

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 IPR2025-01358
U.S. Patent No. 10,529,029

**DECLARATION OF DR. DAVID A. FORSYTH IN SUPPORT OF
PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 10,529,029**

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

ZESTY.AI EX1003
U.S. Patent No. 10,529,029
Aon Ex. 2018, Page 1
Zesty.ai, Inc. v. Aon Re, Inc.
IPR2025-01359

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.....	10
	A. My Understanding of Claim Construction.....	10
	B. My Understanding of Obviousness	11
	C. A Person of Ordinary Skill in the Art.....	13
VI.	OVERVIEW OF THE '029 PATENT	14
	A. Summary of the '029 Patent.....	14
	B. Prosecution History.....	18
	C. Level of Ordinary Skill in the Art	19
	D. Claim Construction	19
VII.	STATE OF THE ART	20
	A. Aerial imagery analysis and property assessment techniques have been used for decades, well before the '029 patent.	20
	B. Machine learning and computer vision techniques for image analysis were well-known before the '029 patent.....	22
VIII.	OVERVIEW OF THE APPLIED REFERENCES	23
	A. Gross	23
	B. Davis	28
	C. Furukawa	29
IX.	GROUND 1: THE COMBINATION OF GROSS AND DAVIS RENDERS OBVIOUS CLAIMS 1-7 AND 15-20	30
	A. A POSA would have found it obvious Gross's techniques were intended to be combined.	30
	B. A POSA would have found it obvious to combine Gross and Davis.	33
	C. Independent Claim 1	36

1. [1.P]: A method for automatically categorizing a repair condition of a property characteristic, comprising:..... 36
2. [1.1]: receiving, from a user at a remote computing device, a request for a property condition classification, wherein the property classification request includes identification of a property and at least one property characteristic; 40
3. [1.2]: obtaining, by processing circuitry of a computing system responsive to receiving the request, an aerial image of a geographic region including the property;..... 46
4. [1.3]: extracting, by the processing circuitry, one or more of a plurality of features from the aerial image corresponding to the property characteristic, wherein the extracted features include pixel groupings representing the property characteristic;..... 47
5. [1.4]: determining, by the processing circuitry from the extracted features, a property characteristic classification for the property characteristic, wherein determining the property characteristic classification includes applying the pixel groupings for the property characteristic to a first machine learning classifier trained to identify property characteristics from a set of pixel groupings; 50
6. [1.5]: determining, by the processing circuitry based on the identified property characteristic and the extracted features, a condition classification for the property characteristic, wherein identifying the condition classification includes applying the pixel groupings for the property characteristic to a second machine learning classifier trained to identify property characteristic conditions from a set of pixel groupings; 52
7. [1.6]: determining, by the processing circuitry based in part on the property characteristic classification and the condition classification, a risk estimate of damage to the property due to one or more disasters; and 57
8. [1.7]: returning, to the user at the remote computing device via a graphical user interface responsive to receiving the request, a condition assessment of the property characteristic including the condition classification and

the risk estimate of damage to the property due to the one or more disasters.	61
D. Independent Claim 15	63
1. [15.P]: A non-transitory computer readable medium having instructions stored thereon, wherein the instructions, when executed by processing circuitry, cause the processing circuitry to:	63
2. [15.1]: receive, from a user at a remote computing device, a property condition classification request, wherein the property classification request includes identification of a property and at least one property characteristic;	63
3. [15.2]: obtain, from a remote data source responsive to receiving the property classification request, an aerial image of a geographic region including the property;	63
4. [15.3]: extract from one or more of a plurality of features from the aerial image corresponding to each property characteristic of the at least one property characteristic, wherein the extracted features include pixel groupings representing the respective property characteristic;	64
5. [15.4]: for each of the at least one property characteristic, ...	64
6. [15.5]: determine, from the extracted features for the respective property characteristic, a respective property characteristic classification, wherein determining the respective property characteristic classification includes applying the pixel groupings for the respective property characteristic to a first machine learning classifier trained to identify property characteristics from a set of pixel groupings, and	64
7. [15.6]: determine, based on the respective identified property characteristic and the respective extracted features, a respective condition classification for the respective property characteristic, wherein identifying the respective condition classification includes applying the respective pixel groupings for the respective property characteristic to a second machine learning classifier trained to identify property characteristic conditions from a set of pixel groupings; and	64

8.	[15.7]: determine, in real-time responsive to receiving the property classification request and using the property characteristic classification of each property characteristic and the condition classification of each property characteristic, at least one risk estimate representing risk of damage due to disaster.....	65
E.	Claim 2	65
F.	Claim 3	66
G.	Claim 4	67
H.	Claim 5	68
I.	Claim 6	71
J.	Claim 7	72
1.	[7.1]: The method of claim 1, wherein identifying the condition classification for the property characteristic comprises: generating, from the pixel groupings, a pixel intensity distribution of the property characteristic; and.....	72
2.	[7.2]: determining, by applying the second machine learning classifier, the condition classification based on the pixel intensity distribution of the property characteristic.....	73
K.	Claim 16	74
L.	Claim 17	76
1.	[17.1]: The non-transitory computer readable medium of claim 16, wherein: the at least one property characteristic comprises two or more property characteristics; and.....	76
2.	[17.2]: the instructions, when executed by the processing circuitry, cause the processing circuitry to, responsive to determining the preferred image type corresponding to a first property characteristic of the at least one property characteristic is a terrestrial image, obtain a terrestrial image of the geographic region including the property.	76
M.	Claim 18	78
N.	Claim 19	80
O.	Claim 20	81

1.	[20.1]: The non-transitory computer readable medium of claim 15, wherein: receiving identification of the property and the at least one property characteristic comprises receiving, via a network from a remote computing device, the identification of the property; and	81
2.	[20.2]: the instructions, when executed by the processing circuitry, cause the processing circuitry to provide, via the network to a graphical user interface at the remote computing device in real-time responsive to receiving the at least one property characteristic, the at least one risk estimate.....	82
X.	GROUND 2: GROSS RENDERS OBVIOUS CLAIMS 8 AND 11-14.	82
A.	Independent Claim 8	82
1.	[8.P]: A system for automatically categorizing a repair condition of a property characteristic, comprising:.....	82
2.	[8.1] processing circuitry; and a non-transitory computer-readable medium having instructions stored thereon; wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to	82
3.	[8.2]: receive, from a user at a remote computing device, a property condition classification request, wherein the property classification request includes identification of a property and at least one property characteristic;	83
4.	[8.3]: obtain, from a remote data source responsive to receiving the property classification request, an aerial image of a geographic region including the property;	83
5.	[8.4]: extract one or more of a plurality of features from the aerial image corresponding to the property characteristic, wherein the extracted features include pixel groupings representing the property characteristic;	84
6.	[8.5]: determine, from the extracted features, a property characteristic classification, wherein determining the property characteristic classification includes applying the pixel groupings for the property characteristic to a first	

	machine learning classifier trained to identify property characteristics from a set of pixel groupings;	84
7.	[8.6]: determine, based on the identified property characteristic and the extracted features, a condition classification for the property characteristic, wherein identifying the condition classification includes applying the pixel groupings for the property characteristic to a second machine learning classifier trained to identify property characteristic conditions from a set of pixel groupings; and	84
8.	[8.7]: determine, in real-time responsive to receiving the property classification request and using the property characteristic classification and the condition classification, a replacement cost for replacing the property characteristic.....	84
B.	Claim 11	87
C.	Claim 12	88
D.	Claim 13	90
E.	Claim 14	92
XI.	GROUND 3: THE COMBINATION OF GROSS AND FURUKAWA RENDERS OBVIOUS CLAIMS 9 AND 10.....	94
A.	Claim 9	94
1.	[9.1]: The system of claim 8, wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to, prior to identifying the one or more of the plurality of features: obtain a shape map image including the property;.....	94
2.	[9.2]: overlay the aerial image with the shape map image; and.....	102
3.	[9.3]: determine whether a boundary of the property as identified by the shape map matches a boundary of the property as illustrated in the aerial image.....	104
B.	Claim 10	105
XII.	CONCLUSION	109

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. 10,529,029 (“the ’029 patent”) (EX1001), which resulted from U.S. Patent Application No. 15/714,376 (“the ’376 application”), filed on September 25, 2017. I understand that the ’029 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 ’029 patent and each of the documents cited herein, in light of general knowledge in the art before September 23, 2016. In formulating my opinions, I have relied upon my experience, education, and knowledge in the relevant art. In formulating my opinions, I have also considered the viewpoint of a person of ordinary skill in the art (“POSA”) (i.e., a person of ordinary skill in the field of image processing and machine learning) prior to September 23, 2016.

II. MY BACKGROUND AND QUALIFICATIONS

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

III. SUMMARY OF OPINIONS

14. In forming my opinions about the '029 patent, I have considered the following grounds of unpatentability. Based on my review of the prior art references that form the basis of these grounds, it is my opinion that claims 1-20 of the '029 patent would have been obvious to a POSA prior to September 23, 2016.

Ground	Basis	Claims	References
1	§103	1-7, 15-20	Gross and Davis
2	§103	8, 11-14	Gross
3	§103	9, 10	Gross and Furukawa

15. I have been asked to consider how a POSA would have understood the challenged claims in light of the disclosures of the '029 patent. I have also been asked to consider how a POSA would have understood the prior art references that form the basis of the above-listed grounds.

16. Further, I have been asked to consider and provide my technical review, analysis, insights, and opinions regarding whether a POSA would have understood that the prior art reference listed in the table above render obvious claims 1-20 of the '029 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 '029 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. 10,529,029 to Okazaki ("’029 patent")
1002	Prosecution History of U.S. Patent No. 10,529,029 ("Prosecution History")

Exhibit No.	Description
1004	<i>Curriculum Vitae</i> of Dr. David A. Forsyth
1005	U.S. Publication No. 2015/0186953 to Gross et al. (“Gross”)
1006	U.S. Patent No. 10,755,357 to Davis et al. (“Davis”)
1007	U.S. Patent No. 6,970,593 to Furukawa (“Furukawa”)
1008	Jafri, <i>et al.</i> , “A Survey of Face Recognition Techniques,” <i>Journal of Information Processing Systems</i> , Vol. 5, No. 2, June 2009, pp. 41-68 (“Jafri”)
1009	U.S. Patent No. 8,078,436 to Pershing et al. (“Pershing”)
1010	U.S. Patent No. 6,885,771 to Takahashi (“Takahashi”)
1011	U.S. Publication No. 2009/0265193 to Collins et al. (“Collins”)
1012	U.S. Publication No. 2015/0302529 to Jagannathan (“Jagannathan”)
1013	U.S. Patent No. 4,941,122 to Weideman (“Weideman”)
1014	U.S. Patent No. 8,995,757 to Ciarcia et al. (“Ciarcia”)
1015	Redistricting Data Hub, “TIGER Boundary Files,” accessible at https://redistrictingdatahub.org/data/about-our-data/tiger-boundary-files/ (last accessed July 10, 2025)
1016	Norman, J., “Origins of Google Earth,” History of Information, accessible at https://www.historyofinformation.com/detail.php?entryid=3145 (last accessed July 13, 2025)
1017	Krizhevsky, <i>et al.</i> , “ImageNet Classification with Deep Convolutional Neural Networks,” <i>Communications of the ACM</i> , Vol. 60, No. 6, May 2017
1018	Lin, <i>et al.</i> , “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, <i>et al.</i> , “scikit-image: image processing in Python,” <i>PeerJ</i> , June 2014

Exhibit No.	Description
1021	“Cartographic Work in Japan 2000 – 2005,” United Nations Economic and Social Council, June 2005
1022	Usery, <i>et al.</i> , “Mapping Developments and GIS in the USGS, 1884-2009,” International Cartographic Association, January 2009
1024	U.S. Patent No. 8,401,222 to Thornberry et al. (“Thornberry”)
1025	U.S. Patent No. 9,158,869 to Labrie et al. (“Labrie”)

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 '029 patent. However, my opinions herein would not be affected if the '029 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 '029 PATENT

A. Summary of the '029 Patent

34. The '029 patent discloses techniques for analyzing aerial imagery to evaluate property characteristics, condition classifications, and risk estimates of damage to a property. EX1001, Abstract. For example, this might include analyzing images to determine the maintenance levels of certain property features, such as the roof. EX1001, 1:31-36, 2:24-28. 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, 2:28-35, 11:51-55 (“The types of disasters can include, in some examples, earthquake, hurricane, tornado, storm surge, straight line winds, or explosion. The types and estimated severity of disasters, in some embodiments, may depend upon the particular property location.”).

