

PUBLIC VERSION

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

ARASHI VISION (U.S.) LLC (D/B/A INSTA360),
Petitioners,

v.

GOPRO, INC.,
Patent Owner

Case IPR2024-01165
U.S. Patent No. 10,958,840

**DECLARATION OF HASSAN FOROOSH PH.D.
IN SUPPORT OF PATENT OWNER'S RESPONSE**

PUBLIC VERSION

1. I have been retained on behalf of GoPro, Inc. (“Patent Owner”) in connection with the above-captioned *inter partes* review (IPR). I have been retained to provide my opinions in support of Patent Owner’s Response. I am being compensated for my time at my standard rate of \$650 per hour. I have no interest in the outcome of this proceeding.

2. In preparing this declaration, I have reviewed and am familiar with the Petition for IPR2024-01165, U.S. Patent No. 10,958,840 (“the ’840 patent”) and its file history, and all other materials cited and discussed in the Petition (including the declaration of Petitioner’s expert Dr. Aggelos Katsaggelos), and cited and discussed in this Declaration.

3. The statements made herein are based upon my own knowledge and opinion. This Declaration represents only the opinions I have formed to date. I may consider additional documents as they become available or other documents that are necessary to form my opinions. I reserve the right to revise, supplement, or amend my opinions based on new information and on my continuing analysis.

I. QUALIFICATIONS

4. I have been retained to serve as a technical expert in this matter based upon my educational and research experience, and specifically my work in the fields of computer vision, image processing, and optics. My curriculum vitae

PUBLIC VERSION

(“CV”) submitted as Exhibit 2002, which contains details about my education, experience, publications, and other qualifications to render an expert opinion in this matter.

5. I am a Professor of Computer Science at the University of Central Florida. My primary research interests include computer vision, image/video processing, and machine learning, including deep learning and large language models.

6. I received my Ph.D. in Computer Science from INRIA (French National Institute for Research on Informatics and Control), University of Nice Sophia Antipolis in 1996. I also earned a Master of Science in Computer Science from the University of Nice Sophia Antipolis in 1993, a Master of Science in Electrical Engineering from Kingston University, U.K., in 1990, and a Bachelor of Science in Electrical Engineering from Kingston University, U.K., in 1989.

7. Prior to research and teaching, I worked as an R&D Engineer at Innovative Engineering Solutions Ltd. in London from 1990 to 1992. Between 1993 and 1996, I served as a Research Assistant in the Scene Analysis and Symbolic Image Processing Group of INRIA, where I performed research in areas such as 3D super-resolution, sub-pixel image registration, and image-based depth and shape recovery and analysis. From 1997 to 2000, I served as an Assistant

PUBLIC VERSION

Research Scientist at the University of Maryland, College Park, where I investigated topics such as image registration (2D and 3D), super-resolution of video data for unmanned aerial vehicles, super-resolution of Synthetic Aperture Radar (SAR) and Foliage Penetrating SAR (FOPEN-SAR), and Automatic Target Detection, Tracking, and Recognition (ATD/R). During this time, I authored one of the most cited works on sub-pixel registration and alignment based on phase correlation: (Hassan Foroosh (Shekarforoush), Josiane Zerubia and Marc Berthod, Extension of Phase Correlation to Sub-pixel Registration, IEEE Transactions on Image Processing, vol. 11, Issue 3, pp. 188-200, 2002). Between 2000 to 2002, I served as Senior Research Scientist at the University of California, Berkeley where I also served as the Technical Coordinator and Principal Research Scientist of a Multi-University Research Initiative (MURI) funded by the Office of Naval Research. As part of the project, I conducted research on multi-sensor acquisition of 3D data for autonomous 3D modeling and visualization of urban environments.

8. I joined the Department of Computer Science at the University of Central Florida (UCF) as an Assistant Professor in 2002. Since then, I have served as an Associate Professor, Professor, and CAE Link Professor of Computer Science in the same department. I also served in various other capacities within the department, such as the Graduate Program Coordinator, chair of various leadership

PUBLIC VERSION

committees, and most recently in the Faculty Senate Strategic Planning Council for UCF.

9. I have taught a number of courses directed to computer vision and 3D modeling, including 3D Computer Vision, Image Processing, Special Topics in 3D Computer Vision, and Visual Simulation. Since joining UCF, I have also advised and chaired the dissertation of over 30 doctoral students, 28 of whom have already successfully defended their doctoral dissertation, and secured academic positions or industry jobs with major high-tech corporations. As one example, I served as the advisor for a dissertation directed to “*Improving Matching and Classification Through Deep Learning of Structure and Varying Illumination*”. This dissertation involved also working with GoPro camera systems for mobile aerial platforms.

10. I have been the principal investigator or co-principal investigator of numerous research projects funded by the National Science Foundation (NSF), NASA, DoD, and industry partners, with total funding exceeding \$14M since joining UCF, and ongoing funded projects for an additional \$3M. As one related example, I recently served as the principal investigator in two projects funded by the National Science Foundation (NSF) on 3D modeling and camera systems, and the role of 3D information in machine understanding of digitized 3D world.

PUBLIC VERSION

11. I have also conducted extensive research in partnership with industry in the field of computer vision, in particular 3D computer vision and camera systems, including body-worn cameras and cameras on agile mobile platforms, or cameras operating in various different modalities (EO, FLIR, polarized, etc.). Below, I highlight some of the projects that were funded by industry and state-sponsored technology transfer initiatives that specifically involved using GoPro camera systems or similar platforms. For example, from 2012 to 2018, I received funding from Jackson Technologies related to Airborne Multisensor Visual Data Processing, which studied automated methods of reconstructing 3D panoramic models around a flying aircraft that used the GoPro camera system and software tools extensively; From 2013 to 2018, I received related funding from the Florida High Tech Corridor (“FHTC”) for Rapid Automated 3D Terrain Modeling from Airborne Data, which similarly studied automated methods of reconstructing 3D models of a terrain from a flying aircraft, again using extensively the GoPro cameras and software tools; From 2008 to 2009, I received funding from Electronic Arts – Tiburon related to Modeling and Tracking Human Activities in Sport Events. The goal of this project was to develop algorithms for tracking and recognition of human activities in sport events, using body-worn cameras similar

PUBLIC VERSION

to GoPro systems, as well as surrounding cameras, network of cell phone cameras, and other sensors.

12. I have authored or co-authored over 200 articles in peer reviewed journals or conference proceedings, relating to the areas of Computer Vision (CVPR, ICCV, ECCV, etc.), Machine Learning (ICLR, AAAI, ACL, EMNLP, etc.), and Image Processing and Pattern Recognition (ICIP, ICPR). Some related example publications include: Yang Zhang, Zixiang Zhou, Philip David, Xiangyu Yue, Zerong Xi, Boqing Gong, Hassan Foroosh, PolarNet: An Improved Grid Representation for Online LiDAR Point Clouds Semantic Segmentation, Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020; Chuan Sun, Imran Junejo, Marshall Tappen, Hassan Foroosh, *Exploring Sparseness and Self-Similarity for Action Recognition*, IEEE Transactions on Image Processing, vol. 24, issue 8, pp. 2488-2501, 2015; Chen Shu, Luming Liang, Wenzhang Liang, Hassan Foroosh, *3D Pose Tracking with Multi-Templates Warping and SIFT Correspondences*, IEEE Trans. on Circuits and Systems for Video Technology, 2015; Nazim Ashraf, Chuan Sun, and Hassan Foroosh, *View-Invariant Action Recognition Using Projective Depth*, Journal of Computer Vision and Image Understanding (CVIU), vol. 123, pp. 41-52, 2014; Yuping Shen and Hassan Foroosh, View-invariant Action Recognition from Point

PUBLIC VERSION

Triplets, IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), vol. 31, no. 10, pp. 1898-1905, 2009; Nazim Ashraf, Yuping Shen, Xiaochun Cao, and Hassan Foroosh, *View-Invariant Action recognition using Weighted Fundamental Ratios*, Journal of Computer Vision and Image Understanding (CVIU), vol. 117, pp. 587-602, 2013.

13. Based on my work, I have also applied for and received several patents in the field of image processing, including U.S. Pat. No. 8,755,569 entitled Methods for Recognizing Pose and Action of Articulated Objects with Collection of Planes and Motion, which was issued in 2014.

14. I have served as the editor and/or committee member on numerous professional organizations. For instance, I have been on the editorial board of the IEEE Transactions on Image Processing (Associate Editor, 2003-2008, 2011-2016); IEEE Conference on Computer Vision and Pattern Recognition (CVPR) serving as area chairs, session chairs and Technical Committee member over many years (2004-2012, 2015-2016, 2019-2023); Organizing Committee and sponsorship chair of ACM Multimedia Conference in 2014; and frequent reviewer on various IEEE transactions, including the IEEE Transactions on Pattern Analysis and Machine Intelligence.

PUBLIC VERSION

15. For my work, I have been honored with numerous awards, including the *Piero Zamperoni Award*, International Association for Pattern Recognition (IAPR) (2004). I have also received several Best Paper awards, including “*The Best Scientific Paper Award*” in the IAPR International Conference on Pattern Recognition in 2008, “*The Best Paper Award*” in the IEEE International Conference on Image Processing (ICIP) in 2018, and the *Distinguished Researcher*, UCF College of Engineering and Computer Science (2004-2005, 2005-2006).

16. In sum, I have studied, taught, practiced, and conducted research in the areas of computer vision, image processing, machine learning (deep learning and large language models), and optics.

II. LEGAL UNDERSTANDING**A. The Person of Ordinary Skill in the Art (“POSITA”)**

17. I understand that a person of ordinary skill in the relevant art at the time of the invention (also referred to herein as “ordinary artisan”) 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.

18. 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:

PUBLIC VERSION

- 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.

19. I have reviewed the declaration of Dr. Katsaggelos, including his opinions regarding the Person of Ordinary Skill in the Art, which is a person with “a bachelor’s degree in engineering or computer science, or equivalent field, and two years of experience working on image or video processing or digital cameras” or alternatively “four years of experience working on image or video processing, or digital cameras.” Ex. 1003, ¶¶31-33.

