

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

TESLA, INC.

Petitioner,

v.

INTELLECTUAL VENTURES II LLC

Patent Owner.

Case No. IPR2025-00342

U.S. Patent No. 7,336,805

PATENT OWNER'S SUR-REPLY

TABLE OF CONTENTS

PATENT OWNER’S EXHIBIT LIST	iii
I. Introduction.....	1
II. Tesla’s Reply Still Fails to Show Claim 1’s Staged Acceptance-Analysis Workflow in Broggi.....	1
A. Figure 1 and the Specification Disclose That The Additional Acceptance Step Is Applied To The Objects Identified In The Preceding Analyzing Steps.....	1
B. The Intrinsic Evidence Demonstrates That The Additional Acceptance Step Is Applied To The Objects Identified In The Preceding Analyzing Steps Recited In The Claim.....	4
C. Tesla’s Reading Of The Extracting Step Renders The Steps Of Edge Detection, Analysis Of The Detected Edges To Identify Potential Geometric Objects, And Analysis Using The Matching Algorithm Completely Irrelevant To The Output Of The Extracting Step.....	6
D. Tesla’s Reading Of “Additional” Renders The Term Superfluous.....	8
E. The Reply Ignores Claim 1’s Ordered Filtering Pipeline.....	10
F. Relabeling Broggi’s Bounding-Box Steps Does Not Cure the Petition’s Collapse of Distinct Claim Stages.....	12
III. The Reply Still Fails to Show the Claimed Tree Structure Being Traversed and Used in the Required Downstream Workflow	17
IV. Broggi’s Moving Lead Vehicle Is Not the Claimed “Potential Destination,” “Most Proximate Destination,” or Endpoint for Optimized Path Computation	21
V. Conclusion	24

PATENT OWNER'S EXHIBIT LIST

EX2001	Petitioner's Answer to Complaint and Counterclaim, dated June 20, 2024, in <i>Intellectual Ventures II LLC v. Tesla, Inc.</i> , C.A. 1:24-cv-00884-ADA (W.D. Tex.)
EX2002	Docket from District Court Litigation, <i>Intellectual Ventures II LLC v. Tesla, Inc.</i> , C.A. 1:24-cv-00884-ADA (W.D. Tex.)
EX2003	Scheduling Order from District Court Litigation, <i>Intellectual Ventures II LLC v. Tesla, Inc.</i> , C.A. 1:24-cv-00884-ADA (W.D. Tex.)
EX2004	Defendant's Second Supplemental Preliminary Invalidity Contentions, dated January 9, 2025, in District Court Litigation, <i>Intellectual Ventures II LLC v. Tesla, Inc.</i> , C.A. 1:24-cv-00884-ADA (W.D. Tex.)
EX2005	Letter from Patent Owner to Associate General Counsel for Petitioner, dated August 20, 2020, enclosing copy of U.S. Pat. No. 7,181,143 and Claim Chart Presentation
EX2006	R. Sedgewick, Algorithms in C, Addison-Wesley Pub. Comp. Inc., 1990
EX2007	Declaration of Dr. Majumder In Support Of Patent Owner's Preliminary Response.
EX2008	CV of Dr. Majumder
EX2009	Ninomiya et al., U.S. Patent No. 5,220,508
EX2010	Declaration of Dr. Majumder In Support Of Patent Owner's Response.

I. Introduction

Petitioner's Reply does not cure the defects in its original obviousness case. The challenged claims recite an ordered technical workflow, but the Reply still does not show that workflow in the prior art. Instead, Petitioner continues to fragment the claims into isolated phrases, subdivide single routines in Broggi into multiple supposed claim stages, and treat generalized disclosures in Huttenlocher as though they satisfied the claims' required downstream structure and use.

The same mismatch remains throughout Petitioner's theories. Broggi's lead-vehicle tracking is not the claimed progression from candidate object, to plausibility screening, to later acceptance analysis, nor is it the claimed use of an accepted destination object for optimized path computation to an endpoint. Huttenlocher likewise is not shown to provide the stored relational structure later traversed and used in the claimed downstream workflow.

II. Tesla's Reply Still Fails to Show Claim 1's Staged Acceptance-Analysis Workflow in Broggi

A. Figure 1 and the Specification Disclose That The Additional Acceptance Step Is Applied To The Objects Identified In The Preceding Analyzing Steps

Tesla's Reply is wrong to characterize Patent Owner's position as an effort to "limit the claims to a specific embodiment" or to "read" a pipeline into claim 1. Reply, 5-7. Figure 1 and the specification do not merely disclose one optional implementation among many. They illustrate the very staged workflow that the

claim language recites and that Dr. Majumder confirmed at her deposition. When asked directly what Figure 1 depicts, Dr. Majumder testified: “They are depicting the pipeline.” EX1073, 99:6-7. Patent Owner is not manufacturing a sequential architecture out of whole cloth. The patent itself presents the extraction process as a sequence of downstream units, and Figure 1 visually reflects that sequence.

