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

DECLARATION OF CHRISTOPHER THOMPSON

TABLE OF CONTENTS

I. INTRODUCTION 1
II. BACKGROUND AND QUALIFICATIONS 2
III. SCOPE OF ASSIGNMENT AND MATERIALS CONSIDERED 6
IV. SUMMARY OF OPINIONS 7
V. LEGAL STANDARDS 8
A. Claim Construction 8
B. Patentable Subject Matter (35 U.S.C. § 101) 9
C. Obviousness (35 U.S.C. § 103) 14
D. Dependent Claims (35 U.S.C. § 112(d)) 18
VI. VANILLA'S '917 PATENT 19
A. Overview of the '917 Claims 19
B. Overview of the '917 Patent Disclosure 25
C. Overview of the Prosecution History 35
VII. LEVEL OF ORDINARY SKILL IN THE ART 38
VIII. CLAIM CONSTRUCTION 38
IX. OPINIONS RELATED TO WEALTH'S § 101 CHALLENGE 40
A. The '917 Claims Recite a Specifically Arranged Implementation of Machine Learning 41
B. The '917 Claims Recite a Non-Conventional and Non-Generic Combination of Limitations 49
X. OPINIONS RELATED TO WEALTH'S § 103 CHALLENGES 51
A. OVERVIEW OF CITED REFERENCES 52
1. Racanelli 52
2. Bateman 56
3. Dintenfass 63
4. McMillen 65
5. Kwak 68

B. The ’917 Claims are Not Rendered Obvious by the Combination of Racanelli, Bateman and Dintenfass (Ground 2) 69

1. The combination of Racanelli, Bateman and Dintenfass fails to disclose or suggest the claimed sequential executions of the “*one or more machine-learning models*” 70

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*” 76

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*” 90

C. The ’917 Claims are Not Rendered Obvious by the Combination of Racanelli, Bateman, Dintenfass, and McMillen (Ground 3) 93

D. The ’917 Claims are Not Rendered Obvious by the Combination of Racanelli, Bateman, Dintenfass, and Kwak (Ground 4) 94

XI. OPINIONS RELATED TO WEALTH’S § 112(d) CHALLENGE 95

I, Christopher Thompson, a resident of Nolensville, TN over 18 years of age, hereby declare as follows:

I. INTRODUCTION

1. I have personal knowledge of the matters stated in this Declaration, except where stated on information and belief, and if called as a witness, I could and would testify competently thereto.

2. I understand that on April 7, 2026, Wealth, Inc. ("Wealth" or "Petitioner") filed a Petition for Post-Grant Review challenging claims 1-20 of U.S. Patent No. 12,353,917 (the "'917 Patent"), which is owned by Vanilla Technologies Inc. ("Vanilla" or "Patent Owner").

3. On behalf of Vanilla, I have been retained to provide my technical opinions in connection with certain challenges to the claims of the '917 Patent that are raised in Wealth's Petition.

4. I am being compensated for my work in this proceeding at my standard consulting rate of $650 per hour, and I am also being reimbursed for any expenses incurred in the course of my work. My compensation is not contingent in any way on my opinions or on the outcome of this matter.

II. BACKGROUND AND QUALIFICATIONS

5. My complete qualifications and professional experiences are described in my curriculum vitae, provided as Appendix A. Appendix A also includes a list of my publications.

6. The following is a brief summary of my qualifications and professional experiences that are relevant to this matter.

7. I received my Bachelor of Engineering in Computer Engineering from Vanderbilt University in 2010.

8. My professional focus is analyzing and explaining the operation of computer-enabled technologies including online platforms, mobile operating systems (including Google Android and iOS), data analytics platforms, online and mobile gaming, digital advertising, mobile baseband, network communications, network storage and file systems, and geolocation services, among others. With respect to such systems, I examine the design and architecture of software-based systems, including evaluating, mapping, and explaining their operation. I scrutinize and parse the source code that controls the behavior of these computer systems, sometimes seeking the code-based answer to specific questions, and sometimes independently investigating how the source code operates and the functions it performs. I also analyze the data generated by software-based systems including

business record analysis, large-scale data productions, web traffic, and big-data analytics.

9. For eight years, I was the Chief Technology Officer at NewsBreak, a company that provides digital shopper marketing solutions to the convenience store industry in the form of digital video and multimedia displayed in brick-and-mortar retail locations as well as fuel dispensing terminals. I am now a member of the Board of Directors and a technical advisor to NewsBreak. NewsBreak leverages consumer sales data as well as contextual variables to optimize the advertisements displayed in stores and at fueling locations at convenience stores. As Chief Technology Officer, I was responsible for NewsBreak's overall technology strategy and have personally developed the majority of NewsBreak's custom media distribution and orchestration software. I was also responsible for implementing the software that provides NewsBreak's suite of data analysis, media processing tools, data consumption analytics, content delivery latency, and other key performance metrics. Moreover, during my time leading technology at NewsBreak, I was personally responsible for developing a product recommendation platform based on machine-learning algorithms.

10. During my time as Chief Technology Officer at the Firebolt Group, I oversaw the acquisition and deployment of the Glimpse computer vision and analytics product, which was based on machine-learning and artificial intelligence.

This effort involved significant improvements to the scalability, data organization, and data governance of the product.

11. I am currently working as a software developer and/or technology consultant for Sports Vault, LLC, JMG Marketing, Vanderbilt University, and University of Kentucky. Additionally, I have previously worked as a Chief Product Officer for Nationwide Studios and as a software developer and/or a technology consultant for various companies dating back to 2012, including OptioLabs, LLC, WellPath, TalentWell, LLC, Savantis Solutions, Braze, Inc., and Macy's Department Stores.

12. Based on my experience as a Chief Technology Officer, Chief Product Officer, software developer, and technology consultant over the past 14+ years, I have extensive first-hand experience building and analyzing software architectures that incorporate artificial intelligence and machine learning, among other types of software architectures that perform advanced analytics operations.

13. I am currently a Partner at 233 Analytics, LLC, which provides technical analyses in technologically complex litigation and acquisition-related matters to law firms and other entities. In my role at 233 Analytics, I also provide software development, project management, and technology leadership services to clients in a variety of industries including the development of mobile applications, API integration, backend services, data retention and processing, machine learning,

computer vision, bioinformatics, medical research, and ERP (Enterprise Resource Planning) systems.

14. I have also previously worked on numerous legal matters involving artificial intelligence technologies and was responsible for performing detailed analysis of these platforms. Some of these cases involved analysis of genetic samples to identify anomalous conditions, such as cancer or chromosomal abnormalities. Other matters involved the use of machine learning to match likely data elements within disparate databases as well as the processing of images for the purpose of establishing likely commercial value of the depicted subject.

15. I have served as an expert in software and/or source code in more than 100 cases. These cases have involved a variety of different types of code, including Java, C++, C#, lisp, and other computer languages. I have provided analysis of source code, scripts, and software development in many of these cases. The source code at issue in these cases has pertained to a wide range of technologies and use cases including: machine learning, big-data analytics, mobile baseband, network communications, mobile OS (including iOS and Google Android), web technologies, digital advertising, computer security, distributed storage, network storage and file systems, online and mobile gaming, geolocation services, biotechnology, and others.

III. SCOPE OF ASSIGNMENT AND MATERIALS CONSIDERED

16. As noted above, I have been retained to provide my technical opinions in connection with certain challenges to the claims of the '917 Patent that are raised in Wealth's Petition, including:

(1) Wealth's contention that the claims of the '917 Patent are allegedly ineligible under § 101 [Petition Ground 1];

(2) Wealth's contention that claims 1, 2, 6-7, 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") [Petition Ground 2];

(3) Wealth's contention that dependent 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") [Petition Ground 3];

(4) Wealth's contention that dependent 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") [Petition Ground 4]; and

(5) Wealth's contention that dependent claim 13 is allegedly an improper dependent claim under § 112(d) [Petition Ground 6].[1]

17. My opinions set forth herein are based on my own knowledge and expertise as well as my analysis of the following materials:

[1] I have been informed that Vanilla will be filing a statutory disclaimer of dependent claims 4-5, and as such, I have not been asked to provide my opinions regarding any of Wealth's challenges to dependent claims 4-5.

- The ’917 Patent (Ex. 1001), its priority application (Ex. 1002), and its file history (Ex. 1003);
- Wealth’s Petition for Post-Grant Review (“Petition” or “Pet.”);
- The Declaration of Jules White, Ph.D. (Ex. 1004) (the “White Declaration” or “White Decl.”);
- The alleged prior art references that form the basis for Wealth’s obviousness challenges;
- All other exhibits submitted in support of Wealth’s Petition; and
- Any other documents discussed within this Declaration.

IV. SUMMARY OF OPINIONS

18. With respect to Wealth’s § 101 challenge, it is my opinion that (1) the claims of the ’917 Patent recite a specifically arranged implementation of machine learning rather than mere application of generic machine learning to a new data environment, which I understand to qualify as a patent-eligible improvement to technology, and (2) the ordered combination of claim limitations required by the claims of the ’917 Patent is non-conventional and non-generic.

19. With respect to Wealth’s § 103 challenges, it is my opinion that the ’917 Claims are not rendered obvious by Wealth’s alleged obviousness combinations because those combinations fail to disclose or suggest multiple independent limitations that are required by all of the ’917 Claims, including (1) the two sequential execution stages of the “*one or more machine-learning models*,” (2) the inputting of both “*one or more [prompt] responses*” and “*rule data comprising*

one or more rules associated with the resources" into the claimed "*one or more machine-learning models*" during the second execution stage, and (3) the requirement that the claimed "*one or more machine-learning models*" be trained based on *"training data comprising training resource data and training rule data*."

20. With respect to Wealth's § 112(d) challenge, it is my opinion that a person of ordinary skill in the art would consider dependent claim 13 to be a proper dependent claim that specifies a further limitation of independent claim 1.

V. LEGAL STANDARDS

21. Counsel for Vanilla has explained the relevant legal standards to me. I describe my understanding of those legal standards below. I have used my understanding of those legal standards in reaching my opinions in this matter.

A. Claim Construction

22. I understand that in this proceeding, the claims of a patent are to be construed in accordance with the ordinary and customary meaning of such claims as understood by one of ordinary skill in the context of the entire patent and its file history.

23. I understand that this is the same claim construction standard used in district court litigation.

24. I understand that, in some instances, the plain and ordinary meaning of a claim term may be readily apparent and claim construction in such cases involves

little more than the application of the widely accepted meaning of commonly understood words.

25. I understand that, in other instances, it may be necessary to look to the intrinsic evidence – which I understand to include the claims, the specification, and the prosecution/file history of the patent at issue – to construe a claim term. When looking at the intrinsic evidence, I understand that the context in which a term is used in a claim can be highly instructive. I further understand that the specification is highly relevant to claim construction and is generally considered to be the single best guide in determining the meaning of a claim term. In this respect, I understand that a claim construction that stays true to the claim language and most naturally aligns with the specification will be the correct construction. I also understand that one must refrain from importing limitations into the claims that are not required by the intrinsic evidence.

26. I understand that extrinsic evidence such as dictionaries and treatises may also be used to assist with claim construction. However, I understand that intrinsic evidence is often more reliable than the extrinsic evidence.

B. Patentable Subject Matter (35 U.S.C. § 101)

27. I understand that, in general, a patent is considered to be patent eligible under § 101 if it claims any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof. However, I

further understand that there are certain judicial exceptions to patent eligibility, which are abstract ideas, laws of nature, and natural phenomena. As such, if a claim is directed to a judicial exception itself, the claim will not be patent eligible under § 101.

28. I understand that there is a two-part test for evaluating whether a patent claim is patent eligible under § 101, which is sometimes referred to as the *Alice* test or the *Alice/Mayo* test. I further understand that the USPTO's Manual of Patent Examining Procedure ("M.P.E.P.") describes a framework for evaluating patent eligibility under § 101 that corresponds to the two-part *Alice* test.

29. I understand that the first step of the *Alice* test involves looking to the character of the claims as a whole to determine whether they are "directed to" a judicial exception.

30. I understand that *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. With respect to Prong Two, I further understand that the USPTO's guidance explains as follows:

> 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.

31. I understand that if a claim as a whole reflects an improvement to the functioning of a computer or another technology or technical field, that claim integrates a judicial exception into a practical application and thus is not "directed to" an abstract idea under *Alice* step one / USPTO Step 2A.

32. I understand that the USPTO's guidance provides the following two-part framework for evaluating when a claim as a whole reflects an improvement to the functioning of a computer or another technology or technical field:

> (1) First evaluate the specification 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.
>
> (2) If the specification sets forth an improvement in technology or a technical field, the claim must then 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.

33. I understand that claims involving artificial intelligence (AI) technology can qualify as a patent-eligible improvement to technology under *Alice* step one / USPTO Step 2A.

34. For instance, I understand there are at least two different routes by which claims involving machine-learning technology can satisfy *Alice* step one:

(1) by reciting specific improvements to the machine-learning models themselves; or

(2) by reciting a specifically arranged implementation of machine learning.

35. On the other hand, I understand that if a patent does no more than claim the application of generic machine learning to a new data environment, then it is not a patent-eligible improvement to technology under *Alice* step one / USPTO Step 2A.

36. I understand that if a claim is not "directed to" a judicial exception under *Alice* step one / USPTO Step 2A, then the claim is patent eligible under § 101. On the other hand, if a claim is "directed to" a judicial exception under *Alice* step one / USPTO Step 2A, then the claim may be further evaluated under the second step of the *Alice* test.

37. I understand that the second step of the *Alice* test requires a consideration of the elements of the claim, both individually and as an ordered combination, in order to determine whether the claim contains an "inventive

concept" that transforms the nature of the claim into a patent-eligible application of the judicial exception. I further understand that *Alice* step two corresponds to Step 2B of the USPTO's framework for evaluating patent eligibility

38. I understand that under *Alice* step two / USPTO Step 2B, consideration 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.

39. I understand that one factor to consider when evaluating whether a claim contains 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. I further understand that even if the individual limitations of a claim 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.

40. I understand that another factor to consider when evaluating whether a claim contains an inventive concept is whether the claim limitations provide improvements to the functioning of a computer or improvements to any other technology or technical field.

C. Obviousness (35 U.S.C. § 103)

41. I understand that a patent claim is "obvious" and therefore unpatentable under § 103 if the differences between the claimed subject matter and the prior art are such that the subject matter as a whole would have been obvious to a person of ordinary skill in the art at the time the invention was made.

42. I understand that the obviousness inquiry involves consideration of the following factors: (1) the scope and content of the prior art; (2) the differences between the claimed invention and the prior art; (3) the level of ordinary skill in the art at the time the invention was made; and (4) objective evidence of non-obviousness (which is often referred to as "secondary considerations").

43. I understand that prior art used to show that a claimed invention is obvious must be "analogous art." Prior art is analogous art to the claimed invention if (1) it is from the same field of endeavor as the claimed invention or (2) it is reasonably pertinent to the problem faced by the inventor.