20. My opinions set forth in this Declaration apply Dr. Katsaggelos’s definition of a POSITA.

21. I am well-qualified to determine the level of ordinary skill in the art. I was a person of at least ordinary skill in the art at the time of the September 2018 provisional application priority date of the ’840 patent under Dr. Katsaggelos’s definition of a POSITA. If I do not explicitly state that my statements below are based on this timeframe, all of my statements are to be understood as an ordinary artisan would have understood something as of the priority date of the ’840 patent.

PUBLIC VERSION**B. Legal Principles**

22. I am not a lawyer and will not provide any legal opinions. Though I am not a lawyer, I have been advised that certain legal standards are to be applied by technical experts in forming opinions regarding the meaning and validity of patent claims.

1. Anticipation

23. I understand that to obtain a patent, a claimed invention must have, as of the priority date, been novel. I understand that to prove that a prior art reference anticipates a patent, the four corners of the prior art reference must describe element of the claimed invention, either expressly or inherently. I also understand that the prior art reference must describe the elements arranged as in the claim and that it is improper to combine disparate embodiments to establish anticipation unless a POSITA when reading the reference could at once envisage such a combination.

2. Obviousness

24. I understand that to obtain a patent, a claimed invention must also have, as of the priority date, been nonobvious in view of 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

PUBLIC VERSION

whole would have been obvious at the time the invention was made to a person having ordinary skill in the art.

25. 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, make 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 to create the inventions claimed in the asserted claim, including how there would have been a reasonable expectation of success.

26. I understand that a patent composed of several elements is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art, and that obviousness cannot be based on the hindsight combination of components selectively culled from the prior art to fit the parameters of the patented invention.

27. I also understand that a reference may be said to teach away when a person of ordinary skill, upon reading the reference, would be discouraged from following the path set out in the reference, or would be led in a direction divergent from the path that was taken by the applicant. Even if a reference is not found to teach away, I understand its statements regarding preferences are relevant to a

PUBLIC VERSION

finding regarding whether a skilled artisan would be motivated to combine that reference with another reference.

3. *My Understanding of Claim Construction Law*

28. I have been informed that patent claims are construed from the viewpoint of a person of an ordinary artisan and that patent claims generally should be understood consistent with their ordinary and customary meaning at the time of the invention. A review of the patent claim language, the patent specification, and its prosecution history are also necessary to determine the proper meaning and scope of the term at issue.

29. I have further been informed that in the specification and prosecution history the patentee may define a claim term in a way that differs from the ordinary and customary meaning. I understand that during prosecution of the patent before the U.S. Patent and Trademark Office, the Applicant may make representations or provide definitions of terms that may affect the scope of the patent claims. In particular, the Applicant may, during the course of prosecution, limit the scope of the claims to overcome prior art and/or disavow claim coverage by making clear and unambiguous statements to that effect.

30. I have been informed that an ordinary artisan may, among other things, consider dictionaries, publications, other patents, and treatises (*i.e.*,

PUBLIC VERSION

“extrinsic evidence”). I understand that extrinsic evidence may generally not be relied on if it contradicts the meaning of claim language provided by the intrinsic evidence, such as express definitions given for terms in the specification.

4. *Secondary Considerations of Non-Obviousness*

31. I understand that one of the factual inquiries when determining the obviousness of a patent is to consider objective evidence related to the issue of obviousness, sometimes called “secondary considerations of non-obviousness” or “objective indicia of non-obviousness.” I understand that such evidence can often be probative of the issues of obviousness.

32. I understand that examples of objective evidence include (1) commercial success, (2) copying, (3) industry praise, (4) skepticism, (5) long-felt but unsolved need, and (6) failure of others.

33. I understand that there must be a nexus between the patented invention and the evidence of secondary considerations. I further understand that there is a presumption of such a nexus if the claims of the patent are coextensive with a product or products that are the subject of the secondary considerations evidence. I further understand that in order for a claim to be coextensive with a product, one must look at the claims at the whole and not just the new claimed

PUBLIC VERSION

feature. In addition, the claimed invention should not be merely a subcomponent of the products sold but should encompass the entire apparatus.

III. CLAIM CONSTRUCTION

34. I understand that: in an *inter partes* review, claim terms receive their plain and ordinary meaning in light of the specification of the patent in which they appear. Petitioner states that no constructions are necessary. See Pet. at 6-7. However, a dispute has arisen regarding the meaning of the “low-light high-pass parameter.”

35. Based on the preliminary arguments in the related IPR2024-01164, the Board understood a “low-light high-pass parameter” as “a parameter that controls the amount or strength of smoothing used in motion blur compensation.” (I.D. at 8-9). While the Board correctly notes that a low-light high-pass parameter is a parameter that controls the amount or strength of smoothing used in motion blur compensation, I understand that it is also important that the claim terms “low-light” and “high-pass” be accorded meaning.

36. In my opinion, Patent Owner’s proposed claim construction set forth below is consistent with the meaning of those terms in view of the specification and the understanding of a POSITA at the time of the invention:

PUBLIC VERSION

Claims term	Patent Owner’s Proposed Construction
“low-light high-pass parameter”	A parameter that controls the amount or strength of smoothing used in motion blur compensation by allowing more high frequencies to pass through based on low-light conditions.

37. In my opinion, consistent with the plain claim language, GoPro’s “low-light high-pass” parameter is a targeted parameter that selectively addresses blur based on “low-light” conditions while otherwise maintaining smoothing strength. GoPro’s parameter also controls smoothing in a specific manner, *i.e.* as a “high-pass” parameter which allows more high frequencies to pass as opposed to simply attenuating the smoothing strength across the entire range of frequencies.

38. In my opinion, GoPro’s proposed construction is supported by how the “low-light high-pass” parameter is described in the intrinsic evidence. The specification explains how the “low-light high-pass” parameter addresses an issue that occurs specifically in low-light conditions due to longer exposure times. Ex. 1001 at 18:1-19:59; *see also* at 19:18-21 (“[L]onger exposure time may be used with low light conditions and shorter exposure time may be used with high light conditions.”). As the patent explains, camera motion during the exposure period—termed “intra-frame” motion—can become “encrusted” in the video frame, leading to undesirable blur artifacts. *Id.* at 18:1-24; 21:66-22:7. Strong smoothing in this

PUBLIC VERSION

situation can cause the motion blur to become more evident, creating what the patent calls “motion blur pulsations” and “poor visual impression.” *Id.* at 22:2-5.

39. To address this particular cause of motion blur, the specification describes when and how the “low-light high-pass” parameter is utilized. For example, the amount of smoothing/stabilization may vary based on the frequency of motion as well as the exposure time. *Id.* at 18:43-19:26; Figure 7. As a result, “[c]ertain portion[s] of high frequencies in the input may be kept based on the image capture configuration, such as exposure start time and exposure duration time . . .” *Id.* at 17:61-18:1. Thus, the output motion can be controlled to “include high frequencies that mask motion blur within the visual content.” *Id.* at 22:16-21. By controlling smoothing in this targeted manner, the claimed “low-light high-pass” parameter addresses motion blur arising from long exposure times, *id.* at 22:8-21, while still maintaining smoothing strength for other types of motion subject only to the “margin constraint,” *id.* at 21:64-66.

40. GoPro’s proposed construction is also consistent with the extrinsic evidence. There is no dispute that “low-light” conditions, as indicated by longer exposure times, is a possible cause of blur. In his declaration, Petitioner’s expert Dr. Katsaggelos states: “A POSITA at the time of the ’840 patent filing would understand that blurs can be caused by low-light conditions because a shutter is

PUBLIC VERSION

open for a longer time in video capture under low-light conditions, causing movement to add more blur to a frame.” Ex. 1003 at ¶ 111; ¶ 237 (same quote); and ¶ 123 (“as explained above, blur may be caused by longer device shutter integration time under low light conditions.”).

41. “Low-light,” however is just one of many potential causes of blur, as Dr. Katsaggelos admitted at his deposition. Ex. 2020 at 103:22-104:8 (testifying that motion blur can be caused “due to many different reasons”). Given that the term specifically recites a “*low-light*” high pass parameter, GoPro’s proposed construction appropriately distinguishes the claimed parameter from techniques for mitigating blur due to other factors.

42. Patent Owner’s proposed construction is also consistent with the well-understood meaning of “high-pass.” As Dr. Katsaggelos confirmed at his deposition, a “high-pass” system is one that allows high frequencies to pass through the smoothing process while still attenuating the low frequencies. *Id.* at 94:6-15 (Q. “[W]ould a high pass system be one that allows most of the information of the signal at high frequencies to go through while attenuating the information at the low frequency of the signal? . . . THE WITNESS: That would be consistent with the first definition. It’s a matter of what is attenuated and what is allows to go through with some smaller distortion, yeah.”) Further, when teaching

PUBLIC VERSION

his students, Dr. Katsaggelos has described “high-pass” systems as those that “allow the high frequencies to go through the system unchanged.” Ex. 2022 (Coursera screenshot).

43. For all these reasons, in my opinion, the term “low-light high-pass parameter” should be construed to mean: “a parameter that controls the amount or strength of smoothing used in motion blur compensation by allowing more high frequencies to pass through based on low-light conditions.”

I. OVERVIEW OF THE '840 PATENT

44. The '840 patent, titled “Systems and Methods for Stabilizing Videos,” discloses an advanced system for stabilizing video content in diverse circumstances and conditions. The system measures position information characterizing the rotational positions of an image capture device during the capture duration. This allows the system to generate a “trajectory of the image capture device during the capture duration,” which is then used to “determine a smoothed trajectory of the housing” in accordance with the further claim limitations. *E.g.*, Ex. 1001, claim 1.

45. Among the '840 patent’s key innovations is the introduction of tunable parameters that allow the system to adaptively modulate the stabilization process based on environmental conditions and/or the characteristics of the

PUBLIC VERSION

captured motion. Ex. 1001, 21:17-25. These tunable parameters include the “low-light high-pass parameter” (*id.*, 21:59-22:21), “weight-balance parameter” (Ex. 1001, 17:32-60, 21:17-58), and “stickiness parameter” (*id.*, 22:22-41), which enable fine tuning of the stabilization process under diverse circumstances and conditions. *Id.*, 30:35-38 (claim 1), 32:44-48 (claim 16).

Low-light high pass parameter

46. The low-light high-pass parameter is targeted to the specific challenge of stabilizing videos captured in low-light environments by controlling the extent to which the system compensates for motion blur that occurs during frame exposure. Ex. 1001, 21:59-22:21.

