

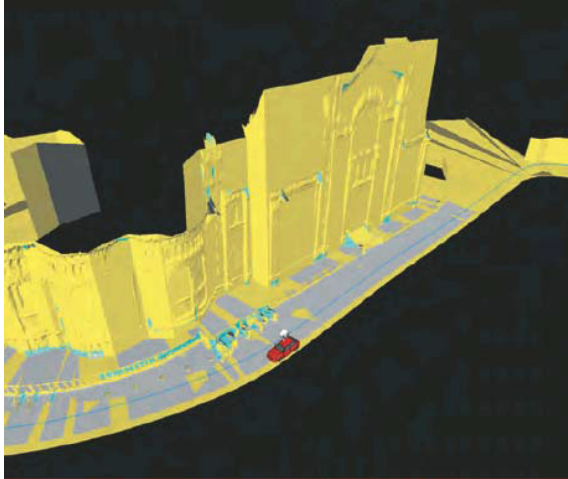
# EXHIBIT C

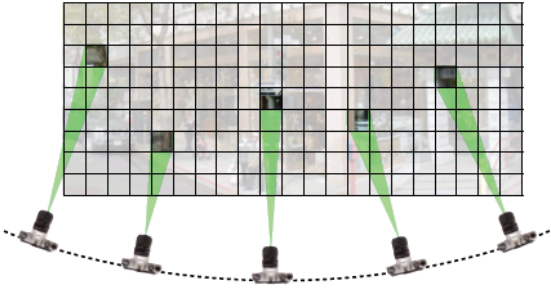
**EXHIBIT C**  
**to VRE's Disclosure of Asserted Claims and Infringement Contentions for U.S. Patent No. 8,078,396**

The claim chart below applies to Google Street View technology found in Google Earth products and all products incorporating the same or similar function detailed and illustrated herein. References in the claim chart below to specific Google products should not be construed to limit the products accused.

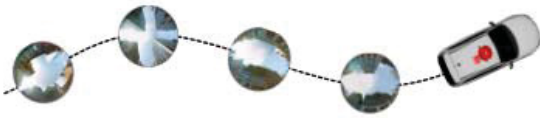
Discovery is ongoing, and these contentions may be supplemented as information specific to products incorporating Google Street View technology becomes available.

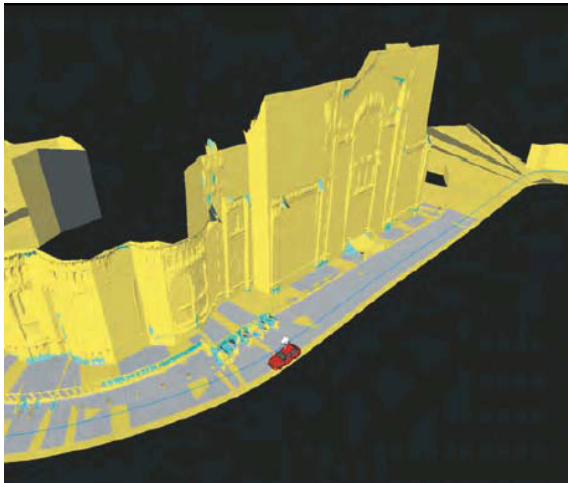
<b>Claim limitation</b>	<b>Practice by Accused Instrumentality</b>
1. A method for creating a continuum of image data, the method comprising:	<p>Google practices a method for creating a continuum of image data. Google creates a continuum of image data, such as a photorealistic 3D model, using Google Street View data.</p> <p style="text-align: center;">Taking extraction of 3D information even further, we use Street View data to create photorealistic 3D models for Google Earth. Traditionally, that service has created 3D city models from nadir or oblique airborne imagery, resulting in low-resolution facades with little detail—while suitable for fly-throughs, they didn't provide a pleasant walk-through experience. In contrast, 3D facade models reconstructed from Street View's laser scans and imagery are high resolution.</p> <p>Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 36.</p>

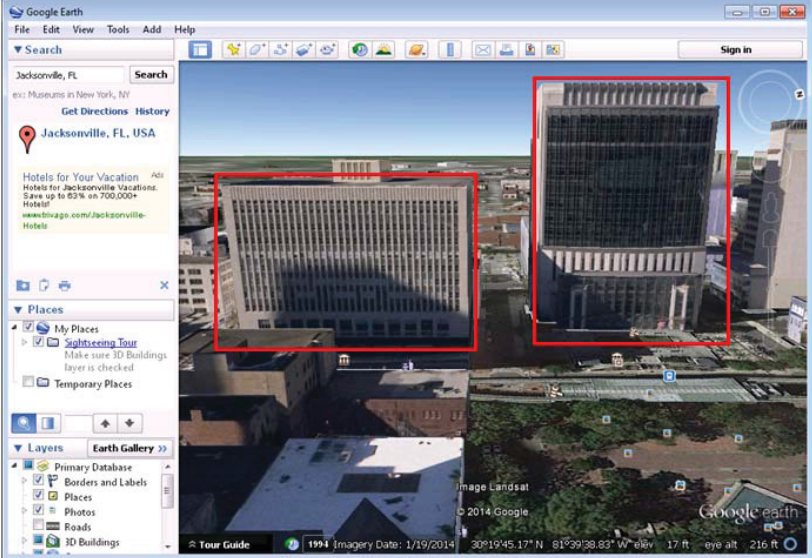
Claim limitation	Practice by Accused Instrumentality
	 <p data-bbox="737 1014 1252 1108"><b>Figure 4.</b> Imagery from new Street View vehicles is accompanied by laser range data, which is aggregated and simplified by robustly fitting it in a coarse mesh that models the dominant scene surfaces.</p> <p data-bbox="526 1140 613 1167"><i>Id.</i> at 35.</p>
<p>a) receiving multiple two-dimensional image data sets of a subject, wherein each image data set is captured from a disparate point on a continuum and each image data set comprises a plurality of features;</p>	<p>Google practices the step of receiving multiple two-dimensional image data sets of a subject. Each image data set is captured from a disparate point on a continuum and each image data set comprises a plurality of features. Google receives multiple two-dimensional images of a subject, such as a business or residential area containing buildings and/or structures, captured using a Google Street View capture vehicle. The images are captured from disparate points on a continuum formed by the capture vehicle's path as it traverses a path adjacent to the subject. Each image is associated with position, trajectory and range information and includes data describing a plurality of features.</p> <p>After filtering out noisy foreground objects, we synthesize a single consistent facade texture by aligning, blending, and mosaicking multiple individual camera</p>