44. I understand that a patent claim composed of several elements is not proved obvious merely by demonstrating that each of its elements was independently known in the prior art, whether those references were set forth in a single reference or in a combination of references. Instead, it must have been obvious to combine those elements in the same way as the claimed invention does.

45. I understand that either a single reference or a combination of references may be the basis for finding that a claim is obvious. Regardless of whether the basis is a single reference or a combination of two or more references, the patent challenger must show that there is a reason, suggestion, or motivation that would lead a person of ordinary skill in the art to modify or combine prior art references.

46. I understand that, when relying on a combination of prior art references, the party asserting invalidity must submit objective evidence as to how and why the prior art references would have been combined to produce the claimed invention. This typically includes evidence that a person of ordinary skill in the art would have had an apparent reason or motivation to combine or modify the prior art in a manner that yields the claimed invention (e.g., a teaching or suggestion in the prior art), as well as evidence that a person of ordinary skill in the art would have had a reasonable expectation that the combination would have succeeded.

47. I further understand that an obviousness inquiry requires consideration of not only whether a person of ordinary skill in the art "could" have made but also whether a person of ordinary skill in the art "would" have been motivated to make a combination or modification of prior art to arrive at the claimed invention.

48. I understand that when prior art teaches away from the claimed invention, that prior art ordinarily cannot be used to render the claimed invention

obvious. In this regard, I understand that prior art teaches away from the claimed invention when a person of ordinary skill in the art would be discouraged from following the path leading to the invention because of the prior art.

49. I understand that when separate prior art references teach away from one another or present conflicting teachings, these separate prior art references ordinarily cannot be combined to render the claimed invention obvious. In this respect, I understand that separate prior art references teach away from one another when their combination would negate the original intent of the prior art or would produce a seemingly inoperative device.

50. I understand that it is improper to rely on hindsight when assessing obviousness, because many true inventions might seem obvious with the benefit of hindsight. For this reason, the obviousness inquiry must be conducted from the standpoint of a person of ordinary skill in the art at the time the claimed invention was made. In this regard, I understand that defining a problem in terms of its solution reveals improper hindsight in the selection of the prior art relevant to obviousness. What is known today, and what is learned from the teachings and disclosures of the patent itself containing the claim under analysis, should not be considered. Nor should one use the patent claim as a roadmap to pick out elements of the prior art for combination.

51. I understand that if a claim element is not obvious, then the claims that depend from it are also not obvious.

52. I understand that various secondary considerations (sometimes referred to as objective indicia of non-obviousness) may support a determination of non-obviousness and that such secondary considerations must be considered as part of an obviousness analysis. I understand that secondary considerations of non-obviousness may include:

- commercial success of a product due to the merits of the claimed invention;
- a long felt need for the solution provided by the claimed invention;
- unsuccessful attempts by others to find the solution provided by the claimed invention;
- initial skepticism of the claimed invention by those skilled in the art;
- copying of the claimed invention by others;
- unexpected and superior results from the claimed invention; and
- acceptance by others of the claimed invention as shown by praise from others in the field or from the licensing or request for license of the claimed invention.

53. I understand that the evidence of any secondary consideration must have a "nexus" to the claimed invention in order for the secondary consideration of non-obviousness to be given substantial weight, which means there must be a sufficient connection between the evidence and the claimed invention.

D. Dependent Claims (35 U.S.C. § 112(d))

54. I understand that, under § 112(d), a dependent claim must contain a reference to a claim previously set forth and then specify a further limitation of the subject matter claimed.

55. I further understand that, under § 112(d), a dependent claim shall be construed to incorporate by reference all the limitations of the claim to which it refers.

56. I understand that, because § 112(d) requires a dependent claim to 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.

57. I understand that, when referring to a specific limitation from a preceding claim, a dependent claim need not repeat the language of that limitation verbatim. Instead, I understand that it is permissible for the dependent claim to use abbreviated language when referring back to the limitation from the preceding claim as long as that abbreviated language is not contrary to the limitation from the preceding claim.

58. On the other hand, I understand that if a dependent claim includes a limitation that directly contradicts a limitation from a preceding claim or fails to

specify a further limitation of the subject matter of the preceding claim, the dependent claim is invalid under § 112(d).

59. I understand that the question of whether a dependent claim is proper under § 112(d) is to be interpreted from the perspective of a person of ordinary skill in the art.

VI. VANILLA'S '917 PATENT

A. Overview of the '917 Claims

60. The '917 Patent issued with claims 1-20, of which claims 1, 17, and 19 are independent and the remainder are dependent. Ex. 1001 ['917 Patent] at 39:1-44:19.

61. The independent claims are reproduced below with the same claim limitation numbering used by Wealth's Petition:

Claim 1

1[p] *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.*

Claim 17

17[p] *A computing system, comprising:*

17[a] *one or more processors;*

17[b] *one or more non-transitory computer-readable media storing instructions that, when executed by the one or more processors, cause the computing system to perform operations, the operations comprising:*

> ***17[c]*** *generating, 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 associated with 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;

17[d] *receiving one or more responses to the one or more prompts;*

17[e] *generating, 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;*

17[f] *generating, 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*

17[g] *updating the one or more visualizations in real-time based on changes in one or more values of the one or more assets.*

Claim 19

19[p] *One or more non-transitory computer-readable media that collectively store instructions that, when executed by one or more computing devices, cause the one or more computing devices to perform operations, the operations comprising:*

19[a] *generating, 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 associated with 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;*

19[b] *receiving one or more responses to the one or more prompts;*

19[c] *generating, 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;*

19[d] *generating, 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*

19[e] *updating the one or more visualizations in real-time based on changes in one or more values of the one or more assets.*

Id. at 39:1-42, 41:57-42:34, 42:49-44:3.

62. The functional limitations 1[a]-1[e] of independent claim 1 are representative of functional limitations 17[c]-17[g] of independent claim 17 and functional limitations 19[a]-19[e] of independent claim 19. As such, my discussion of the claim language in this Declaration will be with reference to the representative limitations of independent claim 1 unless specifically noted otherwise.

63. As demonstrated above, the independent claims recite a specific machine-learning ("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*." The ML-

based architecture required by the independent claims involves five different stages of functionality, 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*."

64. This five-stage ML-based architecture required by the independent claims is shown in the following illustrative diagram:

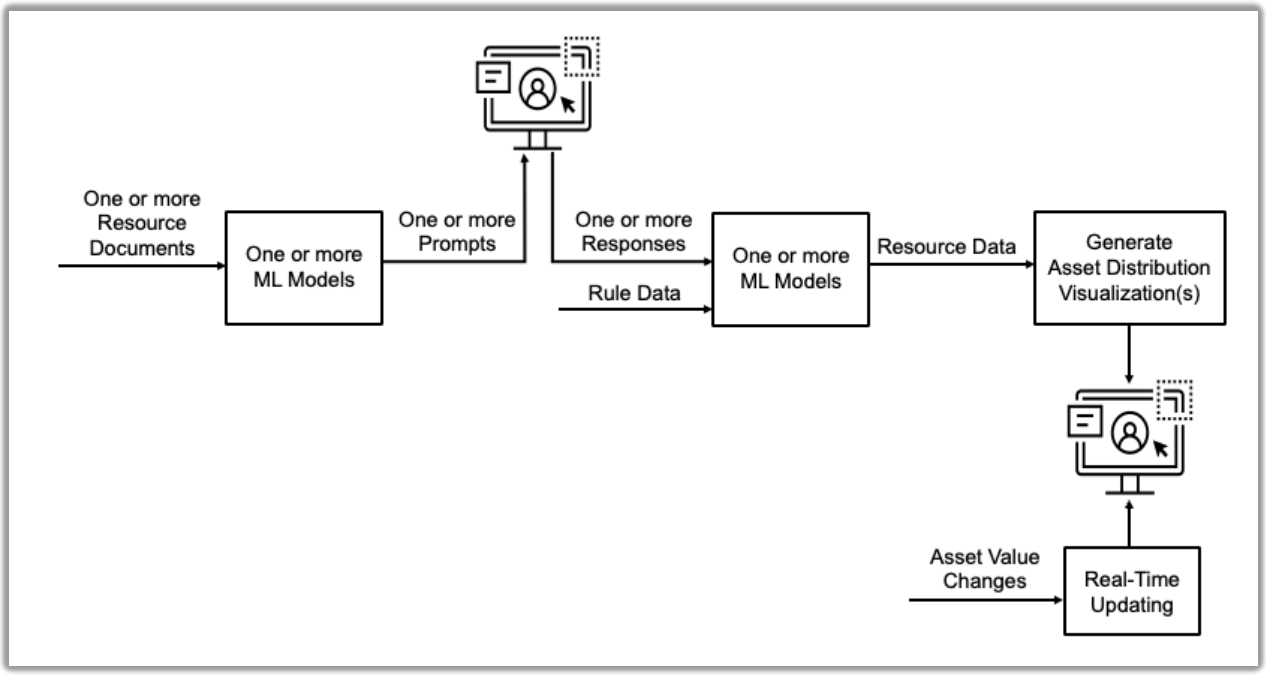

65. The foregoing ML-based architecture involves a specifically-arranged implementation of machine learning for generating the claimed "*resource data*," which involves two distinct execution stages of the claimed "*one or more machine-learning models*" that perform separate and clearly defined analytical tasks along with an intervening receipt of "*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*." These two execution stages are highlighted in the following annotated version of the above illustrative diagram:

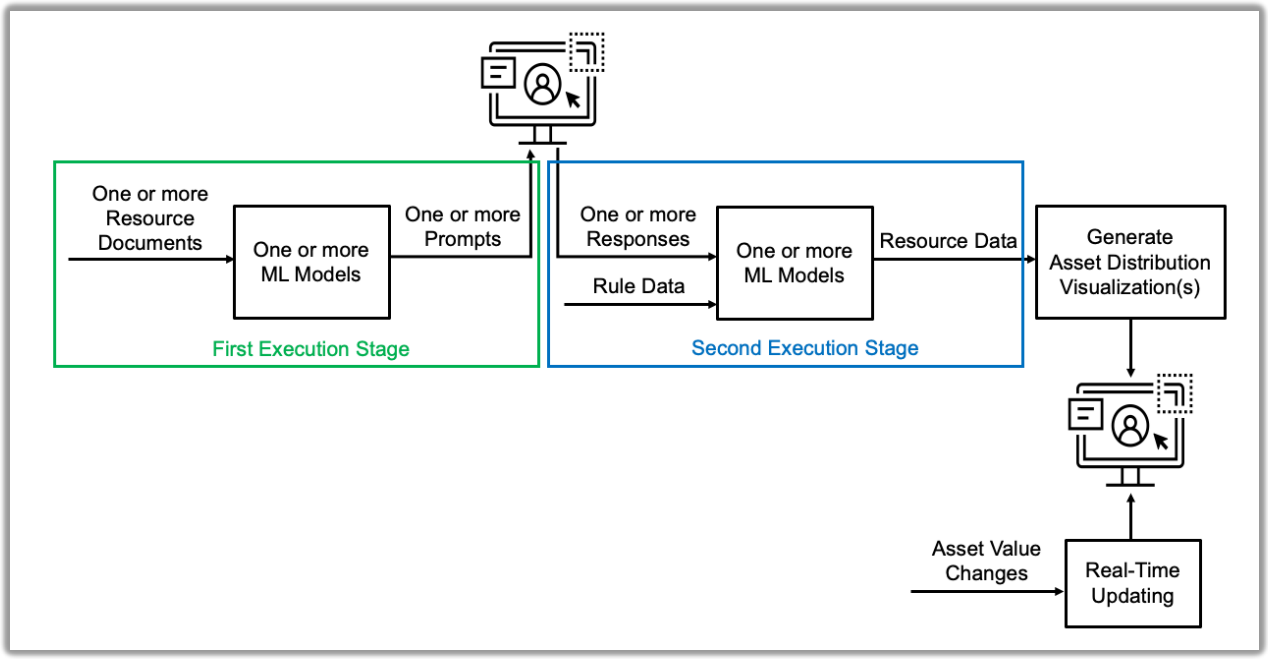

66. The dependent claims of the '917 Patent also add various other limitations to the independent claims, including but not limited to additional execution stages of the claimed "*one or more machine learning models*" and generation of additional types of "*indications*" based on the structured "*resource data*." *See* Ex. 1001 ['917 Patent] at 39:43-41:56. However, for purposes of this Declaration, I have primarily focused on the limitations of the independent claims, which are also incorporated into their entirety into all of the dependent claims.

67. For purposes of this Declaration, I will refer to the claims of the '917 Patent collectively as "the '917 Claims."

B. Overview of the '917 Patent Disclosure

68. 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 that there were problems with this existing technology. Ex. 1001 ['917 Patent] at 1:14-34. First, "the use of multiple software applications to manage resources can be burdensome and demand a significant amount of manual input." *Id*. at 1:26-31. Second, "generalized software applications may not be fully suited to meeting the often specialized requirements of resource allocation." *Id*. at 1:31-34. The disclosed technology is intended to overcome these problems with existing software technology.

69. 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.

70. One specific aspect of the disclosed technology involves a computing system that is configured with the ML-based architecture recited in the independent claims and illustrated in the above figure, 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 5:53-56; *see also id*. at 5:66-6:18 (explaining that "the disclosed technology can comprise a computing system (e.g., a resource data processing computing system) that can comprise one or more computing devices . . . that can process, send, receive, generate, and/or modify data (e.g., data associated with processing resources and/or resource allocation)" and that "the computing system can comprise and/or be associated with one or more machine-learning models that include one or more features of the one or more machine-learning models 304 that are described with respect to FIG. 3").

71. With respect to the first stage of the claimed ML-based architecture, 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," and further discloses that "the one or more resource documents can be inputted into one or more machine-learning models that are configured and/or trained to generate the one or more prompts based on the one or more resource documents." *Id*. at 4:58-5:19, 7:12-20. In turn, the '917 Patent discloses that:

> [T]he one or more machine-learning models can be 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. For example, the one or more machine-learning models can recognize portions of the one or more resource documents that are associated with one or more entities, one or more assets, and/or one or more rules (e.g., computing device policies, network policies, clauses of a will, insurance policy provisions, laws, and/or bank account policies) associated with distributing the one or more assets to the one or more entities.

Id. at 7:20-31.

72. As the '917 Patent explains, the "one or more resource documents" involved in the first stage of the claimed ML-based architecture 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, [or] healthcare directives," and may 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 5:3-16, 7:32-63.

73. Moreover, the "one or more prompts" that are generated during the first stage of the claimed ML-based architecture 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.

74. With respect to the second stage of the claimed ML-based architecture, the '917 Patent discloses that the "one or more prompts" generated by the first stage of the disclosed 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 the user may in turn 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)." *Id*. at 5:20-24, 6:19-25, 8:10-21. As one example to illustrate, the computing system may receive 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.