47. A central concern addressed by the patent is how to stabilize video in “low-light” conditions, where longer exposure times increase the likelihood of motion blur. The specification describes how camera motion during the exposure period—termed “intra-frame” motion—can become “encrusted” in the video frame, leading to undesirable blur artifacts. *Id.* at 18:1–24; 21:66-22:7. In other words, strong smoothing in this situation creates what the patent calls a “poor visual impression.” *Id.* at 22:2–5. The goal of the invention, therefore, is to stabilize in a way where “intra-frame motion is taken into account so that the

PUBLIC VERSION

smoothed trajectory is smooth for inter-frame phases while corresponding to the original motion for intra-frame phases.” *Id.* at 18:18-20.

48. To mitigate the low-light blur problem, the specification describes modifying the smoothing strength based on exposure time and other factors such as motion of the image capture device. *Id.* at 18:25–46. For example, in some implementations, a certain portion of “***high frequencies in the input may be kept based on the image capture configuration, such as exposure start time and exposure duration time***” to compensate for “***motion blur and/or lowlight image capture conditions***.” *Id.* at 17:61–18:9.

49. Figure 7 and its accompanying discussion further elaborate on this by illustrating how smoothing values (S1, S2, S3, S4) can be modulated depending on exposure time and frequency of motion within the frame. *Id.* at 18:43-19:59. In particular, the specification explains that “***longer exposure time may be used with low light conditions***,” and that “***higher exposure time may result in greater amount of motion blur***.” *Id.* at 19:18-21; 19:25–26. Additionally, intra-frame motion “higher than or the same as [a] high motion” threshold is considered high-frequency motion to further tune strength of smoothing. *Id.* at 18:64-19:2.

50. Within this context of adaptive smoothing in low-light conditions, the specification introduces the claimed “low-light high-pass parameter” as a

PUBLIC VERSION

parameter that “controls the amount/strength of smoothing performed to generate the smoothed trajectory.” *Id.* at 21:59-61. The specification reiterates the distinction between “inter-frame motion,” which is smoothed subject only to a margin constraint, and “intra-frame motion” which is addressed by a higher value of the low-light high-pass parameter to let the smoothed trajectory follow the observed trajectory more closely. *Id.* at 21:61-22:21. To do so, the parameter reduces the strength of smoothing in a “high-pass” manner—*i.e.* a manner that maintains more high frequency motion while continuing to attenuate low frequency motion. *Id.* at 22:8–25. (“[T]he output motion may include high frequencies that mask motion blur within the visual content.”).

51. Thus, as described in the specification, the low-light high-pass parameter is targeted to a specific issue—stabilizing video captured in low-light conditions, where longer exposure durations create intra-frame blur. It does so by modifying the strength of smoothing in a “high-pass” manner that seeks to preserves high frequency motion, thereby reducing the appearance of blur in this specific low-light circumstance, while attempting to maintain the stabilization strength for low frequency motion in other circumstances.

52. The result is a more natural and visually pleasing video output for GoPro’s specific use cases. As used in the claims, the “low-light high-pass

PUBLIC VERSION

parameter” is therefore best understood as “a parameter that controls the amount or strength of smoothing used in motion blur compensation by allowing more high frequencies to pass through based on low-light conditions.”

Weight-balance parameter

53. The weight-balance parameter provides a mechanism for stabilizing video by determining how much emphasis to place on minimizing rotational velocity versus rotational acceleration. Ex. 1001, 17:32-60, 21:17-58. This allows the system to adapt to different types of motion and produce smoother, more natural video. *Id.*

54. The specification explains that “the smoothed trajectory may be determined based on minimization of a rotational velocity . . . and a rotational acceleration of the image capture device . . . while respecting a set of constraints.” *Id.* at 17:32–37. For example, a “smooth path” of yaw, pitch, or roll angle may be generated by “finding a path of the image capture device . . . that minimizes a combination of a time derivative, a second time derivative, and/or other time derivative(s)” of these angles. *Id.* at 17:37–53.

55. The specific combination that is minimized is set by the weight-balance parameter, which “may range between 0 and 1, with 0 corresponding to minimization of angular velocity and 1 corresponding to minimization of angular

PUBLIC VERSION

acceleration.” *Id.* at 21:36–39. Intermediate values allow the system to minimize both types of motion, “with minimization of different types of motion weighed according to the value.” *Id.* at 21:39–42. The patent explains the practical effect of different values:

- Minimizing **only angular velocity** (value = 0) may lead to trajectories with sharp, piecewise-linear corners that, as the patent explains, “prefer[] to keep a steady position when possible” and “minimize apparent motion between video frames”—but the resulting motion “may appear unnatural.” *Id.* at 21:44–50.
- Minimizing **only angular acceleration** (value = 1) may reduce corner points and result in smoother acceleration, but it “may produce more motion than [s]ole minimization of angular velocity.” *Id.* at 21:50–54.
- Minimizing **a combination** (value between 0 and 1) produces a trajectory that blends these two effects, with weights determined by the parameter value. *Id.* at 17:54–58.

56. In this way, the weight-balance parameter gives the system the ability to fine-tune stabilization based on the type of motion encountered, enabling more natural and visually appealing video.

Stickiness parameter

57. The stickiness parameter “control[s] the extent by which preceding portion(s) of the (observed) trajectory impacts determination of the smoothed trajectory.” *Id.* at 22:22–25. The parameter may range between 0 and 1—where a value of 1 may result in a “lazy” stabilizer that “maximizes the impact of the preceding motion/position,” and a value of 0 may result in the smoothed trajectory

PUBLIC VERSION

being generated without taking the preceding motion into account. *Id.* at 22:25–33.

Values between 0 and 1 provide a flexible weighting of the preceding motion’s influence, allowing for a dynamic balance between responsiveness and stability.

The parameter may also “impact bootstrapping”—that is, the way the system initially estimates the smoothed trajectory based on prior motion. *Id.* at 22:36–42.

58. Together, these parameters reflect a sophisticated, configurable approach to digital video stabilization—one capable of adapting to varied motion types and lighting conditions.

Temporal horizon enabling approximation of intentional motion

59. In addition, unlike conventional approaches that distinguish between stationary and moving cameras to infer the presence of jitter, the ’840 patent’s “***temporal horizon enabling approximation of intentional motion***” is directed to leveraging a duration of future motion to estimate user intent at a given moment, regardless of whether the camera is stationary or in motion.

60. The ’840 patent explains that “intentional motion may include a forward motion to follow an object of interest, a panning motion to shoot a wide view of a scenery, and/or other motion,” while “unintentional motion may result in noise . . . caused by shaking hands, vibrations, bumps.” Ex. 1001 at 9:8-16. To distinguish between these types of motion, the ’840 patent teaches that “intentional

PUBLIC VERSION

motion may be characterized by lower frequency motion than unintentional motion,” and that “a sufficiently long temporal horizon of experienced motion may need to be used” to identify it. *Id.* at 9:17-23.

61. The specification explains that “positions/motions of the image capture device in the ‘future’ may be analyzed (look ahead) to determine whether particular position(s)/motion(s) of the image capture device at a moment is an intended motion or an unintended motion.” *Id.* at 15:1-6.

62. Figure 5C provides a useful illustration of how the temporal horizon enables approximation of whether a given movement is intentional or not—regardless of the state of the camera (e.g., stationary or in motion).




FIG. 5C

PUBLIC VERSION

63. In processing durations A, B, and C for example, the system's ability to look-ahead across a sufficiently long temporal horizon of the observed trajectory (400) allows the system to infer that measured motion within these durations was, in fact, unintentional and process the smoothed trajectory (532) accordingly. In processing duration D, by contrast, the temporal horizon allows the system to understand that the detected motion persists into the future and is therefore intentional. Thus, the system can immediately begin processing a trajectory at the beginning of duration D that smoothly tracks the path of intended motion. The description in the specification is instructive:

Based on the subsequent portion(s) (look ahead) of the observed yaw trajectory 400, the smoothed trajectory component 104 may determine that the staggered rotation of the image capture device to the right during durations 528, 530 included unintended motions (non-continuous rotation to the right), and may determine the smoothed yaw trajectory 532 to include continuous rotation to the right during durations 528, 530.

64. Thus, the “temporal horizon *enabling approximation of intentional motion*” requires inferring intentional versus unintentional motion in a specific manner. The claim requires more than simply comparing detected motion in a current frame to a motion threshold, or even comparing the difference in motion between two successive frames to a motion threshold. The claim requires looking

PUBLIC VERSION

ahead over a temporal horizon that is sufficiently long to enable an approximation of intentional motion in producing the smoothed trajectory.

65. In my opinion, these innovations are integral to GoPro's HyperSmooth technology, which has played a significant role in the commercial success and industry recognition of GoPro's cameras.

II. OVERVIEW OF THE CITED ART

66. The petition challenges the '840 patent's claims 1 to 21 on four grounds. Ground 1 asserts that the combination of Zhou and Cai renders independent claims 1 and 16, along with certain dependent claims, obvious. Grounds 2 through 4 address the remaining dependent claims by asserting the combination of Zhou and Cai along with one other reference per ground. Below are short summaries of the primary references .

B. Ground 1 Cited Art**1. Zhou (Ex. 1006)**

67. While Zhou's background states that its disclosure "relates to techniques for improved stabilization of video frames captured in low light conditions by taking into account motion blur within the video frames," Zhou accomplishes this in a materially different way than the '840 patent. Specifically, Zhou does not involve a parameter for controlling the amount or strength of smoothing by allowing more high frequencies to pass through based on low-light

PUBLIC VERSION

conditions. Rather, in its “Summary,” Zhou discloses methods that adjust a stabilization strength value based on an “estimated amount of blur” detected in an image. *Cf.* 1:14-17 with 2:1-34.