Claim limitation	Practice by Accused Instrumentality
	<p data-bbox="721 533 1273 678">images. To avoid duplication caused by multiple passes, and to create a single consistent facade model for an entire city, we resolve residual pose inaccuracies between acquisition runs and determine the most suitable set of final 3D facade models. We register these models with existing</p> <p data-bbox="526 709 1360 764">Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 36-7.</p> <p data-bbox="721 795 1273 877">We compute the depth of various points in the scene using laser range scans where available and image motion (optical flow) when laser data isn't available.</p> <p data-bbox="526 909 634 932"><i>Id.</i> at 35-6.</p> <p data-bbox="721 966 1273 1047">Google Street View has a repository of billions of 2D images captured with rolling-shutter camera rigs along vehicle trajectories. Although we use GPS and inertial sensors</p> <p data-bbox="526 1079 1442 1134">Bryan Klinger et al., <u>Street View Motion-from-Structure-from-Motion</u>, IEEE Proceedings of the International Conference on Computer Vision, Apr. 14, 2013, at 1.</p>  <p data-bbox="721 1478 1273 1547">Figure 7. We establish a virtual, linear “feature camera” at each feature whose optical axis is the projection of the feature center through the generalized camera model.</p>

Claim limitation	Practice by Accused Instrumentality
	<p><i>Id.</i> at 4.</p> <p>The scale of SfM is limited by the complexity of matching features between images to establish tracks. For <math>n</math> unposed proximal images, <math>O(n^2)</math> image pairs generally must be compared to identify matches. As the density of imagery increases, this superlinear factor dominates computation, so we must sidestep it to achieve planet scale.</p> <p><i>Id.</i> at 5.</p> <p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
b) generating a composite image of the subject from portions of two or more of the multiple two-dimensional image data sets, wherein the portions are aligned in a dimension consistent with the continuum;	Google generates a composite image of the subject from portions of two or more of the multiple two-dimensional image data sets. The portions of the two-dimensional image data sets are aligned in a dimension consistent with the continuum. Google generates a composite image of the subject, such as a business or residential area containing buildings and/or structures, using portions of the captured two-dimensional image data aligned in a dimension consistent with the capture vehicle trajectory.

Claim limitation	Practice by Accused Instrumentality
	<p data-bbox="719 531 1271 947">Google Street View has a repository of billions of 2D images captured with rolling-shutter camera rigs along vehicle trajectories. Although we use GPS and inertial sensors to estimate the pose of this imagery, it still contains low-frequency error due to challenging GPS environments in cities and elsewhere. To improve the pose of these images, we have extended traditional Structure from Motion (SfM) techniques to construct a point-based model of the street-level world where each point carries both its geometric position as well as its local appearance from several views (see Figure 1). We use the appearance information from this model to find corresponding 3D points viewed from nearby images, and the geometric information to align the cameras that view them, thereby globally correcting the imagery's pose: motion-from-structure-from-motion.</p> <p data-bbox="527 974 1442 1031">Bryan Klinger et al., <a href="#">Street View Motion-from-Structure-from-Motion</a>, IEEE Proceedings of the International Conference on Computer Vision, Apr. 14, 2013, at 1.</p> <div data-bbox="724 1062 1260 1178">  </div> <p data-bbox="719 1205 1271 1257">Figure 2. The vehicle's path establishes good relative pose and natural connectivity for the panoramic imagery we capture.</p> <p data-bbox="527 1285 602 1312"><i>Id.</i> at 2.</p> <p data-bbox="719 1344 1271 1509">our SfM reconstruction. Similar to SfM with video, the natural linear path of the vehicle trajectory establishes potential connectivity between images: only images within a fixed window of each other along the path are considered for joint participation in image tracks. This puts a constant bound on matching and tracking effort per image.</p> <p data-bbox="527 1537 602 1564"><i>Id.</i> at 5.</p>

Claim limitation	Practice by Accused Instrumentality
	<div data-bbox="712 552 1276 1142">  <p data-bbox="737 1045 1252 1136"><b>Figure 4.</b> Imagery from new Street View vehicles is accompanied by laser range data, which is aggregated and simplified by robustly fitting it in a coarse mesh that models the dominant scene surfaces.</p> </div> <p data-bbox="526 1171 1360 1228">Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 36-7.</p> <p data-bbox="526 1255 1464 1478">For example, the image below is a composite image because it includes portions of at least two distinct two dimensional images. The composite image below is formed from a portion of a two dimensional image of first building (shown in the red box annotation on the left) and from a portion of a two dimensional image of a second building (shown in the red box annotation on the right). Each two dimensional image portion is contained within a two-dimensional image set. The portions of the two-dimensional image data sets (shown in the red box annotations) are aligned in a dimension consistent with the continuum formed by the trajectory of a Google Street View capture vehicle.</p>

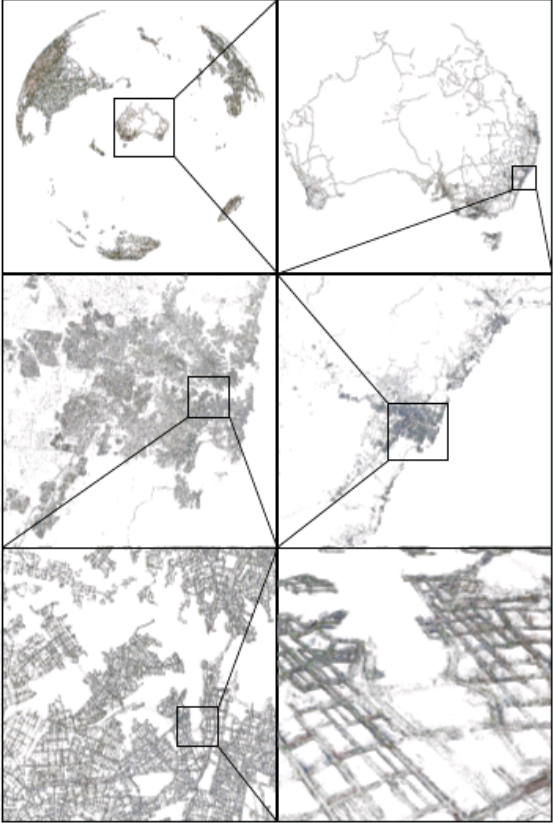
Claim limitation	Practice by Accused Instrumentality
	 <p>Screen shot of Google Earth (version 7.1.2.2041) showing the vicinity of 300 N Hogan St, Jacksonville, FL.</p> <p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
c) receiving data descriptive of a location of the plurality of features;	Google receives data descriptive of a location of the plurality of features. Google receives GPS data and orientation data that describes the location of the plurality of features shown in the images captured by the capture vehicle.



