

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

WEALTH, INC.

Petitioner,

v.

VANILLA TECHNOLOGIES INC.,

Patent Owner.

PGR2026-00042

U.S. Patent No. 12,353,917

**PATENT OWNER VANILLA TECHNOLOGIES INC.'S
PRELIMINARY RESPONSE**

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	VANILLA’S ’917 PATENT.....	4
A.	Overview of the ’917 Claims.....	4
B.	Overview of the ’917 Patent Disclosure	8
C.	Overview of the Prosecution History.....	17
III.	LEVEL OF ORDINARY SKILL IN THE ART	19
IV.	CLAIM CONSTRUCTION	19
V.	WEALTH’S § 101 CHALLENGE FAILS	20
A.	Wealth Fails to Establish That the ’917 Claims Are Directed to an Abstract Idea under <i>Alice</i> Step One / USPTO Step 2A	21
B.	Wealth Fails to Establish That the ’917 Claims Lack an Inventive Concept Under <i>Alice</i> Step 2 / USPTO Step 2B	40
C.	Summary.....	44
VI.	WEALTH’S § 103 CHALLENGES FAIL	44
A.	Wealth Fails to Establish that Claims 1, 2, 6-7, 12-14, 17, and 19 are Rendered Obvious by the Combination of Racanelli, Bateman, and Dintenfass (Ground 2)	46
1.	The combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest the claimed sequential execution stages of the “ <i>one or more machine-learning models</i> ”	46
2.	The combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest generating “ <i>resource data</i> ” based on “ <i>inputting the one or more responses and rule data comprising one or more rules associated with the resources into the one or more machine-learning models</i> ”	51
3.	The combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest that “ <i>the one or more machine- learning models are trained based on training data comprising training resource data and training rule data</i> ”	60

4.	Summary.....	63
B.	Wealth Fails to Establish that Claims 3 and 8-11 are Rendered Obvious by the Combination of Racanelli, Bateman, Dintenfass, and McMillen (Ground 3).....	63
C.	Wealth Fails to Establish that Claims 15, 16, 18, and 20 are Rendered Obvious by the Combination of Racanelli, Bateman, Dintenfass, and Kwak (Ground 4).....	64
VII.	WEALTH’S § 112 CHALLENGE OF DEPENDENT CLAIMS 4-5 IS MOOT IN VIEW OF VANILLA’S STATUTORY DISCLAIMER.....	64
VIII.	WEALTH’S § 112 CHALLENGE OF DEPENDENT CLAIM 13 FAILS.....	65
IX.	CONCLUSION.....	70

TABLE OF AUTHORITIES

Cases

<i>Aon Re, Inc. v. Zesty.ai, Inc.</i> , 791 F. Supp. 3d 531 (D. Del. 2025)	passim
<i>Bascom Global Internet Services, Inc. v. AT&T Mobility LLC</i> , 827 F.3d 1341 (Fed. Cir. 2016)	41
<i>DDR Holdings, LLC v. Hotels.com, L.P.</i> , 773 F.3d 1245 (Fed. Cir. 2014)	40
<i>Enfish, LLC v. Microsoft Corp.</i> , 822 F.3d 1327 (Fed. Cir. 2016)	22, 23, 35
<i>Hutchins v. Zoll Medical Corp.</i> , 492 F.3d 1377 (Fed. Cir. 2007)	66, 67
<i>McRO, Inc. v. Bandai Namco Games America, Inc.</i> , 837 F.3d 1299 (Fed. Cir. 2016)	23
<i>Multilayer Stretch Cling Film Holdings, Inc. v. Berry Plastics Corp.</i> , 831 F.3d 1350 (Fed. Cir. 2016)	68, 69
<i>Pfizer, Inc. v. Ranbaxy Laboratories Ltd.</i> , 457 F.3d 1284 (Fed. Cir. 2006)	69
<i>Ravo v. Covidien LP</i> , No. 11-1637, 2014 WL 198551 (W.D. Pa. Jan. 16, 2014)	66, 67
<i>Raytheon Technologies Corp. v. General Electric Company</i> , 993 F.3d 1374 (Fed. Cir. 2021)	65
<i>Recentive Analytics, Inc. v. Fox Corp.</i> , 134 F.4th 1205 (Fed. Cir. 2025)	25, 37

Statutes

35 U.S.C. § 112(d).....	65, 67, 68
35 U.S.C. § 282(b).....	19

Other Authorities

37 C.F.R. § 42.200(b).....	20
37 C.F.R. § 42.207(e).....	64
M.P.E.P. § 2106.04.....	22, 23, 24, 33, 35
M.P.E.P. § 2106.05.....	41
Notice, <i>2024 Guidance Update on Patent Subject Matter Eligibility, Including on Artificial Intelligence</i> , 89 Fed. Reg. 58,128 (July 17, 2024).....	2, 20-21, 25, 34
U.S. Patent & Trademark Office, Advance Notice of Change to the MPEP in Light of <i>Ex Parte Desjardins</i> , 3 (Dec. 5, 2025).....	22

EXHIBIT LIST

Exhibit	Description
2001	Expert Declaration of Christopher Thompson
2002	Statutory disclaimer of dependent claims 4-5 of U.S. Pat. No. 12,353,917
2003	U.S. Pat. No. 10,911,811
2004	U.S. Pat. No. 11,386,367

I. INTRODUCTION

Patent Owner Vanilla Technologies Inc. (“Vanilla” or “Patent Owner”) submits this Preliminary Response to the Petition for Post-Grant Review No. PGR2026-00042 (“Petition” or “Pet.”) filed on April 7, 2026 by Petitioner Wealth, Inc. (“Wealth” or “Petitioner”), which challenges claims 1-20 of U.S. Patent No. 12,353,917 (the “’917 Patent”) on the following grounds:

Ground	Alleged Basis for Unpatentability
1	Claims 1-20 are allegedly ineligible under § 101
2	Claims 1, 2, 4-7, 12-14, 17, 19 are allegedly rendered obvious under § 103 by a combination of Int’l Patent App. Pub. No. WO2020/028252 (“Racanelli”), U.S. Patent App. Pub. No. 2015/0286630 (“Bateman”), and U.S. Patent App. Pub. No. 2017/0076381 (“Dintenfass”)
3	Claims 3 and 8-11 are allegedly rendered obvious under § 103 by a combination of Racanelli, Bateman, Dintenfass, and U.S. Patent App. Pub. No. 2025/0245763 (“McMillen”)
4	Claims 15, 16, 18, and 20 are allegedly rendered obvious under § 103 by a combination of Racanelli, Bateman, Dintenfass, and a document entitled “Validity Assessment of Legal Will Statements as Natural Language Inference” by Alice Saebom Kwak et al. (“Kwak”)
5	Claims 4 and 5 are allegedly indefinite under § 112(b)
6	Claim 13 is allegedly an improper dependent claim under § 112(d)

Wealth’s Petition fails to establish that at least one claim of the ’917 Patent is more likely than not to be unpatentable under any of these grounds.¹

¹ As explained further below, Vanilla has filed a statutory disclaimer of dependent claims 4-5, which moots Wealth’s § 112(b) challenge.

The '917 Patent claims a machine-learning (“ML”)-based architecture comprising five separate and distinct stages of functionality that work together to generate structured “*resource data*” and then generate and update “*one or more visualizations of the distribution of the one or more assets to the one or more resource recipients.*” That claimed ML-based architecture is built on a specifically arranged implementation of machine learning involving a sequence of two distinct execution stages of “*one or more machine-learning models*” that perform separate and clearly defined analytical tasks, along with an intervening receipt of “*one or more responses to the one or more prompts*” that interconnects the output of the first execution stage with the input of the second execution stage of the “*one or more machine-learning models.*”

Wealth’s § 101 challenge fails for two independent reasons, each of which serves as a separate basis for rejecting Wealth’s § 101 challenge. First, the claimed five-stage ML-based architecture is a specifically arranged implementation of machine learning and therefore is not directed to an abstract idea under *Alice* step one / USPTO Step 2A. See *Aon Re, Inc. v. Zesty.ai, Inc.*, 791 F. Supp. 3d 531, 540 (D. Del. 2025); Notice, *2024 Guidance Update on Patent Subject Matter Eligibility, Including on Artificial Intelligence*, 89 Fed. Reg. 58,128, 58,137 (July 17, 2024). Second, the claimed five-stage ML-based architecture is a non-conventional and

non-generic arrangement that supplies an inventive concept under *Alice* step two / USPTO Step 2B.

The § 101 eligibility of the '917 Patent is further supported by the fact that the Office allowed these claims under the applicable § 101 framework (including *Recentive*, which was decided prior to allowance) less than 15 months ago. While Wealth now asks the Office to revisit that decision, it fails to identify any overlooked fact or intervening change in law that supports its § 101 challenge. Instead, Wealth's § 101 challenge rests on a series of conclusory allegations that overgeneralize the '917 Claims and fail to apply the governing legal framework for § 101 eligibility. For these reasons, Wealth's § 101 challenge should be rejected.

Wealth's § 103 challenges likewise fail. As established below, Wealth's alleged § 103 combinations fail to render obvious at least three key requirements of the claimed five-stage ML-based architecture required by the '917 Patent, which are (1) the claimed sequence of two separate and distinct execution stages of the claimed "*one or more machine-learning models*," (2) the claimed requirement of generating the claimed "*resource data*" based on "*inputting the one or more responses and rule data comprising one or more rules associated with the resources into the one or more machine-learning models*," and (3) the requirement that "*the one or more machine-learning models are trained based on training data comprising training resource data and training rule data*."

Lastly, Wealth's § 112(d) challenge of dependent claim 13 fails. Under § 112(d), dependent claim 13 statutorily incorporates all of the limitations of independent claim 1. The claim's abbreviated reference to one of the limitations of independent claim 1 does not change this fact, and dependent claim 13 is otherwise fully compliant with the requirements of § 112(d).

Thus, because Wealth fails to establish that at least one claim is more likely than not to be unpatentable under any ground, Wealth's Petition should be denied.

In support of this Preliminary Response, Vanilla is also submitting a declaration from technical expert Christopher Thompson. *See* Ex. 2001.

II. VANILLA'S '917 PATENT

A. Overview of the '917 Claims

The '917 Patent issued with claims 1-20, of which claims 1, 17, and 19 are independent and the remainder are dependent. Vanilla has now filed a statutory disclaimer of dependent claims 4-5, which leaves claims 1-3 and 6-20. Ex. 2002. These claims of the '917 Patent will be referred to herein as the "'917 Claims."

Independent claim 1, which is representative of the independent limitations that are included in each of the '917 Claims, is reproduced below with the same claim limitation numbering utilized by Wealth's Petition:

1. A computer-implemented method comprising:

1[a] generating, by a computing system comprising one or more processors, based on inputting one or more resource documents into one or more machine-learning models, one or more prompts comprising requests for

information associated with resource allocation instructions for distribution of resources, wherein the one or more machine-learning models are configured to generate the one or more prompts based on recognition of textual features of the one or more resource documents associated with the resources, wherein the one or more resource documents comprise one or more resource document fields, and wherein the resource allocation instructions are associated with distribution of one or more assets of the resources to one or more entities comprising one or more resource recipients;

1[b] receiving, by the computing system, one or more responses to the one or more prompts;

1[c] generating, by the computing system, based on inputting the one or more responses and rule data comprising one or more rules associated with the resources into the one or more machine-learning models, resource data comprising one or more resource data fields and one or more resource data field values that correspond to the one or more resource data fields, wherein the one or more machine-learning models are trained based on training data comprising training resource data and training rule data, and wherein the one or more resource data fields are based on the one or more resource document fields, and wherein the one or more resource data field values are based on the one or more responses;

1[d] generating, by the computing system, based on the resource data, one or more indications associated with the resource allocation instructions, wherein the one or more indications comprise one or more visualizations of the distribution of the one or more assets to the one or more resource recipients; and

1[e] updating, by the computing system, the one or more visualizations in real-time based on changes in one or more values of the one or more assets.

Ex. 1001 [’917 Patent] at 39:1-42. In this Preliminary Response, Vanilla’s discussion of the claim language is with reference to the representative independent limitations of independent claim 1 unless specifically noted otherwise.


