IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

ZESTY.AI, INC.,
Petitioner

v.

AON RE, INC.,
Patent Owner

Case IPR2025-01360
U.S. Patent No. 11,030,491

DECLARATION OF DR. DAVID A. FORSYTH IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 11,030,491

Mail Stop "PATENT BOARD"
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

I, David A. Forsyth, hereby declare as follows.

I. INTRODUCTION

1. I have been retained as an expert witness on behalf of Petitioner Zesty.ai, Inc. ("Petitioner") for the above-captioned *inter partes* review (IPR). I am being compensated for my time in connection with this IPR at my standard consulting rate, which is $500 per hour. My compensation is not contingent on either my expert opinion or the outcome of this litigation. I have no other interest in this proceeding.

2. I understand that this declaration accompanies a petition for IPR involving U.S. Patent No. 11,030,491 ("the '491 patent") (EX1001), which resulted from U.S. Patent Application No. 17/155,633 ("the '633 application"), filed on January 22, 2021. I understand that the '491 patent alleges a priority date of September 23, 2016. I refer to this date throughout this declaration.

3. In preparing this declaration, I have reviewed the '491 patent and each of the documents cited herein, in light of general knowledge in the art before September 23, 2016. In formulating my opinions, I have relied upon my experience, education, and knowledge in the relevant art. In formulating my opinions, I have also considered the viewpoint of a person of ordinary skill in the art ("POSA") (i.e., a person of ordinary skill in the field of image processing and machine learning) prior to September 23, 2016.

II. MY BACKGROUND AND QUALIFICATIONS

4. My qualifications and credentials are fully set forth in my *Curriculum Vitae*, which is attached as Exhibit 1004. A summary of my qualifications is included below.

5. I am currently the Fulton-Watson-Copp Chair in Computer Science at the University of Illinois at Urbana-Champaign. I believe that my background and expertise qualify me as an expert in the technical issues in this matter.

6. My education began in Cape Town, South Africa. I was an undergraduate at the University of the Witwatersrand, Johannesburg, and I hold a Bachelor of Science (1984) and Master of Science in Electrical Engineering from that University. I was awarded the Diocesan College Rhodes Scholarship in 1984 to attend Oxford University. I hold a Master of Arts by special election from Oxford University and a Doctorate of Philosophy in Engineering Science from Balliol College, Oxford. I was appointed Fellow by Examination of Magdalen College, Oxford in 1989. I was then appointed Assistant Professor of Computer Science at the University of Iowa in 1991; in 1994, I was promoted to Associate Professor of Computer Science, and went on leave of absence to take up an appointment as Assistant Professor of Computer Science at U.C. Berkeley. I was promoted to Associate Professor of Computer Science at U.C. Berkeley in 1996, and to Professor of Computer Science in 2002. I went on leave of absence from

U.C. Berkeley to take up an appointment as Professor of Computer Science at the University of Illinois at Urbana-Champaign in 2004. I am currently the Fulton-Watson-Copp Chair in Computer Science at the University of Illinois at Urbana-Champaign.

7. I have studied Computer Vision since my final year as an undergraduate in 1984, when I engaged in an undergraduate project on the topic. Computer vision is a field of computer science that uses computing methods to interpret images and video, as well as image-like data, in ways that are useful. My Ph.D. thesis treats a traditional problem in computer vision, known as "color constancy," and the paper that resulted from this work is still quite regularly cited. As a Fellow by Examination, I studied illumination and shading effects in vision, and geometric methods to recognize objects. I have published papers on a wide range of topics in computer vision, computer graphics, machine learning and human-computer interfaces. I have published about 220 refereed papers and 10 book chapters. I have published one textbook on computer vision in two editions and four languages, one research monograph, and two recent textbooks on machine learning. I have edited six volumes of collected papers. A reasonably complete list of my publications is contained in my curriculum vitae in Exhibit 1004.

8. In addition, I have occupied various leadership posts within my discipline. I was program co-chair for the Institute of Electrical and Electronics

Engineers ("IEEE") Conference on Computer Vision and Pattern Recognition ("CVPR") in 2000, 2011, 2018 and 2021, and for the European Conference on Computer Vision ("ECCV") in 2008. I was general co-chair for CVPR in 2006 and again held the position for CVPR 2015. I was general co-chair for the IEEE International Conference on Computer Vision in 2019. I am an Associate Editor of the International Journal of Computer Vision; of the Association for Computing Machinery ("ACM") Transactions on Graphics; and of the Journal of the ACM. I was the Editor in Chief of the IEEE Transactions on Pattern Analysis and Machine Intelligence ("TPAMI") for four years. TPAMI is a journal which receives approximately 1000 submissions per year, and which is usually seen as the leading journal in the discipline. I also served for three years on a committee of the National Research Council, convened to study methods to protect children from pornography and other inappropriate material on the internet.

9. Throughout my career I have received a variety of awards. In 1993, the Marr prize for Best Paper at the International Conference on Computer Vision was awarded to "Extracting Projective Structure from Single Perspective Views of 3D Point Sets," a paper I wrote with Charles A. Rothwell, Andrew Zisserman, and Joseph L. Mundy. In 2002, the award for Best Paper in Cognitive Computer Vision at the European Conference in Computer Vision was awarded to "Object recognition as machine translation: Learning a lexicon for a fixed image

vocabulary," a paper I wrote with P. Duygulu, K. Barnard, and J.F.G. de Freitas. I was awarded an IEEE Technical Achievement Award in 2005 for my work on Computer Vision. I became an IEEE Fellow in 2009, and an ACM Fellow in 2014. In 2024, I received the Mark Everingham Prize at the European Conference in Computer Vision.

10. My work is extensively cited in computer vision literature, with 44,249 citations recorded by Google scholar in July, 2025.

11. My appended CV lists all publications authored in the past 10 years, with perhaps a few exceptions that I have forgotten.

12. I have made extensive contributions to computer vision education. I am lead co-author of a widely used textbook in the discipline, "Computer Vision - A Modern Approach", which has appeared in two editions and four languages. I have written two other textbooks, and am the lead co-author of the computer vision chapter in the leading Artificial Intelligence textbook. I have trained 35 Ph.D. students. I have trained people who are very well known academics in computer vision. Tamara Berg was associate professor of computer science at the University of North Carolina at Chapel Hill and is now an entrepreneur. Ali Farhadi is professor of computer science at the University of Washington and CEO of the Allen Institute for Artificial Intelligence. Deva Ramanan is a tenured professor of computer science at Carnegie Mellon University. Gang Wang was assistant

professor of computer science at Nanyang Technological University in Singapore, and moved to become the senior lead in AI research at Alibaba. Zicheng Liao is a professor of computer science at Zhejiang University. Kobus Barnard is professor of computer science at University of Arizona. Each is seen as a young leader in the field.

13. Based on my above-described 35 years of experience in the field of computer vision and the acceptance of my publications and professional recognition by societies in my field, I believe that I am considered to be an expert in the field of computer vision (including image analysis). For example, I have published numerous papers on interpreting images, including: a highly cited paper on identifying object attributes from images (A. Farhadi, I. Endres, D. Hoiem, and D.A. Forsyth, "Describing Objects by their Attributes," CVPR 2009); the first paper demonstrating a method that can write sentences that describe an image (A Farhadi, M Hejrati, M A Sadeghi, P Young, C Rashtchian, J Hockenmaier, D. A. Forsyth, "Every Picture Tells a Story: Generating Sentences from Images," Proc ECCV 2010); the first paper demonstrating frame-rate object detection (M.A. Sadeghi and D.A. Forsyth, "30Hz Object Detection with DPM V5," Proc ECCV 2014); methods to recover surface material descriptions from images (D.A. Forsyth and Jason Rock, "Intrinsic Image Decomposition using Paradigms," IEEE TPAMI, November, 2022); and papers on interpreting the layout of rooms (V. Hedau, D.

Hoiem and D.A. Forsyth, "Understanding the layout of cluttered rooms," Int. Conf. Computer Vision (ICCV) 2009). I have published papers on transforming images, including: methods to relight images (Anand Bhattad, D.A. Forsyth, "StyLitGAN: Prompting StyleGAN to Produce New Illumination Conditions," CVPR 2024); methods to restore face images (Min Jin Chong, Dejia Xu, Yi Zhang, Zhangyang Wang, David Forsyth, Gurunandan Krishnan, Yicheng Wu, Jian Wang, "Copy or not? Reference-based Face Image Restoration with Focus on Fine Details," Winter Conference on Applications of Computer Vision, 2025); and methods to identify synthetic images (Ayush Sarkar, Hanlin Mai, Amitabh Mahapatra, Svetlana Lazebnik, D.A. Forsyth, Anand Bhattad, "Shadows Don't Lie and Lines Can't Bend! Generative Models don't know Projective Geometry…for now," CVPR 2024).

III. SUMMARY OF OPINIONS

14. In forming my opinions about the '491 patent, I have considered the following grounds of unpatentability. Based on my review of the prior art references that form the basis of these grounds, it is my opinion that claims 1–20 of the '491 patent would have been obvious to a POSA prior to September 23, 2016.

Ground	Basis	Claims	References
1	§103	1, 6–11, and 16–20	Gross
2	§103	2–5	Gross and Furukuwa
3	§103	12–15	Gross and Davis

15. I have been asked to consider how a POSA would have understood the challenged claims in light of the disclosures of the ’491 patent. I have also been asked to consider how a POSA would have understood the prior art references that form the basis of the above-listed grounds.

16. Further, I have been asked to consider and provide my technical review, analysis, insights, and opinions regarding whether a POSA would have understood that the prior art reference listed in the table above render obvious claims 1–20 of the ’491 patent.

IV. LIST OF DOCUMENTS CONSIDERED

17. In formulating my opinions, I have relied upon my training, knowledge, and experience that are relevant to the ’491 patent. I have also reviewed and am familiar with the following documents and materials in addition to any other documents cited in this declaration:

Exhibit No.	Description
1001	U.S. Patent No. 11,030,491 to Okazaki (“’491 patent”)
1002	Prosecution History of U.S. Patent No. 11,030,491 (“Prosecution History”)

Exhibit No.	Description
1004	*Curriculum Vitae* of Dr. David A. Forsyth
1005	U.S. Publication No. 2015/0186953 to Gross et al. ("Gross")
1006	U.S. Patent No. 10,755,357 to Davis et al. ("Davis")
1007	U.S. Patent No. 6,790,593 to Furukawa ("Furukawa")
1008	Jafri, *et al.*, "A Survey of Face Recognition Techniques," *Journal of Information Processing Systems*, Vol. 5, No. 2, June 2009, pp. 41–68 ("Jafri")
1009	U.S. Patent No. 8,078,436 to Pershing et al. ("Pershing")
1010	U.S. Patent No. 6,885,771 to Takahashi ("Takahashi")
1011	U.S. Publication No. 2009/0265193 to Collins et al. ("Collins")
1012	U.S. Publication No. 2015/0302529 to Jagannathan ("Jagannathan")
1013	U.S. Patent No. 4,941,122 to Weideman ("Weideman")
1014	U.S. Patent No. 8,995,757 to Ciarcia et al. ("Ciarcia")
1015	Redistricting Data Hub, "TIGER Boundary Files," accessible at https://redistrictingdatahub.org/data/about-our-data/tiger-boundary-files/ (last accessed July 10, 2025)
1016	Norman, J., "Origins of Google Earth," History of Information, accessible at https://www.historyofinformation.com/detail.php?entryid=3145 (last accessed July 13, 2025)
1017	Krizhevsky, *et al.*, "ImageNet Classification with Deep Convolutional Neural Networks," Communications of the ACM, Vol. 60, No. 6, May 2017
1018	Lin, *et al.*, "Network In Network," 2013
1019	Bradski, "The OpenCV Library," accessible at https://www.drdobbs.com/open-source/the-opencv-library/184404319 (last accessed July 10, 2025)
1020	van der Walt, *et al.*, "scikit-image: image processing in Python," *PeerJ*, June 2014

Exhibit No.	Description
1021	"Cartographic Work in Japan 2000 – 2005," United Nations Economic and Social Council, June 2005
1022	Usery, *et al*., "Mapping Developments and GIS in the USGS, 1884–2009," International Cartographic Association, January 2009

18. To the best of my knowledge, the above-mentioned documents and materials are true and accurate copies of what they purport to be. An expert in the field would reasonably rely on them to formulate opinions such as those set forth in this declaration.

V. LEGAL UNDERSTANDING

19. I have also relied upon various legal principles (as explained to me by Zesty.ai's counsel) in formulating my opinions. My understanding of these principles is summarized below.

A. My Understanding of Claim Construction

20. I understand that during an *inter partes* review proceeding, claims are to be construed in light of the specification as would be read by a person of ordinary skill in the relevant art at the time the application was filed. I understand that claim terms are given their ordinary and customary meaning as would be understood by a person of ordinary skill in the relevant art in the context of the entire disclosure. A claim term, however, will not receive its ordinary meaning if the patentee acted as his own lexicographer and clearly set forth a definition of the

claim term in the specification. In this case, the claim term will receive the definition set forth in the patent.

B. My Understanding of Obviousness

21. I understand that a patent claim is invalid if the claimed invention would have been obvious to a POSA at the time the application was filed. This means that even if all of the requirements of the claim cannot be found in a single prior art reference that would anticipate the claim, the claim can still be invalid.

22. To obtain a patent, a claimed invention must have, as of the priority date, been nonobvious in view of the prior art in the field. I understand that an invention is obvious when the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious to a POSA at the time the invention was made.

23. I understand that to prove that prior art or a combination of prior art renders a patent obvious, it is necessary to:

(1) identify the particular references that, singly or in combination, render the patent obvious;

(2) specifically identify which elements of the patent claim appear in each of the asserted references; and

(3) explain how the prior art references could have been combined in order to create the inventions claimed in the asserted claim.

24. I also understand that prior art references can be combined under several different circumstances. For example, it is my understanding that one such circumstance is when a proposed combination of prior art references results in a system that represents a predictable variation, which is achieved using prior art elements according to their established functions.

25. I also understand that when considering the obviousness of a patent claim, one should consider whether a teaching, suggestion, or motivation to combine the references exists so as to avoid impermissibly applying hindsight when considering the prior art. I understand this test should not be rigidly applied, but that the test can be important to avoid such hindsight.

26. I understand that certain objective indicia can be important evidence as to whether a patent is obvious or nonobvious. Such indicia include: (1) commercial success of products covered by the patent claims; (2) a long-felt need for the invention; (3) failed attempts by others to make the invention; (4) copying of the invention by others in the field; (5) unexpected results achieved by the invention as compared to the closest prior art; (6) praise of the invention by the infringer or others in the field; (7) the taking of licenses under the patent by others; (8) expressions of surprise by experts and those skilled in the art at the making of the invention; and (9) the patentee proceeded contrary to the accepted wisdom of the prior art.

27. At this point, I am not aware of any secondary indicia of non-obviousness. But, I reserve the right to review and opine on any evidence of objective indicia of nonobvious that may be presented during this proceeding.

28. I also understand that "obviousness" is a legal conclusion based on the underlying factual issues of the scope and content of the prior art, the differences between the claimed invention and the prior art, the level of ordinary skill in the prior art, and any objective indicia of non-obviousness.

29. For that reason, I am not rendering a legal opinion on the ultimate legal question of obviousness. Rather, my testimony addresses the underlying facts and factual analysis that would support a legal conclusion of obviousness or non-obviousness, and when I use the term obvious, I am referring to the perspective of one of ordinary skill at the time of invention.

C. A Person of Ordinary Skill in the Art

30. I understand that a person of ordinary skill in the relevant art ("POSA") is presumed to be aware of all pertinent art, thinks along conventional wisdom in the art, and is a person of ordinary creativity—not an automaton.

31. I have been asked to consider the level of ordinary skill in the field that someone would have had at the time the claimed invention was made. In deciding the level of ordinary skill, I considered the following:

- the levels of education and experience of persons working in the field;
- the types of problems encountered in the field; and
- the sophistication of the technology.

32. My opinion below explains how a POSA would have understood the technology described in the references I have identified herein around the September 23, 2016 timeframe, which I have been advised is the earliest possible effective filing date for the ’491 patent. However, my opinions herein would not be affected if the ’491 patent is found to only be entitled to a later effective filing date.

33. Regardless of whether I use “I” or a “POSA” during my technical analysis below, all of my statements and opinions are always to be understood to be based on how a POSA would have understood or read a document at the time of the alleged invention.