Claim limitation	Practice by Accused Instrumentality
	<p data-bbox="719 533 1271 747">To create our SfM model, we must overcome two important challenges not encountered in typical SfM problems: rolling-shutter cameras and planet-wide scale. We use the same tool to tackle both: a good initial estimate of the local vehicle trajectory. Although GPS may introduce global error of many meters in our estimate of the trajectory, the local shape is determined almost entirely by integration of inertial sensors. These sensors are trustworthy over short distances</p> <p data-bbox="526 779 1445 831">Bryan Klinger et al., <u>Street View Motion-from-Structure-from-Motion</u>, IEEE Proceedings of the International Conference on Computer Vision, Apr. 14, 2013, at 1.</p> <p data-bbox="719 863 1271 1003">We capture images using vehicles outfitted with sensors in addition to cameras: accelerometers, gyroscopes, wheel speed sensors and GPS receivers. We fuse this sensor data to establish an initial trajectory for the vehicle uninformed by the imagery. This initial pose has two interesting prop-</p> <p data-bbox="526 1035 602 1056"><i>Id.</i> at 5.</p> <p data-bbox="526 1087 1453 1255">To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
d) tracking the position of at least one of the plurality of features in two or more of the two-dimensional images;	Google tracks the position of at least one of the plurality of features in two or more of the two-dimensional images. Google tracks the position of features found within distinct two-dimensional images captured along the capture vehicle's trajectory.

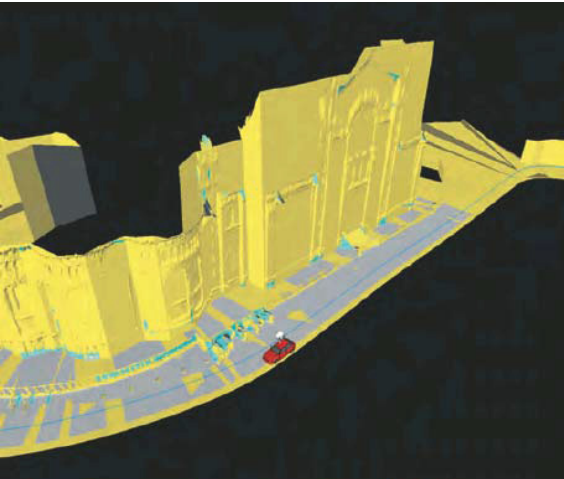
Claim limitation	Practice by Accused Instrumentality
	<p>Google Street View has a repository of billions of 2D images captured with rolling-shutter camera rigs along vehicle trajectories. Although we use GPS and inertial sensors to estimate the pose of this imagery, it still contains low-frequency error due to challenging GPS environments in cities and elsewhere. To improve the pose of these images, we have extended traditional Structure from Motion (SfM) techniques to construct a point-based model of the street-level world where each point carries both its geometric position as well as its local appearance from several views (see Figure 1). We use the appearance information from this model to find corresponding 3D points viewed from nearby images, and the geometric information to align the cameras that view them, thereby globally correcting the imagery's pose: motion-from-structure-from-motion.</p> <p>Bryan Klinger et al., <a href="#">Street View Motion-from-Structure-from-Motion</a>, IEEE Proceedings of the International Conference on Computer Vision, Apr. 14, 2013, at 1.</p> <p><b>5.2. Tracking</b></p> <p>A strong relative pose prior lets us avoid visual odometry via RANSAC-based relative pose estimation, which is often the first step in an SfM system. Instead, we get odometry from independent sensors. The relative pose is good enough for triangulation along this trajectory, so we can bootstrap the system by <i>tracking</i> features rather than by matching images.</p> <p>First, we extract invariant local features similar to SIFT [19], SURF [6], and HOG [8] from each image and construct feature cameras for each feature. We link features into tracks—a track represents a 3D world point—along the trajectory incrementally. Given incoming tracks from previous panoramas in a trailing window, and given the set of features in the current panorama, a track is extended only if we have both visual and geometric confirmation:</p>

Claim limitation	Practice by Accused Instrumentality
	<ul style="list-style-type: none"> <li>• <b>Visual confirmation:</b> The feature visually matches a feature in the track better than any other track. Matching is strict: the two features must be mutual 1-NN matches between the two images, and they must also pass a Lowe threshold test [19] of 0.7.<sup>2</sup></li> <li>• <b>Geometric confirmation:</b> The track triangulates with a 3D triangulation error below 1 meter and a reprojection error below 1 degree. Triangulation is done using standard methods with the track's set of feature cameras. The pose prior permits geometric verification of matches without RANSAC-based model estimation or heuristic outlier rejection.</li> </ul> <p><i>Id.</i> at 5-6.</p> <p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
e) generating a point cloud array for the at least one of the plurality of features;	Google generates a point cloud array for the at least one of the plurality of features found in Street View imagery.

Claim limitation	Practice by Accused Instrumentality
	 <p data-bbox="712 1381 1274 1514">Figure 1. An appearance-augmented point cloud comprising 404 billion tracked features, computed from street-level imagery. Every point in the cloud carries with it local appearance descriptors from at least three different viewpoints. Frames zoom successively closer to a detail of Sydney harbor.</p> <p data-bbox="526 1541 1446 1566">Bryan Klinger et al., <u>Street View Motion-from-Structure-from-Motion</u>, IEEE Proceedings of</p>