In her deposition, Dr. Majumder walked step-by-step through the method sequence disclosed in Figure 1. The specification explains that “an edge detector and segmenter” extracts edges and segments that are stored in a hierarchically organized tree, and then a downstream unit 12 analyzes “if I’m looking for a rectangular destination, I would look for objects with rectangular shape or at least square shape or partially rectangle shape, but anything that doesn’t have corners, like a circle, would be rejected.” EX1073, 101:7-12. Dr. Majumder explained that the hierarchical tree is processed by a “[c]onnected downstream” unit and that “downstream means it is taking its input” from the prior unit. EX1073, 100:4-14, 100:19-21; 101:1-12; 100:19-21. That testimony directly confirms that the patent discloses a first stage that stores extracted edge data in a hierarchical tree, and a downstream stage that analyzes the tree to determine if the extracted edges can be pieced together to form geometric objects that describe typical destinations.

Dr. Majumder then addressed the comparator/plausibility stage in which a comparator unit “analyz[es] the detected geometric objects” to “check for

plausibility using a matching algorithm,” and explained that this stage compares the geometric objects output from the prior stage against stored object patterns to determine which candidates remain plausible. EX1073, 101:15-21; 102:4-21. In her words, reference numeral 13 receives the output from 12 and then determines “which of these objects stored are actually matching” the possible destination. EX1073, 102:4-12. The specification explains that the purpose of this analysis is so that “the subsequent outlay for processing may be limited already at this stage of the inventive method to relatively few plausible objects which are actually similar to a potential destination.” EX1001, 5:10-13. Thus, the patent discloses a distinct plausibility screen, that uses a matching (or comparison) process to narrow the previously “pieced together” candidates in order to limit the resources spent on subsequent processing.

Most importantly, Dr. Majumder then addressed the acceptance analysis unit (14) that operates on the narrowed candidates identified by the matching (or comparison) and “based on the comparison of selected geometric objects,” analyzes the shape of the image formation based on knowledge of imaging properties of the sensor relative to the surrounding field. EX1073, 103:6-13. As explained in the specification:

... Within the framework of the inventive method, the detected geometric objects that correspond to the typical geometric form are then analyzed for plausibility using a matching algorithm, these objects, which are classified as plausible, undergo[] an additional acceptance

analysis to the effect that, based on the knowledge of the imaging properties of the imaging sensor in relation to its surrounding field ... EX1001, 3:47-49 (emphasis added).

Dr. Majumder explained that, from a POSITA's perspective, a plausible candidate object has been identified by the matching, and the system asks whether the image of that object appears as it should appear given the camera's position and orientation. EX1073, 103:16-22; 104:1-8. Thus, the disclosed system analyzes a geometric form that was previously identified as plausible in the context of the surrounding field, in order to determine if a shape of that formation in the context of the surrounding field is consistent with how such a destination should appear. EX1073, 105:1-17. In this regard, when directed to the passage stating that the objects "scaled and oriented in [the] comparator unit by the matching algorithm" "subsequently undergo a further acceptance analysis," Dr. Majumder explained that this later stage asks whether the image "from above" conforms to how the destination would appear "given that I know the position and orientation of the image sensor." EX1073, 106:1-10; 107:1-11. That downstream imaging-property-based gate necessarily follows the comparator-unit matching; it does not collapse into it.

B. The Intrinsic Evidence Demonstrates That The Additional Acceptance Step Is Applied To The Objects Identified In The Preceding Analyzing Steps Recited In The Claim

Claim 1 recites the same multi-stage analysis process disclosed in the specification and described above, i.e., (1) an edge detection step where edge

segments are stored in a hierarchical tree, (2) a first analyzing step to determine if the segments are able to be pieced together to form geometric objects, (3) a second analyzing step that performs a plausibility screen using a matching process to narrow the previously “pieced together” candidates, and (4) an additional acceptance step that analyzes a geometric form that was previously identified as plausible in the context of the surrounding field. In this regard, claim 1 recites in pertinent part:

- performing an edge detection and edge segmentation on the image data so as to break down the image data into a plurality of edge segments and storing interrelationships of the plurality of edge segments in a mathematical tree structure;
- analyzing the plurality of edge segments for the presence of a geometric object associated with a geometrical form that may at least partially describe a potential destination so as to identify one or more potential geometric objects;
- analyzing the one or more potential geometric objects for plausibility using a matching algorithm so as to identify one or more plausible objects;
- performing an additional acceptance analysis so as to identify at least one accepted object, the additional acceptance analysis including analyzing a shape of the image formation of each object in the image data based on knowledge of at least one imaging property of the imaging sensor relative to the surrounding field; ...