As demonstrated above, the ’917 Claims recite a specific ML-based architecture for generating “*resource data*” that is then utilized to generate “*one or*

more visualizations of the distribution of the one or more assets to the one or more resource recipients,” which are also updated “in real-time based on changes in one or more values of the one or more assets.” Id. The claimed ML-based architecture involves five separate and distinct stages of functionality:

- (1) a first execution stage of the “*one or more machine-learning models*” that generates “*one or more prompts*” based on “*recognition of textual features of the one or more resource documents*” that are input into the first execution stage;
- (2) an intervening stage for receiving “*one or more responses to the one or more prompts*” that are generated during the first execution stage of the “*one or more machine-learning models*”;
- (3) a second execution stage of the “*one or more machine-learning models*” that generates structured “*resource data*” based on a combination of the “*one or more [prompt] responses*” and “*rule data comprising one or more rules associated with the resources*” (where the “*one or more machine-learning models*” are trained based on “*training data comprising training resource data and training rule data*,” the “*resource data fields*” are based on “*resource document fields*,” and the “*resource data field values*” are based on the “*one or more [prompt] responses*”);
- (4) a generation stage for “*one or more visualizations of the distribution of the one or more assets to the one or more resource recipients*” based on the structured “*resource data*”; and
- (5) a real-time updating stage for the “*one or more visualizations*” based on “*changes in one or more values of the one or more assets.*”


Id.; see also Ex. 2001 [Thompson Decl.], ¶ 63.

This five-stage ML-based architecture required by each of the ’917 Claims is shown in the following illustrative diagram from Mr. Thompson’s Declaration:



Ex. 2001 [Thompson Decl.], ¶ 64.

Notably, this claimed ML-based architecture involves a specifically-arranged implementation of machine learning for generating “*resource data*,” which involves two distinct execution stages of the claimed “*one or more machine-learning models*” that perform separate and clearly defined analytical tasks along with an intervening stage for receiving “*one or more responses to the one or more prompts*” that are generated during the first execution stage of the “*one or more machine-learning models*.” *Id.*, ¶ 65. These two execution stages are highlighted in the following annotated version of the above illustrative diagram:



Id.

The dependent limitations of the '917 Claims also add other details to the claimed ML-based architecture, including additional execution stages of the claimed “*one or more machine learning models*” and generation of additional types of “*indications*” based on the “*resource data*.” See Ex. 1001 ['917 Patent] at 39:43-41:56; Ex. 2001 [Thompson Decl.], ¶ 66.

B. Overview of the '917 Patent Disclosure

The background of the '917 Patent begins by explaining that the existing technology for managing resource processing and/or allocation involved “different types of software [that] can be used to collect, store, and disseminate the resource information to others,” but there were problems with this existing software technology. Ex. 1001 ['917 Patent] at 1:26-34. First, “the use of multiple software

applications to manage resources can be burdensome and demand a significant amount of manual input.” *Id.* at 1:29-31. Second, “generalized software applications may not be fully suited to meeting the often specialized requirements of resource allocation.” *Id.* at 1:32-34. The disclosed technology is intended to overcome these problems with the existing software technology.

The ’917 Patent introduces the disclosed technology as follows:

In general, the present disclosure is directed to generating and processing data including resource data associated with resource processing and allocation. In accordance with this disclosure, a personalized analysis of resource allocation instructions including data relating to the assets and entities associated with the resources can be generated and used to automatically deal with various aspects of resources planning. Further, the resource data that is generated can be based on resource allocation instructions of an entity associated with the resources (e.g., a network administrator, a computing systems architect, and/or an asset administrator) and can include information associated with assets of the resources and the distribution of those assets to other entities including resource recipients of the resources. In some embodiments, the disclosed technology can receive information from various sources and generate resource data that can be used to generate indications including visualizations of different aspects of the resources. The disclosed technology can also implement machine-learning models that have been configured and/or trained to generate various projections (e.g., asset valuation projections) or predictions that can be used to improve resource processing and allocation.

Additionally, the present disclosure is directed to improved computer systems, computer applications, computer-implemented methods, user interfaces, and/or services for resource processing and/or allocation. In some instances, the resource allocation systems and methods described herein can be used to structure and visualize an existing resources plan, strategize potential opportunities, offer online document creation, and create a collaborative online space for accessing current resources plans and documents. Further, the systems and methods described herein can

account for complex fact patterns, multiple types of trust documents, state, and federal tax laws, and have a robust reporting engine.

Id. at 3:66-4:33.

One aspect of the disclosed technology involves a computing system that is configured to implement the claimed ML-based architecture that is illustrated in the above diagram, which includes multiple different stages of functionality that work together to “generate resource data and visualizations of the resource data that are based on specific user responses to prompts associated with resources.” *Id.* at 5:53-6:18; Ex. 2001 [Thompson Decl.], ¶ 70. A summary of ’917 Patent’s disclosure related to each of the five stages of the claimed ML-based architecture is set forth below. *See also* Ex. 2001 [Thompson Decl.], ¶¶ 71-83.

Stage (1): generating “one or more prompts”

The ’917 Patent discloses that “a computing system can generate one or more prompts that are associated with requests for information associated with resource allocation instructions for distribution of resources” based on “one or more resource documents associated with the resources,” which can be accomplished by inputting the “one or more resource documents” into “one or more machine-learning models” that are configured and/or trained to generate the one or more prompts based on “recognition of textual and/or semantic features of the one or more resource documents.” *Id.* at 4:58-5:19, 7:12-24.

In connection with this first-stage functionality, the '917 Patent further explains that:

- The “one or more resource documents” may take the form of “one or more documents associated with the resources and/or the resource allocation instructions,” examples of which may include “a network policy, authentication policy, legal documents, estate documents, a will, deeds to property, beneficiary designations, a power of attorney, insurance policy, healthcare directives and/or other types of documents that are associated with resources.” *Id.* at 5:5-11, 7:32-63.
- The “one or more resource documents” comprise “one or more resource document fields,” examples of which may include “a document field associated with the name of a resource recipient and/or a document field associated with the address of real property, a network address (e.g., IP address), and/or the name of tangible personal property comprising a piece of artwork (e.g., a painting or sculpture).” *Id.* at 5:11-16, 7:47-63.
- The “one or more prompts” may take the form of “questions that can be used to direct a user to enter responses that can be used to generate resource data that is then processed and used to provide information, updates, and notifications to a user,” examples of which may include a request for information about an “asset” (e.g., “the state, address, estimated value, and owner of an asset comprising real property”), “a person (e.g., a name, date of birth, address, and/or contact information of a person),” “a device (e.g., the software, processor, storage, and/or memory configuration of a computing device), and/or “an organization (e.g., the name of the organization, the geographic location in which the organization is registered and/or incorporated).” *Id.* at 4:64-5:3, 6:26-53.

See also Ex. 2001 [Thompson Decl.], ¶¶ 71-73.

Stage (2): receiving “one or more [prompt] responses”

The '917 Patent discloses that the “one or more prompts” generated by the first stage of the ML-based architecture may be presented to a user via “an interface (e.g., a graphical user interface) of a display component associated with the

computing system” and/or “an audio output device (e.g., loudspeakers) associated with the computing system,” and that in turn, the user may provide “one or more responses” to the “one or more prompts” via “one or more input devices (e.g., a touchscreen, a mouse, a keyboard, microphone, stylus, and/or a numeric input pad).” Ex. 1001 [’917 Patent] at 5:20-24, 6:19-25, 8:10-21. For example, a user may provide a response to a “prompt requesting the name of a resource recipient which can be associated with a resource recipient field.” *Id.* at 5:20-24.

This intervening functionality of prompting a user to supplement the data in the “one or more resource documents” prior to generating the “resource data” is one specific aspect of the ML-based architecture that improves over the existing software technology. *See id.* at 5:56-65 (explaining that “the disclosed technology can generate resource data and visualizations of the resource data that are based on specific user responses to prompts associated with resources” and thereby “assist a user in more effectively performing the technical task of resource data processing by means of a continued and/or guided human-machine interaction process in which responses to generated prompts based on resource documents are received and the disclosed technology generates real-time visualizations based on continuously updated resource data”); *see also* Ex. 2001 [Thompson Decl.], ¶ 74.

Stage (3): generating “resource data”

The '917 Patent discloses that “[t]he computing system can generate resource data that is based on one or more responses and rule data that includes one or more rules associated with the resources (e.g., access rules and/or permissions),” and further discloses that “machine-learning models can be used in the process of generating personalized resource data.” Ex. 1001 ['917 Patent] at 5:27-31, 8:22-29, 14:56-58.

In connection with this third-stage functionality, the '917 Patent explains that:

- The “rule data can comprise one or more rules associated with the resources” and the “one or more rules can be associated with” (i) “one or more access policies, one or more authorization policies, one or more authentication policies, one or more computing device permissions, one or more file access permissions, and/or assignment of one or more assets (e.g., computing devices) to one or more entities (e.g., computing device users)” and/or (ii) “one or more laws (e.g., one or more federal laws and/or one or more state laws) that can be associated with various aspects of the resources including probate law, resources law, tax law, inheritance law, and/or property law.” *Id.* at 8:24-45.
- The “resource data” that is generated “comprise[s] one or more resource data fields and/or one or more resource data field values.” *Id.* at 8:46-47.
- The “resource data fields” may be “based on the one or more resource document fields” of the “one or more resource documents,” and illustrative examples of these fields may include a “name of an asset,” a “name of [an] intended resource recipient,” “an appraised value of [a] house,” and a “maturity date of a debt instrument.” *Id.* at 8:47-8:57.
- The “one or more resource data field values can correspond to the one or more resource data fields” and “can be based on the one or more responses.” For example, “if a response to a prompt requesting information about the intended resource recipient of an asset is a particular

name (e.g., ‘JOHN DOE’), the one or more resource data field value[s] associated with the name of the intended resource recipient may match the response (e.g., the resource data field can indicate ‘JOHN DOE’).” *Id.* at 8:51-52, 8:61-9:7.

See also Ex. 2001 [Thompson Decl.], ¶¶ 78-79.

Moreover, the ’917 Patent discloses that “one or more machine-learning models” can be trained with “training data” that includes (i) “training resource data that can comprise information associated with resources, processing resources, allocating resources, and/or generating resource allocation instructions” and (ii) “training rule data that can comprise information associated with one or more rules (e.g., one or more regulations) that determine the way in which assets are treated.” Ex. 1001 [’917 Patent] at 21:9-19; *see also id.* at 22:54-23:25 (describing the training process for the “one or more machine-learning models”); Ex. 2001 [Thompson Decl.], ¶¶ 76-77.

Stage (4): generating “one or more visualizations”

The ’917 Patent discloses that “[t]he computing system can generate indications that can include visualizations of the distribution of the one or more assets to the one or more resource recipients” based on the generated “resource data.” Ex. 1001 [’917 Patent] at 5:42-52, 9:42-49. For example, “based on resource data that indicates the specific assets (e.g., computing devices, documents, houses, vehicles, securities, and/or cash) to be distributed to various resource recipients, the computing system can generate visualizations that indicate which resource recipients

receive which assets and the amounts of the assets that each of the resource recipients receives.” *Id.* at 5:45-52. Such a “visualization” may take the form of “a chart that indicates the names of the one or more resource recipients and the assets associated with each resource recipient.” *Id.* at 9:44-49.

Stage (5): real-time updating “one or more visualizations”

The ’917 Patent further discloses that the “one or more visualizations” can be “updated in real-time or near real-time based on changes in the values of the one or more assets and/or the one or more resource recipients (e.g., resource recipients that are added or removed).” Ex. 1001 [’917 Patent] at 9:49-53.²

Technological Improvements

As the ’917 Patent explains, the disclosed technology improves over the existing software technology for resource allocation in a number of different ways:

The systems, methods, devices, computer-readable media (e.g., tangible non-transitory computer-readable media) in the disclosed technology can provide a ***variety of technical effects and benefits*** including an ***improvement in the generation of resource data*** that can be used in the analysis and implementation of resource allocation instructions. Further, the resource data can be personalized based on information specific to each individual entity associated with the resource allocation instructions. The disclosed technology can provide the ***technical effect of improving the effectiveness with which resources planning tasks are performed***. For example, the computing system can be continuously updated based on user responses to

² The ’917 Patent discloses various other details regarding the claimed ML-based architecture as well, including but not limited to details regarding additional stages of the ML-based architecture and details regarding generation of additional types of “indications” based on the structured “resource data.” *See id.* at 9:54-14:32.

prompts, changes in rules that govern resources, and/or the occurrence of key events that may change the way in which the assets associated with resources are distributed. In some embodiments, machine-learning models can be used in the process of generating personalized resource data. Further, the machine-learning models can be continuously trained and/or updated in response to updated resource data that is specific to the resources and/or resources plan. As a result, ***more relevant and timely resource data can be provided to a user*** based on specific information associated with that particular user's resource data. Further, the disclosed technology can generate visualizations that can include projected asset valuations and/or asset distributions that can be used to plan different scenarios associated with resources.