75. With respect to the third stage of the claimed ML-based architecture, 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." *Id*. at 5:20-39, 14:56-58.

76. In this respect, the '917 Patent explains that the "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," which "can be based on historical data (e.g., data based on actual historical resources instructions and/or actual historical resources) that can comprise the historical values of assets," 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." *Id*. at 21:8-19.

77. The '917 Patent also describes the training process for the "one or more machine-learning models" as follows:

> Configuring and/or training the one or more machine-learning models 304 can comprise the use of a cost function that can be used to minimize the error (e.g., inaccuracy) of the output of the one or more machine-learning models 304 with respect to a set of ground truth values corresponding to accurate output. For example, the training data can comprise a plurality of asset distributions associated with a plurality of resources and resources plans. The ground-truth data may indicate values associated with actual asset distributions. Accurate output by the one or more machine-learning models 304 can comprise accurately determining the asset distributions associated with the corresponding plurality of resources and resources plans. Inaccurate output by the one or more machine-learning models 304 can comprise not accurately determining the asset distributions associated with the corresponding plurality of resources and resources plans. As the one or more machine-learning models 304 are configured and/or trained, the weighting of the plurality of parameters of the one or more machine-learning models 304 can be modified until the error associated with the output of the one or more machine-learning models 304 is minimized to a predetermined level (e.g., a level associated with generating output that is at least 95% accurate). Configuring and/or training the one or more machine-learning models 304 can be performed over a plurality of rounds and/or iterations. Further, configuring and/or training the one or more machine-learning models 304 can end when a predetermined level of accuracy of the one or more machine-learning models 304 is achieved. Additionally, the one or more machine-learning models 304 can be periodically retrained based on updated training data. For example, with the passage of time additional training data based on new historical data associated with resources and/or resources plans can be added to the training data. Further, synthetic training data based on historical training data can be generated using one or more synthetic data generation techniques in which the synthetic training data can have

statistical properties that are similar to the statistical properties of actual historical data.

Id. at 22:54-23:25.

78. As the '917 Patent explains, the "one or more rules associated with the resources" that are encoded by the "rule data" used as input to the third stage of the claimed ML-based architecture can be associated with (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 (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," among other possible examples. *Id*. at 8:24-45.

79. Moreover, the "resource data" that is generated by the third stage of the claimed ML-based architecture is described as a structured dataset that "comprise[s] one or more resource data fields and/or one or more resource data field values." *Id*. at 8:46-47. These "resource data fields" may be "based on the one or more resource document fields" of the "one or more resource documents" that were input into the first stage of the claimed ML-based architecture, 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 8:46-9:7. In turn, 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." *Id*. As one example to illustrate, "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*.

80. With respect to the fourth stage of the claimed ML-based architecture, 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." *Id*. at 5:42-52, 9:42-48. As one example to illustrate, "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; *see also id*. at 9:44-49 (explaining that "the one or more visualizations of the distribution of the one or more assets to the one or more resource recipients can comprise a chart that indicates the names of the one or more resource recipients and the assets associated with each resource recipient").

81. Lastly, with respect to the fifth stage of the claimed ML-based architecture, the '917 Patent 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)." *Id*. at 9:49-52.

82. 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.

83. As the '917 Patent explains, the disclosed technology improves over the existing technology 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 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.

Id. at 14:41-15:12.

C. Overview of the Prosecution History

84. 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.

85. 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.

86. 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.

87. 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.

88. 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.

89. The issue fee was paid on June 5, 2025, and the '917 Patent was issued on July 8, 2025. *Id*. at 182, 187.

VII. LEVEL OF ORDINARY SKILL IN THE ART

90. 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).

91. For purposes of this Declaration, I have applied this level of ordinary skill in the art.

92. Based on my education and professional experience, I have at least the level of ordinary skill in the art under Wealth's proposed definition. Additionally, based on my education and professional experience, I am also familiar with the knowledge and abilities that a person of ordinary skill in the art would have had at the time of the invention of the '917 Patent, which is no later than November 29, 2023.

VIII. CLAIM CONSTRUCTION

93. In line with the Legal Standards set forth above, I understand that the '917 Claims are to be construed in accordance with the ordinary and customary meaning of such claims as understood by one of ordinary skill in the context of the entire '917 Patent and its file history.

94. In this Declaration, I have applied the ordinary and customary meaning of the terms of the '917 Claims as understood by one of ordinary skill in the context of the entire '917 Patent and its file history.

95. Wealth's Petition and the White Declaration also purport to use the ordinary and customary meaning of the terms of the '917 Claims as understood by one of ordinary skill in the context of the entire '917 Patent and its file history. Pet. 8; White Decl. ¶ 46. [2]

96. However, based on my review of Wealth's Petition and the White Declaration, it is my opinion that Wealth and Dr. White have relied on interpretations of at least two claim terms that are not consistent with the ordinary and customary meaning of those terms as understood by one of ordinary skill in the art in the context of the '917 Patent, which are: (1) "*rule data comprising one or more rules associated with the resources*" and (2) "*training rule data*." I have addressed the proper interpretation of these claim terms below in the context of my opinions related to Wealth's § 103 challenges.

[2] The lone exception is that, in Dr. White's "Claim Construction" section of his Declaration, he addresses the phrase "*one or more relationships between the one or more entities*" from dependent claims 4-5. However, I understand that this issue is moot in view of Vanilla's filing of a statutory disclaimer of dependent claims 4-5.

IX. OPINIONS RELATED TO WEALTH'S § 101 CHALLENGE

97. Wealth's Petition contends that the '917 Claims are not patent eligible under § 101 on the basis that they are (1) "directed to" an abstract idea under *Alice* step one (which corresponds to USPTO Step 2A) and (2) lack an "inventive concept" under *Alice* step two (which corresponds to USPTO Step 2B). Pet. 11-27.

98. In the sub-sections below, I explain my opinions related to certain aspects of the *Alice* two-step test and the corresponding USPTO framework for evaluating patent eligibility under § 101, which are that (1) each of the '917 Claims recites a specifically arranged implementation of machine learning rather than mere application of generic machine learning to a new data environment, and (2) the ordered combination of claim limitations required by each of the '917 Claims is non-conventional and non-generic.

99. For purposes of this Declaration, I focus my analysis on the independent limitations of the '917 Claims, because those limitations are common across all of the '917 Claims. However, my focus here on the independent limitations shall not be interpreted to mean that I agree with Wealth's or Dr. White's analysis of the dependent limitations of the '917 Claims, or that I have made any concession, admission, or waiver with respect to the dependent limitations of the '917 Claims. I understand that I may later be asked to present additional opinions

regarding Wealth's § 101 challenges that are based on the dependent limitations of the '917 Claims in a future declaration.

100. Additionally, while the sub-sections below primarily focus on my opinions that the '917 Claims are directed to an improvement to technology and recite an ordered combination of limitations that is non-conventional and non-generic, this does not mean that I agree with any of Wealth's other contentions or Dr. White's other opinions that are related to Wealth's § 101 challenge of the '917 Claims, including but not limited Wealth's contentions or Dr. White's opinions regarding whether the individual limitations of the '917 Claims were well-understood, routine and conventional to a person or ordinary skill in the art. I understand that I may be asked to provide additional opinions that are related to Wealth's § 101 challenge of the '917 Claims in a future declaration.

A. The '917 Claims Recite a Specifically Arranged Implementation of Machine Learning

101. In line with the Legal Standards set forth above, I understand that if a claim as a whole reflects an improvement to the functioning of a computer or another technology or technical field, then the claim integrates any alleged abstract idea into a practical application and thus is not "directed to" an abstract idea under *Alice* step one / USPTO Step 2A.

102. I further understand that if a claim as a whole recites a specifically arranged implementation of machine learning, that qualifies as a patent-eligible

improvement to technology under *Alice* step one / USPTO Step 2A, whereas if a claim as a whole does no more than recite the application of generic machine learning to a new data environment, that does not qualify as a patent-eligible improvement to technology under *Alice* step one / USPTO Step 2A.

103. In my opinion, a person of ordinary skill in the art would recognize that each of the '917 Claims as a whole recites a specifically arranged implementation of machine learning (and thus an improvement to technology) rather than mere application of generic machine learning to a new data environment.

104. As explained above, the '917 Claims require a specific ML-based architecture comprising 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*," 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 39:1-42.

105. This five-stage ML-based architecture is based on a specifically arranged implementation of machine learning involving two distinct execution stages of the claimed "*one or more machine-learning models*" that operate to 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 responses to the one or more prompts*" that serves as a bridge between the two distinct execution stages of the claimed "*one or more machine-learning models*."

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*," whereas 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 execution stage along with "*rule*

data comprising one or more rules associated with the resources." Additionally, in connection with the second execution stage, the '917 Claims recite that 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)*," and the "*resource data field value(s)*" are based on the "one *or more [prompt] responses*," which adds further specificity to the second execution stage.

107. This is a structured arrangement of machine learning that combines two separate, distinct, and specific execution stages of the claimed "*one or more machine-learning models*," where the "*one or more prompts*" output by the first execution stage drives one of the two specifically-defined inputs for the second execution stage by prompting a user to supplement the data that is extracted from the "*one or more resource documents*" provided as input to the first execution stage.

108. I disagree that a person of ordinary skill in the art would consider this structured, specific arrangement of machine learning to be a generalized use of machine learning.

109. It is also my opinion, under the USPTO's two-part framework for evaluating when a claim as a whole reflects an improvement to the functioning of a computer or another technology or technical field, each of the '917 Claims would be understood to be an improvement to technology.

110. In particular, under the first part of the USPTO's two-part framework, it is my opinion that the '917 Patent's disclosure provides sufficient details such that a person of ordinary skill in the art would recognize the claimed invention of the '917 Claims as providing an improvement in existing software technology for resource allocation.

111. 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.

112. Next, the '917 Patent introduces and discusses the disclosed technology. *See id*. at 3:66-9:53. As part of that discussion, 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 describes the multi-stage ML-based architecture required by the '917 Claims in extensive detail. *Id*. (emphasis added).

113. Moreover, the '917 Patent 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, 14:41-15:12.

114. It is my opinion that, in the context the '917 Patent's disclosure, a person of ordinary skill in the art would recognize these statements regarding the "improvement[s]" and "technical effect[s]" of the disclosed technology to be describing an improvement in the existing software technology for resource allocation that was previously introduced in the background section and described

as being "burdensome" and not "fully suited to meeting the often specialized requirements of resource allocation."

115. Turning to the second part of the USPTO's two-part framework, it is my opinion that the '917 Claims reflect the disclosed improvements to the existing software technology for resource allocation. In particular, as explained above, the '917 Claims include a detailed recitation of the multi-stage ML-based architecture for generating "*resource data*" and then generating (and updating) "*one or more visualizations*," which is tied to the '917 Patent's description of the improvements over existing technology.

116. I also disagree with several of the contentions that Wealth makes in support of its position that the '917 Claims are directed an abstract idea under *Alice* step one.

117. First, I disagree with Wealth's contention that the '917 Claims "merely apply generic machine learning to visualizing and making determination about resource allocations such as estate plans." Pet. 1. For the reasons explained above, it is my opinion that a person of ordinary skill in the art would consider the five-stage ML-based architecture required by the '917 Claims to be a specifically arranged implementation of machine learning, rather than mere application of generic machine learning.

118. Second, I disagree with Wealth's contention that the '917 Claims "merely gather data, process it, and display the results." Pet. 1-2. As explained above, the '917 Claims require 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 receipt of "*one or more responses to the one or more prompts*" that serves as a bridge between the two distinct execution stages of the claimed "*one or more machine-learning models*." In addition, the '917 Claims require real-time updating of "*one or more visualizations*" based on "*changes in one or more values of the one or more assets*." This functionality does not "merely gather data, process it, and display the results."

119. Third, I disagree with Wealth's contention that the functional limitations of the '917 Claims "describe results, but not how to achieve those results." In my opinion, a person of ordinary skill in the art would recognize that the functional limitations of the '917 Claims do include details as to how the claimed outputs are achieved, including by 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*," 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 "*resource data field value(s)*" are based on the "*one or more [prompt] responses*"), generating the "*one or more visualizations*" based on the structured "*resource data*," and updating the "*one or more visualizations*" "*in real-time based on changes in one or more values of the one or more assets*."

B. The '917 Claims Recite a Non-Conventional and Non-Generic Combination of Limitations

120. In line with the Legal Standards set forth above, I understand that a claim contains an "inventive concept" under *Alice* step two / USPTO step 2B if the limitations of the claim, taken both individually and as an ordered combination, recite something more than well-understood, routine, and conventional activity.

121. In my opinion, a person of ordinary skill in the art would have understood the limitations of the '917 Claims, when taken together as an ordered combination, to recite something more than well-understood, routine, and conventional activity.

122. In particular, the '917 Claims require the specifically-arranged, ML-based architecture discussed above, which comprises 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*."

123. The ordered combination and interactions between the claimed stages of the ML-based architecture required by the '917 Claims would be understood by

a person of ordinary skill in the art to be a non-conventional and non-generic arrangement that goes beyond well-understood, routine, and conventional activity. For instance, as explained above, the claimed ML-based architecture involves a sequence of two different execution stages of the claimed "*one or more machine-learning models*" that operate to perform different analytical tasks – including receiving different inputs and generating different outputs – that are configured to work together and are interconnected through an intervening stage for receiving the "*one or more responses*" to the "*one or more prompts*" output by the first execution stage. In my opinion, a person of ordinary skill in the art would not consider this ordered combination of limitations to be well-understood, routine, and conventional activity.

124. I further note that neither Wealth nor Dr. White appear to address the question of whether the ordered combination of limitations required by the '917 Claims was well-understood, routine, and conventional. *See* Pet. 21-22; White Decl. ¶¶ 74-86. Instead, Wealth and Dr. White only appear to focus on whether certain of the individual limitations of the '917 Claims, in isolation, were well-understood, routine, and conventional. However, my understanding is that the analysis under *Alice* step two / USPTO step 2B requires consideration of whether the limitations of

the claims were well-understood, routine, and conventional both individually and as an ordered combination.[3]

X. OPINIONS RELATED TO WEALTH'S § 103 CHALLENGES

125. Wealth's Petition contends that the independent claims and certain of the dependent claims of the '917 Patent are rendered obvious by a combination of Racanelli, Bateman, and Dintenfass, and that other dependent claims of the '917 Patent are rendered obvious by a combination of Racanelli, Bateman, Dintenfass, and one of McMillen or Kwak.

126. In the sub-sections below, I provide an overview of the cited references followed by my opinions related to Wealth's obviousness combinations.