VI. OVERVIEW OF THE ’491 PATENT

A. Summary of the ’491 Patent

34. The ’491 patent discloses techniques for analyzing aerial imagery to evaluate property characteristics and condition classifications. EX1001, Abstract. For example, this might include analyzing images to determine the maintenance levels of certain property features, such as the roof. EX1001, 1:46–51, 2:39–54. This information can then be used, e.g., to estimate the risk of damage in the event

of "disaster conditions, such as severe storms," to calculate repair costs, or to confirm a property has been repaired. EX1001, 2:47–54; 12:29–33 ("The types of disasters can include, in some examples, earthquake, hurricane, tornado, storm surge, fire, straight line winds, or explosion. The types and estimated severity of disasters, in some embodiments, may depend upon the particular property location.").

35. The '491 patent teaches applying "machine learning analysis" to "analyze aerial imagery and automatically extract characteristics of individual properties, providing *fast and efficient automated classification* of building styles and repair conditions," in order to perform risk evaluations. EX1001, 2:21–30, 39–47 (emphasis added). Figure 1, reproduced below, illustrates an "example system operational flow 100 for matching aerial image features to corresponding property maintenance levels." EX1001, 6:21–23:

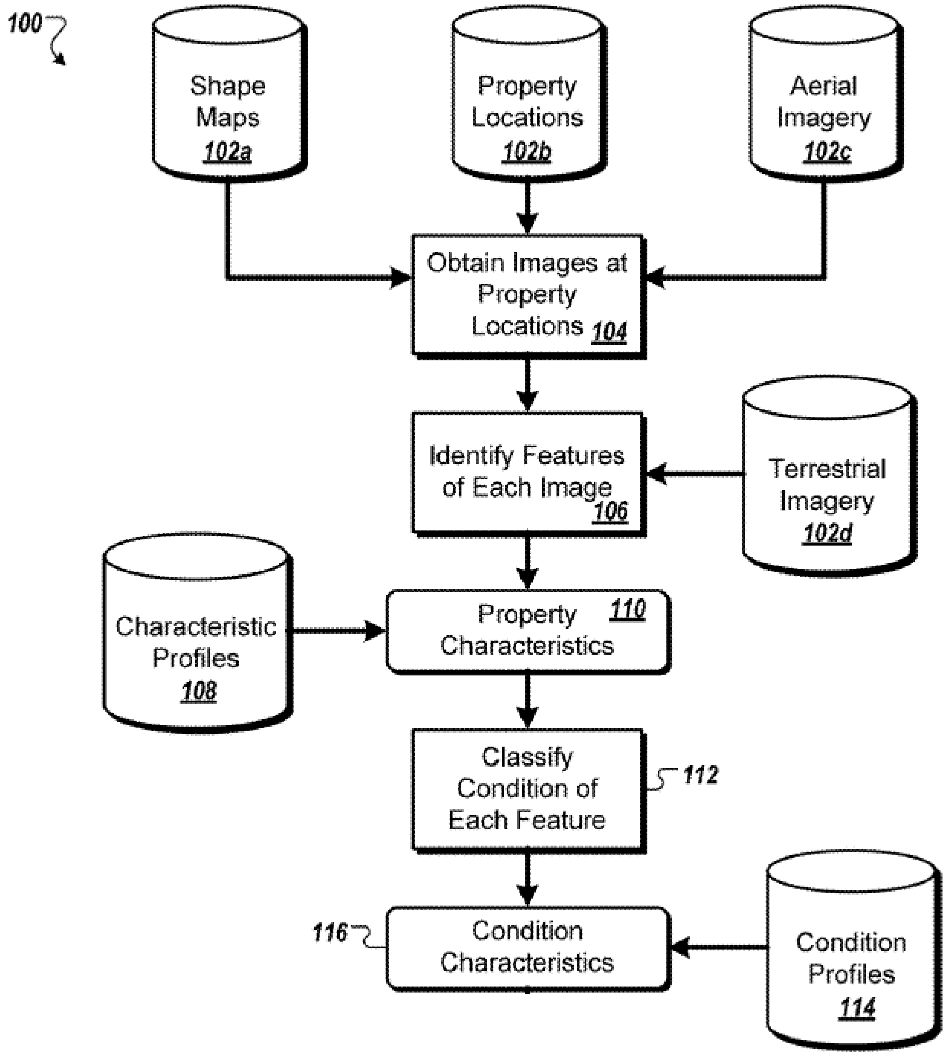

EX1001, FIG. 1 (cropped).

36. As shown in Figure 1, the system first obtains images, including e.g., "aerial imagery" and "shape map images," at property locations at step 104. Such "shape map images" can "include a block shape layout of existing properties" (e.g., "urban planning maps"). EX1001, 7:46–49. The shape map images can be used to confirm the location of a particular property in the aerial image. EX1001, 7:58–59, FIG. 1. In the embodiment of claim 1, "boundary information" is applied to the

aerial image to “isolate the property location.” EX1001, Claim 1. At step 106, the system then identifies (or “extracts”) “features” of each aerial image, such as “substantially homogenous fields” and “pixel groupings” (groups of pixels). EX1001, 8:30–40. These “features” represent a characteristic of the property, such as a “rooftop” or a “tree.” EX1001, 8:34–44; *see also* EX1001, 2:66–3:9 (“The property characteristic of a structure such as a house, factory, or barn, in some examples, can include a rooftop, porch, chimney, or skylights. Property characteristics of manmade structures within a property location, in some examples, can include a deck, swimming pool, shed, gazebo, detached garage, tennis court, fence, retaining wall, dock, playground equipment, equipment or vehicles, or greenhouse. Property characteristics of natural features within a property location, in some examples, include trees, ponds, marshes, rivers, lakes, grasses, cliffs, or ocean shore.”).

37. Then, at step 110, the system determines a “characteristic classification” for the identified property characteristic from the image. EX1001, 8:40–50. For example, this may include determining the “type of rooftop” (e.g., gambrel, gable, hipped, square, flat). EX1001, 8:40–50, 9:23–43, FIG. 2A; *see also* EX1001, 3:9–16 (“Further, property characteristic classifications can include shapes, materials, size (breadth and/or height, relative or actual), and/or distance of the property characteristic from other features. In some embodiments, the property

characteristic classification may include existing versus not existing (e.g., in the event of determining replacement or removal of a manmade feature at a property location such as a fence).”). This classification is done using “characteristic profiles 108,” which include machine learning analysis algorithms trained to classify property characteristics. EX1001, 8:50–9:22.

38. At step 112, the system determines a “condition classification” of the property characteristic, such as, e.g., a “good” condition or a “bad” condition. EX1001, 3:27–28, 10:49–57. To do so, “condition profiles 114” (e.g., machine learning analysis algorithms trained to classify conditions of property characteristics) are applied to the extracted image features from step 106. EX1001, 10:60–11:20.

39. The ’491 patent does not purport to have invented any new machine learning or image analysis techniques. Rather, the classification steps 110 and 112 described by the ’491 patent apply existing machine learning and image analysis techniques. For example, the ’491 patent identifies convolutional neural network (CNN) technology as one example, which the ’491 patent describes as “a well-established and popular machine-learning methodology.” EX1001, 8:50–9:8.

B. Prosecution History

40. I understand that the ’491 patent issued from U.S. Application No. 17/155,633 (“the ’633 application”) without any substantive prior art examination.

The ’491 patent is a continuation of U.S. Application No. 16/868,113 (now U.S. Patent No. 11,195,058), which is a continuation of U.S. Application No. 16/733,888 (now U.S. Patent No. 10,650,285), which is a continuation-in-part of U.S. Application No. 15/714,376 (now U.S. Patent No. 10,529,029 (“the ’029 patent”). EX1002, 66. The prosecution history reveals that the claims never underwent meaningful analysis against the prior art that renders them obvious. *See generally* EX1002.

41. During prosecution, the ’633 application received a double-patenting rejection over the ’029 patent. EX1002, 127–133. In response, the applicant filed a terminal disclaimer, after which the Examiner issued a Notice of Allowance. EX1002, 165–166, 193–199.[1] This limited prior art consideration occurred despite the robust prior art landscape that existed at the time of filing.

---

[1] The Notice of Allowance references the reasons for allowance in the 15/714,376 application (U.S. Patent 10,529,029), where the Examiner distinguished three references (Swanson, Downey, and Cuttell) without having issued a single rejection against the allowed claims of the ’029 patent and without the detailed analysis typically conducted during examination. The “Reasons for Allowance” for the ’029 patent also discusses an *Alice* rejection, but again there were no substantive rejections made of the claims during prosecution of the ’029 patent.

C. Level of Ordinary Skill in the Art

42. In my opinion, a POSA at the time of the ’491 patent would have had a bachelor’s degree in Computer or Electrical Engineering or equivalent discipline, and approximately two years’ experience in image processing and machine learning. Additional education could substitute for professional experience.

43. I am well-qualified to determine the level of ordinary skill in the art. I am very familiar with the technology of the ’491 patent during the September 2016 timeframe. As mentioned above, by 2016, I had obtained a Bachelors and Masters (in Electrical Engineering) from the University of the Witwatersrand, Johannesburg and a Ph. D. (in Engineering Science) from Balliol College, Oxford and had acquired extensive experience relating to image processing and computer vision.

D. Claim Construction

44. I have been informed that the challenged claims of the ’491 patent are to be given their ordinary and customary meaning as understood by one of ordinary skill in the art at the time of the invention in light of the specification and the prosecution history pertaining to the patent. Based on my review of the challenged claims, it is my opinion that none of the claim terms require explicit constructions. Accordingly, all claim terms should receive their ordinary and customary meaning, in the context of the ’491 patent specification.

VII. STATE OF THE ART

A. Aerial imagery analysis and property assessment techniques have been used for decades, well before the ’491 patent.

45. Computerized analysis of aerial imagery for property assessment was well-established before the September 23, 2016 priority date of the ’491 patent. By the 1990s, Geographic Information Systems (GIS) were commonly used by government agencies and private companies to analyze aerial imagery for property assessment and urban planning. For example, the first version of the TIGER (Topologically Integrated Geographic Encoding and Referencing) system was released by the US Census Bureau in 1989. EX1015. TIGER boundary files “contain information about many different features such as roads, rivers, lakes, political boundaries (state, county), and census boundaries (Block, Block Group, Census Tract).” EX1015, 1. Commercial aerial imagery databases also became widely accessible via digital platforms by the early 2000s. For example, the technology behind Google Earth was launched in 2001, providing high-resolution aerial imagery to the public. EX1016.

46. Property assessors routinely accessed these types of databases to evaluate property characteristics without requiring physical site visits well before 2016. For instance, Takahashi described an “image recognition method” for analyzing aerial photographs and satellite images. EX1010, 1:14–22. In particular, Takahashi discloses an “image recognition apparatus for use in an image

recognition system, for extracting from a color image the shapes of objects which are to be recognized." EX1010, 1:14–22. Takahashi also discloses its image processing techniques are "to be applied to an image such as satellite images or aerial photograph, for example to accurately and reliably extract the shapes of specific objects such as roads, buildings, etc., from the image contents." EX1010, 2:61–63. In addition, Takahashi discloses analysis of specific building elements from aerial photographs, such as "extracting only the shape of the building roof" EX1010, 31:50.

47. Techniques for analyzing property structures and conditions from aerial imagery were also established and well-known before the September 23, 2016 priority date of the '491 patent. For example, Collins discloses systems for "automated property insurance inspection" using "video images of property," which "may be performed on the exterior or interior of the property, and may be performed periodically, on-demand or continuously." EX1011, Abstract. Jagannathan similarly describes property assessment techniques that analyze aerial images "to determine an indicator of condition risk of a roof of … a building." EX1012, ¶¶5, 35.

B. Machine learning and computer vision techniques for image analysis were well known before the '491 patent

48. Artificial neural networks were a class of machine learning techniques that was first developed in the 1940s and 1950s. Weideman, filed in January 1989,

describes early neural network architectures, such as "a neural-simulating system for an image processing system [that] includes a plurality of networks arranged in a plurality of layers" for image processing applications. EX1013, Abstract. As machine learning technology developed, convolutional neural networks (CNNs) made an appearance in the 1980s and became widely adopted for image recognition by the early 2000s. For example, the AlexNet CNN, which is referenced in the '491 patent as an example of a CNN that can be used in the machine learning analysis of the disclosed embodiments (*see* EX1001, 8:50–55, 60–66), became regarded as the first widely recognized application for image classification tasks after it won the ImageNet 2012 challenge in September 30, 2012, making CNN-based approaches standard in computer vision applications well before the 2016 priority date. EX1017. The Network in Network (NIN) model, also referenced in the '491 patent as an example machine learning classifier that can be used in the disclosed embodiments (*see* EX1001, 1:65–2:3, 8:66–9:2), was published in December 2013 by Lin et al. EX1018. The NIN architecture was publicly available and well-known to practitioners in machine learning for image classification tasks, offering improvements over traditional CNN approaches.

49. Automated machine learning-based extraction of features such as edges, shapes, textures, and color distributions from digital images for property analysis was also routine practice before 2016. Computer vision libraries and

frameworks such as OpenCV (which first released in 2000) and scikit-image (which first released in 2009) provided standard tools for preprocessing, feature extraction, and classification of image contents. EX1019, EX1020. Using these accessible tools, machine learning classifiers were commonly applied to extract features from images for property element identification. For instance, Ciarcia describes "automated roof detection" techniques on "aerial images of [a] building" that use "an artificial intelligence ('AI') system" that trains "on multiple historical images for which the roof and non-roof areas are known." EX1014, 2:66–3:36, 4:33–43, 7:16–26.

VIII. OVERVIEW OF THE APPLIED REFERENCES

A. Gross

50. Gross is directed to "tools, methods, and systems which assess the condition of living structures and other appurtenant real property features." EX1005, ¶2. Like the '491 patent, Gross teaches that these tools are useful in the area of insurance, among others. EX1005, ¶¶2, 5, 186, 237–238. In my opinion, Gross is analogous art to the '491 patent, because it is in the same field of endeavor and reasonably pertinent to the problems of the '491 patent.

51. Gross's property assessment computing system involves training a "Classifier Engine" and using it to assess and rate different properties. EX1005, ¶57. To train the classifier engine, Gross's system first collects reference images of

"building elements" (e.g., roof, tree, etc.) along with their "types" (e.g., pitched roof, double pitched roof, shingle roof, etc.). EX1005, ¶¶57–58. Gross's system then defines and associates the building elements within each reference image with a corresponding "condition," such as "new," "damaged," "missing," and more. EX1005, ¶¶60–70. Gross's system trains the Classifier Engine "using a variety of techniques, including statistical processing, neural networks, etc., which can be used to detect and extract features." EX1005, ¶¶77–78.

52. After training the Classifier Engine, Gross's system uses it to assess new images of properties. EX1005, ¶¶57, 123. These images of a property to be assessed are first retrieved from a database. EX1005, ¶121. The property images may also be obtained via a camera of a mobile device. EX1005, ¶¶108, 109. Gross's Figure 5, reproduced below, illustrates one embodiment of the assessment process.

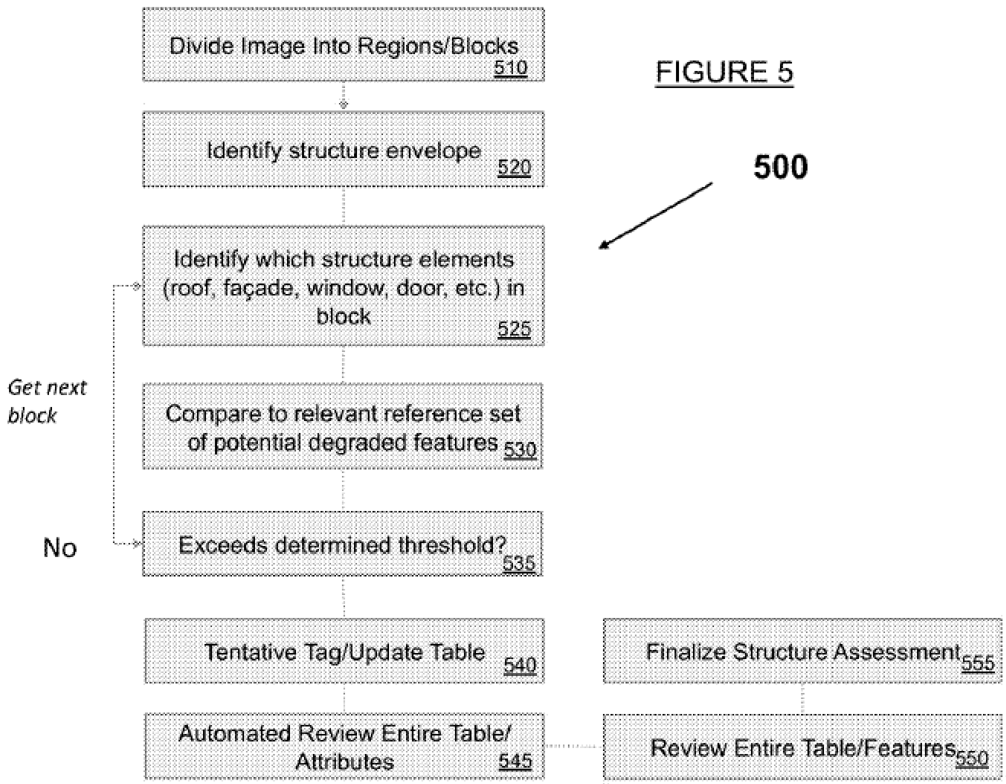

EX1005, FIG. 5.

