

SYSTEMS AND METHODS FOR GENERATING ANNOTATIONS OF
STRUCTURED, STATIC OBJECTS IN AERIAL IMAGERY USING GEOMETRIC
TRANSFER LEARNING AND PROBABILISTIC LOCALIZATION

5 CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is based on U.S. Application No. 16/669,309, filed October 30, 2019, the contents of which are hereby incorporated by reference in its entirety.

10 TECHNICAL FIELD

[0002] This disclosure relates generally to computer vision, and in particular but not exclusively, relates to detecting objects in aerial imagery.

BACKGROUND

15 [0003] Structured objects are objects that have a specific structure, either geometric or semantic. Examples of structured objects in outdoor environments include, but are not limited to, power poles, power lines, swimming pools, and traffic lights. Detecting such objects in aerial images is an important step in various applications such as mapping, image analysis, image stitching, and route planning.

20

SUMMARY

[0003a] According to a first aspect of the present invention, there is provided a non-transitory computer-readable storage medium having instructions stored thereon that, in response to execution by one or more processors of a computing system, cause the
25 computing system to perform actions comprising:

determining geographical locations of one or more structures in a geographic area;

obtaining an aerial image of the geographic area; and

annotating the aerial image based at least on the determined geographical

30 locations of the one or more structures in the geographic area;

wherein determining the geographical locations of one or more structures in the geographic area includes determining annotations of the one or more structures within aerial images of a set of aerial images of the geographic area; and

5 wherein determining annotations of the one or more structures within the aerial images of the set of aerial images of the geographic area includes:

receiving a first set of aerial images of the geographic area;

creating manual annotations of one or more structures within the first set of aerial images;

10 training a bootstrap machine learning model based on the first set of aerial images and the manual annotations;

receiving a second set of aerial images of the geographic area; and

creating model-generated annotations of the one or more structures within the second set of aerial images using the bootstrap machine learning model.

[0003b] According to a second aspect of the present invention, there is provided a
15 computer-implemented method for annotating aerial images, the method comprising:

determining annotations of one or more structures within aerial images of a set of aerial images of a geographic area;

determining geographical locations associated with the annotations;

20 updating a probabilistic localizer model based on the determined geographical locations; and

using the probabilistic localizer model to determine annotations for new aerial images of the geographic area;

wherein determining the geographical locations associated with the annotations includes, for each annotation and the associated aerial image:

25 determining a pose of a camera that captured the associated aerial image; and

determining the geographical location based on an image location of the annotation and the pose of the camera; and

30 wherein determining the geographical location based on the annotation and the pose of the camera includes:

determining a terrain model of the geographic area; and
determining an intersection between the terrain model and a ray
extending from the camera to the annotation.

[0003b] According to a third aspect of the present invention, there is provided a non-

5 transitory computer-readable storage medium having instructions stored thereon that, in response to execution by one or more processors of a computing system, cause the computing system to perform actions comprising:

determining geographical locations of one or more structures in a geographic area;

10 obtaining an aerial image of the geographic area; and

annotating the aerial image based at least on the determined geographical locations of the one or more structures in the geographic area;

wherein determining the geographical locations of one or more structures in the geographic area includes determining annotations of the one or more structures within
15 aerial images of a set of aerial images of the geographic area, and determining geographical locations associated with the annotations;

wherein determining the geographical locations associated with the annotations includes, for each annotation and the associated aerial image: determining a pose of a camera that captured the aerial image; and determining the geographical location based
20 on an image location of the annotation and the pose of the camera; and

wherein determining the geographical location based on the annotation and the pose of the camera includes:

determining a terrain model of the geographic area; and

25 determining an intersection between the terrain model and a ray extending from the camera to the annotation.

[0003d] According to a fourth aspect of the present invention, there is provided a non-transitory computer-readable medium having logic stored thereon that, in response to execution by one or more processors of a computing system, cause the computing system to perform actions comprising:

30 determining annotations of one or more structures within aerial images of a set of aerial images of a geographic area;

determining geographical locations associated with the annotations;
updating a probabilistic localizer model based on the determined geographical
locations; and

5 using the probabilistic localizer model to determine annotations for new aerial
images of the geographic area;

wherein determining the geographical locations associated with the annotations
includes, for each annotation and the associated aerial image:

determining a pose of a camera that captured the associated aerial image;

and

10 determining the geographical location based on an image location of the
annotation and the pose of the camera; and

wherein determining the geographical location based on the annotation and the
pose of the camera includes:

determining a terrain model of the geographic area; and

15 determining an intersection between the terrain model and a ray extending
from the camera to the annotation.

