology fails to be recognized as a practical application, then this inherently puts AI claims at a significant disadvantage as compared to other fields of computer-implemented technology without clear rationale. Similar issues arise within the context of inventive concept under Step 2B.

9. *Example 48 fails to recognize inventive concepts and technological improvements within the field of AI itself.*

The Example 48 analysis at Step 2B provides examples of inventive concepts being derived from other fields of use but not from AI technology, implicitly dismissing legitimate technological improvements that may exist for an AI claim. AIPLA is concerned that this reductive approach ignores the technological environment of, and concrete innovations within, AI technology. While at least Claim 1 of Example 48 is likely unpatentable, in its current form,[23] for reasons of anticipation or obviousness, AIPLA disagrees with the assessment that the respective claim elements identified as additional elements recite WURC activity. The Step 2B analysis of Claim 1, for example, alleges that using mixed speech signal data as input to a trained DNN is well-understood, routine, conventional activity ("WURC activity") because the Background discloses that smartphones and other devices with integrated microphones are able to receive a mixed speech signal. Oversimplification of the claimed technology disregards the technical operations performed by a trained DNN and seemingly implies that providing data as input to a trained AI model is synonymous with the collecting of the data itself. Input data must first be collected by a system component such as a microphone prior to being received by a trained DNN for processing. This concept is readily apparent to a user of ordinary skill in the art and is further supported by the Background disclosure, which explicitly states that the trained DNN receives the mixed speech signal from an audio recording device such as a microphone.

---

[23] We emphasize that adding temporal features as model inputs alongside the spectrogram features could be novel and inventive if the §112 problems with this claim element were overcome.

The analysis fails to recognize that a combination of WURC activity elements may amount to an inventive concept when analysis is applied to the claim as a whole.[24]

The search for inventive concept within Step 2B is distinct from a novelty or obviousness determination.[25] Yet, searching for inventive concept in practice is often difficult to distinguish from a novelty or obviousness determination. The subject matter of a claim may be anticipated by, or obvious in view of, the prior art (without amounting to WURC activity. This distinction could benefit from more clarity to avoid scenarios like Example 48. It is instead more accurate from a holistic perspective to say that the claimed invention in Example 48 is non-novel and obvious in light of the prior art, given lack of definition of temporal features, rather than providing a meaningful example of WURC activity.

10.   *Example 48 Summary*

While AIPLA again appreciates the Office's efforts to provide useful guidance for the subject matter eligibility of deep learning technology, Example 48 is hampered by technical inaccuracies reflected within the claims, misidentification of abstract ideas, and a failure to fully recognize the potential teaching opportunities for practical applications and improvements. Similar to other examples, the value of Example 48 would increase substantially by claim revisions that more accurately reflect the technical nature of the invention and provide clearer, more consistent guidance on the eligibility of AI-related inventions.

AIPLA suggests that Example 48 is reviewed with an elevated level of scrutiny to ensure that the analysis is performed in compliance with both MPEP Chapter 2100 and judicial precedent. AIPLA also suggests that the Office consider modifying the claiming strategy to be more realistic to maximize the potential practical benefit for applicants and examiners. In its present form, it is difficult to understand the driving factors that led the Office to formulate Example 48 in its current form and what is to be learned from the example analyses of Claims 1-3. This deficiency could be cured with revision to the claims and a clarified rationale.

---

[24] See *Diamond v. Diehr,* 450 U.S. at 188, 209 USPQ at 9 (1981).

[25] "As made clear by the courts, the "'novelty' of any element or steps in a process, or even of the process itself, is of *no relevance* in determining whether the subject matter of a claim falls within the § 101 categories of possibly patentable subject matter.'" MPEP §2106.05, quoting *Intellectual Ventures I v. Symantec Corp.,* 838 F.3d 1307, 1315, 120 USPQ2d 1353, 1358 (Fed. Cir. 2016) (quoting *Diamond v. Diehr,* 450 U.S. at 188–89, 209 USPQ at 9).

DETAILED COMMENTS ON EXAMPLE 49

Example 49 explores the critical domain of leveraging AI to inform treatment decisions. AIPLA thanks the Office for addressing an intersection that significantly amplifies the importance of both AI innovation and strategic treatment decision-making. Given the complexities involved in eligibility assessments and claim drafting priorities, it is very important that the Office provide applicants and examiners useful feedback about how innovations in this area can be protected using patents.