The disclosed technology can ***improve the operation of a resource processing device by more effectively performing a variety of tasks*** with the specific benefits of providing more accurate resources planning information. Further, the disclosed technology can use various rules-based techniques and/or machine-learning models to ***more efficiently process resource documents*** that would otherwise require time consuming and burdensome manual review and handling. Accordingly, the improvements offered by the disclosed technology can result in tangible benefits to a variety of devices and/or systems comprising computing systems, electronic systems, and/or mechanical systems associated with processing resource data.

Ex. 1001 ['917 Patent] at 14:41-15:13 (emphasis added); *see also id.* at 5:56-65 (explaining that “the disclosed technology can assist a user in more effectively performing the technical task of resource data processing by means of a continued and/or guided human-machine interaction process in which responses to generated prompts based on resource documents are received and the disclosed technology generates real-time visualizations based on continuously updated resource data,”

which “can allow for the generation of personalized resource data provision that can more effectively address a user’s questions”); Ex. 2001 [Thompson Decl.], ¶ 83.

C. Overview of the Prosecution History

The ’917 Patent was originally filed as U.S. Non-Provisional App. No. 18/963,034 (the “’034 Application”) on November 27, 2024, and it claims priority to U.S. Provisional App. No. 63/603,895 filed on November 29, 2023. Ex. 1001 [’917 Patent] at 1.

On January 27, 2025, the U.S. Patent and Trademark Office (“USPTO”) issued a first Non-Final Office Action that rejected original claims 1-20 of the ’034 Application as allegedly being directed to a judicial exception without significantly more under § 101. Ex. 1003 [File History of ’917 Patent] at 128-29.

On April 7, 2025, an Examiner interview was held to discuss the § 101 rejections. During that interview, “Applicant's representative presented proposed claim amendments to overcome the existing rejection under 35 U.S.C. 101” and the “Examiner expressed that the amendments would likely overcome the existing rejection, subject to further search and consideration.” *Id.* at 132.

On April 17, 2025, Vanilla filed a written response to the Non-Final Office Action, which included amendments to independent claims 1, 17, and 19 (as well as dependent claims 3, 5, and 8) along with corresponding remarks. *Id.* at 134-144.

Vanilla's amendment to independent claim 1, which is representative of the amendments made to independent claims 17 and 19, is reproduced below:

1. (Currently Amended) A computer-implemented method comprising:
generating, by a computing system comprising one or more processors, based on inputting one or more resource documents into one or more machine-learning models, one or more prompts comprising requests for information associated with resource allocation instructions for distribution of resources, wherein the one or more machine-learning models are configured to generate the one or more prompts ~~[[are]]~~ based on recognition of textual features of the one or more resource documents associated with the resources, wherein the one or more resource documents comprise one or more resource document fields, and wherein the resource allocation instructions are associated with distribution of one or more assets of the resources to one or more entities comprising one or more resource recipients;
receiving, by the computing system, one or more responses to the one or more prompts;
generating, by the computing system, based on inputting the one or more responses and rule data comprising one or more rules associated with the resources into the one or more machine-learning models, resource data comprising one or more resource data fields and one or more resource data field values that correspond to the one or more resource data fields, wherein the one or more machine-learning models are trained based on training data comprising training resource data and training rule data, and wherein the one or more resource data fields are based on the one or more resource document fields, and wherein the one or more resource data field values are based on the one or more responses; ~~and~~
generating, by the computing system, based on the resource data, one or more indications associated with the resource allocation instructions, wherein the one or more indications comprise one or more visualizations of the distribution of the one or more assets to the one or more resource recipients; and
updating, by the computing system, the one or more visualizations in real-time based on changes in one or more values of the one or more assets.

Id. at 134. As shown, Vanilla added technical details throughout the claims, including by specifying that:

- (1) the “*one or more prompts*” are generated “*based on inputting one or more resource documents into one or more machine-learning models*” that are “*configured to generate the one or more prompts based on recognition of textual features of the one or more resource documents*”;

- (2) the “*resource data*” is generated by “*inputting the one or more responses and rule data . . . into the one or more machine-learning models*” that are “*trained based on training data comprising training resource data and training rule data*”; and
- (3) the “*one or more visualizations*” are updated “*in real-time based on changes in one or more values of the one or more assets.*”

Id.

On May 2, 2025, the USPTO then mailed a Notice of Allowance that was responsive to Vanilla’s April 17, 2025 submission and allowed claims 1-20 as amended. *Id.* at 147-51. Vanilla subsequently paid the issue fee on June 5, 2025, and the ’917 Patent was issued on July 8, 2025. *Id.* at 182, 187.

III. LEVEL OF ORDINARY SKILL IN THE ART

In its Petition, Wealth advances the following level of ordinary skill in the art:

[A]n undergraduate degree in computer science or a related field, and three to five years of industry experience with artificial intelligence and machine learning. This level of skill is approximate, and more education could have compensated for less experience, and vice versa.

Pet. 7 (citing White Decl. ¶ 41).

For purposes of this Preliminary Response only, Vanilla applies this level of ordinary skill in the art.

IV. CLAIM CONSTRUCTION

In this proceeding, the claims are to be construed “using the same claim construction standard that would be used to construe the claim in a civil action under 35 U.S.C. § 282(b) including construing the claim in accordance with the ordinary

and customary meaning of such claim as understood by one of ordinary skill in the art.” 37 C.F.R. § 42.200(b).

Wealth’s Petition purports to use the “ordinary and customary meaning” of the ’917 Claims and states that “[t]he Board need not expressly construe any claim term at this stage.” Pet. 8. However, in reality, Wealth’s Petition relies on interpretations of several claim terms that are not consistent with the ordinary and customary meaning of those terms as understood by a person of ordinary skill in the art (“POSITA”) in the context of the ’917 Patent, including “*rule data comprising one or more rules associated with the resources*” and “*training rule data.*” Vanilla addresses the proper interpretation of these claim terms in the context of its response to Wealth’s § 103 challenges below.

V. WEALTH’S § 101 CHALLENGE FAILS

Wealth’s § 101 challenge fails for two independent reasons, each of which serves as a separate basis for rejecting Wealth’s § 101 challenge.

First, the ’917 Claims as a whole are directed to a structured ML-based architecture that is a specific implementation of machine learning and is therefore not directed to an abstract idea under *Alice* step one / USPTO step 2A. *See Aon*, 791 F. Supp. 3d at 540 (finding that “claiming a specifically arranged implementation of machine learning” is “another route to eligibility” under *Alice* step one); *2024*

Guidance Update, 89 Fed. Reg. at 58,137 (recognizing “a specific application of AI to a particular technological field” as eligible subject matter under USPTO Step 2A).

Second, when the limitations of the ’917 Claims are properly evaluated as an ordered combination, it is clear that the ’917 Claims recite a non-conventional and non-generic arrangement that supplies an inventive concept under *Alice* step two / USPTO Step 2B.

Wealth’s position to the contrary is based on a series of conclusory allegations that overgeneralize the ’917 Claims and fail to apply the governing legal framework for § 101 eligibility.

Thus, Wealth’s Petition fails to establish that any of the ’917 Claims is more likely than not to be patent ineligible under § 101.

For purposes of this Preliminary Response to Wealth’s § 101 challenge, Vanilla focuses on the subject matter required by the independent limitations that are common to all of the ’917 Claims. However, Vanilla expressly reserves its right to later present further arguments regarding the § 101 eligibility of the dependent claims based on the additional limitations of those claims.

A. Wealth Fails to Establish That the ’917 Claims Are Directed to an Abstract Idea under *Alice* Step One / USPTO Step 2A

As an initial matter, Wealth’s Petition fails to establish that the ’917 Claims are directed to an abstract idea under *Alice* step one / USPTO step 2A.

The *Alice* step one analysis involves “look[ing] to the character of the claims as a whole to determine whether they are ‘directed to’ patent-ineligible subject matter.” *Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327, 1335 (Fed. Cir. 2016). *Alice* step one corresponds to Step 2A of the USPTO’s framework for evaluating patent eligibility, which involves a two-pronged inquiry for determining whether a claim is “directed to” a judicial exception: (1) “whether the claim recites a judicial exception,” and (2) if so, “whether the claim as a whole integrates the exception into a practical application of that exception.” M.P.E.P. § 2106.04.II. With respect to Prong Two, the USPTO’s guidance further explains that:

The Prong Two analysis considers ***the claim as a whole***. That is, the limitations containing the judicial exception as well as the additional elements in the claim besides the judicial exception need to be evaluated together to determine whether the claim integrates the judicial exception into a practical application. Because a judicial exception alone is not eligible subject matter, if there are no additional claim elements besides the judicial exception, or if the additional claim elements merely recite another judicial exception, that is insufficient to integrate the judicial exception into a practical application. However, ***the way in which the additional elements use or interact with the exception may integrate it into a practical application***. Accordingly, the additional limitations should not be evaluated in a vacuum, completely separate from the recited judicial exception. Instead, the analysis should take into consideration ***all the claim limitations and how those limitations interact and impact each other*** when evaluating whether the exception is integrated into a practical application.

Id. at § 2106.04(d)(III) (emphasis added); *see also* U.S. Patent & Trademark Office, Advance Notice of Change to the MPEP in Light of *Ex Parte Desjardins*, 3 (Dec. 5, 2025) (emphasizing that claims should be evaluated “as an ordered combination,

without ignoring the requirements of the individual steps” and that the claims should not be evaluated “at such a high level of generality” that potentially meaningful technical limitations are dismissed without adequate explanation).

One well-established way that a claim as a whole can be found to integrate an abstract idea into a practical application under *Alice* step one / USPTO Step 2A is by reciting a claimed invention that improves the functioning of a computer or improves another technology or technical field. *See, e.g., McRO, Inc. v. Bandai Namco Games America, Inc.*, 837 F.3d 1299, 1313-1316 (Fed. Cir. 2016); *Enfish*, 822 F.3d at 1338-39; M.P.E.P. § 2106.04(d)(1). And while “[t]he courts have not provided an explicit test for this consideration,” the USPTO’s guidance provides a two-part framework for performing this analysis:

[F]irst the specification should be evaluated to determine if the disclosure provides sufficient details such that one of ordinary skill in the art would recognize the claimed invention as providing an improvement in the functioning of a computer, or an improvement to other technology or a technical field. The specification need not explicitly set forth the improvement, but it must describe the invention such that the improvement would be apparent to one of ordinary skill in the art. . . .

Second, if the specification sets forth an improvement in technology or a technical field, the claim must be evaluated to ensure that the claim itself reflects the disclosed improvement, i.e., that the claim includes the components or steps of the invention that provide the improvement described in the specification. The claim itself does not need to explicitly recite the improvement described in the specification (e.g., “thereby increasing the bandwidth of the channel”).

M.P.E.P. § 2106.04(d)(1) (as revised by U.S. Patent & Trademark Office, Advance Notice of Change to the MPEP in Light of *Ex Parte Desjardins* 2-3 (Dec. 5, 2025)).