53. Once the image is retrieved, it is divided into distinct regions or "blocks" 701 at step 510. EX1005, ¶121. This is shown, for example, in Figure 7B, reproduced below.

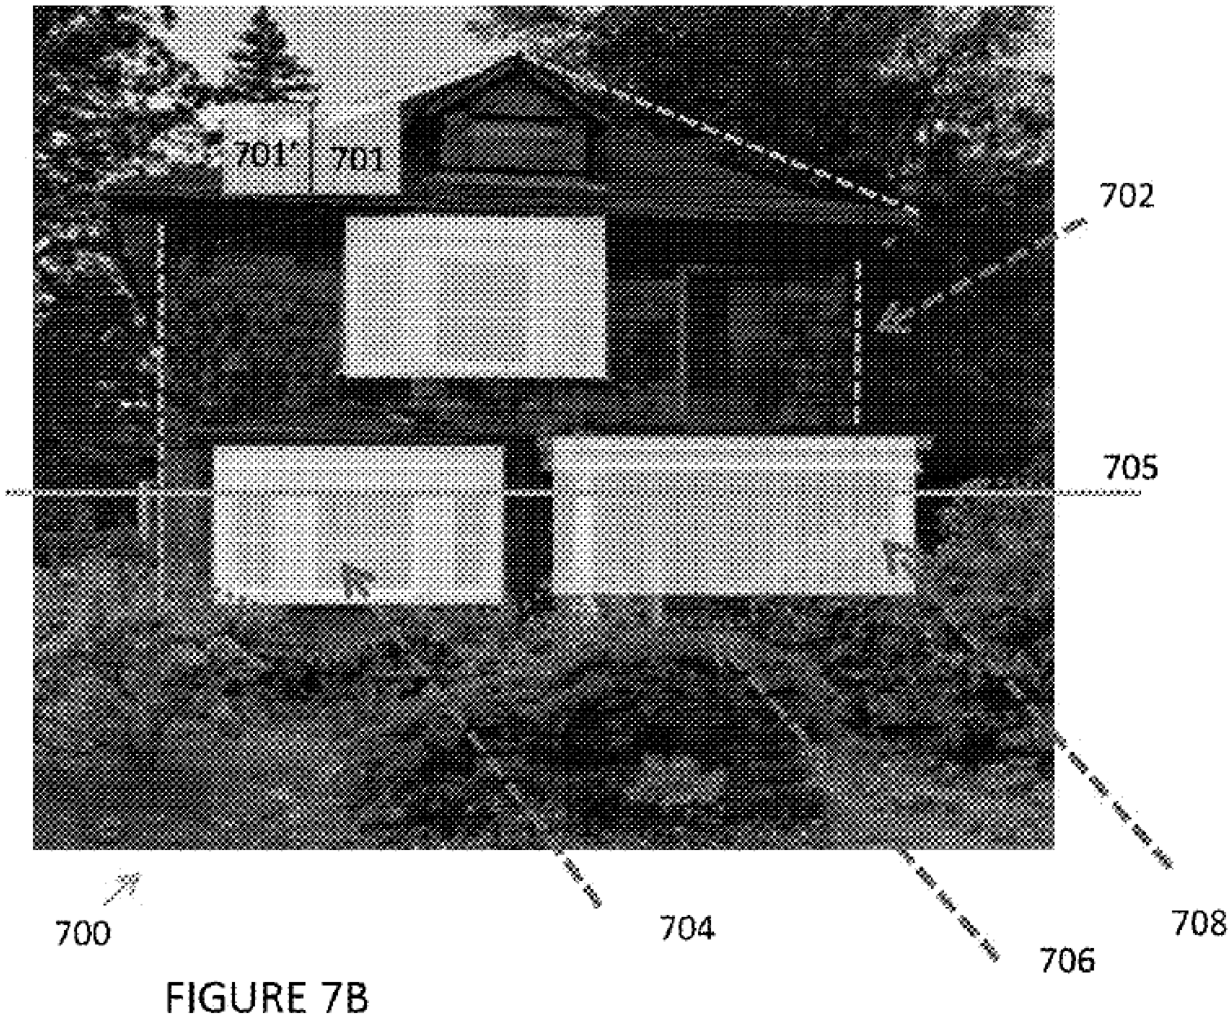

EX1005, FIG. 7B (excerpted).

At step 520 of the method of Figure 5, Gross's system identifies a "building envelope 702" (which as seen in Figure 7B is essentially an outline of the building) by detecting edges of structures in the image. EX1005, ¶122, FIG. 7B. Such "boundaries" (e.g., edges) of other objects in the property (e.g., landscaping, trees, etc.) can also be determined. EX1005, ¶122. The image blocks include "structural elements" (e.g., a roof, tree, etc.), which are determined (alongside a "type" of the structural element (e.g., pitched roof, shingle roof, etc.) at step 525. EX1005, ¶123. As one example, "the first block 701 may be coded to denote that an attribute {roof} with type {pitched} and {shingle} is contained" therein. EX1005, ¶123. As Gross explains, and similar to the '491 patent, identifying the structural element and its classification "can be performed in any number of ways by Classifier Engine…including by pattern recognition techniques and statistical image processing using models." EX1005, ¶123.

54. Again similar to the '491 patent, Gross's system then employs Classifier Engine to determine a "condition" of the determined structural element at steps 530 and 535, such as a "good" condition or a "poor" condition. EX1005, ¶¶126–127, 141. The condition classification is then stored into memory at step 540. EX1005, ¶127. This assessment process is repeated for one or more extracted image blocks for the image and the final assessment data is reviewed and finalized at steps 545, 550, and 555. EX1005, ¶¶129–133.

B. Davis

55. Davis discloses a computer-implemented method for determining property features from aerial images and analyzing the property features to determine an estimated risk of damage to the property. EX1006, Abstract. In my opinion, Davis is analogous art to the '491 patent, because it is in the same field of endeavor and reasonably pertinent to the problems of the '491 patent. Davis's

system "allow[s] an insurance provider to more accurately and efficiently underwrite risk for a home insurance policy" and "reduc[e] the number, frequency and/or cost of … risk mitigation requests." EX1006, 1:64–2:3. A policyholder user (e.g. "the owner of [a] property") may first request "a risk mitigation program." EX1006, 7:53–60. In response, Davis's system, like the '491 patent, obtains and processes an aerial image of the property associated with the request. EX1006, 7:53–60. Davis's system then processes the aerial image to "determine one or more features of the property," using techniques such as feature extraction. EX1006, 18:27–34.

56. Davis's system also performs "object detection/identification," for example, to "identify a particular tree" of a property or "a type of roof" on a house on a property from the aerial image. EX1006, 18:36–39, 48–52; 9:24–29. Subsequently, Davis's system employs "other image analysis techniques" (e.g., image classification) to determine classifications of the detected object, such as "a size of at least a portion of the tree[,] a species of the tree, an orientation of a portion of the tree with respect to the structure or object," "distances of tree trunks and/or branches from homes or other structures located on the properties," or "a type of roof on a structure." EX1006, 18:27–39; 9:22–29, 3:38–46. Based on these determinations, Davis's system then analyzes a risk of damage to a structure on the property (e.g., a house). EX1006, 19:6–16. Davis's system then generates "a risk

output … based on the risk of damage" and transmits "[i]nformation indicative of the generated risk output" to "an employee (e.g., agent) of an insurance provider … and/or to the policyholder." EX1006, 20:40–42, 59–64.

C. Furukawa

57. Furukawa describes techniques for processing image data (e.g., satellite image data) and correcting errors using map data that, for example, contain "information of external shape of buildings." EX1007, Abstract. In my opinion, Furukawa is analogous art to the '491 patent, because it is in the same field of endeavor and reasonably pertinent to the problems of the '491 patent. Furukawa's system "extracts shape information of a building from images by image processing means" (e.g., via "building external shape information detecting means 6") or "obtains the same from … existing map data." EX1007, 6:43–46, 32:65–33:7. Furukawa's system then compares the extracted "building external shape information … [with] the map data" to determine whether the map data and image data match. EX1007, 35:17–31, 21:66–22:22. If the map data and image data do not match, Furukawa's system then modifies the map data to obtain an alternate, updated raster image. *See, e.g.*, EX1007, 35:50–57, 30:56–31:53.

IX. GROUND 1: GROSS RENDERS OBVIOUS CLAIMS 1, 6–11, 16–20

58. In my opinion, claims 1, 6–11, and 16–20 would have been obvious in view of Gross.

A. A POSA would have found it obvious Gross's techniques were intended to be combined.

59. As I discuss above, Gross discloses training a Classifier Engine 150 and using it to assess and rate properties. EX1005, ¶57, FIG. 2. Gross describes specific details on how its Classifier Engine 150 is trained. *See* EX1005, ¶¶77–93. For example, Gross describes how "a set of reference images" is first "collected" and "tagged … to identify each building attribute, an associated condition, a location in the image, etc." EX1005, ¶58. Gross provides a list of exemplary structure/building attributes, attribute types, and conditions/impairments that Classifier Engine 150 may be trained to identify from an image, including roof {new, missing/damage tile, shingle metal, holes/cracks, unevenness, brightness (moss mold)}; façade (new, missing, shingles/siding, breaks, holes, discontinuities, discoloration, warping, paint bubbles/blistering/peeling—aging, weathering, water damage}; and other exemplary attributes/types/conditions. EX1005, ¶¶60–70.

60. Gross's training process then proceeds to analyze the identified attributes "to break them down and reduce them into (or extract) smaller distinct features that can be more readily detected." EX1005, ¶77. Gross describes how "[t]he training results in extraction and optimization of specific features most appropriate for a particular element, so as to minimize classification errors for the attribute in question." EX1005, ¶78.

61. Gross further describes that, after it has been trained, Classifier

Engine 150 may be used to identify attributes (and their types and conditions) contained in an image of a property to be assessed. *See* EX1005, ¶¶120–129, FIG. 5. Gross describes that "the identification of such attributes (and their types) can be performed in any number of ways by Classifier Engine 150," "including by pattern recognition techniques and statistical image processing using models, templates, etc." EX1005, ¶123. For example, Classifier Engine 150 may be trained to assess whether the following exemplary structure attributes, attribute types, and conditions/impairments are contained in an image: roof {new, missing/damage tile, shingle metal, holes/cracks, unevenness, brightness (moss mold)}; façade (new, missing, shingles/siding, breaks, holes, discontinuities, discoloration, warping, paint bubbles/blistering/peeling—aging, weathering, water damage}; and other exemplary attributes/types/conditions. EX1005, ¶¶60–70.

62. A POSA reading Gross would have understood that the image processing and attribute classification techniques used to train Classifier Engine (e.g. Figure 2) were intended to also be utilized in its building attribute/condition assessment process (e.g., Figure 5). Indeed, that is the whole point of Gross—to use its trained classifier to determine building attributes and condition assessments. Both processes involve analyzing property images to identify attributes (such as roof elements, window elements, etc.), their types (e.g. for a roof – shingle, metal, tile), and conditions (e.g., for a roof – missing/damaged tile, holes/cracks). A

POSA would have recognized that Classifier Engine 150 would employ the same, or very similar, techniques as those used during training in order to process images to identify and classify attributes contained therein. *See, e.g.*, EX1005, ¶57. For example, Gross describes how "conventional steps that would be included … have been omitted." EX1005, ¶245. A POSA would have understood such implementation would amount to combining (as intended) Gross's disclosure of the image processing elements used to train the Classifier Engine with its disclosure of image processing used in its building attribute/condition assessment process according to known methods to yield predictable results with a reasonable expectation of success because both processes employ similar image processing techniques—i.e., image processing techniques used to train the Classifier to identify attributes/conditions would in fact also have been used when performing a property assessment.

63. Therefore, it is my opinion that a POSA would have found it obvious to combine different embodiments of Gross, would have had a reasonable expectation of success in doing so, and that the results of the combination would have been predictable (i.e., using the trained Classifier Engine for attribute classification and condition assessment).

B. Independent Claim 1

1. [1.P]: A system for automatically assessing features of a property location comprising a structure, the system comprising:

64. I understand that in certain circumstances the preamble of a claim may be limiting. To the extent the preamble is determined to be limiting, it is my opinion that Gross discloses the preamble.

65. It should first be noted that the word *feature* can have multiple meanings, depending on the context, which can affect how the claim phrase *features of a property location* is interpreted. In computer vision contexts, the word *feature* (or "image feature") typically refers to pieces of information that can be derived from the pixel content of an image (as opposed to descriptive aspects and qualities for the image such as metadata). In standard contexts, the word *feature* can more broadly refer to an aspect, quality, or facet of something. When referring to the claimed *features of a property location*, *property features*, and *respective feature from each image*, the '491 patent uses the word *feature* in the second sense, describing "property features" as physical aspects, qualities, or facets of a property (e.g., "rooftop, swimming pool, chimney, and sky lights"). EX1001, 8:36–38, *see also* EX1001, 2:42–47 (explaining how "property features" include "structural features, manmade features included within a vicinity of a structure … and/or natural features included within the vicinity of the structure).

66. Gross discloses “training [a] Classifier Engine … and using it to assess and rate” properties. EX1005, ¶57. In assessing and rating properties, Gross explains that “an identification is performed” by Classifier Engine “to determine structural elements presented in a first block” of an image, “including such attributes as roof elements, window elements, façade elements, etc.” EX1005, ¶123. These attributes are *features of a property location comprising a structure*, as claimed, because they are physical parts of a structure, such as a house or building.

67. Gross also discloses that Classifier Engine identifies conditions associated with a particular attribute (i.e., *feature of a property location*). EX1005, ¶60. For example, the Classifier Engine may automatically assess the condition of a roof (*feature of a property location*) as “new,” “missing/damage tile,” and/or “holes/cracks.” EX1005, ¶61. In some instances, “severity weightings (i.e., a scaling factor of any convenient range (1-10, or the like) for the identified condition” may also be assigned by Classifier Engine. EX1005, ¶70. The Classifier Engine may also compute a “condition score … for the detected attribute.” EX1005, ¶126. Gross further discloses that “with sufficient training an automated classifier could participate in the process, if not perform the entire process of classifying structural features,” i.e., *automatically*. EX1005, ¶145. Gross also refers to its Classifier Engine in various sections of its disclosure as an “automated

classifier." EX1005, ¶¶145, 174, 176. Therefore, it is my opinion that Gross discloses *automatically assessing feature of a property location*.

2. [1.1]: processing circuitry; and a non-transitory computer-readable medium having instructions stored thereon, wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to

68. In my opinion, Gross teaches element [1.1]. Gross discloses "the critical steps identified in this process can be implemented into executable software routines and modules using any number of ways by skilled artisans." EX1005, ¶161. These "executable software routines" are executed via "*processing circuitry*." Gross further discloses that its "applications can be embodied in tangible, machine readable form for causing a computing system to execute appropriate operations in accordance" with the teachings. EX1005, ¶56. Therefore, it is my opinion that Gross teaches element [1.1].