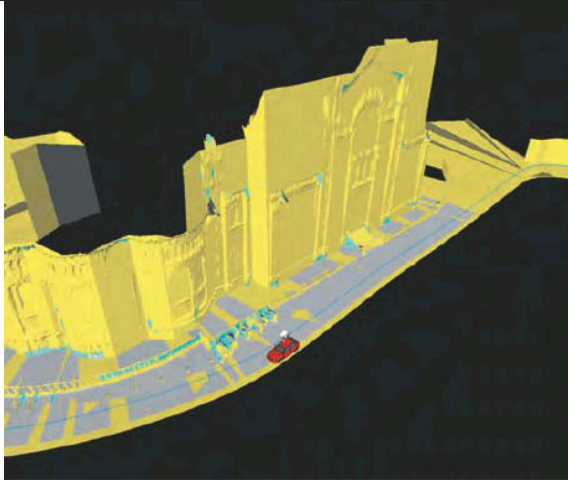
Claim limitation	Practice by Accused Instrumentality
	<p data-bbox="527 533 1219 560">the International Conference on Computer Vision, Apr. 14, 2013, at 1.</p> <p data-bbox="719 583 1273 863">With this approach, all the feature cameras within each panorama form a rigid assembly (porcupine) given by the relative pose prior. This assembly may move rigidly during bundling by changing the nominal rosette pose. Thus, the highest frequencies of pose are baked into the camera model; the medium frequency errors are reduced by bundling; the lowest frequency errors remain and must be addressed via loop closing. Figure 8 shows an example point cloud before and after bundle adjustment. Section 6 discusses loop closing.</p> <p data-bbox="527 890 602 917"><i>Id.</i> at 5.</p>

Claim limitation	Practice by Accused Instrumentality
	<div data-bbox="722 535 1263 1092"> <p data-bbox="714 1123 1271 1228">Figure 8. A comparison of an unbundled (top) to bundled (bottom) SfM point cloud, including top-down, block-level details of the clouds and inset histograms of reprojection error for all tracked feature points.</p> </div> <p data-bbox="527 1255 555 1283"><i>Id.</i></p> <p data-bbox="527 1310 1453 1480">To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
f) converting the point cloud array to a polygon based model; and	Google converts the point cloud array to a polygon based model. For example, a point cloud array is converted to a coarse mesh simplified by plane fitting.

Claim limitation	Practice by Accused Instrumentality
	<p data-bbox="716 554 1273 758">Imagery from new Street View vehicles is accompanied by laser range scans, which accurately measure the depth of a vertical fan of points on the two sides and the front of the vehicle. We aggregate this range data and simplify it by robustly fitting it in a coarse mesh that models the dominant scene surfaces. Figure 4 shows an example of a mesh simplified by plane fitting.</p> <p data-bbox="526 787 1360 842">Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 36.</p>  <p data-bbox="737 1356 1255 1451"><b>Figure 4.</b> Imagery from new Street View vehicles is accompanied by laser range data, which is aggregated and simplified by robustly fitting it in a coarse mesh that models the dominant scene surfaces.</p> <p data-bbox="526 1486 613 1514"><i>Id.</i> at 35.</p> <p data-bbox="526 1541 1414 1566">To the extent that the products incorporating Google Street View are not found to literally</p>


Claim limitation	Practice by Accused Instrumentality
	<p>infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
<p>g) associating a location for the polygon based model relative to the portions of the image data sets and based upon the location of the plurality of features.</p>	<p>Google associates a location for the polygon based model relative to the portions of the image data sets and based upon the location of the plurality of features. Google uses image data, range data and triangulation to associate a location for the polygon based model relative to the portions of the image data sets and based upon the location of the plurality of features.</p> <p style="text-align: center;">Imagery from new Street View vehicles is accompanied by laser range scans, which accurately measure the depth of a vertical fan of points on the two sides and the front of the vehicle. We aggregate this range data and simplify it by robustly fitting it in a coarse mesh that models the dominant scene surfaces. Figure 4 shows an example of a mesh simplified by plane fitting.</p> <p>Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 36.</p>

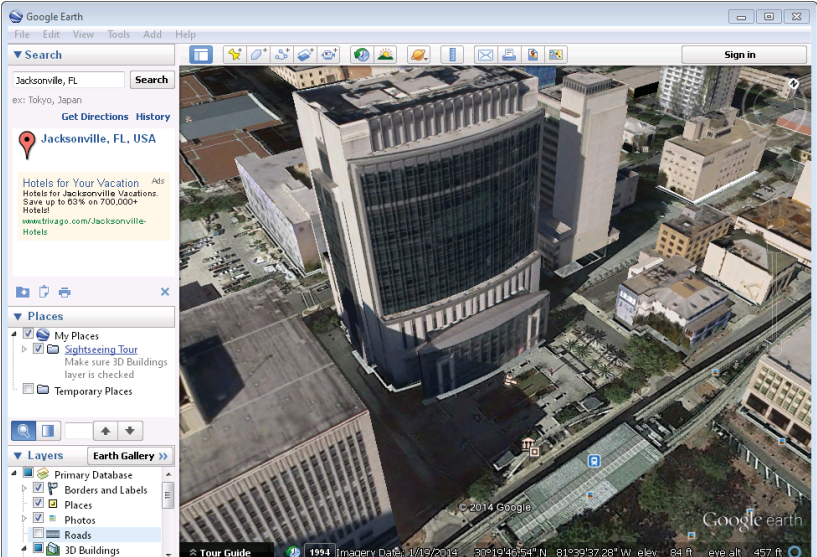


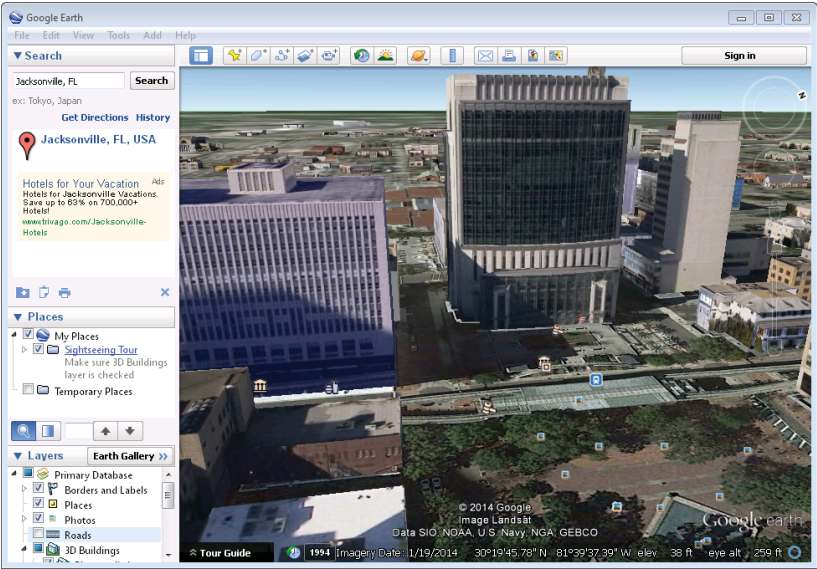
Claim limitation	Practice by Accused Instrumentality
	<div data-bbox="712 527 1276 1110">  <p data-bbox="737 1014 1252 1110"><b>Figure 4.</b> Imagery from new Street View vehicles is accompanied by laser range data, which is aggregated and simplified by robustly fitting it in a coarse mesh that models the dominant scene surfaces.</p> </div> <p data-bbox="524 1140 613 1169"><i>Id.</i> at 35.</p>