68. As Zhou teaches, “motion blur” can be caused by a number of different factors independent of low-light, such as by “hand shaking,” “streaking,” or “rapid movement.” *Id.* 1:46-55. Zhou’s blur vector metric is designed to address blur as a result of any one of these causes, without regard to lighting conditions. In fact, the calculation used by Zhou to apply blur correction confirms that the approach is not based on low-light conditions. Rather, to estimate “image motion blur,” Zhou uses gyroscope samples and calculates an “image row blur vector” for each row of pixels in an image. *Id.* at 5:26-29. The magnitude of the image row blur vector (the “image row blur vector metric”) is calculated based simply on how much the image translated in the horizontal (“tx”) and vertical (“ty”) directions. *Id.* at 5:65-6:1. Zhou then calculates an overall image blur metric based on a weighted average of the image row blur vectors for the image. *Id.* at 6:11-14. This overall image blur metric may be averaged across multiple frames to avoid abrupt changes. *Id.* at 6:27-30. Thus, Zhou’s overall image blur metric is not targeted to blur caused specifically by low-light conditions. Rather Zhou calculates blur based

PUBLIC VERSION

on image translation, independent of lighting conditions, and thus addresses a different problem to that solved by the '840 patent.

69. Zhou also does not address blur in a “high-pass” manner as required by the claim language. Zhou’s overall image blur metric is used to calculate a stabilization strength “S” for each frame entry in the buffer. *Id.* at 6:38-7:16, 7:57-60. This stabilization strength parameter is applied as part of a calculation to determine the extent to which any given frame should move in accordance with the stabilization operation. *Id.* at 7:24-38. Importantly, however, the stabilization strength parameter is applied without regard to the frequency of motion in a given frame. In other words, the blur-related attenuation of smoothing Zhou—based on the stabilization strength parameter calculated from the image blur metric—is applied across-the-board to both low-frequency and high-frequency motion. The stabilization strength parameter therefore does not act as a “high-pass” parameter that allows more high frequency motion to pass through.

70. Zhou also does not disclose a “stickiness parameter” that enables weighting of past frames to control the extent by which preceding portion(s) of the trajectory affect determination of the smoothed trajectory, as claimed in the '840 patent. Zhou describes the use of buffering to obtain data from past and future frames “to adjust the strength of a stabilization metric” for a current frame. *Id.* at

PUBLIC VERSION

Abstract, 3:63-4:4. However, the mere fact that Zhou considers some symmetric number of past and future frames, based on buffer size, in implementing its stabilization technique does not teach or suggest a stickiness parameter for weighting the *extent* to which those past frames are considered.

71. Zhou’s lack of disclosure of the low-light high-pass, weight balance, and stickiness parameters is reflected in the fact that Zhou was cited during prosecution, and the Examiner allowed the ’840 patent over it. See Appl. No. 16/713,798, List of References Cited by Examiner (Nov. 30, 2020).

2. Cai (Ex. 1009)


72. Cai relates to stabilizing video and “in particular to detecting motion in successive video frames to stabilize video that has ***no motion***.” *Id.* at 1:14-16. (Emphasis added). Video that has “no motion” refers to video taken with a stationary camera, in contrast with a camera in motion. The system is designed for handheld portable cameras, where—unlike professional setups with stable platforms— “the photographer may not be able to provide a stable platform.” *Id.* at 1:20-34. In such a case, when a camera is stationary even small hand movements can result in noticeable jitter.

73. To address the issue of jitter, Cai distinguishes “zero-motion” (stationary camera) situations from non-zero motion (camera in motion) situations

PUBLIC VERSION

by applying a low-pass filter followed by a motion threshold comparison. *Id.* at 4:12-5:28. When the inter-frame movement falls below the threshold, Cai treats the video segment as a “zero-motion” scenario. *Id.* at 4:36-39.

74. All of Cai’s methods for zero-motion detection are limited to comparing motion in past and current frames to the threshold. This is shown in Figure 5, whose “Zero Motion Detection” block (522) takes only “Previous Frames & Current Frame” as its input:




Ex. 1009 at Fig. 5 (annotation added).

75. Cai’s trajectory planning block *only* considers future frames as part of comparing them to the output of the zero motion detection block to detect the possibility of what Cai refers to as a “jitter-clamped” frame. *Id.* at 8:11-25. This

PUBLIC VERSION

clamping check has nothing to do with approximating intentional motion, but rather is a teaching in Cai for predicting where the difference between the original trajectory and the filtered trajectory may exceed the compensation limit. To address this issue, Cai introduces an “anchor point” to avoid a discontinuity or “sudden jump” in the smoothed trajectory. *Id.* at 3:16-29, 7:58-8:32, Fig. 6. Thus, Cai does **not** consider future frames for its zero motion detection algorithm, and instead limits its consideration of future frames for a completely different functionality that is unrelated to the ’840 patent claims.


76. That Cai’s zero-motion detection process does not contemplate a consideration of future frames across a temporal horizon is made abundantly clear by the “false alarm” embodiment shown in Figure 4:



PUBLIC VERSION

Ex. 1009 at Fig. 4. In Figure 4, the displacement between any two coordinates of the low-pass filtered trajectory are small, which is initially registered by Cai as zero-motion. *Id.* at 6:4-12. As Cai teaches, however, its “false alarm” logic eventually detects that these small displacements have accumulated which causes the planned trajectory to begin following the low pass filtered path. Critically, however, notice how Cai’s planned trajectory initially fails to identify the intentional motion, and only begins to follow the intended path after a “lag” period.

77. This confirms that Cai’s approach fails to capture the key benefit of the ’840 patent’s temporal horizon limitation. Compare, for example, Cai’s false alarm embodiment in which the planned trajectory lags the intentional motion with Figure 5C of the ’840 patent, where the smoothed trajectory immediately begins following the intended path of motion:



PUBLIC VERSION

78. Relatedly, Cai’s approach of distinguishing between camera states based on a motion threshold merely makes a binary determination about whether a current frame was intended to be zero-motion or not. All detected motion for a frame above the threshold is deemed intentional, while all detected motion below the threshold is deemed unintentional. Although this technique may crudely identify unintentional jitter for frames when the camera is stationary, it does not enable a smoothed trajectory that approximates intentional motion as taught and claimed by the ’840 patent. Cai has no mechanism for distinguishing between intentional and unintentional motion above the threshold, and thus no ability to approximate the user’s intended trajectory for a camera in motion.

3. Bell Paper (Ex. 1004)

79. The Bell Paper discloses “a method for video stabilization . . . for videos captured on mobile devices.” Ex. 1004 at 294 (Abstract). The disclosed method uses “data from an on-board gyroscope to track the camera’s *angular velocity*.” *Id.* (Emphasis added). The Bell Paper explains: “We remove small motions and rolling-shutter distortions due to hand shake,” and for “larger motions, we filter the camera’s *angular velocity* to produce a smooth output.” *Id.* (Emphasis added).

PUBLIC VERSION

80. The Bell Paper applies a “nonlinear filtering method which completely flattens small motions regardless of frequency, and performs low-pass smoothing when the virtual camera must move to keep the crop window inside the input frame.” *Id.* at 296. It explains that the “gyroscope provides a series of discrete ***angular velocity*** measurements with time-stamps, which we integrate to produce a function of time that describes the camera orientation.” *Id.* at 297. (Emphasis added).

81. The Bell Paper does not perform any optimization or minimization of any rotational quantity. Rather, the Bell Paper’s algorithm calculates the orientation of a virtual camera frame based on the prior orientation and angular velocity of the virtual camera frame. *Id.* at 299. A “decay factor” is used to dampen movement of the virtual camera when its boundaries fall within an inner region of a padding space, and a mixing weight is used to accelerate the virtual camera when its boundaries fall outside the inner region. There is no mechanism in Bell for determining a path that minimizes any rotational quantity.

82. Notably, in a parallel case IPR2024-01164, Petitioner relied upon the Bell patent (U.S. 10,027,893), whose algorithms are substantially the same as those in the Bell Paper, but Petitioner did *not* contend that the Bell patent disclosed or

PUBLIC VERSION

suggested the weight-balance parameter, relying instead on a combination with a different reference (Kwatra) in an effort to demonstrate this limitation.

83. Also, while the Bell Paper uses angular velocity as part of its algorithm, it expressly discourages factoring in acceleration measurements as being “difficult and of limited value”:

In theory we could be more precise by also measuring translation with the device’s accelerometer, but in practice this is difficult and of limited value.

Id. at 297.

84. In its “Results” section, the Bell Paper discusses a single example (a walking video) and provides plots of the angular velocity versus frame number. *Id.* at 303-304. The Bell Paper observes that its method removed “most of the angular acceleration,” resulting in a “smoothed camera path [that] is nearly free of rotational motions.” *Id.* at 304. Notwithstanding this generalized observation, the Bell Paper nowhere discloses measuring, calculating, or minimizing rotational acceleration for a smoothed path. Nor does the Bell Paper teach or suggest any parameter or method directed to balancing the extent to which rotational velocity and rotational acceleration are minimized.

PUBLIC VERSION

III. GROUND 1: PETITIONER FAILS TO ESTABLISH THAT ZHOU AND CAI RENDER CLAIMS 1, 14 AND 16 OBVIOUS**A. Claims 1, 14, and 16: Zhou Does Not Disclose Nor Render Obvious “Determine a Smoothed Trajectory of the Housing Based On” Either A “Low-Light High Pass Parameter” Or A “Stickiness Parameter”**

85. Claims 1 and 16 require the stabilization system to “determine a smoothed trajectory of the housing based on a look-ahead of the trajectory and one or more of a weight-balance parameter, *a low-light high-pass parameter, and/or a stickiness parameter.*” (Emphasis added). Claim 14 further requires that “*the low-light high-pass parameter* controls [the] amount or strength of smoothing performed in smoothed trajectory generation and an extent of motion blur compensation applied in determination of the smoothed trajectory.” (Emphasis added).

PUBLIC VERSION

1. *Zhou Does Not Disclose the Claimed Low-Light High-Pass Parameter*

86. As discussed above, the claimed “low-light high-pass parameter” means “a parameter that controls the amount or strength of smoothing used in motion blur compensation by allowing more high frequencies to pass through based on low-light conditions.” This construction gives meaning to the terms “low-light” and “high-pass,” and is consonant with both the intrinsic and extrinsic evidence.

87. The Petition asserts that Zhou’s “stabilization strength parameter” controls the amount of smoothing and the extent of motion blur compensation and is therefore the claimed “low-light high-pass parameter.” However, under the proper construction of “low-light high-pass parameter,” Zhou discloses no such parameter. This is evident for at least two reasons.