127. For purposes of this Declaration, I focus my analysis on the independent limitations of the '917 Claims, because those limitations are common across all of the '917 Claims. However, my focus here on the independent limitations shall not be interpreted to mean that I agree with Wealth's or Dr. White's analysis of the dependent limitations of the '917 Claims, or that I have made any

[3] I also find certain of Wealth's contentions and Dr. White's opinions regarding the individual limitations being well-understood, routine, and conventional activity to be flawed as well. For instance, Dr. White's offers an opinion that "[m]achine-learning models trained on 'rule data' were also well-understood, routine, and conventional as of November 29, 2023" (White Decl. ¶ 83), but that opinion is premised on disclosure from the Bateman reference that does not amount to the "*training rule data*" required by the '917 Claims. I have addressed this distinction in further detail below when discussing Wealth's ¶ 103 challenges.

concession, admission, or waiver with respect to the dependent limitations of the '917 Claims. I understand that I may be asked to provide additional opinions regarding Wealth's § 103 challenges that are based on the dependent limitations of the '917 Claims in a future declaration.

A. OVERVIEW OF CITED REFERENCES

1. Racanelli

128. Racanelli is titled "Estate Disposition Modeling" and describes "a system and computer-implemented method for modeling estate disposition." Ex. 1005 [Racanelli] at Abstract.

129. Racanelli discloses functionality that involves (1) "receiv[ing] balance sheet inputs, estate planning document inputs and timing/event inputs," (2) applying "conveyance and tax logic" to the "inputs," (3) generating "[f]lowcharts illustrating the disposition of the estate based on the inputs," and also (4) generating "[t]ax and liquidity analysis based on the inputs." *Id*.

130. Figure 1 of Racanelli illustrates a disclosed embodiment that functions to "receive balance sheet 103 and estate document inputs 105 corresponding to an estate planning client, apply conveyance 109, taxation 111 and other 107 logic to the client's inputs to generate one or more customizable disposition flowcharts 121 visually depicting the flow of estate components to beneficiaries and/or an estate of the estate's future liquidity status 123":

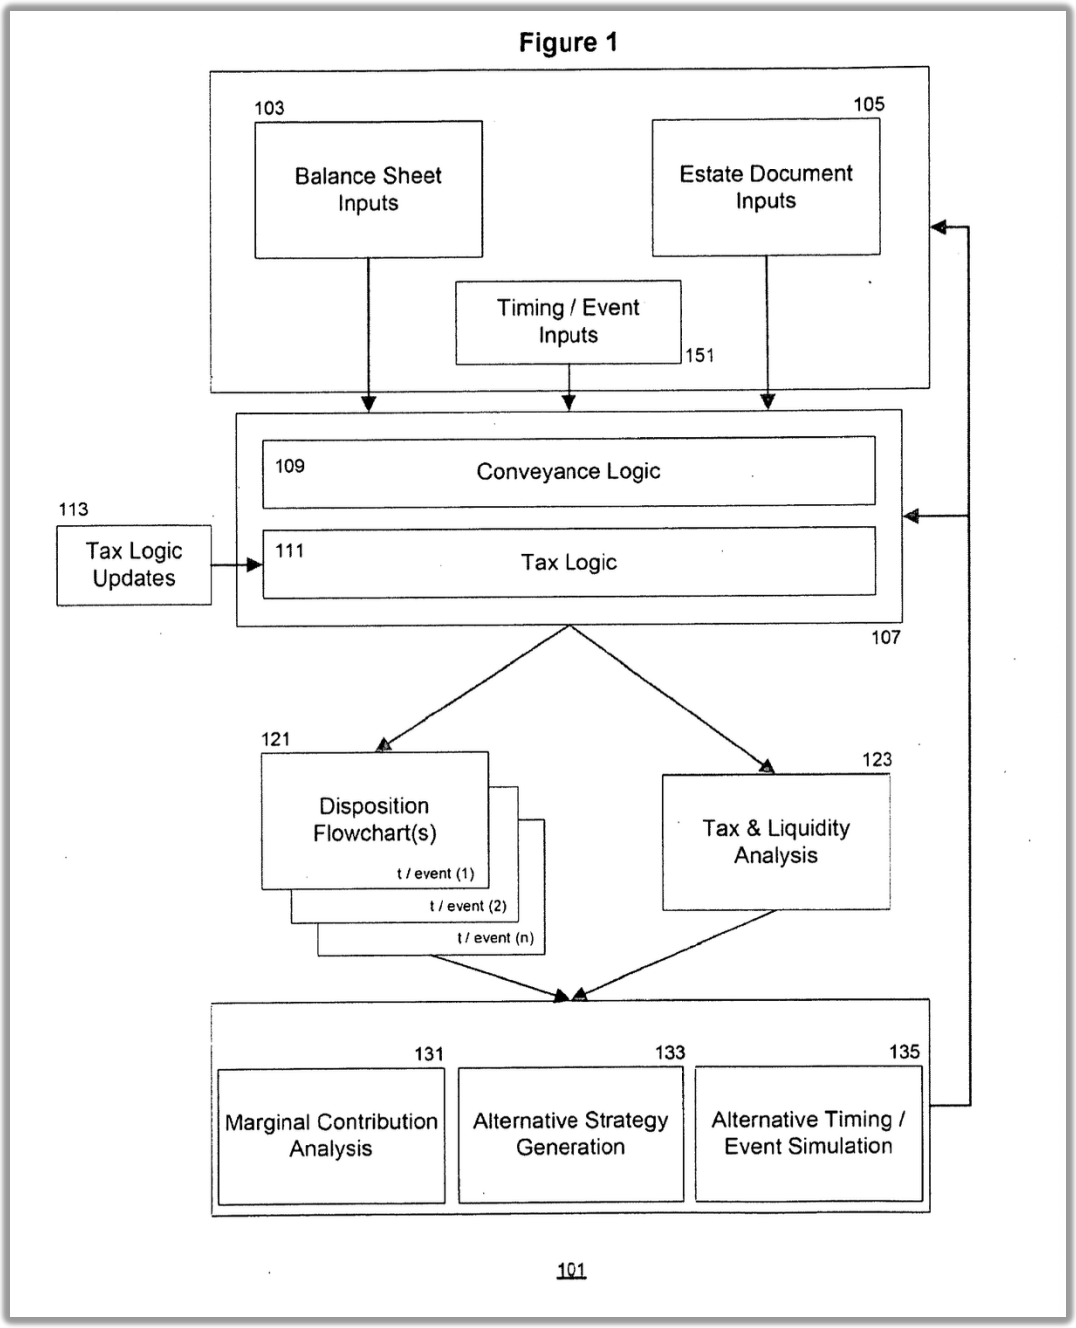

Id. at [0029], FIG. 1.

131. Figure 2 of Racanelli the illustrates "a flowchart representing an exemplary method for modeling estate disposition" that involves the following functionality:

> In one embodiment, balance sheet inputs are received S201. Estate document inputs are also received S203. The conveyance logic contained within the estate documents is applied to the balance sheet and estate document inputs S205. Application of the conveyance logic ascertains which assets, or portions thereof, from the client's balance sheet are conveyed to a particular spouse, child or other beneficiary. Additionally, tax logic, such as applicable federal and state income, death, estate, state specific inheritance taxes, and gift taxes, are applied S207. Application of the tax logic ascertains how much of the gross estate is actually conveyed to the intended beneficiaries. Upon applying this logic, a disposition flowchart representing the flow of assets among parties or financial entities is generated S211. Alternatively or in combination with the disposition flowchart, the projected liquidity status of the estate is estimated S221. Additional iterations of the disposition flowcharts, tax estimates and liquidity analysis may optionally be performed based upon subsequent events or changes in circumstances (for example, the death of a spouse or other beneficiary) or for multiple generations of conveyances (for example, to model the estate disposition after the death of children and grand-children).

Id. at [0033].

132. Racanelli describes the "balance sheet inputs" with reference to Figure 3 and explains that such "balance sheet inputs" may include "a listing of an individual's or an estate's assets 353 and liabilities 355, as well as how these assets and liabilities are held 361, 363, 365, 371" and "may be retrieved from a variety of tangible or electronic records including, but not limited to, bank statements, investment account statements, mortgage documents, deeds, titles, certificates, and stock option plan agreements." *Id*. at [0034]-[0042].

133. Further, Racanelli describes that the "estate document inputs" generally "reflect[] which assets are to be transferred to a particular survivor or beneficiary, and in what amount," and "may be retrieved from a variety of tangible or electronic records including, but not limited to, wills, codicils, revisions, amendments, revocable and irrevocable trusts, living trusts, assignments and other estate planning documents." *Id*. at [0043]-[0048]. More specifically, Racanelli explains that:

> In one embodiment, these inputs are entered into the estate documents input interface. In another, one or more estate document inputs is retrieved electronically through, for example, natural language processing or reading fields or tags in an electronic estate document format. For example, an algorithm can analyze the text of the will (for example, parsing various instructions or conveyances within the will using templates or masks to recognize common language structures contained in wills such as "I [conveyor] leave my car [asset to be conveyed] to my nephew [heir (apparent) or intended beneficiary]."") to extract or approximate the conveyances contained therein. In yet another embodiment, the balance sheet information is drawn automatically from industry-standard or proprietary electronic formats.

Id. at [0048].

134. Further yet, Racanelli describes several examples of "disposition flowcharts" with reference to Figures 6-9. *Id*. at [0053]-[0060].

135. Racanelli concludes by summarizing the advantages of its disclosed functionality as follows:

> Various disclosed embodiments advantageously allow for a user, such as a wealth advisor or estate planner, to efficiently and easily customize visual representations of the disposition of a client's estate based on various customizable circumstances. Further, disclosed embodiments allow the client to view the allocation or distribution of their assets and

> liabilities in their estate plan at a given time and under a variety of simulated circumstances.
>
> Various embodiments also provide estate planners, wealth advisors and clients with customizable and annotatable flowcharts illustrating future dispositions of the estate that adjust automatically to changes in a client's balance sheet and estate plan. Further, detecting and estimating a potential liquidity shortfall and other issues upon disposition of the estate at a future time provides a client with an opportunity to prepare for or adjust the estate plan accordingly, thereby avoiding or reducing the likelihood of such a potentially disastrous outcome for the client's heirs and beneficiaries.

Id. at [0084]-[0085].

2. Bateman

136. Bateman is titled "Natural Language Processing for Extracting Conveyance Graphs" and describes “a process for extracting conveyance records from unstructured text documents” that involves (1) obtaining “a plurality of documents describing, in unstructured form, one or more conveyances of interest in real property,” (2) “determining, . . . for each of the documents, a respective jurisdiction,” (3) “selecting, . . . from a plurality of language processing models for the English language, a respective language processing model for each of the documents based on the respective determined jurisdiction,” (4) “extracting, . . . for each of the documents, a plurality of structured conveyance records from each of the plurality of documents by applying the language processing model selected for the respective document based on the jurisdiction associated with the document,” and

"storing . . . the extracted, structured conveyance record in memory." Ex. 1007 [Bateman] at Abstract.

137. FIG. 1 of Bateman shows "a mineral-rights management system 108 configured, in some cases, to facilitate re-use and sharing of analyses of mineral rights and, in some cases, to expedite such analyses by programmatically extracting some or all of conveyance records from unstructured text documents":

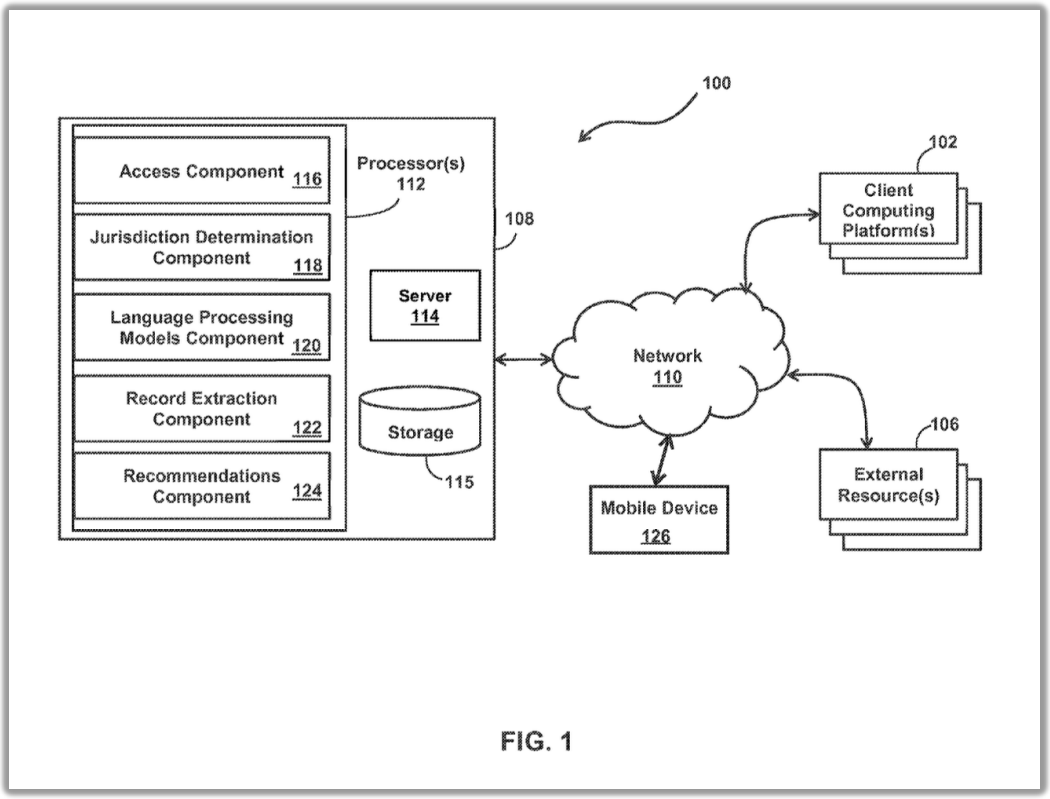

FIG. 1

Id. at [0023], FIG. 1.

138. With respect to the "mineral-rights management system 108" of FIG. 1, Bateman explains that:

> In some cases, such extractions may be based on a natural language processing (NLP) language model selected according to a jurisdiction

> in which land at issue is disposed, a selection which is expected to yield relatively accurate extractions according to models trained on terminology used in various jurisdictions. Extracted information may be presented in a manner that expedites review and revision by a human reviewer, e.g., with links to excerpts or relevant portions of governing documents. Further, some embodiments facilitate re-use of analyses and collaboration between system users by translating conveyance records into (and in some cases from) a standardized format, such that different, otherwise inoperable systems may exchange data about mineral rights.

Id. at [0023].

139. Bateman further explains that "[e]mbodiments are not limited to tracking mineral rights and include tracking components of mineral rights. For example, embodiments may track real estate generally (e.g., fee simple titles), air rights, water rights, wind rights, surface rights, beach access rights, easements, coal rights, and the like." *Id*. at [0048] (also explaining that "[t]o 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").