The claim language when in light of the specification demonstrates that the “additional acceptance” step is sequentially tethered to the previous analysis steps. More specifically, “a shape of the image formation of each object in the image data” in the “additional acceptance step” is linguistically tethered to “a geometric form that may at least partially describe a potential destination” in the first analyzing step. Both “form” and “formation” have the same root, and as explained in the

specification, the “shape of the formation” (in the “additional acceptance” step) corresponds to the “geometric form” (from the first analyzing step) when viewed in the surrounding field:

.. Within the framework of the inventive method, the detected geometric objects that correspond to the typical geometric form are **then analyzed** for plausibility using a matching algorithm, these objects, which are classified as plausible, undergo[] an additional acceptance analysis to the effect that, based on the knowledge of the imaging properties of the imaging sensor in relation to its surrounding field, the shape of the image formation of the objects in the image data is analyzed (thus, an image of a typical rectangular geometric form viewed from an elevated location is formed in the image data as a trapezoid that tapers toward the top). EX1001, 3:47-57 (emphasis added).

Based on the intrinsic evidence, it is clear that “a shape of the image formation” (in the additional acceptance analysis) corresponds to “a geometric form” (from the first analyzing step) viewed relative to the surrounding field, and it would be an error to interpret the “shape of the formation” (in the “additional acceptance” step) as lacking dependency upon the “geometric form” (from the first analyzing step).

C. Tesla’s Reading Of The Extracting Step Renders The Steps Of Edge Detection, Analysis Of The Detected Edges To Identify Potential Geometric Objects, And Analysis Using The Matching Algorithm Completely Irrelevant To The Output Of The Extracting Step

The claimed “extracting” step [1.2] must result in “positional parameters of at least one potential destination.” That extracting step [1.2] requires five substeps,

namely, step [1.4.0] which recites “performing an edge detection and edge segmentation ...,” step [1.4.1] which recites “analysis of the edge segments ... to identify one or more potential geometric objects ...,” step [1.4.2] which recites “analyzing the one or more potential geometric object for plausibility ...,” step [1.4.3] which recites “an additional acceptance analysis so as to identify at least one accepted object ...,” and step [1.4.4] which recites “storing at least one accepted object ... together with corresponding relative positional data” The positional parameters recited in [1.2] are the output of the “extracting” map to the “relative positional data” stored in step [1.4.4], which corresponds to “at least one accepted object” identified in step 1.4.3. Under Tesla’s interpretation, the output recited in step [1.4.4] is only dependent on the identification of “at least one accepted object” in step [1.4.3], and can be completely unrelated to the results of steps [1.4.0], [1.4.1] and [1.4.2]. This reading renders steps [1.4.0], [1.4.1] and [1.4.2] completely superfluous by severing any requirement that those steps affect the output of the extracting step and disconnecting those steps from any other aspect of the claim language.

In *Mformation Techs., Inc. v. Research in Motion Ltd.*, 764 F.3d 1392, 1399 (Fed. Cir. 2014), the Federal Circuit expressly rejected a claim interpretation when “the separate sub-step ... would become ‘superfluous’.” Similarly, the Court explained in *Sound View Innovations, LLC v. Hulu, LLC* that an “ordering exists

when there are inherent logical dependencies or functional relationships between the recited steps of a method claim.” 166 F.4th 958 (Fed. Cir. 2026). Here, a POSITA would not interpret claim 1 in a way that rendered steps [1.4.0], [1.4.1] and [1.4.2] superfluous to the output of the extracting step and disconnected from every other aspect of the claim language.

D. Tesla’s Reading Of “Additional” Renders The Term Superfluous

Tesla contends that “additional” does not impose any sequencing because the claim does not separately recite another “acceptance analysis,” and because words like “then” and “subsequently” appear in the specification rather than the claims. Reply, 7-8. However, the fact that the earlier “analyzing” steps do not recite the exact term “acceptance analysis,” does not diminish the fact that they are performing an “acceptance analysis” within the parlance of the patent. The specification uses “an additional acceptance analysis” in the following context:

.. Within the framework of the inventive method, the detected geometric objects that correspond to the typical geometric form are then analyzed for plausibility using a matching algorithm, these objects, which are classified as plausible, undergoing an additional acceptance analysis to the effect that, based on the knowledge of the imaging properties of the imaging sensor in relation to its surrounding field, the shape of the image formation of the objects in the image data is analyzed. EX1001, 3:47-55 (emphasis added).

