Summarizing and Searching Video: Geometry-Aware Visual Surveillance

Introduction

Tracking moving objects in a video is a fundamental problem in surveillance. This problem is made even more challenging if the camera is constantly moving, as it is in drone surveillance. In this work, we develop a pipeline that estimates camera motion on-the-fly while tracking, allowing us to "cancel" the camera motion before deploying our tracking algorithm. The tracker works by matching detected objects against future stabilized frames, with the help of a trajectory forecaster.

Method

We operate on images from pairs of frames:

1. Using the frames and the drone's approximate pose (height and angle from ground) estimate the depth of each pixel.
2. Stabilize for egomotion by re-projecting the 3D scene from a virtual camera placed in a stationary position.
3. Featurize the stabilized images with a learned Encoder-Decoder.
4. Crop the object to track from the first feature map; call this the template. Crop a relatively larger region from the second feature map; call this the search region.
5. Cross-correlate the template with the search region, obtaining a map of possible matches.
6. Refine the probability map using a velocity or trajectory estimate.
7. Take the best unique match.

Results

Our tracker achieves an average Intersection Over Union (IOU) score of 0.7958 on our synthetic dataset, whereas a similar un-stabilized tracker achieves 0.5185.

Our geometry-aware object tracker **stabilizes the 3D scene** during tracking, improving its ability to **match** and **forecast**.

Depth estimation

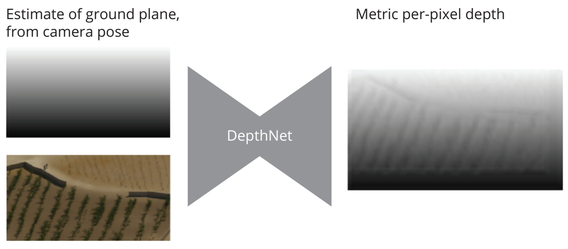

Egomotion Stabilization

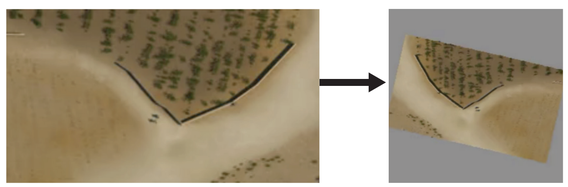

Tracking pipeline

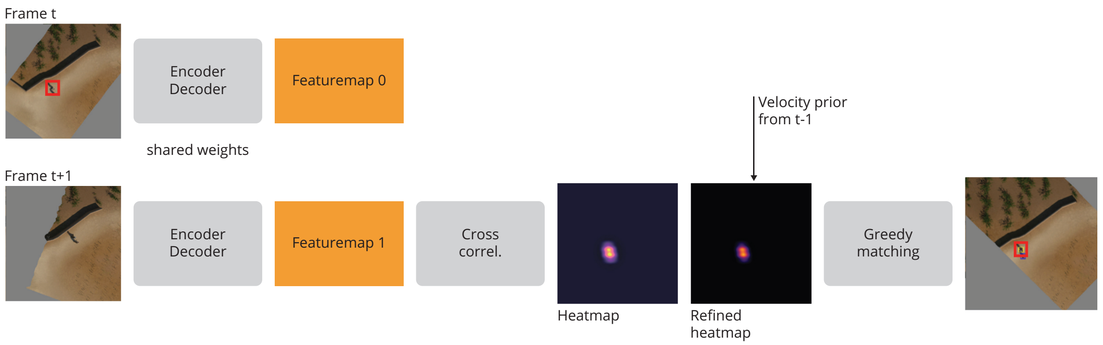

Standard RNN

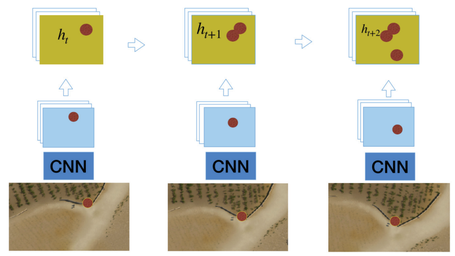

Geometry-Aware RNN

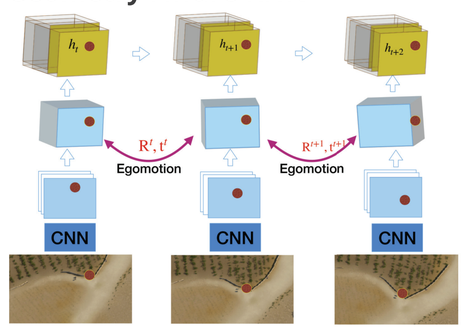

Stabilization boosts IOU by over 10%

Method	Mean IOU
Unstabilized	0.5185
Unstabilized+DataAug	0.5567
Unstabilized+DataAug+Hungarian	0.6240
Unstabilized+DataAug+Hungarian+VelocityPrior	0.6428
Stabilized+DataAug+Hungarian+VelocityPrior	**0.7958**

Trajectory forecasting is possible only in stabilized space

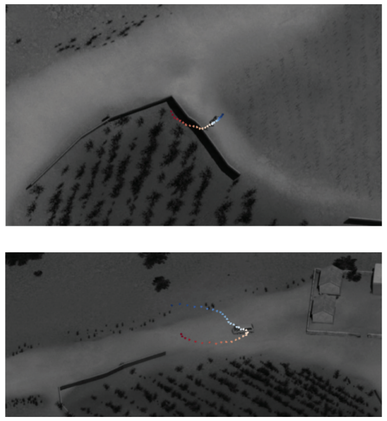

Raw Images and Trajectories

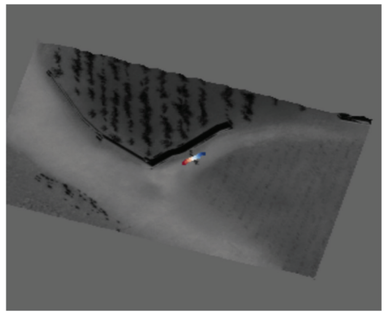

Stabilized Images and Trajectories

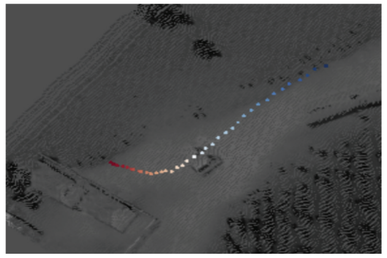

Carnegie Mellon University
Software Engineering Institute

Edwin Morris | ejm@sei.cmu.edu

DM19-1031