3. [1.2]: access a plurality of images of the property location including the structure;

69. In my opinion, Gross teaches element [1.2]. In particular, Gross describes "aerial drones, satellite, balloon, and similar technology" can be used that "capture[s] structure image data from a variety of perspectives" (i.e., *a plurality of images of the property location including the structure*) and "can facilitate building out a comprehensive image database." EX1005, ¶164. Gross explains further how obtaining images "at different angles, elevations, and profiles

"helps increase building coverage and features currency/accuracy." EX1005, ¶148. The building image data is stored in a "structure image database," which Gross teaches contains primary (e.g., original) images for a structure, along with "structure sub-images which are based on automatically dividing the original image into separate blocks, or separate areas corresponding to distinct building elements." EX1005, ¶165.

70. Gross further describes "a data acquisition process" used by a classifier (e.g., Classifier Engine 150) and that "the purpose "of this specialized computing module is to acquire appropriate image data for structures within a target area." EX1005, ¶161. Classifier engine retrieves (*accesses*) an image 700 from a database to begin the building attribute/condition assessment process and the image data that is retrieved is for the target address. EX1005, ¶¶120–21, 164. In light of Gross's teachings of a plurality of images for a property location (as discussed above), a POSA would have understood or found obvious that Gross's "comprehensive image database" (containing different types of images for a property location) teaches, or renders obvious, that Classifier Engine *access[es] a plurality of images of the property that includes the structure* to increase classification "accuracy" because otherwise there would be little reason for Gross to store multiple images of the same building structure. For example, an image that would "increase building coverage" of the roof of a property and therefore increase

the "accuracy" of the roof condition classification would be a top-down image of the property that exposes a higher amount of rooftop surface area, where a different elevation view would increase the coverage and improve accuracy of a condition of a façade.

71. Therefore, it is my opinion that Gross teaches this element.

4. [1.3]: apply boundary information to one or more images of the plurality of images to isolate the property location or the structure thereon within each respective image; and

72. In my opinion, Gross teaches element [1.3]. Gross describes identifying a "building envelope 702" by image processing adapted to detect edges of structures. EX1005, ¶122. The building envelope functions as *boundary information*, which is applied *to isolate the structure* within the image because it detects the edge of the structure and isolates the structure within the image. I note that other than in independent claims 1 and 11, the '491 patent does not describe applying boundary information to isolate a structure within an image. Figure 7B of Gross (depicted below) illustrates an example of boundary information (building envelope 702), which isolates the structure.

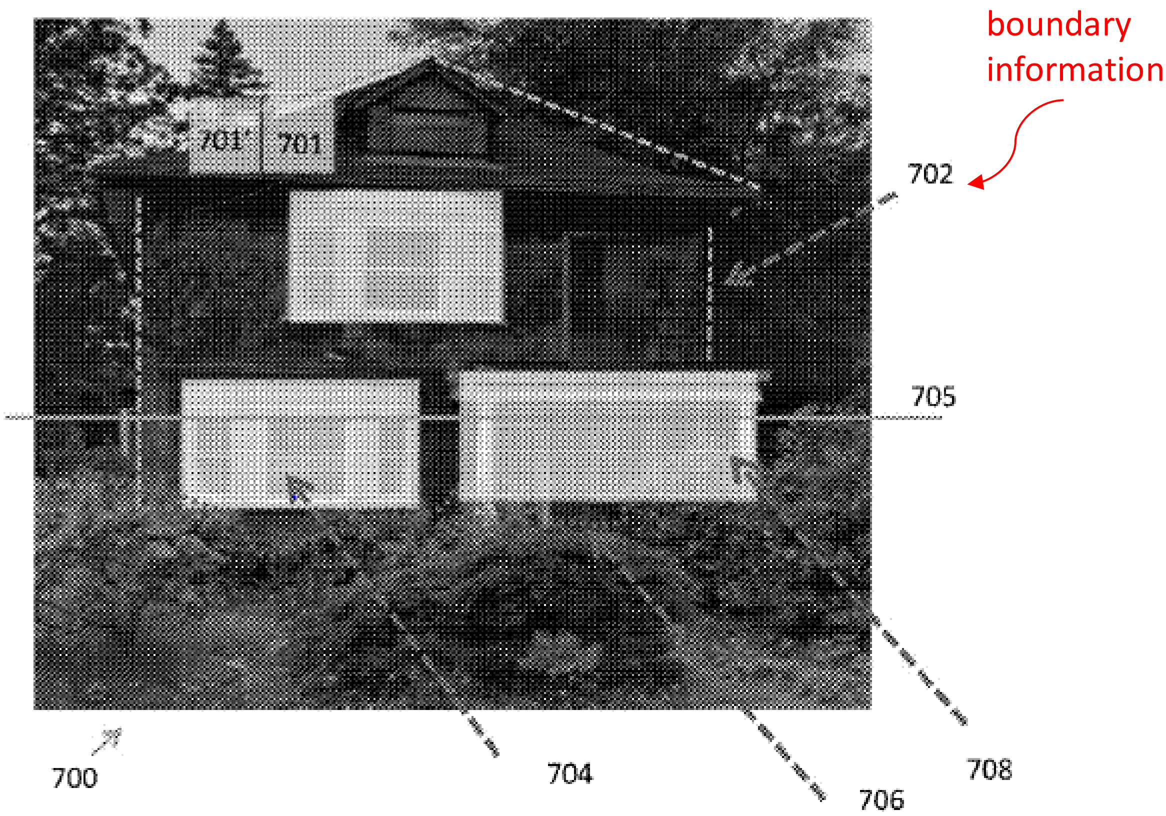

EX1005, FIG. 7B (annotated, excerpted).

73. Therefore, it is my opinion that Gross teaches *apply[ing] boundary information to one or more images of the plurality of images to isolate … the structure thereon*.

74. I understand from counsel that this claim element is satisfied if the boundary information is applied to isolate *either* "the property location" *or* the "structure thereon." However, Gross also teaches *apply[ing] boundary information to isolate the property location* for at least the reasons I described above. In addition, Gross discloses "a data acquisition process 1600 used by a classifier [e.g., Classifier Engine] as it would be implemented on a customized structure

assessment." EX1005, ¶161. Gross explains the purpose of this process "is to acquire appropriate image data for structures within a target area." EX1005, ¶161.

75. Gross describes additional data could be gleaned "from public records which identify the layout of building envelopes on each lot." EX1005, ¶143. Gross further discloses using "[i]nformation identifying a beginning and end of each individual street, road, alley" "from any convenient database or similar source" to retrieve "image data" for a "target address in question." EX1005, ¶¶162, 164. Figure 10 of Gross illustrates an example of such "[i]nformation":

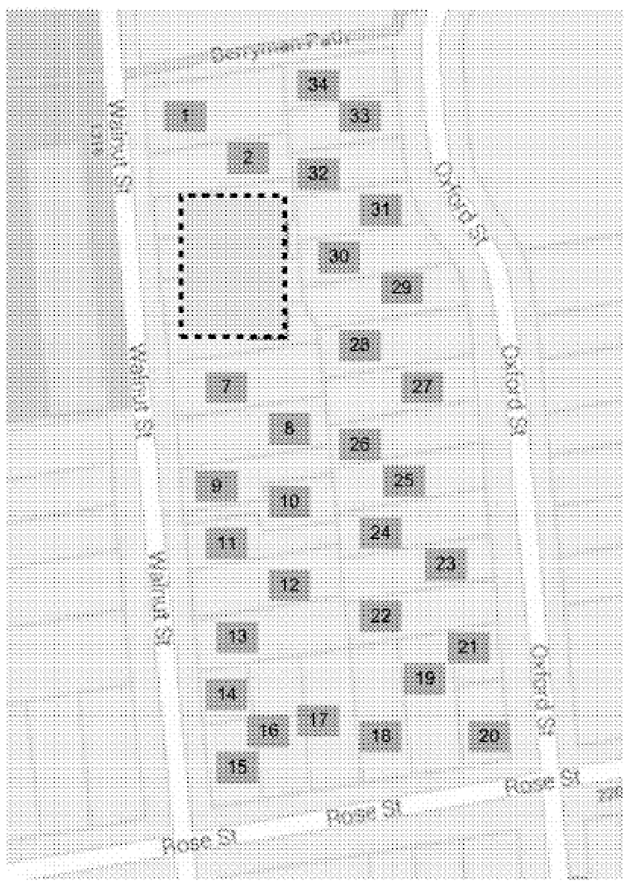

EX1005, FIG. 10.

Both the data in public records that identifies the layout of building envelopes on each lot (i.e., identifies lot locations) and the information identifying beginning and end of streets/roads/alleys constitutes *boundary information* because it is used to isolate the location of the target address from other properties in the area. Therefore, it is my opinion that Gross teaches *applying boundary information to* isolate the location of the "target address in question."

5. [1.4]: classify a condition of one or more property features of the property location, wherein classifying comprises, for each feature of the one or more features,

76. In my opinion, Gross teaches element [1.4]. Gross discloses classifying "conditions/impairments … associated with a particular building [structure] attribute." EX1005, ¶60. The structure attributes are *property features of the property location* because, like the '491 patent, they are characteristics located on the property (*see*, *e.g.*, EX1001, 2:39–45, 8:36–38, and my discussion above for element [1.P]). Gross describes examples of structure attributes (i.e., *property features*) and related *conditions*, including: roof (new, missing/damage tile, holes/cracks); façade (new, missing, holes); and support (leaning/damaged columns). EX1005, ¶¶60–69. The examples provided are just a partial list and other individual structure attribute/condition pairings can be identified by Classifier Engine. EX1005, ¶70.

77. For the reasons I provide below for elements [1.5]–[1.7], it is my

opinion that Gross teaches *wherein classifying comprises, for each feature of the one or more features*, the limitations set forth in [1.5]–[1.7] are performed for at least one feature, which I understand from counsel only requires elements [1.5]–[1.7] to be executed a single time because one feature is "one or more".

6. [1.5]: extracting the respective feature from each image of at least one image of the one or more images,

78. In my opinion, Gross teaches or renders obvious element [1.5]. I understand from counsel that "*the respective feature*" in this limitation is a reference to "*each feature of the one or more* [*property*] *features*" from element [1.4]. I agree with this reading of the claim language. As I discuss above in element [1.P], the '491 patent refers to the claimed *respective* [*property*] *feature* (that is extracted from the one or more images) as being physical aspects, qualities, or facets of a property, such as a "rooftop, swimming pool, chimney, and sky lights" EX1001, 8:36–38. *See also* EX1001, 8:36–38. The '491 patent describes that "extracted features may be analyzed to determine one or more property characteristics of that feature" and how "[i]n one example, a swimming pool may be identified as the property feature of a pool … then characterized as a particular shape[,] type[,] or size." EX1001, 8:40–50.

79. Similarly, Gross discloses objects can be analyzed in image patterns, using a variety of techniques (including statistical processing and neural networks) that can be used to detect and *extract* features of attributes. EX1005, ¶77. During

this process, an image is obtained and "an identification is performed to determine structural elements presented in a first block 701, including such attributes as roof elements, window elements, façade elements, etc." EX1005, ¶¶121, 123. Gross also explains that "the identification of such attributes … can be performed in any number of ways by Classifier Engine[,] including by pattern recognition techniques and statistical image processing using models, templates, etc." EX1005, ¶123. A POSA would have understood or found obvious that, similar to the '491 patent, this identification *extracts the respective feature* from the image because image features representing a specific structural element (i.e., *property features* such as a roof, window, etc.) are extracted from the image (using the disclosed image analysis techniques, such as neural networks).

7. [1.6]: applying a set of machine learning analysis algorithms to at least a portion of each image of the at least one image to determine a set of characteristics of the respective property feature, wherein each algorithm of the first set of machine learning analysis algorithms is trained to identify one or more characteristics of the set of characteristics, and

80. In my opinion, Gross teaches, or renders obvious, element [1.6]. Gross discloses that, after it has been trained, Classifier Engine 150 (which includes the claimed "*machine learning analysis algorithms*") identifies "structural elements" (i.e., *property features*) presented in an image. EX1005, ¶123. The specific types of the attributes (i.e., *characteristics of the property feature*) are also captured "(i.e., façade (shingle, siding, brick), … roof {pitched, double pitched,

hipped, flat, metal, tile, shingle, slate, parapet, dormer, mansard, fascia, brackets, eaves, pent, pediment, etc.})." EX1005, ¶57. For instance, Gross describes how a first block of the image "may be coded [by Classifier Engine] to denote that an attribute {roof} with type {pitched} and {shingle} is contained" in the image. EX1005, ¶123. To obtain the attribute type classifications, Gross explains that Classifier Engine 150 "employs image processing" and the "identification of such attributes (and their types) can be performed in any number of ways by Classifier Engine 150 after it has been trained." EX1005, ¶¶120, 123. Therefore, it is my opinion that Gross teaches applying its Classifier Engine to at least a portion of an image *to determine a set of characteristics* (e.g., "pitched") *of the respective property feature* (e.g., roof).

81. Gross's Classifier Engine includes *a set of machine learning analysis algorithms*, *wherein each algorithm is trained to identify one or more characteristics of the set of characteristics*. In particular, Gross discloses its Classifier Engine is trained with reference images and that objects in the image are analyzed using neural networks to detect and extract features. EX1005, ¶77. For example, Classifier Engine is trained to identify a roof feature and a set of characteristics associated with the roof ("pitched, double pitched, hipped, flat, metal, tile, shingle, slate, parapet, dormer, mansard"). EX1005, ¶57. Gross further describes that "[f]or better accuracy it may be useful to employ multiple classifiers

trained with different algorithms." EX1005, ¶80. In light of Gross's teachings, a POSA would, therefore, have understood, or found obvious, Classifier Engine includes *a set of machine learning analysis algorithms* (and a POSA would have understood a set could be a set of one) that is *trained to identify one or more characteristics of the set of characteristics* (e.g., trained to identify a roof is pitched, double pitched, hipped, etc.). EX1005, ¶¶77, 123.

8. [1.7]: applying another machine learning analysis algorithm to at least a portion of each image of the at least one image to classify a condition of the respective property feature, wherein the condition is classified as a respective classification of a set of possible classifications, and the another machine learning analysis algorithm is trained to identify one or more conditions of property locations and/or structures thereon having the set of characteristics of the respective property feature.

82. In my opinion, Gross teaches, or at least renders obvious, element [1.7]. Gross discloses identifying and classifying "conditions/impairments … associated with a particular building attribute" (i.e., *classifies a condition of the respective property feature*). EX1005, ¶60. Gross describes examples of structure attributes (i.e., *property features*) and related *condition* classification, including: roof (new, missing/damage tile, holes/cracks); façade (new, missing, holes); and support (leaning/damaged columns). EX1005, ¶¶60–69. Gross explains that the examples provided are "just a partial list" and other individual structure attribute/condition pairings can be identified by Classifier Engine. EX1005, ¶70.

83. After the Classifier Engine is trained, it can identify indicators in attributes "of an irregularity, defect, damage, or the like" and compute a condition score for the detected attribute. EX1005, ¶¶123, 126. Therefore, it is my opinion that Gross teaches applying its Classifier Engine (i.e., *another machine learning analysis algorithm* for reasons explained below) to at least a portion of an image *to classify a condition of the respective property feature* (e.g., roof has a "missing tile"), *wherein the condition is classified as a respective classification of a set of possible classifications* (e.g., new, missing/damage tile, holes/cracks).

84. Additionally, for the reasons I discuss above for element [1.6], Gross teaches that its Classifier Engine includes *a set of machine learning analysis algorithms*. EX1005, ¶77. And although Gross does not explicitly disclose that its conditions of property conditions are classified via *another* machine learning algorithm than that used to identify attributes, Gross does explicitly disclose that "it may be useful to employ multiple classifiers … to identify the attributes and classify the structure." EX1005, ¶80. Gross further discloses examples of using a classifier (*machine learning analysis algorithm*) to identify conditions, such as "damaged/exposed roof elements" and "broken—façade," EX1005, ¶¶85, 91, which correspond to the attributes (*property feature*) that I discuss above for Element [1.4], such as a "roof" or "façade." Gross also discloses that different classifiers (i.e., *another machine learning analysis algorithm*) could be customized

for different purposes by claiming, "an electronic classifier ***customized*** for a vendor … as corresponding to said individual building ***elements***, ***types*** and/or ***associated conditions***." EX1005, claim 26 (emphasis added). Additionally, Gross explains that "each attribute (and related condition qualifier) may be reduced in dimensionality and characterized by a distinct set of pixels, shapes, proportions, curves, textures, or other mathematically derivable feature from the image," which teaches that different types of features (e.g., attributes and conditions) require different image processing approaches. EX1005, ¶78. In light of these teachings, a POSA would have understood that Gross teaches, or renders obvious, the use of *another machine learning algorithm*, as claimed, to classify a condition of the property feature because customized algorithms are more accurate or more appropriate for different vendors.