For this sentence to make sense, the step of analyzing “the typical geometric form ... for plausibility using a matching algorithm” must necessarily be an “acceptance analysis”; otherwise, it would have been nonsensical for the patent to refer in the

next phrase to undergoing “an additional acceptance analysis.” Petitioner’s reading strips “additional” of any meaning.

Tesla’s reliance on the open-ended nature of “comprising” and “including” is misplaced. Reply, 4-5. Given the open-ended nature that those phrases already imparted to the claim, it would have been redundant and superfluous for the patentee to have used “additional” to modify “acceptance analysis” simply to convey that the “acceptance analysis” was an “additional” step included in the claim. Contrary to Tesla’s position, the patentee did not place the modifier “additional” in front of “acceptance analysis” for no reason. Rather, the patentee placed “additional” in front of “acceptance analysis” to signal that an “additional” (or further) “acceptance analysis” was performed based on earlier analyzing steps, which themselves included an “acceptance analysis.”

While “comprising” and “including” may permit additional unrecited steps, that does not diminish the requirement that the claims be read in light of the specification and that a POSITA doing so would have understood the claims to require the storage of edge segments in a tree, followed by an analysis of those edge segments for potential geometric objects, followed by an “acceptance” analysis assessing those potential geometric objects for plausibility, followed by the “additional” acceptance analysis. EX1073, 83:20-22; 84:1-21; 85:1-16.

E. The Reply Ignores Claim 1's Ordered Filtering Pipeline

The Reply does not identify where Broggi discloses a comparator-stage plausibility screen followed by a distinct, imaging-property-based, additional acceptance analysis. Instead, it continues to rely on Broggi's bounding-box routine, particularly the search for the "top part of the bounding box" in a region determined by perspective and size constraints, along with stereo-based distance refinement, and asserts that those features satisfy limitation [1.4.3]. Reply, 18-21. But those cited features remain part of the same vehicle-detection process in Broggi. Tesla itself describes Broggi as first finding the bottom corners and then, as a "second, separate, and subsequent analysis step," detecting the top of the bounding box after locating the bottom corners, with that second step using perspective and size constraints. Reply, 19-21. Even taking Tesla's description at face value, that is still only a continued bounding-box confirmation sequence. It is not a later acceptance-analysis unit that, after comparator-based plausibility screening, tests whether a candidate object's image formation matches expected projective distortion and camera-scene geometry. The Reply points to region-constrained top-of-box search and stereo refinement, not to anything analogous to the patent's unit 14.

That mismatch is precisely what Dr. Majumder confirmed in the deposition. She testified that, in Broggi, "all of them have been compressed into one," that "there is no pipeline where I am sieving through different things and rejecting," and that

Broggi instead identifies the box to look for and then applies target-search and refinement operations within that same tracking routine. EX1073, 95:9-21; 96:1-15. She further confirmed that it remained her position that the Petition improperly relies on Broggi's "top-of-box detection for both 1.4.1 and 1.4.3." EX1073, 97:21-98:3. Tesla's Reply attempts to relabel the mappings to avoid that criticism, but the underlying problem persists: Petitioner still relies on the same Broggi vehicle-box process, with the same search-window logic and associated refinement, to do the work of both early candidate identification and later additional acceptance analysis. Reply, 18-21.

The question is whether the Petition has identified in the prior art the particular kind of "additional acceptance analysis" the claim recites and Figure 1 depicts. Dr. Majumder's testimony shows that the answer is no. In the patent, unit 12 locates candidate geometry, unit 13 applies matching/plausibility, and unit 14 then asks a different question: whether the image of the candidate object is consistent with the expected appearance given the imaging sensor's relationship to the scene. EX1073, 100:4-14; 101:15-21; 103:6-13; 107:3-11. Broggi's top-of-box search does not ask or answer that later question. It remains part of a box-detection heuristic aimed at confirming and tracking a lead vehicle. Reply, 20-21.

Accordingly, the Board should reject Tesla's attempt to flatten the patent's specific, stepwise process into one generic detection routine. Figure 1, the

specification passages tied to reference numerals 12, 13, and 14, and Dr. Majumder's deposition testimony all point the same way: the '805 patent claims a downstream acceptance gate that necessarily follows earlier geometric-object and plausibility stages and that evaluates candidate image formation against known sensor-relative imaging properties. EX1073, 99:6-7; 100:19-21; 102:4-12; 103:16-22; 107:3-11. Tesla's Reply still identifies no such stage in Broggi. It identifies only operations within the same bounding-box/vehicle-detection sequence. Reply, 18-21.

