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 IPR2025-01357
U.S. Patent No. 10,650,285

DECLARATION OF DR. DAVID A. FORSYTH IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 10,650,285

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. INTRODUCTION 1
II. MY BACKGROUND AND QUALIFICATIONS 2
III. SUMMARY OF OPINIONS 7
IV. LIST OF DOCUMENTS CONSIDERED 8
V. LEGAL UNDERSTANDING 9
A. My Understanding of Claim Construction 10
B. My Understanding of Obviousness 10
C. A Person of Ordinary Skill in the Art 13
VI. OVERVIEW OF THE '285 PATENT 14
A. Summary of the '285 Patent 14
B. Prosecution History 17
C. Level of Ordinary Skill in the Art 18
D. Claim Construction 19
VII. STATE OF THE ART 19
A. Aerial imagery analysis and property assessment techniques have been used for decades, well before the '285 patent. 19
B. Machine learning and computer vision techniques for image analysis were well-known before the '285 patent 21
VIII. OVERVIEW OF THE APPLIED REFERENCES 23
A. Gross 23
B. Davis 26
IX. GROUND 1: GROSS RENDERS OBVIOUS CLAIMS 1-7 AND 9-22 28
A. A POSA would have found it obvious Gross's techniques were intended to be combined. 28
B. Independent Claim 1 30
1. [1.P]: A method for automatically categorizing a condition of a characteristic of a property, the method comprising: 30

2. [1.1]: obtaining, by processing circuitry, an aerial image of a geographic region including the property; 34

3. [1.2]: extracting, by the processing circuitry, one or more features of a plurality of features from the aerial image, wherein the one or more features are each represented by a set of pixel groupings, and the one or more features represent the characteristic of the property, ... 35

4. [1.3]: wherein each pixel grouping of the set of pixel groupings comprises at least one of angles, outlines, or substantially homogenous pixel fields; 40

5. [1.4]: applying, by the processing circuitry, a first portion of the set of pixel groupings to a first machine learning classifier to determine a characteristic classification of the characteristic from a plurality of potential classifications of the characteristic, wherein the first machine learning classifier is trained to identify at least a portion of the one or more features from the first portion of the set of pixel groupings; and 41

6. [1.5]: applying, by the processing circuitry, a second portion of the set of pixel groupings to a second machine learning classifier to determine a condition classification of the characteristic from a plurality of potential condition classifications of the characteristic, wherein the second machine learning classifier is trained to identify, from the second portion of the set of pixel groupings, property characteristic conditions of the characteristic having the characteristic classification. 44

C. Independent Claim 9 49

1. [9.P]: A system for automatically categorizing a condition of a property feature of a property, the system comprising: 49

2. [9.1]: a non-transitory computer readable storage region storing a first machine learning analysis model trained to identify one or more property characteristics, and 50

3. [9.2]: a second machine learning analysis model trained to identify one or more property conditions; 50

4. [9.3]: 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 extract a set of image-related features from an aerial image, wherein the set of image-related features represent the property feature;..... 51

5. [9.4]: apply a first portion of the set of image-related features to the first machine learning analysis model to determine a characteristic classification of the property feature from a plurality of potential characteristic classifications of the propery feature; and 51

6. [9.5]: apply a second portion of the image-related features to the second machine learning analysis model to determine a condition classification of the property feature from a plurality of potential condition classifications of the property feature, wherein the second machine learning analysis model is trained to identify, from the second portion of the set of image-related features, property characteristic conditions of the property feature having the characteristic classification. 52

D. Independent Claim 16 52

1. [16.P]: A non-transitory computer readable medium having instructions stored thereon, wherein the instructions, when executed by processing circuitry, cause the processing circuitry to: 52

2. [16.1]: obtain aerial and/or terrestrial imagery including a property; 52

3. [16.2]: extract a plurality of pixel groupings from the aerial and/or terrestrial imagery, wherein the plurality of pixel groupings includes a respective set of pixel groupings representative of each feature of one or more features of the property, wherein..... 54

4. [16.3]: each pixel grouping of the set of pixel groupings comprises at least one of angles, outlines, or substantially homogenous pixel fields; and..... 55

5. [16.4]: determine at least one characteristic and at least one condition of the property, wherein determining comprises, for each feature of the one or more features of the property,..... 55

6. [16.5]: apply at least a portion of the one or more sets of pixel groupings corresponding to the respective feature to one or more first machine learning classifiers trained to identify one or more characteristics of the respective feature, and 55

7. [16.6]: apply at least a portion of the one or more sets of pixel groupings corresponding to the respective feature to one or more second machine learning classifiers trained to identify a condition of the respective feature in view of at least a first characteristic of the one or more characteristics. 56

E. Claim 2 56

F. Claim 3 58

G. Claim 4 60

H. Claim 5 61

I. Claim 6 61

J. Claim 7 62

K. Claim 10 63

L. Claim 11 64

M. Claim 12 66

N. Claim 13 67

O. Claim 14 68

P. Claim 15 70

Q. Claim 17 72

R. Claim 18 72

1. [18.1]: The non-transitory computer readable medium of claim 17, wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to access one or more known property characteristics,..... 72

2. [18.2]: wherein determining the one or more of the damage risk estimate, the repair estimate, or the replacement estimate comprises applying, to the determining, the one or more known property characteristics. 73

S. Claim 19 74

T. Claim 20 74

U. Claim 21 76

V. Claim 22 77

X. GROUND 2: THE COMBINATION OF GROSS AND DAVIS RENDERS OBVIOUS CLAIM 8 80

A. Claim 8 80

XI. CONCLUSION 85

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. 10,650,285 ("the '285 patent") (EX1001), which resulted from U.S. Patent Application No. 16/733,888 ("the '888 application"), filed on January 3, 2020. I understand that the '285 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 '285 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 Rashtchian1, 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 '285 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 1-22 of the '285 patent would have been obvious to a POSA prior to September 23, 2016.

Ground	Basis	Claims	References
1	§103	1-7 and 9-22	Gross
2	§103	8	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 '285 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 1-22 of the '285 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 '285 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. 10,650,285 to Okazaki ("'285 patent")
1002	Prosecution History of U.S. Patent No. 10,650,285 ("Prosecution History")
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")
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")

Exhibit No.	Description
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

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 '285 patent. However, my opinions herein would not be affected if the '285 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.

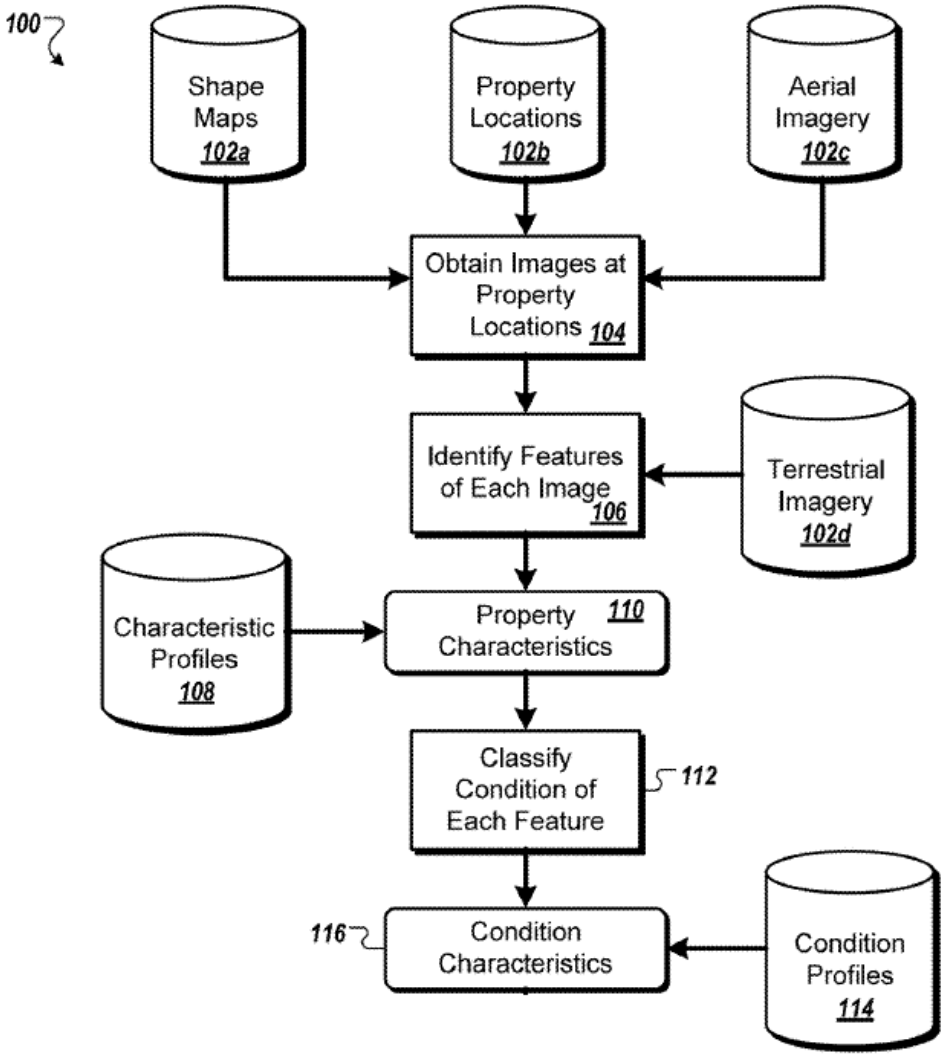
VI. OVERVIEW OF THE '285 PATENT

A. Summary of the '285 Patent

34. The '285 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, 1:36-41, 2:29-37. 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, 2:37-44, 12:23-27 ("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 '285 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, 2:11–20, 29–37 (emphasis added). Figure 1, reproduced below, illustrates an "example system

operational flow 100 for matching aerial image features to corresponding property maintenance levels," EX1001, 6:12-14:



EX1001, FIG. 1 (excerpted).

36. As shown in Figure 1, the system first obtains images at property locations at step 104. At step 106, the system then identifies (or "extracts") "features" of each image, such as "substantially homogenous fields" and "pixel groupings" (groups of pixels). EX1001, 8:22–32. These "features" represent a

characteristic of the property, such as a "rooftop" or a "tree." EX1001, 8:26–36; *see also* EX1001, 2:56–66 ("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, 8:32–42. For example, this may include determining the "type of rooftop" (e.g., gambrel, gable, hipped, square, flat). EX1001, 8:32–42, 9:15-35, FIG. 2A; *see also* EX1001, 2:66-3:6 ("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 models trained to classify property

characteristics. EX1001, 8:43–9:14.

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, 3:17–18, 10:41–49. To do so, "condition profiles 114" (e.g., machine learning models trained to classify conditions of property characteristics) are applied to the extracted image features from step 106. EX1001, 10:52–12:15.

39. The '285 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 '285 patent apply existing machine learning and image analysis techniques. For example, the '285 patent identifies convolutional neural network (CNN) technology as one example, which the '285 patent describes as "a well-established and popular machine-learning methodology." EX1001, 8:42–67, Claim 5.