85. A POSA would have found it obvious to separate the classifiers and would have been motivated to employ another machine learning analysis algorithm to identify a condition classification because Gross explicitly suggests that different classifiers (i.e., different machine learning analysis algorithms) can be used and that classifiers can be customized for different purposes. EX1005, ¶80, claim 26. Moreover, a POSA would recognize that the reason Gross suggests using different classifiers is so that they would then be trained with specific data that optimize and improve the accuracy of their classifications, such as in this case the

classification of a condition of property features. A POSA would have been motivated to employ another machine learning analysis algorithm to identify a condition classification because Gross explicitly suggests that different classifiers (i.e., different machine learning analysis algorithms) can be used and that classifiers can be customized for different purposes. EX1005, ¶80, claim 26. A POSA would have understood that the two classification tasks of identifying what object is present versus determining the condition of that object are fundamentally different in nature and could require analysis of different types of image features. In addition, a POSA would have understood that separating these classification tasks allows for independent optimization of each algorithm's training parameters and training datasets. Doing so would lead to increased overall system accuracy compared to a single algorithm attempting to perform two types of classification simultaneously.

86. For example, a first algorithm may focus on pixels and image features that delineate object boundaries and structural elements to identify *property characteristics* (e.g., detecting edges, corners, and geometric shapes that define a roof structure), while a second algorithm may focus on pixels and image features within the interior regions of identified objects to assess condition (e.g., detecting texture, color, surface artifacts). This differentiation is further supported by Gross's teaching that "[t]he training results in extraction and optimization of specific

features most appropriate for a particular element, so as to minimize classification errors." EX1005, ¶78. A POSA would have readily implemented the predictable variation of Gross's Classifier Engine to use multiple machine learning analysis algorithms—a first set of algorithms used to determine characteristic of the property feature and another machine learning analysis algorithm used to identify conditions of those features/characteristics.

87. Therefore, it is my opinion that Gross teaches, or at least renders obvious, the claimed *another machine learning analysis algorithm* that is *trained to identify one or more conditions of property locations and/or structures thereon* (e.g., trained to identify the condition of a roof, which is a condition of a structure on the property) *having the set of characteristics of the respective property feature* (e.g., a roof with characteristics "tile" and "parapet."). EX1005, ¶¶60, 61, 123, 126.

C. Independent Claim 11

1. [11.P]: A method for automatically assessing risk damage to a property, the method comprising:

88. To the extent the preamble is determined to be limiting, it is my opinion that Gross discloses element [11.P] for the reasons I discuss above for the preamble of claim 1 (element [1.P]). Gross further discloses its assessment system can be used for insurance risk assessments and that correlations can "be developed between property condition ratings, occupancy estimates and number of claims filed, type of claim, severity etc." EX1005, ¶248 In other words, Gross teaches its

methods can *automatically assess risk damage to a property*.

2. [11.1]: accessing, by processing circuitry, a plurality of images of a property location including a structure

89. In my opinion, Gross teaches element [11.1] for the reasons I discuss above for element [1.1] and [1.2].

3. [11.2]: applying, by the processing circuitry, boundary information to one or more images of the plurality of images to isolate the property location or the structure thereon within each respective image; and

90. In my opinion, Gross teaches element [11.2] for the reasons I discuss above for element [1.1] and [1.3].

4. [11.3]: classifying, by the processing circuitry, a condition of one or more property features of the property location, wherein classifying comprises, for each feature of the one or more property features,

91. In my opinion, Gross teaches element [11.3] for the reasons I discuss above for element [1.1] and [1.4].

5. [11.4]: adjusting each image of the at least one image to obtain an adjusted image for analyzing the respective feature,

92. In my opinion, Gross teaches element [11.4]. Gross discloses that "customized image processing operations (rotations, noise clean up, etc.) may also be performed" (i.e., image adjustment operations) as part of an image pre-processing process (i.e., before the image analysis). EX1005, ¶76. The incorporated Jafri reference also describes that "appropriate preprocessing of the

images can compensate for noise and slight variations in orientation, scale, and illumination." EX1008, 41. These image processing operations adjust the image *to obtain an adjusted image for analyzing the respective feature*, because they would change the underlying pixel content of the image.

6. [11.5]: applying a set of machine learning analysis algorithms to at least a portion of each image of the at least one adjusted image to determine a set of characteristics of the respective property feature, wherein each algorithm of the first set of machine learning analysis algorithms is trained to identify one or more characteristics of the set of characteristics, and

93. In my opinion, Gross teaches, or renders obvious, element [11.5] for the reasons I discuss above for element [1.6].

7. [11.6]: applying another machine learning analysis algorithm to at least a portion of each image of the at least one image to classify a condition of the respective property feature, wherein the condition is classified as a respective classification of a set of possible classifications, and the another machine learning analysis algorithm is trained to identify one or more conditions of property locations and/or structures thereon having the set of characteristics of the respective property feature.

94. In my opinion, Gross teaches, or renders obvious, element [11.6] for the reasons I discuss above for element [1.7].

D. Claim 6

1. [6.1] The system of claim 1, wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to determine at least one property characteristic of the structure including one or more of an occupancy type, a number of stories, or a siding type,

95. In my opinion, Gross teaches element [6.1]. Gross describes that further information that can be determined by its Classifier Engine includes "[i]information on the … number of stories" (i.e., *property characteristic of the structure including … a number of stories*). EX1005, ¶146. Gross also describes that its Classifier Engine can determine the type of façade of the structure (i.e., *siding type*), such as shingle, brick, and stone. EX1005, ¶57. And Gross describes its Classifier Engine can identify whether structures are inhabited or not, such as "confirmed vacant," "confirmed occupied," "likely vacant, uncertain, likely occupied, and so on" (i.e., *occupancy type*). EX1005, ¶¶118, 131.

2. [6.2] wherein: determining the at least one property characteristic comprises, for each characteristic of the at least one property characteristic, selecting a given image of the at least one image for determining the respective characteristic, and

96. In my opinion, Gross teaches, or renders obvious, element [6.2]. For the reasons I discuss above for elements [1.2] and [1.5], Gross describes that its attribute/condition assessment process retrieves an image from the database for the target address. Thus, Gross teaches *selecting a given image … for determining the respective characteristic* through its description of selecting an image for the target

address. EX1005, ¶164.

97. Gross further describes images can be obtained "to increase building coverage and feature currency/accuracy." EX1005, ¶148. And Gross describes that "[w]hile some images may allow only for one/two dimensional measurements … additional data could be gleaned from other perspectives obtained by other image capturing systems" (i.e., other images). EX1005, ¶143. In other words, Gross discloses that multiple images of the same structure exist, and that different data would be "gleaned" from each of these different images. EX1005, ¶143. And although Gross doesn't explicitly state how its system would know which data would be gleaned from which image, a POSA would have understood or found obvious that regardless of how Gross's system made this determination, Gross's system would have some mechanism for determining which specific images to analyze among its various "perspectives" of a structure to "glean" the aforementioned "additional data." EX1005, ¶143. In this manner, a POSA would have understood or found obvious that Gross's comprehensive image database of images from various perspectives would allow the classifier to use certain, preferred types of images for certain property classification tasks. EX1005, ¶164. For example, an image that would "increase building coverage" of the roof of a property and increase the "accuracy" of the roof condition classification could be a top-down image of the property that exposes a higher amount of rooftop surface

area (e.g., as opposed to an image of the property at street elevation). Similarly, a POSA would understand Classifier Engine would select an alternative image than a top-down image of the property to increase the accuracy when classifying the "siding type" or "number of stories" of a property characteristic, such as an image of the property at street elevation.

3. [6.3] applying a machine learning analysis algorithm trained to identify the respective property characteristic to the given image to determine the respective characteristic.

98. In my opinion, Gross teaches element [6.3]. For the reasons I discuss above in element [1.6], Gross teaches applying a machine learning analysis algorithm trained to determine a set of characteristics of the property feature. Gross similarly describes its Classifier Engine (*machine learning analysis algorithm*) is trained to identify *property characteristics*, such as *occupancy type*, *number of stories*, and façade (*siding type*) (*see* element [6.1]). *See* EX1005, ¶¶57, 63, 118, 146.

E. Claim 7

99. Claim 7 depends from claim 6 and further recites, "*wherein the given image corresponding to a first property characteristic of the at least one property characteristic is a terrestrial image*."

100. In my opinion, Gross teaches, or renders obvious, this claim. Gross discloses that images "at different angles, elevations and profiles" help "increase

building coverage and feature currency/accuracy.” EX1005, ¶148. Gross also illustrates several *terrestrial images* of properties, such as in Figure 7B (reproduced below).

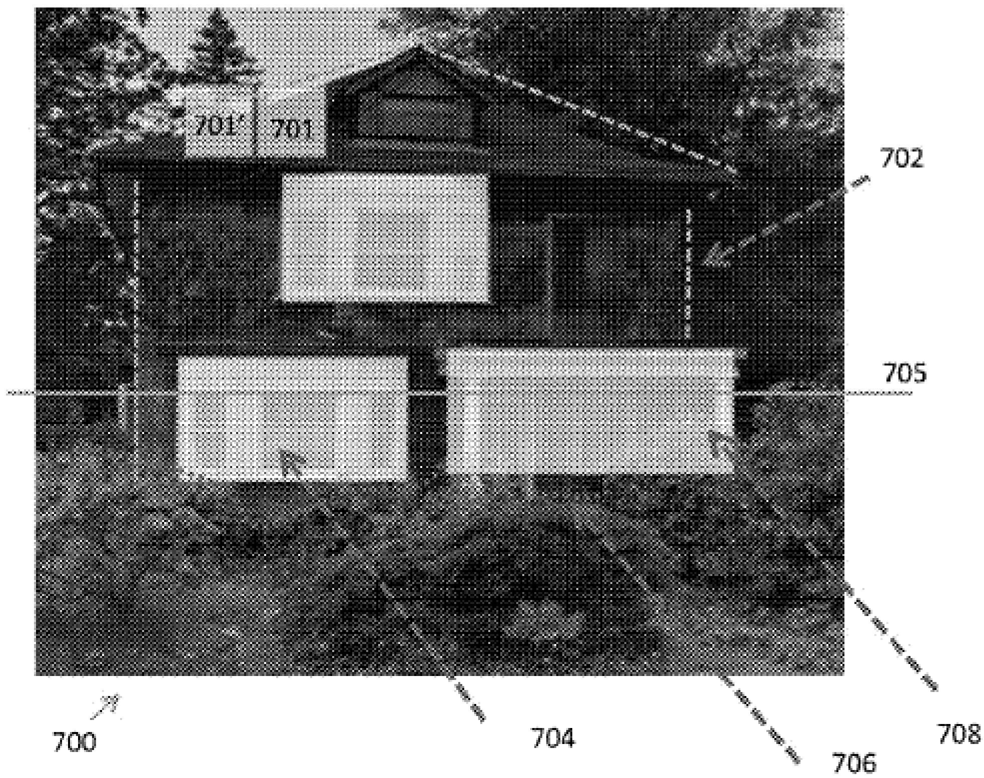

FIGURE 7B

EX1005, FIG. 7B.

101. For the reasons I discuss above in element [6.1], Gross teaches its Classifier Engine determines *a number of stories* and *a siding type*. In light of Gross’s explicit teachings, a POSA would have understood or found obvious that a “terrestrial image” (obtained, for example, by “crowdsourcing”) would “increase building coverage” and “increase feature accuracy” for the classification of *siding type* and *a number of stories* because such image type would have a higher chance

of containing pixels necessary to identify the siding (façade) elements and the number of stories of a structure. EX1005, ¶148.

F. Claim 8

102. Claim 8 depends from claim 1 and further recites, "*wherein a first feature of the at least one feature of the property location comprises a natural feature*."

103. In my opinion, Gross teaches this claim. Gross describes that "the presence, type and condition of … landscaping" (i.e., *natural feature*) can be collected and assessed as part of its classification process. EX1005, ¶141. Gross also discloses "[l]andscaping/grounds, such as "[t]rees, shrubs, hedge, grass," and others are examples of structure attributes and related condition pairs that are identified and classified by its system. EX1005, ¶¶68, 70.

G. Claim 9

104. Claim 9 depends from claim 8 and further recites, "*wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to determine a distance between the natural feature and the property location*."

105. In my opinion, Gross teaches this claim. Specifically, Gross discloses its classification system can be used by "property insurer[s]." EX1005, ¶238. Gross teaches data captured and processed by its assessment system (Classifier Engine)

can include hazards, such as "trees that are too close." EX1005, ¶238. Trees that have a distance which is too close to the property structure are hazards. In order to determine if a tree is "too close" and therefore a hazard, the distance between the tree (i.e., *natural feature*) and the property must be assessed. Therefore, it is my opinion that Gross teaches or renders obvious determining *a distance between the natural feature* (tree) *and the property location*. EX1005, ¶238.

H. Claim 10

106. Claim 10 depends from claim 1 and further recites, "*wherein assessing a first feature of the at least one feature of the property location comprises comparing the classification of the first feature to a historic condition classification of the first feature*."

107. In my opinion, Gross teaches this claim. For instance, Gross discloses using "long term evaluations over defined time periods … to identify changes in a property condition." EX1005, ¶134. Gross accomplishes this by storing the "final structure assessment data" (i.e., *historic condition classification*) "in database 142" during the feature coding (e.g., condition assessment) process. EX1005, ¶133. For example, Gross describes how "the presence of [a building] feature, type and condition is coded and collected and stored in digital form" and that "[i]nformation on … the presence, type and condition of facades, roofs, awnings, porticos, landscaping, flags, exterior features, vehicles, yards, articles, garages…and even

colors of objects can be collected." EX1005, ¶141. Gross's system can "include and assign severity weightings (i.e., a scaling factor of any convenient range (1–10, or the like) for the identified condition)," where "a positive or high end of a condition scale" can correspond to "being in top condition and lacking any noticeable weathering or aging." EX1005, ¶¶70, 135.

108. Gross also explains "the coding process" that obtains the property assessments for the "long term evaluations" described above "may be targeted primarily or solely for identifying structural improvements, such as the presence of a new fence, new roof, new paint job." EX1005, ¶¶134, 146. Given the recorded "structure assessment data," a POSA would have understood or found obvious that the way one would identify these "structural improvements" (e.g., that a new roof was added) would be through a comparison between assessment data of current, updated images of a property and prior, stored historic assessment data of the same property. EX1005, ¶146. For example, the system could compare a new assessment indicating a roof in "good" condition or having a "10" scaling factor with a prior assessment indicating a roof in "damaged" condition or having a "1" scaling factor. EX1005, ¶¶70, 133. These comparisons would indicate the change in condition of the property structure (e.g., from "damaged" to "good," scaling factor difference of 9, etc.) and whether it was a positive change (i.e., improvement) or otherwise.

109. Therefore, it is my opinion that Gross teaches *comparing the*

classification of the first feature (e.g., "good" condition roof) *to a historic condition classification* (e.g., "damaged" roof) as part of *assessing a first feature.* EX1005, ¶¶134, 146

I. Claim 16

110. Claim 16 depends from claim 11 and further recites, "*wherein adjusting each image comprises adjusting color of the at least one image*."

111. In my opinion, Gross teaches, or at least renders obvious, this claim.

112. Gross teaches that Classifier Engine may analyze color in images to assess property conditions through its description that "dark or irregular color [of a group of pixels representing a roof] may cause a roof element to score poorly" during condition determination. EX1005, ¶126. As I discuss above in element [11.4], Gross further describes that customized image processing operations may be applied to images, including "rotations, noise clean up, etc." EX1005, ¶76. Performing "noise clean up" (EX1005, ¶76) (i.e., removing noise from an image) is fundamentally an operation that modifies pixel values. This is because image "noise" manifests as incorrect pixel intensity values, that is, they are pixels that display color levels that are different from what they should represent in the final, noise-free image. For example, a pixel in a blue sky might be expected to have an RGB value of around (100, 150, 200) but may have been corrupted to a different value, such as (180, 120, 150). Then, to remove this noise, the system would

change the pixel's color information from (180, 120, 150) to (100, 150, 200), which would also change a color in the image in the process. Additionally, the incorporated Jafri reference also describes that "appropriate preprocessing of the images can compensate for noise and slight variations in orientation, scale, and illumination." EX1008, 41. A POSA would have known that illumination variation include variations in illumination color and that recognition algorithms can preprocess images to reduce those color effects.