Claim 1 recites an ordered filtering pipeline. Tesla's Reply still does not identify in Broggi a later "additional acceptance analysis" applied to the narrowed candidate set produced by the earlier stages. Instead, it treats the "additional acceptance" step ([1.4.3]) as a free-floating reanalysis of any object in the image, which is inconsistent with both the intrinsic evidence and with Dr. Majumder's explanation that each step consumes the output of the prior step. Reply, 4-5, 7-8; EX1073, 84:10-21; 85:3-16. The Reply fails to identify a single passage in the specification supporting its position that the "additional acceptance" step ([1.4.3]) is a free-floating independent reanalysis of any object in the image, untethered to the prior steps that analyze and identify geometric forms.

F. Relabeling Broggi's Bounding-Box Steps Does Not Cure the Petition's Collapse of Distinct Claim Stages

The Reply does not identify distinct claim-corresponding stages in Broggi. It subdivides one vehicle-box localization routine—bottom-corner detection, top-of-

box confirmation, and stereo refinement—and assigns those sub-operations to multiple claim limitations. But relabeling sub-steps within a single lead-vehicle detection sequence does not create the claimed progression from candidate object, to plausibility screening, to a later acceptance analysis based on sensor-relative image formation. Reply, 19-21.

The Reply itself confirms that Petitioner’s [1.4.3] theory remains tied to the same bounding-box process. Tesla argues that Broggi’s “second, separate, and subsequent analysis step” is “detecting the top part of the bounding box” after locating the bottom corners, and that this second step operates using perspective and size constraints. Reply, 20. It then points to the very same box-completion logic and associated distance refinement as showing an “accepted object” suitable for trajectory computation. Reply, 21. But those are all still sub-operations within Broggi’s effort to detect and confirm the presence and distance of a lead vehicle. They are not different claim-corresponding gates with different technical inputs, outputs, and decision criteria. In Broggi, the search for the top of the box exists to complete and confirm the same vehicle localization task initiated by the bottom-corner search. Reply, 20-21. The Reply therefore fails to identify any later stage that takes already-plausible candidate destination objects and separately evaluates the shape of their image formation against known imaging properties of the sensor relative to the scene as claimed.

When asked about Broggi, Dr. Majumder testified that Broggi's analysis "compress[es]" the claimed stages together, explaining that "all of them have been compressed into one," rather than being implemented as the staged pipeline claimed in the '805 patent. EX1073, 95:9-21. She further explained that, in Broggi, there is "no pipeline where I am sieving through different things and rejecting," but instead a single box-detection flow in which the system determines where to look and then refines that same target. EX1073, 96:1-15. Thus, even if one accepts Petitioner's new mapping of bottom corners with [1.4.1] and top-of-box detection with [1.4.3], Dr. Majumder explained why those features still belong to one continuous vehicle-detection routine, not to the separate claim stages Tesla needs.

Dr. Majumder also confirmed that Patent Owner's criticism was not based on a misreading of isolated labels, but on the substantive overlap in the Petition's mapping. She testified that it remained her view that the Petition relies on Broggi's "top-of-box detection for both 1.4.1 and 1.4.3." EX1073, 97:21-98:3. Tesla's Reply answers by insisting that the Petition's formal mapping assigned the top-of-box step only to [1.4.3]. Reply, 19-20. But that does not resolve the underlying problem she identified. Her point was that Petitioner's Broggi theory still uses one bounding-box-detection methodology, with its interrelated bottom-corner, top-of-box, and refinement operations, as though it disclosed the '805 patent's distinct candidate-identification, plausibility, and later acceptance stages. EX1073, 95:9-21; 97:21-

98:3. The Board should focus on that substantive mismatch, not on Petitioner’s post hoc effort to relabel which sentence in the Petition corresponds to which sub-limitation.

Tesla’s own expert underscores the same defect. Dr. Rodriguez now says his prior declaration relied on Broggi’s “separate and subsequent step of detecting the top part of the bounding box” for [1.4.3], not [1.4.1], and that this “does not collapse distinct steps into a single step.” Reply, 19-20. But in the same breath, he explains that the [1.4.3] theory is still based on Broggi’s detecting the top of the preceding vehicle using the same pattern-matching framework applied to the bottom corners, with perspective and size constraints and later distance refinement. Reply, 20-21. Thus, the Petitioner has not identified a distinct downstream acceptance stage; it has merely subdivided one bounding-box routine into smaller pieces and assigned those pieces to different claim limitations.