B. Prosecution History

40. I understand that the '285 patent issued from U.S. Application No. 16/733,888 ("the '888 application"), without any substantive prior art examination. The '285 patent is a continuation-in-part of U.S. Application No. 15/714,376 ("the '376 application") that is now U.S. Patent No. 10,529,029. EX1002, 0011. The prosecution history reveals that the claims never underwent meaningful analysis against the prior art that renders them obvious. *See generally* EX1002. I also

understand that during prosecution, the '888 application received an objection to claim 5 for including an acronym ("CNN") and a double-patenting rejection over the '029 patent. EX1002, 0204-09. In response, the applicant amended the claims to spell out the acronym and filed a terminal disclaimer. EX1002, 0246-47, 0254-60. The Examiner never applied any prior references in 35 U.S.C. §§ 102, 103 rejections, despite the robust prior art landscape that existed at the time of filing and allowed the '888 application following the applicant's amendment and response. EX1002, 0271-73.

C. Level of Ordinary Skill in the Art

41. In my opinion, a POSA at the time of the '285 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.

42. I am well qualified to determine the level of ordinary skill in the art. I am very familiar with the technology of the '285 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

43. I have been informed that the challenged claims of the ’285 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 ’285 patent specification.

VII. STATE OF THE ART

A. Aerial imagery analysis and property assessment techniques have been used for decades, well before the ’285 patent.

44. Computerized analysis of aerial imagery for property assessment was well-established before the September 23, 2016 priority date of the ’285 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. 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.

45. 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.

46. 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 '285 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 '285 patent

47. 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 '285 patent as an example of a CNN that can be used in the machine learning analysis of the disclosed embodiments (*see* EX1001, 8:42-47, 8:52-58), 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 '285 patent as an example machine learning classifier that can be used in the disclosed embodiments (*see* EX1001, 1:55-60, 8:58-61), 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.

48. 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

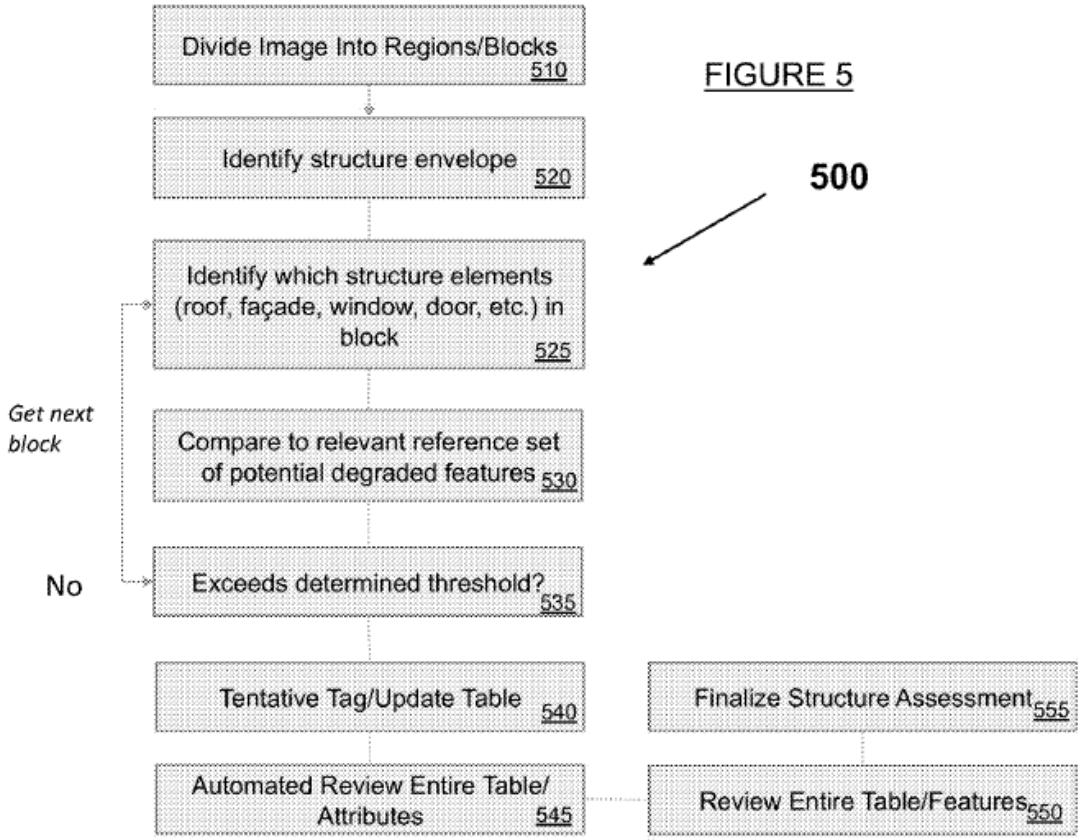
A. Gross

49. 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 '285 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 '285 patent, because it is in the same field of endeavor and reasonably pertinent to the problems of the '285 patent.

50. 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.

51. 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 images may be aerial images captured by "aerial drones, satellite, balloon, and similar technology." EX1005, ¶164. 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.

52. 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.

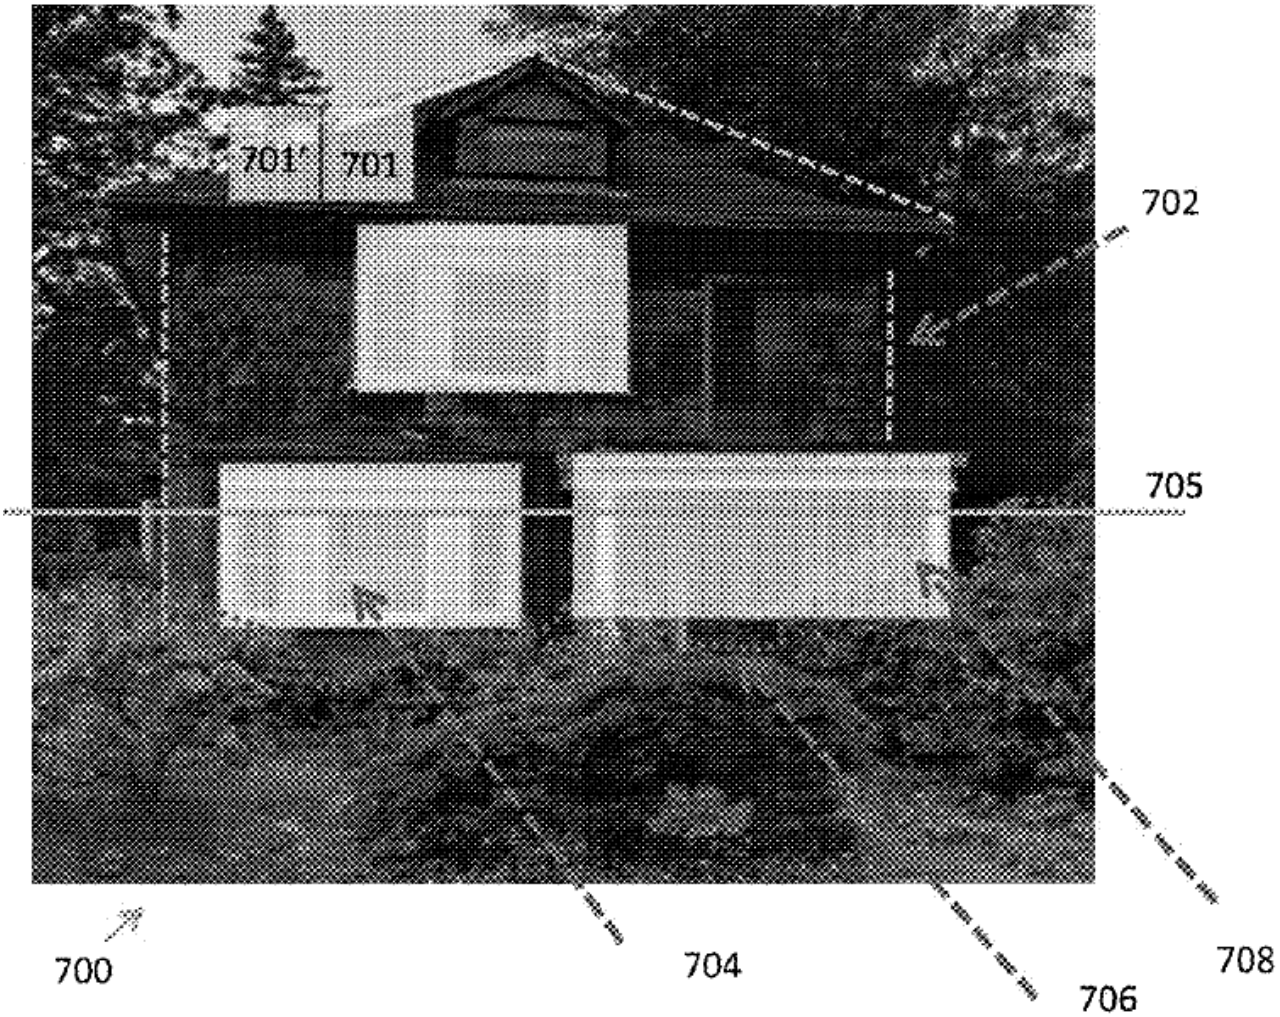


FIGURE 7B

EX1005, FIG. 7B.

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 '285 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.

53. Again similar to the '285 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

54. 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 '285 patent, because it is in the same field of

endeavor and reasonably pertinent to the problems of the '285 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 '285 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.

55. 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, etc." or "a type of roof on a structure." EX1006, 18:27–39; 9:22–29. Based on these determinations, Davis's system then, similar to the '285 patent, analyzes a risk of damage to a structure on the property (e.g., a house). EX1006, 19:6–16.

IX. GROUND 1: GROSS RENDERS OBVIOUS CLAIMS 1-7 AND 9-22

56. In my opinion, claims 1–7 and 9-22 would have been obvious in view of Gross.

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

57. 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.

58. 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.

59. 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.

60. 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.

61. 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 method for automatically categorizing a condition of a characteristic of a property, the method comprising:

62. 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.

63. It should first be noted that the word *property* can have multiple meanings, which can affect how the claim phrases *characteristic of a property* and *property feature* are interpreted. The word *property* can refer to either: 1) a physical building or piece of real estate or 2) an aspect, quality, or facet of something. In some sections, the '285 patent uses the word *property* in the second sense. *See, e.g.*, EX1001, 1:23–26 (describing how a "risk exposure database contains a compilation of as many building properties or characteristics relevant to insurance as possible," and that "[t]hese properties can include characteristics like location coordinates, address, slope, and elevation").

64. However, the '285 patent primarily uses the term *property* in the first sense. For example, the '285 patent describes that a "property may be a single family home." EX1001, 3:51–52. Using this meaning, the claimed *characteristic of a property* would refer to an aspect or quality of a physical building or piece of real estate. This is consistent with how the '285 patent exemplifies *characteristics of a property* or property characteristics:

> 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.

EX1001, 2:56–66. I note that the '285 patent does not specifically differentiate between the terms *characteristic of a property* and *property feature*, and uses them interchangeably across its disclosure. For example, various examples in the above list for *characteristic of a property* ("rooftop," "swimming pool," "chimney," and "sky lights") overlap with the examples that the '285 patent provides for *property feature*. *Compare* EX1001, 2:56–66, *with* EX1001, 8:24–30. The '285 patent also refers to "property features" as examples of *property characteristics*:

> "The property characteristic(s), for example, may be classified by the property characteristic classification engine 322, as described in relation to FIG. 3. In a particular example, a rooftop property feature may be classified as one of a variety of rooftop shapes."

EX1001, 21:4–8 (emphasis added).

65. Gross discloses "training [a] Classifier Engine … and using it to assess and rate" properties. EX1005, ¶57. "[A]n 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 *characteristics of a property*, as claimed, because they are physical parts of a structure, such as a house or building.