88. First, while Zhou recognizes that problems with blur are exacerbated in low-light conditions, Zhou’s “stabilization strength parameter” is not calculated or adjusted based on low-light conditions. Rather, Zhou discloses that it adjusts its stabilization strength parameter in response to its estimate of the amount of blur incurred (*i.e.*, an image blur metric). *Id.* at 6:38-7:16, 7:57-60. Zhou’s method of calculating the image blur metric is independent of lighting conditions, and instead based on the degree the image translates within a frame. Ex. 1004 at 5:65-6:1.

PUBLIC VERSION

89. At deposition, Petitioner’s expert Dr. Katsaggelos confirmed that Zhou’s stabilization strength metric is not directly tied to low-light conditions. Specifically, Dr. Katsaggelos admitted that blur can be caused by many different things including fast moving objects and that Zhou’s method for calculating blur is based on the amount the image translates rather than on the existence of low-light conditions directly. *See* Ex. 2020 at 121:12-22 (A. “So by measuring the blur, you *indirectly* address the low lighting . . .”); 122:23-123:3 (“fast moving objects or this or that can cause blur . . .”); 124:5-15 (Q “Do you see the calculation that it provides for blur? So what Zhou is teaching is that the amount of blur is calculated to be simply a function of the amount the image translates in each of the horizontal and vertical directions; correct? A. This is one embodiment that he teaches, yes.”).

90. Indeed, as Zhou itself recognizes, “motion blur” can be caused by many different factors, including “hand shaking,” “streaking,” or “rapid movement.” *Id.* 1:46-55. And, given that Zhou’s blur estimation logic is based entirely on the translation of an image within a frame, it is possible for Zhou to calculate a high value for the blur vector metric (resulting in a low value for the stabilization strength parameter) in well-lit conditions involving fast moving objects. Similarly, Zhou may calculate a low value for the image blur metric

PUBLIC VERSION

(resulting in a high value for the stabilization strength parameter) in poorly lit conditions involving little to no motion. This illustrates how Zhou’s stabilization strength parameter is calculated independently of lighting conditions, and thus does not operate as a “low-light high-pass parameter” as claimed.

91. Second, Zhou’s stabilization strength metric does not function as a “high-pass” system—*i.e.*, one allowing more high frequencies to pass through while continuing to attenuate low frequencies. Rather, Equations 1A and 1B in Zhou provide that for any given frame, Zhou’s stabilization strength parameter is applied to impact smoothing equally across the entire spectrum of motion frequencies. Ex. 1004 at 7:24-38. In other words, Zhou’s stabilization strength parameter makes no distinction between high-frequency and low-frequency motion in controlling the strength of smoothing. A POSITA would recognize that this is inconsistent with the stabilization strength parameter acting as a “high-pass” parameter that allows higher frequencies to pass through. Ex. 2020 at 93:23-94:15.

92. The differences between Zhou’s blur mitigation techniques and the those enabled by the low-light high-pass parameter of the claimed invention have real world implications. GoPro’s invention is directed to cameras designed for fast moving sporting activities like cycling or running where camera shake is a major issue. In these circumstances, it makes sense to maintain the strength of smoothing,

PUBLIC VERSION

even at the expense of some motion blur especially in fast-paced, high-motion scenes. Thus, GoPro’s low-light high-pass parameter is tailored to address a specific type of motion blur—*i.e.*, one that occurs in a “low-light” scenario characterized by longer exposure times—and does so in a targeted way—by allowing high frequencies to pass while attempting to maintain attenuation of lower frequency motion. For other scenarios, GoPro’s invention prioritizes smoothing camera shake over concerns about blur.

93. On the other hand, Zhou (an Apple patent) is directed at a solution where camera shake is less of an issue than clarity of the images. Accordingly, Zhou prioritizes blur correction, even if it comes at the expense of smoothing strength. Zhou thus reduces the smoothing strength regardless of the cause of motion blur and does so across the entire range of motion frequencies.

94. These are two different technologies and approaches for addressing blur with different trade-offs and utilities. One, GoPro’s claimed approach, is targeted to addressing a specific low-light, high frequency circumstance indicated by longer exposure times—and otherwise prioritizes smoothing, not blur correction. The other, Zhou’s approach, is a general blur detection approach that prioritizes removing blur irrespective of its cause.

PUBLIC VERSION

95. To be sure Zhou's disclosure teaches that its techniques can be useful in low-light conditions. But it does not teach or suggest a parameter directed specifically at this issue. For example, because Zhou's approach is based entirely on the amount of image translation in a frame, a POSITA recognizes that there are possible examples where Zhou may fail to recognize blur and maintain smoothing strength for very short, high-frequency motions even occurring in low-light conditions. Likewise, Zhou's approach would generally register blur and attenuate smoothing for low-frequency, fast moving objects even occurring in well-lit conditions. Zhou's parameter is simply configured differently from GoPro's "low-light high pass parameter," making a very different trade-off between smoothing and blur correction than the claimed "low-light high-pass" parameter. The fundamental difference between Zhou's estimated stabilization strength metric and the claimed "low-light high-pass" parameter is further highlighted in the way Zhou's metric is estimated and how it works. Zhou calculates an overall image blur metric based on a weighted average of the image row blur vectors for the image, and then further averages across multiple frames to avoid abrupt changes. *Id.* at 6:11-14, 6:27-30. Both steps of weighted averaging across image rows and also averaging across multiple frames are low-pass filters. This is at odds with the claimed "low-light high-pass" parameter. Consider, for example, a scenario where

PUBLIC VERSION

lighting changes between a frame and the adjacent ones so that the frame under consideration for stabilization becomes low-light and also blurred. Zhou clearly teaches that it would average (i.e. low-pass filter) its stabilization strength metric for this low-light blurred frame by averaging with high-light frames to avoid an abrupt change. *Id.* at 6:27-30. But this is not what the claimed “low-light high-pass” parameter is teaching. In fact, quite the opposite, the ‘840 patent would reduce the strength of blur compensation for the low-light blurred frame to let high frequencies pass. The opposite scenario may also occur, i.e. in Zhou a high-light blur-free frame considered for stabilization may get averaged with multiple low-light blurred frames. By averaging these frames (low-pass filtering), Zhou teaches that it would increase blur compensation for the high-light frame to avoid an abrupt change, whereas the claimed “low-light high-pass” parameter is designed to reduce blur compensation for a high-light frame. Accordingly, Zhou’s “stabilization strength metric” is not a “low-light high-pass parameter” as required by the claims.

2. *Zhou Does Not Disclose the Claimed Stickiness Parameter*

96. The ‘840 patent teaches that the claimed “stickiness parameter” is a parameter “that control[s] the extent by which preceding portion(s) of the (observed) trajectory impacts determination of the smoothed trajectory.” Ex. 1001 at 22:22-25. Zhou does not disclose any such parameter.

PUBLIC VERSION

97. Petitioner asserts that Zhou’s Gaussian weighting functions—denoted $\omega_x()$ and $\omega_y()$ —constitute the claimed “stickiness parameter,” arguing that these functions determine the extent to which preceding portions of the trajectory influence smoothing. *See* Pet. 27–28 (citing Ex. 1006 at 7:32–8:2). However, this mischaracterizes the role of these weighting functions in Zhou. The weighting functions are not adjustable parameters for controlling the influence of preceding portions on determination of the smoothed trajectory. Rather, they are downstream functions used to implement Zhou’s stabilization strength metric “S.” Ex. 1006 at 7:55-8:3 (“[S]tabilization strength S may vary from buffer entry-to-buffer entry and, as a result, weighting values $\omega_x()$ and $\omega_y()$ may also vary from buffer entry-to-buffer entry.”); *see also* Fig. 4.

98. As Zhou explains, the stabilization strength metric S may define the standard deviation of the weighting functions. *Id.* at 7:61-64 (“stabilization strength S may be the standard deviation of Gaussian weighting function . . . $\omega_x()$ and $\omega_y()$.”). Recall from the prior section that Zhou’s stabilization strength parameter is calculated from the amount of blur estimated in a particular frame. In other words, the spread of the weighting functions is determined by the value S, which in turn is a function of detected motion blur. These functions are merely a

PUBLIC VERSION

mechanism for implementing Zhou’s blur mitigation technique—not some independent input used to control the influence of past motion samples.

99. To understand this, a Gaussian weighting function can be thought of as a bell curve centered on the current video frame. The standard deviation S determines how wide or narrow the curve is. A larger value of S (*i.e.*, less detected blur and more smoothing) produces a wider curve that places greater emphasis on both past and future frames in the buffer, resulting in a trajectory that is more of an average of all frames across the buffer. By contrast, a smaller value of S (*i.e.*, more detected blur and less smoothing) narrows the curve, concentrating influence on the current frame and causing the final trajectory to more closely track the actual motion in the frame.

100. Thus, the weighting functions and their output should not be viewed as some independent parameter for controlling the extent to which past frames can influence the smoothed trajectory. The weighting functions are instead how Zhou implements its stabilization strength parameter to modulate the strength of smoothing in response to detected blur. The weighting functions $\omega_x()$ and $\omega_y()$ are not themselves adjustable parameters, but rather distribution functions whose “spread” depends on the value of Zhou’s stabilization strength parameter rather than any desire to emphasize past frames.

PUBLIC VERSION

101. Even if Zhou’s weighting functions could be viewed as independent parameters for controlling smoothing, they are not “stickiness” parameters as taught and claimed by the ’840 patent. The stickiness parameter is a parameter that “control[s] the extent by which preceding portion(s) of the (observed) trajectory impacts determination of the smoothed trajectory.” Ex. 1001 at 22:22-25. In Zhou, however, the weighting functions are centered on the current frame and apply symmetrically to both past and future motion data. Dr. Katsaggelos confirmed this at deposition, testifying that “the Gaussian curve also known as the Bell curve . . . is symmetric around its mean.” *See* Ex. 2020 at 95:6-13.

102. Thus, Zhou’s weighting functions are incapable of controlling the extent to which preceding portions impact determination of the smoothed trajectory as the contemplated by the stickiness parameter. This is because in Zhou, regardless of the weighting function output, any emphasis placed on past frames (*i.e.*, frame -N) will be exactly ***equal*** to the offsetting emphasis placed on a future frame at an equivalent position in the buffer (*i.e.*, frame +N). Zhou’s weighting functions therefore have no mechanism for expressing a preference for past frames. If anything, the weighting functions control the degree to which Zhou’s system prefers recency (*i.e.*, an equal emphasis on motion in the near future and past), as opposed to the degree the system prefers the past.