35. The ’029 patent discusses 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, 2:6–15, 2:24–28.¹ Figure 3 (annotated below) illustrates an exemplary embodiment of the ’029 patent’s alleged invention.

¹ Emphasis added throughout, unless otherwise noted.




FIG. 3

EX1001, FIG. 3 (annotated).

36. The '029 patent's system operates in a series of steps, corresponding to numeric annotations in Figure 3 above. First, the system receives a request by "a user of a particular client computing system 306" to classify a property. EX1001, 13:21–24. The request "may include at least one property identifier 340" (e.g., address, geolocation coordinates) "as well as one or more property characteristics 342" (e.g., a "type of rooftop") (step "1"). EX1001, 13:24–36, 8:3–5. The system then obtains "imagery of the property from imagery source computing system(s) 304 using an image acquisition engine 336" (step "2"). EX1001, 14:24–26.

37. Upon obtaining the image(s), the system “extracts features from the property image(s) related to the identified property characteristics 342” via a “feature identification engine 320.” EX1001, 15:27–30. “Features” (or “image-related features”) can include, for example, angles, outlines, substantially homogenous fields (step “3). EX1001, 7:64–8:1.

38. The system “classifies the property characteristic” via “a property characteristic classification engine 322,” for example, into “a grouping, type, or other characterization of a particular property feature” (step “4”). EX1001, 15:34–43. For example, the property characteristic may be a roof and the property characteristic classification may be “hipped rooftop.” EX1001, 11:32–37. Property characteristic classification engine 322 may use “one or more machine learning techniques” to classify the property characteristic. EX1001, 15:37–40.

39. The system “classifies the condition of each property characteristic” (e.g., a “good” condition or a “bad” condition) via a “characteristic condition classification engine 324” (step “5”). EX1001, 16:12–27. To classify the conditions of property characteristics, characteristic condition classification engine 324 may process the image(s) through “a machine learning classifier” (e.g., “deep learning analysis such as CNN or NIN”). EX1001, 16:32–39; 10:28–42. The system “estimate[s] risk of damage based upon one or more types of disaster” based on the property characteristic classification and the condition classification

via a “risk calculation engine 316” (step “6”). EX1001, 16:50–57. The system returns the risk estimate data to a user of client 306 via a graphical interface (step “7”). EX1001, 18:19–31.

B. Prosecution History

40. I understand that the '029 patent issued from U.S. Application No. 15/714,376 (“the '376 application”), without any substantive prior art examination. 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 '376 application received a single substantive rejection—a 35 U.S.C. § 101 rejection for being directed to an abstract idea. EX1002, 92–96. In response to the § 101 rejection, the Applicant made only minor amendments to independent claims 1, 8, and 15, adding the limitation “applying the pixel groupings for the property characteristic” to a first and second machine learning classifier. EX1002, 107–127. Applicant argued these features provided “an improvement in feature detection and image processing techniques that improve the accuracy and processing efficiency of automated building condition assessments.” EX1002, 107–127.

41. In response, the Office allowed the claims—without conducting any prior art analysis in any 35 U.S.C. §§ 102, 103 rejections. The only prior art discussion appears in the Notice of Allowance’s “Reasons for Allowance” section,

where the Examiner distinguished two references (Swanson and Downey) without having issued a single rejection against the claims of the '029 patent and without the detailed analysis typically conducted during examination. EX1002, 1048-1050. This limited prior art consideration occurred despite the robust prior art landscape that existed at the time of filing. I understand that the Office did not consider the prior art applied herein, which as I explain below teaches all claim limitations.

C. Level of Ordinary Skill in the Art

42. In my opinion, a POSA at the time of the '029 patent would have had a bachelor's degree in Computer or Electrical Engineering or equivalent discipline, and approximately two years' experience in image processing and machine learning. Additional education could substitute for professional experience.

43. I am well qualified to determine the level of ordinary skill in the art. I am very familiar with the technology of the '029 patent during the September 2016 timeframe. As mentioned above, by 2016, I had obtained a Bachelors and Masters (in Electrical Engineering) from the University of the Witwatersrand, Johannesburg and a Ph. D. (in Engineering Science) from Balliol College, Oxford and had acquired extensive experience relating to image processing and computer vision.

D. Claim Construction

44. I have been informed that the challenged claims of the '029 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 '029 patent specification.

VII. STATE OF THE ART

A. Aerial imagery analysis and property assessment techniques have been used for decades, well before the '029 patent.

45. Computerized analysis of aerial imagery for property assessment was well-established before the September 23, 2016 priority date of the '029 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.

46. Property assessors routinely accessed these types of databases to evaluate property characteristics without requiring physical site visits well before 2016. For instance, Takahashi described an “image recognition method” for analyzing aerial photographs and satellite images. EX1010, 1:14–22. In particular, Takahashi discloses an “image recognition apparatus for use in an image recognition system, for extracting from a color image the shapes of objects which are to be recognized.” EX1010, 1:14–22. Takahashi also discloses its image processing techniques are “to be applied to an image such as satellite images or aerial photograph, for example to accurately and reliably extract the shapes of specific objects such as roads, buildings, etc., from the image contents.” EX1010, 2:61–64. In addition, Takahashi discloses analysis of specific building elements from aerial photographs, such as “extracting only the shape of the building roof” EX1010, 31:50.

47. Techniques for analyzing property structures and conditions from aerial imagery were also established and well-known before the September 23, 2016 priority date of the '029 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 '029 patent.

48. Artificial neural networks were a class of machine learning techniques that was first developed in the 1940s and 1950s. Weideman, filed in January 1989, describes early neural network architectures, such as “a neural-simulating system for an image processing system [that] includes a plurality of networks arranged in a plurality of layers” for image processing applications. EX1013, Abstract. As machine learning technology developed, convolutional neural networks (CNNs) made an appearance in the 1980s and became widely adopted for image recognition by the early 2000s. For example, the AlexNet CNN, which is referenced in the '029 patent as an example of a CNN that can be used in the machine learning analysis of the disclosed embodiments (*see* EX1001, 8:11-15, 20-26), became regarded as the first widely recognized application for image classification tasks after it won the ImageNet 2012 challenge on 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 '029 patent as an example machine learning classifier that can be used in the disclosed embodiments (*see* EX1001, 1:50-55, 8:26-29),

was published in December 2013 by Lin et al. EX1018. The NIN architecture was publicly available and well-known to practitioners in machine learning for image classification tasks, offering improvements over traditional CNN approaches.

49. Automated machine learning-based extraction of features such as edges, shapes, textures, and color distributions from digital images for property analysis was also routine practice before 2016. Computer vision libraries and frameworks such as OpenCV (which first released in 2000) and scikit-image (which first released in 2009) provided standard tools for preprocessing, feature extraction, and classification of image contents. EX1019, EX1020. Using these accessible tools, machine learning classifiers were commonly applied to extract features from images for property element identification. For instance, Ciarcia describes “automated roof detection” techniques on “aerial images of [a] building” that use “an artificial intelligence (‘AI’) system” that trains “on multiple historical images for which the roof and non-roof areas are known.” EX1014, 2:66-3:36, 4:33-43, 7:16–26.

VIII. OVERVIEW OF THE APPLIED REFERENCES


A. Gross

50. Gross discloses “tools, methods, and systems which assess the condition of living structures and other appurtenant real property features.” EX1005, ¶2. Like the ’029 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 '029 patent, because it is in the same field of endeavor and reasonably pertinent to the problems of the '029 patent.

51. Gross's property assessment computing system involves training a "Classifier Engine" and using it to assess and rate different properties. EX1005, ¶57. To train the classifier engine, Gross's system first collects reference images of "building elements" (e.g., roof, tree) along with their "types" (e.g., pitched roof, double pitched roof, shingle roof). EX1005, ¶¶57–58. Gross's system then defines and associates the building elements within each reference image with a corresponding "condition," such as "new," "damaged," "missing," and more. EX1005, ¶¶60–70. Gross's system trains the Classifier Engine "using a variety of techniques, including statistical processing, neural networks, etc., which can be used to detect and extract features." EX1005, ¶¶77–78.

52. After training the Classifier Engine, Gross's system uses it to assess new images of properties. EX1005, ¶¶57, 123. These images of a property to be assessed are first retrieved from a database. EX1005, ¶121. The 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.

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




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) can also be determined. EX1005, ¶122. The image blocks include "structural elements" (e.g., a roof, tree), which are determined (alongside a "type" of the

structural element (e.g., pitched roof, shingle roof)) 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 ’029 patent, identifying the structural element and its classification “can be performed in any number of ways by Classifier Engine ... including by pattern recognition techniques and statistical image processing using models.” EX1005, ¶123.

54. Again similar to the ’029 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.

55. Like the ’029 patent, these condition determinations can then be used for assessing risk. EX1005, ¶238. For example, a “property insurer” may be interested in knowing if a building is vacant or has other potential hazards. *See* EX1005, ¶238. Based on an evaluation of the property, the system can assess risk, such as whether a property is “more likely to be vandalized or have a higher risk of arson” and notify a property insurer accordingly. EX1005, ¶238.

B. Davis

56. 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 '029 patent, because it is in the same field of endeavor and reasonably pertinent to the problems of the '029 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 '029 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.

57. 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, 18:32-56; 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” 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 ’029 patent, analyzes a risk of damage to a structure on the property (e.g., a house). EX1006, 19:6–16. Davis’s system then generates “a risk output based on the risk of damage” and transmits “[i]nformation indicative of the generated risk output” to “an employee (e.g., agent) of an insurance provider ... and/or to the policyholder.” EX1006, 20:40–42, 59–64.