AIPLA's concerns with Example 49 are as follows:

- **1) Example 49 currently seems to imply that a new compound is the only patentable aspect of the Example claims and disregards how to protect novel AI innovations itself in the biology context.** Composition claims have presented a solid avenue to specifically protect new compounds for many decades. Meanwhile, the only eligible aspect identified in Example 49 is merely the identification of a new compound within Claim 2. AIPLA understands that the purpose of the new Office AI-focused guidance and example is to illustrate how to protect AI innovations (not just new compounds).

- **2) Both claims of Example 49 present <u>divided-infringement issues</u> and include <u>steps that would be performed by people</u>** (in addition to the software aspects). Applicants and examiners would benefit from identifying types of claims that do not introduce divided-infringement issues and that do not include human-performed actions (that are eligible) in the bioinformatics/artificial-intelligence space.

- **3) Proposed alternative claims are provided that may be useful to include in the eligibility example.**

11.   Novelty Focuses and AI Eligibility Analysis in Example 49

In the provided Background section of Example 49, at least five potential innovations can be discerned, each with implications for how the claims might cover these innovations. The five potential innovations that we have identified are provided below.

As illustrated in the list below, the present claims fail to clearly focus on any of the potential five innovations. At best, it appears that Claim 2 of Example 49 attempts to weave in compound novelty with other unfocused details that may pertain to AI. This makes it very difficult to apply proposed Example 49 to a given real-world AI claim.

1. **A New Anti-Fibrotic Drug (Compound X):** The claims fail to explicitly cover the compound. Claim 2 only mentions administration as involving the compound, potentially excluding compositions that incorporate other elements alongside the compound. This narrows the scope significantly if Compound X is novel and non-obvious.

2. **A New Method of Using Compound X Post-Microstent Surgery to Reduce Scarring:** The claims do not specifically cover the administration of Compound X in the defined circumstances of post-surgery or for the stated outcome of reducing scarring.

3. **Identification of Informative SNPs:** While the background mentions 37 single nucleotide polymorphisms ("SNPs") identified as genetic markers for inflammation post-implantation, the claims lack specificity in identifying these SNPs, relying instead on a potentially vague definition in the specification.

4. **A Machine Learning Model (ezAI) for SNP Data Transformation:** Neither the claims nor the specification detail the model's architecture, input data, or any processing involved, which is crucial given the claim that post-implantation inflammation is a polygenic condition. The absence of details on any bespoke model enhancements is a significant oversight.

5. **A Criterion for Predicting High-Risk Post-Implantation Inflammation (PI) Based on PRS:** Claim 1 sketches a high-level process for how the ezAI model generates the polygenic risk score ("PRS") but lacks details on how the model's output leads to a high-risk determination, leaving the practical application of this determination unclear.

AIPLA believes that each of potential innovations #3-5 are pertinent to inform eligibility guidance in the AI space. While potential innovations #1-2 are not individually likely to inform eligibility guidance in the AI space, if AI was used to identify the new compound/method, an example focused on such an instance may be informative in the AI context. Further, it appears as though the proposed Example 49 is attempting to avoid presenting a fact pattern where a novel AI processing technique is used to identify the informative SNPs or the polygenic risk score. Such a novel technique may include a new model architecture, new loss function, new

embedding technique, new generation of, or use of, confidence metrics of one or more SNPs, etc.

Thus, proposed Example 49 misses an opportunity to focus on many different AI specifics that occur in the bioinformatics context (e.g., a new AI approach for identifying informative SNPs, a new AI technique for generating the PRS, a new AI technique for predicting high-risk subjects based on the PRS).

12. *Divided Infringement Concerns within Example 49*

When Examples are published by the Office, both examiners and applicants consider this information to inform how to draft and examine claims. Therefore, it is important that the claims are of a format and type that would likely be considered and filed by applicants. For at least the reasons below, AIPLA believes that claims with other structure would be more attempting to protect AI technology in the biology context.

Claim 1 includes many steps, each of which would likely be performed by a different entity. For example:

- "collecting … a sample from a glaucoma patient" would likely be performed by a doctor or nurse.
- "genotyping [the] sample" would likely be performed by a laboratory;
- "identifying the glaucoma patient as high risk … based on a weighted polygenic risk score that is generated from informative … SNPs in the genotype dataset by an ezAI model" is performed by software; and
- "administering an appropriate treatment" would likely be performed by the patient, given that the specification and Claim 2 indicates that the treatment may be eye drops.

Practitioners and applicants strive to avoid a claim with this type of divided-infringement potential. The AI component of the claim is limited to that from Claim 1 step (b) ("identifying the glaucoma patient …"). Inclusion of the collecting, genotyping and administering steps leads to confusion in terms of what type of single-actor and/or AI-focused claim would be allowable. Even the preamble, which indicates that the claimed method is a "treatment method," provides divided-infringement issues.