PUBLIC VERSION

103. Accordingly, Zhou’s disclosure of Gaussian weight functions does not disclose a stickiness parameter—i.e., a parameter “that control[s] the extent by which preceding portion(s) of the (observed) trajectory impacts determination of the smoothed trajectory.” Ex. 1001 at 22:22–25.

* * *

104. Zhou’s lack of disclosure of the low-light high-pass or stickiness parameters is reflected in the fact that Zhou was cited during prosecution, and the Examiner allowed the ’840 patent over it. See Appl. No. 16/713,798, List of References Cited by Examiner (Nov. 30, 2020).

B. Claims 1 and 16: Zhou In View of Cai Does Not Disclose Nor Render Obvious the Claimed “Temporal Horizon Enabling Approximation of Intentional Motion”

105. Claims 1 and 16 require: “the determination of the smoothed trajectory based on the look-ahead of the trajectory includes use of a temporal horizon of motion experienced by the image capture device, the *temporal horizon enabling approximation of intentional motion* of the image capture device while causing a delay in the determination of the smoothed trajectory.”

106. The Petition relies on Cai’s disclosure of a method for distinguishing stationary (“zero-motion”) camera states from moving (“non-zero motion”) camera states. Pet. at 32 (citing Ex. 1009 at 3:58-59). Cai, however, does not teach the

PUBLIC VERSION

claimed “temporal horizon enabling approximation of intentional motion.” First, Cai’s zero-motion detection process relied upon by Petitioner does not involve the use of a temporal horizon at all, but rather is based only the Cai’s past and current frames. Second, Cai’s zero motion detection merely makes a binary distinction between a stationary camera and one in motion—it has no ability to distinguish between intentional and unintentional motion for a non-stationary camera, and thus is incapable of enabling approximation of intentional motion as claimed. Third, a POSITA would not have been motivated to combine Cai’s zero-motion detection techniques with the teachings of Zhou.


1. *Cai Does Not Disclose Or Suggest The Use Of A Temporal Horizon As Claimed*

107. As the claim indicates, the “temporal horizon of motion experienced by the image capture device” is part of the “look-ahead” upon which the smoothed trajectory is determined. This is supported by the specification which explains how “intentional motion may be characterized by lower frequency motion than unintentional motion,” and that “a sufficiently long temporal horizon of experienced motion may need to be used” to identify it. *Id.* at 9:17-23. The specification explains that “positions/motions of the image capture device in the ‘future’ may be analyzed (look ahead) to determine whether particular position(s)/motion(s) of the image capture device at a moment is an intended motion or an unintended motion (e.g.,

PUBLIC VERSION

shaking due to vibration, rotation due to bumping/mishandling the image capture device).” *Id.* at 15:1-6. Thus, the claimed “temporal horizon” requires a consideration of future frames.

108. Cai’s zero-detection motion logic fails to disclose or suggest any temporal horizon, because it does not consider future frames. *See* Ex. 1009 at 4:26-5:28. As Figure 5 indicates, the only input to Cai’s “zero motion detection” block 522 are previous and current frames of the smoothed trajectory:



Ex. 1009 at Figure 5 (annotated). Raw trajectory future frames are only considered by Cai *after* the zero-motion detection process is complete as part of a separate process that has nothing to do with approximating intentional motion. *See* Section VIII.A.2.

PUBLIC VERSION

109. To understand this more clearly, a full description of Cai’s zero-motion detection logic is instructive. Cai first applies a low-pass filter to obtain what Cai calls its “predicted trajectory.” Ex. 1009 at 4:4-19. Cai also conceptualizes “what the trajectory coordinates would be if there were no motion,” which Cai calls the “planned trajectory.” *Id.* at 4:18-25. To determine whether a given frame in the predicted trajectory is “zero motion” or not, Cai compares the motion associated with the frame to a pre-defined motion “threshold.” The motion associated with the frame can be determined, for example, by comparing “the gap between the last low pass filter coordinate and the current low pass filter coordinate.” *Id.* at 4:30-46. Alternatively, Cai teaches that the motion associated with a given frame can be determined by comparing the gap between the predicted trajectory and a corresponding frame of the planned trajectory. *Id.* at 4:45-56; *see also* Fig. 2.

110. Other approaches involving a “mean” or “variance” across a short sequence of frames of the predicted trajectory are also disclosed. *Id.* at 4:57-5:41. But in all of these approaches the zero motion detection is based on either an analysis of past predicted trajectory frames or a comparison of those past frames to the ***corresponding*** frames of the planned trajectory. Importantly, Cai discloses no buffer or consideration of future frames in its zero-motion detection process, and thus fails to teach the claimed “temporal horizon” of motion.

PUBLIC VERSION

111. When the detected motion in a given frame falls below the predefined threshold, Cai treats the segment as “zero motion” and infers that any motion is jitter resulting from minor hand movements. *Id.* By contrast, detected motion above the threshold is treated as a camera in motion. *Id.* at 4:37-39. Cai refers to this process as “zero motion detection.” *See id.* at 5:40-49 & Fig. 5 at 522; *see also id.* at 9:8-9 & Fig. 7 at 706.

112. If Cai detects that a frame is a “zero motion” frame, the planned (*i.e.*, zero-motion) trajectory is used for video stabilization instead of the low-pass filter output. *Id.* at 4:54-56. Cai refers to this separate step as its “zero motion stabilization” process. *Id.* at 9:9-12 (“This eliminates even small movements that the Gaussian filter cannot remove.”); *see also id.* Fig. 7 at 708. Essentially, the low-pass filter frame is rejected, and the zero motion frame is adopted. Cai then teaches that “the smoothed trajectory from the zero motion stabilization is projected into future frames,” so that it may serve as a further comparison for corresponding future frames of the low-pass filtered frames. *Id.* at 9:13-18.

113. In the Institution Decision, the Board cited Cai’s disclosure at column 9, lines 12-17, providing that “the smoothed trajectory from the zero motion stabilization is projected into future frames, for example, the next five frames.” *I.D.* at 12. However, this disclosure does not teach or suggest that the future frames are


PUBLIC VERSION

used as part of the zero detection logic. To the contrary, this passage teaches that for a given frame “zero motion detection” (706) is performed first, *before* Cai considers future frames as part of the separate “zero motion stabilization” process (708). Ex. 1009 at 9:8-17; Fig. 7. Only if zero motion in a current frame is detected does Cai project the assumed zero motion trajectory into the future, *e.g.* the next five frames, so that those next five zero-motion trajectory frames can each be compared to the “Gaussian filtered next five frames to characterize the video as zero motion or not.” *Id.* at 9:15-17. This does not teach the use of future frame data to detect zero motion for any given frame; rather, this passage simply teaches that once zero motion is detected, the assumed zero-motion trajectory is imposed for that frame and projected forward to provide further comparison points for the next five Gaussian filtered frames.

114. Confirming that Cai’s zero-motion detection process does not consider future frames is the fact that its system suffers from a lag in detecting intentional motion. Figure 4 of Cai highlights this problem. As shown, the predicted trajectory initially appears to follow a straight path (suggesting zero motion), and only after several frames does the cumulative deviation grow large enough to trigger a transition to a sloped trajectory. During this lag period, the planned trajectory fails to follow the actual, intended camera motion—that is, it fails to approximate the

PUBLIC VERSION

intended turn. The adjustment to the planned trajectory comes several frames later, *only after* a false zero-motion classification is made at which point the planned trajectory begins following the intended trajectory.



Ex. 1009 at Figure 4.

115. This behavior stands in sharp contrast to the '840 patent, which uses a “temporal horizon” of motion to look-ahead to future frames so that the system can infer and approximate the intended trajectory without any lag, even where the early parts of the intentional motion are small and subtle. Notice in Figure 5C of the '840 patent how the smoothed trajectory is able to track the intended trajectory in duration D even before the intended motion reaches the threshold that the system previously rejected as unintended motion in durations A, B, and C.

PUBLIC VERSION




FIG. 5C

Ex. 1001 at Figure 5C. This benefit is wholly absent in Cai, reflecting its reliance on current and past frames and its failure to look-ahead across a “temporal horizon” as claimed.

116. Cai’s detection of a “false alarm” as shown in Figure 4 does not reflect a system that enables approximation of intentional motion. Rather, it is a delayed correction of an earlier misclassification. The adjustment is not based on modeling the nature of the motion itself or analyzing its continuity across a temporal horizon. Instead, it is triggered only once the displacement over time exceeds a fixed threshold. As such, Cai’s approach fails to disclose either the use of a *temporal horizon*.

PUBLIC VERSION*1. Cai Does Not Enable Approximation Of Intentional Motion*

117. Even if Cai did teach a “temporal horizon,” its zero-motion detection logic merely distinguishes stationary from moving cameras and thus does not satisfy the “*enabling approximation of intentional motion*” requirement. Specifically, Cai’s approach has no ability to distinguish intentional movements from unintentional movements once the threshold is exceeded and the camera is deemed to be non-stationary. This is inconsistent with the ’840 patent.

118. The ’840 patent teaches that “intentional motion may include a forward motion to follow an object of interest, a panning motion to shoot a wide view of a scenery, and/or other motion,” while “unintentional motion may result in noise . . . caused by shaking hands, vibrations, bumps.” Ex. 1001 at 9:8-16. Moreover, the ’840 patent teaches that intentional motion (e.g., panning or tracking a subject) and unintentional motion (e.g., jitter or bumps) can both occur when the camera is in motion. Figure 5C, for example, illustrates how, in duration D and E, the system continues to use a temporal horizon to infer that intentional motion is underway, and approximate the intended trajectory while still smoothing unintentional perturbations.

PUBLIC VERSION




FIG. 5C

119. Cai lacks the ability to make further distinctions between intentional and unintentional motion once the initial motion threshold is exceeded and the camera is determined to be in motion. For example in Figure 3, Cai teaches its approach for transitioning from a zero motion to non-zero motion mode. Ex. 1009 at 5:43-48. As Cai explains, once the motion between two successive frames becomes larger than the threshold, “[t]his corresponds to some motion of the camera other than jitter,” and “[a]s a result, there is no modification of the low-pass filter result.” In other words, when a non-zero motion frame is adopted, Cai simply adopts the low-pass filtered result without any further ability to identify unintentional motion components. Thus, Cai fails to teach an approach that is capable of *enabling approximation of intentional motion* as claimed.