Claim limitation	Practice by Accused Instrumentality
	<p data-bbox="716 527 915 556"><b>3.2. Triangulation</b></p> <p data-bbox="716 573 1271 709">A fundamental operation in bundle adjustment is triangulation from multiple views. Given multiple image observations <math>\hat{x}_{im}^k</math> (superscripts hereafter dropped for clarity) of an unknown 3D world point, we wish to find the world point <math>x_w</math> that minimizes reprojection error:</p> $\operatorname{argmin}_{x_w} \sum_{\text{views}} \ x_{im} - \hat{x}_{im}\ ^2 \text{ s.t. Equation 4} \quad (5)$ <p data-bbox="716 806 1271 1003">The constraint is implicit in <math>x_{im}</math>, but we can make the constraint explicit with a simple assumption. Because rolling shutters tend to be fast, <math>t(x_{im})</math> is a slowly changing function and therefore <math>t(x_{im}) \approx t(\hat{x}_{im})</math>. Then we may avoid computing <math>t(x_{im})</math> entirely: the projection of <math>\hat{x}_w</math> into each camera may be done at the <i>a priori</i> known exposure times <math>t(\hat{x}_{im})</math> for each feature location <math>\hat{x}_{im}</math>. The optimization problem is now:</p> $\operatorname{argmin}_{x_w} \sum_{\text{views}} \left\  \underbrace{\overbrace{\hat{x}_{im}^T \cdot {}_cT_r \cdot {}_rT_w(t(\hat{x}_{im}))}^{\text{Lens Camera Pose}} \cdot x_w}_{x_{im}} - \hat{x}_{im} \right\ ^2 \quad (6)$ <p data-bbox="716 1115 1271 1220">which is a standard triangulation problem with known camera poses. As the rolling shutter approaches an instantaneous shutter, <math>t(x_{im})</math> becomes constant, and the equation simplifies to that of standard multi-view triangulation.</p> <p data-bbox="716 1220 1271 1388">This approximation degrades as <math>\hat{x}_{im}</math> and <math>x_{im}</math> diverge, but, critically for optimization, <i>the approximation is exact at the feature location</i>. Moreover, the approximation diverges slowly in practice: For our cameras, over a very large feature of diameter 100 pixels, lens distortion varies by a few percent and pose varies by ~1 cm at 30 kph.</p> <p data-bbox="526 1419 1443 1476">Bryan Klinger et al., <a href="#">Street View Motion-from-Structure-from-Motion</a>, IEEE Proceedings of the International Conference on Computer Vision, Apr. 14, 2013, at 3-4.</p> <p data-bbox="526 1503 1443 1558">To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person</p>

Claim limitation	Practice by Accused Instrumentality
	<p>having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
<p>2. The method of claim 1 additionally comprising the step of spraying image data onto the polygon based model, wherein the polygon based model comprises at least one three dimensional polygon based model and wherein the image data sprayed onto the polygon based model is based upon two-dimensional image data sets and the location of the polygon based model.</p>	<p>The polygon based model is a three dimensional polygon based model. Google sprays image data onto the polygon based model. The image data sprayed onto the polygon based model is based upon two-dimensional image data sets and the location of the polygon based model.</p> <p style="text-align: center;">Taking extraction of 3D information even further, we use Street View data to create photorealistic 3D models for Google Earth. Traditionally, that service has created 3D city models from nadir or oblique airborne imagery, resulting in low-resolution facades with little detail—while suitable for fly-throughs, they didn't provide a pleasant walk-through experience. In contrast, 3D facade models reconstructed from Street View's laser scans and imagery are high resolution.</p> <p>Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 36.</p>

Claim limitation	Practice by Accused Instrumentality
	<div data-bbox="716 527 1274 1430">  <p data-bbox="716 871 743 892">(a)</p> <p data-bbox="716 1234 743 1255">(b)</p> <p data-bbox="743 1276 1274 1409"><b>Figure 8.</b> Using Street View data to enhance user walk-through experiences in Google Earth. (a) Original 3D models of a New York City scene from airborne data only. (b) Fused 3D model with high-resolution facades. The visual quality is considerably higher, and many storefronts and signs can now be easily identified and recognized.</p> </div> <p data-bbox="524 1459 613 1486"><i>Id.</i> at 37.</p> <p data-bbox="524 1516 1474 1568">To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person</p>

Claim limitation	Practice by Accused Instrumentality
	having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.
[5pre] The method of claim 1 additionally comprising the steps of:	Google practices the method of claim 1. <i>See</i> claim 1.
[5a] repeating steps a) through f) multiple times each from disparate points along the continuum; and spraying image data onto the each polygon based model, wherein the image data sprayed onto the each polygon based model is based upon two-dimensional image data sets and the location of each respective polygon based model.	<p>Google repeats steps a) through f) multiple times each from disparate points along the continuum. <i>See</i> claims 1a) through 1f) above. Google sprays image data onto each polygon model. The image data sprayed onto each polygon based model is based upon two-dimensional image data sets captured at disparate points and the location of each respective polygon based model. Google repeats steps a) through f) and sprays image data onto each polygon based model of a building to create its three dimensional Google Earth database.</p> 