113. Additionally, a POSA would have understood or found obvious that adjusting color of the image is one of a finite number of identified, predictable image adjustment operations that could be used and that many (if not most) image operations (e.g., adjusting contrast, brightness, hue, saturation, color balance, exposure, and sharpness) adjust the color of an image in some fashion. A POSA could have readily implemented well-known color adjustment techniques and would have been motivated to do so to adjust color of the image before determining characteristics and conditions of property features to provide better accuracy in the assessment process. Therefore, it is my opinion that Gross teaches, or at least renders obvious, this claim.

J. Claim 17

114. Claim 17 depends from claim 11 and further recites, "*wherein the set of machine learning analysis algorithms are part of a neural network*."

115. In my opinion, Gross teaches or renders obvious this claim.

116. The ’491 patent describes the machine learning classifier includes deep learning analysis, such as convolutional neural networks (CNN). EX1001, 11:13–19. For the reasons I discuss above in elements [1.6] and [11.5], Gross teaches its Classifier Engine includes *a set of machine learning analysis algorithms*. Similar to the ’491 patent, Gross describes Classifier Engine uses “neural networks” “to detect and extract features.” EX1005, ¶77. Gross notes that feature extraction algorithms are well-known in the art, and identifies a suitable example to be used in its embodiments is described in “*A Survey of Face Recognition Techniques*,” by Jafri (EX1008), which I understand from counsel is incorporated by reference into Gross. EX1005, ¶79. I understand from counsel that material that is incorporated by reference into another document is effectively part of the host document as if it were explicitly contained therein. Jafri, as incorporated in Gross, describes using “a hybrid neural network solution” to analyze image patterns. EX1008, 49. Therefore, it is my opinion that Gross teaches or renders obvious this claim.

K. Claim 18

117. Claim 18 depends from claim 11 and further recites, “*presenting, to a user at a computing device via a network, a graphical user interface comprising information regarding the set of characteristics and the condition of each feature*

of the at least one property feature."

118. In my opinion, Gross teaches this claim.

119. Specifically, Gross describes its system 110 engages with users employing computing devices 112 and, in a preferred embodiment, system 110 is part of a website (i.e., *a graphical user interface*) which can be accessed through a conventional browser. EX1005, ¶43. Gross describes that a report of the results is *present[ed] to a user* (*at a computing device via a network*). EX1005, ¶96. For example, the user can be provided with "classifying details about the building structure such as highlighted, tagged or annotated version of the building structure identified in FIG. 7E … to quickly see the assessment of the property as performed by Building Classifier Engine 150" (i.e., *information regarding the set of characteristics and the condition of each feature*). EX1005, ¶103. Figure 7E (depicted below) presents an embodiment of this visual assessment report and illustrates exemplary *set of characteristics*, such as a "dormer" roof, "shingle" façade, and *condition of property features*, such as a "damaged" roof and "damaged/defective façade."

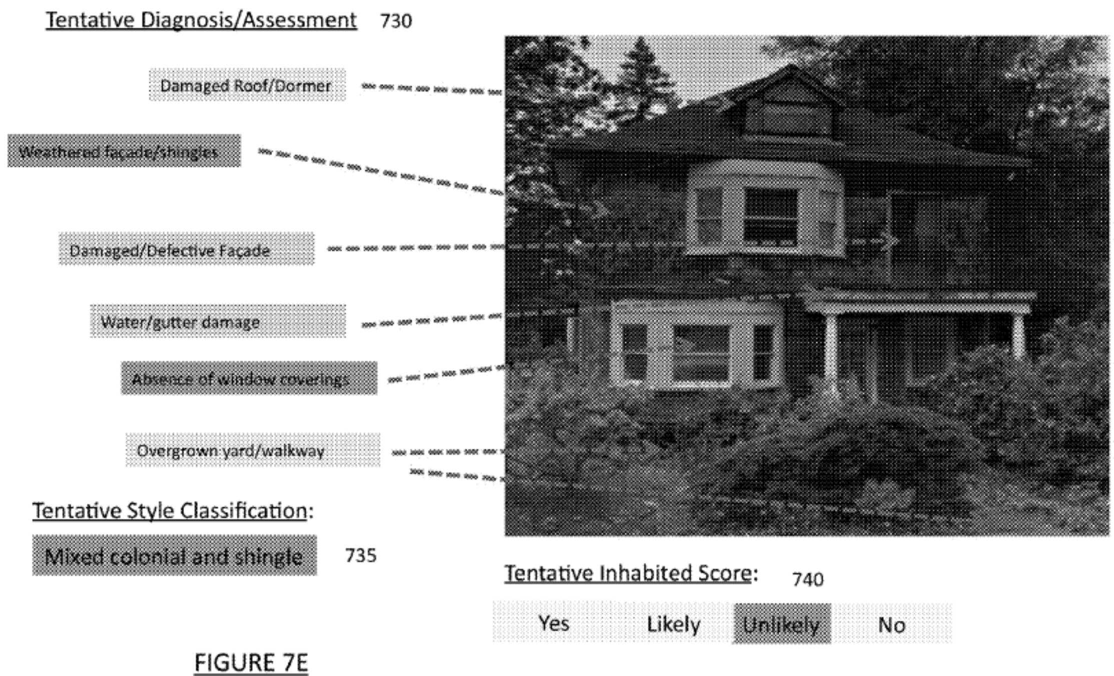

EX1005, FIG. 7E.

120. Gross further discloses that "[t]his data is presented in a list and visual form for a human reviewer." EX1005, ¶130. This "visual assessment report" is *a graphical user interface comprising information regarding the plurality of characteristics and the condition of each feature of the at least one property feature*. Furthermore, although in the example shown in Figure 7E, the visual assessment report may not depict a plurality of characteristics for a feature, this is just an example of a report. Gross teaches *a set of characteristics* may be assessed for a feature for reasons I discuss above in element [1.6] (e.g., for a roof, "dormer" and "shingle"). And Gross describes "the user can be provided with additional classifying details about the building structure … identified in FIG. 7E" to allow

the user "to quickly see the assessment of the property" (EX1005, ¶103)—i.e., Gross teaches its report can include the *set of characteristics* identified in its assessment process.

L. Claim 19

121. Claim 19 depends from claim 11 and further recites, "*presenting, to a user at a computing device via a network, a graphical user interface comprising a map view presenting the structure and at least a portion of the property location, wherein the map view comprises information regarding risk in view of the condition of each feature of the at least one property feature and/or cost in view of the condition of each feature*."

122. In my opinion, Gross renders obvious this claim. Gross discloses that "risk assessments, etc., can be based on an evaluation of an upkeep/maintenance evidenced for a particular property." EX1005, ¶238. Gross further discloses "correlations may be developed between property condition ratings" (i.e., the *condition of a property feature* described above in element [1.7]) and insurance metrics such as "type of claim, severity." EX1005, ¶238. Therefore, it is my opinion that Gross teaches *information regarding the risk* (its risk assessment) *in view of the condition of at least one property feature* (property condition rating, such as a roof in "poor" condition; or trees that are too close to a structure as I discuss above in claim 9).

123. And for the reasons I discuss above in claim 18, it is my opinion that Gross teaches *presenting, a graphical user interface comprising information regarding the set of characteristics and the condition of each feature of the at least one property feature*. Gross does not explicitly disclose that the information presented in the graphical user interface includes the *risk in view of the condition of at least one property feature.* A POSA would have been motivated to modify the visual assessment report (*graphical user interface*) disclosed by Gross to include *information regarding the risk in view of the condition of the property feature*—taught by Gross's insurance risk assessment—in order to provide one report to convey all relevant information to the user. Conveying the information in one consolidated report would have been a simple design choice and would allow a user (e.g., an insurance agent) to see all relevant information in one place, rather than viewing two separate reports. And such a modification to the report to include the risk assessment information would have been a predictable variation and easily implemented by a POSA.

124. Gross also does not explicitly disclose its visual assessment report 730 graphical user interface comprises a *map view*. However, Gross does illustrate a *map view* in Figure 10 (reproduced below) that "shows a typical city sized block in a residential neighborhood which can be assessed … in accordance with embodiments of the present teachings." EX1005, ¶137.

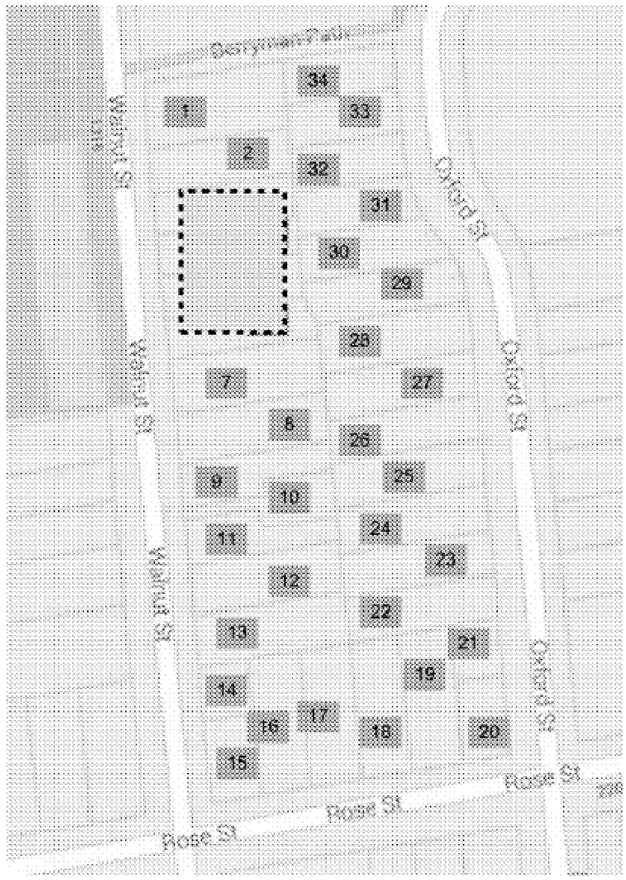

EX1005, FIG. 10.

125. Gross also describes another "interface 350" that "can take any convenient form known in the art, including by identifying the locations of listings on a map (virtual pins)." EX1005, ¶103. With respect to this interface, Gross discloses that "the user can be provided with additional classifying details about the building structure such as highlighted, tagged or annotated version of the building structure," which "allows the user to quickly see the assessment of the property as performed by Building Classifier Engine 150." EX1005, ¶103. A POSA would have been motivated to further modify Gross's visual assessment

report to display in a *map view*, such as shown in Figure 10 and taught by "interface 350," so that the information for all the properties being assessed is displayed in a single report for ease of viewing by the user. EX1005, ¶103. Using a "map view" to convey the information would have been a simple design choice and a predictable variation easily implemented by a POSA.

M. Claim 20

126. Claim 20 depends from claim 11 and further recites, "*wherein the another machine learning analysis algorithm is trained to identify a first condition of one or more conditions of the property feature having at least two characteristics, wherein the at least two characteristics comprises at least two of the following: a type of material of the property feature, a distance between the property feature and another feature at the property location, and a shape of the property feature*."

127. In my opinion, Gross teaches this claim. For the reasons I discuss above in element [1.7] and [11.6], Gross teaches *another machine learning analysis algorithm trained to identify a first condition of one or more conditions of the property feature.* Gross also discloses additional attributes and characteristics that can be identified, "such as façade, eaves, windows, balcony, porch, arch, piers, columns, lattices, false timbering, ornamental, etc. along with specific types (i.e., façade (shingle, siding, brick, stone, horizontal boards, vertical boards, etc.) roof

{pitched, double pitched, hipped, flat, metal, tile, shingle, slate, parapet, dormer, mansard, fascia, brackets, eaves, pent, pediment, etc.} and so on)." EX1005, ¶57. The examples include *a type of material of the property feature and a shape of the property feature*—e.g., characteristics of a roof (*property feature*) can be identified as flat (*a shape of the property feature*) and metal (*a type of material of the property feature*). EX1005, ¶57. Therefore, it is my opinion that Gross teaches this claim.

X. GROUND 2: THE COMBINATION OF GROSS AND FURUKAWA RENDERS OBVIOUS CLAIMS 2–5

128. In my opinion, the combination of Gross and Furukawa renders obvious claim 2–5.

A. Claim 2

129. Claims 2 depends from claim 1 and further recites, "*wherein the boundary information comprises a shape map*."

130. In my opinion, the combination of Gross and Furukawa teaches, or at least renders obvious, this claim. The '491 patent describes a "shape map" as being "used to confirm [a] location of a particular property," such as being "overlaid with a corresponding aerial image 102c to match properties with images." EX1001, 7:58–62. Figures 5A and 5C of the '491 patent (shown below) respectively depict an example "shape map" and an example result of overlaying shape outlines with an aerial image.

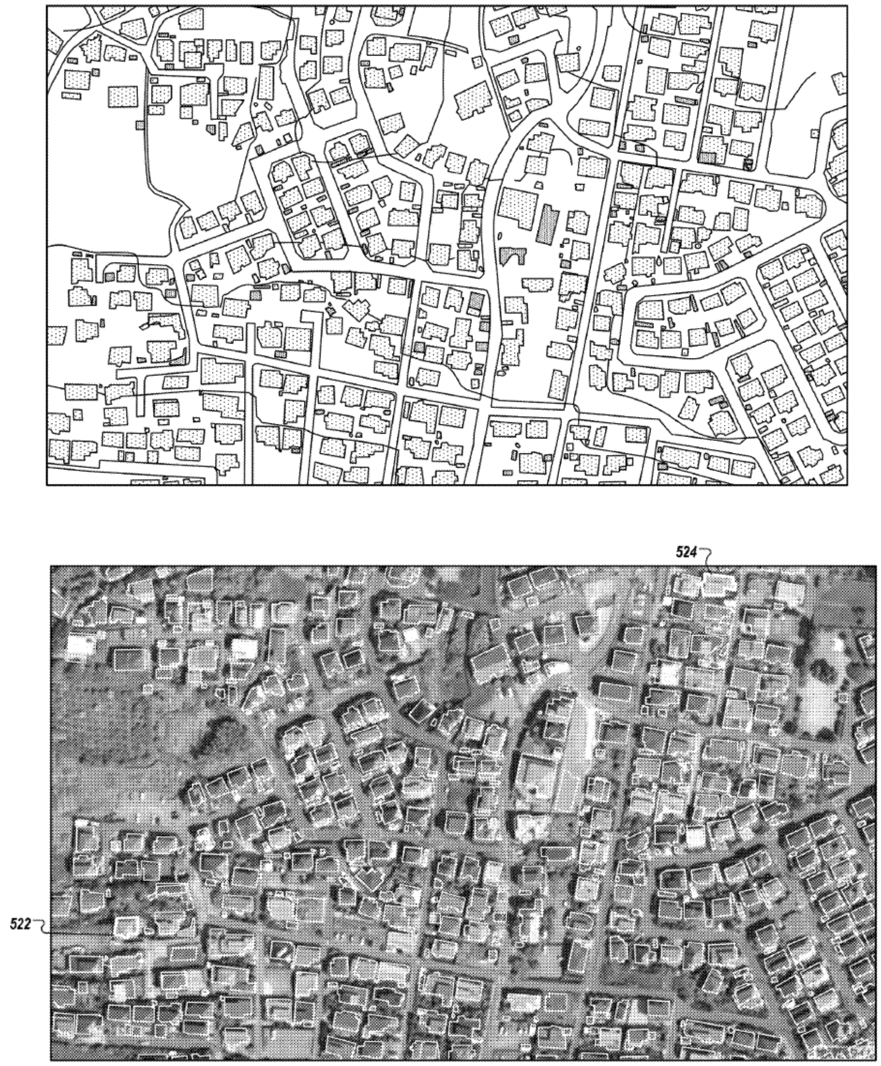

FIG. 5A (top), FIG. 5C (bottom).

131. Figure 10 of Gross (reproduced below) shows a similar image depicting a map view containing various shapes corresponding to "a typical city sized block … which can be assessed … in accordance with embodiments of [Gross's] teachings," on which various buildings have been "identified with unique numbers." EX1005, ¶137.