66. Gross discloses a "condition score is computed [by Classifier Engine] for the detected attribute" (i.e., *a condition of a characteristic of the property*). EX1005, ¶126. For example, the Classifier Engine may automatically categorize the condition of a roof (characteristic of a property) as "new," "missing/damage tile," and/or "holes/cracks." EX1005, ¶61. In some instances, "severity weightings … for the identified condition" may also be assigned by Classifier Engine. EX1005, ¶70. 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 *a method for automatically categorizing a condition of a characteristic of a property*.

67. Gross's Figure 11 "identifies examples of structural features, parameters, conditions … for a particular building structure," which illustrates the results of the Classifier Engine categorizing conditions of a characteristic of a property, as claimed. EX1005, ¶21.

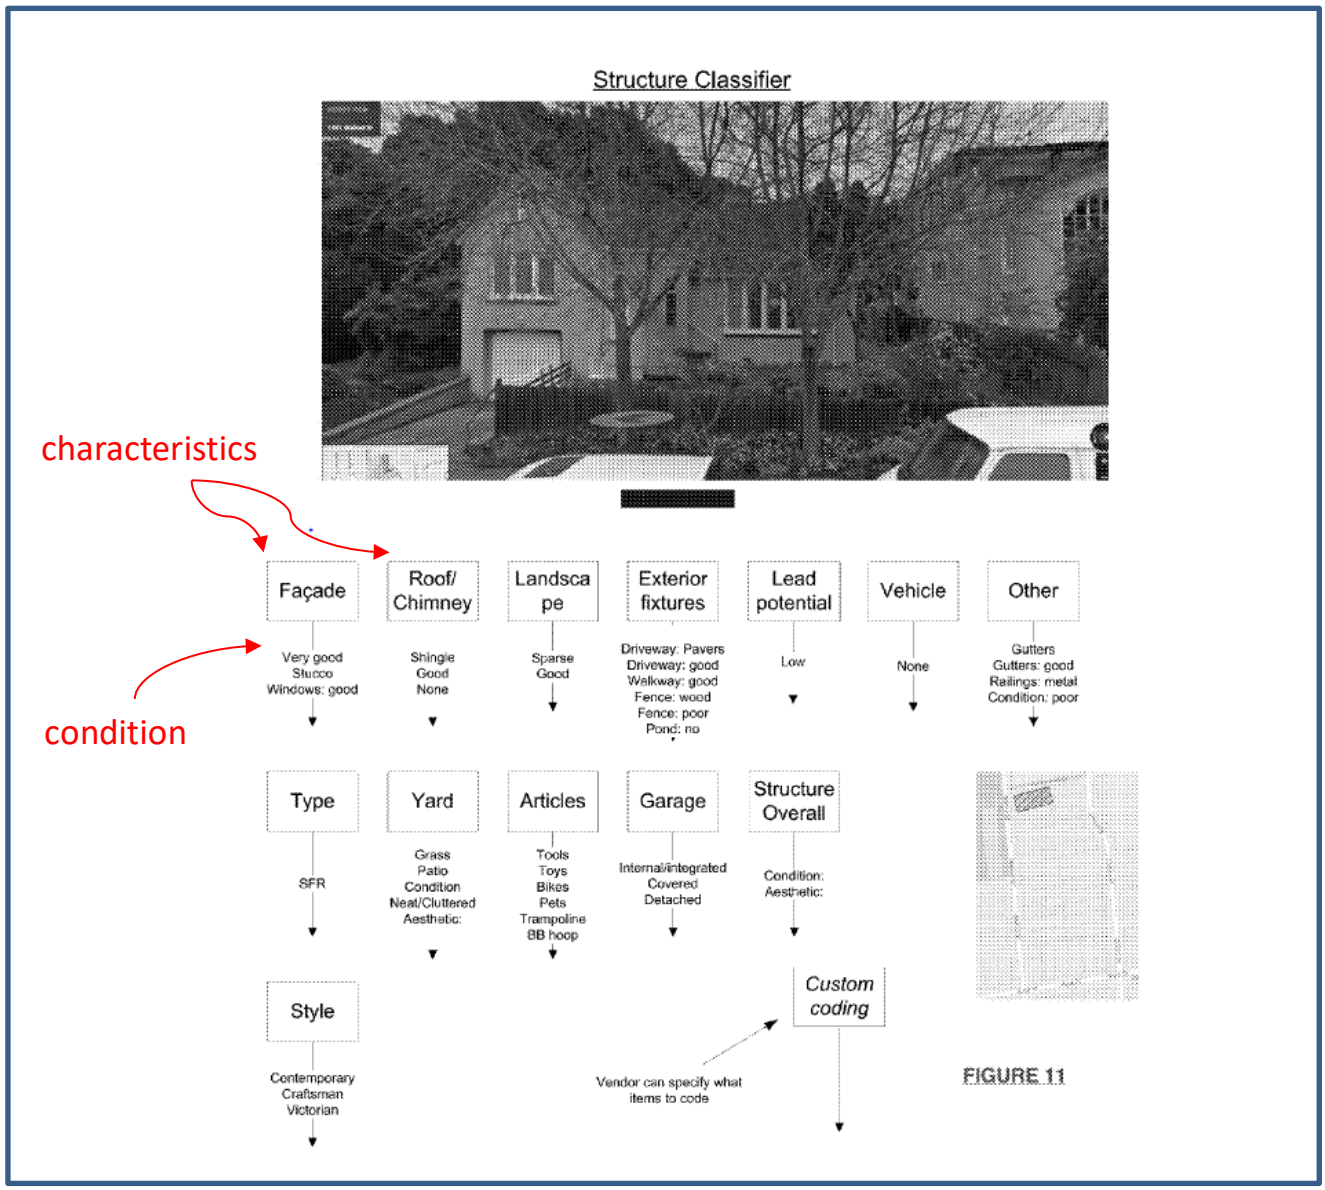


FIGURE 11

EX1005, FIG. 11 (annotated).

2. [1.1]: obtaining, by processing circuitry, an aerial image of a geographic region including the property;

68. In my opinion, Gross teaches element [1.1]. In particular, Gross 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," such as a city. EX1005,

¶161. Gross also discloses that "aerial drones, satellite, balloon, and similar technology can be used … to easily capture structure image data." EX1005, ¶164. These images captured by the aerial drones, satellite, and balloon (i.e., "image data … for the building structure … for the target address in question") are *aerial images of a geographic region that includes the property*. EX1005, ¶164. This building image data is then stored in "a structure image database 1652." EX1005, ¶165.

69. Gross further describes "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. A POSA would have understood or found obvious that these "executable software routines" are executed via "processing circuitry." Accordingly, it is my opinion that Gross teaches *obtaining by processing circuitry* the aerial image.

3. [1.2]: extracting, by the processing circuitry, one or more features of a plurality of features from the aerial image, wherein the one or more features are each represented by a set of pixel groupings, and the one or more features represent the characteristic of the property,

70. In my opinion, Gross teaches element [1.2]. Element [1.2] includes three parts: (i) extracting, by the processing circuitry, one or more features of a plurality of features from the aerial imagine; (ii) the one or more features are each represented by a set of pixel groupings (and it should be noted that a set

encompasses a set of one grouping of pixels); and (iii) the one or more features represent the characteristic of the property. In my opinion, Gross teaches all three parts.

71. Similar to my discussion above with respect to the word *property*, it should be noted that the word *feature* can also have multiple meanings, depending on the context. In standard contexts, the word *feature* can refer to an aspect, quality, or facet of something (like meaning 2, above, of the word *property*). For example, the '285 patent describes "property features" as being physical aspects, qualities, or facets (e.g., "rooftop, swimming pool, chimney, and sky lights") of a property. EX1001, 8:24–30. When applied to 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). For example, the '285 patent describes that its "extract[ed] image-related features" can include "angles, outlines, substantially homogenous fields, etc." EX1001, 8:26–36. Accordingly, the claimed *features*, here, appear to follow the computer vision meaning of the term. *See also, e.g.*, EX1001, claim 9 (clarifying that it is *image-related features* that are *extract[ed]*).

72. First, Gross discloses objects can be analyzed in image patterns, using "a variety of techniques" that can be used to detect and *extract* features of

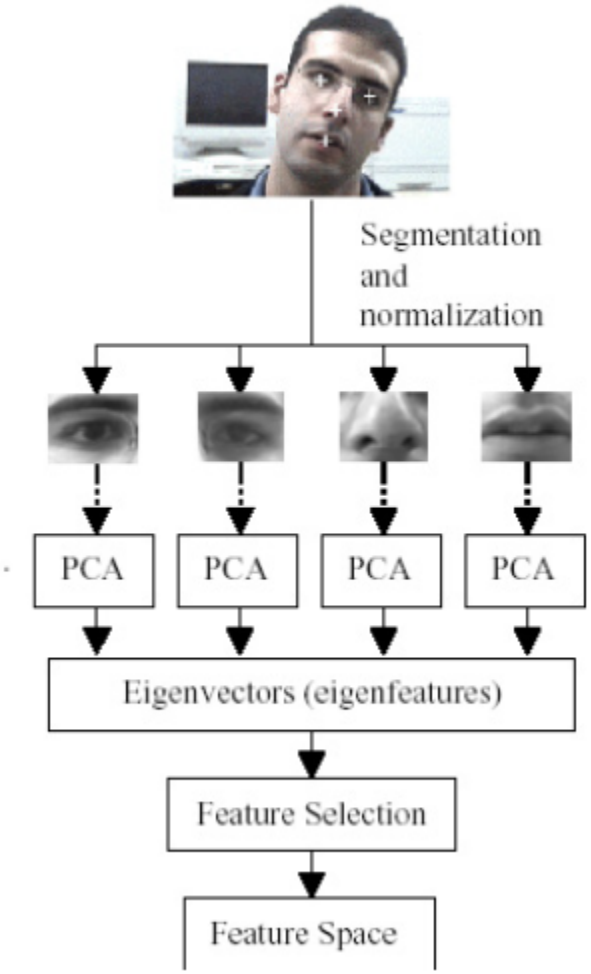
attributes. EX1005, ¶77. An image (e.g., an *aerial image* as I discussed above in element [1.1]) is obtained and "an identification is performed to determine structural elements presented in a first block **701**, including such attributes as roof elements, window elements, façade elements, etc." EX1005, ¶¶120, 123. A POSA would have understood or found obvious this identification is performed *by processing circuitry* that *extracts features from the aerial image*. In fact, Gross specifically discloses that its Classifier Engine is trained by "extract[ing] smaller distinct (image) features" from reference images and further notes that "[f]eature extraction algorithms are well-known in the art." EX1005, ¶¶77, 79.

73. Second, Gross explicitly teaches that the *one or more features are represented by a set of pixel groupings*. Gross discloses that "[u]sing conventional image processing, each attribute" (i.e., *characteristic of the property*) "may be … characterized by a distinct set of pixels, shapes, proportions, curves, textures, or otherwise mathematically derivable feature from the image." EX1005, ¶78. It should be noted that from an image processing perspective, digital images are made up of pixels arranged in a grid shape, where each pixel represents a single point of color and intensity data. As such, any perceivable visual element of an image such as geometric shapes, curved boundaries, text patterns, etc. are necessarily composed of and represented by groups of constituent pixels.