140. With respect to the unstructured text documents that serve as input to the disclosed functionality, Bateman explains as follows:

> Access component 116 may be configured to obtain one or more unstructured, machine-readable sets of data describing one or more conveyances of interest in real property. In some embodiments, the one or more sets of data may be a plurality of scanned documents. In some embodiments, the plurality of scanned documents may be optical-character-recognized (OCR) documents, each having OCR-produced

> English language text describing, in unstructured form, one or more conveyances of interest in real property.
>
> In some cases, the documents may be multi-page printable document format (PDF) documents that depict legally operative conveyances of interests in real property. Examples include wills, leases, deeds, easements, deeds of trust, and the like. In some cases, such documents may include text defining parties, such as prose describing a "grantor" conveying an interest (e.g., "John Doe (hereinafter 'Seller')"; prose describing a grantee (e.g., "Mary Smith (hereinafter 'Buyer')"); prose describing a date of conveyance (e.g., "effective Mar. 23, 2001"); prose describing what is being conveyed, like an interest in land (e.g., a fee simple, leasehold, mineral rights, easement, etc.); and prose describing the land at issue (e.g., the metes and bounds, lot number, or latitude and longitude coordinates of vertices of a bounding polygon). . . .
>
> The documents may contain unstructured content, e.g., English language text, that either does not have a pre-defined data model or is not organized in a pre-defined manner. For instance, information relevant to the analysis of land rights (e.g., some or all of such information) may be expressed in the documents in human-composed paragraphs written for purposes of being interpreted by human parties and a court, rather than in labeled key-value pairs or other formats more suitable for extraction by a computer. . . .

Id. at [0051]-[0053].

141. Bateman explains that the unstructured text documents may be input into "language processing models" that are configured to extract structured information from the unstructured text documents:

> The one or more language processing models may be configured to automatically (which is not to imply that other steps herein are not automatic, or that manual steps may not be interspersed with the extraction, e.g., to clear errors, select inputs, etc.) extract structured information from the one or more unstructured, machine-readable sets of data.

* * * *

> In some embodiments, language processing models may include natural language processing models configured to extract structured information from human or natural language contained in the one or more sets of data. Natural language processing (NLP) models may be configured for named entity recognition. In some embodiments, an NLP model may be configured to extract proper names from the unstructured sets of data and the type of the extracted names (e.g., person, location, and/or organization). In some embodiments, an NLP model may be configured to identify the relationships among named entities (e.g., who is the wife of whom). In some embodiments, an NLP model may be configured for co-reference resolution (e.g., determine which words refer to the same objects in a set of unstructured data).

Id. at [0061], [0066]; *see also id.* at [0063], [0074], [0107]-[0108].

142. Bateman's "language processing models" are trained based on one or both of "hand-coded rules" and/or "machine-learning algorithms":

> In some embodiments, the language processing models may include a language processing model based, at least in part, on hand-coded rules. Examples of such rules may include regular expressions (e.g., look for the text "hereby referred to as Grantor," extract the closest preceding proper noun, and designate the field "grantor" as equal to the extracted value in a record for the instrument), nested if-then-else statements (e.g., decision trees), and/or other rules.
>
> In some embodiments, the one or more language processing models may include a language processing model trained by machine-learning algorithms. A machine-learning algorithm, applied to a language processing model, may be trained, at least in part, on unstructured text documents and corresponding structured conveyance records (e.g., a collection of different instruments from a given jurisdiction and their corresponding runsheets as previously analyzed manually by a human reviewer). . . .
>
> * * * *
>
> In some embodiments, training a given language processing model may be based on a combination of machine-learning algorithms and hand-coded rules. For example, a training set for a language processing

> model may include features that are based on hand-coded rules and weights that are learned based on training data. Combining machine-learning algorithms and hand-coded rules is described in Michael Curtotti and Eric Mccreath, Corpus Based Classification of Text in Australian Contracts, Proceedings of Australasian Language Technology Association Workshop, 18-26 (2010), http://www.aclweb.org/anthology/U10-1005, which is incorporated by reference in its entirety.

Id. at [0069]-[0070], [0072].

143. The structured information that is extracted from the unstructured text documents by Bateman's "language processing models" is then encoded as parameter values for structured "conveyance records." *See id*. at [0063], [0071], [0073]-[0074], [0108].

144. In connection with the functionality for extracting structured information from the unstructured text documents to encode the parameter values for "conveyance records," Bateman further explains that multiple "candidate values" for a parameter could be extracted, in which case the "candidate values" may be presented to a "human reviewer" for selection:

> In cases where multiple candidate names [or dates] are detected, some embodiments may associate the candidates with a corresponding conveyance record field, and a human reviewer may select among the candidates during human review. The received selection, upon receipt, may be applied to the corresponding conveyance record.
>
> * * * *
>
> Extracting conveyance records may include extracting only some parameters of such records and extracting likely candidate values for such parameters. In some cases, natural language processing may be sufficient to extract, for example, four of eight parameters of a

> conveyance record, and other parameters may be supplied by a human reviewer. In some cases, extracting conveyance record parameters may include identifying a set of candidate values, ranking the candidate values according to likelihood of being the correct value as scored by the language processing model, and associating the ranking with the conveyance record, such that a human reviewer may select among the same. In some cases, the extracted values may be incorrect, but even extracting values that are correct 40% of the time is expected to result in time savings for human reviewers, who often manually enter such values 100% of the time with traditional systems.

Id. at [0067]-[0068], [0071].

145. I have prepared the following diagram that illustrates Bateman's functionality for extracting "conveyance records":

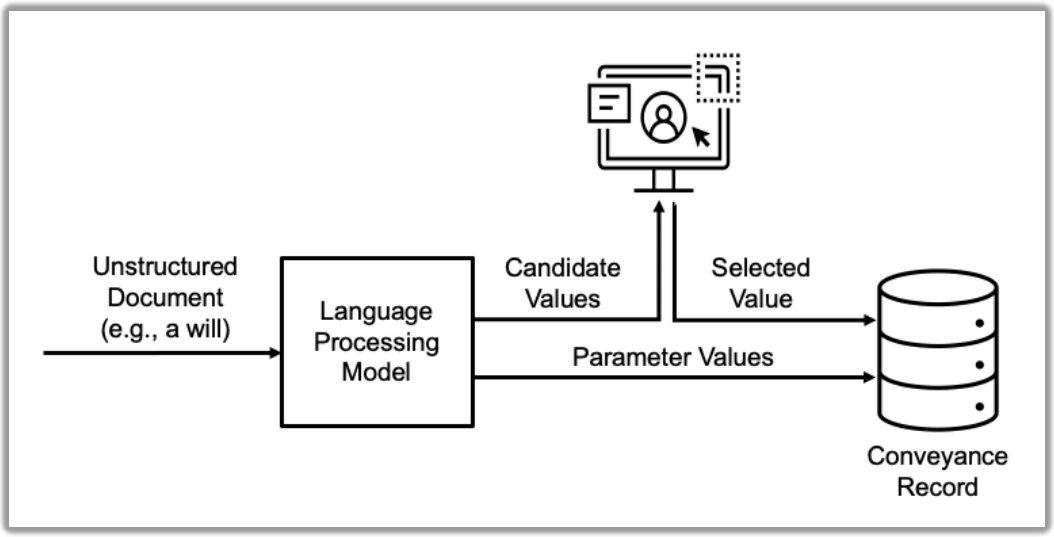

146. The "conveyance records" that are extracted may be stored in memory, and "representations" of the "conveyance records" may also be presented to a user. *See id*. at [0031], [0037], [0090]-[0092], [0095], [0109].

147. Bateman describes other functionality that may be carried out based on the "conveyance records" as well, such as scoring and/or ranking the "language

processing models," generating "recommendations and/or alerts," or interrogating a "directed graph data structure" to "detect compliance with various graph validity rules" and/or "track ownership of different types of interests in the land at issue." *Id*. at [0096]-[0100].

3. Dintenfass

148. Dintenfass is titled "System for Assessment of Allocated Assets" and describes systems configured to (1) "monitor user finances and/or a financial status associated with one or more beneficiaries," (2) "receive a user selection of a source of user asset," (3) "retrieve information associated with the selected source of user asset," (4) "transform the retrieved information into a first data format," (5) "receive a user selection of an account associated with a beneficiary," and (6) "reallocate funds from the source of user asset to the account of the beneficiary." Ex. 1009 [Dintenfass] at Abstract; *see also id*. at [0001] (explaining that the present application "relates generally to allocating assets" and more specifically to "assessment of allocated assets by transforming information received from one or more sources of user assets in real-time"); [0045]-[0055], FIG. 2.

149. With respect to the monitoring of user finances and/or a financial status associated with one or more beneficiaries, Dintenfass discloses that "the system may be configured for selectively monitoring one or more real-time data feeds from the one or more sources of user assets and financial status of the beneficiaries" and that,

"[i]n response to continuously monitoring the real-time data feeds, the system may be configured to update the asset value associated with the one or more sources of user assets and/or financial status of the beneficiaries in real-time for display on the estate portal application." *Id*. at [0049].

150. Dintenfass describes the foregoing functionality in the context of an "estate portal application." *See id*. at [0045]-[0055].

151. Dintenfass further discloses that "[i]n addition to the estate portal application, in some embodiments, the system may be configured to dynamically prepare and generate a will draft and/or a text of complex documents with alternative text provisions, insertable text and/or multiple provisions that are selected in response to information relevant to the preparation of the will draft." *Id*. at [0056]. FIG. 3 of Dintenfass "illustrates a high level process flow for dynamically retrieving information for drafting a will 300," which includes (1) "providing a will generation application," (2) "receiving an indication from a user to initiate the will generation application," (3) "receiving information associated with a state residency of the user," (4) "retrieving one or more state laws associated with the residency of the user," (5) "establishing a communication link between the user input system . . . , one or more sources of user assets, and one or more accounts associated with one or more beneficiaries," (6) "retrieving information associated with allocation of assets from the one or more sources to the one or more beneficiaries and/or to the one or

more accounts associated with the one or more beneficiaries," and (7) "dynamically generating a will draft based on at least the retrieved information, the answers to the questions posed to the user, and the one or more state laws." *Id*. at [0057]-[0059].

152. Dintenfass also describes that "the system may be configured to transmit a notification to the user indicating that the user may have to update the estate portal application based on new information," such as "a death of a beneficiary, change in state or regulatory laws since the generation of the previous will, change in the residency of the user, or the like." *Id*. at [0064].

4. McMillen

153. McMillen is titled "Digital Estate Planning System" and describes an "online system for managing beneficiaries and assets of a will" that is configured to (1) "access transaction data related to a user," (2) "identify a person or a property based on the transaction data," (3) "present a confirmation request to the user to confirm that the person is a beneficiary of the will or that the property is an asset of the will," and (4) "add, in response to receiving a confirmation from the user, the person to a list of beneficiaries of the will, or the property to a list of assets of the will." Ex. 1010 [McMillen] at Abstract.

154. One aspect of McMillen's disclosure is a "control and data flow" for "identifying potential assets or beneficiaries to include in a will," which involves analyzing "transaction data from one or more data sources" to "identify trigger

events," and using the "triggering event" to "initiate one or more processes" including "adding an asset to a will," "removing an asset from a will," adding a beneficiary to a will," and/or "removing a beneficiary from a will." *Id*. at FIG. 2, [0039]-[0040] (also explaining that these operations "may be implemented in various ways, such as with a trained machine learning model, conditional programming, directed graphs, decision making systems, or the like"); *see also id*. at [0028] (explaining that "[m]any life events may cause the need to change a will" and that such "life events may be detected by the online system 106 using systems and techniques described herein").

155. McMillen explains that the "[t]rigger events may be classified into various categories, such as 'potential asset acquisition', 'potential asset disposition', 'potential addition of beneficiary', or 'potential removal of beneficiary'." *Id*. at [0052].

156. McMillen further explains that the "trigger event" may be identified using a "model" that is "trained based on transaction data" and "may be specific to detect or infer assets or beneficiaries from the transaction data." *Id*. at [0050]-[0056].

157. Another aspect of McMillen's disclosure is a method for "assigning asset valuation" in which (1) "a trigger event indicating a potential asset purchase is received," (2) "a baseline valuation is determined," (3) "a user is prompted for an

image of the asset," and (4) "the baseline valuation is revised based on the image." *Id*. at [0057]-[0060]; *see also id*. at [0013] (stating that "the digital will platform may use machine learning techniques (e.g., convolutional neural network) to identify and classify uploaded digital content and automatically select an asset with which to associate the digital content"), [0015] (stating that "the digital will platform may include an asset valuation backend engine configured to process data associated with an asset in order to generate an estimated asset value for the asset" and that "[a] valuation engine may be configured with one or more machine learning models (e.g., CNNs, recurrent neural networks (RNNs), etc.) to process the data associated with the asset to generate a valuation estimate for the asset").

158. McMillen explains that "[i]mages may be used periodically (e.g., requested annually) to update the valuation," and that an "image-based valuation may be adjusted as time continues." *Id*. at [0062].

159. Yet another aspect of McMillen's disclosure is a method for "managing beneficiaries and assets of a will" in which (1) "transaction data related to a user is accessed," (2) "a person or a property is identified based on the transaction data," (3) "a confirmation request is presented to the user to confirm that the person is a beneficiary of the will or that the property is an asset of the will," and, (4) "in response to receiving a confirmation from the user," the person is added to a list of beneficiaries or the property is added to a list of assets." *Id*. at [0067]-[0072].

5. Kwak

160. Kwak is titled "Validity Assessment of Legal Will Statements as Natural Language Inference" and the abstract is as follows:

> This work introduces a natural language inference (NLI) dataset that focuses on the validity of statements in legal wills. This dataset is unique because: (a) each entailment decision requires three inputs: the statement from the will, the law, and the conditions that hold at the time of the testator's death; and (b) the included texts are longer than the ones in current NLI datasets. We trained eight neural NLI models in this dataset. All the models achieve more than 80% macro F1 and accuracy, which indicates that neural approaches can handle this task reasonably well. However, group accuracy, a stricter evaluation measure that is calculated with a group of positive and negative examples generated from the same statement as a unit, is in mid 80s at best, which suggests that the models' understanding of the task remains superficial. Further ablative analyses and explanation experiments indicate that all three text segments are used for prediction, but some decisions rely on semantically irrelevant tokens. This indicates that overfitting on these longer texts likely happens, and that additional research is required for this task to be solved.

Ex. 1011 [Kwak] at Abstract.

161. In the Introduction section, Kwak states that:

- "[O]ur dataset statements from wills (as hypotheses), laws (as premises), and conditions, which are circumstances at the time of probate," and that this adaptation is "necessary as will statements' validity cannot be evaluated without considering both circumstances at the time of probate ('conditions') and relevant legal rules ('laws')."
- "We create an open-access annotated dataset with 1,014 data points, generated from 23 publicly available wills"
- "We demonstrated that validity assessment of legal will statements can be handled reasonably well with state-of-the-art NLI models when trained on our dataset."