C. Furukawa

58. Furukawa describes techniques for processing image data (e.g., satellite image data) and correcting errors using map data that, for example, contain “information of external shape of buildings.” EX1007, Abstract. In my opinion, Furukawa is analogous art to the ’029 patent, because it is in the same field of endeavor and reasonably pertinent to the problems of the ’029 patent. Furukawa’s system “extracts shape information of a building from images by image processing means” (e.g., via “building external shape information detecting means 6”) or “obtains the same from ... existing map data.” EX1007, 6:43–46, 32:65–33:7. Furukawa’s system then compares the extracted “building external shape information ... with the map data” to determine whether the map data and image data match. EX1007, 35:17–31, 21:66–22:22. If the map data and image

data do not match, Furukawa's system then modifies the map data to obtain an alternate, updated raster image. *See, e.g.*, EX1007, 35:50–57, 30:56–31:53.

**IX. GROUND 1: THE COMBINATION OF GROSS AND DAVIS
RENDERS OBVIOUS CLAIMS 1-7 AND 15-20**

59. In my opinion, the combination of Gross and Davis renders obvious claims 1-7 and 15-20.

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

60. As I discuss above, Gross discloses training a Classifier Engine 150 and using the trained engine 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.

61. 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.

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

63. 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), 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.

64. 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/condition).

B. A POSA would have found it obvious to combine Gross and Davis.

65. With respect to the claimed *risk estimate of damage to the property due to one or more disasters*, Gross discloses that its system and processes may be used for insurance risk assessments. EX1005, ¶238. For example, Gross discloses that “correlations may be developed between property condition ratings” and “type of claim, severity, etc.” EX1005, ¶238. These risk assessments can be risk estimates due to one or more potential disasters, such as due to “arson” 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. Gross also describes a prior art risk assessment used to assess a risk of damage to a house as a result of physical disasters, such as wind and earthquakes (EX1005, ¶6), but does not explicitly discuss use of its property image assessment system to perform risk assessments due to such natural disasters.

66. Davis teaches a risk assessment of damage to property due to “disasters.” Like Gross, Davis also describes analyzing digital images and features

of a property to determine risk of damage to a property structure. EX1006, 2:7-16.

For example, Davis discloses “receiving ... one or more digital aerial images of a property,” “processing ... the one or more digital images according to one or more image analysis techniques to determine one or more features of the property,” and “analyzing ... the one or more features of the property to determine a risk of damage to a structure or object located on the property.” EX1006, 2:7–16.

67. In particular, Davis describes analyzing features contained in an aerial image of a property “to determine whether particular homes are at high risk of fire damage, for example.” EX1006, 11:18-25. Davis also discloses determining risk of damage based on a hurricane forecast, “(e.g., whether a hurricane ... is expected).” EX1006, 4:9-18. To determine the risk of damage, 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. Thus, Davis more explicitly discloses determining risk assessment of damage to the property based on disasters (e.g., fires, hurricanes).

68. 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 overall risk of property

damage. When calculating an overall or cumulative risk of damage, each source of risk can contribute independently to the cumulative probability of property damage occurring. As such, knowing the risk of property damage from one source—a risk of property damage due to disasters—is beneficial when calculating the overall or cumulative risk of property damage.

69. Excluding this component could create a potential gap in the risk model and lead to an underestimated result. As I discussed above, Gross discloses that performing “risk assessments” may involve identifying “[o]ther potential hazards (trees that are too close or overgrown, dilapidated ancillary structures adjacent to a structure, undesired and dangerous fixtures (trampolines etc.).” EX1005, ¶238. Calculating the risk of damage from each of these hazards would factor into overall “risk assessments” for a property, and adding Davis’s disaster-specific risk consideration to Gross’s “risk assessment” would provide another data point to rely on. EX1005, ¶238.

70. This 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 and machine learning to analyze features of properties to determine insurance risk. *See, e.g.*, EX1005, ¶¶57, 120, 238; EX1006, 2:7-16.

C. Independent Claim 1

1. [1.P]: A method for automatically categorizing a repair condition of a property characteristic, comprising:

71. 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.

72. It should first be noted that the word *property* can have multiple meanings, which can affect how the claim phrase *property characteristic* is 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 '029 patent uses the word *property* in the second sense. *See, e.g.*, EX1001, 1:18–21 (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”).

73. However, the '029 patent primarily uses the term *property* in the first sense. For example, the '029 patent describes that a “property may be a single family home.” EX1001, 3:22–25. Using this meaning, the claimed *property*

characteristic would refer to an aspect or quality of a physical building or piece of real estate—i.e., a property (meaning 1, above). This is consistent with how the '029 patent exemplifies *property characteristics* or characteristics of properties:

[Property] characteristics [can include] location coordinates, address, slope, and elevation, Other characteristics include construction type, occupancy type, year built and/or year of renovation, building height, soft stories, number of stories, and floor area. Further characteristics can include roof condition, roof shape, roof covering, roof anchors, roof equipment, cladding, and pounding (distance to closest building). Some of these characteristics can only be assessed by on-site inspections or by official documentation, but others can be measured using visual imagery.

Characteristics addressed in this disclosure include roof shape and roof condition.

EX1001, 1:20–32. The '029 patent also uses the phrase “property feature” in various places, which it describes as being physical structures, “such as rooftop, swimming pool, chimney, and sky lights,” located on a property. EX1001, 7:64–8:1. And though such physical structures are not explicitly recited in the list of example *property characteristics* above, the '029 patent does also sometimes refer to these physical “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, 20:25–29.


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

75. Gross also discloses that Classifier Engine identifies repair conditions associated with “a particular attribute” (i.e., *a repair condition of a property characteristic*). EX1005, ¶60. For example, the Classifier Engine may automatically categorize the condition of a roof (*property characteristic*) as “new,” “missing/damage tile,” and/or “holes/cracks.” EX1005, ¶61. Other examples of such condition classifications disclosed by Gross include “[f]ixtures” with a “new,” or “damaged” condition and “[w]indows/[s]kylights/[d]oors” with a “new,” “breaks,” or “warping” condition. EX1005, ¶¶62, 65.

76. In some instances, “severity weightings (i.e., a scaling factor of any convenient range (1-10, or the like) for the identified condition” may also be

assigned by Classifier Engine. EX1005, ¶70. The Classifier Engine may also compute a “condition score ... for the detected attribute.” EX1005, ¶126. These condition categorization and condition scores/severity are *a repair condition* and they are associated with an attribute, i.e., *a property characteristic*. 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. Thus, Gross discloses *a method for automatically categorizing a repair condition of a property characteristic*.

77. Figure 11 of Gross “identifies examples of structural features, parameters, conditions ... for a particular building structure,” which illustrates the results of the Classifier Engine categorizing repair conditions of property characteristic, as claimed. EX1005, ¶21.





EX1005, FIG. 11 (annotated).

2. [1.1]: receiving, from a user at a remote computing device, a request for a property condition classification, wherein the property classification request includes identification of a property and at least one property characteristic;

78. In my opinion, Gross renders obvious element [1.1].

79. In my opinion, Gross teaches *receiving, from a user at a remote computing device* a property-related *request* that includes identification of a geographical region of the property and property structural elements (*property*

characteristic). In particular, Gross describes its system 110 engages with users employing computing devices 112, such as smartphones, tablets, laptops, and desktops (i.e., *remote computing devices*) through interface logic 120. EX1005, ¶43. For example, Figure 1 of Gross illustrates system 110 connecting through interface logic 120 with external (*remote*) data sources 113, 114 and user's computing devices 112). Gross also discloses that “[i]n a preferred embodiment system 110 is part of a website which can be accessed through a conventional browser.” EX1005, ¶43. Gross discloses various interactive user interface, which in a preferred embodiment is implemented as part of an interactive web page utilized by users to search for real estate/building stock that meets particular criteria of interest (i.e., a request for information about property). *See, e.g.*, EX1005, ¶98. The interactive user interfaces described by Gross depict *receiving, from a user at a remote computing device, a request*, as that is what user interfaces do. For example, Figures 3B and 3C of Gross (annotated below) depict interactive user interfaces that receive requests from users related to “selection of search parameters” and “identification and review of property details” via “a number of query selection boxes,” including query box 360 to specify a location for a property search (e.g., a geographic region, such as a city, state, or neighborhood) and query box 366 to specify additional property structural elements (such as specific attribute styles for facades). EX1005, ¶¶98–99, 102, 108.



EX1005, FIG. 3B (top, annotated), FIG. 3C (bottom, annotated).

80. Gross further describes a property condition classification process that, as described in element [1.7] below, returns to a user (via a graphical user interface) a condition assessment of a property characteristic. EX1005, ¶¶120, 130, FIGs. 5, 7E. The output includes “additional classifying details about the building structure” to “allow the user to quickly see the assessment of the property as performed by Building Classifier Engine 150.” EX1005, ¶103, FIG. 7E. In a

preferred embodiment, building structures in “a target city or other convenient population unit” are assessed. EX1005, ¶¶161-162. Gross does not explicitly disclose how its assessment process of Figure 5 is initiated. *See* EX1005, ¶¶120-121 (stating Figure 5 depicts a building attribute/condition assessment process; beginning the process with retrieving an image).

81. A POSA would have understood, or found obvious, that a user interface—just like the ones Gross explicitly contemplates in Figures 3B and 3C—was an identified, predictable solution that could have been used to initiate Gross’s assessment process. This is because although Gross describes its interactive web user interface 350 is used to search for real estate/building stock that meet certain criteria of interest (EX1005, ¶95), a POSA would have found obvious that a similar interactive user interface could have been used to allow a user to submit the request for performance of the attribute/condition assessment process (i.e., *a property condition classification*) described above and would have been motivated to modify Gross to use such an interface, because interactive user interfaces were a well-known way in which to receive a request for information from a user. For example, Thornberry discloses an interactive user interface that allows a user to input a target address and a number of planes of the roof (property characteristic) to request a roof assessment. EX1024, FIG. 4D. A POSA would have been motivated to allow the user to request performance of the attribute/condition

assessment process in order to provide the information when needed for a particular property. Additionally, a POSA would have understood performing the assessment at the request of a user would be more efficient and reduce processing costs associated with unnecessary processing of property images.

82. In other words, this interactive user interface described in Gross that receives requests from user could be modified to initiate the assessment for a specific *identification of a property*. Furthermore, a POSA would have been further motivated to perform this modification because such a modification would allow the user to request Gross's property assessment to be performed for a specific target property of interest. This would also reduce processing costs that would have otherwise been associated with the batch processing of properties for an entire city or regional area, which is described in Figure 16 of Gross. *See, e.g.*, EX1005, ¶¶162–163.

83. And, in light of Gross's teachings that its interactive user interface includes query boxes for property characteristics (e.g., façade types) (EX1005, ¶107), a POSA would have been further motivated to allow a user to identify in its assessment request *at least one property characteristic*, allowing a user to request an assessment only for specific characteristics of interest (e.g., a roof company may only wish to perform the assessment for a roof). This modification also would further reduce processing costs by only performing the assessment for relevant

property characteristics of interest to a user, rather than for the entire property.

84. The described modifications to Gross's interactive web interface would have been a predictable variation of an interactive web user interface and readily implemented by a POSA. Indeed, user requests to perform assessments for a particular property location were well known. For example, Pershing discloses a customer can request an estimate for repair or replacement of a roof and the assessment is initiated by entering the customer's address. EX1009, 2:13-24. Jagannathan similarly discloses a user can request a roof assessment by providing the target roof input (e.g. street address). EX1012, ¶32. Thornberry provides a feature for a user to input an address to identify the location of a roof structure to be assessed. EX1024, 9:43-10:39. And Labrie discloses digital inspection data, that includes geocoding information (e.g., address of the location of the property) identifying the property for which to perform a property repair analysis. EX1025, 13:45-14:39.