74. Thus, the "distinct set of pixels, shapes, proportions, curves, textures,

or otherwise mathematically derivable feature from the image" in Gross that characterize building attributes are all examples of *pixel groupings* that represent one or more features (of the image). EX1005, ¶78. Each exemplary type of characterization is an underlying group of pixels in the image (e.g., an extracted feature characterized by a "shape" is a pixel group that represents the shape). For example, a curved roofline exists in the digital image as a sequence of pixels whose color and intensity values differ from surrounding pixels in a pattern that processing circuitry (and the human eye) can recognize as a curve.

75. Gross further notes that feature extraction algorithms are well-known in the art and a suitable example that can be used in its embodiments is set forth 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. Gross, through the incorporated Jafri material, additionally discloses *features represented by a set of pixel groupings* through Jafri's description of a recognition system that detects "prominent facial landmarks (i.e., the eyes, nose, and mouth)" to track those *features.* EX1008, 43, 51. For example, Figure 19 of Jafri (depicted below) illustrates pixel groupings around the nose, eye, and lips.



EX1008, FIG. 19.

76. Third, Gross discloses the attributes (i.e., characteristics of the property) "are analyzed in the images to break them down and reduce them into (or extract) smaller distinct features that can be more readily detected." EX1005, ¶76. In other words, the extracted image features are used to detect the attributes (*characteristics of the property*) during classification. Thus, the extracted image feature(s) *represent the characteristic of the property*. Exemplary characteristics of the property that may be represented by the image features include roof elements,

window elements and façade elements. *See* EX1005, ¶¶60-68, 123.

4. [1.3]: wherein each pixel grouping of the set of pixel groupings comprises at least one of angles, outlines, or substantially homogenous pixel fields;

77. In my opinion, Gross teaches element [1.3]. As I discussed above, Gross discloses the claimed *pixel groupings* through its description that each attribute in the image may be "characterized by a distinct set of pixels, shapes, proportions, curves, textures, or other mathematically derivable feature from the image." EX1005, ¶78. Additionally, Gross, through the incorporated Jafri material, discloses extracting features and computing geometric relationships (e.g., "distances, areas and angles") to represent pixel groupings of facial features, both of which are feature extraction techniques that Gross notes "are suitable and/or can be adapted easily for use in [its] present embodiments." EX1008, 43; EX1005, ¶79. Gross further discloses identification of a building envelope (i.e., *outline* of the building). EX1005, ¶122. Therefore, it is my opinion that Gross teaches *each pixel grouping comprises at least one of angles* ("angles" or "other mathematically derivable feature"); *outlines* (shapes, curves, building envelope); *or substantially homogenous pixel fields*.

5. [1.4]: applying, by the processing circuitry, a first portion of the set of pixel groupings to a first machine learning classifier to determine a characteristic classification of the characteristic from a plurality of potential classifications of the characteristic, wherein the first machine learning classifier is trained to identify at least a portion of the one or more features from the first portion of the set of pixel groupings; and

78. In my opinion, Gross teaches element [1.4].

79. I note that in computer vision contexts, with respect to the classification process (e.g., running pixel inputs through a machine learning classifier to obtain a classification result), it is typically the machine learning classifier that is "applied" to a set of pixels. This phrasing is consistent with how the '285 patent describes the classification process in various places. *See, e.g.*, EX1001, 3:8–10 (describing how "[a]nalyzing the features to determine a property characteristic may include applying a deep learning analysis model to the features"), 3:12–15 (describing how "analyzing the region of the aerial image including the property characteristic to determine the condition classification includes applying a machine learning analysis model to image pixels within the region"). However, it appears that the claimed *applying the set of pixel groupings to a first machine learning classifier* is also referring to the same process of running pixel inputs through a machine learning classifier *to determine a characteristic classification*, a process that Gross, in my opinion, teaches for the reasons I discuss below.

80. Gross discloses that, after Classifier Engine 150 (the claimed *first machine learning classifier*) is trained, Gross's system uses Classifier Engine 150 to perform identification of attributes (e.g., roof elements) presented in an image. EX1005, ¶123. The specific types of the attributes (i.e., *characteristic classification(s)*) are also captured during property assessment "(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." EX1005, ¶120.

81. For instance, Gross describes that 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," which can include "pattern recognition techniques and statistical image processing using models, templates, etc." EX1005, ¶123. In other words, Classifier Engine *applies a first portion of the set of pixel groupings* extracted in element [1.2] (e.g., a portion of the set of pixel groupings that represent a roof) *to determine a characteristic classification* (e.g., a portion of the pixel groupings that identifies the roof type as "shingle"). For the reasons I discuss

above for element [1.1], a POSA would have understood or found obvious that processing circuitry executes the assessment process performed by Classifier Engine. *See also* EX1005, ¶161 (describing that the critical steps in the process can be implemented into executable software routines and modules). Therefore, it is my opinion that Gross teaches that the processing circuitry, *applies a first portion of the set of pixel groupings to determine a characteristic classification* (e.g., shingle) *of the characteristic* (e.g., roof) *from a plurality of potential classifications of the characteristic* (e.g., metal, tile, shingle, slate).

82. Gross's Classifier Engine 150 is *a machine learning classifier trained to identify at least a portion of the one or more features from the first portion of the set of pixel groupings*. In particular, Gross discloses its Classifier Engine is trained "with the template consisting of the training set of reference images 144 and tags," and that "[o]bjects (and their reduced feature set) can be analyzed in image patterns using a variety of techniques, including statistical processing, neural networks, etc., which can be used to detect and extract features." EX1005, ¶77. Gross, through the incorporated Jafri material, also describes using "a hybrid neural network solution" (i.e., a form of classifier) to analyze image patterns. EX1008, 49. As I explained in Section VII.B, a neural network is *a machine learning classifier* that was well-known and widely used for image classification tasks by 2016.

83. Gross also discloses that Classifier Engine analyzes images to extract

"smaller distinct features that can be readily detected." EX1005, ¶77. Gross describes that the training results in extraction of specific features *characterized* by a "distinct set of pixels, shapes … or otherwise mathematically derivable feature" (i.e., pixel groupings). EX1005, ¶78. Therefore, it is my opinion that Classifier Engine 150 is *a machine learning classifier that is trained to identify a portion of the one or more features from the first portion of pixel groupings* (features characterized by the pixel groupings). EX1005, ¶¶77, 123.

6. [1.5]: applying, by the processing circuitry, a second portion of the set of pixel groupings to a second machine learning classifier to determine a condition classification of the characteristic from a plurality of potential condition classifications of the characteristic, wherein the second machine learning classifier is trained to identify, from the second portion of the set of pixel groupings, property characteristic conditions of the characteristic having the characteristic classification.

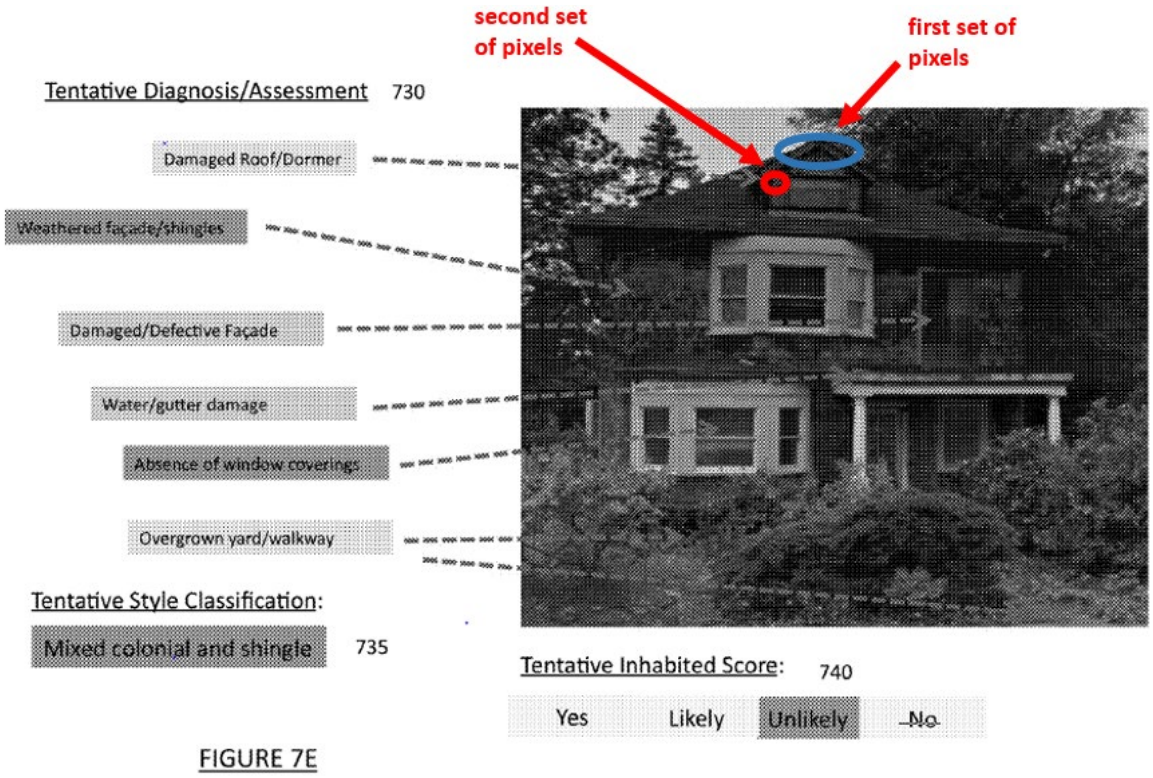
84. In my opinion, Gross teaches, or at least renders obvious, element [1.5]. Gross discloses identifying and classifying "conditions/impairments … associated with a particular building attribute" (i.e., *determines a condition classification of the characteristic*). EX1005, ¶60. Gross describes examples of structure attributes (characteristics) and related conditions (condition classifications), 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.

85. 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. To calculate this condition score, Classifier Engine (i.e., *second machine learning classifier* for reasons I explain in further detail below) "employs image processing" and each "related condition qualifier" "may be … characterized by a distinct set of pixels, shapes, sizes, proportions, curves, textures, or otherwise mathematically derivable feature from the image" (i.e., *applies pixel grouping to determine the condition classification*). EX1005, ¶78, 120.

86. A POSA would have understood or found obvious the pixel groupings that characterize the condition qualifier are *a second portion of the pixel groupings* as compared to the first portion used to classify the characteristic. For example, a POSA would have understood that the set of pixel groupings that Classifier Engine 150 relies upon to identify that the *characteristic classification* is a roof would generally be distinct from the set of pixel groupings that Classifier Engine 150 relies upon to identify that the roof (that was first identified by the first portion of pixels in element [1.4]) is in a damaged condition. This is because, for example, the pixels required to delineate a roof will be those on or near the boundary of the

roof, while the pixels required to determine the condition of the roof will tend to be those in the interior regions of the roof. I have annotated Figure 7E of Gross below to illustrate an example *first portion of the set of pixel groupings* (identified with the blue circle) that Classifier Engine may apply to identify that the characteristic is a roof and an example *second portion of the set of pixel groupings* (identified with a red circle) that may be applied by Classifier Engine to identify a damaged section of the roof. In this arrangement, Classifier Engine 150 may apply the first portion of pixels capturing an angled building element, for example, to identify that the characteristic is a roof and apply the second portion of pixels capturing a missing roof tile for example, to identify a damaged condition of the roof.



EX1005, FIG. 7E (annotated).