PUBLIC VERSION

2. *A POSITA Would Not Have Been Motivated To Combine Cai's Zero-Motion Detection Logic With Zhou To Arrive At The Claimed Approach*

120. Even setting aside the deficiencies in Cai, Petitioner fails to adequately establish a motivation to combine Cai's zero-motion detection approach with Zhou. A POSITA recognizes that when Cai detects a zero-motion frame, its "zero motion stabilization" is a "maximalist" approach to stabilization that eliminates all motion below the predefined threshold. This is the strongest form of smoothing which eliminates even the "small" motions that a low-pass filter cannot remove. Ex. 1009 at 9:11-12.

121. However, because Cai's approach imposes a zero-motion constraint on frames that actually have some motion, a POSITA recognizes that it can exacerbate the blur and shimmering artifacts that Zhou seeks to avoid. Ex. 1006 at 3:3-10, 6:48-54 ("[S]maller stabilization strengths may be employed in an effort to mitigate the appearance of the aforementioned unwanted 'shimmering' artifacts in the stabilized video (i.e., by letting the image sequence jitter slightly, rather than holding it perfectly stable, the motion blur will follow the natural trail of movement in the image sequence and not appear as anomalous to the viewer)."). Indeed, Zhou is primarily focused on addressing the issue of blur that occurs when a frame having motion is over-stabilized. For example, Zhou's stabilization

PUBLIC VERSION

strength parameter is always bounded by a maximum value, Ex. 1006 at 6:54-62 & Fig. 5C, and even in situations involving strong smoothing, Zhou's Gaussian weighting functions will interpolate motion across a wide range of buffered frames, *id.* at 7:24-31. A POSITA recognizes that Zhou specifically avoids Cai's maximalist approach to stabilization, as such an approach would exacerbate the very blur condition that Zhou seeks to address. Thus, Cai's zero-motion detection techniques are fundamentally at odds with Zhou's blur management teachings, and a POSITA seeking to implement Zhou would not have been motivated to look to the teachings of Cai to improve Zhou.

122. Furthermore a POSITA implementing Zhou would not be motivated to modify Zhou itself to arrive at the claimed "temporal horizon enabling approximation of intentional notion" feature. Zhou expresses no motivation for distinguishing between intentional and unintentional motion, focusing instead on the general blur problem that it seeks to address. Zhou, in fact, discourages the type of "sufficiently long" temporal horizon that would be required to enable approximation of intentional motion, teaching instead how buffering frames "can be expensive in terms of memory (e.g., one frame at 1080p video resolution can be between 2 and 3 megabytes, MB)." Ex. 1006 at 4:62-5:2.

PUBLIC VERSION

123. Thus, Petitioner fails to establish that its purported modification of Zhou and/or combination with Cai would have been obvious. The Petition fails for this independent reason.

IV. GROUND 2-3 FAIL ON THE MERITS

124. Petitioner asserts that the combination of Zhou and Cai in combination with various other references render the remaining dependent claims of the '840 patent obvious. However, except for claims 13 and 17 (addressed below), these combinations fail for the same reasons as Ground 1.

V. GROUND 4: PETITIONER FAILS TO ESTABLISH THAT ZHOU, CAI, AND THE BELL PAPER RENDER CLAIMS 13 AND 17 OBVIOUS

125. Claims 13 and 17 further require that the “*weight-balance parameter controls types of motion that are minimized in smoothed trajectory generation, the types of motion including rotational velocity and rotational acceleration.*” Petitioner relies on the Bell Paper for this limitation—but the Bell Paper does not disclose a weight-balance parameter that controls the minimization of *both* rotational velocity *and rotational acceleration*.

126. The Bell Paper describes “a method for video stabilization . . . for videos captured on mobile devices.” Ex. 1004 at 294 (Abstract). It uses “data from an on-board gyroscope to track the camera’s *angular velocity*.” *Id.* (Emphasis added). The Bell Paper explains: “We remove small motions and rolling-shutter

PUBLIC VERSION

distortions due to hand shake,” and for “larger motions, we filter the camera’s *angular velocity* to produce a smooth output.” *Id.* (emphasis added). But filtering angular velocity to smooth the camera path is not the same as *minimizing* rotational velocity. And the Bell Paper says nothing about measuring, calculating, or minimizing the camera’s *rotational acceleration*—i.e., how quickly rotation speeds up or slows down—let alone using a weight-balance parameter to control minimization of both.

127. The Bell Paper applies a “nonlinear filtering method which completely flattens small motions regardless of frequency, and performs low-pass smoothing when the virtual camera must move to keep the crop window inside the input frame.” *Id.* at 296. It explains that the “gyroscope provides a series of discrete *angular velocity* measurements with time-stamps, which we integrate to produce a function of time that describes the camera orientation.” *Id.* at 297. (Emphasis added). Again, this is filtering—not minimization—and while it uses angular velocity, the Bell Paper never even measures *rotational acceleration*.

128. Indeed, the Bell Paper discourages the use of acceleration measurements, even for translational motion, as being “difficult and of limited value”—a rationale that logically applies with even more force to rotational acceleration, given its greater complexity:

PUBLIC VERSION

In theory we could be more precise by also measuring translation with the device's accelerometer, but in practice this is difficult and of limited value.

Id.

129. In describing its “motion model and smoothing algorithm,” the Bell Paper explains: “We parametrize the camera path with the camera’s orientation and ***angular velocity*** of each frame . . . represent[ing] the physical and virtual camera orientations” as quaternions . . .” *Id.* at 299. (Emphasis added). The smoothing algorithm “creates a new camera path that keeps the virtual camera static when the measured motion is small enough to suggest that the actual intention is to keep the camera static, and that otherwise follows the intention of the measured motion with smooth changes in ***angular velocity***.” *Id.* at 299. Again, the Bell Paper smooths angular velocity—but does not minimize it, and does not even reference ***rotational acceleration***, much less a parameter for minimizing both.

130. In its “Results” section, the Bell Paper mentions a single video (“the walking video”) where the method removed “most of the angular acceleration,” resulting in a “smoothed camera path [that] is nearly free of rotational motions.” *Id.* at 304. But this isolated reference is a general, qualitative description of a perceived result—it does not teach or disclose any rigorous method or even desire to optimize the rotational acceleration of the camera housing in any rigorous or

PUBLIC VERSION

quantitative way. Indeed, the Bell Paper provides no method for measuring or minimizing rotational acceleration. It provides no indication that the authors recognized or appreciated the tradeoffs between minimizing rotational velocity and rotational acceleration or the benefits of striking a balance between these two goals. And, the Bell Paper certainly discloses no weight-balance parameter for actually controlling the extent to which an angular velocity and acceleration are minimized, as required by the '840 patent.

III. SECONDARY CONSIDERATIONS CONFIRM THE NON-OBVIOUSNESS OF THE '840 PATENT

131. In my opinion, there is strong, objective evidence that the '840 patent is not obvious. Patent Owner GoPro, Inc. ("GoPro") successfully manufactured and sold action cameras that include the HyperSmooth functionality, which embodies the '840 patent.

A. HyperSmooth Practices the '840 Patent

132. In my opinion, GoPro action cameras with HyperSmooth practice at least claims 1, 13 and 14 of the '840 patent. My analysis below focuses on the GoPro HERO7 Black, [REDACTED]

[REDACTED]

[REDACTED] [REDACTED]

[REDACTED]

PUBLIC VERSION

[REDACTED]

[REDACTED]

[REDACTED]

1[pre]: An image capture system that stabilizes videos, the image capture system comprising:

133. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

1[a]: a housing of an image capture device:

135. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

PUBLIC VERSION

[REDACTED]

[REDACTED]

1[b]: an optical element carried by the housing and configured to guide light within a field of view to an image sensor:

136. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

1[c]: the image sensor carried by the housing and configured to generate a visual output signal conveying visual information based on light that becomes incident thereon during a capture duration, the visual information defining visual content having the field of view, the field of view being larger than a size of a viewing window for generating stabilized visual content:

138. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

PUBLIC VERSION

[REDACTED]

[REDACTED]

1[d]: a position sensor carried by the housing and configured to generate a position output signal based on positions of the housing during the capture duration, the position output signal conveying position information that characterizes rotational positions of the housing at different moments within the capture duration:

140. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

PUBLIC VERSION

1[e]: one or more physical processors configured by machine-readable instructions to:

142. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

1[f]: determine a trajectory of the housing during the capture duration based on the position information, the trajectory reflecting the rotational positions of the housing at the different moments within the capture duration, the trajectory including a first portion corresponding to a first moment within the capture duration and a second portion corresponding to a second moment subsequent to the first moment within the capture duration:

144. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

PUBLIC VERSION

[REDACTED]

[REDACTED]

[REDACTED]

1[g(i)]: determine a smoothed trajectory of the housing based on a look-ahead of the trajectory and one or more of a weight-balance parameter, a low-light high-pass parameter, and/or a stickiness parameter:

146. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

147. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

PUBLIC VERSION

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

1[g(ii)]: the look-ahead of the trajectory including use of a subsequent portion of the trajectory to determine a preceding portion of the smoothed trajectory such that a portion of the smoothed trajectory corresponding to the first portion of the trajectory is determined based on the second portion of the trajectory, the smoothed trajectory having smoother changes in the rotational positions of the housing than the trajectory:

PUBLIC VERSION

151. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

l[g(iii)]: wherein the determination of the smoothed trajectory based on the look-ahead of the trajectory includes use of a temporal horizon of motion experienced by the image capture device, the temporal horizon enabling approximation of intentional motion of the image capture device while causing a delay in the determination of the smoothed trajectory:

153. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

PUBLIC VERSION

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

156. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

motion.