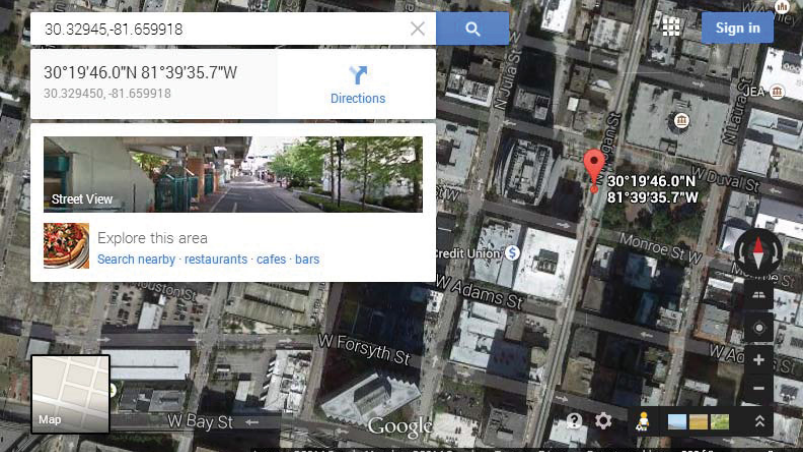
Claim limitation	Practice by Accused Instrumentality
	<p data-bbox="526 556 1417 611">Screen shot of Google Earth (version 7.1.2.2041) showing the vicinity of 300 N Hogan St, Jacksonville, FL.</p>  <p data-bbox="526 1230 553 1260"><i>Id.</i></p> <p data-bbox="719 1287 1271 1373">After filtering out noisy foreground objects, we synthesize a single consistent facade texture by aligning, blending, and mosaicking multiple individual camera</p>

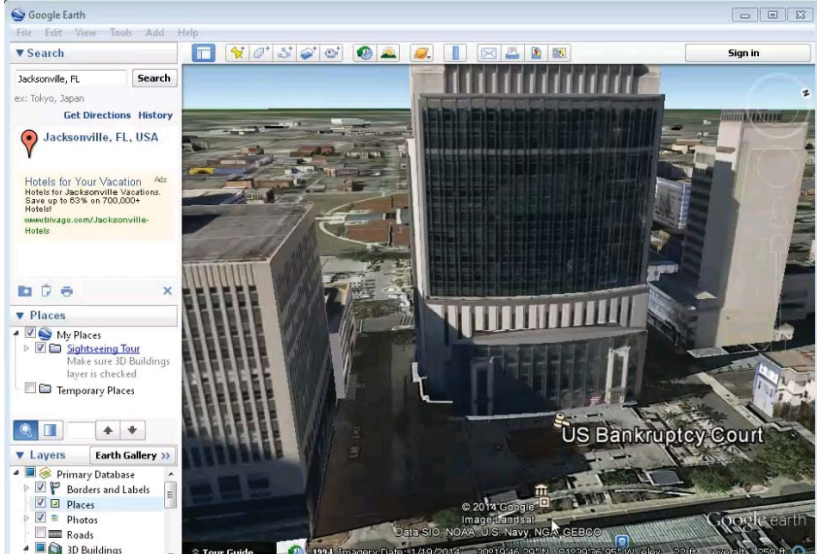
Claim limitation	Practice by Accused Instrumentality
	<p>images. To avoid duplication caused by multiple passes, and to create a single consistent facade model for an entire city, we resolve residual pose inaccuracies between acquisition runs and determine the most suitable set of final 3D facade models. We register these models with existing airborne models and then fuse them into a single model that has high-resolution facades as well as rooftops and back sides from an airborne view. Figure 8 shows how</p> <p>Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 36-7.</p> <p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
8. The method of claim 1 wherein the data descriptive of a location comprises a Cartesian coordinate.	Google practices the method of claim 1 wherein the data descriptive of a location comprises a Cartesian coordinate. Google receives position data formed by pairs of numerical coordinates that describe the location of the plurality of features shown in the images captured by the capture vehicle.

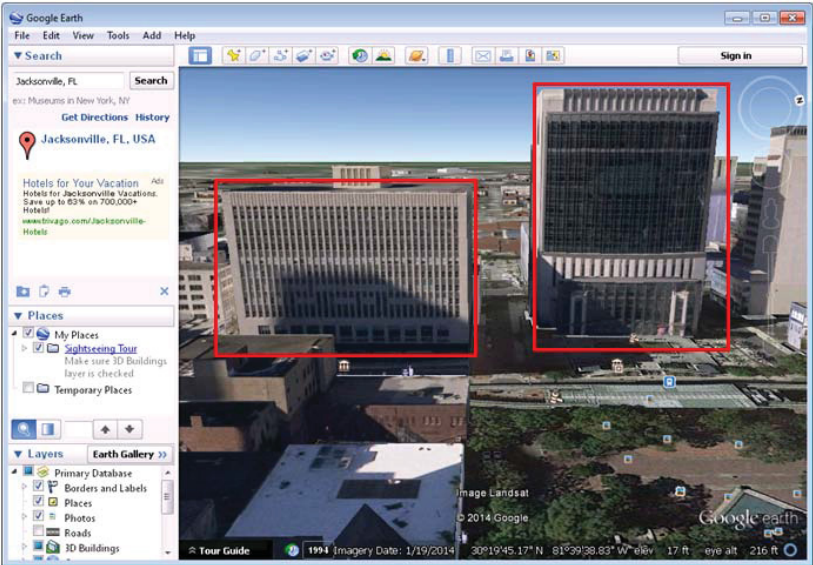


Claim limitation	Practice by Accused Instrumentality
	<p><b>POSE OPTIMIZATION</b></p> <p>Accurate position estimates of Street View vehicles are essential for associating our high-resolution panoramas with a street map and for enabling an intuitive navigation experience. We use the GPS, wheel encoder, and inertial navigation sensor data logged by the vehicles to obtain these estimates. An online Kalman-filter-based algorithm is deployed on the vehicles to provide real-time navigation information to the drivers. As part of the Street View processing pipeline, we use a batch algorithm open sourced by Google (<a href="http://code.google.com/p/gpo/wiki/GPO">http://code.google.com/p/gpo/wiki/GPO</a>) to achieve a smoother and locally accurate solution for the pose trajectory. This trajectory is computed at a resolution of 100 Hz, which is necessary to process laser data and to accurately correspond camera pixels to 3D rays in the presence of a rolling shutter.</p> <p>Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 34.</p> <p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
<p>9. The method of claim 8 wherein the data descriptive of a location of the plurality of features comprises a latitude and longitude coordinate.</p>	<p>Google practices the method of claim 8 wherein the data descriptive of a location of the plurality of features comprises a latitude and longitude coordinate. Google receives GPS data, including latitude and longitude coordinates, that describe the location of the plurality of features shown in the images captured by the capture vehicle.</p>