[0004] In some embodiments, a non-transitory computer-readable storage
medium is provided. The storage medium has instructions stored thereon that, in
response to execution by one or more processors of a computing system, cause the
20 computing system to perform actions comprising: determining geographical locations
of one or more structures in a geographic area; obtaining an aerial image of the
geographic area; and annotating the aerial image based at least on the determined
geographical locations of the one or more structures in the geographic area.

[0005] In some embodiments, a computer-implemented method for annotating
25 aerial images is provided. Annotations of one or more structures within aerial

CLAIMS

What is claimed is:

1. A non-transitory computer-readable storage medium having instructions stored
5 thereon that, in response to execution by one or more processors of a computing system, cause the computing system to perform actions comprising:
 - determining geographical locations of one or more structures in a geographic area;
 - obtaining an aerial image of the geographic area; and
 - 10 annotating the aerial image based at least on the determined geographical locations of the one or more structures in the geographic area;
 - wherein determining the geographical locations of one or more structures in the geographic area includes determining annotations of the one or more structures within aerial images of a set of aerial images of the geographic area; and
 - 15 wherein determining annotations of the one or more structures within the aerial images of the set of aerial images of the geographic area includes:
 - receiving a first set of aerial images of the geographic area;
 - creating manual annotations of one or more structures within the first set of aerial images;
 - 20 training a bootstrap machine learning model based on the first set of aerial images and the manual annotations;
 - receiving a second set of aerial images of the geographic area; and
 - creating model-generated annotations of the one or more structures within the second set of aerial images using the bootstrap machine learning model.
- 25 2. The computer-readable storage medium of claim 1, wherein annotating the aerial image based at least on the determined geographical locations of the one or more structures in the geographic area includes, for each structure of the one or more structures:
 - 30 determining, using a probabilistic localizer model, a most likely geographical location for the structure;

determining a pose of a camera that captured the aerial image; and
determining an image location for an annotation by determining an
intersection between an image plane of the aerial image and a ray extending from
the pose of the camera through a terrain model at the most likely geographical
5 location.

3. The computer-readable storage medium of claim 1, wherein the actions further
comprise:

adding the annotated aerial image to a training data set; and
10 training a machine learning model using the training data set.

4. The computer-readable storage medium of claim 1, wherein determining the
geographical locations of the one or more structures in the geographic area includes:
determining geographical locations associated with the annotations; and
15 updating a probabilistic localizer model based on the determined
geographical locations.

5. The computer-readable storage medium of claim 4, wherein updating the probabilistic
localizer model based on the determined geographical locations includes updating one or
20 more segments of the probabilistic localizer model based on the determined geographical
locations using Bayesian inferencing.

6. The computer-readable storage medium of claim 4, wherein determining the
geographical locations associated with the annotations includes, for each annotation and
25 the associated aerial image:

determining a pose of a camera that captured the aerial image; and
determining the geographical location based on an image location of the
annotation and the pose of the camera.

7. The computer-readable storage medium of claim 6, wherein determining the
geographical location based on the annotation and the pose of the camera includes:

determining a terrain model of the geographic area; and
determining an intersection between the terrain model and a ray
extending from the camera to the annotation.

5 8. The computer-readable storage medium of claim 1, wherein the actions further
comprise planning a route for an unmanned vehicle to avoid the determined
geographical locations of the one or more structures.

9. A computer-implemented method for annotating aerial images, the method
10 comprising:
determining annotations of one or more structures within aerial images of
a set of aerial images of a geographic area;
determining geographical locations associated with the annotations;
updating a probabilistic localizer model based on the determined
15 geographical locations; and
using the probabilistic localizer model to determine annotations for new
aerial images of the geographic area;
wherein determining the geographical locations associated with the
annotations includes, for each annotation and the associated aerial image:
20 determining a pose of a camera that captured the associated aerial
image; and
determining the geographical location based on an image location
of the annotation and the pose of the camera; and
wherein determining the geographical location based on the annotation
25 and the pose of the camera includes:
determining a terrain model of the geographic area; and
determining an intersection between the terrain model and a ray
extending from the camera to the annotation.

30 10. The computer-implemented method of claim 9, wherein updating the probabilistic
localizer model based on the determined geographical locations includes updating one or

more segments of the probabilistic localizer model based on the determined geographical locations using Bayesian inferencing.