87. Therefore, it is my opinion that Gross teaches *applying, by the processing circuitry, a second portion of the set of pixel groupings to determine a condition classification* (e.g., missing tile) *of the characteristic* (e.g., roof) *from a plurality of potential condition classifications* (e.g., new, missing/damage tile, holes/cracks) *of the characteristic*.

88. Additionally, for the reasons I discussed for element [1.4], Gross teaches that its Classifier Engine is *a machine learning classifier*. EX1005, ¶77. And although Gross does not explicitly disclose that its condition qualifiers are classified via a different machine learning classifier than was used to identify attributes, Gross explicitly discloses 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 to identify conditions, such as "damaged/exposed roof elements" and "broken—façade," EX1005, ¶¶85, 91, which correspond to the attributes (*characteristic of the property*) that I discuss above for Element [1.4], such as a "roof" or "façade." Gross also discloses that different classifiers 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. Gross further teaches that "[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," demonstrating that specialized training data and minimizing error is important towards improving classification accuracy. EX1005, ¶78.

89. In light of these teachings, a POSA would have understood that Gross teaches, or renders obvious, the use of *a second machine learning classifier*, as claimed, to classify the property characteristic conditions because customized classifiers are more accurate or more appropriate for different vendors. A POSA would have found it obvious to separate the classifiers and would have recognized that characteristic classification and condition assessment represent different classification tasks that would benefit from tailored training regiments by different classifiers. Characteristic classification, such as identifying between a "tile roof" and a "shingle roof," focuses on recognizing structural or material aspects (e.g., shape, texture, color, etc.). In contrast, condition classification, such as identifying between a "good condition" roof and a roof that "needs repair," focuses on recognizing indicators of degradation (e.g., weathering patterns, missing elements, damage artifacts, etc.).

90. A POSA would been motivated to employ a second machine learning classifier to identify a condition classification because Gross explicitly suggests that different classifiers 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 different algorithms" and "additional data" that optimize and focus on identification of property characteristic conditions, which would improve accuracy of the classification. In addition, a POSA would have understood that separating these classification tasks allows for independent optimization of each classifier's training parameters and training datasets. Doing so would lead to increased overall system accuracy compared to a single classifier attempting to perform two types of classification simultaneously. A POSA could have readily implemented the predictable variation of implementing Gross's Classifier Engine as two machine learning classifiers—a first classifier trained to identify features that represent property characteristics and a second classifier trained to identify conditions of those characteristics. Therefore, it is my opinion that Gross teaches, or at least renders obvious, *a second machine learning classifier trained to identify* (using image processing techniques), *from the second portion of the set of pixel groupings, property characteristic conditions* (e.g., missing tile) *of the characteristic* (e.g., roof) *having the characteristic classification* (e.g., "tile roof"). EX1005, ¶¶60, 61, 123, 126.

C. Independent Claim 9

1. [9.P]: A system for automatically categorizing a condition of a property feature of a property, the system comprising:

91. To the extent the preamble is determined to be limiting, it is my

opinion that Gross discloses the preamble of claim 9 for the reasons I discuss above for the preamble of claim 1 (element [1.P]). Gross further describes that its property assessment computing system may be "an online collection of computing machines and accompanying software modules suitable and configured particularly for performing the operations described," i.e., *a system*. EX1005, ¶42.

2. [9.1]: a non-transitory computer readable storage region storing a first machine learning analysis model trained to identify one or more property characteristics, and

92. In my opinion, Gross teaches element [9.1] for the reasons I discuss above for element [1.4]. Gross's Classifier Engine, which is trained and uses neural networks and a "hybrid neural network solution" to analyze images, is a *machine learning analysis model*. EX1005, ¶¶77, 79; EX1008, 49. Gross further discloses that its "applications," which would include its Classifier Engine, "can be embodied in tangible, machine readable form for causing a computing system to execute appropriate operations in accordance" with the teachings, i.e., *a non-transitory computer readable storage region*. EX1005, ¶56.

3. [9.2]: a second machine learning analysis model trained to identify one or more property conditions;

93. In my opinion, Gross teaches element [9.2] for the reasons I discuss above for element [1.5] and [9.1].

4. [9.3]: 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 extract a set of image-related features from an aerial image, wherein the set of image-related features represent the property feature;

94. In my opinion, Gross teaches, or renders obvious, element [9.3] for the reasons I discuss above for element [1.1] and [1.2]. The *image-related features* of this element correspond to the "features represented by a set of pixel groupings" in [1.2] and *property feature* corresponds to *the characteristic of the property* in [1.2]. 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, i.e., *a non-transitory computer readable medium having instructions stored thereon*. EX1005, ¶56.

5. [9.4]: apply a first portion of the set of image-related features to the first machine learning analysis model to determine a characteristic classification of the property feature from a plurality of potential characteristic classifications of the property feature; and

95. In my opinion, Gross teaches element [9.4] for the reasons I discuss above for element [1.4] and [9.1].

6. [9.5]: apply a second portion of the image-related features to the second machine learning analysis model to determine a condition classification of the property feature from a plurality of potential condition classifications of the property feature, wherein the second machine learning analysis model is trained to identify, from the second portion of the set of image-related features, property characteristic conditions of the property feature having the characteristic classification.

96. In my opinion, Gross teaches, or renders obvious, element [9.5] for the reasons I discuss above for element [1.5] and [9.2].

D. Independent Claim 16

1. [16.P]: A non-transitory computer readable medium having instructions stored thereon, wherein the instructions, when executed by processing circuitry, cause the processing circuitry to:

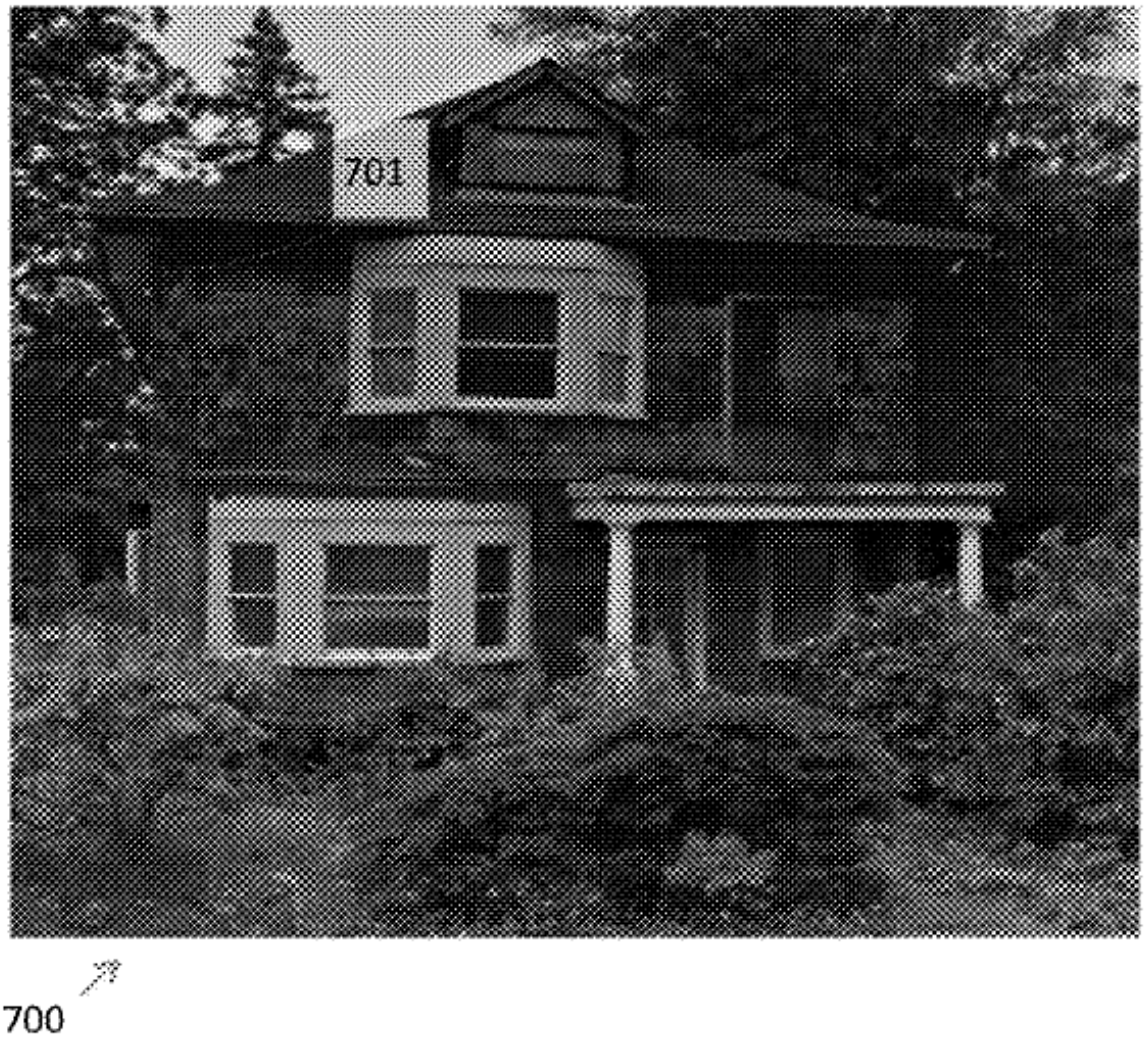
97. To the extent the preamble is determined to be limiting, it is my opinion that Gross discloses the preamble of claim 16. In particular, Gross teaches that its "applications can be embodied in tangible, machine readable form for causing a computing system to execute appropriate operations in accordance" with the teachings, i.e., *a non-transitory computer readable medium having instructions stored thereon, wherein the instructions, when executed by processing circuitry, cause the processing circuitry to* perform the operations discussed as follows. EX1005, ¶56.

2. [16.1]: obtain aerial and/or terrestrial imagery including a property;

98. For the reasons I discuss above for element [1.1], it is my opinion that

Gross teaches instructions that cause the processing circuitry to obtain aerial imagery including a property.

99. I understand from counsel that this claim element is satisfied if the processing circuitry obtains *either* aerial *or* terrestrial imagery. Although not required by the claim, Gross also teaches obtaining *terrestrial images*. Gross discloses images may be obtained directly from a user's camera and the "user's image data of the property can easily be broken down and analyzed for attributes, features, etc. by Classifier Engine [] in a similar manner as discussed above, even if it takes more time." EX1005, ¶¶108, 109. Figure 7A of Gross (depicted below), for example, illustrates an example image taken from the perspective of a ground location (*a terrestrial image*) that Gross's system would be "retrieved from database 144 (or derived from [an]other source)." EX1005, ¶121.



EX1005, FIG. 7A.

100. These user-captured images are *terrestrial imagery* because they are images of the property taken from a location on the ground. Therefore, Gross further teaches *instructions that cause the processing circuitry to obtain terrestrial imagery including a property*.

3. [16.2]: extract a plurality of pixel groupings from the aerial and/or terrestrial imagery, wherein the plurality of pixel groupings includes a respective set of pixel groupings

representative of each feature of one or more features of the property, wherein

101. In my opinion, Gross teaches element [16.2] for the reasons I discuss above for element [1.2] and [16.1]. As I described in element [1.2], Gross discloses an image is obtained and "an identification is performed to determine structural elements presented." EX1005, ¶¶120, 123. The image may be an "aerial image" as I discussed in element [1.2] or the "terrestrial image" obtained from a user camera as I described in element [16.1]. After the image (aerial and/or terrestrial) is obtained, the operation to *extract a plurality of pixel groupings* is performed, for the reasons I discuss above in element [1.2].