85. Therefore, it is my opinion that Gross renders obvious *receiving a request* (through a modified interactive user interface) *for a property condition classification* (the result of which is output of Classifier Engine's assessment of the condition(s) of property attribute(s)). *The request includes identification of a property and at least one property characteristic* (allowing the user to perform a targeted assessment), thereby reducing processing costs.

3. [1.2]: obtaining, by processing circuitry of a computing system responsive to receiving the request, an aerial image of a geographic region including the property;

86. In my opinion, Gross teaches element [1.2].

87. Gross describes “aerial drones, satellite, balloon, and similar technology can be used ... to easily capture structure image data” “for the building structure” and can “facilitate building out a comprehensive image database.”

EX1005, ¶164. These images for a building are captured by aerial drones, satellite, and balloon (e.g., aerial objects), and are thus *aerial images of a geographic region that includes the property*. The building image data is stored in a structure image database. EX1005, ¶165.

88. Gross discloses that its “data acquisition process” is used by a classifier (e.g., Classifier Engine 150) “to acquire” (i.e., *obtain[]*) “appropriate image data for structures within a target area,” such as a city. EX1005, ¶161. Classifier Engine retrieves an image 700 from a database to begin the “building attribute/condition assessment process.” EX1005, ¶¶120-21. Gross obtains the aerial image *responsive to receiving the request* in limitation [1.1], because the data acquisition process described in Gross *obtains the aerial image* of the property from the database so that Classifier Engine can perform the assessment process and provide to the user the requested property condition classification.

89. Gross additionally 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. These “executable software routines” are executed via “processing circuitry.” Accordingly, Gross discloses *obtaining by processing circuitry* the aerial image.

4. **[1.3]: extracting, by the processing circuitry, one or more of a plurality of features from the aerial image corresponding to the property characteristic, wherein the extracted features include pixel groupings representing the property characteristic;**

90. In my opinion, Gross teaches element [1.3].

91. 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, as I discussed above, the '029 patent describes “property features” as being physical aspects, qualities, or facets (e.g., “rooftop, swimming pool, chimney, and sky lights”) of a property. EX1001, 7:64–8:1. 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 '029 patent describes that its “extract[ed] image-related features” can include “angles, outlines, substantially homogenous fields, etc.” EX1001, 7:64–

8:4. Accordingly, the claimed *extracted features*, here, appear to follow the computer vision meaning of the term.

92. Gross discloses “objects ... can be analyzed in image patterns, using a variety of techniques ... which can be used to detect and *extract* features of attributes.” EX1005, ¶77. An image, such as an aerial image as I discussed 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.” EX1005, ¶¶120, 123. A POSA would have understood or found obvious that this identification is performed *by processing circuitry that extracts one or more of a plurality of features from the aerial image*. In fact, Gross specifically discloses that its Classifier Engine is trained by extracting smaller distinct (image) features from reference images and further notes that such image “[f]eature extraction algorithms are well-known in the art.” EX1005, ¶¶77, 79.


93. Gross discloses that exemplary attributes (“structural elements” of a property) identified by its Classifier Engine include “roof elements, window elements,” and “façade elements.” EX1005, ¶123. As explained above in paragraph 74, these elements represent *property characteristics*. The image feature(s) extracted by Classifier Engine correspond to (represent) “the property characteristic” identified in limitation [1.1]. Gross additionally teaches *the*

extracted features include pixel groupings representing the property characteristic.

Gross discloses that “[u]sing conventional image processing, each attribute” (i.e., *property characteristic*) “may be ... characterized by a distinct set of pixels, shapes, sizes, proportions, curves, textures, or otherwise mathematically derivable feature from the image.” EX1005, ¶78. These are all examples of *pixel groupings representing the property characteristic* because 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 a digital image as (and is represented by) a corresponding sequence of pixels whose color and intensity values differ from surrounding pixels to form a curve pattern that processing circuitry (and the human eye) can recognize.

94. Gross further notes that feature extraction algorithms are well-known in the art, and identifies a suitable example to be used in its embodiments is described in “*A Survey of Face Recognition Techniques* by Jafri” (EX1008), which I understand from counsel is incorporated by reference into Gross. EX1005, ¶79. I understand from counsel that material that is incorporated by reference into another document is effectively part of the host document as if it were explicitly contained therein. Jafri, as incorporated in Gross, describes a recognition system that detects “prominent facial landmarks (i.e., the eyes, nose, and mouth)” using *pixel*

groupings to identify those image *features*. EX1008, 43, 51. For example, Figure 19 of Jafri (depicted below) illustrates pixel groupings around the nose, eye, and lips.



EX1008, FIG. 19.

5. [1.4]: determining, by the processing circuitry from the extracted features, a property characteristic classification for the property characteristic, wherein determining the property characteristic classification includes applying the pixel groupings for the property characteristic to a first machine learning classifier trained to identify property characteristics from a set of pixel groupings;

95. In my opinion, Gross teaches element [1.4]. Gross discloses that, after

it has been trained, Classifier Engine 150 (the claimed *first machine learning classifier*) performs an identification of attributes (e.g., roof elements) presented in an image. EX1005, ¶123. The specific types of the attributes (i.e., *property characteristic classification(s)*) are also captured “(i.e., façade (shingle, siding, brick), ... roof {pitched, double pitched, hipped, flat, metal, tile, shingle, slate, parapet, dormer, mansard, fascia, brackets, eaves, pent, pediment, etc.})”. EX1005, ¶57. For example, 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 do this, Classifier Engine “employs image processing” and the “identification of such attributes (and their types) can be performed in any number of ways by Classifier Engine 150 after it has been trained.” EX1005, ¶¶120, 123. Put differently, Gross takes the pixels of features (i.e., *pixel groupings*) extracted from images (described in [1.3]), feeds them through its Classifier Engine (i.e., *applies them to a first machine learning classifier*) to identify the property attribute (*property characteristic*) and its type (*property characteristic classification*). EX1005, ¶¶120, 123. As discussed for element [1.2], “processing circuitry” executes the 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).

96. Therefore, it is my opinion that Gross teaches determining, by the

processing circuitry from the extracted features (extracted in limitation [1.3]), a property characteristic classification (e.g., pitched) for the property characteristic (e.g., roof), wherein determining the property characteristics includes applying the pixel groupings (extracted in limitation [1.2]) to Classifier Engine 150.

97. Gross's Classifier Engine 150 is *a first machine learning classifier trained to identify property characteristics from a set of pixel groupings.*" In particular, Gross discloses its Classifier Engine is trained with reference images and that objects in the image are analyzed using neural networks to detect and extract features. EX1005, ¶77. Classifier Engine analyzes images to extract smaller distinct features that can be readily detected. EX1005, ¶77. Gross, through the incorporated Jafri material, describes using "a hybrid neural network solution," which is a form of classifier, to analyze image patterns. EX1008, 49. 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., a set of pixel groupings). EX1005, ¶78. Thus, Classifier Engine is *a machine learning classifier that is trained to identify property characteristics from a set of pixel groupings* (features characterized by the pixel groupings). EX1005, ¶¶77, 123.

6. **[1.5]: determining, by the processing circuitry based on the identified property characteristic and the extracted features, a condition classification for the property characteristic, wherein identifying the condition classification includes**

applying the pixel groupings for the property characteristic to a second machine learning classifier trained to identify property characteristic conditions from a set of pixel groupings;

98. 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., *determining a condition classification for the property characteristic*). EX1005, ¶60. Gross describes examples of structure attributes (property characteristics) and related conditions (condition classifications), including: roof (new, missing/damage tile, holes/cracks); façade (new, missing shingles/siding, holes); and support (leaning/damaged columns). EX1005, ¶¶60-69. The examples provided are just a partial list and other individual structure attribute/condition pairings can be identified by Classifier Engine. EX1005, ¶70.

99. 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/rating for the detected attribute, such as a “poor” condition rating. EX1005, ¶¶123, 126. To do this, Classifier Engine (i.e., *second machine learning classifier* for reasons I explain 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.” EX1005, ¶¶78, 120. In other words, the *pixel groupings for the*

(identified) *property characteristic* are further analyzed to *determin[e]* the *condition classification* for that identified characteristic. For the reasons I provided above for element [1.2] and [1.4], “processing circuitry” executes the process(es) performed by Classifier Engine. Therefore, it is my opinion that Gross teaches *determining, by the processing circuitry based on the identified property characteristic and the extracted features, a condition classification (e.g., missing tile) for the property characteristic (e.g., roof), wherein identifying ... includes applying the pixel groupings for the property characteristic to a second machine classifier.*

100. Gross teaches that its Classifier Engine is *a machine learning classifier*. EX1005, ¶77; *see* element [1.4]. 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 does explicitly disclose that “it may be useful to employ multiple classifiers ... to identify the attributes and classify the structure.” EX1005, ¶80. Gross further discloses examples of using a classifier to identify conditions, such as “damaged/exposed roof elements” and “exposed—broken façade,” EX1005, ¶¶85, 91, which correspond to the attributes (*property characteristics*) discussed 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. Additionally, Gross explains that “each attribute (and related condition qualifier) may be reduced in dimensionality and characterized by a distinct set of pixels, shapes, sizes, proportions, curves, textures, or other mathematically derivable feature from the image,” which teaches that different types of features (e.g., attributes and conditions) require different image processing approaches. EX1005, ¶78. In light of these teachings, a POSA would have understood that Gross teaches, or renders obvious, the use of *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.

101. A POSA would have found it obvious to separate the classifiers and would have been motivated to employ a second machine 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 improve the accuracy of their classifications, such as in this case the identification of property characteristic conditions. A POSA would have understood that the two classification tasks of

identifying what object is present versus determining the condition of that object are fundamentally different in nature and require analysis of different types of image features. In addition, a POSA would have understood that separating these classification tasks allows for independent optimization of each 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.

102. For example, a first classifier may focus on pixels and image features that delineate object boundaries and structural elements to identify *property characteristics* (e.g., detecting edges, corners, and geometric shapes that define a roof structure), while a second classifier may focus on pixels and image features within the interior regions of identified objects to assess condition (e.g., detecting texture, color, surface artifacts). This differentiation is further supported by Gross's teaching that "[t]he training results in extraction and optimization of specific features most appropriate for a particular element, so as to minimize classification errors." EX1005, ¶78. A POSA would have readily implemented the predictable variation of 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.

103. 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) property characteristic conditions (e.g., missing tile) from a set of pixel groupings*. EX1005, ¶¶60-66. For the reasons provided for element [1.4], this second machine learning classifier would be *trained to identify property characteristic conditions from a set of pixel groupings*. EX1005, ¶¶77 (describing training Classifier Engine with reference images), 126–127; EX1008, 43, 51, FIG. 19 (Jafri’s description of a classifier using pixel groupings to identify features). A POSA would have understood or found obvious that this second machine learning classifier employs image processing after it has been trained and that determining the condition classification includes *applying the pixel groupings for the property characteristic classification* to Classifier Engine 150. EX1005, ¶126.