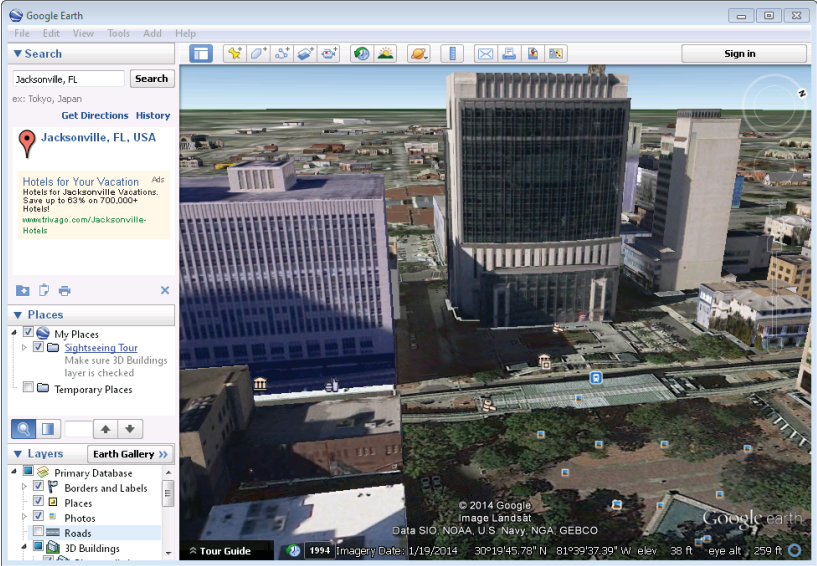


Claim limitation	Practice by Accused Instrumentality
	<p><b>POSE OPTIMIZATION</b></p> <p>Accurate position estimates of Street View vehicles are essential for associating our high-resolution panoramas with a street map and for enabling an intuitive navigation experience. We use the GPS, wheel encoder, and inertial navigation sensor data logged by the vehicles to obtain these estimates. An online Kalman-filter-based algorithm is deployed on the vehicles to provide real-time navigation information to the drivers. As part of the Street View processing pipeline, we use a batch algorithm open sourced by Google (<a href="http://code.google.com/p/gpo/wiki/GPO">http://code.google.com/p/gpo/wiki/GPO</a>) to achieve a smoother and locally accurate solution for the pose trajectory. This trajectory is computed at a resolution of 100 Hz, which is necessary to process laser data and to accurately correspond camera pixels to 3D rays in the presence of a rolling shutter.</p> <p>Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 34.</p> 

Claim limitation	Practice by Accused Instrumentality
	<p>30°19'46.0"N 81°39'35.7"W - Google Maps, <a href="https://www.google.com/maps/place/30%C2%B019'46.0%22N+81%C2%B039'35.7%22W/@30.3290881,-81.6619508,470m/data=!3m1!1e3!4m2!3m1!1s0x0:0x0">https://www.google.com/maps/place/30%C2%B019'46.0%22N+81%C2%B039'35.7%22W/@30.3290881,-81.6619508,470m/data=!3m1!1e3!4m2!3m1!1s0x0:0x0</a> (last visited June 15, 2014).</p> <p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
<p>10. The method of claim 1 additionally comprising the step of associating metadata with a polygon based model.</p>	<p>Google associates metadata, such as the name of a represented location, with a polygon based model. As shown below, metadata consisting of the name of a location (“US Bankruptcy Court”) is associated with a polygon based model representing an office building.</p> 

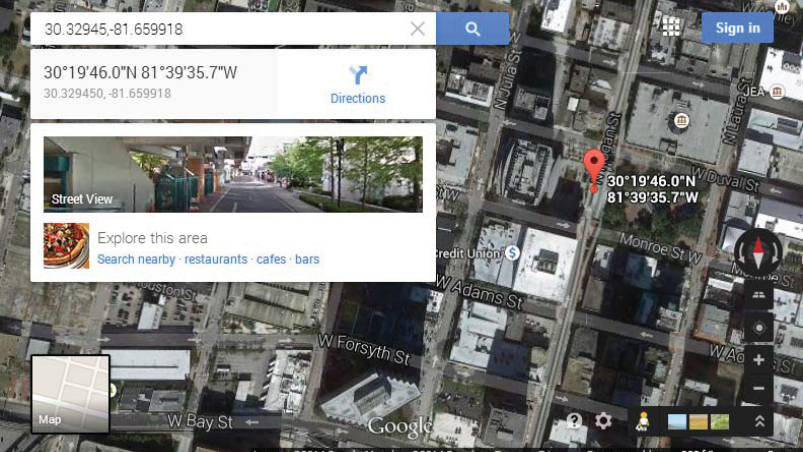
Claim limitation	Practice by Accused Instrumentality
	<p>Screen shot of Google Earth (version 7.1.2.2041) showing the vicinity of 300 N Hogan St, Jacksonville, FL.</p> <p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
<p>11. The method of claim 1, wherein the composite image comprises a single vertical slice of each of the two or more images.</p>	<p>Google practices the method of claim 1, wherein the composite image comprises a single vertical slice of each of the two or more images. As shown below in red annotations, the composite image comprises a single vertical slice of each of two two-dimensional images.</p>  <p>The screenshot shows the Google Earth interface with a search for Jacksonville, FL. Two buildings are highlighted with red rectangular boxes, each containing a single vertical slice of the building's facade. The interface includes a search bar, a list of places, and a layers panel on the left. The status bar at the bottom shows the date 1/19/2014 and coordinates 30°19'45.17" N, 81°39'30.83" W.</p>

Claim limitation	Practice by Accused Instrumentality
	<p>Screen shot of Google Earth (version 7.1.2.2041) showing the vicinity of 300 N Hogan St, Jacksonville, FL.</p> <p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
[12pre] The method of claim 11, additionally comprising the steps of:	Google practices the method of claim 11. <i>See</i> claim 11.
[12a] recording positional data descriptive of a respective location of each of the multiple two-dimensional image data sets of a subject; and	<p>Google records positional data descriptive of a respective location of each of the multiple two-dimensional image data sets of a subject. For each captured image, Google records positional data descriptive of the image location based at least upon one or more of GPS, wheel encoder and inertial navigation sensor data.</p> <p><b>POSE OPTIMIZATION</b></p> <p>Accurate position estimates of Street View vehicles are essential for associating our high-resolution panoramas with a street map and for enabling an intuitive navigation experience. We use the GPS, wheel encoder, and inertial navigation sensor data logged by the vehicles to obtain these estimates. An online Kalman-filter-based algorithm is deployed on the vehicles to provide real-time navigation information to the drivers. As part of the Street View processing pipeline, we use a batch algorithm open sourced by Google (<a href="http://code.google.com/p/gpo/wiki/GPO">http://code.google.com/p/gpo/wiki/GPO</a>) to achieve a smoother and locally accurate solution for the pose trajectory. This trajectory is computed at a resolution of 100 Hz, which is necessary to process laser data and to accurately correspond camera pixels to 3D rays in the presence of a rolling shutter.</p>