Nor does Tesla’s reliance on “multi-step” rhetoric solve the problem. Patent Owner does not dispute that Broggi has multiple operations. The issue is whether those operations correspond to the specific ordered stages the claims require. Reply, 19-20. The ’805 patent claims a staged extraction architecture in which earlier steps identify and narrow candidate destination-describing geometry, and a later acceptance that stage then evaluates whether the image formation of those candidate objects is consistent with known sensor-relative imaging properties. Tesla’s Broggi

theory still does not show that kind of later gate. Instead, it points to the completion and confirmation of a vehicle bounding box. Reply, 20-21. As Dr. Majumder explained, that is still part of detecting and refining the vehicle box, not a separate downstream acceptance analysis of the sort the patent describes. EX1073, 95:9-21; 96:1-15.

The same is true of Tesla's invocation of stereo-based distance refinement. The Reply argues that only objects that pass the "second analysis" and complete a bounding box with associated distance refinement are treated as accepted and stored. Reply, 21. But that simply adds another confirmation sub-step within the same lead-vehicle localization routine. It still does not identify the claim's separate acceptance analysis based on "a shape of the image formation" and "knowledge of at least one imaging property of the imaging sensor relative to the surrounding field." Reply, 21.

Accordingly, even accepting Tesla's revised characterization of the Petition, the same defect remains: Broggi does not disclose distinct claim stages with distinct technical roles. It discloses one bounding-box detection pipeline in which bottom-corner identification, top-of-box confirmation, and stereo refinement all serve the same immediate purpose of localizing and tracking a preceding vehicle. Reply, 19-21.

III. The Reply Still Fails to Show the Claimed Tree Structure Being Traversed and Used in the Required Downstream Workflow

Even under Tesla's broader reading of "mathematical tree structure," the Reply still fails because it never shows the stored relational structure later traversed and used in the claimed downstream workflow. The claims do not recite a tree as a label in the abstract; they recite a structure that stores interrelationships of the identified edge segments and then carries those relationships into subsequent candidate-object, plausibility, and acceptance analyses. Reply, 9-10, 23-24. Dr. Majumder's testimony reinforces that functional point: the tree is created to store those interrelationships and later traversed so the downstream stages can retrieve them. The annotated figures underscore the same point by showing disconnected components rather than one stored-and-traversed hierarchy.

That is where Tesla's merits showing fails. The Reply asserts that Huttenlocher "expressly teaches" storing interrelationships of edge segments in a "scale-space segment tree," and further argues that the combination uses Huttenlocher's tree structure in both limitations [1.4.1] and [1.4.3]. Reply, 23-24. But that does not show that the same stored structure is the one later traversed and used to assemble and evaluate destination-describing geometry. Dr. Majumder testified that claim 1 requires traversal "when it is going down the pipeline to figure out which are plausible objects," and that traversal is used downstream to retrieve the stored relationships needed for the later filtering steps. EX1073, 40:21-41:5.

Tesla's Reply, by contrast, simply jumps from "Huttenlocher has a multiscale tree" to "therefore the limitation is met." Reply, 23-24. That leap is especially problematic because Petitioner's own theory repurposes Huttenlocher's general feature hierarchy as part of Broggi's vehicle-detection routine, rather than showing the claim-specific workflow in which a single stored relational structure feeds the subsequent candidate-object, plausibility, and later analysis stages.

Claim 2 makes that defect even clearer. Tesla is correct that Patent Owner preserves its position that a "mathematical tree structure" is a single connected tree rather than a forest. Reply, 9-10. But the more important point for the Sur-Reply is that claim 2 requires "processing the mathematical tree structure using a tree traversal algorithm," and Tesla still does not show that what Huttenlocher discloses is processed as that same claimed structure. Reply, 23-24. Dr. Majumder testified that the tree is first created and then "traversed as a unit by the next stage," and she explained that "traversed as a unit" means the traversal remains within that tree because "you cannot through a tree have traversal to a node which is not connected to the tree." EX1073, 36:2-6; 41:8-15. She further testified that a tree means "a single tree," not "multiple trees," and that, in her view, "computer science excludes forest." EX1073, 42:2-3; 42:18-19. Tesla's Reply does not answer that functional point. Even if Tesla is correct that the word "structure" is broad enough to include more than one tree in some abstract sense, Reply, 9-10, claim 2 still requires processing

“the mathematical tree structure” with “a tree traversal algorithm,” not merely pointing to a figure with multiple levels and calling the coarse-to-fine comparison a traversal. Dr. Majumder’s testimony underscores why the claims require more: the stored structure must support later traversal to extract the correct relational information needed in the next stage. EX1073, 38:9-18.