PUBLIC VERSION

l[h]: determine placement of the viewing window for the visual content with respect to the field of view of the visual content based on the smoothed trajectory of the housing, the viewing window defining one or more extents of the visual content to be included within a punchout of the visual content:

157. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

158. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

l[i]: generate the stabilized visual content of a video based on the viewing window, the stabilized visual content including the

PUBLIC VERSION

punchout of the one or more extents of the visual content within the viewing window, wherein inclusion in the stabilized visual content of the one or more extents of the visual content within the viewing window effectuates stabilization of the visual content via selective cropping:

159. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Claim 13: The image capture system of claim 1, wherein the weight-balance parameter controls types of motion that are minimized in smoothed trajectory generation, the types of motion including rotational velocity and rotational acceleration:

161. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

162. [REDACTED]

[REDACTED]

PUBLIC VERSION

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Claim 14: The image capture system of claim 1, wherein the low-light high-pass parameter controls amount or strength of smoothing performed in smoothed trajectory generation and an extent of motion blur compensation applied in determination of the smoothed trajectory:

163. In my opinion, the HERO7 Black practices this limitation. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

PUBLIC VERSION**B. Nexus**

165. Here, secondary considerations—commercial success, industry praise, long-felt need, failure of others, unexpected results, and copying—are directly attributable to the unique stabilization approach claimed in the '840 patent. GoPro's HERO7 Black's HyperSmooth system leverages a weight-balance parameter to adjust the emphasis between rotational velocity and acceleration, a low-light high-pass parameter to preserve high-frequency motion in low-light conditions, and a temporal horizon to approximate intentional motion. This multifaceted optimization represented a substantial departure from prior approaches like Shi's blur-detection-based blending or Kwatra's linear image-driven smoothing, enabling GoPro to achieve stabilization praised by consumers and industry experts alike.

166. HERO7 Black embodies the inventions claimed in the '840 patent, including claims 1, 13, and 14. While Petitioner disputes that HERO7 Black meets limitation 1[g(iii)]—requiring a temporal horizon approximating intentional

PUBLIC VERSION

motion and introducing a processing delay—that argument mischaracterizes the HERO7’s stabilization.

167. I have confirmed that HERO7 Black includes a temporal horizon as understood by a POSITA. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

168. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

169. The HERO7 Black was the first GoPro product to commercialize HyperSmooth and embody the claimed inventions of the ’840 patent. *See* Ex. 2024 (Lema Decl.), ¶¶ 6, 10. Its launch marked a paradigm shift for GoPro and the industry, driven by HyperSmooth’s capabilities. *See id.* at ¶ 10. The resulting commercial success and industry recognition—including HERO7 Black being

PUBLIC VERSION

called one of GoPro’s most important releases—were directly tied to the stabilization enabled by the ’840 patent. *See id.* at ¶ 11.

170. Accordingly, the HERO7 Black’s commercial success, industry praise, and copying by competitors are specifically tied to the patented features. The record thus establishes a clear nexus between the ’840 patent and the secondary considerations supporting non-obviousness.

C. Commercial Success

171. I understand that GoPro’s HyperSmooth technology, which practices the ’840 patent, has driven significant commercial success. According to GoPro’s head of product, Mr. Pablo Lema:

The release of HyperSmooth was a “shock to the market.” HERO7 Black became the fastest-selling production launch in GoPro’s history at the time. Our sales jumped 170% compared to HERO6 in the same post-launch window, driven almost entirely by the introduction of HyperSmooth. From a feature standpoint, the HERO6 camera is nearly indistinguishable from the HERO7 Black with the exception of HyperSmooth.

Ex. 2024 at ¶ 11.

172. Mr. Lema also cited internal survey data confirming that HyperSmooth and immersive lens options were among “the most important and valued features driving customer demand across our HERO customer base,” and

PUBLIC VERSION

that GoPro’s “product planning decisions—including resource allocation and feature prioritization—relied heavily on this data.” *Id.* at ¶19.

173. Specifically, in a [REDACTED] (Ex. 2040, also Ex. A to Lema Dec., CX-590C), “Advanced video stabilization” ranked as the top feature for the Core HEROes segment and among the top three for the market segments representing GoPro’s historical customer base. *Id.* at ¶¶ 14-19.

174. In a separate [REDACTED] (Ex. 2041, also Ex. B to Lema Dec., CX-649C), HERO11 Black buyers selected “Advanced video stabilization (HyperSmooth 5.0 with AutoBoost)” as the #1 driver of purchase—above all other features. *Id.* at ¶ 20.

175. These results, according to Mr. Lema, “reflect the role that GoPro’s patented stabilization and lens features play in driving consumer demand and purchasing behavior.” *Id.* at ¶ 21. He concluded that “the patented technologies in the ’894, ’840, ’832, and ’413 patents were directly responsible for GoPro’s sustained commercial success with HERO cameras, especially from the HERO7 onward.” *Id.*

176. GoPro’s dominant position in the U.S. action camera market—holding more share than all competitors combined—further illustrates the commercial strength of its patented technologies. Partnerships with brands like Red Bull, NHL,

PUBLIC VERSION

and Adobe underscore this market confidence. *See* Ex. 2005 (NHL Senior VP of Programming and Production said, “[O]nly GoPro has the technology to help us and our media partners showcase the beauty and intensity of hockey in new and deeper ways.”); Ex. 2008 (GoPro “named the exclusive action camera of Red Bull” mountain bike competition; GoPro athletes will capture their runs wearing HERO11 featuring “GoPro’s Emmy Award-winning HyperSmooth video stabilization”); Ex. 2009 (Vice President of Adobe Stock said, “Adding GoPro’s unique footage to Adobe Stock gives our community access to a library of immersive ‘wow’ moments”).

D. Industry Praise and Recognition

177. I understand that GoPro’s HyperSmooth technology has received widespread industry recognition, including a 2019 CES Innovation Award and a 2022 Technology and Engineering Emmy. *See* Ex. 2010 (““In Camera Sensor and Software Stabilization – Awarded to GoPro and Socionext” . . . “The technology has created the ability for camera operators and producers to obtain stabilized action footage when a subject is rapidly moving easily and economically in unstable situation.”); Ex. 2011; Ex. 2012 (“This Emmy is in recognition of the technology powering GoPro’s industry-leading HyperSmooth video stabilization”).

PUBLIC VERSION

178. Mr. Lema confirms, “GoPro won the 73rd Technology and Engineering Emmy award specifically for its in-camera sensor and software image stabilization, introduced with HERO7.” Ex. 2024 at ¶¶ 13, 22. He also notes “a vast amount of industry praise for HyperSmooth and HERO7” in the form of customer reviews, publications, and tech blogs. *Id.* at ¶ 23. These technologies enabled “a step-change” in user experience that drove demand, growth and recognition for GoPro’s products. *Id.* at ¶ 24.

179. Third-party reviews echo this view. Analysts at JP Morgan called GoPro’s “gimbal like hypersmooth stabilization (developed in-house . . . a significant upgrade” and “key selling point” for the HERO7. Ex. 2013.; Tech reviewers wrote: “GoPro’s HyperSmooth image stabilization is the best we’ve seen on any action camera.” Ex. 2014 at 5.

180. Technical reviews of the subsequently released HERO8 Black note how HyperSmooth had already “eliminated the need for using a motorized gimbal,” and how the electronic image stabilization of HyperSmooth 2.0 was “unfathomably good.” Ex. 2042 (CX-1388) at 3.

E. Long-Felt but Unmet Need

181. Before HyperSmooth, users of action cameras struggled to capture stable footage during high-vibration activities like mountain biking, hiking, or

PUBLIC VERSION

surfing. Mechanical gimbals were often used but were bulky, fragile, and poorly suited to these dynamic conditions. The '840 patent's digital stabilization filled this long-felt need.

182. According to GoPro's Mr. Lema, HyperSmooth was a "shock to the market." Ex. 2024 at ¶ 10. Public reaction was "one of disbelief"—some users assumed GoPro had faked footage using mechanical stabilizers. But once they tried the product, "the response was overwhelmingly positive." *Id.* at ¶ 12.

183. These outcomes reflect how HyperSmooth filled a long-felt need—further demonstrating the non-obviousness of the claimed inventions.

F. Failure of Others

184. GoPro was not the only camera maker seeking to perfect digital image stabilization in 2018. Some of the world's greatest technology leaders, including Sony, Google, Samsung, and Apple, were all intently pursuing their own solutions. Many of these approaches involved hybrid electronic image stabilization ("EIS") and optical image stabilization ("OIS") solutions—as disclosed in Shi, a Google patent, Ex. 1006 at 1:36-44 and Zhou, an Apple patent, Ex. 1004 at 1:34-45—which were widely believed to be necessary given the computational requirements of pure EIS solutions and the limitations of mobile devices. [REDACTED]

PUBLIC VERSION

[REDACTED]

[REDACTED] :

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

G. Unexpected Results

185. GoPro’s system consistently outperformed traditional stabilization techniques, which often failed under rapidly changing motion. HyperSmooth maintained stabilization where others did not. Ex. 2014 (“GoPro’s HyperSmooth image stabilization is the best we’ve seen on any action camera”); Ex. 2024 at ¶ 11 (HyperSmooth was a “shock to the market”).

PUBLIC VERSION

186. These unexpected results further reinforce the non-obviousness of the '840 patent.

H. Copying by Competitors

187. I understand that in this case, Petitioner Insta360, a direct competitor, has been accused of copying GoPro's HyperSmooth technology, underscoring the value and innovation of the '840 patent. Ex. 2006 (Complaint).

188. I understand that beginning in 2021, Insta360 began releasing its own version of "in-camera" electronic image stabilization technology that relied upon look-ahead smoothing, optimization of the smooth camera trajectory path, and specialized tunable parameters to control the strength of smoothing. Ex. 2048 (Insta360 Speeds Up Workflow with New 360 Reframing Tool, Better Stabilization & Quick Reader). [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

- [REDACTED]

- [REDACTED]

- [REDACTED]

PUBLIC VERSION

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]

189. In my opinion, these secondary considerations—commercial success, industry praise, long-felt need, unexpected results, and copying— confirm that the '840 patent's inventions were not obvious. HyperSmooth revolutionized the action camera market, drove demand, earned recognition, and prompted imitation. In my opinion, the claims should be upheld.

IV. CONCLUSION

190. In signing this declaration, I recognize that the declaration will be filed as evidence in a contested case 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

PUBLIC VERSION

United States. If cross-examination is required of me, I will appear for cross-examination within the United States during the time allotted for cross-examination.

191. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on the 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: April 19, 2025

/s/ Hassan Foroosh

Hassan Foroosh, Ph.D.