13. *Legal and Practical Concerns with the Claim Drafting of Example 49*

Under 35 U.S.C. §287(c), protections are limited when actions are performed by a medical practitioner ("any natural person who is licensed by a State to provide the medical activity described in subsection (c)(1) or who is acting under the direction or such person in performance of the medical activity"). Both claims of Example 49 include steps likely to be performed by medical professionals (e.g., the "collecting," "genotyping" and potentially the "administering" steps identified above). It would be helpful if Example 49 instead focused on providing guidance as to what type of software and computer-implemented AI innovations could be protected in the AI space, without confounding the example with actions that may fall under the 287(c) limitation. Again, even the preamble (indicating that the claim is "treatment method" presents concerns).

14.   *Proposed Alternative Claims and Example Revisions*

Example 49 could be revised to replace the two claims in the Example 49 with the three claims below. The Background section of the example should also be updated to include additional detail about the loss function and training of the machine learning model (in accordance with claim 3 below), additional detail about a second machine learning model used to predict the 37 SNPs (in accordance with claim 2 below), and Table 1 that identifies the 37 SNPs (in accordance with claims 1 and 2 below).

[Claim 1] A computer-implemented method comprising:

(a)   accessing a DNA sequence data set corresponding to a subject that represents a plurality of DNA sequence data sets spanning a plurality of DNA loci;

(b)   generating, based on the DNA sequence data set, a SNP-summary data set that predicts, for each locus of a set of preidentified loci, whether DNA of the subject includes a preidentified single nucleotide polymorphism (SNP) at the locus;

(c)   generating, by processing the SNP-summary data set using a machine learning model, a polygenic risk score.

(d)   classifying the subject as being at high risk of post-implantation inflammation (PI) based on the polygenic risk score; and

(e)   outputting a recommendation that the subject receive an anti-fibrotic treatment after any microstent implant surgery.

[Claim 2] The computer-implemented method of claim 1, wherein:

    (a)   the preidentified loci were identified using a second machine learning model comprising a deep convolutional neural network (CNN) with residual connections and dilated convolutions configured to capture non-additive genetic interactions and epistasis;

    (b)   the set of preidentified loci includes each locus identified in Table 1;

    (c)   wherein the preidentified SNPs comprise each SNP identified in Table 1.

[Claim 3] The computer-implemented method of claim 1, wherein the machine learning model was trained using a loss function and a training data set, wherein:

    (a)   the training data set comprises, for each other subject of a set of other subjects:

    a.   another SNP-summary data set; and

    b.   label data that indicates whether the other subject experienced post-implantation inflammation (PI), and, if so:

    i.   a delay between onset of the PI and completion of a corresponding microstent surgery, wherein the training data set indicates that at least one other subject of the set of other subjects experienced PI; and

    ii.   a severity of the PI; and

    (b)   the loss function is configured to minimize a weighted combination of: (i) a binary cross-entropy loss that penalizes incorrect predictions of whether PI occurred; (ii) a mean squared error loss that penalizes deviations between the predicted and actual delays of onset of PI for subjects who experienced PI; and (iii) a ranking loss that penalizes incorrect predictions of inflammation severity by prioritizing more severe cases in the training process.

The proposed alternative claims are drafted in an attempt to provide example claims that aim to avoid the above-described claim drafting concerns and that additionally provide elements of different types of innovations arising at the intersection of AI and biology. For example, the

claims are drafted to attempt to avoid any divided-infringement concern and to avoid actions performed by a medical practitioner.

15.   *Example 49 Summary*

Though biological innovations have been tied to the laboratory for many decades now, there are many instances where the innovations are (additionally or alternatively) tied to new computational techniques and results. Example 49 clearly demonstrates that the Office acknowledges the importance of providing guidance and examples that pertain to the era where many of the most important advancements in biology and medicine are rooted in and/or made possible due to innovations that bridge AI and biology.

However, AIPLA believes that Example 49 currently fails to illustrate how to protect AI inventions and instead, focuses on innovation in other fields that are not directly related to the AI itself. Example 49 has the potential to provide needed guidance focused on AI inventions relating to bioinformatics but lacks sufficient focus on the AI elements themselves. Furthermore, claim drafting concerns related to divided infringement and limited protections under 35 U.S.C. §287(c) limit the practical benefit of the example claims. Example 49 introduces an opportunity for the Office to provide guidance about the eligibility of many types of innovation arising in the bioinformatics (beyond merely focusing on novelty of a new compound or use thereof).