More recently, the USPTO has confirmed that claims directed to artificial intelligence (AI) technology can qualify as a patent-eligible improvement to technology under *Alice* step one / USPTO Step 2A. Indeed, the USPTO’s 2024 AI eligibility guidance states that “[m]any claims to AI inventions are eligible as improvements to the functioning of a computer or improvements to another technology or technical field,” and provides the following framework for evaluating claims involving AI inventions:

A key point of distinction for AI inventions is between a claim that reflects an improvement to a computer or other technology described in the specification, which is eligible, and a claim in which the additional elements amount to no more than reciting “apply it,” instructing implementation of a judicial exception on a computer, or generally linking the judicial exception to a particular technological environment or field of use, which is ineligible. An important consideration in determining whether a claim improves technology is the extent to which the claim covers a particular solution to a problem or a particular way to achieve a desired outcome, as opposed to merely claiming the idea of a solution or outcome. AI inventions may provide a particular way to achieve a desired outcome when they claim, for example, ***a specific application of AI to a particular technological field*** (*i.e.*, a particular solution to a problem). In those circumstances, the claim is not merely directed to the idea of a solution or outcome and amounts to more than merely “applying” the judicial exception or generally linking the judicial exception to a field of use or technological environment. In other words, the claim reflects an improvement in a computer or other technology.

2024 Guidance Update, 89 Fed. Reg. at 58,137 (internal quotations and citations omitted) (emphasis added).

Along similar lines, a recent federal court decision recognized that there are at least two different routes by which claims involving machine-learning technology can satisfy *Alice* step one:

- (1) by claiming “specific improvements to the machine-learning models” themselves; or
- (2) by claiming “a specifically arranged implementation of machine learning, analogous to the patent-eligible arrangements of *McRO* and *Koninklijke*.”

Aon, 791 F. Supp. 3d at 540 (distinguishing *Recentive Analytics, Inc. v. Fox Corp.*, 134 F.4th 1205, 1213-16 (Fed. Cir. 2025)) (rejecting the patent challenger’s argument that ML-based claims must “offer a new twist on machine learning itself” in order to satisfy *Alice* step one).

On the other hand, claims that “do no more than claim the application of **generic** machine learning to new data environments” are not analogous to the patent-eligible arrangements of *McRO* and *Koninklijke* and thus do not qualify as a patent-eligible improvement to technology under *Alice* step one. *Recentive*, 134 F.4th at 1215-16 (emphasis added).

When evaluated under this legal framework, it is clear that the ’917 Claims as a whole are directed to a patent-eligible improvement to technology because they recite a specifically arranged implementation of machine learning rather than mere


application of generic machine learning to a new data environment. See Ex. 2001 [Thompson Decl.], ¶¶ 103-108, 117.

Indeed, as explained above, the '917 Claims require a structured ML-based architecture comprising a sequence of five stages of functionality that interact with one another to generate the claimed “*resource data*” and then generate and update the claimed “*one or more visualizations*,” which are:

- (1) a first execution stage of the “*one or more machine-learning models*” that generates “*one or more prompts*” based on “*recognition of textual features of the one or more resource documents*” that are input into the first execution stage;
- (2) an intervening stage for receiving “*one or more responses to the one or more prompts*” that are generated during the first execution stage of the “*one or more machine-learning models*”;
- (3) a second execution stage of the “*one or more machine-learning models*” that generates structured “*resource data*” based on a combination of the “*one or more [prompt] responses*” and “*rule data comprising one or more rules associated with the resources*” (where the “*one or more machine-learning models*” are trained based on “*training data comprising training resource data and training rule data*,” the “*resource data fields*” are based on “*resource document fields*,” and the “*resource data field values*” are based on the “*one or more [prompt] responses*”);
- (4) a generation stage for “*one or more visualizations of the distribution of the one or more assets to the one or more resource recipients*” based on the structured “*resource data*”; and
- (5) a real-time updating stage for the “*one or more visualizations*” based on “*changes in one or more values of the one or more assets*.”

See Ex. 1001 ['917 Patent] at 39:3-42; Ex. 2001 [Thompson Decl.], ¶ 104.

The annotated version of the illustrative diagram from above that shows this five-stage ML-based architecture is reproduced here:



As shown, the claimed five-stage ML-based architecture is built on a specifically arranged implementation of machine learning involving a sequence of two distinct execution stages of the claimed “*one or more machine-learning models*” that perform separate and clearly defined analytical tasks – including receiving separate and distinct inputs and generating separate and distinct outputs – along with an intervening receipt of “*one or more [prompt] responses*” that interconnects the sequence of two distinct execution stages of the claimed “*one or more machine-learning models*.” Ex. 2001 [Thompson Decl.], ¶¶ 105-106. More specifically:

- the first execution stage of “*one or more machine-learning models*” performs a first analytical task of generating “*one or more prompts*” based

on “*recognition of textual features of the one or more resource documents*”;

- the intervening stage receives “*one or more responses*” to the “*one or more prompts*” generated during the first execution stage;
- the second execution stage of the “*one or more machine-learning models*” performs a second analytical task of generating structured “*resource data*” based on a combination of the received “*one or more responses*” to the “*one or more prompts*” generated during the first execution stage along with “*rule data comprising one or more rules associated with the resources*,” where the “*one or more machine-learning models*” are trained based on “*training data comprising training resource data and training rule data*,” the “*resource data field(s)*” are based on “*resource document field(s)*,” the “*resource data field value(s)*” are based on the “*one or more [prompt] responses*,” and the generated “*resource data*” is thereafter utilized to generate the claimed “*one or more visualizations*.”

Id.

This is a structured arrangement of machine learning in which two separate, distinct, and specific execution stages of the claimed “*one or more machine-learning models*” are combined together into a defined sequence and the “*one or more prompts*” output by the first execution stage drives the second execution stage’s defined inputs by prompting a user to supplement the data extracted from the “*one or more resource documents*” during the first execution stage. *Id.*, ¶ 107.

The ’917 Claims share many similarities with the claim that was found to be a patent-eligible implementation of machine learning in *Aon v. Zesty.ai*. That claim recited “two separate and independently trained machine-learning classifiers that operate independently” – where “[o]ne identifies a physical characteristic of a property (such as a roof shape), and the other evaluates the condition of that same

characteristic (such as signs of deterioration)” – in order to provide “a risk estimate of damage to the property due to the one or more disasters.” 791 F. Supp. 3d at 538. The court found this architecture to be a “specifically arranged implementation of machine learning” that is “analogous to the patent-eligible arrangements of *McRO* and *Koninklijke*,” rather than an “application of generic machine learning to new data environments” as in *Recentive*, based on the following reasoning:

[C]laim 1 of the 029 patent recites a concrete and specific technological solution. Rather than referring to the use of machine learning in general, the claim requires individually trained classifiers that operate independently, ***performing separate and clearly defined analytical tasks***: one classifier identifies a specific physical characteristic of a property from aerial imagery, and the other independently assesses the condition of that characteristic. ***This division of labor***, together with the requirement to analyze defined pixel groupings and produce risk assessment outputs based on the classifiers’ assessments, confirms that the claim represents ***more than a generalized use of machine learning***. Although the patent does not purport to invent machine learning algorithms, computer hardware, or the general concept of aerial imagery analysis or risk assessment, it does ***specifically define how these known components must be structured*** to achieve a supposed technological improvement in accuracy and efficiency to the combination of those components.

* * * *

Here, the ***independent training and distinct analytical roles*** assigned to each of the two machine learning classifiers provide a ***structured technological method*** specific to analyzing pixel groups in aerial imagery for property risk assessments. Like the claims upheld in *McRO* and *Koninklijke*, claim 1 clearly defines how known technological components must be ***specifically structured and combined***, rather than broadly claiming a purported abstract idea of applying machine learning to aerial images to assess risk. By ***assigning separate and clearly defined analytical tasks*** to the two independently trained

classifiers, claim 1 provides a particular structured solution that the patent says is accurate and efficient. Thus, claim 1 is not directed to the abstract idea advanced by Zesty.

Id. at 539-540 (emphasis added).

The '917 Claims likewise recite a “specifically arranged implementation of machine learning” that is “analogous to the patent-eligible arrangements of *McRO* and *Koninklijke*.” *Id.* Indeed, similar to the *Aon* claim:

- The '917 Claims involve multiple ML-model execution stages that are specifically configured to perform ***separate and clearly defined analytical tasks***: (1) a first ML-model execution stage that generates “*one or more prompts*” based on “*recognition of textual features of the one or more resource documents*” and (2) a second ML-model execution stage that generates structured “*resource data*” based on a combination of “*one or more responses to the one or more prompts*” generated during the first execution stage and “*rule data comprising one or more rules associated with the resources*.”
- There is ***division of labor*** between the two claimed ML-model execution stages along with an intervening requirement to receive the “*one or more [prompt] responses*” that are fed as input to the second ML-model execution stage, which confirms that the claims represent more than a generalized use of machine learning.
- The '917 Claims define how known technological components must be specifically structured and combined – namely, through the sequential stages described above – rather than broadly claiming a purported abstract idea of applying machine learning to resource allocation.

Thus, as with the *Aon* claim, the '917 Claims are a patent-eligible improvement to technology under *Alice* step one / USPTO Step 2A.

The conclusion that the '917 Claims are a patent-eligible improvement to technology is confirmed by the USPTO’s two-part framework for evaluating that

question. Indeed, as an initial matter, the '917 Patent's disclosure provides sufficient details such that a POSITA would recognize the claimed invention as providing an improvement in existing software technology for resource allocation. Ex. 2001 [Thompson Decl.], ¶¶ 110-114.

For instance, as explained above, the '917 Patent begins by introducing the problems with existing software technology for resource allocation, including that (1) “the use of multiple software applications to manage resources can be burdensome and demand a significant amount of manual input” and (2) “generalized software applications may not be fully suited to meeting the often specialized requirements of resource allocation.” Ex. 1001 ['917 Patent] at 1:26-34.

In turn, the '917 Patent states that “the present disclosure is directed to *improved* computer systems, computer applications, computer-implemented methods, user interfaces, and/or services for resource processing and/or allocation,” and explains that “the disclosed technology can provide a variety of technical effects and benefits” that improve over the existing software technology for resource allocation, including:

- “[A]ssist[ing] a user in *more effectively performing the technical task of resource data processing* by means of a *continued and/or guided human-machine interaction process* in which responses to generated prompts based on resource documents are received and the disclosed technology generates real-time visualizations based on continuously updated resource data,” which enables “generation of personalized resource data provision that can more effectively address a user’s questions.”

- “A[n] ***improvement in the generation of resource data*** that can be used in the analysis and implementation of resource allocation instructions,” where “machine-learning models can be used in the process of generating personalized resource data” such that “more relevant and timely resource data can be provided to a user based on specific information associated with that particular user’s resource data.”
- “[I]mproving the effectiveness with which resources planning tasks are performed,” such as by enabling a computing system to be “continuously updated based on user responses to prompts, changes in rules that govern resources, and/or the occurrence of key events that may change the way in which the assets associated with resources are distributed.”
- “[I]mprov[ing] the operation of a resource processing device by more effectively performing a variety of tasks with the specific benefits of providing more accurate resources planning information.”
- “[U]s[ing] various rules-based techniques and/or machine-learning models to ***more efficiently process resource documents*** that would otherwise require time consuming and burdensome manual review and handling.”

Id. at 5:56-65, 14:41-15:13 (emphasis added).

Based on this disclosure, a POSITA would recognize these statements to be describing an improvement in the existing software technology for resource allocation that was previously introduced in the background section of the '917 Patent, which was described as being “burdensome” and not “fully suited to meeting the often specialized requirements of resource allocation.” Ex. 2001 [Thompson Decl.], ¶¶ 113-114.

Moreover, the '917 Claims reflect the disclosed improvements to the existing software technology for resource allocation through their recitation of the structured,

five-stage ML-based architecture for generating “*resource data*” and then generating and updating “*one or more visualizations.*” *Id.*, ¶ 115.

Thus, because (i) the ’917 Patent’s disclosure provides sufficient details such that a POSITA would recognize the claimed invention as providing an improvement in existing software technology for resource allocation and (ii) the ’917 Claims reflect that described improvement, the ’917 Claims are directed to a patent-eligible improvement to technology under *Alice* step one / USPTO Step 2A. *See* M.P.E.P. § 2106.04(d)(III).

Wealth’s position to the contrary is based on a number of legal errors in Wealth’s characterization and application of the legal framework for evaluating whether claims are directed to an abstract idea under *Alice* step one / USPTO Step 2A, each of which serves as an independent basis for rejecting Wealth’s § 101 challenge under *Alice* step one / USPTO Step 2A.