7. **[1.6]: determining, by the processing circuitry based in part on the property characteristic classification and the condition classification, a risk estimate of damage to the property due to one or more disasters; and**

104. In my opinion, Gross, in combination with Davis, renders obvious element [1.6]. 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 instance, Gross discloses that “risk assessments, etc., can be based on an evaluation of an upkeep/maintenance evidenced for a particular property” (i.e., *condition classification*). For example, a “missing/damage tile” or

“holes/cracks” roof condition evidences poor upkeep and maintenance. EX1005, ¶¶60-70, 238. Gross further discloses that “correlations may be developed between property condition ratings” (i.e., the *condition classifications* described above in element [1.5]) and insurance metrics such as “number of claims filed, type of claim, severity.” 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.

105. Gross does not explicitly disclose basing its risk assessment on the property characteristic classification, but Davis describes that “risk analysis unit 34” “may analyze the feature(s) to determine a risk of damage to at least a portion of the property” and the rules may, for example, depend “on the type of roof construction,” (i.e., risk assessment *based on the property characteristic classification*). EX1006, 10:18-38. In light of Davis’s teachings, a POSA would have understood or found obvious to base Gross’s insurance risk assessment on the property characteristic classification (e.g., type of roof), in addition to the condition classifications (e.g., an old roof) because it is common sense that the risk varies depending on the specific property characteristic classification (e.g., a metal roof would have less risk of damage for certain types of disasters than a slate roof).

106. Gross also teaches these risk assessments can be *risk estimates of*

damage 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.*, ¶63, (listing “water damage” as an example condition assessment). Other 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.

107. Davis also teaches risk assessments based on “disasters.” In particular, 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, 2:7–15. 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.

108. Davis describes its risk analysis unit 34 “may generate different kinds of risk information based upon the analyzed feature(s) and other information” (i.e., *property 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.

109. A POSA would have found it obvious to modify Gross to incorporate Davis's teachings of determining a risk of property damage due to disasters for the reasons I provided above in Section IX.B. As I explained, a POSA would have recognized the benefits in the combined Gross-Davis system to provide more specific information for an insurer to determine the risk of property damage based on property characteristic classifications and conditions and would have been readily able to implement this predictable variation.

110. The Gross-Davis combination would, therefore, in my opinion render obvious *determining by the processing circuitry* (e.g., Gross's Classifier Engine 150) *based in part on the property 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).

8. **[1.7]: returning, to the user at the remote computing device via a graphical user interface responsive to receiving the request, a condition assessment of the property characteristic including the condition classification and the risk estimate of damage to the property due to the one or more disasters.**

111. In my opinion, Gross, and the Gross-Davis combination, renders obvious element [1.7]. As I described above in limitation [1.1], Gross discloses its user interface is preferably an interactive web page, which is *a graphical user interface*. EX1005, ¶98. Gross describes that a report of the results is presented to the user (*at the remote computing device*) in response to receiving the request. EX1005, ¶96. For example, the user can be provided with “classifying details about the building structure such as highlighted, tagged or annotated version of the building structure identified in FIG. 7E ... to quickly see the assessment of the property as performed by Building Classifier Engine 150” (i.e., *a condition assessment of the property characteristic including the condition classification*). EX1005, ¶103. Figure 7E of Gross (depicted below) illustrates a “visual assessment report 730” that “identifies at least those attributes identified by the system as having some measure of damage, impairment, aging, weathering, etc.” EX1005, ¶130.




FIGURE 7E

EX1005, FIG. 7E.

This report includes *a condition assessment of the property characteristic including the condition classification* because it includes the condition classifications of property characteristics (e.g., damaged/defective façade).

112. A POSA would have been motivated to modify the visual assessment report disclosed by Gross to include *the risk estimate of damage to property*—taught by Gross’s insurance risk assessment and also by the Gross-Davis combination—to provide one report to convey all relevant information to the user. Conveying the information in one report would have been a simple design choice and would have allowed a user (e.g., an insurance agent) to see all relevant information in one place, rather than viewing two separate reports. And such a

modification to the report to include the risk assessment information would have been a predictable variation and easily implemented by a POSA. Therefore, it is my opinion that Gross, and the Gross-Davis combination, teaches limitation [1.7], *i.e., returning, to the user at the remote computing device via a graphical user interface responsive to receiving the request, a condition assessment of the property characteristic including the condition classification and the risk estimate of damage to the property due to the one or more disasters.*

D. Independent Claim 15

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

113. To the extent the preamble is determined to be limiting, it is my opinion that Gross discloses the preamble of claim 15 for the reasons I discuss below for element [8.1].

2. **[15.1]: receive, from a user at a remote computing device, a property condition classification request, wherein the property classification request includes identification of a property and at least one property characteristic;**

114. In my opinion, Gross teaches element [15.1] for the reasons I discuss above for element [1.1].

3. **[15.2]: obtain, from a remote data source responsive to receiving the property classification request, an aerial image of a geographic region including the property;**

115. In my opinion, Gross teaches element [15.2] for the reasons I discuss

for element [1.2] and [8.3].

4. **[15.3]: extract from one or more of a plurality of features from the aerial image corresponding to each property characteristic of the at least one property characteristic, wherein the extracted features include pixel groupings representing the respective property characteristic;**

116. In my opinion, Gross teaches element [15.3] for the reasons I discuss above for element [1.3].

5. **[15.4]: for each of the at least one property characteristic,**

117. In my opinion, Gross teaches element [15.4]. I understand from counsel that this element requires elements [15.5] and [15.6] be executed only a single time because one property characteristic is “at least one.” Although not required by the claim, it is my opinion that Gross teaches executing elements [15.5] and [15.6] for multiple property characteristics, such as a “roof” as a first property characteristic and a “façade” as a second property characteristic.

6. **[15.5]: determine, from the extracted features for the respective property characteristic, a respective property characteristic classification, wherein determining the respective property characteristic classification includes applying the pixel groupings for the respective property characteristic to a first machine learning classifier trained to identify property characteristics from a set of pixel groupings, and**

118. In my opinion, Gross teaches element [15.5] for the reasons I discuss above for element [1.4].

7. **[15.6]: determine, based on the respective identified property characteristic and the respective extracted features, a**

respective condition classification for the respective property characteristic, wherein identifying the respective condition classification includes applying the respective pixel groupings for the respective property characteristic to a second machine learning classifier trained to identify property characteristic conditions from a set of pixel groupings; and

119. In my opinion, Gross renders obvious element [15.6] for the reasons I discuss above for element [1.5].


8. [15.7]: determine, in real-time responsive to receiving the property classification request and using the property characteristic classification of each property characteristic and the condition classification of each property characteristic, at least one risk estimate representing risk of damage due to disaster.

120. In my opinion, the Gross-Davis combination renders obvious element [15.7] for the reasons I discuss for element [1.6] and [8.7].

E. Claim 2

121. Claim 2 depends from claim 1 and further recites, “wherein the property characteristic classification is rooftop shape.”

122. 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,” such as roof shapes, e.g., a “flat” shape in Figure 12. EX1005, ¶21; *see also* EX1001, 20:28-29, FIG. 2A (identifying exemplary rooftop shapes as gambrel, gable, hipped, square, flat).



EX1005, FIG. 12.

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.

F. Claim 3

123. Claim 3 depends from claim 2 and further recites, “wherein the first machine learning classifier is trained with a set of aerial training images having known rooftop shapes such that the machine learning classifier is configured to

determine the rooftop shape from a plurality of rooftop shapes.”

124. In my opinion, Gross teaches this claim. Gross discloses that “the Classifier Engine 150 is trained with the template consisting of the training set of reference images 144 and tags as created and logged in database 152 ... using a variety of techniques, including statistical processing, neural networks, etc.”

EX1005, ¶77. These “reference images 144” include “conditions/impairments ... associated with a particular building attribute,” and with respect to roofs, Gross provides various examples of roof shapes and conditions for training Classifier Engine 150, such as “pitched, double pitched, hipped, flat, metal, tile, shingle, slate, parapet, dormer, mansard, fascia, brackets, eaves, pent, pediments, etc.”

EX1005, ¶¶57, 60.

125. Therefore, it is my opinion that Gross teaches *the first machine learning classifier* (e.g., Classifier Engine) *is trained with a set of aerial training images* (e.g., “reference images 144”) *having known rooftop shapes* (e.g., “pitched,” “flat,” “shingle”) *such that the machine learning classifier is configured to determine the rooftop shape* (e.g., “flat”) *from a plurality of rooftop shapes*.

G. Claim 4

126. Claim 4 depends from claim 1 and further recites, “wherein at least one of the first machine learning classifier and the second machine learning classifier is a Network in Network (NIN) classifier.”

127. In my opinion, Gross teaches this claim. Gross discloses that “[f]eature extraction algorithms are well-known in the art” and that “[u]sing conventional image processing, each attribute (and related condition qualifier) may be reduced in dimensionality and characterized by a distinct set of pixels, shapes, sizes, proportions, curves, textures, or other mathematically derivable feature from the image.” EX1005, ¶¶78, 79. These “conventional image processing” techniques and “well-known” “feature extraction algorithms” included a Network in Network classifier. As I described above in Section VII.B, this is further supported by the ’029 patent’s own Background section, which notes the Network in Network model as being “[o]ne model used for deep learning.” EX1001, 1:50–55.

H. Claim 5

128. Claim 5 depends from claim 1 and further recites, “wherein determining the risk estimate comprises applying a disaster risk profile corresponding to a) a first disaster of the at least one disaster and b) the property characteristic.”

129. As I explained above, Gross discloses that “risk assessments, etc., can be based on an evaluation of an upkeep/maintenance evidenced for a particular property; in this respect correlations may be developed between property condition ratings ... and ... severity.” EX1005, ¶238. These risk assessments can be risk estimates due to one or more potential disasters, such as due to “arson” 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.,* ¶63, (listing “water damage” as an example condition assessment). In other words, these correlations may associate maintenance condition ratings of a property attribute (e.g., a roof in “poor” condition) with higher risk of damage (e.g., from arson). 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.

130. Gross, therefore, teaches this element in my opinion. In particular, obtaining a “risk assessment” for a disaster using “correlations ... developed between property condition ratings ... and ... severity” would involve *applying a disaster risk profile* (e.g., the “correlations”) to *a first disaster of the at least one disaster* (e.g. “arson”) and *the property characteristic* (e.g., the attributes associated with the “property condition ratings”), because the correlations are correlating the risk of damage for a specific disaster with the analyzed “condition rating” of a property attribute/structure. EX1005, ¶¶238, 78. In other words, to obtain the risk assessment result, the system would apply the “correlation” between 1) “risk of arson” damage and 2) a “building” or other “fixture” on the property. EX1005, ¶¶238, 62. *See also* EX1005, ¶¶61–70 (of other structure attributes of a property whose “condition ratings” for which risk assessment correlations with

disasters in Gross can be developed).