Tesla’s reliance on Sedgewick’s statement that a forest can be represented as a binary tree does not cure the problem. Reply, 9-10. Even under Tesla’s broader reading, the question is still what the claimed structure stores and how that stored information is later used. Dr. Majumder testified that the patent’s use of a tree is tied to connectivity, unique-path structure, and downstream traversal of the same stored relationships. EX1073, 34:21-35:8; 36:2-6. She also explained that the reason for the structure is to support “hierarchical search by traversal over relations,” so the system can assemble higher-order shapes such as rectangles from lower-level edge primitives. EX1073, 31:19-32:4. A representation argument does not establish that Huttenlocher stores the interrelationships of the downstream-used plurality in a structure that can be traversed in that way. Reply, 9-10, 23-24. At most, Tesla argues that some subset of segments in some portion of Huttenlocher’s multiscale display could qualify. But that is precisely the point: Petitioner is not showing a single claimed structural workflow; it is selecting convenient pieces from Huttenlocher’s display and then assigning them to the claim by hindsight.

The deposition also makes clear why this is not just a construction dispute. Dr. Majumder was asked directly whether storing interrelationships itself requires traversal, and she answered that storing happens during creation of the tree, while traversal is the later process used downstream to access the stored relationships. EX1073, 37:13-38:4; 39:5-13. She further testified that, if edge segments were stored in a tree but the tree were never traversed, that might satisfy mere storage, but it “will not satisfy doing anything with the tree.” EX1073, 39:14-21. That testimony is fatal to Tesla’s Reply on the merits. Tesla repeatedly says Huttenlocher teaches a multiscale tree and that coarse-to-fine matching amounts to traversal. Reply, 23-24. But the claim requires more than a generic multilevel analysis. It requires a stored relational structure that is then actually used as the operative tree in the downstream analysis pipeline. Tesla never bridges that gap.

Nor does Tesla solve the problem by arguing that the claims require only “a plurality” of edge segments and only some unspecified quantity of “interrelationships.” Reply, 24. Even accepting that premise for argument’s sake, Tesla still has to show a structure that stores the interrelationships of the plurality it is relying on and that the same structure is then processed in the later steps. Dr. Majumder testified that the patent’s pipeline proceeds by carrying forward outputs from one stage to the next, and that the tree is the organized repository of relational information that the later stages access through traversal. EX1073, 36:2-6; 38:9-18;

40:21-41:5. Tesla's Reply does not identify where Huttenlocher stores the interrelationships of the very plurality it later uses to assemble and validate destination geometry in Broggi. Instead, it points to a red-box subset in Figure 4 and says that subset is enough. Reply, 24. But that only underscores the deficiency: Petitioner has narrowed to an after-the-fact subset rather than showing the claimed storing-and-processing sequence over a relational structure actually used downstream as the claims require.

Accordingly, the Board need not decide the entire forest-versus-tree question in Patent Owner's favor to reject Tesla's showing. Patent Owner preserves its position that a "mathematical tree structure" in the '805 patent is a single connected hierarchy, not a forest. But even under Tesla's broader framing, Petitioner still does not show the required stepwise workflow: first storing the interrelationships of the identified plurality in a mathematical tree structure and then using that same stored structure in the downstream analytical sequence, including claim 2's traversal step. Reply, 23-24. The Reply, therefore, still fails to prove this limitation as actually claimed.

IV. Broggi's Moving Lead Vehicle Is Not the Claimed "Potential Destination," "Most Proximate Destination," or Endpoint for Optimized Path Computation

Claim 1 requires an accepted destination object whose positional parameters are stored and then used as the endpoint of optimized path computation. Broggi does

not select and store such an endpoint. It continuously tracks a moving lead vehicle and updates steering to maintain formation. That is reactive platooning, not destination-based trajectory computation. Broggi's target is tracked, not selected and used as a destination in the claim's sense. Reply, 14-17, 25-27.

Dr. Majumder's testimony explains that distinction. She testified that a destination in the patent is "a fixed, stationary location" and a "fixed, physically meaningful endpoint" where the vehicle will "go and stop." EX1073, 46:10-21; 52:5-7. She also explained that "transient" means something continuously moving and that such objects are handled by "tracking," while the '805 patent is directed to stationary fixed objects where the vehicle identifies the destination and stops. EX1073, 48:3-19. And she confirmed that her understanding excludes "a moving vehicle from being a potential destination." EX1073, 49:21-50:2. That is enough to distinguish Broggi's reactive following of a moving leader.