4. [16.3]: each pixel grouping of the set of pixel groupings comprises at least one of angles, outlines, or substantially homogenous pixel fields; and

102. In my opinion, Gross teaches element [16.3] for the reasons I discuss above for element [1.3].

5. [16.4]: determine at least one characteristic and at least one condition of the property, wherein determining comprises, for each feature of the one or more features of the property,

103. In my opinion, Gross teaches element [16.4] for the reasons I discuss above for element [1.4] and [1.5].

6. [16.5]: apply at least a portion of the one or more sets of pixel groupings corresponding to the respective feature to one or more first machine learning classifiers trained to identify one or more characteristics of the respective feature, and

104. In my opinion, Gross teaches element [16.5] for the reasons I discuss

above for element [1.4].

7. [16.6]: apply at least a portion of the one or more sets of pixel groupings corresponding to the respective feature to one or more second machine learning classifiers trained to identify a condition of the respective feature in view of at least a first characteristic of the one or more characteristics.

105. In my opinion, Gross teaches, or renders obvious, element [16.6] for the reasons I discuss above for element [1.5].

E. Claim 2

106. Claim 2 depends from claim 1 and further recites, "*comparing, by the processing circuitry, the condition classification to a historic condition classification to identify a change in condition of the characteristic of the property*."

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 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 that "the coding process," which 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, etc." EX1005, ¶¶146, 134. Given the recorded "structure assessment data," a POSA would have understood or found obvious that to identify these "structural improvements," a comparison is made 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., an "improvement") or otherwise.

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

processing circuitry, the condition classification (e.g., "good" condition roof) *to a historic condition classification* (e.g., "damaged" roof) *to identify a change in condition of the characteristic of the property* (e.g., "new" roof).

F. Claim 3

110. Claim 3 depends from claim 1 and further recites, "*preparing, by the processing circuitry for presentation to a user, a graphical report comprising condition information regarding the condition classificatio*n."

111. In my opinion, Gross teaches this claim. Specifically, Gross discloses that a "visual assessment report 730 can be generated [] which preferably identifies at least those attributes identified by the system as having some measure of damage, impairment, aging, weathering, etc." EX1005, ¶130. 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., *condition information regarding the condition classification*) EX1005, ¶103. Figure 7E (depicted below) presents an embodiment of this visual assessment report.

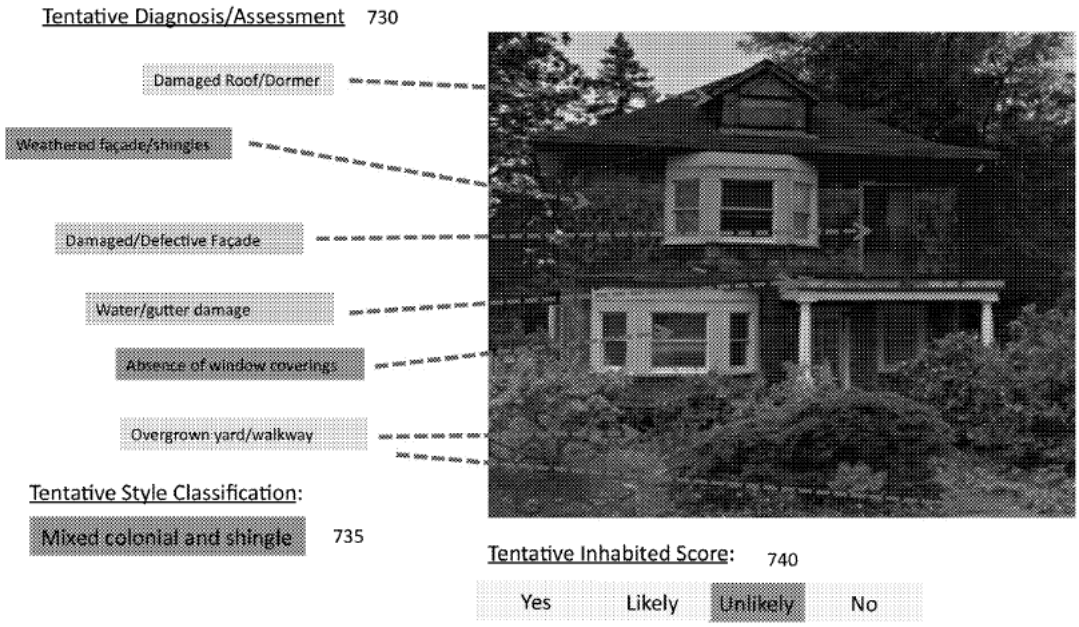


FIGURE 7E

EX1005, FIG. 7E.

112. Gross further discloses that "[t]his data is presented in a list and visual form for a human reviewer." EX1005, ¶130. And Gross describes its system 110 engages with users employing computing devices 112 and, in a preferred embodiment, system 110 is part of a website which can be accessed through a conventional browser. EX1005, ¶43. The '285 patent does not provide any details about what is required of a "graphical report." In light of Gross's disclosure of a website interface, a POSA would have understood or found obvious that a "visual assessment report" is a *graphical report* that is *for presentation to a user* because it provides assessment report data in graphical form (i.e., in a graphical user interface, such as a website).

113. Gross also discloses that "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*.

114. Therefore, it is my opinion that Gross teaches *preparing, by the processing circuitry for presentation to a user, a graphical report* (e.g., visual assessment report 730) *comprising condition information regarding the condition classification* (e.g., "damaged" roof).

G. Claim 4

115. Claim 4 depends from claim 3 and further recites, "*wherein the preparing comprises preparing the graphical report as an interactive user interface*."

116. In my opinion, Gross teaches this claim. Gross discloses that the assessment data in its "visual assessment report" is prepared as an "interactive user interface." Specifically, "visual assessment report 730 [] is presented in a list and visual form for a human reviewer, and within an interface may have *interactive* tags so that the tagged elements/conditions may be further reviewed, modified, etc. by the human operator." EX1005, ¶130 (emphasis added). By "reviewed, modified, etc. by the human operator," Gross teaches that a user can interact with *the graphical report*. EX1005, ¶130.

H. Claim 5

117. Claim 5 depends from claim 1 and further recites, "*wherein at least one of the first machine learning classifier or the second machine learning classifier comprises a convolutional neural network (CNN) model*."

118. In my opinion, Gross teaches, or renders obvious, this claim. Gross discloses that its feature extraction can be accomplished "using a variety of techniques, including … neural networks, etc." EX1005, ¶77. Gross further discloses that these "neural networks" can be "convolutional neural network models." Specifically, Gross describes that suitable examples of well-known feature extraction algorithms that can be adapted for use in its embodiments are set forth in Jafri, which I understand from counsel is incorporated by reference into Gross. EX1005, ¶79. Gross, through the incorporated Jafri material, expressly describes using "a hybrid neural network solution which combines [] a convolutional neural network." EX1008, 49. Thus, a POSA would have understood or found obvious that the disclosed "neural networks" included a *convolutional neural network (CNN) model*, which, regardless, was a well-known neural network model for image processing, as I explained above in Section VII.B.

I. Claim 6

119. Claim 6 depends from claim 1 and further recites, "*wherein the plurality of potential condition classifications comprises a numeric scale of*

condition classifications ranging between a needs replacement condition and a like new condition."

120. In my opinion, Gross teaches, or renders obvious, this claim. For instance, Gross discloses that "[t]he attributes of reference property 830 [] can be classified as corresponding to a positive or high end of a condition scale, i.e. being in top condition and lacking any noticeable weathering or aging," i.e., *like new condition*. EX1005, ¶135. Gross discloses that this condition scale can also be a numeric scale, such as 1–10. For example, Gross discloses "severity weightings (i.e., a scaling factor of any convenient range (1-10, or the like) for the identified condition)" and "types of conditions are associated with a particular feature across a spectrum of classified values–i.e., what constitutes a poor condition shingle (level 1) an excellent condition shingle (level 10) and so on". EX1005, ¶¶70, 179. A POSA would have understood or found obvious that the level 10 "top" condition would be a "like new" condition for the property characteristic and that a level 1 "poor" condition would be a "needs replacement" condition for the property characteristic.

J. Claim 7

121. Claim 7 depends from claim 1 and further recites, "*determining, by the processing circuitry based in part on the characteristic classification and the condition classification, a cost for repair or replacement of the characteristic of*

the property."

122. In my opinion, Gross teaches, or renders obvious, this claim. For instance, Gross discloses generating "an overall estimate [that] identif[ies] a cost to restore the building structure to a nominal target state," which is calculated by "correlating each of the impairments to repair or improvement figures, and summing over all the attribute conditions." EX1005, ¶134. These "improvement figures" teach, or at least render obvious, *a cost for repair or replacement of the characteristic of the property* (e.g., attribute, such as a roof) because it was both known and is common sense that one way to improve the attribute is to replace the attribute. Gross teaches its determination of the replacement cost uses: *the condition classification* ("summing over all the attribute *conditions*"); and *the characteristic classification* because repair and improvement figures depend on the classification (e.g., a tile roof replacement has a different cost than a slate roof). EX1005, ¶134.

K. Claim 10

123. Claim 10 depends from claim 9 and further recites, "*wherein each image-related feature of the set of image-related features comprises at least one of angles, outlines, or substantially homogenous pixel fields*."

124. In my opinion, Gross teaches this claim for at least the reasons I discuss above for element [1.3].

L. Claim 11

125. Claim 11 depends from claim 9 and further recites, "*wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to, prior to extracting the set of image-related features, assess orthogonality of the aerial image*."

126. In my opinion, Gross teaches, or renders obvious, this claim. Gross teaches *assess[ing] orthogonality of the aerial image*. Orthogonality of an image refers to the degree to which a camera's optical axis is perpendicular to the imaged surface (e.g., its degree of alignment with the true vertical axis of the captured image). Gross describes how various tools can be used for assessing buildings via satellite aerial inspections of roof, such as the techniques set forth in U.S. Patent No. 8,078,436 to Pershing et al. (EX1009), which I understand from counsel is incorporated by reference into Gross. EX1005, ¶6. Gross, through the incorporated Pershing material, describes a system that assesses orthogonal characteristics of an aerial image, specifically addressing the technical problem that arises when aerial images are captured at various "distances from the building, or at a different pitch, roll, or yaw angles of the aircraft." EX1009, 8:59–64.

127. The incorporated Pershing material explains that such assessment involves determining whether "the vertical axis of a particular aerial image sometimes will not substantially match the actual vertical axis of its scene" and

then implementing corrective measures to "adjust the scale and/or relative angle of the aerial image to correct for such errors." EX1009, 8:59–9:1. In other words, this orthogonality assessment process evaluates the geometric relationship between the captured image orientation and the true ground reference frame to identify when correction is needed, for example, to achieve proper orthophoto form.