First, Wealth’s Petition is premised on an assumption that the only way for patents involving machine learning to be eligible under *Alice* step one / USPTO Step 2A is by claiming an improvement to machine learning itself. Pet. 2, 17. That is incorrect. As explained above, both the courts and the USPTO have recognized that patents claiming a specific implementation of machine learning are also eligible under *Alice* step one / USPTO Step 2A. *Aon*, 791 F. Supp. 3d at 540 (distinguishing *Recentive* and holding that “another route to eligibility . . . is claiming a specifically

arranged implementation of machine learning”); *2024 Guidance Update*, 89 Fed. Reg. at 58,137 (noting that “a specific application of AI to a particular technological field” is a patent-eligible improvement to technology).

That is exactly the situation presented here. In particular, as established above, the ’917 Claims recite a specifically arranged implementation of machine learning that is analogous to the patent-eligible arrangements of *McRO* and *Koninklijke*, rather than an “application of generic machine learning to new data environments” as in *Recentive*. See *Aon*, 791 F. Supp. 3d at 539-40.

Second, Wealth’s Step 2A, Prong Two analysis merely addresses certain individual limitations of the ’917 Claims in a vacuum and never considers the claims as a whole. Pet. 17-20. In fact, Wealth’s entire analysis of the ’917 Claims under Step 2A, Prong Two boils down to three allegations:

- (1) The “generic computing components” of the ’917 Claims do not integrate the alleged abstract idea into a practical application.
- (2) Training the “*one or more machine-learning models*” based on “*training data comprising training resource data and training rule data*” does not integrate the alleged abstract idea into a practical application.
- (3) The claims “do[] not explain how to perform its method limitations with its generic components.”

Id. This is legally improper – the analysis under *Alice* step one and USPTO Step 2A, Prong Two must consider the claims **as a whole**, including **all** the claim

limitations and how those limitations interact with and impact each other. *See Enfish*, 822 F.3d at 1335; M.P.E.P. § 2106.04(d)(III).

As a result of this failure to evaluate the '917 Claims as a whole, Wealth's Petition overlooks that the limitations of the '917 Claims interact with one another in a specific way in order to integrate the alleged abstract idea into a practical application. In particular, as explained above, the '917 Claims require a structured, ML-based architecture involving five different stages of functionality that interact with one another to generate the claimed "*resource data*" and then generate and update the claimed "*one or more visualizations*." When properly evaluating this structured, five-stage sequence ML-based architecture ***as a whole*** – including the interactions between the different stages of functionality – it is clear that the '917 Claims integrate any alleged abstract idea into a practical application under *Alice* step one / USPTO Step 2A.

Third, Wealth's Petition fails to give proper weight to the '917 Patent's disclosure. Pet. 18-19. Wealth concedes that the background of the '917 Patent frames the problem in terms of the limitations of "software applications for managing resource information" – which is a problem with ***technology*** – and further concedes that the '917 Patent expressly describes the disclosed technology as an "improvement in the generation of resource data" that provides "technical effect[s]," including "improving the effectiveness with which resource planning tasks are

performed” and “improv[ing] the operation of a resource planning device by more effectively performing a variety of tasks with the specific benefits of providing more accurate resources planning information.” *Id.* Nevertheless, Wealth attempts to dismiss these improvements by characterizing them as improvements to mental processes, rather than improvements to technology. This is not correct.

In the context of the ’917 Patent’s disclosure – and in particular the background section describing the problems with existing software technology for resource allocation – a POSITA would recognize that the ’917 Patent is describing improvements to that existing software technology for resource allocation. Ex. 2001 [Thompson Decl.], ¶¶ 113-114. For instance, when the ’917 Patent states that the disclosed technology can “provide the technical effect of improving the effectiveness with which resources planning tasks are performed” and “improve the operation of a resource processing device by more effectively performing a variety of tasks with the specific benefits of providing more accurate resources planning information,” a POSITA would recognize this to be describing improvements in the previously-introduced software technology for resource allocation that was “burdensome” and not “fully suited to meeting the often specialized requirements of resource allocation.” *Id.* The other statements from the ’917 Patent identified above support this conclusion as well.

Fourth, while Wealth’s Petition cites to various cases that purportedly support its *Alice* step one / USPTO Step 2A analysis, Wealth fails to provide any meaningful analysis of how or why the ’917 Claims are analogous to the claims found to be ineligible in those cases.

For instance, Wealth’s *Alice* step one / USPTO Step 2A analysis relies heavily on the *Recentive* decision, but Wealth never once actually compares the ’917 Claims to the *Recentive* claims that were found to be an ineligible “application of generic machine learning.” Pet. 17-20. Given this failure to establish the applicability of *Recentive* to the ’917 Claims, Wealth’s arguments under *Recentive* are insufficient to meet its burden. And to be sure, the ’917 Claims are materially different from the *Recentive* claims, which did not recite a specifically arranged implementation of machine learning. For instance:

- In the *Recentive* “Network Map” patents, the only claimed reference to machine learning was a “*wherein*” clause reciting that “*generating the network map comprises using a machine learning technique to optimize an overall television rating across the first plurality of live events and the second plurality of live events*”; and
- In the *Recentive* “Machine Learning Training” patents, the claims recite “*iteratively training*” and using a single “*machine learning (ML) model*” that is configured to perform a single task – namely, generating “*a schedule for [a] future series of live events.*”

See Ex. 2003 [U.S. Pat. No. 10,911,811] at 9:66-10:32; Ex. 2004 [U.S. Pat. No. 11,386,367] at 14:2-49; *see also Recentive*, 134 F.4th at 1208-15.

Thus, unlike either the patent-eligible *Aon* claim or the '917 Claims, the *Recentive* claims fail to recite a structured ML-based architecture involving multiple ML-model execution stages that are specifically configured to perform separate and clearly defined analytical tasks within an ordered workflow. *See Aon*, 791 F. Supp. 3d at 538 (explaining that the *Recentive* claims “failed to identify a specific technical implementation” because they only “broadly referenced iterative training and real-time updating without specifying the steps through which the machine learning technology achieves an improvement”) (internal quotation omitted).

Wealth also cites to the *In re Killian*, *FairWarning IP*, and *Electric Power Group* cases in connection with its contention that the '917 Claims “merely gather data, process it, and display the results,” but Wealth’s Petition fails to provide any analysis whatsoever in support of its contention that the multi-stage ML-based architecture required by the '917 Claims is analogous to the highly-generic claims found to be patent-ineligible in the *In re Killian*, *FairWarning IP*, and *Electric Power Group* cases. To be sure, the '917 Claims recite details that go well beyond highly-generic functionality for gathering data, processing it, and displaying the results. Ex. 2001 [Thompson Decl.], ¶ 118.

Lastly, Wealth cites to *Int’l Bus. Machs. Corp. v. Zillow Grp., Inc.* in connection with its contention that the '917 Claims are “written in result-based functional language that does not sufficiently describe how to achieve these results

in a non-abstract way,” but once again, Wealth’s Petition provides no explanation of how the ’917 Claims compare to the claims found to be ineligible in *Zillow*. Pet. 20 (internal quotations omitted). In any event, Wealth’s unsupported allegation that the “method limitations” of the ’917 Claims are “result-based functional language” is contradicted by the plain language of the ’917 Claims, which does not merely recite the results of generating the “*one or more prompts*,” generating the “*resource data*,” generating “*one or more visualizations*,” or updating the “*one or more visualizations*.” Ex. 2001 [Thompson Decl.], ¶ 119.

To the contrary, the ’917 Claims expressly recite how those results are achieved – including by (1) inputting “*one or more resource documents*” into “*one or more machine-learning models*” that are configured to “*generate the one or more prompts based on recognition of textual features of the one or more resource documents*,” (2) inputting “*one or more [prompt] responses*” and “*rule data*” into “*one or more machine-learning models*” that are “*trained based on training data comprising training resource data and training rule data*” (where the “*resource data field(s)*” are based on the “*resource document field(s)*” and the “*resource data field value(s)*” are based on the “*one or more [prompt] responses*”), (3) generating the “*one or more visualizations*” based on the structured “*resource data*,” and (4) updating the “*one or more visualizations*” “*in real-time based on changes in one or more values of the one or more assets*.” *Id.*

Accordingly, for at least the foregoing reasons, Wealth fails to meet its burden of establishing that the '917 Claims are directed to an abstract idea under *Alice* step one / USPTO step 2A.

B. Wealth Fails to Establish That the '917 Claims Lack an Inventive Concept Under *Alice* Step 2 / USPTO Step 2B

Because the '917 Claims are not directed to an abstract idea under *Alice* step one / USPTO step 2A, the Board need not even reach *Alice* step two / USPTO step 2B of the § 101 framework. However, to the extent that the Board finds that the '917 Claims are directed to an abstract idea under *Alice* step one / USPTO step 2A, Wealth's Petition also fails to establish that the '917 Claims lack an inventive concept under *Alice* step two / USPTO step 2B.

The *Alice* step two analysis involves a consideration of “the elements of each claim – both individually and *as an ordered combination* – to determine whether the additional elements transform the nature of the claim into a patent-eligible application of that abstract idea.” *DDR Holdings, LLC v. Hotels.com, L.P.*, 773 F.3d 1245, 1255 (Fed. Cir. 2014) (emphasis added). This *Alice* step two analysis is often referred to as “a search for an ‘inventive concept,’ or some element or combination of elements sufficient to ensure that the claim in practice amounts to ‘significantly more’ than a patent on an ineligible concept.” *Id.*

Alice step two corresponds to Step 2B of the USPTO's framework for evaluating patent eligibility, and the USPTO's guidance on patent eligibility

emphasizes that “[c]onsideration of the elements in combination is particularly important, because even if an additional element does not amount to significantly more on its own, it can still amount to significantly more when considered in combination with the other elements of the claim.” M.P.E.P. § 2106.05.

One factor that both courts and the USPTO evaluate when looking for an inventive concept is whether the claim limitations – when considered both individually and as an ordered combination – recite something other than “well-understood, routine, and conventional activities.” *Bascom Global Internet Services, Inc. v. AT&T Mobility LLC*, 827 F.3d 1341, 1350 (Fed. Cir. 2016). In this respect, even if the individual limitations are found to be “well-understood, routine, and conventional activities,” an “inventive concept” can still be found in the “non-conventional and non-generic arrangement of known, conventional pieces.” *Id.* This is particularly the case when the ordered combination provides a “technical improvement” over prior art. *Id.*; *see also* M.P.E.P. § 2106.05 (noting that limitations found to qualify as “significantly more” include “[i]mprovements to the functioning of a computer,” “[i]mprovements to any other technology or technical field,” “[a]pplying a specific limitation other than what is well-understood, routine, conventional activity in the field, or adding unconventional steps that confine the claim to a particular useful application”).

Wealth's Petition relies on this factor to argue that the '917 Claims "recite[] an abstract idea performed by well-understood, routine, and conventional computing components and machine learning," but Wealth fails to meet its burden of establishing that the limitations of the '917 Claims are well-understood, routine, and conventional when considered both individually and *as an ordered combination*. In fact, the *only* aspect of the '917 Claims that Wealth's Petition specifically challenges as being well-understood, routine, and conventional activities is the claimed "one or more machine-learning models' trained based on training resource data and training rule data." Pet. 21-22. However, this is not the only non-abstract limitation of the '917 Claims. To the contrary, the '917 Claims include several other non-abstract limitations that, when taken together as an ordered combination, supply an inventive concept under *Alice* step two / USPTO step 2B. *See* Ex. 2001 [Thompson Decl.], ¶¶ 121-123. Wealth's Petition and Dr. White's Declaration wholly fail to address this ordered combination.