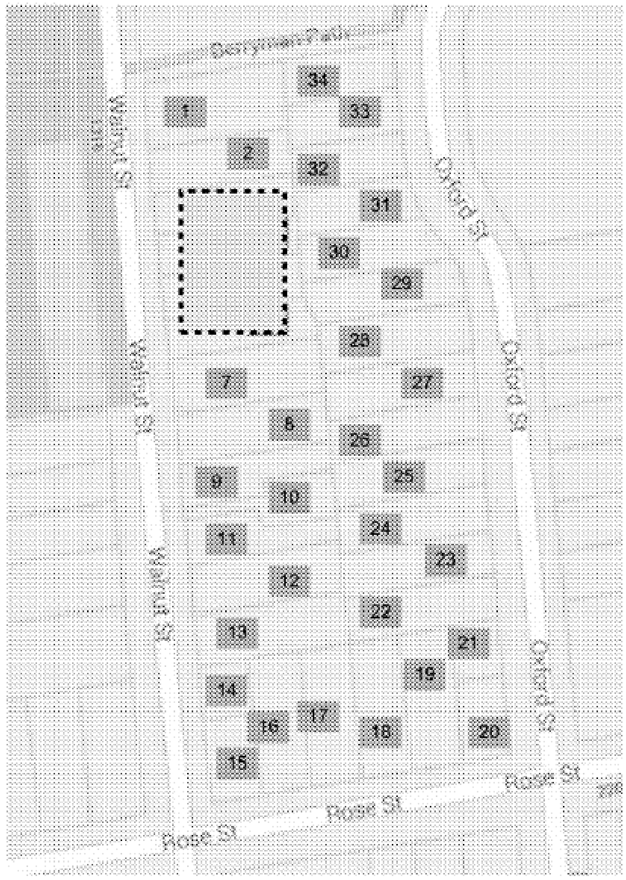

EX1005, FIG. 10.

Gross describes that the images of "any particular city, town, etc., can be" "identified" "in this fashion … based on computer records, address data, etc. … for purposes of achieving the objectives" set forth in Gross (e.g., property assessment). EX1005, ¶137.

132. Figure 4 of the incorporated Pershing material (reproduced below) also shows a similar image containing various shapes "showing the building 92 and the surrounding areas 93 around the building 92." EX1009, 4:12–14.

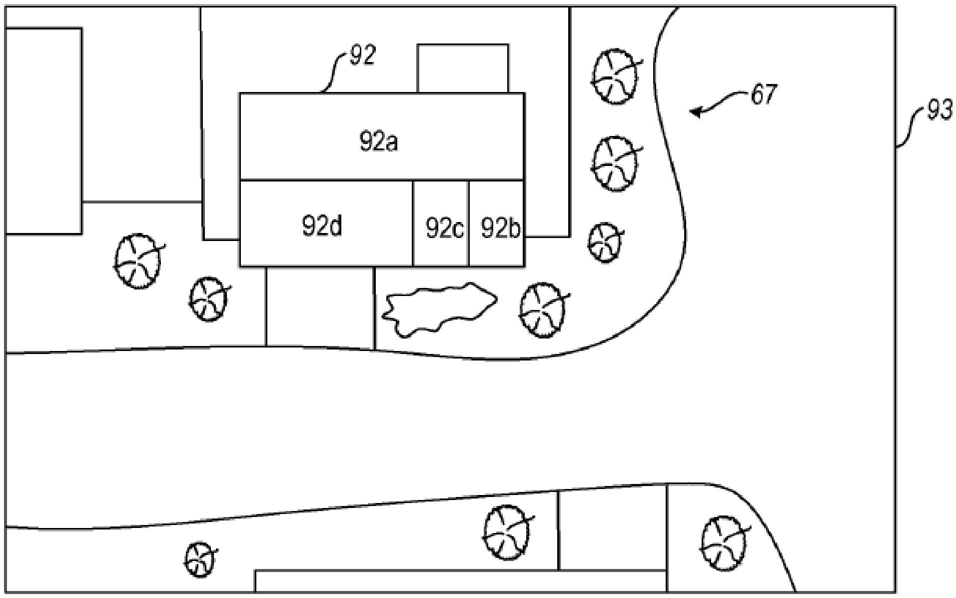

EX1009, FIG. 4.

Both Figure 10 of Gross and Figure 4 of the incorporated Pershing material are images of shapes that are mapped together to delineate properties and their surrounding objects, which can serve as the "shape map" that the '491 patent "use[s] to confirm [a] location of a particular property." EX1001, 7:58–62.

133. Additionally, as I discussed above in element [1.3], Gross describes a "building envelope" (the claimed *boundary information*) that *isolate[s] a property structure*. And for the reasons I discuss above in element [1.3], Gross further teaches applying different *boundary information to isolate the property location*. Gross does not explicitly disclose that its "[i]nformation" used to isolate the location of the "target address in question" is *a shape map*. EX1005, ¶¶162, 164.

134. In my opinion, Furukawa teaches *boundary information compris[ing]*

a shape map. In particular, Furukawa teaches obtaining map data that includes "information of the external shape of buildings" which is stored and provided by a "map data storing means 4." EX1007, 15:19–31. Figure 7 of Furukawa depicts an example of such map data:

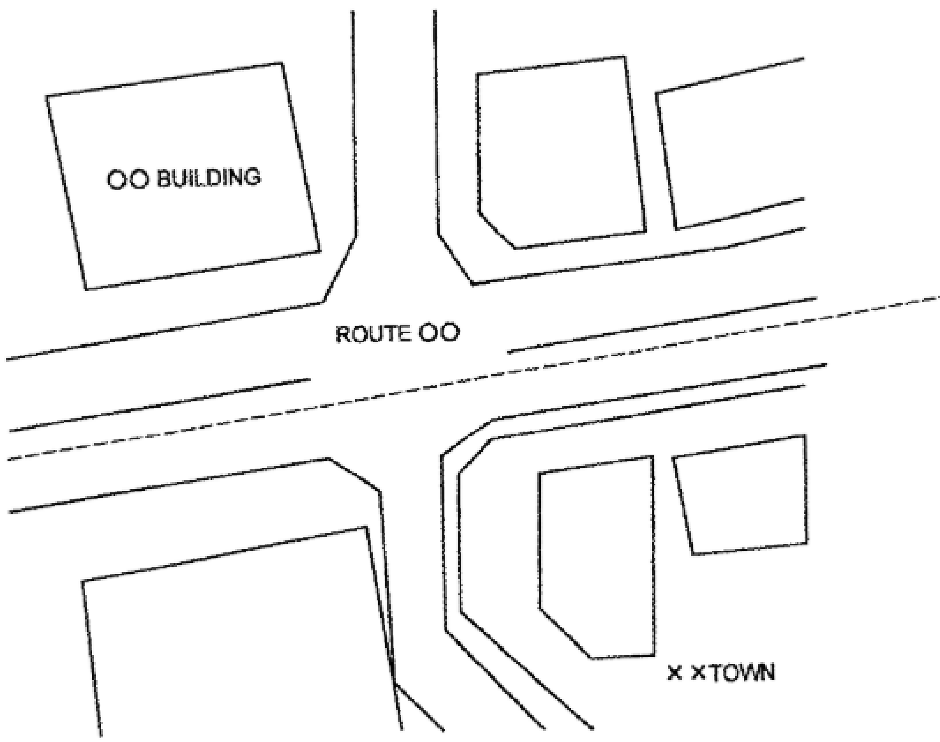

EX1007, FIG. 7.

135. Furukawa further explains how the "map data" depicted in Figure 7 includes "shapes of buildings" as "raster data … of those data showing shapes of buildings," and can also include identification information, such as "names of administrative districts/roads/buildings and the like" and location information, such as "latitude and longitude." EX1007, 12:26–39. This map data is *boundary information compris[ing] a shape map*, because, similar to the '491 patent it plots

out the boundaries and shapes of various structures in an area and includes their location information. And, similar to the ’491 patent, which describes that the shape map data may be obtained from external sources (e.g., Geospatial Information Authority of Japan or Zenrin Co. Ltd. of Japan), Furukawa also describes its map data may be obtained from many products that are commercially available. *Compare* EX1001, 7:51–57, *with* EX1007, 8:62–66.

136. Such *shape maps* (e.g., urban planning maps, cadastral maps) were well-known before the September 23, 2016 priority date of the ’491 patent. For example, the Geospatial Information Authority of Japan launched its National Large Scale Map Data Base project in 1989 to compile digital map data for urban planning, utility management, navigation, etc., and opened the project to the public in 1993. EX1021, 9. The United States Geological Survey (USGS) released *The National Map* in 2001, which compiled topographic information including boundary, transportation, and structure data for the U.S. EX1022, 8.

137. A POSA would have been motivated to modify Gross to incorporate Furukawa’s teachings of using map data, because Gross and Furukawa provide a teaching, suggestion, and motivation to do so. Gross teaches the need for address confirmation, describing how “building image[s] and tentative address[es] are tentatively tagged” during the classification process, and how “this address information” should be “confirmed” to ensure the system “contain[s] accurate

information for a particular address." EX1005, ¶¶165, 170. Gross's teaching that tentative address information should require confirmation would have suggested to a POSA the benefit of incorporating additional data sources for such verification purposes. Indeed, the process of confirming a piece of data involves comparing that piece of data to one or more independent reference points to check or prove its validity. But Gross doesn't describe how the address confirmation would occur, and that would have motivated a POSA to seek out such solutions.

138. A solution to this need is supplied by Furukawa. Furukawa teaches the use of "map data" containing, for each of its various buildings: building shape information such as "the external shape of buildings," identification information such as "names of …buildings," and "location information such as latitude and longitude" data, each of which is an additional stored reference data point that serves as the confirmation mechanism Gross contemplates. EX1007, 8:55–58, 12:26–34. Furukawa also teaches using this "map data" for verification purposes through the validation and correction of satellite image analysis of buildings. EX1007, Abstract. Thus, the prior art provided both Gross's problem of building address confirmation and Furukawa's map data validation solution, creating a teaching and motivation for their combination. A POSA would have had a reasonable expectation of success in combining known elements (Furukawa's shape map with Gross's building envelope and boundary information) according to

known methods to yield predictable results (boundary information comprising a shape map).

139. Therefore, it is my opinion that the Gross-Furukawa combination teaches *boundary information compris[ing] a shape map* (e.g., map data including "information of the external shape of buildings"). EX1007, 15:19–31.

B. Claim 3

140. Claim 3 depends from claim 1 and further recites, "*wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to select the one or more images as a subset of the plurality of images responsive based at least in part on the boundary information*."

141. In my opinion, the Gross-Furukawa combination teaches, or at least renders obvious, this claim. Gross teaches *select[ing] the one or more images as a subset of the plurality of images*. As I discuss above in element [1.2], Gross discloses maintaining and accessing multiple images of structures in its database (i.e., a *plurality of images*). For example, Gross discloses how "image data … for the building structure is retrieved for the target address in question" and can be "obtained from a third party vendor" or "generated as needed" using various methods such as "aerial drones, satellite, balloon, and similar technology." EX1005, ¶164. These methods "can obtain image data [at] different elevations" and "at different times." EX1005, ¶164. Gross further teaches that its "structure

image database" contains primary images for a structure, as well as "structure sub-images which are based on automatically dividing the original image into separate blocks, or separate areas corresponding to distinct building elements." EX1005, ¶165.

142. Gross further discloses how during the structure classification process, "image data for a particular target address is obtained at step 1725, along with a tentative address tag" and how the address tag is "confirmed at step 1730" to ensure the system "contain[s] accurate information for a particular address." EX1005, ¶170. Then at step 1750, the system proceeds to "identify, classify and rate a condition of features" in the image. EX1005, ¶172. This process is illustrated in Figure 17 of Gross (annotated below).

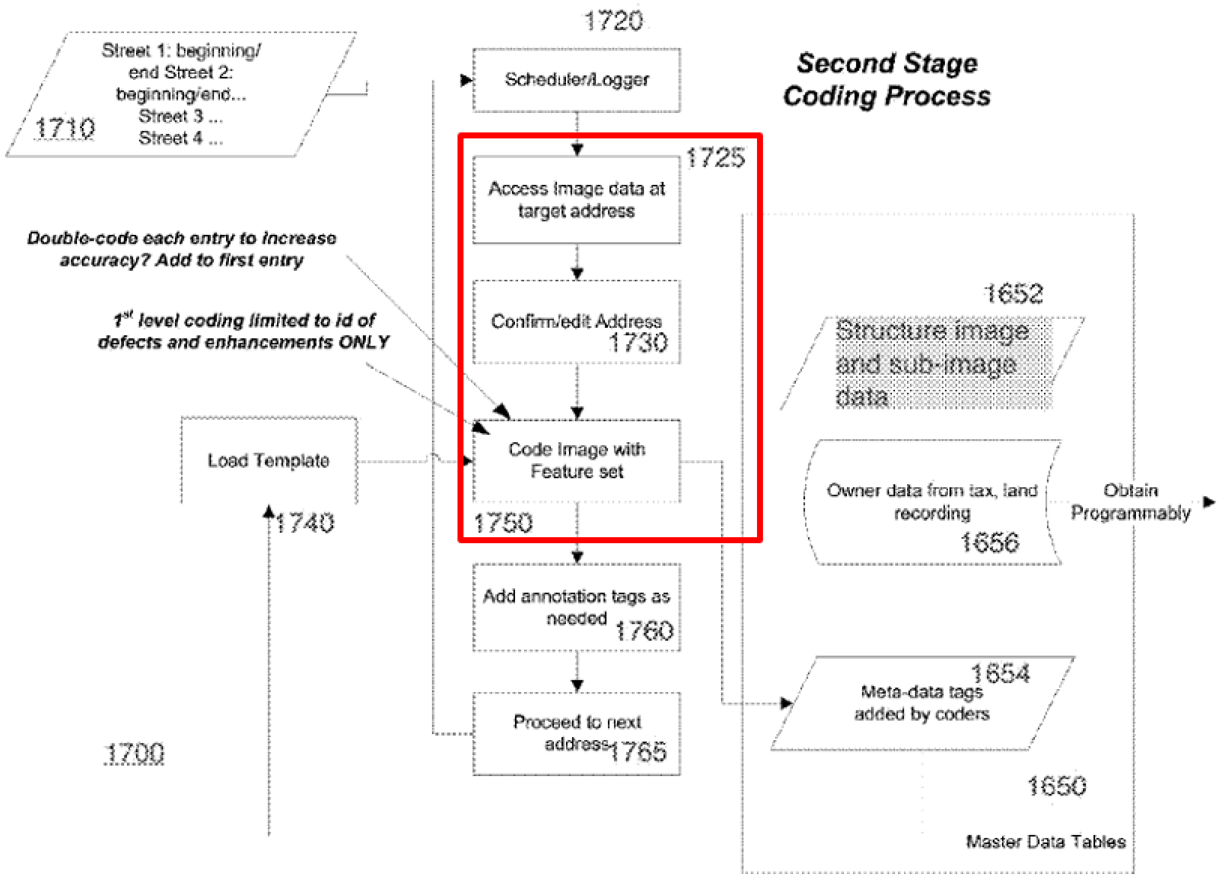

EX1005, FIG. 17A (annotated).

This confirmation process is *select[ing] the one or more images as a subset of the plurality of images*, because in the combination of Gross and Furukawa, the images with confirmed address tags would be used for classification at step 1750 (as opposed to images whose address tags cannot be confirmed). Put more simply, the images with confirmed address tags are selected for assessment.

143. As described above, Furukawa provides a solution for confirming addresses associated with an image of a property using map data that includes "information of the external shape of buildings" (*boundary information*) and

"identification information." EX1007, 15:19–31. In particular, Furukawa's system first "analyzes … obtained satellite images, and extracts vector information showing the external shape for the photographed building," which Furukawa describes is performed by a "building external shape information detecting means 6." EX1007, 32:65–33:7.

144. Furukawa's system then determines whether this boundary (extracted from the satellite images) matches the boundary of the same building in the stored map data. Specifically, Furukawa teaches that a "building external shape information comparing/combining means 7 compares and combines the information of the external shape of buildings of the map data … with the information of the external shape of buildings extracted by the building external shape information detecting means 6." EX1007, 32:17–24. During this process, "a pair of the information of the external shape of buildings, which exist at locations only a little way from each other and of which difference in areas is smaller than a predetermined value and ratio of the common part is larger than a predetermined value, is made to correspond to each other as information indicating the same building." EX1007, 23:56–24:2.