131. In my opinion, Davis also teaches this claim. Davis describes that its risk analysis unit 34 may “utilize certain types of a priori knowledge to determine risk,” which “may be incorporated as rules stored in data storage 40.” EX1006, 10:22–24. Davis describes how “risk analysis unit 34 may access rules specifying that certain tree species are more prone to failure than others ... with respect to particular conditions such as high winds, wet soils, lightning, insect infestations.” These rules may also specify whether trees of certain species more likely to catch fire. *See* EX1006, 10:25–38. Risk analysis unit 34 then applies the “access[ed] rules” to a disaster of interest (e.g., wildfire, arson, hurricane, high winds) and a property characteristic (e.g., tree, roof) to “estimate a likelihood that damage will occur to some portion of the property.” EX1006, 10:25–26, 11:30–34. Applying “access[ed] rules” specifying a likelihood of disaster based on a property characteristic is *applying a disaster risk profile*.

132. It would have been obvious to modify Gross to incorporate Davis’s teachings of applying rules to a disaster and a property characteristic when determining a risk of property damage due to disasters for the reasons I discuss above in Section IX.B. As I explained above, a POSA would have recognized the benefits in the combined Gross-Davis system to provide more specific information for a an insurer to determine the risk of property damage based on property

characteristic classifications and conditions and would have been readily able to implement this predictable variation.

133. The Gross-Davis combination would, therefore, in my opinion render obvious *applying a disaster risk profile* (e.g., the “access[ed] rules,” “correlations”) *to a first disaster of the at least one disaster* (e.g., wildfire) *and the property characteristic* (e.g., tree). EX1006, 10:25–26; EX1005, ¶238.

I. Claim 6

134. Claim 6 depends from claim 1 and further recites, “wherein the pixel groupings comprise at least one of angles, outlines, and substantially homogenous pixel fields defining the property characteristic.”

135. 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, sizes, proportions, curves, textures, or other mathematically derivable feature from the image.” EX1005, ¶78. Additionally, Gross, through the incorporated Jafri material, expressly discloses extracting features and computing geometric relationships (e.g., “distances, areas and angles”) to represent pixel groupings of facial features. EX1008, 43; *see also* EX1005, ¶79 (noting that Jafri’s feature extraction algorithms are suitable or can easily be adapted for use in Gross’s embodiments). 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), *and substantially homogenous pixel fields*.

EX1005, ¶78; EX1008, 43.

J. Claim 7

- 1. [7.1]: The method of claim 1, wherein identifying the condition classification for the property characteristic comprises: generating, from the pixel groupings, a pixel intensity distribution of the property characteristic; and**

136. In my opinion, Gross renders obvious this element. Gross expressly teaches that “the dark or irregular color [of a group of pixels representing a roof] may cause a roof element to score poorly” during condition determination.

EX1005, ¶126. This disclosure demonstrates that Gross’s system analyzes pixel color and brightness characteristics—i.e., pixel intensity—to assess property conditions. The incorporated Jafri material also discloses using a pixel intensity grouping to classify features (e.g., eyes, nose). EX1008, FIG. 19.

137. Gross further teaches that this analysis involves “[u]sing conventional image processing,” where “each attribute (and related condition qualifier) may be reduced in dimensionality and characterized by ... proportions ... or other mathematically derivable feature from the image.” EX1005, ¶78. A POSA would have understood or found obvious that *generating, from the pixel groupings, a pixel intensity distribution* constitutes exactly this type of “conventional image

processing” that produces “mathematically derivable feature[s]” for condition assessment.

138. Pixel intensity refers to the brightness level of pixels in an image and is a fundamental parameter for characterizing and comparing grayscale and colored images. Given Gross’s teaching that “dark or irregular color” affects condition scoring, a POSA would have found it obvious to systematically generate a pixel intensity distribution from pixel groupings of an image to quantify these color/brightness characteristics for input into Gross’s Classifier Engine. In other words, a pixel intensity distribution would be the means by which Gross’s system could learn that the “color” of pixel groupings in an image was “dark or irregular.” This represents a predictable variation that a POSA could have readily implemented by feeding a pixel intensity distribution from extracted pixel groupings into Gross’s Classifier Engine. Therefore, it is my opinion that Gross renders obvious, *generating, from the pixel groupings* (e.g., extracted image features), *a pixel intensity distribution of the property characteristic* (e.g., roof). EX1005, ¶¶78, 126.

2. [7.2]: determining, by applying the second machine learning classifier, the condition classification based on the pixel intensity distribution of the property characteristic.

139. In my opinion, Gross teaches this element. As explained, Gross discloses that during the condition determination, “the dark or irregular color [of a

group of pixels representing a roof] may cause a roof element to score poorly.”

EX1005, ¶126. Determining that the condition (of a property attribute such as a roof) is poor based on finding a “dark or irregular color” is determining the condition classification *based on the pixel intensity distribution* of the property characteristic. Therefore, it is my opinion that Gross further teaches *determining, by applying the second machine learning classifier* (e.g., Classification Engine), *the condition classification* (e.g., poor condition) *based on the pixel intensity distribution of the property characteristic* (e.g., dark or irregular color).

K. Claim 16

140. Claim 16 depends from claim 15 and further recites, “wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to, prior to obtaining the aerial image of the geographic region, determine, based upon the at least one property characteristic, a preferred image type corresponding to each property characteristic of the at least one property characteristic.”

141. 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.

142. In light of these teachings, a POSA would have understood or found obvious that this “comprehensive image database” (containing different types of images) would allow the classifier to use certain, preferred types of images for certain property classification tasks. For example, an image that would “increase building coverage” of the roof of a property and increase the “accuracy” of the roof condition classification could be a top-down image of the property that exposes a higher amount of rooftop surface area (e.g., as opposed to an image of the property at street elevation). 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 characteristic 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 of the geographic region*.

143. 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 upon the at least one property characteristic (e.g., roof), a preferred image type corresponding to each property characteristic of the at least one property characteristic (e.g., a top-down image, an image that “increase[s] feature currency/accuracy”).

L. Claim 17

- 1. [17.1]: The non-transitory computer readable medium of claim 16, wherein: the at least one property characteristic comprises two or more property characteristics; and**

144. In my opinion, Gross teaches this element. For instance, Gross discloses that the condition classification process—in addition to identifying a first element (e.g., “roof elements”)—“is done repeatedly ... to identify any number of elements 712 (siding, dormers, fascias, overgrowth, water stains, missing windows, etc.) in poor condition.” EX1005, ¶¶123, 127.

145. Therefore, it is my opinion that Gross teaches *the at least one property characteristic comprises two or more property characteristics* (e.g., a roof and a window of a property).

- 2. [17.2]: the instructions, when executed by the processing circuitry, cause the processing circuitry to, responsive to determining the preferred image type corresponding to a first property characteristic of the at least one property**

characteristic is a terrestrial image, obtain a terrestrial image of the geographic region including the property.

146. In my opinion, Gross teaches or renders obvious this element for at least the reasons I discuss above for element [17.1]. As I explained above for claim 16, Gross discloses how obtaining images “at different angles, elevations and profiles” helps “increase building coverage and feature currency/accuracy.” EX1005, ¶148. Gross also discloses identifying and classifying multiple structural elements of a property, including “window elements.” EX1005, ¶123. In light of these teachings, a POSA would have understood or found obvious that an image that would “increase building coverage” of a window element of a property and increase the “accuracy” of the window condition classification was a side-profile image of the property at street elevation, for example, as opposed to an aerial image of the property shot from above, since such an image type would have a higher chance of containing pixels corresponding to a window element. EX1005, ¶148. Additionally, a POSA would have understood or found obvious that a side-profile image of a property at street elevation is “a terrestrial image,” since such an image is taken at street elevation.

147. Therefore, it is my opinion that Gross teaches instructions (e.g., instructions on “the computing system”) that cause the processing circuitry to, responsive to determining the preferred image type corresponding to a first property characteristic of the at least one property characteristic (e.g., a side-profile

image for window elements) is a terrestrial image (e.g., a side-profile image of a property at street elevation), obtain a terrestrial image of the geographic region including the property (e.g., “obtained directly” or from “an existing database of images” (EX1005, ¶148)).

M. Claim 18

148. Claim 18 depends from claim 15 and further recites, “wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to access one or more known property characteristics, wherein determining the at least one risk estimate comprises determining the at least one risk estimate further based on the one or more known property characteristics.”

149. In my opinion, Gross teaches this element for at least the reasons I discuss above for element [1.6]. 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.” 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. Obtaining occupancy information from “public records” requires the system to access, at least in part, the “public records” for the occupancy information, which is *accessing one or more known property characteristics* (e.g., the occupancy

information). EX1005, ¶118.

150. 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.” EX1005, ¶238. As I explained above, property occupancy is an example of a *known property characteristic* that can be accessed by Gross’s system. EX1005, ¶¶118, 115. Then, developing a risk estimate correlation value from property occupancy would be *applying a known property characteristic* when determining a risk estimate, as doing so would involve applying occupancy information accessed from “public records” to a mathematical model, thereby obtaining a mathematical correlation between the property occupancy characteristic and potential risk.

151. Therefore, it is my opinion that Gross teaches *instructions* (e.g., instructions on “the computing system”) that *cause the processing circuitry to access one or more known property characteristics* (e.g., occupancy information), *wherein determining the at least one risk estimate* (e.g., risk of damage to property due to fire) *comprises determining the at least one risk estimate further based on the one or more known property characteristics* (e.g., higher risk of damage

associated with “confirmed vacant” occupancy information).

N. Claim 19

152. Claim 19 depends from claim 18 and further recites, “wherein the one or more known property characteristics comprise at least one of a property age, a property elevation, a property slope, a year built, a year renovated, and a building height.”

153. In my opinion, Gross renders this element obvious. As I discuss above in claim 18, 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. In addition to “occupancy data,” Gross discloses accessing and correlating property characteristics, including “internal details (size, # of bedrooms, # rooms, etc.).” EX1005, ¶81.

154. The claimed *property characteristics* represent obvious variations from a finite set of well-known property details. Gross’s disclosure of “internal details” such as “size, # of bedrooms, # rooms” indicates that the system was designed to incorporate various standard property characteristics. EX1005, ¶81. There are only a limited number of fundamental property characteristics routinely used in property assessment systems, including the basic physical characteristics (e.g., size, building height, property elevation, property slope), temporal

characteristics (e.g., property age, year build, year renovated), and structural characteristics (e.g., # of bedrooms and # rooms) disclosed by Gross.

155. Given Gross’s teaching to use “internal details” and other property data for risk correlation, it would have been obvious to incorporate any of these well-known characteristics from the finite universe of standard property assessment parameters. Each of the claimed characteristics—*property age*, *property elevation*, *property slope*, *year built*, *year renovated*, and *building height*—represents predictable variations that a POSA would have been motivated to include based on their known relevance to property risk assessment. For example, an older property could have a higher risk of property damage, due to various factors such as material wear, outdated construction technology, hazardous materials, or reduced upkeep. The selection of any particular combination of these characteristics would have been obvious to try given the limited number of fundamental property details available and Gross’s teachings.

O. Claim 20

1. **[20.1]: The non-transitory computer readable medium of claim 15, wherein: receiving identification of the property and the at least one property characteristic comprises receiving, via a network from a remote computing device, the identification of the property; and**

156. In my opinion, Gross teaches element [20.1] for the reasons I discuss above for element [1.1]. Additionally, Gross discloses that its “software modules

... can be implemented ... using ... servers connected to a network (such as the Internet).” EX1005, ¶56. Therefore, it is my opinion that Gross discloses receiving identification of the property and the at least one property characteristic *via a network* (e.g., the Internet).