- "Our analyses indicate that in most cases the NLI classifiers use meaningful phrases from all three pieces of text, indicating the models do capture useful information" but "in some situations, the models use features that are not intuitive for humans, indicating that the task is incomplete."

Id. at 1-2.

162. Moreover, in the Conclusion section, Kwak states that "[t]he NLI models trained on our dataset showed reasonable performance in assessing the validity of will statements" and that "[a]blative experiments demonstrated that the models' performances worsen if any input type (condition, law, or both) is omitted." *Id*. at 9.

B. The '917 Claims are Not Rendered Obvious by the Combination of Racanelli, Bateman and Dintenfass (Ground 2)

163. Wealth's Petition contends that independent claims 1, 17, and 19 as well as dependent claims 2, 6-7, 12-14, 17, 19 of the '917 Patent are allegedly rendered obvious under § 103 by a combination of Racanelli, Bateman, and Dintenfass. Pet. 27-60. I disagree.

164. In my opinion, none of the '917 Claims are rendered obvious by the combination of Racanelli, Bateman, and Dintenfass because those references fail to disclose or suggest several independent limitations that are required by all of the '917 Claims.

165. In the sub-sections below, I address three clear deficiencies in Wealth's obviousness combination of Racanelli, Bateman, and Dintenfass, each of which

serves as an independent basis for my conclusion that the '917 Claims are not rendered obvious by the combination of Racanelli, Bateman, and Dintenfass. However, my focus here on these three clear deficiencies does not mean that I agree with any of Wealth's other contentions or Dr. White's other opinions that are related to Wealth's § 103 challenges of the '917 Claims, including but not limited Wealth's contentions or Dr. White's opinions regarding analogous art, rationale for combination, reasonable expectation of success, or the mapping of any other aspect of the '917 Claims. I understand that I may be asked to provide additional opinions that are related to Wealth's § 103 challenges of the '917 Claims in a future declaration.

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

166. As explained above, the '917 Claims require a specifically-arranged ML-based architecture comprising 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*."

167. As further explained above, the claimed ML-based architecture involves two execution stages of the claimed "*one or more machine-learning models* that are required to be performed in sequence, which are:

- 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. This aspect of the '917 Claims is shown in the annotated version of the illustrative diagram that I introduced above, which I have also reproduced here:

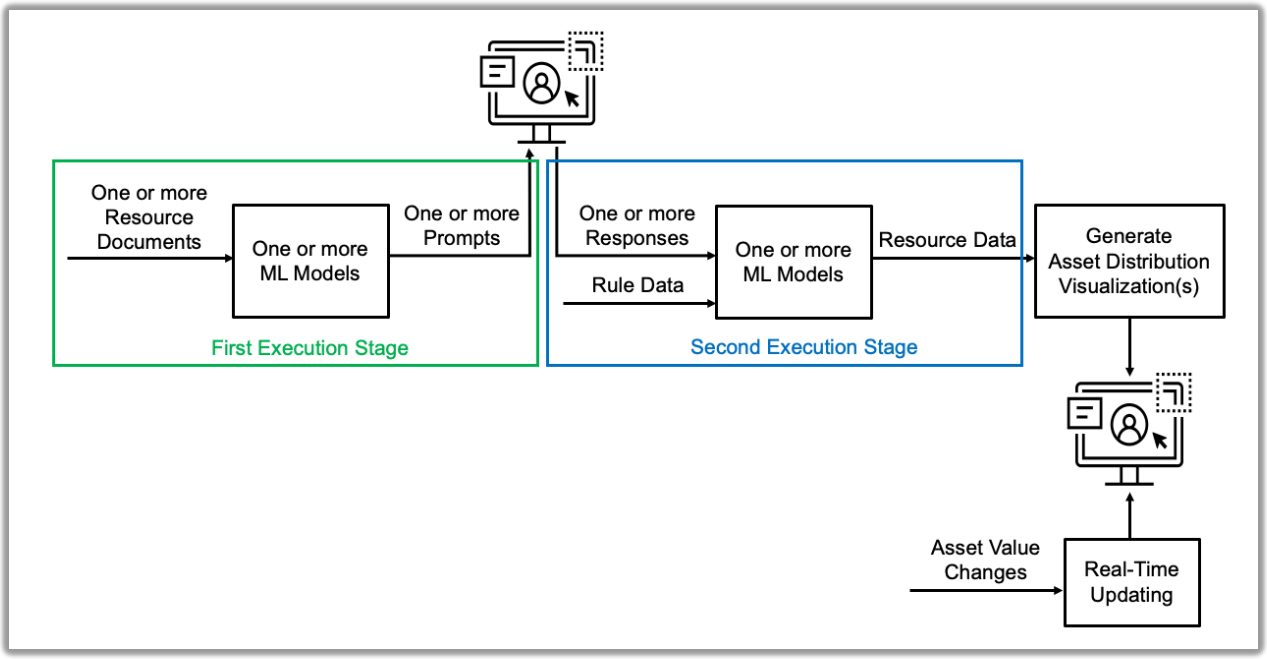

168. Wealth's Petition relies on Bateman for both of the claimed sequential execution stages of the "*one or more machine-learning models*." In particular, Wealth's Petition maps Bateman's language processing model to the claimed "*one or more machine-learning models*" and contends that the language processing model

used to both (1) to generate multiple candidate values for a parameter of a conveyance record that are presented to a human reviewer for selection and (2) "appl[y]" the reviewer's selection of a candidate value to the conveyance record. Pet. 34-35, 37. I disagree.

169. 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].

170. Separately, Bateman discloses that if the language processing model extracts multiple candidate values for a conveyance-record parameter that are presented to a human reviewer who then selects from among the candidate values, "[t]he received selection, upon receipt, may be applied to the corresponding conveyance record." *Id*. at [0067].

171. There is no teaching or suggestion in Bateman that the reviewer's selection is 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. *See id*. at [0067]-[0071].

172. A person of ordinary skill in the art would understand Bateman's disclosure that the reviewer's selection is "applied to the corresponding conveyance record" to mean that the selected candidate value is inserted into the conveyance

record, not that the selected candidate value is run back through the language processing model.

173. As such, Bateman discloses only a single execution stage of its language processing model, and the output of language processing model along with the human reviewer's selection is then applied to the conveyance record without any further use of language processing model.

174. The actual functionality of Bateman's architecture is shown in the following annotated version of the illustrative diagram I have prepared based on my review and analysis of Bateman, which confirms that Bateman's architecture only involves a single execution stage of its language processing model:

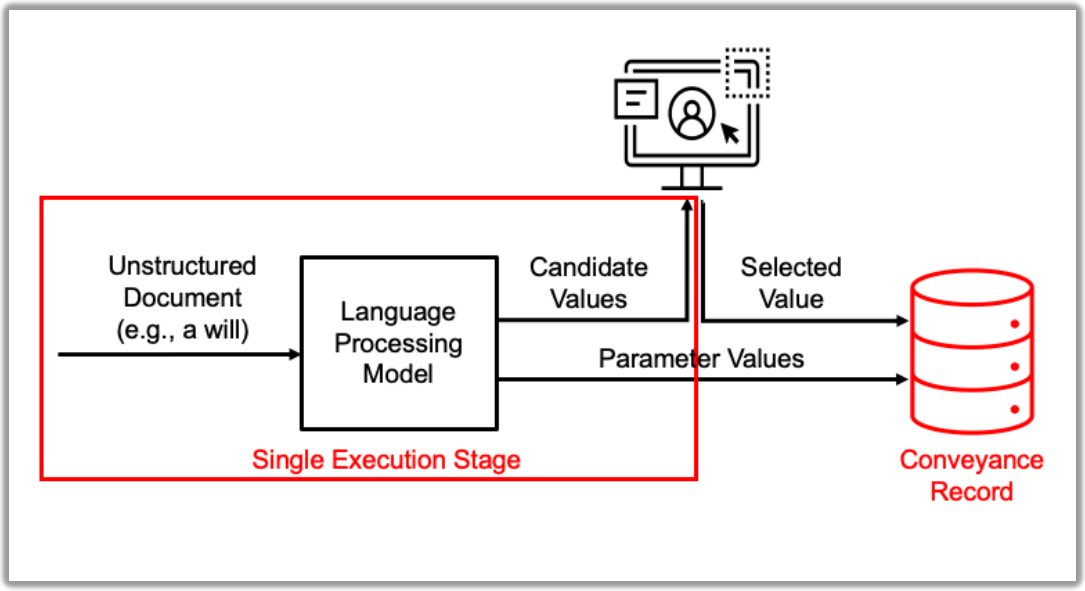

175. Notwithstanding the foregoing, Wealth's Petition includes the following contention:

> 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.).

176. However, as noted above, Bateman at [0067] 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 indicating that the reviewer's selection is "input into the language processing model," and I disagree that a person of ordinary skill in the art would interpret [0067] of Bateman in this manner.

177. I also note that, in the context of Bateman's disclosure, there is no apparent technical purpose for running the selected candidate value back through Bateman's language processing model. That is because once the human reviewer selects a particular candidate value for a parameter of the conveyance record, the selected candidate value can be "applied" to the conveyance record by inserting it directly into the conveyance record without any further need to involve the language processing model. Inputting the selected candidate value back into the language processing model a second time would serve no technical purpose in the context of Bateman's architecture as there is no need to further interpret that selected candidate value before inserting it into the conveyance record.

178. This is further confirmed by Bateman's disclosure that the purpose of its language processing model is to extract structured information from unstructured textual documents. Ex. 1007 [Bateman] at [0063]-[0066], [0070]. But there is no need to perform that function on the selected candidate value, because it was already previously extracted by the language processing model and thus is already structured information. In this respect, Wealth's contention that the language processing model is applied to selected candidate value, which is already structured information, is contrary to the described purpose of the language processing model.

179. For these reasons, it is my opinion that Bateman fails to disclose or suggest the sequential execution stages of the "*one or more machine-learning models*" that are required by the '917 Claims.

180. Wealth's Petition does not otherwise rely on Racanelli or Dintenfass in connection with the claimed requirement of having two sequential execution stages of the "*one or more machine-learning models*." *See* Pet. 34-35 (relying exclusively on Bateman for the execution of "*one or more machine-learning models*" required by Limitation 1[a]), 37-40 (relying exclusively on Bateman for the execution of the "*one or more machine-learning models*" required by Limitation 1[c]). In any event, based on my review of the Racanelli and Dintenfass references, it is my opinion that those references fail to cure this deficiency of the Bateman reference.

181. Accordingly, in my opinion, the combination of Racanelli, Bateman and Dintenfass fails to render obvious the two sequential execution stages of the "*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*"

182. As explained above, the specific details of the second execution stage of the claimed "*one or more machine-learning models*" are found in Limitation 1[c], which recites:

> *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]).

183. As further explained above, Wealth's Petition relies on Bateman for this second execution stage of the claimed "*one or more machine-learning models*." *See* Pet. 37-41.

184. However, in addition to its failure to disclose a sequence of two execution stages of the claimed "*one or more machine-learning models*," Bateman

fails to disclose or suggest any execution stage of the claimed "*one or more machine-learning models*" that involves inputting both "*one or more [prompt] responses*" and "*rule data comprising one or more rules associated with the resources*" into the claimed "*one or more machine-learning models*" as required by Limitation 1[c] of the '917 Claims.

"one or more [prompt] responses"

185. In its Petition, Wealth maps the claimed "*one or more [prompt] responses*" to the human reviewer's selection of a candidate value for a parameter of a conveyance record and contends that the human reviewer's selected candidate value is input back through Bateman's language processing model. Pet. 37.

186. However, as discussed above, Bateman fails to disclose or suggest that the human reviewer's selection candidate value is input back into the language processing model. Ex. 1007 [Bateman] at [0067]. Instead, Bateman discloses that the human reviewer's selection is "applied to the corresponding conveyance record," which a person of ordinary skill in the art would understand to mean that the selected candidate value is inserted into the conveyance record rather than run back through the language processing model. *Id*. For the reasons already explained, there is no apparent technical purpose for involving Bateman's language processing model in this operation.

187. This means that, even if one were to apply Wealth's 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 human 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*" during a second execution stage as required by Limitation 1[c] of the '917 Claims.

188. For this reason alone, Bateman fails to disclose or suggest Limitation 1[c] of the '917 Claims.

"rule data"

189. In its Petition, Wealth maps the claimed "*rule data*" on "three types of rule data" disclosed in Bateman that are allegedly input into Bateman's language processing model, which are (1) "provisions of a will," (2) "a jurisdiction-specific ontology," and (3) "hand-coded rules." Pet. 37-40.

190. However, I disagree that Bateman's disclosure with respect to any of these three types of alleged "rule data" amounts to inputting the claimed "*rule data*" together with "*one or more [prompt] responses*" (which Wealth maps to the reviewer's selected candidate value) into the claimed "*one or more machine-learning models*" (which Wealth maps to the language processing model) during a second execution stage as required by Limitation 1[c] of the '917 Claims.

191. I will address Wealth's reliance on each of these three types of alleged "rule data" in turn below.

(1) provisions of a will

192. With respect to the alleged "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" as that term is used in the '917 Patent. Pet. 37 (citing Bateman at [0052], [0063]; White Decl. ¶ 158); *see also id*. at 39 (contending that "Bateman describes inputting a will (including provisions conveying property) into a language processing model to generate structured conveyance records").

193. However, I disagree that Bateman's disclosure of inputting an unstructured will document into its language processing model amounts to the claimed requirement of inputting the claimed "*rule data*" together with "*one or more [prompt] responses*" into the claimed "*one or more machine-learning models*."

194. As an initial matter, I note that Wealth's Petition already relies on Bateman's disclosure of inputting an unstructured will document into the language processing model with respect to the first execution stage of the claimed "*one or more machine-learning models*" required by Limitation 1[a], which recites

"*inputting one or more resource documents into one or more machine-learning models*" as a basis for generating the "*one or more prompts*."

195. This is a separate and distinct function of the '917 Claims that necessarily occurs prior to 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*," because 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*."

196. Because Wealth is relying on Bateman's disclosure of inputting an unstructured will document 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], and because that function only occurs one time during each iteration of Bateman's architecture, it cannot also meet the later-occurring second execution stage of the claimed "*one or more machine-learning models*" required by Limitation 1[c].

197. Moreover, 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*," but even under Wealth's own mapping of the "*one or more machine learning models*" to the language processing model and the "*one or more [prompt] responses*" to the reviewer's selected candidate value, Bateman fails to disclose or

suggest that an unstructured will document is input into the language processing model together with the reviewer's selected candidate value.

198. In Bateman, the set of candidate values that Wealth has mapped to the "*one or more prompts*" do not even come into existence until after an unstructured will document is input into the language processing model and the language processing model extracts the candidate values from the unstructured document. This means that, at the time that an unstructured will document is input into the language processing model, there has been no extraction of the candidate values that Wealth maps to the "*prompt*," which makes it logically impossible for the unstructured will document to be input into the language processing model together with *"one or more [prompt] responses*" as required by Limitation 1[c] of the '917 Claims.