145. In other words, the system determines whether the boundaries match by evaluating factors including location proximity, area differences, and overlap ratios. For example, boundaries that matched up with each other would produce

low area differences and high overlap ratios. This would also mean that any matching building identified from the satellite image would share the same identification information (e.g., "building" "name") and "location information" (e.g., "latitude and longitude") with its map data counterpart, since they are linked to the same building with the matching boundaries. When these calculations fall within "predetermined value[s]," the system considers the depicted boundaries as matching and "indicating the same building." EX1007, 23:56–24:2. Inversely, boundaries that do not match up with each other would produce high area differences and low overlap ratios.

146. A POSA would have found it obvious to modify Gross to incorporate Furukawa's boundary comparison teachings for at least the same reasons I discuss above for claim 2. As I explained for claim 2, the prior art provided both Gross's problem of building address confirmation and Furukawa's map data validation solution, creating a teaching and motivation for their combination and would have had a reasonable expectation of success.

147. Therefore, it is my opinion that the Gross-Furukawa combination teaches, or at least renders obvious, *select[ing]* (e.g., proceeding with classification) *the one or more images* (e.g., images with confirmed address tags) *as a subset of the plurality of images* (e.g., images of a property location) *responsive based at least in part on the boundary information* (e.g., boundary in

"satellite image" matching a boundary from "map data"). EX1007, 15:19–31, 32:65–33:7.

C. Claim 4

148. Claim 4 depends from claim 3 and further recites, "*wherein selecting the one or more images comprises identifying the one or more images as most closely matching the boundary information*."

149. In my opinion, the Gross-Furukawa combination teaches or at least renders obvious this claim. As I discuss above in claim 3, Furukawa's system determines whether a boundary extracted from the satellite images matches the boundary of the same building in the stored map data by evaluating factors including location proximity, area differences, and overlap ratios. When these evaluation factors match within "predetermined value[s]," the system considers the depicted boundaries as matching and "indicating the same building." EX1007, 32:17–24, 23:56–24:2.

150. A POSA would have found it obvious and been motivated to modify Gross to incorporate Furukawa's boundary matching teachings for at least the same reasons I discuss above for claims 2 and 3. As I explained above, the prior art provided both Gross's problem of building address confirmation and Furukawa's map data validation solution, creating a teaching and motivation for their combination.

151. Therefore, is it my opinion that the Gross-Furukawa combination teaches, or at least renders obvious, *wherein selecting the one or more images* (e.g., proceeding with classification on an image) *comprises identifying the one or more images as most closely matching the boundary information* (e.g., boundary in "satellite image" matching a boundary from "map data"). EX1007, 15:19–31, 32:65–33:7.

D. Claim 5

152. Claim 5 depends from claim 1 and further recites, "*wherein applying the boundary information comprises cropping each respective image*."

153. In my opinion, the Gross-Furukawa combination teaches this claim. For instance, Gross discloses how "[d]uring or after the coding [e.g., classification] process, the image for the structure can be automatically divided by an image processing system into different sub-images of different size," "which correspond to distinct building elements" and can be "selected by the coder" "using a conventional image cropping tool, or automatically identified and bounded by an automated classifier" (i.e., *cropping*). EX1005, ¶174; *see also* EX1005, ¶76 (describing how Gross's images "may be cropped" as part of an image pre-processing or normalization process). Gross explains that this cropping process is performed "to identify the element and its condition" and "be used in reports [and] responding to queries." EX1005, ¶174. Gross does not explicitly describe whether

these teachings necessarily involve *applying boundary information*.

154. Furukawa teaches *cropping* when *applying boundary information*. For instance, Furukawa describes obtaining "three-dimensional data images" (e.g., "DEM (Digital Elevation Map) Data") and "superimposing the map data on the three-dimensional data in such a manner that the same points match each other." EX1007, 13:1:3, 1:17–18, 7:5. As I discuss above for claim 2, the "map data" in Furukawa includes "information of the external shape of buildings," which is *boundary information* that defines the perimeter of a building structure. EX1007, 15:19–31. Furukawa further describes how "[a]fter the map data is superimposed on the three-dimensional [data], the three-dimensional data included in a closed area surrounded by vector data is set for each building as an area corresponding to the building by the area setting means 32." EX1007, 13:4–7. For each building in the map data, the three-dimensional data is set to the area within the external shape of the building as the candidate area of each building. EX1007, 9:13–17 (describing how "area setting means 32 sets the three-dimensional data included in the area within the external shape of the building as the candidate area of each building"). In Furukawa, "[t]he area setting means 32 sets the three-dimensional data included inside as a building candidate area for each building area in the superimposed map data (step S12 in FIG. 4). For each set building candidate area, the in-area histogram analysis means 33 performs statistic analysis to obtain a

histogram, and so on for three-dimensional data values included in the area (step S13 in FIG. 4)." EX1007, 10:18–24.

155. Figure 6A of Furukawa depicts an "example three-dimensional data image" (EX1007, 7:5), and Figure 8A depicts an example "three-dimensional data set by area setting means" (EX1007, 7:10), both of which are reproduced below:

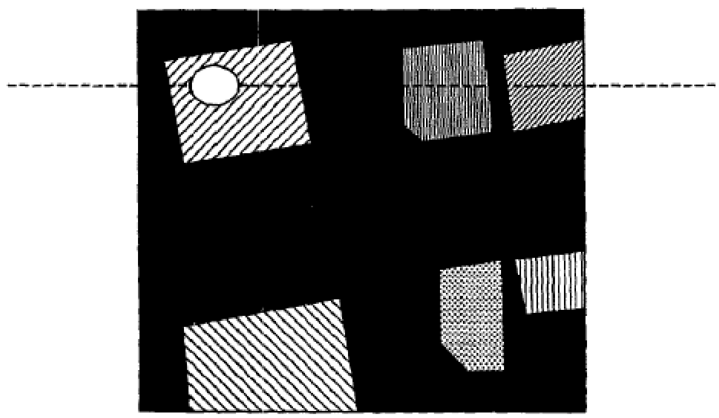

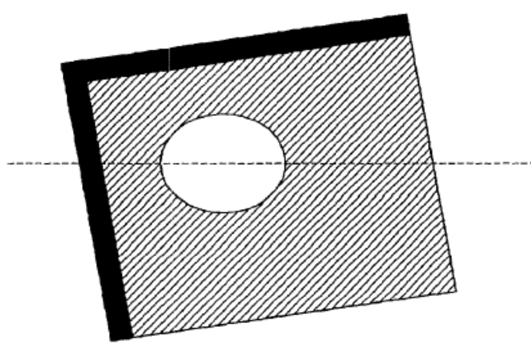

EX1007, FIG. 6A (top), FIG. 8A (bottom).

Here, the "area setting means" is *cropping* the *image* when *applying the boundary information*, because it is using (i.e., *applying*) the "closed area surrounded by vector data" (i.e., the map data's *boundary information*) to obtain "the candidate area of each building" from the "three-dimensional data image." EX1007, 13:4–7, 7:5, 9:13–17. In other words, Figure 8A is a cropped version of Figure 6A based on the boundary information in the map data.

156. A POSA would have found it obvious and been motivated to modify Gross's classification system to incorporate Furukawa's boundary-based cropping techniques because the combined Gross-Furukawa system would provide more precise and accurate identification of individual building structures within images. The modification would have been a predictable variation of Gross, because Gross explicitly contemplates using "automated classifier[s]" to "identif[y] and bound" building elements, and Furukawa's superimposition and area setting techniques represent straightforward applications of known image processing and geographic data integration methods that a POSA would routinely use to enhance the automated classification and bounding capabilities contemplated by Gross. EX1005, ¶174; EX1007, 15:19–31, 13:1:3. A POSA would also have readily been able to implement Furukawa's techniques in Gross because both systems similarly use automated image processing to analyze building structures, with Gross providing the framework for structure classification, and Furukawa providing complementary techniques for isolating structures using established boundary data.

157. Therefore, it is my opinion that the Gross-Furukawa combination teaches, or at least renders obvious, *wherein applying the boundary information* (e.g., map data including "information of the external shape of buildings")

comprises cropping each respective image (e.g., "set[ting] the three-dimensional data included in the area within the external shape of the building as the candidate area of each building"). EX1007, 13:1–3, 9:13–17, 15:19–31.

XI. GROUND 3: THE COMBINATION OF GROSS AND DAVIS RENDERS OBVIOUS CLAIMS 12–15

158. In my opinion, the combination of Gross and Davis renders obvious claims 12–15.

A. Claim 12

159. Claim 12 depends on claim 11 and further recites, "*applying, by the processing circuitry, a risk profile corresponding to at least one risk to the property location or the structure thereon in view of the classification of the condition of each of the one or more property features to determine a damage risk to the property location.*"

160. In my opinion, the combination of Gross and Davis renders obvious this claim. Gross discloses that "risk assessments, etc., can be based on an evaluation of an upkeep/maintenance evidenced for a particular property; in this respect correlations may be developed between property condition ratings" and insurance metrics such as "type of claim, severity." EX1005, ¶238. These risk assessments are determinations of *a damage risk to the property location*. And Gross teaches risk profiles through its description that these correlations may associate certain maintenance condition ratings of a property attribute (i.e., the

condition of a property feature I describe above in element [1.7], such as a roof in "poor" condition or trees that are too close to a structure as I describe above in claim 9) with higher risk of damage (e.g., from "arson"). EX1005, ¶238. Obtaining a "risk assessment" for a disaster using "correlations … developed between property condition ratings … and … severity" would involve *applying a risk profile* (e.g., the "correlations") *corresponding to at least one risk* (e.g. "arson") and *the classification of the condition of property features* (e.g., the attributes associated with the "property condition ratings") because the correlations are correlating the risk of damage for a specific disaster with the analyzed "condition rating" of a property attribute/structure. EX1005, ¶¶238, 78.

161. Davis discloses *risk profiles*. Like Gross, Davis also describes analyzing digital images and features of a property to determine risk of damage to a property structure. EX1006, 2:7–16. To determine risk, Davis describes a risk analysis unit 34 that "utilize[s] certain types of a priori knowledge to determine risk," which "may be incorporated as rules stored in data storage 40." EX1006, 10:22–24. Davis describes "risk analysis unit 34 may access rules specifying that dead branches … are prone to cause at least a certain amount of damage." EX1006, 10:31–34. Davis further describes other risk information based on a chance that a dead or dying tree will fall over and a branch in close proximity to a chimney will cause damage due to backdraft. EX1006, 11:30–38. These rules may also specify

whether trees of certain species more likely to catch fire. *See* EX1006, 11:25–39. Risk analysis unit 34 then applies the "access[ed] rules" to "estimate a likelihood that damage will occur to some portion of the property." EX1006, 11:30–34. In other words, Davis more explicitly teaches its risk analysis unit *appl[ies]* rules (*a risk profile*) *corresponding to a condition of a property feature* (tree location/condition) *to determine a damage risk to the property location*.

162. A POSA would have been motivated to modify Gross's teachings of a classification system used to rate the condition of a property for insurance risk purposes with Davis's specific teachings of applying a risk profile to determine a damage risk because the combined Gross-Davis system would provide an efficient method for the insurer to determine the risk of property damage and to adjust risk profiles as needed. The modification would have been a predictable variation of Gross, because Gross explicitly contemplates developing "correlations," and Davis's specific risk correlation techniques—implemented through "rules stored in [a] data storage"—represent straightforward applications of known data correlation and rule-based processing methods that a POSA would routinely use to establish the correlations contemplated by Gross. EX1006, 10:23–25. And a POSA would have readily been able to implement Davis's techniques in Gross with a reasonable expectation of success because both systems similarly use image analysis and machine learning to analyze features of properties to determine insurance risk.

163. The Gross-Davis combination would, therefore, in my opinion render obvious *applying a disaster risk profile* (e.g., the "access[ed] rules") *corresponding to at least one risk to the property location* (e.g., risk from wildfire) *in view of the classification of the condition of the one or more property features* (e.g., a roof in poor condition, a tree that is too close to the structure) *to determine a damage risk to the property location*.

B. Claim 13

164. Claim 13 depends on claim 12 and further recites, "*wherein the at least one risk comprises a fire or an explosion*."

165. In my opinion, the combination of Gross and Davis teaches this claim. Gross describes risk assessment due to a "risk of arson" (*a fire*). EX1005, ¶238. And Davis describes analyzing features of property "to determine whether particular homes are at high risk of fire damage." EX1006, 11:18–25. And for the reasons I discuss above in claim 12, the Gross-Davis combination teaches that the risk assessment is based on *the classification condition of the one or more property features* (e.g., whether property contains dead or dying trees or trees that are too close to a structure). *See* EX1006, 11:25–36. Therefore, it is my opinion that the combination of Gross and Davis teaches this claim.

C. Claim 14

166. Claim 14 depends on claim 12 and further recites, "*wherein applying*

the risk profile comprises modeling vulnerability of the property location to damage in view of the at least one risk."

167. In my opinion, the combination of Gross and Davis renders obvious this claim. For the reasons I discuss above in claim 12, Gross-Davis, in my opinion, renders obvious *applying a risk profile corresponding to at least one risk to the property location or the structure thereon*. As discussed, Gross describes its risk correlations are "developed between property condition ratings … and type of claim, severity, etc." EX1005, ¶238. Gross's development of risk correlations teaches *modeling vulnerability of the property location to damage in view of the at least one risk* (e.g., a risk correlation modeling vulnerability of the property in view of arson because the property condition is not occupied). *See* EX1005, ¶238.

168. Davis also teaches *modeling vulnerability of the property location to damage in view of the at least one risk*. For example, Davis describes its risk analysis unit may estimate a likelihood that damage will occur to the property (i.e., *modeling vulnerability of the property location to damage*), such as a percent chance that a branch in close proximity to a chimney will cause damage due to backdraft (*in view of the at least one risk*). EX1006, 11:32–29.

169. Therefore, it is my opinion that the Gross-Davis combination renders obvious this claim.

D. Claim 15

1. [15.1] The method of claim 12, further comprising: determining, by the processing circuitry, a characteristic of a neighboring structure at the property location in addition to the structure;

170. In my opinion, Gross teaches element [15.1]. Gross describes that structural element fields (classified by Classifier Engine as part of its assessment process) for a property include "an identification of each structural element presented in the property" and "a condition of such element." EX1005, ¶¶115–16. Gross discloses a particular property "may have multiple structural elements" and the "same information" (classification/assessment information) "is made for appurtenant elements … including for example ancillary buildings, fences, garages, vehicles, foliage, etc." (i.e., *neighboring structures*). EX1005, ¶116. In other words, Gross teaches *determining a characteristic of a neighboring structure.* Examples of these *neighboring structures* taught by Gross include the appurtenant elements (e.g., garage, retaining wall, fence), and natural structures, such as trees and landscaping. EX1005, ¶¶67–68.

2. [15.2] wherein applying the risk profile comprises applying the risk profile in view of the classification of the condition of a first feature of the one or more property features and the characteristic of the neighboring structure.

171. In my opinion, the Gross-Davis combination renders obvious element [15.2]. As I discuss above for claim 12, Gross discloses that "correlations may be developed between property condition ratings" (i.e., the *conditions* I discuss above

in element [1.6]) and insurance metrics such as "number of claims filed, type of claim, severity, etc." EX1005, ¶238. And Davis describes that "risk analysis unit 34" "may analyze the feature(s) to determine a risk of damage to at least a portion of the property" and its rules may, for example, depend "on the type of roof construction." EX1006, 10:19–38. That is, risk analysis unit 34 may *apply a risk profile in view of the characteristic of the structure*.

172. In light of Davis's teachings, a POSA would have understood or found obvious that in addition to the *applying a risk profile in view of the classification of the condition of the first feature of the structure*, such as a tree feature that is "too close" or "overgrown" as I discuss above in claim 9 (EX1005, ¶238), the insurance risk assessment taught by Gross-Davis (e.g., *applying the risk profile*) would have also been based, at least in part, on *the characteristic of the neighboring structure*, such as a characteristic of a roof of a structure because it was both known and is common sense the risk varies depending on the characteristics of a structure (e.g., a metal roof would have less risk of damage than a slate roof for certain types of disasters, e.g., from trees that are too close, etc.).

XII. CONCLUSION

173. In signing this declaration, I recognize that the declaration will be filed as evidence in an *inter partes* review before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in the case and that cross-examination will take place within the United States. If such cross-examination is required, I will appear for cross-examination within the United States during the time allotted.

174. I hereby declare that all statements made herein of my own 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.

Executed on this 29 day of July 2025, in Traverse City, MI.

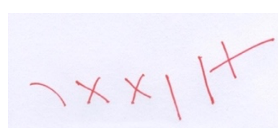

David A. Forsyth, Ph.D.