2. **[20.2]: the instructions, when executed by the processing circuitry, cause the processing circuitry to provide, via the network to a graphical user interface at the remote computing device in real-time responsive to receiving the at least one property characteristic, the at least one risk estimate.**

157. In my opinion, the Gross-Davis combination renders obvious element [20.2] for the reasons I discuss above for elements [1.7] and [20.1].

X. GROUND 2: GROSS RENDERS OBVIOUS CLAIMS 8 AND 11-14.

158. In my opinion, Claims 8 and 11-14 would have been obvious over Gross. A POSA would have found it obvious to combine different embodiments of Gross for the reasons I discuss in Ground 1.

A. Independent Claim 8

1. **[8.P]: A system for automatically categorizing a repair condition of a property characteristic, comprising:**

159. To the extent the preamble is determined to be limiting, it is my opinion that Gross discloses the preamble of claim 8 for the reasons I discuss above for the preamble of claim 1 (element [1.P]).

2. **[8.1] processing circuitry; and a non-transitory computer-readable medium having instructions stored thereon;**

wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to

160. Gross discloses “the critical steps identified in this process can be implemented into executable software routines and modules using any number of ways by skilled artisans.” EX1005, ¶161. A POSA would have understood or found obvious that these “executable software routines” are executed via *processing circuitry*. Gross further discloses that its “applications can be embodied in tangible, machine readable form for causing a computing system to execute appropriate operations in accordance” with the teachings, i.e., *a system*. EX1005, ¶56. Therefore, it is my opinion that Gross teaches element [8.1].

3. **[8.2]: receive, from a user at a remote computing device, a property condition classification request, wherein the property classification request includes identification of a property and at least one property characteristic;**

161. In my opinion, Gross teaches element [8.2] for the reasons I discuss above for element [1.1].

4. **[8.3]: obtain, from a remote data source responsive to receiving the property classification request, an aerial image of a geographic region including the property;**

162. In my opinion, Gross teaches element [8.3] for the reasons I discuss above for element [1.2]. Gross further teaches the aerial image is obtained *from a remote data source* through its description that aerial image data is retrieved for a target address and preferably “obtained from a third party vendor.” EX1005, ¶164.

5. **[8.4]: extract one or more of a plurality of features from the aerial image corresponding to the property characteristic, wherein the extracted features include pixel groupings representing the property characteristic;**

163. In my opinion, Gross teaches element [8.4] for the reasons I discuss above for element [1.3].

6. **[8.5]: determine, from the extracted features, a property characteristic classification, wherein determining the property characteristic classification includes applying the pixel groupings for the property characteristic to a first machine learning classifier trained to identify property characteristics from a set of pixel groupings;**

164. In my opinion, Gross teaches element [8.5] for the reasons I discuss above for element [1.4].

7. **[8.6]: determine, based on the identified property characteristic and the extracted features, a condition classification for the property characteristic, wherein identifying the condition classification includes applying the pixel groupings for the property characteristic to a second machine learning classifier trained to identify property characteristic conditions from a set of pixel groupings; and**

165. In my opinion, Gross renders obvious element [8.6] for the reasons I discuss above for element [1.5].

8. **[8.7]: determine, in real-time responsive to receiving the property classification request and using the property characteristic classification and the condition classification, a replacement cost for replacing the property characteristic.**

166. In my opinion, Gross teaches or renders obvious element [8.7]. Gross discloses “an overall estimate can ... be generated to identify a cost to restore the

building structure to a nominal target state.” EX1005, ¶134. This overall estimate is calculated by “correlating each of the impairments to repair or improvement figures, and summing over all the attribute conditions.” EX1005, ¶134. These “repair or improvement figures” teach, or at least render obvious, *a replacement cost for replacing the property characteristic* (e.g., attribute, such as a roof or window) because it was both known and is common sense that one way to improve the attribute is to replace the attribute. For example, Gross teaches that one identified condition is a window that has holes, EX1005, ¶65, which would require replacement.

167. Gross teaches its determination of the replacement cost uses: *the condition classification* (“summing over all the attribute conditions”); and *the property characteristic classification* because repair and improvement figures depend on the classification of that attribute. EX1005, ¶134. For example, a tile roof replacement has a different cost than a slate roof (e.g., due to material cost, supply and demand factors, etc.).

168. Further, Gross teaches, or renders obvious, its determination *is in real-time responsive to receiving the property classification request*. For the reasons I discuss above in element [1.1] and [1.7], Gross renders obvious an interactive user interface (e.g., an interactive web page), which generates a report responsive to a property condition classification request. And, for the reasons I

discuss above in element [1.1], Gross teaches, or renders obvious, *a property condition classification request* received from a user that allows the user to initiate the performance of the attribute/condition assessment process, and the system then returns the graphical user interface discussed in element [1.7].

169. Because the user's request initiates the assessment process, the process occurs *responsive to receiving the request*. Thus, Gross teaches, or renders obvious, that its Classifier Engine 150 performs its analysis *in real-time* because Gross depicts that its building attribute/condition assessment process retrieves the images, performs the assessment, and returns the assessment report to a human operator (which is a user) for immediate review. *See* EX1005, ¶¶120-134, FIG. 5 (illustrating the real-time steps of the property condition assessment process). Gross confirms its Classifier Engine 150 can perform its analysis *in real-time* (which includes the “generated” determination of replacement costs that is based on the condition assessment), for example, when it describes that a user may use a smartphone to capture a structure image for “rapid, on the spot identification and review of property details.” EX1005, ¶108. Gross also explains that even without “explicit geolocation data the user's image data of the property can easily be broken down and analyzed for attributes, features, etc. by Classifier Engine to identify likely matches, even if it takes more time.” EX1005, ¶109. Therefore, it is my opinion Gross teaches or renders obvious element [8.7].

B. Claim 11

170. Claim 11 depends from claim 8 and further recites, “wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to compare a current condition classification of the property characteristic to a stored historic condition classification of the property characteristic.”

171. 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., stored historic condition classification) “in database 142” during the feature coding (e.g., condition assessment) process. EX1005, ¶133. For example, Gross describes how “the presence of [a building] feature, type and condition is coded and collected and stored in digital form,” including “the presence, type and condition of facades, roofs, awnings,” and more. 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.

172. Gross also explains that “the coding process” that obtains the property assessments for the “long term evaluations” described above “may be targeted

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

173. Therefore, it is my opinion that Gross teaches instructions (e.g., instructions on “the computing system) that cause the processing circuitry to compare a current condition classification of the property characteristic (e.g., “good” condition roof) to a stored historic condition classification of the property characteristic (e.g., “damaged” roof). EX1005, ¶¶134, 146, 61.

C. Claim 12

174. Claim 12 depends from claim 11 and further recites, “wherein the instructions, when executed by the processing circuitry, cause the processing

circuitry to, in real-time responsive to receiving the property condition classification request, determine whether repairs have been made to the property based on the comparison of the current condition classification to the historic condition classification of the property characteristic.”

175. In my opinion, Gross teaches this claim. As I explained above for claim 11, a POSA would have understood or found obvious that performing “long term evaluations ... to identify changes” would involve comparisons between assessment data of current, updated images of a property and prior, historic assessment data of the same property. For example, a current assessment of a roof indicating a roof in “good” or “new” condition or having a “10” scaling factor can be compared to a historic assessment of a roof indicating a roof in “damaged” condition or having a “1” scaling factor to identify “the presence of a ... new roof” (i.e., a “repair[] ... to the property”). EX1005, ¶¶61, 146. Gross also discloses how its “coding process may be targeted primarily or solely for identifying [such] structural improvements.” EX1005, ¶146. Given these teachings and the reasons discussed in element [8.7], a POSA would have understood or found obvious that these structural improvement identifications occurred in real-time responsive to receiving the property condition classification request. A POSA could have readily implemented the predictable variation of determining structural repairs in real-time in response to a condition classification request to Gross’s Classifier Engine.

176. Therefore, it is my opinion that Gross teaches instructions (e.g., instructions on “the computing system) that cause the processing circuitry to, in real-time responsive to receiving the property condition classification request, (e.g., in response to a classification request made at Gross’s interactive user interface) determine whether repairs (e.g., “structural improvements”) have been made to the property based on the comparison of the current condition classification (e.g., “good” condition roof) to the historic condition classification of the property characteristic (e.g., “damaged” roof). EX1005, ¶¶134, 146, 61, 98.

D. Claim 13

177. Claim 13 depends from claim 8 and further recites, “wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to, prior to identifying the one or more of the plurality of features, assess orthogonality of the aerial imagery.”

178. In my opinion, Gross teaches 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.

179. 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.

180. 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. The incorporated Jafri reference also describes that “appropriate preprocessing of the images can compensate for noise and slight variations in orientation, scale, and illumination.” EX1008, 41. A POSA would have 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 identifying the one or more of the plurality of features. EX1005, ¶76.

181. Therefore, it is my opinion that Gross teaches instructions (e.g., instructions on “the computing system”) that cause the processing circuitry to, prior to identifying the one or more of the plurality of features (e.g., prior to using Gross’s Classifier Engine), assess orthogonality of the aerial imagery (e.g., assessing orthogonality during an image pre-processing/normalization process).

E. Claim 14

182. Claim 14 depends from claim 13 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 imagery does not correspond to a true orthophoto form, cause correction of the aerial imagery to the true orthophoto form.”

183. In my opinion, Gross teaches this claim for at least the reasons I discuss above for claim 13. 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 ’029 patent does not provide a definition for, or describe, true orthophoto form. A POSA would have understood that 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 ’029 patent. For example, the ’029 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, 20:3–7, 20:10–16.

184. As described 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.

185. Therefore, it is my opinion that Gross teaches instructions (e.g., instructions on “the computing system”) that cause the processing circuitry to, responsive to assessing that the orthogonality of the aerial imagery does not correspond to a true orthophoto form (e.g., determining that the vertical axis of a particular aerial image does not match the actual vertical axis of its scene), cause correction of the aerial imagery to the true orthophoto form (e.g., adjusting scale or relative angle of the aerial image).

XI. GROUND 3: THE COMBINATION OF GROSS AND FURUKAWA RENDERS OBVIOUS CLAIMS 9 AND 10.


186. In my opinion, the combination of Gross and Furukawa renders obvious claims 9 and 10.

A. Claim 9

- 1. [9.1]: The system of claim 8, wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to, prior to identifying the one or more of the plurality of features: obtain a shape map image including the property;**


187. In my opinion, Gross and Furukawa render obvious this element. The

'029 patent describes “shape map images” as being “used to confirm [a] location of a particular property,” such as being “overlaid with a corresponding aerial image 102c to match properties with images.” EX1001, 7:21–25. Figures 5A and 5C of the '029 patent (reproduced below) respectively depict an example “shape map image” and an example result of overlaying shape outlines with an aerial image.



EX1001, FIG. 5A (top), FIG. 5C (bottom).

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




EX1005, FIG. 10.

Gross describes that the images of “any particular city, town, etc., can be” “identified” “in this fashion ... based on computer records, address data, etc. ... for

purposes of achieving the objectives” set forth in Gross (e.g., property assessment).

EX1005, ¶137.

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



EX1009, FIG. 4.


