Fusion

Overview

The fusion node is responsible for taking the outputs of our SLAM system and the output of our object detection system, combining them, and publishing the results to the /tf topic.

Structure

Subscribes to:

- /vision/left - Object detections for left camera
- /SLAMpoints - Points for objects detected in SLAM

Publishes to:

- /tf - Global coordinates provided by SLAM combined with the labels from object detection

How to Run

Relies on the orb_slam node and the object detection node, as well as all their dependencies. To run, execute the following command from any directory:

```
$ rosrun robosub fusion
```

How it Works

The basic concept of the node is simple: take the bounding box points provided by the SLAM system and attach the labels from the object detection system so we get an identified object in our map. Both SLAM and object detection treat the cameras as 2D images when finding objects, so by calculating the closest x,y coordinate pairs within a single image we can reasonably determine which labels belong to which bounding boxes. Since SLAM provides both a 3D vector to the bounding box and the x,y coordinate within the image, the end result is a 3D vector to a labeled object in the map. The final step of the fusion node is to transform the object from its current frame to the "world" frame and publish that to /tf.