METHOD FOR SEMANTIC LOCALIZATION OF AN UNMANNED AERIAL VEHICLE

ABSTRACT

A computer-implemented method comprises receiving an image captured by a camera on an unmanned aerial vehicle (UAV). The image depicts an environment below the UAV. A feature mask associated with the image is generated via a machine learning model that is trained to identify and semantically label pixels representing the environment depicted in the image. One or more reference tiles associated with the environment are retrieved. The reference tiles are associated with particular geographic locations and specify semantically labeled pixels representing the geographic locations. The semantically labeled pixels of the feature mask are correlated with the semantically labeled pixels of at least one of the one or more reference tiles to determine the geographic location of the UAV in the environment.

Fig. 1A