In particular, as discussed above, the '917 Claims require a structured ML-based architecture comprising a sequence of five stages of functionality that interact with one another to generate the claimed "*resource data*" and then generate and update the claimed "*one or more visualizations*," which are:

- (1) a first execution stage of the "*one or more machine-learning models*" that generates "*one or more prompts*" based on "*recognition of textual features of the one or more resource documents*" that are input into the first execution stage;

- (2) an intervening stage for receiving “*one or more responses to the one or more prompts*” that are generated during the first execution stage of the “*one or more machine-learning models*”;
- (3) a second execution stage of the “*one or more machine-learning models*” that generates structured “*resource data*” based on a combination of the “*one or more [prompt] responses*” and “*rule data comprising one or more rules associated with the resources*” (where the “*one or more machine-learning models*” are trained based on “*training data comprising training resource data and training rule data,*” the “*resource data fields*” are based on “*resource document fields,*” and the “*resource data field values*” are based on the “*one or more [prompt] responses*”);
- (4) a generation stage for “*one or more visualizations of the distribution of the one or more assets to the one or more resource recipients*” based on the structured “*resource data*”; and
- (5) a real-time updating stage for the “*one or more visualizations*” based on “*changes in one or more values of the one or more assets.*”

Ex. 1001 [’917 Patent] at 39:3-42; Ex. 2001 [Thompson Decl.], ¶ 122.

Wealth’s failure to address these limitations as an ordered combination is fatal to its *Alice* step two / USPTO Step 2B analysis, and when properly evaluated, it is clear that the ’917 Claims recite a non-conventional and non-generic arrangement that supplies an inventive concept under *Alice* step two / USPTO Step 2B. Ex. 2001 [Thompson Decl.], ¶ 123.

For instance, the claimed ML-based architecture combines two separate and distinct execution stages of the claimed “*one or more machine-learning models*” that are configured to perform separate and clearly defined analytical tasks and are interconnected with one another by virtue of the fact that the “*one or more prompts*”

output by the first execution stage serves as the basis for the “*one or more responses*” that are provided as input to the second execution stage. This ordered combination of limitations requires a specifically arranged implementation of machine learning that is not well-understood, routine, and conventional activity. *Id.*

Thus, even to the extent that the Board finds that the ’917 Claims are directed to an abstract idea under *Alice* step one / USPTO step 2A and proceeds to *Alice* step two / USPTO step 2B, Wealth’s Petition fails to establish that the ’917 Claims lack an inventive concept that amounts to significantly more than the alleged abstract idea under *Alice* step two / USPTO step 2B.

C. Summary

For the reasons explained above, Wealth’s § 101 challenge to the ’917 Claims fails to establish either that (1) any of the ’917 Claims is directed to an abstract idea under *Alice* step one / USPTO step 2A or that (2) any of the ’917 Claims lacks an inventive concept under *Alice* step two / USPTO step 2B. Accordingly, Wealth fails to meet its burden of establishing that any of the ’917 Claims is more likely than not to be patent ineligible under § 101.

VI. WEALTH’S § 103 CHALLENGES FAIL

Wealth’s § 103 challenges to the ’917 Claims all rely exclusively on the Bateman reference for the limitations related to the claimed “*one or more machine-learning models*,” but the Bateman reference is clearly deficient with respect to at

least three different requirements of the '917 Claims. **First**, Bateman fails to disclose or suggest the claimed sequence of two separate and distinct execution stages of the claimed “*one or more machine-learning models*” that are required by Limitations 1[a]-1[c]. **Second**, Bateman fails to disclose or suggest the additional details of the second execution stage required by Limitation 1[c], including in particular the claimed requirement of generating the claimed “*resource data*” based on “*inputting the one or more responses and rule data comprising one or more rules associated with the resources into the one or more machine-learning models.*” **Third**, Bateman fails to disclose or suggest the additional requirement of Limitation 1[c] that “*the one or more machine-learning models are trained based on training data comprising training resource data and training rule data.*”

Each of these deficiencies of the Bateman reference serves as a separate and independent basis for rejecting Wealth’s § 103 challenges to all of the '917 Claims. Thus, Wealth’s Petition fails to establish that any of the '917 Claims is more likely than not to be unpatentable as obvious under § 103.

For purposes of this Preliminary Response to Wealth’s § 103 challenges, Vanilla focuses on Bateman’s failure to disclose or suggest these three requirements of the independent claims, which is applicable to all of the '917 Claims and all of Wealth’s § 103 grounds. However, Vanilla expressly reserves its right to later present further arguments in response to Wealth’s § 103 challenges, including but

not limited to further arguments in response to Wealth's contentions regarding analogous art, rationale for combination, reasonable expectation of success, or the mapping of any other independent or dependent limitation of the '917 Claims.


A. Wealth Fails to Establish that Claims 1, 2, 6-7, 12-14, 17, and 19 are Rendered Obvious by the Combination of Racanelli, Bateman, and Dintenfass (Ground 2)

1. The combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest the claimed sequential execution stages of the “*one or more machine-learning models*”

The '917 Claims require a structured ML-based architecture comprising a sequence of five stages of functionality that interact with one another to generate the claimed “*resource data*” and then generate and update the claimed “*one or more visualizations*,” which involves a sequence of two separate and distinct execution stages of the claimed “*one or more machine-learning models*”:

- A first execution stage during which the claimed “*one or more machine-learning models*” generate the claimed “*one or more prompts*” based on the claimed “*one or more resource documents*”; and
- A second execution stage – which occurs after “*one or more responses to the one or more prompts*” generated during the first execution stage are received – during which the claimed “*one or more machine-learning models*” generate the claimed “*resource data*” based on the “*one or more [prompt] responses*” along with the claimed “*rule data*.”

Ex. 1001 ['917 Patent] at 39:3-33. These sequential execution stages are shown in the annotated illustrative diagram from above, which is reproduced here:



Ex. 2001 [Thompson Decl.], ¶ 167.


Wealth’s Petition relies on Bateman for both of the claimed execution stages of the “*one or more machine-learning models*” by mapping Bateman’s “language processing model” to the claimed “*one or more machine-learning models*” and contending that the language processing model is used to both (1) generate candidate values for a parameter of a conveyance record that are presented to a human reviewer for selection, which Wealth maps to the first execution stage, and (2) “appl[y]” the reviewer’s selection of a candidate value to the conveyance record, which Wealth maps to the second execution stage. Pet. 34-35, 37.

However, contrary to Wealth’s contentions, Bateman’s architecture only involves a ***single execution stage*** of the language processing model. In particular, Bateman discloses that its language processing models “take as input the text of a

given document and extract parameters of conveyance records.” Ex. 1007 [Bateman] at [0063]. Separately, Bateman discloses that if the language processing model extracts multiple candidate values for a conveyance-record parameter and one of those candidate values is then selected by a human reviewer, “[t]he received selection, upon receipt, may be applied to the corresponding conveyance record.” *Id.* at [0067].

However, Bateman never discloses or suggests that the reviewer’s selection is ever input back into the language processing model or that the language processing model is otherwise used to “appl[y]” the reviewer’s selection to the conveyance record. Ex. 2001 [Thompson Decl.], ¶¶ 171-174 (citing Ex. 1007 [Bateman] at [0067]-[0071]). To the contrary, a POSITA would understand Bateman’s disclosure of the reviewer’s selection being “applied to the corresponding conveyance record” to simply mean that the selected candidate value is inserted into the conveyance record. *Id.*

To illustrate, the functionality of Bateman’s architecture is shown in the following diagram:



Id., ¶ 174. As shown, there is only a *single execution stage* of the language processing model in Bateman’s architecture, and the output of the language processing model along with the reviewer’s selection is then applied to the conveyance record without any further use of the language processing model.

Wealth’s position to the contrary is based solely on this contention from its Petition:

Bateman explains that the human reviewer’s selection is received and applied to the structured conveyance record being extracted, confirming that the selection (response) is input into the language processing model.

Pet. 37 (citing Bateman at [0067]; White Decl. ¶¶ 155-156). However, as just explained, Bateman at [0067] merely states that “[t]he received selection, upon receipt, may be *applied to the corresponding conveyance record.*” Ex. 1007 [Bateman] at [0067] (emphasis added). There is nothing in this sentence (or any other passage of Bateman) suggesting that the reviewer’s selection is “input into the

language processing model.” Ex. 2001 [Thompson Decl.], ¶ 176. Rather, this is an unsupported and conclusory allegation advanced by Wealth and Dr. White that is not grounded in Bateman’s disclosure.

In fact, not only is this alleged functionality of inputting the reviewer’s selection into the language processing model plainly missing from Bateman’s disclosure, but this functionality also would not make technical sense in the context of Bateman’s architecture. Ex. 2001 [Thompson Decl.], ¶ 177. There are at least two reasons for this. First, a reviewer’s selected candidate value can be “applied” to the conveyance record by simply inserting it directly into the conveyance record without any further need to involve the language processing model. *Id.* Second, Bateman repeatedly and consistently discloses that the purpose of its language processing model is to extract ***structured*** information from ***unstructured*** textual documents, and there is no need to perform that function with respect to the reviewer’s selected candidate value because it is already structured information. *Id.*, ¶ 178 (citing to Ex. 1007 [Bateman] at [0063]-[0066], [0070]).

Racanelli and Dintenfass fail to cure this deficiency of Bateman. Ex. 2001 [Thompson Decl.], ¶ 180. Indeed, Wealth’s Petition does not rely on Racanelli or Dintenfass for ***any*** aspect of the ML-based functionality of the ’917 Claims – let alone the claimed requirement of having two sequential execution stages of the claimed “*one or more machine-learning models.*” Pet. 34-35 (relying on Bateman

for the “*one or more machine-learning models*” of Limitation 1[a]), 37-41 (relying on Bateman for the “*one or more machine-learning models*” of Limitation 1[c]).

Thus, the combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest the two sequential execution stages of the claimed “*one or more machine-learning models*” that are required by the ’917 Claims.

2. The combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest generating “*resource data*” based on “*inputting the one or more responses and rule data comprising one or more rules associated with the resources into the one or more machine-learning models*”

The specific details of the second execution stage of the claimed “*one or more machine-learning models*” are recited in Limitation 1[c], which requires:

generating, by the computing system, based on inputting the one or more [prompt] responses and rule data comprising one or more rules associated with the resources into the one or more machine-learning models, resource data comprising one or more resource data fields and one or more resource data field values that correspond to the one or more resource data fields, . . . wherein the one or more resource data fields are based on the one or more resource document fields, and wherein the one or more resource data field values are based on the one or more responses[.]

Ex. 1001 [’917 Patent] at 39:20-33 (Limitation 1[c]).

As noted above, Wealth’s Petition relies solely on Bateman for this second execution stage of the claimed “*one or more machine-learning models.*” Pet. 37-41. However, Bateman fails to disclose or suggest that there is **any** second execution stage of its language processing model – let alone a second execution stage that

involves inputting “*one or more [prompt] responses*” and “*rule data comprising one or more rules associated with the resources*” into the language processing model.

Wealth’s Petition maps the claimed “*one or more machine-learning models*” to Bateman’s language processing model and maps the claimed “*one or more [prompt] responses*” to the reviewer’s selection of a candidate value for a parameter of a conveyance record. Pet. 37. However, as established above, Bateman fails to disclose or suggest that the reviewer’s selected candidate is ever input back into Bateman’s language processing model. Ex. 2001 [Thompson Decl.], ¶ 186. Instead, Bateman discloses that the selected candidate value is “applied to the corresponding conveyance record,” which simply means that the selected candidate value is inserted into the conveyance record without any further use of the language processing model. *Id.*

Thus, even under Wealth’s own mapping of the “*one or more machine-learning models*” to Bateman’s language processing model and the claimed “*one or more [prompt] responses*” to the reviewer’s selection of a candidate value, Bateman fails to disclose or suggest the claimed functionality of generating “*resource data*” based on “inputting *the one or more [prompt] responses . . .* into the one or more machine-learning models.”

Wealth’s position for Limitation 1[c] also fails for the additional reason that Bateman does not disclose or suggest inputting ***both*** the reviewer’s selection (which

it maps to the “*one or more responses*”) and ***also*** the claimed “*rule data*” into its language processing model. In its Petition, Wealth contends that the claimed “*rule data*” can be mapped to “three types of rule data” disclosed in Bateman, which are (1) “provisions of a will,” (2) “a jurisdiction-specific ontology,” and (3) “hand-coded rules.” Pet. 37-40. However, none of these types of alleged “rule data” in Bateman amounts to “*rule data comprising one or more rules associated with the resources*” that is input together with “*one or more [prompt] responses*” into the “*one or more machine-learning models*,” as required by Limitation 1[c] of the ’917 Claims. See Ex. 2001 [Thompson Decl.], ¶¶ 189-218.