199. For at least these reasons, it is my opinion that Bateman's disclosure of inputting of an unstructured will document into the language processing model fails to disclose or suggest the claimed requirement of inputting the claimed "*rule data*" together with "*one or more [prompt] responses*" into the claimed "*one or more machine-learning models*."

(2) jurisdiction-specific ontology

200. With respect to the "jurisdiction-specific ontology," Wealth's Petition contends that (i) "jurisdiction-specific ontology includes rules associated with real

property that identify the rights and subsets of rights in real property that can be conveyed," (ii) "Bateman describes that 'some embodiments may select a corresponding ontology according to the jurisdiction in which land is disposed in the course of performing analyses described herein,'" and (iii) "Bateman's language processing model performs those analyses, which shows that the selected ontology is input into that model." Pet. 37 (Bateman at [0048], [0070], [0072]; White Decl. ¶ 159); *see also id*. at 39 (contending that "Bateman describes tracking interests in real property based on a jurisdiction-specific ontology, so ordinarily skilled artisans would have understood that the structured conveyance records Bateman generates to track those interests in real property are based on that ontology" (citing Bateman at [0048], [0074], [0108]; White Decl. ¶¶ 165-166.).

201. However, even if one were to map Bateman's "ontology" to the claimed "*rule data*," I fail to see any disclosure or suggestion in Bateman of the "ontology" being input into the language processing model.

202. Bateman's disclosure regarding the "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).

203. This passage does not disclose that a selected "ontology" is input into Bateman's language processing model.

204. Wealth's assertion to the contrary appears to be based on the chain of logic that because (i) this passage says that "some embodiments may select a corresponding ontology according to the jurisdiction in which land is disposed in the course of performing analyses described herein," and (ii) the "analyses described herein" are carried out using the language processing model, this necessarily means that "the selected ontology is input into that model." Pet. 37. I disagree.

205. In my opinion, a person of ordinary skill in the art would understand the phrase "in the course of performing analyses described herein" to refer to Bateman's overall conveyance/title-analysis workflow, which involves many different functions – including determining a jurisdiction of an unstructured text document, selecting a language processing model to utilize for processing, extracting structured conveyance information from unstructured data, presenting conveyance records to users for review and/or editing, scoring and/or ranking language processing models, generating recommendations and/or alerts based on conveyance records, and interrogating directed graph data structures that are based on conveyance-record information to determine compliance with graph validity rules. *See, e.g.,* Ex. 1007 [Bateman] at [0063], [0071], [0074], [0092], [0095]-[0100], [0108]. Most of these functions do not involve the language processing model. As

such, a person of ordinary skill in the art would not understand the phrase "in the course of performing analyses described herein" to be disclosing or suggesting that the "ontology" is input into Bateman's language processing model.

206. Moreover, given that (i) Bateman describes the "ontology" as "including records of which rights are by default subsets of other rights unless explicitly reserved or conveyed," which a person of ordinary skill in the art would understand to describe a structured dataset, and (ii) the disclosed purpose of the language processing model is to process unstructured textual documents received as input in order to extract structured information, there is no apparent technical purpose for inputting Bateman's structured "ontology" into the language processing model.

207. Lastly, even if one were to interpret Bateman's disclosure to suggest that the disclosed "ontology" is input into the language processing model, Bateman still fails to disclose or suggest that the "ontology" would be input into the language processing model together with the human reviewer's selected candidate value, which is what Wealth reads on the claimed "*one or more [prompt] responses*." On this point, I further note Wealth's Petition never presents any evidence or explanation as to how or why Bateman's language model would use this particular combination of inputs (i.e., both an "ontology" and the reviewer's selected candidate

value) to produce a conveyance record, which further highlights the deficiencies in Wealth's obviousness theory.

208. For at least these reasons, it is my opinion 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" fails to disclose or suggest the claimed requirement of inputting the claimed "*rule data*" together with "*one or more [prompt] responses*" into the claimed "*one or more machine-learning models*."

(3) hand-coded rules

209. With respect to the "hand-coded rules," Wealth's Petition contends that "Bateman discloses inputting hand-coded rules, such as regular expressions, nested if-then-else statements, and other rules into its language processing model" on the basis that (i) "Bateman describes training its language processing model with those rules" and (ii) "[t]hose hand-coded rules are associated with the resources because the language processing model uses them to identify conveyances of resources." Pet. 38 (citing Bateman at [0069], [0072]; White Decl. ¶ 160); *see also id*. at 39-40 (contending that "Bateman describes training its language processing model to extract structured conveyance records based on those rules, so the structured conveyance records are based on those rules").

210. However, I disagree that Bateman's disclosure of its "hand-coded rules" amounts to the claimed requirement of inputting the claimed "*rule data*" together with "*one or more [prompt] responses*" into the claimed "*one or more machine-learning models*."

211. As an initial matter, Bateman fails to disclose or suggest that the "hand-coded rules" are input into the language processing model during the function of generating a conveyance record. This is confirmed by Bateman's disclosure regarding the "hand-coded rules," which is as follows:

> In some embodiments, the language processing models may include a language processing model based, at least in part, on ***hand-coded rules***. Examples of ***such rules*** may include regular expressions (e.g., look for the text "hereby referred to as Grantor," extract the closest preceding proper noun, and designate the field "grantor" as equal to the extracted value in a record for the instrument), nested if-then-else statements (e.g., decision trees), and/or other rules.
>
> * * * *
>
> In some embodiments, training a given language processing model may be based on a combination of machine-learning algorithms and ***hand-coded rules***. For example, a training set for a language processing model may include features that are based on ***hand-coded rules*** and weights that are learned based on training data. Combining machine-learning algorithms and ***hand-coded rules*** is described in Michael Curtotti and Eric Mccreath, Corpus Based Classification of Text in Australian Contracts, Proceedings of Australasian Language Technology Association Workshop, 18-26 (2010), http://www.aclweb.org/anthology/U10-1005, which is incorporated by reference in its entirety.

Ex. 1007 [Bateman] at [0069], [0072] (emphasis added).

212. These passages disclose 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*. However, using "hand-coded rules" as a basis for "training" the language processing model is a separate and distinct function from "*inputting*" the "hand-coded rules" into the language processing model in order to generate the conveyance record that Wealth maps to the claimed "*resource data*." In the context of the '917 Patent, a person of ordinary skill in the art would understanding that training a model (which is the operation that actually creates the model itself) and inputting data into a model to generate a model output (which is often referred to as execution of the model) are separate and distinct operations carried out at different times. *See, e.g.,* Ex. 1001 ['917 Patent] at 20-33 (reciting separate functions of "*inputting*" data into the "*one or more machine-learning models*" and "*train[ing]*" the "*one or more machine-learning models*" based on training data).

213. As such, these passages fail to disclose or suggest that the "hand-coded rules" are "*input[]*" into the language processing model in the manner required by the '917 Claims – let alone "*input[]*" into the language processing model together with the reviewer's selected candidate value. I further note that Wealth's Petition never presents any evidence or explanation as to how or why Bateman's language processing model would use this particular combination of inputs (i.e., both "hand-

coded rules" and the reviewer's selected candidate value) to produce a conveyance record, which further highlights the deficiencies in Wealth's obviousness theory.

214. In addition to the fact that the "hand-coded rules" are not "*input[]*" into Bateman's language processing model, Wealth's Petition also fails to establish that Bateman's "hand-coded rules" amount to the claimed "*rule data*" as that term would be understood in the context of the '917 Patent.

215. 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; *see also id.* at 8:24-45 (providing examples of the claimed "*one or more rules*" that include (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 (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").

216. However, Wealth fails to identify any disclosure or suggestion that Bateman's "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*," which is how a person of ordinary skill in the art would interpret the '917 Claims.

217. Wealth attempts to overcome this deficiency in Bateman by contending that the "hand-coded rules are associated with the resources because the language processing model uses them to identify conveyances of resources." Pet. 38. I disagree with this contention for two reasons. First, Wealth's statement that the language processing model "uses" the "hand-coded rules" to "identify conveyances of resources" appears to be premised on its theory that the "hand-coded rules" are actually input into language processing model, which is incorrect for the reasons explained above. Second, that the language processing model allegedly identifies "conveyances of resources" based on "hand-coded rules" does not mean that the "hand-coded rules" themselves are "*associated with the resources*."

218. For at least these reasons, it is my opinion that Bateman's disclosure of using "hand-coded rules" to train its language processing model fails to disclose or suggest the claimed requirement of inputting the claimed "*rule data*" together with "*one or more [prompt] responses*" into the claimed "*one or more machine-learning models*."

Summary

219. For all of the reasons explained above, it is my opinion that Bateman fails to disclose or suggest the Limitation 1[c] requirement of inputting both the claimed "*one or more [prompt] responses*" and the claimed "*rule data comprising one or more rules associated with the resources*" into the claimed "*one or more machine-learning models*."

220. Wealth's Petition does not otherwise rely on Racanelli or Dintenfass in connection with this claimed requirement. *See* Pet. 37-40. In any event, based on my review of the Racanelli and Dintenfass references, it is my opinion that those references fail to cure this deficiency of the Bateman reference.

221. Accordingly, in my opinion, the combination of Racanelli, Bateman and Dintenfass fails to render obvious 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*"

222. As set forth 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 39:20-33 (Limitation 1[c]).

223. In its Petition, Wealth relies on Bateman's disclosure of training the language processing model for this aspect of Limitation 1[c]. Pet. 41. However, Bateman fails to disclose or suggest that its language processing model is trained based on both "*training resource data and training rule data*."

224. 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*.

225. However, I disagree that either of these types of alleged "rule data" amounts to the claimed "*training rule data*" as that term would be understood by a person of ordinary skill in the art in the context of the '917 Patent.

226. In my opinion, a person of ordinary skill in the art would not consider the "unstructured text documents" that are used to train Bateman's language processing model to be the claimed "*training rule data*" based on the possibility that some of those documents may include provisions for conveying property. That is because a person of ordinary skill in the art would understand that, in the context of the '917 Patent, the claimed "*training rule data"* refers to training data that

specifically represents, identifies, or encodes "*rules*" associated with "*resources*," and Bateman fails to disclose or suggest that its "unstructured text documents" meet this requirement. Rather, Bateman discloses that the "unstructured text documents" are used as unstructured source text for training a language processing model to extract structured information from unstructured text documents.

227. Moreover, in my opinion, a person of ordinary skill in the art would not consider the "hand-coded rules" that are used to train Bateman's language processing model to be the claimed "*training rule data*." This is because a person of ordinary skill in the art would understand that, in the context of the '917 Patent, the claimed "*training rule data*" is the training-stage counterpart to the claimed "*rule data*," which 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*." *See* Ex. 1001 ['917 Patent] at 39:20-33; *see also id*. at 21:16-19 (stating that "the training data 302 can comprise 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"). However, as explained above, there is no disclosure or suggestion in Bateman that the "hand-coded rules" used to train its language processing model are associated with "*resources*" comprising "*one or more assets*" being distributed to "*one or more entities comprising one or more resource recipients*."

228. For at least these reasons, it is my opinion that Bateman fails to disclose or suggest the Limitation 1[c] requirement that "*the one or more machine-learning models are trained based on training data comprising training resource data and training rule data*."

229. Wealth's Petition does not otherwise rely on Racanelli or Dintenfass in connection with this claimed requirement. *See* Pet. 37-40. In any event, based on my review of the Racanelli and Dintenfass references, it is my opinion that those references fail to cure this deficiency of the Bateman reference.

230. Accordingly, in my opinion, the combination of Racanelli, Bateman and Dintenfass fails to render obvious 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*."

C. The '917 Claims are Not Rendered Obvious by the Combination of Racanelli, Bateman, Dintenfass, and McMillen (Ground 3)

231. Wealth's Petition contends that dependent claims 3 and 8-11 are allegedly rendered obvious under § 103 by a combination of Racanelli, Bateman, Dintenfass, and McMillen. Pet. 60-74. I disagree.

232. Each of these dependent claims depends from independent claim 1, and incorporates all of the limitations of independent claim 1.

233. As explained above, it is my opinion that the combination of Racanelli, Bateman, and Dintenfass fails to render obvious independent claim 1 for at least the reason that it fails to disclose or suggest the limitations discussed above.

234. Wealth's Petition does not rely on the additional McMillen reference in connection with the limitations of independent claim 1 discussed above that are missing from the combination of Racanelli, Bateman, and Dintenfass. *See* Pet. 60-74. In any event, based on my review of the McMillen reference, it is my opinion that it fails to cure the deficiencies discussed above.

235. Accordingly, in my opinion, neither dependent claims 3 and 8-11 nor any of the other '917 Claims are rendered obvious by the combination of Racanelli, Bateman, Dintenfass, and McMillen.

D. The '917 Claims are Not Rendered Obvious by the Combination of Racanelli, Bateman, Dintenfass, and Kwak (Ground 4)

236. Wealth's Petition contends that dependent claims 15, 16, 18, and 20 are allegedly rendered obvious under § 103 by a combination of Racanelli, Bateman, Dintenfass, and Kwak. Pet. 74-85. I disagree.

237. Each of these dependent claims depends from one of the independent claims, and incorporates all of the limitations of the independent claims.

238. As explained above, it is my opinion that the combination of Racanelli, Bateman, and Dintenfass fails to render obvious the independent claims for at least the reason that it fails to disclose or suggest the limitations discussed above.

239. Wealth's Petition does not rely on the additional Kwak reference in connection with the limitations of the independent claims discussed above that are missing from the combination of Racanelli, Bateman, and Dintenfass. *See* Pet. 74-85. In any event, based on my review of the Kwak reference, it is my opinion that it fails to cure the deficiencies discussed above.

240. Accordingly, in my opinion, neither dependent claims 15, 16, 18, and 20 nor any of the other '917 Claims are rendered obvious by the combination of Racanelli, Bateman, Dintenfass, and Kwak.

XI. OPINIONS RELATED TO WEALTH'S § 112(d) CHALLENGE

241. Wealth's Petition challenges dependent claim 13 on the additional ground that it is allegedly an improper dependent claim under § 112(d). Pet. 86-87.

242. Wealth's basis for this challenge is that, while Limitation 13[a] of dependent claim 13 references Limitation 1[c], it does not repeat the verbatim language of Limitation 1[c]. Wealth's Petition includes the following side-by-side comparison of the language of Limitation 1[c] and Limitation 13[a] in support of this challenge:

1[c]	13[a]
"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"	"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"

Id. at 86.

243. According to Wealth, the fact that Limitation 13[a] does not repeat the verbatim language of Limitation 1[c] means that Limitation 13[a] "purports to remove the underlined features" from Limitation 1[c]. *Id*. at 87. I disagree.