128. Gross also teaches assessing orthogonality prior to training Classifier Engine. In particular, Gross discloses that its "customized image processing operations (rotations, noise clean up, etc.) may also be performed" as part of an image pre-processing or "normaliz[ation]" process. EX1005, ¶76. A POSA would have found obvious that these "customized image processing operations"—particularly in view of the incorporated disclosure of Pershing—would have included correcting orthogonality, because normalizing the orthogonality of reference images would improve the accuracy of Classifier Engine and reduce errors that occur, for example, when "the vertical axis of a particular aerial image [does] not substantially match the actual vertical axis of its scene. EX1009, 8:59–9:1. Such implementation would amount to combining Gross's disclosure of pre-processing techniques prior to training Classifier Engine with its disclosure of orthogonality assessment according to known methods to yield predictable results. Additionally, since Gross's pre-processing process occurs "prior to training" (and

using) Classifier Engine, the orthogonality assessment would also occur *prior to extracting the set of image-related features*. EX1005, ¶76.

M. Claim 12

129. Claim 12 depends from claim 11 and further recites, "*wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to, responsive to assessing that the orthogonality of the aerial image does not correspond to a true orthophoto form, cause correction of the aerial image to the true orthophoto form*."

130. In my opinion, Gross teaches, or renders obvious, this claim for at least the reasons provided for Claim 11. Gross teaches *caus[ing] correction of the aerial image to the true orthophoto form*. As described above, orthogonality of an image refers to the degree to which a camera's optical axis is perpendicular to the imaged surface, (e.g., its degree of alignment with the true vertical axis of the captured image). The '285 patent does not provide a definition, or describe, "true orthophoto form." A POSA would understand an image in *true orthophoto form* is one where the camera's optical axis is geometrically corrected to perfectly (or near-perfectly) align with the true vertical axis of the captured image. This understanding of *true orthophoto form* is consistent with how the phrase is used in the '285 patent. For example, the '285 patent describes "generat[ing] a true orthophoto form" of an image "if the aerial imagery is not orthogonal" or to

"compensate for the angle of capture of the aerial image." EX1001, 20:50–54, 20:57–62.

131. As I explained above, Gross, through the incorporated Pershing material, describes a system that assesses orthogonal characteristics of an aerial image, such as when "the vertical axis of a particular aerial image sometimes will not substantially match the actual vertical axis of its scene," and "adjust[s] the scale and/or relative angle of the aerial image to correct for such errors." EX1009, 8:59–9:1. In light of these teachings, a POSA would have understood or found obvious that correcting the orthogonality of an aerial image would produce a *true orthophoto form* of the original image.

N. Claim 13

132. Claim 13 depends from claim 9 and further recites, "*wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to obtain, via a network from a remote computing system, the aerial image of the geographic area*."

133. In my opinion, Gross teaches, or renders obvious, this claim for at least the reasons I discuss above for element [1.1]. Additionally, Gross teaches *obtain*[ing], the images *via a network from a remote computing system*. For instance, Gross discloses that "image data [] is obtained from a third party vendor" (i.e., *remote computing system*) and that "the software modules described [] can be

implemented from the present teachings using [] large scale computing systems, including servers connected to a network (such as the Internet)." EX1005, ¶¶164, 56.

134. Gross, through the incorporated Pershing material, also discloses:

> [A] roof estimation service is contacted by a potential customer requesting an estimate for repair or replacement of a roof on a building. The roof estimation service uses a local computer with an estimating software program loaded into its working memory **to access an image file database loaded on the computer or on a remote server connected via a wide area network** to the local computer. The image file database contains image files of various buildings. When a request for an estimate is received from a potential customer, the roof estimation service enters the customer's address into the software program and **aerial images of the building** are then presented to the roof estimation service.

EX1009, 2:13–25 (emphasis added).

135. In view of these teachings that Gross is implemented using large scale computing systems and servers connected to a network, it would have been obvious that obtaining image data from a third party vendor (the *remote computing system*) would involve obtaining the image data *via a network* (e.g., the "Internet," a "wide area network"). EX1005, ¶56; EX1009, 2:19.

O. Claim 14

136. Claim 14 depends from claim 13 and further recites, "*wherein the*

instructions, when executed by the processing circuitry, cause the processing circuitry to, prior to obtaining the aerial image, determine, based at least in part on the property feature, a preferred image type corresponding to the property feature."

137. In my opinion, Gross teaches, or renders obvious, this claim. Gross discloses "a data acquisition process 1600 … for acquir[ing] appropriate image data for structures within a target area." EX1005, ¶161. This "data acquisition process" involves retrieving "image data" corresponding to a specific property. EX1005, ¶164. Gross discloses that when retrieving this "image data," "it is expected that aerial drones, satellite, balloon and similar technology can be used in certain areas to easily capture structure image data from a variety of perspectives, and at different times," and that "this will also 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 feature currency/accuracy." EX1005, ¶148.

138. In light of these teachings, a POSA would have understood or found obvious that this "comprehensive image database" that contained different types of images for a building would allow the classifier to use certain, preferred types of images for certain property classification tasks with the goal of increased classification "accuracy." EX1005, ¶148. For example, an image that would

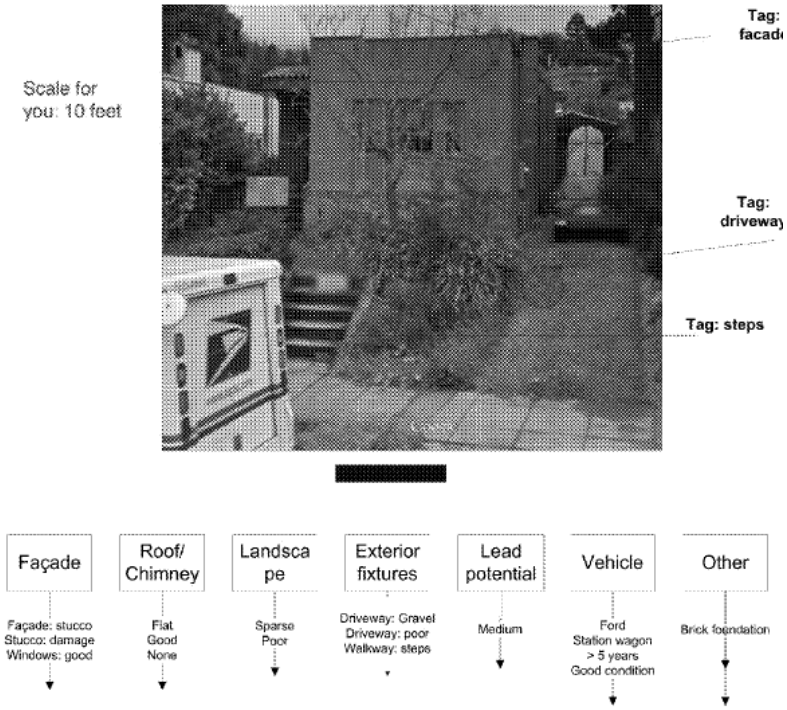
"increase building coverage" of the roof of a property and therefore increase the "accuracy" of the roof condition classification would be a top-down image of the property that exposes a higher amount of rooftop surface area, for example, as opposed to an image of the property taken at street elevation. EX1005, ¶148. A POSA could have readily implemented the predictable variation of determining the preferred type of image (that would provide increased "accuracy") based on the specific property feature being classified. Additionally, since such a determination influences which image type to obtain—and consequently the actual image—the determination would also occur *prior to obtaining the aerial image*.

139. Therefore, it is my opinion that Gross teaches *instructions* (e.g., instructions on "the computing system") that *cause the processing circuitry to, prior to obtaining the aerial image of the geographic region, determine*, *based at least in part on the property feature* (e.g., roof), *a preferred image type corresponding to the feature* (e.g., a top-down image, an image that "increase[s] feature currency/accuracy," etc.).

P. Claim 15

140. Claim 15 depends from claim 9 and further recites, "*wherein the plurality of potential characteristic classifications comprises at least one of a plurality of potential sizes, a plurality of potential shapes, or a plurality of potential materials of the property feature*."

141. In my opinion, Gross teaches this claim. For instance, Figure 12 of Gross depicts examples of "structural features, parameters, conditions, etc. that can be identified, assessed, tagged, coded and stored," which include roof shapes such as a "flat" roof. EX1005, ¶¶21, 146.



EX1005, FIG. 12.

142. Gross also discloses additional attributes, "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 of "specific types" corresponding to each attribute include *a plurality of potential shapes* (e.g.,

pitched, hipped, flat) and *a plurality of potential materials* (e.g., metal, tile, shingle).

Q. Claim 17

143. Claim 17 depends from claim 16 and further recites, "*using the at least one characteristic and the at least one condition to determine one or more of a risk estimate, a repair estimate, or a replacement estimate*."

144. In my opinion, Gross teaches, or renders obvious, this claim for at least the reasons I discuss above for claim 7. As explained, Gross discloses using the property attribute and the attribute condition to determine a cost for repair or replacement. *See also* EX1005, ¶134.

R. Claim 18

1. [18.1]: The non-transitory computer readable medium of claim 17, wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to access one or more known property characteristics,

145. In my opinion, Gross teaches, or renders obvious, this element for at least the reasons I discuss above for claims 7 and 17. Additionally, Gross discloses generating "a score … identifying a likelihood of current occupancy of the structure," which "may be associated or mapped to other content labels, such as "confirmed vacant" or "confirmed occupied," etc." EX1005, ¶118. Gross also describes "where this information is already known with certainty by reference to public records, it can be included" in a record generated by the Classifier Engine.

EX1005, ¶¶118, 115. It would have been obvious that obtaining occupancy information from "public records" would require accessing, at least in part, the "public records" for the occupancy information. Here, the property occupancy information is the *known property characteristic*.

2. [18.2]: wherein determining the one or more of the damage risk estimate, the repair estimate, or the replacement estimate comprises applying, to the determining, the one or more known property characteristics.

146. In my opinion, Gross teaches, or renders obvious, this element for at least the reasons I discuss above for claims 7 and 17 and element [18.1]. Additionally, Gross discloses determining "policy premiums, risk assessments, etc. based on an evaluation of an upkeep/maintenance evidenced for a particular property," and developing correlations for "occupancy estimates." EX1005, ¶238. For example, Gross discloses that "a property insurer is likely to be interested in knowing if a building is vacant and thus more likely to be vandalized or have a higher risk of arson, etc." EX1005, ¶238. As I described above, property occupancy is an example of a *known property characteristic* that can be accessed by Gross's system. EX1005, ¶¶118, 115. A POSA would have understood or found obvious that developing a risk estimate correlation value from property occupancy would involve *applying* a *known property characteristic* when determining a risk estimate, because a risk estimate correlation value calculation involves mathematically modeling (*applying*) the occupancy information accessed from

"public records." A POSA would have also understood or found obvious that in addition to the applying the "property occupancy" characteristic disclosed by Gross, the insurance risk assessment taught by Gross would have also been based, at least in part, on applying other property characteristic classification because it is common sense the risk varies depending on the classification (e.g., a metal roof would have less insurance risk of damage from a fire than a slate roof).

S. Claim 19

147. Claim 19 depends from claim 16 and further recites, "*wherein the one or more characteristics of a first feature of the one or more features comprises a shape characteristic and a material characteristic*."

148. In my opinion, Gross teaches this claim for at least the reasons I discuss above for claim 15. As explained, the Classifier Engine in Gross is trained to identify and assess "structural features, parameters, conditions, etc." such as roof shape and material. See also EX1005, ¶57, FIG. 12.

T. Claim 20

149. Claim 20 depends from claim 16 and further recites, "*wherein extracting the plurality of pixel groupings from the aerial and/or terrestrial imagery comprises extracting a first set of pixel groupings from an aerial image and a second set of pixel groupings from a terrestrial image*."