Both Figure 10 of Gross and Figure 4 of the incorporated Pershing material are images of shapes that are mapped together to delineate properties and their surrounding objects, which can serve as the “shape map images” that the ’029 patent “use[s] to confirm [a] location of a particular property.” EX1001, 7:21–25.

190. Gross further discloses “a data acquisition process 1600 used by a classifier [e.g., Classifier Engine] as it would be implemented on a customized

structure assessment.” EX1005, ¶161. Gross explains the purpose of this process “is to acquire appropriate image data for structures within a target area.” EX1005, ¶161. Gross further discloses how “for [a] target address in question,” the system obtains and uses “[i]nformation identifying a beginning and end of each individual street, road, alley” “from any convenient database or similar source.” EX1005, ¶¶162, 164.

191. Gross also discloses obtaining “boundaries established by street address ranges” as part of this process. EX1005, ¶¶163, 143. As this data acquisition process obtains the image input(s) for the assessment process of the classifier—Gross’s Classifier Engine—this obtaining is also *prior to identifying the one or more of the plurality of features*. EX1005, ¶¶161–163.

192. In my opinion, Furukawa teaches instructions to *obtain a shape map image*. In particular, Furukawa teaches obtaining map data that includes “information of the external shape of buildings” which is stored and provided by a “map data storing means 4.” EX1007, 15:19–31. Figure 7 of Furukawa depicts an example of such map data:



EX1007, FIG. 7.

193. Furukawa further explains how the “map data” depicted in Figure 7 includes “shapes of buildings” as “raster data ... of those data showing shapes of buildings,” and can also include identification information, such as “names of administrative districts/roads/buildings and the like” and location information, such as “latitude and longitude.” EX1007, 12:26–39. This map data is a *shape map image including the property*, because, similar to the ’029 patent, it plots out the shapes of various structures in an area and includes their location information. And, also similar to the ’029 patent, which describes that the shape map data may be obtained from external sources (e.g., Geospatial Information Authority of Japan or Zenrin Co. Ltd. of Japan), Furukawa also describes its map data may be obtained from many products that are commercially available. *Compare* EX1001,

7:14-20, *with* EX1007, 8:62-66.

194. Such *shape map images* (e.g., urban planning maps, cadastral maps) were well-known before the September 23, 2016 priority date of the '029 patent. For example, the Geospatial Information Authority of Japan launched its National Large Scale Map Data Base project in 1989 to compile digital map data for urban planning, utility management, navigation, etc., and opened the project to the public in 1993. EX1021, 9. The United States Geological Survey (USGS) released *The National Map* in 2001, which compiled topographic information including boundary, transportation, and structure data for the U.S. EX1022, 8.

195. A POSA would have been motivated to modify Gross to incorporate Furukawa's teachings of obtaining map data, because Gross and Furukawa provide a teaching, suggestion, and motivation to do so. Gross teaches the need for address confirmation, describing how "building image[s] and tentative address[es] are tentatively tagged" during the classification process, and how "this address information" should be "confirmed" to ensure the system "contain[s] accurate information for a particular address." EX1005, ¶¶165, 170. Gross's teaching that tentative address information should require confirmation would have suggested to a POSA the benefit of incorporating additional data sources for verification purposes. Indeed, the process of confirming a piece of data involves comparing that piece of data to one or more independent reference points to check or prove its

validity. But Gross doesn't describe how the address confirmation would occur, and that would have motivated a POSA to seek out such solutions.

196. This need is supplied by Furukawa. Furukawa teaches the use of “map data” containing, for each of its various buildings: building shape information such as “the external shape of buildings,” identification information such as “names of ...buildings,” and “location information such as latitude and longitude” data, each of which is an additional stored reference data point that serves as the confirmation mechanism Gross contemplates. EX1007, 8:55–58, 12:26–34. Furukawa also teaches using this “map data” for verification purposes through the validation and correction of satellite image analysis of buildings. EX1007, Abstract. Thus, the prior art provided both Gross's problem of building address confirmation and Furukawa's map data validation solution, creating a teaching and motivation for their combination.

197. Therefore, it is my opinion that the Gross-Furukawa combination renders obvious *instructions* (e.g., instructions on “the computing system”) that, *when executed by the processing circuitry, cause the processing circuitry to, prior to identifying the one or more of the plurality of features* (e.g., before the assessment process): *obtain a shape map image including the property* (e.g., map data including “information of the external shape of buildings”). EX1007, 15:19–31.

2. [9.2]: overlay the aerial image with the shape map image; and

198. In my opinion, the Gross-Furukawa combination renders obvious this element. As I explained above in element [9.1], Furukawa's system obtains map data that includes "information of the external shape of buildings" which is stored and provided by a "map data storing means 4." EX1007, 15:19–31. Furukawa's system also "analyzes ... obtained satellite stereo images, and extracts vector information showing the external shape for the photographed building," which Furukawa describes is performed by a "building external shape information detecting means 6." EX1007, 32:65–33:7.

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

200. In other words, the shape outline obtained from the satellite image and the shape outline obtained from the map data are overlaid with each other, that is, so their corresponding points match each other, from which values such as area differences and overlap ratios may be calculated. For example, shape outlines that matched up with each other would produce low area differences and high overlap ratios. This would also mean that any matching building identified from the satellite image would share the same identification information (e.g., “building” “name”) and “location information” (e.g., “latitude and longitude”) with its map data counterpart, since they are linked to the same building with the matching shape outline. Inversely, shape outlines that do not match up with each other would produce high area differences and low overlap ratios.

201. A POSA would have found it obvious to modify Grossto incorporate Furukawa’s teachings of superimposing stored map data with shape information extracted from an aerial image to solve Gross’s need to identify the address of the property structure, as described in element [9.1]. As explained, the prior art provided both Gross’s problem of building address confirmation and Furukawa’s map data validation solution, creating a reason for the combination.

202. The Gross-Furukawa combination would, therefore, in my opinion renders obvious *overlaying* (e.g., “made to correspond to each other”) *the aerial image* (e.g., Gross’s aerial image) *with the shape map image* (e.g., map data

including “information of the external shape of buildings”). EX1007, 23:65, 32:17–24, 15:19–31.

3. [9.3]: determine whether a boundary of the property as identified by the shape map matches a boundary of the property as illustrated in the aerial image.

203. In my opinion, the Gross-Furukawa combination renders obvious, this element. As described above, Furukawa discloses a “building external shape information comparing/combining means 7” that combines and compares extracted shape information for a building with stored map data. EX1007, 32:17–24. The comparison process described in Furukawa determines whether the boundaries—depicted by the extracted shape information and the map data—match by evaluating factors including location proximity, area differences, and overlap ratios. Similar to my discussion above in element [9.2], shape outlines that matched up with each other would, for example, produce lower area differences, higher overlap ratios, and lower location proximity values. When these calculations fall within “predetermined value[s],” the system considers the depicted boundaries as matching and “indicating the same building.” EX1007, 23:56–65. Inversely, if these calculations did not meet the “predetermined value[s],” it follows that the system would consider the depicted boundaries as not matching and possibly not “indicating the same building.” EX1007, 23:56–65.

204. A POSA would have found it obvious to modify Gross to incorporate

Furukawa's boundary comparison teachings for at least the same reasons I discuss above for element [9.1]. As explained, the prior art provided both Gross's problem of building address confirmation and Furukawa's map data validation solution, creating a teaching and motivation for their combination.

205. Therefore, it is my opinion that the Gross-Furukawa combination renders obvious *determining whether a boundary of the property as identified by the shape map* (e.g., “information of the external shape of buildings of the map data”) *matches* (e.g., area difference below a threshold, overlap ratio above a threshold) *a boundary of the property as illustrated in the aerial image* (e.g., extracted “vector information showing the external shape of the photographed building”). EX1007, 21:22–23, 58–59.

B. Claim 10

206. Claim 10 depends from claim 9 and further recites, “wherein the instructions, when executed by the processing circuitry, cause the processing circuitry to, upon determining that the shape map does not match the boundary of the corresponding property, obtain an alternate aerial image of the property.”

207. In my opinion, the Gross-Furukawa combination renders obvious this claim. As I discuss above in claim 9, Furukawa's boundary comparison process determines whether boundaries match by evaluating factors including location proximity, area differences, and overlap ratios. EX1007, 23:56–24:2. And, when

these factors do not meet “predetermined value[s],” the system would recognize a mismatch. EX1007, 23:56–24:2.

208. Gross provides the teaching and motivation for *obtaining an alternative aerial image of the property* when mismatches occur in the combined system. First, Gross maintains multiple images of structures in its database. Gross discloses that “image data ... for the building structure is retrieved for the target address in question” and can be “obtained from a third party vendor” or “generated as needed” using various methods such as “aerial drones, satellite, balloon, and similar technology.” EX1005, ¶164. Gross also discloses that these methods “can obtain image data [at] different elevations” and “at different times.” EX1005, ¶164. Gross further teaches that its “structure image database” contains primary (e.g., original) images for a structure, along with “structure sub-images which are based on automatically dividing the original image into separate blocks, or separate areas corresponding to distinct building elements.” EX1005, ¶165.

209. Second, Gross renders obvious image matching processes that involve accessing multiple images when initial matches fail. For example, Gross explains that to “identify the property,” its system scans through a “range of image data” or uses Classifier Engine 150 “to identify likely matches from the preexisting stock of building structures in database 142.” EX1005, ¶109. Gross further teaches that “image matching for building structures can be accomplished in any number of

ways by adapting existing recognition algorithms to recognize the types of objects of interest.” EX1005, ¶108.

210. A POSA would have found obvious that these matching algorithms involved iterating through available images, that is, *obtaining alternate images*, because the algorithm would evaluate alternate available images from the database until a match was found. For example, in order to identify one “likely match” “from the preexisting stock of building structures,” the system would start by obtaining a first potential image of the property from the stock and determine whether it matched. EX1005, ¶109. If the first obtained image does not match, the system would then obtain an alternate potential image of the property and again determine whether that alternate image matched, repeating this process until a match was found.

211. A POSA would have also been motivated to modify the Gross-Furukawa combination to obtain alternate aerial images upon determining a boundary mismatch, because doing so would allow the Gross-Furukawa system to more frequently perform classification on the building that is actually at the target address. Gross’s teaching that building and address information should be “confirmed” to ensure that the system “contain[s] accurate information for a particular address” naturally suggests trying alternate image sources when an initial verification fails. EX1005, ¶170. For example, if the contemplated address

confirmation turned out unsuccessful, Gross's system would not proceed to assess the corresponding image, because the assessment results would not correspond to the correct property.


212. Additionally, Gross's teachings that multiple images of the same property can be captured "at different times" and "elevations" and that its system maintains a database with multiple image options would have suggested to a POSA that alternate images would be available (and beneficial) for when initial matching fails. EX1005, ¶164. And because Gross teaches that "image matching for building structures can be accomplished in any number of ways," a POSA would have also found it obvious or common sense to try alternative approaches (including using alternate images) when one approach failed, in order to maximize the chances that the overall "image matching" process for a property was successful. EX1005, ¶108.

XII. CONCLUSION

213. 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.

214. 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.

A handwritten signature in red ink, consisting of a series of loops and a final vertical stroke with a crossbar.

David A. Forsyth, Ph.D.