244. As set forth above in the Legal Standards section, I understand that a dependent claim shall be construed to incorporate by reference all the limitations of a preceding claim to which it refers, and that it is permissible for a dependent claim to use abbreviated language when referring back to a limitation from the preceding claim as long as that abbreviated language is not contrary to the limitation from the preceding claim.

245. Dependent claim 13 depends from independent claim 1 and 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 . . .*," which refers to Limitation 1[c] of independent claim 1.

246. In my opinion, a person of ordinary skill in the art would interpret the language of Limitation 13[a] shown above to be abbreviated language that incorporates by reference the entirety of Limitation 1[c] – including the underlined aspects of Limitation 1[c] – regardless of the fact that those underlined aspects are not repeated verbatim by Limitation 13[a]. There is also nothing in the language of Limitation 13[a] that is inconsistent with the requirements of Limitation 1[c] or any other limitation of independent claim 1.

247. Moreover, I note that dependent claim 13 does specify a further limitation of the claimed functionality of Limitation 1[c]. In particular, 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.

This additional functionality further narrows the scope of Limitation 1[c] by requiring that the claimed "*one or more rules*" of Limitation 1[c] 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*."

248. In my opinion, a person of ordinary skill in the art would interpret this additional functionality related to how the claimed "*one or more rules*" of Limitation 1[c] are determined to be a genuine further limitation that is entirely consistent with the language and scope of Limitation 1[c], including the underlined aspects that are not repeated verbatim in Limitation 13[a].

249. For these reasons, it is my opinion that a person of ordinary skill in the art would interpret dependent claim 13 to incorporate the entirety of Limitation 1[c] and then specify a further limitation related to how the claimed "*one or more rules*" of Limitation 1[c] are determined, which meets the requirements of § 112(d).

I hereby declare that all statements made herein of my knowledge are true and that all statements made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

July 9, 2026

Date

Christopher Thompson

APPENDIX A

Christopher Thompson

233 Analytics, LLC
112 Hadley Reserve Ct.
Nolensville, TN 37135
Cell Phone (727) 481-2833
chris@233analytics.com

Education

Vanderbilt University Nashville, Tennessee *Bachelor of Engineering May 2010*
Major: Computer Engineering
Minor: Engineering Management (School of Engineering)

Certifications

Certified Information Systems Security Professional (CISSP)
Valid Through: March 2029
ISC^2 Member Number: 3242605

Publications

Journal Publications

Jules White, Chris Thompson, Hamilton Turner, Brian Dougherty, Douglas C. Schmidt, WreckWatch: Automatic Traffic Accident Detection and Notification with Smartphones, Springer Journal of Mobile Applications and Networks

Jules White, Sibohan Clarke, Brian Dougherty, Chris Thompson, Douglas C. Schmidt, R&D Challenges and Solutions for Mobile Cyber-Physical Applications and Supporting Internet Services, Springer Journal of Internet Services and Applications

Jules White, Brian Dougherty, Chris Thompson, Douglas C. Schmidt, "ScatterD: Spatial Deployment Optimization with Hybrid Heuristic / Evolutionary Algorithms," ACM Transactions on Autonomous and Adaptive Systems Special Issue on Spatial Computing.

Conference Publications

Chris Thompson, Hamilton Turner, Jules White, and Douglas C. Schmidt, Analyzing Mobile Application Software Power Consumption via Model-Driven Engineering, Proceedings of the 1st International Conference on Pervasive and Embedded Computing and Communication Systems, Algarve, Portugal, March 5-7, 2011. **Received Best Student Paper Award**

Chris Thompson, Jules White, and Douglas C. Schmidt, Analyzing Mobile Application Software Power Consumption via Model-Driven Engineering, ACM/IEEE 13th International Conference on Model Driven Engineering Languages and Systems, October 3-8, 2010 Oslo, Norway

Chris Thompson, Jules White, Brian Dougherty and Douglas C. Schmidt, Using Smartphones to Detect Car Accidents and Provide Situational Awareness to Emergency Responders, MobilWare 2010, June 28 – July 2, 2010, Chicago, IL USA. **Best Paper Recipient**

Brian Dougherty, Jules White, Jaiganesh Balasubramanian, Chris Thompson, and Douglas C. Schmidt, Deployment Automation with BLITZ, 31st International Conference on Software Engineering, May 16-24, 2009 Vancouver, Canada.

Brian Dougherty, Jules White, Chris Thompson, and Douglas C. Schmidt, Automating Hardware and Software Evolution Analysis, 16th Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS), April 13-16, 2009 San Francisco, CA USA.

Workshop Publications

Chris Thompson, Jules White, Brian Dougherty, and Douglas C. Schmidt, Optimizing Mobile Application Performance with Model-Driven Engineering, The Seventh Annual IFIP Workshop on Software Technologies for Future Embedded and Ubiquitous Systems, November 16-18, 2009 Newport Beach, California

Book Chapters

Hamilton Turner, Jules White, Chris Thompson, Krzysztof Zienkiewicz, Scott Campbell, Douglas C. Schmidt, Building Mobile Sensor Networks Using Smartphones and Web Services: Ramifications and Development Challenges, *Handbook of Research on Mobility and Computing: Evolving Technologies and Ubiquitous Impacts*, edited by Maria Manuela Cruz-Cunha and Fernando Moreira, IGI Global, Hershey, PA, USA 2009 (Accepted Proposal)

Patents

Issued

8,907,772: *System and method for automatic unsafe driving determination and notification*
9,578,445: *Systems and methods to synchronize data to a mobile device based on a device usage context*
9,363,670: *Systems and methods for restricting access to network resources via in-location access point protocol*
9,773,107: *Systems and methods for enforcing security in mobile computing*
12,450,075: *System and method for rendering a graphical user interface for defining policies to output digital content*

Pending

US20230385862A1: *System and method for targeting digital content at product terminals using reference data*
US20230316309A1: *System and method for targeting digital content at product terminals using reference data*
US20230316310A1: *System and method for targeting digital content at product terminals using reference data*

Testimony

June 2026, submitted an expert report and provided testimony on behalf of Appalachian Underwriters in Slice Insurance Technologies, Inc. v. Applachian Underwriters, Inc.
Southern District of New York, Case No.: 24-cv-09803

March 2026, submitted an expert report and provided testimony on behalf of Valtrus Innovations Ltd in Valtrus Innovations Ltd v. NTT Data Services, LLC, et al. and Vertiv Corporation *and* Vertiv Corporation v. Valtrus Innovations Ltd
Eastern District of Texas, Marshall Division, Case Nos. 2:24-cv-00361-JRG-RSP and 2:24-cv-00907-JRG-RSP

March 2026, provided testimony in court on behalf of AIDigital Operating LLC, Grace Dominique, And Zachary Hescott in Simplifi Holdings, Inc. v. AIDigital Operating LLC, Grace Dominique, And Zachary Hescott
District Court of Dallas County, Texas, 95th Judicial District, Case No. DC-26-00975

October 2025, submitted an expert report and provided testimony on behalf of Donna Crowe, et al., in Donna Crowe, et al., v. Managed Care of North America, INC. d/b/a MCNA Dental, MCNA Insurance Company d/b/a MCNA Dental, and Healthplex, Inc.
Southern District of Florida, Fort Lauderdale Division, Case No. 23-61065-CIV-SINGHAL

July 2025, submitted a declaration and provided testimony on behalf of Coupa in Coupa Software Inc. et al v. Optilogic, Inc. et al
District of Delaware, Case No. 1:24-cv-01275

June 2025, provided testimony in court on behalf of Sun Pharmaceuticals in *In the Matter of Certain Photodynamic Therapy Systems, Components Thereof, And Pharmaceutical Products Used In Combination With The Same*,
U.S. International Trade Commission, Investigation No. 337-TA1411

June 2025, submitted an expert report and provided testimony on behalf of Rule 14 in Rule 14 v. UiPath in Rule 14 v UiPath
Eastern District of Texas; Case No. 2:23-cv-627

VANILLA EXHIBIT 2001

June 2025, provided testimony in court on behalf of Csupo et. al, in Csupo et. al. v. Google LLC
Superior Court of the State of California, County of Santa Clara; Case No. 19CV352557

March 2025, submitted an expert report and provided testimony on behalf of Sun Pharmaceuticals in in *In the Matter of Certain Photodynamic Therapy Systems, Components Thereof, And Pharmaceutical Products Used In Combination With The Same*,
U.S. International Trade Commission, Investigation No. 337-TA1411

February 2025, submitted an affidavit and provided testimony in court on behalf of ZOLL Data Systems, Inc. in ZOLL Data Systems, Inc v. Wave HDC, LLC and Experian Health, Inc
District of Colorado, Case No. 1:24-cv-00584-DDD-KASA

January 2025, submitted a declaration and provided testimony in support of claim construction on behalf of Rule 14 in Rule 14 v. UiPath in Rule 14 v UiPath
Eastern District of Texas; Case No. 2:23-cv-627

January 2025, submitted expert reports and provided testimony on behalf of Csupo et. al, in Csupo et. al. v. Google LLC Superior Court of the State of California, County of Santa Clara; Case No. 19CV352557 and on behalf of Taylor et. al, in Taylor et. al. v. Google LLC
Northern District of California; Case No. 5:20-CV-07956-VK

September 2024, provided testimony in court on behalf of IBM in International Business Machines v. Zynga, Inc
District of Delaware; Case No. 1:22-cv-00590

May 2024, provided testimony on behalf of IBM in International Business Machines v. Rakuten, Inc. et al.
District of Delaware; Case No. 1:21-cv-00461-LPS

March 2024, submitted an expert report/provided testimony on behalf of IBM in International Business Machines v. Zynga, Inc
District of Delaware; Case No. 1:22-cv-00590

March 2024, submitted an expert report/provided testimony on behalf of Ravgen Inc. in Ravgen Inc. v. Ariosa Diagnostics Roche Sequencing Solutions, Inc., Roche Molecular Systems, Inc., and Foundation Medicine, Inc
District of Delaware; Case No. 20-cv-1646-RGA-JLH

January 2024, provided testimony in court on behalf of Ravgen Inc. in Ravgen Inc. v. Natera, Inc.
Western District of Texas; Case No. 6:20-cv-451

September 2023, submitted an expert report/provided testimony on behalf of Ravgen Inc. in Ravgen Inc. v. Natera, Inc.
Western District of Texas; Case No. 6:20-cv-451

June 2023, submitted an expert report/provided testimony on behalf of Guardant Health in Guardant Health in Twinstrand Biosciences Inc., & University of Washington v. Guardant Health
District of Delaware; Case No 1:21-cv-01126-LPS

April 2022, submitted declarations and provided testimony in court on behalf of Brown et. al. in Brown et. al. v. Google, LLC et. al
Northern District of California; Case No. 20-CV-03664-LHK

September 2021, submitted an expert report/provided testimony on behalf of Ravgen Inc. in Ravgen Inc. v. Natera, Inc.
Western District of Texas; Case No. 6:20-cv-451

Technology Consulting and Experience (ongoing)

Sports Vault, LLC Nashville, Tennessee
Technology Consultant May 2026 – Present
- Responsible for the architecture and implementation of a custom fantasy-sports platform

NewsBreak Media Networks Nashville, Tennessee
Chief Technology Officer/Technical Advisor January 2015 – Present
- Responsible for ground-up architecture of marketing automation tools for brick and mortar, convenience retail
- Designed and built critical software components and infrastructure for video distribution to remote and distributed player software
- Tasked with defining and implementing the product roadmap and technology vision

Firebolt Group Wixom, MI
Chief Technology Officer July 2022 – April 2024
- Responsible for overseeing new technology initiatives and software development
- Engaged to provide leadership and structure for early-stage CV and analytics technology acquisition
- Onboarded and grew multinational development and operations team

JMG Marketing
Technology Strategy Consultant and Developer March 2021 - Present
- Provided development resources for print and business process automation

Vanderbilt University
Technology Strategy Consultant and Developer January 2018 – Present
- Research IT consulting and software development for research-related projects.

Vanderbilt University Medical Center
Technology Strategy Consultant and Developer January 2019 – Present
- Research IT consulting and software development for research-related projects. Entity separated from Vanderbilt University.

University of Kentucky
Technology Strategy Consultant and Developer August 2020 – Present
- Research IT consulting and software development for research-related projects.

233 Analytics, LLC
Consultant April 2010 - Present
- Provided development resources for custom enterprise applications and mobile applications along with consulting expert services for software patent litigation.

Technology Consulting and Experience (historical)

Macy's Department Stores
Third-party software developer June 2013 – August 2022
- Software development services for back-office automation of loss prevention division.

Braze, Inc
Third-party software developer May 2019 – December 2021
- Developed a mobile application to support developer and user events demonstrating the Braze user-engagement platform.

VANILLA EXHIBIT 2001

Savantis Solutions
Technology Strategy Consultant September 2019 – December 2019
- Provided technology and software development process guidance for custom product development pertaining to point of sale and e-commerce systems

TalentWell, LLC
Technology Strategy Consultant January 2018 – December 2019
- Provided CTO and technology strategy services for the professional development space.

Nationwide Studios Nashville, Tennessee
Chief Product Officer December 2017 – November 2018
- Responsible for overseeing development and architecture of photography capture and commerce products
- Designed and implemented data analytics platform to leverage historical customer data to increase product efficacy
- Managed migration of legacy back office systems to modern architectures to increase efficiency and agility

WellPath
Third-party software developer January 2017 – August 2019
- Provided mobile app development and data analytics services for the prison healthcare space.

Nationwide Studios, Inc
Chief Product Officer June 2017 – November 2018
- Responsible for overseeing development and architecture of photography capture and commerce products
- Designed and implemented data analytics platform to leverage historical customer data to increase product efficacy
- Managed migration of legacy back office systems to modern architectures to increase efficiency and agility

OptioLabs, LLC Nashville, Tennessee
Senior Software Development Consultant March 2012 – December 2015
- Worked on numerous projects including a secure Android middleware layer, command and control systems for remote devices, custom data analytics packages and other business services
- Architected and managed cloud computing resources including EC2 instances and infrastructure
- Automated deployment and testing frameworks

Cyber Physical Systems, Inc. Nashville, Tennessee
Founder & CEO May 2010 – May 2013
- Raised over $2.1m in venture funding for consumer electronic startup
- Architected and implemented command and control server for distributed network of devices
- Managed product development cycle to bring concept to physical product
- Featured in Engadget, TechCrunch, VentureBeat, FoxNews.com among others

Institute for Software Integrated Systems Vanderbilt University, Nashville, Tennessee
Undergraduate Research Assistant Summer 2009
- Constructed 4 mobile applications based on the Android platform
- Conducted research on mobile performance analysis and power-consumption optimization

VANILLA EXHIBIT 2001

- Presented findings and projects to upper-level administrators within the School of Engineering and Vanderbilt University.

VANILLA EXHIBIT 2001