150. In my opinion, Gross teaches, or renders obvious, this claim. For at

least the reasons I discuss above for elements [1.1] and [1.2], Gross discloses *extracting a first set of pixel groupings from an aerial image*. As I explained above, pixel groupings that represent image features can be extracted from aerial property images captured via "aerial drones, satellite, balloon, and similar technology." EX1005, ¶164. Gross also teaches *extracting a second set of pixel groupings from a terrestrial image* for at least the reasons provided for elements [16.1] and [16.2]. For instance, Gross describes how terrestrial images may be obtained directly from a user's camera and "be broken down and analyzed for attributes, features, etc. by Classifier Engine … in a similar manner as discussed above, even if it takes more time." EX1005, ¶¶108, 109.

151. Gross further teaches, or at least renders obvious, using both types of images in identifying and classifying property features. Gross explains that obtaining images "at different angles, elevations and profiles" (e.g., aerial images and terrestrial images) helps "increase building coverage and feature currency/accuracy." EX1005, ¶148; *see also* EX1005, ¶164 (describing how "image data from a variety of perspectives … will also facilitate building out a comprehensive image database."). For example, an aerial image of a property would "increase building coverage" of a roof of a property, for example, as opposed to a terrestrial image captured at street level, because the aerial image could expose a larger amount of rooftop surface area than the terrestrial image.

Similarly, a terrestrial image would "increase building coverage" of a window of a property, for example, as opposed to an aerial image, because the terrestrial image could contain a higher number of pixels that correspond to the window.

152. Accordingly, in light of Gross's teachings that multiple property characteristics are assessed (e.g., roofs, windows) and images at different angles/elevations/profiles increase building coverage and feature accuracy, a POSA would have understood or found obvious to extract *a first set of pixel groupings* from an *aerial image* (e.g., a group of pixels corresponding to a rooftop) and to extract *a second set of pixel groupings* from an *terrestrial image* (e.g., a group of pixels corresponding to a window), because higher "accuracy" of the respective property characteristic condition classifications would be obtained as a result.

U. Claim 21

153. Claim 21 depends from claim 16 and further recites, "*wherein a first feature of the one or more features comprises a rooftop of a house, and a second feature of the one or more features comprises a structure adjacent to the house*."

154. In my opinion, Gross teaches this claim. Gross discloses *a first feature of the one or more features* (of the property) *comprises a rooftop of a house*. For instance, Gross discloses that during the building attribute/condition assessment process, "an image 700 is retrieved from database 144 … and is divided into

distinct regions or blocks 701." EX1005, ¶121. Gross also discloses that a "building envelope" is then identified to determine which blocks should be classified. EX1005, ¶122. These blocks can include a rooftop of a house. *See* EX1005, ¶123 ("an identification is performed to determine structural elements presented in a first block 701, including such attributes as roof elements").

155. Further, Gross discloses *a second feature of the one or more features comprises a structure adjacent to the house*. In addition to a "building envelope," Gross discloses that "[o]ther boundaries for other objects such as landscaping 706, trees 708, etc., can also be determined and logged." EX1005, ¶122. The image blocks within these determined "boundaries" would include *a structure adjacent to the house*, such as natural structures like landscaping and trees, or ancillary structures. EX1005, ¶122. Gross, for example, describes additional natural and man-made structures adjacent to a house that may be included as features, such as a "fence, wall[,] walkway[,] garage, carport, mailbox, chainlink[,] [t]rees, shrubs, hedge, grass, debris (mail, newspapers)," and more. EX1005, ¶¶67–70. Gross also discloses identifying "trees that are too close or overgrown [and] dilapidated ancillary structures adjacent to a structure" EX1005, ¶238.

V. Claim 22

156. Claim 22 depends from claim 16 and further recites, "*wherein the aerial and/or terrestrial imagery comprises a three-dimensional (3D) aerial*

image."

157. In my opinion, Gross teaches, or renders obvious, this claim. The '285 patent does describe what a *three-dimensional (3D) aerial image* refers to in any detail. The '285 patent describes that "[t]hree-dimensional images, in comparison [to two-dimensional images], can be used to determine location coordinates, street name, construction type, occupancy type, year built, building height, soft stories, number of stories, roof condition, roof shape, roof covering, roof anchors, roof equipment, cladding, and pounding." EX1001, 7:25–30. This characterization is broad and does not meaningfully limit what a *three-dimensional (3D) aerial image* could refer to. For example, a paper map "can be used to determine location coordinates" of a building.

158. Nevertheless, based on this description, Gross teaches, or at least renders obvious, this element. As explained, Gross discloses obtaining aerial images via "aerial drones, satellite, balloon, and similar technology." EX1005, ¶164. From such images, Gross discloses that "[r]elative sizes and areas of objects in the image data can also be collected [] such as "a size of a roof area, amount of landscaping, height/size of trees, hedges, a size of a driveway, size of window openings, sidewalks, etc." EX1005, ¶143. Gross also discloses that "[a]n amount of height displacement from street level can also be measured if desired." EX1005, ¶143. Gross yet further discloses capturing a "street address" or "[t]he number of

distinct visible stories" of a property from such images. EX1005, ¶140. These are the details the '285 patent describes can be determined from a three-dimensional image.

159. In addition, to the extent the claimed *three-dimensional aerial image* refers to a three-dimensional (3D) model of a property, Gross renders obvious this element. For instance, various aerial inspection methods are set forth in Pershing (EX1009), which as noted above I understand from counsel is incorporated by reference into Gross's disclosure. EX1005, ¶6. Gross recognizes that Pershing describes aerial inspections of roofs. EX1005, ¶6. And Pershing describes building a 3D model as part of its roof assessment process. EX1009, 6:14-17, 6:56-67, 7:12-18.

160. Gross, through the incorporated Pershing material, describes a system that obtains a 3D model of a building, e.g., via "photographs taken from two or more view points to 'triangulate' points of interest on the object in three-dimensional ('3D') space [] to build a 3D model of the structure." EX1009, 6:56–67. The 3D model described by Pershing also aligns with the '285 patent's description of three-dimensional images, which "can be used to determine location coordinates, street name, construction type, occupancy type, year built, building height, soft stories, number of stories, roof condition, roof shape, roof covering, roof anchors, roof equipment, cladding, and pounding." For example, Pershing

teaches that "[i]nformation from the reconstructed model may then be used to generate a report containing information" such as "square footage and pitch of each section of roof, linear measurements of all roof segments, identification and measurement of ridges and valleys, and different elevation views rendered from the 3D model." EX1009, 7:12–18.

161. As I described above, Gross explains that obtaining images "at different angles, elevations and profiles" (e.g., aerial images and terrestrial images) helps "increase building coverage and feature currency/accuracy." EX1005, ¶148. In light of these teachings, a POSA would have found obvious to use a 3D model as the input aerial image for Gross's system, because a 3D model would provide "different angles" and "increase building coverage" as compared to a two dimensional image with one viewing angle. EX1005, ¶148. Such implementation would amount to combining Gross's disclosure of the feature extraction elements with its disclosure of obtaining 3D models for a property according to known methods to yield predictable results.

X. GROUND 2: THE COMBINATION OF GROSS AND DAVIS RENDERS OBVIOUS CLAIM 8

162. In my opinion, the combination of Gross and Davis renders obvious claim 8.

A. Claim 8

163. Claim 8 depends from claim 1 and further recites, "*determining, by*

the processing circuitry based in part on the characteristic classification and the condition classification, a risk estimate of damage to the characteristic of the property due to one or more disasters."

164. In my opinion, the combination of Gross and Davis renders obvious this claim. The '285 patent describes a "risk estimate" as "corresponding to a risk of damage." EX1001, 12:19-23. Similarly, Gross describes its property assessment system can be used for "insurance risk assessments" (i.e., *a risk estimate of damage to the property*). EX1005, ¶238. For example, a "missing/damage tile" or "holes/cracks" roof condition evidences poor upkeep and maintenance. EX1005, ¶238. Gross further discloses that "correlations may be developed between property condition ratings," that is, the *condition classifications* described above in element [1.5], and insurance metrics such as "number of claims filed, type of claim, severity, etc." EX1005, ¶238. In other words, Gross teaches, or renders obvious, risk assessments based on the condition classification (evaluation of upkeep and correlation between property condition rating and insurance claims/severity). EX1005, ¶238.

165. Although Gross does not explicitly disclose basing its risk assessment on the characteristic classification, Davis describes that risk assessment may be based on characteristic classifications. For example, 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." EX1006, 10:19-21. Davis also discloses that the rules accessed by risk analysis unit 34 may, for example, depend "on the type of roof construction." EX1006, 10:31-38. In light of Davis's teachings, a POSA would have understood or found obvious that in addition to the condition classifications (e.g., an old roof), the insurance risk assessment taught by Gross would have also been based, at least in part, on the property *characteristic classification* because it is common sense that the risk varies depending on the classification. For example, a metal roof would have less risk of damage for certain types of disasters than a slate roof.

166. Gross also teaches these risk assessments can be risk estimates *due to one or more potential disasters*, such as due to "arson" (i.e., *a disaster*) or "[o]ther potential hazards (trees that are too close or overgrown, dilapidated ancillary structures adjacent to a structure, [and] undesirable and dangerous fixtures." EX1005, ¶238; *see also, e.g.*, EX1005, ¶63 (listing "water damage" as an example condition assessment). Gross also explains how "[o]ther potential hazards … can be identified by insurers and used to adjust premiums on a structure by structure basis." EX1005, ¶238. Gross further describes prior art systems that assess "a risk of damage to a house as a result of physical disturbance such as wind, earthquakes, etc." EX1005, ¶6.

167. Davis also teaches risk assessments based on "disasters." Davis

teaches "analyzing, by one or more processors, the one or more features of the property to determine a risk of damage to a structure or object located on the property." EX1006, 2:7–14. For example, features of property contained in aerial images may be analyzed "to determine whether particular homes are at high risk of fire damage." EX1006, 11:18-25. 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.

168. Risk analysis unit 34 "may generate different kinds of risk information based upon the analyzed feature(s) and other information" (i.e., *characteristic classification* and *condition classification*) and "may estimate a likelihood that damage will occur to some portion of the property." EX1006, 11:30-34. Davis teaches its risk assessment may be based on whether the property contains trees of a species (i.e., tree species of a particular *characteristic classification*) more likely to catch fire (i.e., a *disaster*) or whether property contains dead or dying trees (i.e., *condition classification*). *See* EX1006, 11:25-36. Davis also discloses other types of information may be used to *determine risk estimate of damage to the property*, such as damage based on a hurricane forecast (i.e., another type of *disaster*). EX1006, 4:9-18. Thus, Davis more explicitly discloses determining risk assessment of damage to property based on disasters

(e.g., fires, hurricanes).

169. 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 determining the risk of property damage due to disasters because the combined Gross-Davis system would provide more information to the insurer useful to determine the risk of property damage. 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 because both systems similarly use image analysis to analyze features of properties to determine insurance risk.

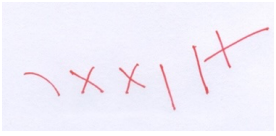
170. It is my opinion that the Gross-Davis combination would, therefore, teach *determining by the processing circuitry* (e.g., Gross's Classifier Engine 150) *based in part on the characteristic classification* (e.g., "tile" roof) *and the condition classification* (e.g., "new"), *a risk estimate of damage to the property due to one or more disasters* (e.g., risk of damage to property due to fire).

XI. CONCLUSION

171. 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.

172. 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.