Tesla points to Dr. Majumder's acknowledgment that a stopped vehicle could sometimes qualify. Reply, 26-27. But she immediately explained that the patent "specifically talks about immovable structures, like docking stations," that is, places where the vehicle will go and stop. EX1073, 50:6-20. The disputed issue is not whether only buildings can ever qualify. Instead, the issue is that Broggi's asserted target is a continuously moving vehicle being tracked in real time, is not a fixed stopping endpoint identified through the patent's geometric filtering process.

The column 7 “new destination” passage does not help Tesla. Dr. Majumder explained that the patent first runs the pipeline to arrive at one destination and then, only if that destination cannot be reached, selects another destination and runs “the whole pipeline” again. EX1073, 55:4-20; 56:2-22. She confirmed that those are “two different guidance attempts,” and that the first pipeline is completed before a new destination is chosen. EX1073, 56:20-22; 58:12-17. That is fundamentally different from Broggi’s continuous tracking of a moving leader.

The same problem appears in the “most proximate destination” and trajectory limitations. Dr. Majumder testified that the most proximate destination is the result of the patent’s filtering process and therefore must also be a fixed, physically meaningful endpoint. EX1073, 64:22-65:22; 66:9-67:6. She also testified that once the destination is selected, the optimized path is fixed until the destination changes, and that the patent does not support continuously readapting the trajectory along the way. EX1073, 68:2-14. Broggi’s continuously updated steering to maintain alignment with a moving vehicle is therefore not the claimed optimized travel path to an accepted destination. Reply, 25-27.

In the end, Tesla’s theory reduces “potential destination” to anything the vehicle happens to be heading toward at a given moment. Reply, 14-17, 25-27. But the ’805 patent claims more than that: a destination selected through the patent’s staged analysis and used as the endpoint for path computation. EX1073, 53:15-54:5;

64:22-65:22; 68:2-14. Broggi's moving lead vehicle is tracked, not selected and used as a destination in that claimed sense.

V. Conclusion

For all of these reasons, Petitioner's Replies do not cure the defects in its original showings. The record still demonstrates that Petitioner has not shown the claimed ordered workflow, the distinct downstream acceptance-analysis stage, the required tree-structure dataflow and traversal, or the claimed destination-based optimized path computation. The supplemental record likewise confirms that Petitioner never adequately disclosed or justified its parallel indefiniteness position before institution. The Board should therefore reject Petitioner's Reply arguments and confirm that Petitioner has not carried its burden to prove unpatentability of the challenged claims.

Respectfully submitted,

Dated: March 17, 2026

By: /Brandon R. Theiss/
Brandon R. Theiss
Registration No. 70,507
Counsel for Patent Owner
Intellectual Ventures II LLC

CERTIFICATE OF COMPLIANCE WITH WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), I certify that this **PATENT OWNER'S SUR-REPLY** complies with the type-volume limits set forth in the Memorandum and 37 C.F.R. § 42.24 because it contains 5,562 words, excluding the parts of the Brief that are exempted by 37 C.F.R. § 42.24(a), according to the word processing system used to prepare this Brief.

Respectfully submitted,

Dated: March 17, 2026

By: /Brandon R. Theiss/
Brandon R. Theiss
Registration No. 70,507

Counsel for Patent Owner
Intellectual Ventures II LLC

CERTIFICATE OF SERVICE

The undersigned certifies that pursuant to 37 C.F.R. § 42.6(e), a copy of the foregoing **PATENT OWNER'S SUR-REPLY** was served via email (as consented to by counsel) on March 17, 2026 to lead and backup counsel of record for Petitioner as follows:

Jonathan R. Bowser
jon.bowser.ipr@haynesboone.com
HAYNES AND BOONE, LLP
2801 N. Harwood Street, Suite 2300
Dallas, TX 75201

Scott T. Jarratt
scott.jarratt.ipr@haynesboone.com
HAYNES AND BOONE, LLP
2801 N. Harwood Street, Suite 2300
Dallas, TX 75201

Gregory P. Webb
greg.webb.ipr@haynesboone.com
HAYNES AND BOONE, LLP
2801 N. Harwood Street, Suite 2300
Dallas, TX 75201

Matthew Beck
matthew.beck.ipr@haynesboone.com
HAYNES AND BOONE, LLP
2801 N. Harwood Street, Suite 2300
Dallas, TX 75201

Michael McCarty
michael.mccarty@haynesboone.com
HAYNES AND BOONE, LLP
2801 N. Harwood Street, Suite 2300
Dallas, TX 75201

Ashraf Fawzy
afawzy@tesla.com
TESLA, INC.
1 Tesla Road
Austin, Texas 78725

Dated: March 17, 2026

By: /Brandon R. Theiss/
Brandon R. Theiss
Registration No. 70,507
Counsel for Patent Owner
Intellectual Ventures II LLC