11. The computer-implemented method of claim 9, wherein determining annotations of
5 the one or more structures within the aerial images of the set of aerial images of the geographic area includes:

receiving a first set of aerial images of the geographic area;
creating manual annotations of one or more structures within the first set of aerial
images;
10 training a bootstrap machine learning model based on the first set of aerial images and the manual annotations;
receiving a second set of aerial images of the geographic area; and
creating model-generated annotations of the one or more structures within the second set of aerial images using the bootstrap machine learning model.

12. The computer-implemented method of claim 11, wherein creating model-generated annotations of the one or more structures within the second set of aerial images includes
15 using the probabilistic localizer model to determine confidence scores for the annotations.

20 13. The computer-implemented method of claim 9, wherein determining the terrain model of the geographic area includes using a visual structure from motion (SFM) technique to determine the terrain model based on a plurality of aerial images of the geographic area.

25 14. The computer-implemented method of claim 9, wherein using the probabilistic localizer model to determine annotations for the new aerial images of the geographic area includes, for each new aerial image and for each structure of the one or more structures:

30 determining, using the probabilistic localizer model, a most likely geographical location for the structure;

determining a pose of a camera that captured the new aerial image; and
determining an image location for an annotation by determining an intersection
between an image plane of the new aerial image and a ray extending from the pose of the
camera through a terrain model at the most likely geographical location.

5

15. The computer-implemented method of claim 14, further comprising:
adding the new aerial image to an updated training data set; and
re-training a machine learning model using the updated training data set.

10

16. The computer-implemented method of claim 15, further comprising:
obtaining an aerial image of a new geographic area; and
annotating the aerial image of the new geographic area using the retrained
machine learning model.

15

17. The computer-implemented method of claim 9, further comprising planning a route
for an unmanned vehicle to avoid the determined geographical locations associated with
the annotations.

20

18. A non-transitory computer-readable storage medium having instructions stored
thereon that, in response to execution by one or more processors of a computing system,
cause the computing system to perform actions comprising:

determining geographical locations of one or more structures in a geographic
area;

25

obtaining an aerial image of the geographic area; and
annotating the aerial image based at least on the determined geographical
locations of the one or more structures in the geographic area;

wherein determining the geographical locations of one or more structures in the
geographic area includes determining annotations of the one or more structures within
aerial images of a set of aerial images of the geographic area, and determining

30

geographical locations associated with the annotations;

wherein determining the geographical locations associated with the annotations includes, for each annotation and the associated aerial image: determining a pose of a camera that captured the aerial image; and determining the geographical location based on an image location of the annotation and the pose of the camera; and

5 wherein determining the geographical location based on the annotation and the pose of the camera includes:

 determining a terrain model of the geographic area; and

 determining an intersection between the terrain model and a ray
 extending from the camera to the annotation.

10

19. The computer-implemented method of claim 18, wherein determining the geographical locations of the one or more structures in the geographic area further includes updating a probabilistic localizer model based on the determined geographical locations associated with the annotations.

15

20. The computer-implemented method of claim 18, wherein determining annotations of the one or more structures within the aerial images of the set of aerial images of the geographic area includes:

 receiving a first set of aerial images of the geographic area;

20 creating manual annotations of one or more structures within the first set of aerial images;

 training a bootstrap machine learning model based on the first set of aerial images and the manual annotations;

 receiving a second set of aerial images of the geographic area; and

25 creating model-generated annotations of the one or more structures within the second set of aerial images using the bootstrap machine learning model.

21. A non-transitory computer-readable medium having logic stored thereon that, in response to execution by one or more processors of a computing system, cause the
30 computing system to perform actions comprising:

- determining annotations of one or more structures within aerial images of a set of aerial images of a geographic area;
- determining geographical locations associated with the annotations;
- updating a probabilistic localizer model based on the determined geographical
- 5 locations; and
- using the probabilistic localizer model to determine annotations for new aerial images of the geographic area;
- wherein determining the geographical locations associated with the annotations includes, for each annotation and the associated aerial image:
- 10 determining a pose of a camera that captured the associated aerial image;
- and
- determining the geographical location based on an image location of the annotation and the pose of the camera; and
- wherein determining the geographical location based on the annotation and the
- 15 pose of the camera includes:
- determining a terrain model of the geographic area; and
- determining an intersection between the terrain model and a ray extending from the camera to the annotation.