Beginning with the “provisions of a will,” Wealth’s Petition contends that “Bateman discloses inputting a will whose provisions reflect conveyances of interests in real property into its language processing model” and that these will provisions “constitute one or more rules associated with the resources.” Pet. 37 (citing Bateman at [0052], [0063]; White Decl. ¶ 158); *see also id.* at 39. However, this contention fails for at least two reasons.

First, Wealth’s Petition already relies on Bateman’s disclosure of inputting a will into the language processing model in connection with the first execution stage of the claimed “*one or more machine-learning models*” recited in Limitation 1[a], which requires “inputting ***one or more resource documents*** into one or more machine-learning models” as a basis for generating the “*one or more prompts*.” Pet.

32-35. This is a separate and distinct function of the '917 Claims that necessarily must occur *before* the Limitation 1[c] functionality of “*inputting the one or more [prompt] responses and rule data comprising one or more rules associated with the resources into the one or more machine-learning models,*” as the “*one or more prompts*” have to be generated before the “*one or more [prompt] responses*” and the “*rule data*” are input into the “*one or more machine-learning models.*” Ex. 2001 [Thompson Decl.], ¶¶ 194-196. Thus, given that Wealth is relying on Bateman’s inputting of a will into the language processing model to meet the first execution stage of the claimed “*one or more machine-learning models*” required by Limitation 1[a], it cannot also rely on that exact same operation to meet the separate, distinct, and later-occurring second execution stage of the claimed “*one or more machine-learning models*” required by Limitation 1[c]. *Id.*

Second, while Limitation 1[c] requires the “*rule data*” to be input into the “*one or more machine learning models*” **together with** the “*one or more [prompt] responses,*” Bateman fails to disclose or suggest that a will is ever input into the language processing model together with the selected candidate value that Wealth maps to the “*one or more [prompt] responses.*” Ex. 2001 [Thompson Decl.], ¶¶ 197-198. Indeed, in Bateman, the candidate values that Wealth maps to the “*one or more prompts*” do not even come into existence until **after** an unstructured document such as a will is input into the language processing model and the candidate values are

extracted from the unstructured document. *Id.* Thus, at the time that a will is input into Bateman’s language processing model, no alleged “*prompt*” has even been generated, which makes it logically impossible for the will to be input into the language processing model *together with* a “*response*” to the “*prompt*” as required by Limitation 1[c]. *Id.*

Turning next to the “jurisdiction-specific ontology,” Wealth’s Petition contends that Bateman’s disclosure of “select[ing] a corresponding ontology according to the jurisdiction in which land is disposed in the course of performing analyses described herein” meets the “*rule data*” aspect of Limitation 1[c]. Pet. 37, 39 (citing Bateman at [0048], [0070], [0072] [0074], [0108]; White Decl. ¶¶ 159, 165-166). However, even if Bateman’s “ontology” were mapped to the claimed “*rule data*,” Bateman fails to disclose or suggest that the “ontology” is ever input into Bateman’s language processing model – let alone input into the language processing model together with a reviewer’s selected candidate value (which is what Wealth maps to the “*one or more [prompt] responses*”). Ex. 2001 [Thompson Decl.], ¶¶ 200-208.

Indeed, the entirety of Bateman’s disclosure regarding the “jurisdiction-specific ontology” is as follows:

To track such rights, some embodiments may maintain in memory an ***ontology*** of real property related rights, including records of which rights are by default subsets of other rights unless explicitly reserved or conveyed. In some cases, different jurisdictions may be associated with

different *ontologies* according to different default rules, and some embodiments may select a corresponding *ontology* according to the jurisdiction in which land is disposed in the course of performing analyses described herein.

Ex. 1007 [Bateman] at [0048] (emphasis added). However, this passage indisputably does not disclose that a selected “ontology” is input into Bateman’s language processing model. *Id.*; Ex. 2001 [Thompson Decl.], ¶ 203.

Wealth attempts to overcome this deficiency by arguing that Bateman’s disclosure of “select[ing] a corresponding ontology according to the jurisdiction in which land is disposed in the course of *performing analyses described herein*” necessarily means that the ontology must be input into the language processing model. Pet. 37. This argument has no merit.

The “analyses described herein” in Bateman involve many different functions that do not involve the language processing model, and a POSITA would not understand this broad statement regarding the “analyses” of Bateman to mean that the “ontology” is input into the language processing model. Ex. 2001 [Thompson Decl.], ¶¶ 204-206. This is particularly the case given that Bateman’s “ontology” is described as a structured dataset and the purpose of the language processing model is to extract *structured* information from *unstructured* textual documents, which contradicts Wealth’s contention that the structured “ontology” is input into the language processing model. *Id.*

Thus, Bateman does not even disclose or suggest that its “ontology” is input into its language processing model, much less discloses that it is input into its language processing model **together with** the reviewer’s selected candidate value that Wealth maps to the claimed “*one or more [prompt] responses.*” *Id.*, ¶ 207.

Turning lastly to the “hand-coded rules,” Wealth’s Petition contends that Bateman’s disclosure of training the language processing model based on such “hand-coded rules” meets the “*rule data*” aspect of Limitation 1[c]. Pet. 38-40 (citing Bateman at [0069], [0072]; White Decl. ¶¶ 160, 167). However, this contention likewise fails for at least two different reasons.

First, Bateman fails to disclose or suggest that those “hand-coded rules” are ever **input** into Bateman’s language processing model – let alone input into the language processing model together with the reviewer-selected candidate value that Wealth maps to the “*one or more [prompt] responses.*” Ex. 2001 [Thompson Decl.], ¶¶ 211-213. To the contrary, Bateman only discloses that the “hand-coded rules” are used during “training” of the language processing model and that the language processing model resulting from that “training” is “based on” the “hand-coded rules.” *Id.* (citing Ex. 1007 [Bateman] at [0069], [0072]). However, in the context of the ’917 Patent, it is clear that **training** a model (which is the operation that actually creates the model itself) is a separate and distinct function from **inputting** data into a model to generate a model output (which is often referred to as execution

of the model). *Id.* Thus, Bateman’s disclosure of using “hand-coded rules” as a basis for “training” the language processing model does not amount to “*inputting*” the “hand-coded rules” into the language processing model as would be required by Limitation 1[c] under Wealth’s own mapping.³

Second, Wealth’s Petition fails to establish that Bateman’s “hand-coded rules” amount to the claimed “*rule data*.” Indeed, the ’917 Claims recite that the “*rule data compris[es] one or more rules associated with the resources*,” and that the claimed “*resources*” comprise “*one or more assets*” being distributed to “*one or more entities comprising one or more resource recipients*.” Ex. 1001 [’917 Patent] at 39:3-17, 39:20-23. Consistent with this claim language, the ’917 Patent explains that the “*one or more rules*” may take the form of (1) “one or more access policies, one or more authorization policies, one or more authentication policies, one or more computing device permissions, one or more file access permissions, and/or assignment of one or more assets (e.g., computing devices) to one or more entities

³ This deficiency of Bateman is further illustrated by Wealth’s improper attempt to rely on the same disclosure of training the language processing model based on “hand-coded rules” for both the claimed “*rule data*” that is “*input[]*” into the “*one or more machine-learning models*” and the claimed “*training rule data*” that is used to “*train[]*” the “*one or more machine-learning models*.” Compare Pet. 38-39 with *id.* at 41. In the context of the ’917 Claims, it is clear that these are separate and distinct functional limitations, and Wealth cannot rely on Bateman’s singular disclosure of using “hand-coded rules” to train the language processing model to meet both of these limitations.

(e.g., computing device users)” and/or (2) “one or more laws (e.g., one or more federal laws and/or one or more state laws) that can be associated with various aspects of the resources including probate law, resources law, tax law, inheritance law, and/or property law.” *Id.* at 8:24-45.

However, Wealth fails to identify any disclosure or suggestion in Bateman that the “hand-coded rules” are associated with “*resources*” comprising “*one or more assets*” being distributed to “*one or more entities comprising one or more resource recipients*,” as required under a proper construction of the ’917 Claims. *See* Pet. 37-40; Ex. 2001 [Thompson Decl.], ¶¶ 214-217. Instead, Wealth contends that the “hand-coded rules are associated with the resources because the language processing model uses them to identify conveyances of resources,” but even if this characterization of Bateman’s disclosure were true, it does not establish that the “hand-coded rules” themselves are “*associated with the resources*” as required by the ’917 Claims. Ex. 2001 [Thompson Decl.], ¶ 217.

For these reasons, none of the three types of alleged “rule data” in Bateman amounts to “*rule data comprising one or more rules associated with the resources*” that is input together with “*one or more [prompt] responses*” into the “*one or more machine-learning models*,” as required by Limitation 1[c] of the ’917 Claims.

Lastly, there is nothing in either Racanelli or Dintenfass that cures these deficiencies of Bateman. *Id.*, ¶ 220. This is confirmed by the fact that, as noted

above, Wealth’s Petition does not rely on Racanelli or Dintenfass for **any** aspect of the ML-based functionality of the ’917 Claims – let alone the specific details of the second execution stage of the claimed “*one or more machine-learning models*.”

Accordingly, the combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest the specific details of the second execution stage of the claimed “*one or more machine-learning models*” recited by Limitation 1[c], including the requirement of “*inputting the one or more [prompt] responses and rule data comprising one or more rules associated with the resources into the one or more machine-learning models*.”

3. The combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest that “*the one or more machine-learning models are trained based on training data comprising training resource data and training rule data*”

In addition to reciting the specific details of the second execution stage of the claimed “*one or more machine-learning models*” discussed above, Limitation 1[c] further recites that:

the one or more machine-learning models are trained based on training data comprising training resource data and training rule data[.]

Ex. 1001 [’917 Patent] at 39:26-29 (Limitation 1[c]).

Wealth’s Petition also relies solely on Bateman for this additional requirement and maps it to Bateman’s disclosure of training its language processing model (Pet.

41), but Bateman fails to disclose or suggest that its language processing model is trained based on both “*training resource data and training rule data.*”

Wealth’s Petition maps the claimed “*training resource data*” to Bateman’s “structured conveyance records” and maps the claimed “*training rule data*” to two types of alleged “rule data” in Bateman, which are (1) “unstructured text documents with provisions for conveying property, such as wills” and (2) “hand-coded rules.” *Id.* However, neither of these types of alleged “rule data” amounts to the claimed “*training rule data*” as a POSITA would understand the ordinary and customary meaning of that term in the context of the ’917 Patent. *See* Ex. 2001 [Thompson Decl.], ¶¶ 224-228.

First, a POSITA would understand the claimed “*training rule data*” to refer to training data that specifically represents, identifies, or encodes “*rules*” associated with “*resources.*” *Id.*, ¶ 226. However, Bateman’s “unstructured text documents” fail to meet this requirement. *Id.* Indeed, in Bateman, the training process is not provided with any indication of whether or not the “unstructured text documents” contain “*rules*” associated with “*resources,*” and the training process merely uses the “unstructured text documents” as unstructured source text rather than as “*rule data*” when training the language processing model. *Id.* This is confirmed by the fact that the model being trained in Bateman is a language processing model that is configured to extract structured information from unstructured text documents,

rather than a model that is configured to generate “*resource data*” based on inputs that specifically include the claimed “*rule data*.”

Second, given that the “*training rule data*” is the training-stage counterpart to the claimed “*rule data*,” a POSITA would understand that the “*training rule data*” likewise requires “*one or more rules*” associated with “*resources*” comprising “*one or more assets*” being distributed to “*one or more entities comprising one or more resource recipients*.” *Id.*, ¶ 227 (citing to Ex. 1001 [’917 Patent] at 21:16-19, 39:20-33). However, as explained above, Wealth fails to identify any disclosure or suggestion in Bateman that the “hand-coded rules” meet this requirement. *Id.*

For these reasons, neither of the types of alleged “rule data” used to train Bateman’s language processing model amounts to the claimed “*training rule data*.”