Claim limitation	Practice by Accused Instrumentality
	<p>Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 34.</p> <p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
<p>[12b] associating the composite image with a particular portion of the subject based upon the positional data.</p>	<p>Google associates the composite image with a particular portion of a subject, such as a business or residential area, based upon the recorded positional data.</p>  <p>Screen shot of Google Earth (version 7.1.2.2041) showing the vicinity of 300 N Hogan St, Jacksonville, FL.</p>

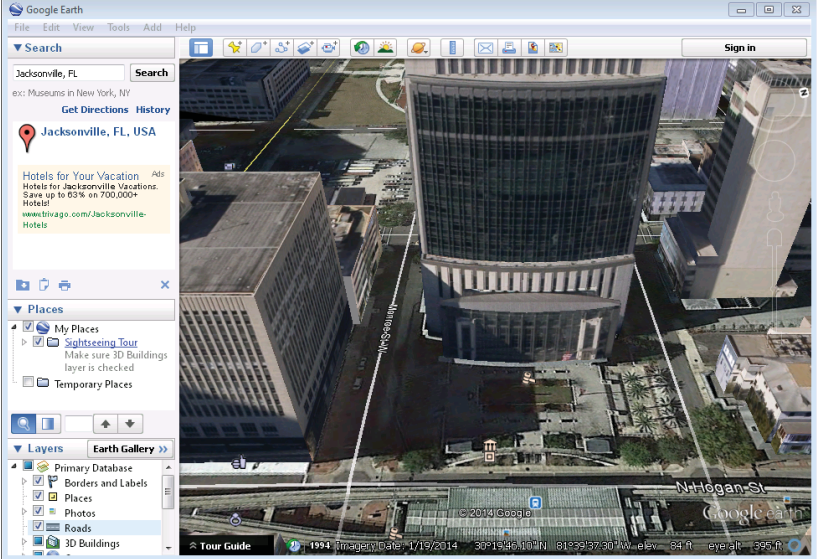


Claim limitation	Practice by Accused Instrumentality
	<p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
<p>13. The method of claim 12 wherein the subject comprises a geographic area and the positional data comprises a geospatial designation.</p>	<p>Google practices the method of claim 12 wherein the subject comprises a geographic area (such as a business or residential area) and the positional data comprises a geospatial designation (such as latitude and longitude coordinates). Google captures images of geographic areas, such as business and residential areas.</p> <div data-bbox="625 919 1318 1058" data-label="Text"> <p>10    8. Google operates vehicles that drive through business and  11    residential areas around the world collecting images.  12    9. The majority of these images are captured using a car with as  13    many as fifteen cameras mounted on top, as shown below.</p> </div> <p>Declaration of David Martin of Google, Inc. in Support of Google's Motion for Summary Judgment of Noninfringement of All Asserted Claims at 2:10-13, Vederi, LLC. v. Google, Inc., No. 10-cv-10774-AK-CW (C.D. Cal. May 2, 2012).</p> <p>Google Street View software records positional data comprising a geospatial designation, such as latitude and longitude coordinates, descriptive of the location of each captured image. For each captured image, Google records latitude and longitude data based at least upon one or more of GPS, wheel encoder and inertial navigation sensor data.</p>

Claim limitation	Practice by Accused Instrumentality
	<p><b>POSE OPTIMIZATION</b></p> <p>Accurate position estimates of Street View vehicles are essential for associating our high-resolution panoramas with a street map and for enabling an intuitive navigation experience. We use the GPS, wheel encoder, and inertial navigation sensor data logged by the vehicles to obtain these estimates. An online Kalman-filter-based algorithm is deployed on the vehicles to provide real-time navigation information to the drivers. As part of the Street View processing pipeline, we use a batch algorithm open sourced by Google (<a href="http://code.google.com/p/gpo/wiki/GPO">http://code.google.com/p/gpo/wiki/GPO</a>) to achieve a smoother and locally accurate solution for the pose trajectory. This trajectory is computed at a resolution of 100 Hz, which is necessary to process laser data and to accurately correspond camera pixels to 3D rays in the presence of a rolling shutter.</p> <p>Dragomir Anguelov et al., <u>Google Street View: Capturing the World at Street Level</u>, Computer, June 2010, at 34.</p> 

Claim limitation	Practice by Accused Instrumentality
	<p>30°19'46.0"N 81°39'35.7"W - Google Maps, <a href="https://www.google.com/maps/place/30%C2%B019'46.0%22N+81%C2%B039'35.7%22W/@30.3290881,-81.6619508,470m/data=!3m1!1e3!4m2!3m1!1s0x0:0x0">https://www.google.com/maps/place/30%C2%B019'46.0%22N+81%C2%B039'35.7%22W/@30.3290881,-81.6619508,470m/data=!3m1!1e3!4m2!3m1!1s0x0:0x0</a> (last visited June 15, 2014).</p> <p>To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>
<p>16. The method of claim 12 wherein the method additionally comprises the steps of overlaying metadata on the composite image descriptive of the composite image.</p>	<p>Google practices the method of claim 12 wherein the method additionally comprises the steps of overlaying metadata on the composite image descriptive of the composite image. For example, Google overlays street name metadata on a composite image. The street name metadata is descriptive of the composite image.</p>



Claim limitation	Practice by Accused Instrumentality
	 <p data-bbox="524 1119 1417 1171">Screen shot of Google Earth (version 7.1.2.2041) showing the vicinity of 300 N Hogan St, Jacksonville, FL.</p> <p data-bbox="524 1203 1453 1371">To the extent that the products incorporating Google Street View are not found to literally infringe this element, they nonetheless infringe under the doctrine of equivalents. A person having ordinary skill in the art would have considered the differences between the claim language and the product features set forth above to be insubstantial or would have found that they perform substantially the same function and work in substantially the same way to achieve substantially the same result as required by the claim.</p>