Lastly, there is nothing in either Racanelli or Dintenfass that cures this deficiency. *Id.*, ¶ 229. This is confirmed by the fact that, as noted above, Wealth’s Petition does not rely on Racanelli or Dintenfass for *any* aspect of the ML-based functionality of the ’917 Claims – let alone the requirement that “*the one or more machine-learning models are trained based on training data comprising training resource data and training rule data*.”

Accordingly, the combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest the claimed requirement that “*the one or more machine-learning*

models are trained based on training data comprising training resource data and training rule data.”

4. Summary

For the reasons explained above, the combination of Racanelli, Bateman, and Dintenfass fails to render any of claims 1, 2, 6-7, 12-14, 17, and 19 of the '917 Patent obvious under § 103. Accordingly, Wealth fails to meet its burden of establishing that any of claims 1, 2, 6-7, 12-14, 17, and 19 of the '917 Patent is more likely than not to be unpatentable as obvious under § 103.

B. Wealth Fails to Establish that Claims 3 and 8-11 are Rendered Obvious by the Combination of Racanelli, Bateman, Dintenfass, and McMillen (Ground 3)

Wealth's § 103 challenge to dependent claims 3 and 8-11 fails at least for the same reasons as its § 103 challenge to independent claim 1. In particular, as explained above, the combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest three required limitations that are common to all of the '917 Claims. Moreover, Wealth's analysis of dependent claims 3 and 8-11 relies exclusively on its prior, flawed analysis for the incorporated independent limitations and then adds McMillen only for the additional limitations of dependent claims 3 and 8-11, which fails to cure the deficiencies discussed above. *See* Ex. 2001 [Thompson Decl.], ¶¶ 233-235.

Thus, for this reason alone, Wealth's Petition fails to establish that dependent claims 3 and 8-11 are more likely than not to be unpatentable under § 103.

C. Wealth Fails to Establish that Claims 15, 16, 18, and 20 are Rendered Obvious by the Combination of Racanelli, Bateman, Dintenfass, and Kwak (Ground 4)

Wealth's § 103 challenge to dependent claims 15, 16, 18, and 20 fails at least for the same reasons as its § 103 challenge to the independent claims. In particular, as explained above, the combination of Racanelli, Bateman, and Dintenfass fails to disclose or suggest three required limitations that are common to all of the '917 Claims. Moreover, Wealth's analysis of dependent claims 15, 16, 18, and 20 relies exclusively on its prior, flawed analysis for the incorporated independent limitations and then adds Kwak only for the additional limitations of dependent claims 15, 16, 18, and 20, which fails to cure the deficiencies discussed above. *See* Ex. 2001 [Thompson Decl.], ¶¶ 238-240.

Thus, for this reason alone, Wealth's Petition fails to establish that dependent claims 15, 16, 18, and 20 are more likely than not to be unpatentable under § 103.

VII. WEALTH'S § 112 CHALLENGE OF DEPENDENT CLAIMS 4-5 IS MOOT IN VIEW OF VANILLA'S STATUTORY DISCLAIMER

Vanilla has filed a statutory disclaimer with the Patent Office for claims 4-5 of the '917 Patent. *See* Ex. 2002. Pursuant to 37 C.F.R. § 42.207(e), "[n]o post-grant review will be instituted based on disclaimed claims." 37 C.F.R. § 42.207(e). Accordingly, in view of Vanilla's statutory disclaimer, Wealth's § 112(b) challenge

of dependent claims 4-5 is moot and cannot serve as a basis for instituting any post-grant review.

While Vanilla will not otherwise address Wealth's § 112(b) challenge of dependent claims 4-5 in this Preliminary Response, Vanilla is not making any admission as to the validity of disclaimed dependent claims 4-5. Thus, no adverse inference shall result from the disclaimer, and the disclaimer shall have no effect on the non-disclaimed claims of the '917 Patent. *See Raytheon Technologies Corp. v. General Electric Company*, 993 F.3d 1374, 1379, n.4 (Fed. Cir. 2021) (internal citations and quotations omitted) ("Disclaimed claims are treated as if they never existed, ... and disclaimer does not legally constitute an admission that the subject of the disclaimer appears in the prior art.").

VIII. WEALTH'S § 112 CHALLENGE OF DEPENDENT CLAIM 13 FAILS

Wealth's Petition also challenges dependent claim 13 under § 112(d) on the ground that it is allegedly an improper dependent claim. Pet. 86-87. More specifically, Wealth contends that because Limitation 13[a] does not repeat the verbatim language of Limitation 1[c], this means that Limitation 13[a] "purports to remove" and "does not include all" of the requirements of Limitation 1[c]. *Id.* at 87. This § 112(d) challenge has no merit.

Section 112(d) requires that (1) "a claim in dependent form shall contain a reference to a claim previously set forth and then specify a further limitation of the

subject matter claimed” and that (2) “a claim in dependent form *shall be construed to incorporate by reference all the limitations of the claim to which it refers.*” 35 U.S.C. § 112(d) (emphasis added).

Given this requirement that a dependent claim be construed to incorporate by reference all the limitations of the claim to which it refers, the dependent claim need not repeat all of the language of a preceding claim in order to properly reference and incorporate the limitations of the preceding claim. *See, e.g., Hutchins v. Zoll Medical Corp.*, 492 F.3d 1377, 1382 (Fed. Cir. 2007) (finding that the term “general purpose computer” recited in each independent claim is necessarily “incorporated into every dependent claim” under § 112).

Rather, it is entirely permissible (and common) for a dependent claim to use abbreviated language when referring back to a limitation from a preceding claim as long as that abbreviated language is not in conflict with the limitation from the preceding claim. *See, e.g., Ravo v. Covidien LP*, No. 11-1637, 2014 WL 198551, at *2 (W.D. Pa. Jan. 16, 2014) (finding that a dependent claim’s use of the “shorthand” phrase “*the surgical device of claim 1*” to refer back to the “*surgical intra luminal resection and reconstruction device*” of claim 1 was “typical of the structure used when drafting dependent patent claims” and “follows the common and acceptable format and structure for incorporating by reference all the features of a prior claim into a dependent claim”).

Applying this legal framework here, it is clear that dependent claim 13 is a proper dependent claim under § 112(d). *See* Ex. 2001 [Thompson Decl.], ¶¶ 242-249. Indeed, Limitation 13[a] recites:

*wherein **the** generating, by the computing system, based on one or more responses and rule data comprising one or more rules associated with the resources, resource data comprising one or more resource data fields and one or more resource data field values that correspond to the one or more resource data fields, wherein the one or more resource data fields are based on the one or more resource document fields, and wherein the one or more resource data field values are based on the one or more responses comprises:*

Ex. 1001 ['917 Patent] at 41:14-23 (emphasis added). This limitation indisputably refers back to Limitation 1[c] of independent claim 1. *See* Pet. 86-87; Ex. 2001 [Thompson Decl.], ¶ 245. And while it does not repeat the verbatim language of Limitation 1[c], the clear reference back to Limitation 1[c] nevertheless shall be statutorily construed to incorporate the **entirety** of Limitation 1[c] into dependent claim 13 under § 112(d). *See* 35 U.S.C. § 112(d); *Hutchins*, 492 F.3d at 1382.

The fact that Limitation 13[a] uses abbreviated language to refer to Limitation 1[c] does not alter this conclusion, because as noted, it is entirely permissible for a dependent claim to use abbreviated language when referring back to a limitation from a preceding claim as long as that abbreviated language is not in conflict with the limitation from the preceding claim. *Ravo*, 2014 WL 198551, at *2. There is no such conflict here – Limitation 13[a] is entirely consistent with the requirements of Limitation 1[c]. Ex. 2001 [Thompson Decl.], ¶ 246.

While Wealth's Petition does not contend otherwise, dependent claim 13 also meets the primary requirement of § 112(d), which is that a dependent claim must specify a further limitation of the subject matter claimed in the claim to which it refers. 35 U.S.C. § 112(d). In particular, after Limitation 13[a] recites the reference back to Limitation 1[c], dependent claim 13 further specifies that the claimed functionality of Limitation 1[c] comprises:

13[b] determining, by the computing system, one or more locations associated with the one or more assets, wherein the one or more locations comprise one or more administrative divisions; and

13[c] determining, by the computing system, the one or more rules associated with the one or more locations.

Ex. 1001 ['917 Patent] at 41:24-29. These additional limitations further narrow the scope of Limitation 1[c] by requiring that the claimed “*one or more rules*” be associated with “*one or more locations associated with the one or more assets, wherein the one or more locations comprise one or more administrative divisions.*”

Ex. 2001 [Thompson Decl.], ¶¶ 247-248. And as with Limitation 13[a], these additional limitations are entirely consistent with the requirements of Limitation 1[c]. *Id.* As such, dependent claim 13 is in full compliance with § 112(d).

Wealth's attempt to rely on the *Multilayer* and *Pfizer* cases also fails, as those cases dealt with dependent claims that are materially different than dependent claim 13 of the '917 Patent. *See Multilayer Stretch Cling Film Holdings, Inc. v. Berry*

Plastics Corp., 831 F.3d 1350, 1361–62 (Fed. Cir. 2016); *Pfizer, Inc. v. Ranbaxy Laboratories Ltd.*, 457 F.3d 1284, 1291-92 (Fed. Cir. 2006).

For instance, in *Multilayer*, the dependent claim at issue attempted to add a requirement (use of LDPE in an inner layer) that was expressly excluded from the independent claim through its use of “consisting of” language, and as a result, that dependent claim was found to “contradict” the independent claim. *Multilayer*, 831 F.3d at 1361-62. In contrast, dependent claim 13 here plainly does not add a limitation that was expressly excluded from independent claim 1, so there is no such conflict between dependent claim 13 and independent claim 1.

Moreover, in *Pfizer*, the dependent claim at issue recited “non-overlapping subject matter” that was “completely outside the scope” of the claim to which it referred, and as a result, it failed to “specify a further limitation of the subject matter of the claim to which it refers” as required under § 112. *Pfizer*, 457 F.3d at 1291-92 (internal quotation omitted). In contrast, dependent claim 13 here is not outside the scope of independent claim 1 and does specify a further limitation of the claimed subject matter of independent claim 1.

Accordingly, for these reasons, Wealth’s Petition fails to establish that dependent claim 13 is more likely than not to be unpatentable under § 112(d).

IX. CONCLUSION

For the reasons explained above, Wealth's Petition fails to establish that at least one claim is more likely than not to be unpatentable under any of the grounds set forth in Wealth's Petition. As such, Wealth's Petition should be denied.

Respectfully submitted,

LEE, SULLIVAN, SHEA & SMITH LLP

Dated: July 10, 2026

By: /Rory P. Shea/

Rory P. Shea, Reg. No. 60,529
LEE SULLIVAN SHEA & SMITH LLP
656 West Randolph Street, Suite 5W
Chicago, IL 60661
(312) 754-9602
shea@ls3ip.com

*Counsel for Patent Owner
Vanilla Technologies, Inc.*

CERTIFICATE OF WORD COUNT

The undersigned certifies that the foregoing PATENT OWNER'S PRELIMINARY RESPONSE complies with the type-volume limits in 37 C.F.R. § 42.24(b)(1). According to the utilized word-processing system's word count, this paper—excluding the caption, table of contents, table of authorities, table of exhibits, certificate of word count, and certificate of service—contains 16,642 words.

Dated: July 10, 2026

By: /Rory P. Shea/

Rory P. Shea, Reg. No. 60,529
LEE SULLIVAN SHEA & SMITH LLP
656 West Randolph Street, Suite 5W
Chicago, IL 60661
(312) 754-9602
shea@ls3ip.com

*Counsel for Patent Owner
Vanilla Technologies, Inc.*

CERTIFICATE OF SERVICE

The undersigned hereby confirms that the foregoing Patent Owner's Preliminary Response was served on July 10, 2026 via email directed to counsel of record for the Petitioner at the following:

Scott Seeley (Reg. No. 76,800)
scott@eastgateip.com

Tim Seeley (Reg. No. 53,575)
tim@eastgateip.com

Patrick G. Maloney (Reg. No. 67,247)
pmaloney@corkmaloney.com

/Rory P. Shea/

Rory P. Shea, Reg. No. 60,529
LEE SULLIVAN SHEA & SMITH LLP
656 West Randolph Street, Suite 5W
Chicago, IL 60661
(312) 754-9602
shea@ls3ip.com