Supplementary Material for: Planning Paths Through Unknown Space by Imagining What Lies Therein

Yutao Han, Jacopo Banfi, and Mark Campbell
Department of Mechanical and Aerospace Engineering
Cornell University
United States
{yh675, jb2639, mc288}@cornell.edu

4th Conference on Robot Learning (CoRL 2020), Cambridge MA, USA.

1 Contents

The supplementary material includes randomly chosen visual examples of inpainting performance and path planner performance.

2 More Inpainting Examples

Figure 1 shows ten randomly chosen inpainting examples for both the lidar input and stereo input. The color key is shown in (Figure 4 of the paper). In rows 1-3 the network overpredicts the amount of road (pink pixels). In row 4, the road is cutoff early although the network fills in vegetation (dark green) well. In rows 5-10 the network does a good job of predicting the general structure of the scene and fills in the unobserved spaces with reasonable predictions of road (pink), sidewalk (purple), vegetation (dark green), building (yellow), and fence (orange).

3 More Planning Examples

Randomly chosen planning examples taken from the obstacle maps shown in (Figure 5 of the paper). Five different different examples are taken from images 45, 150, and 285 each. Refer to (Figure 6 in the paper) for a statistical analysis.

Figure 2 shows randomly chosen planning results for image 45. Due to the simplicity of the scene the paths are all fairly similar and the inpainting is not much of a factor in planner performance. Note that in rows 3 and 5 no path could be found.

Figure 3 shows randomly chosen planning results for image 150. In the first example (row 1), the lidar input image has a gap in the center shown by the light blue circle which is filled in as vegetation on the inpainted image. This causes the planner using the lidar input obstacle map to plan paths which run into deadends and eventually take much longer (441 steps) to reach the goal location compared to using the inpainted obstacle map (141 steps). For the other examples, planner performance is fairly similar. Note that in rows 4 and 5 no path could be found.

Figure 4 shows randomly chosen planning results for image 285. In the second example (row 2), the lidar input image has several gaps in the bottom center shown by the light blue oval (on row 1 to reduce clutter) which is filled in as building in the inpainted image. The planner using the lidar input obstacle maps attempts to plan paths through these gaps causing it to run into dead ends and take an unnecessarily long path (497 steps) before reaching the goal. The inpainted image has filled in those gaps and is able to plan a much shorter path (197 steps) similar to the ground truth path. For the other examples the inpainted image did not make a significant difference.

The inpainted obstacle map makes a significant difference for the planner when planning paths through occluded regions, as the lidar input obstacle map will frequently plan through occluded regions which contain dead ends. The inpainted obstacle maps are extremely helpful for complex urban scenes with structures which cause occlusions for a lidar or stereo sensor. For simpler scenes or paths through mostly visible regions the inpainted obstacle map does not make as much of a difference.

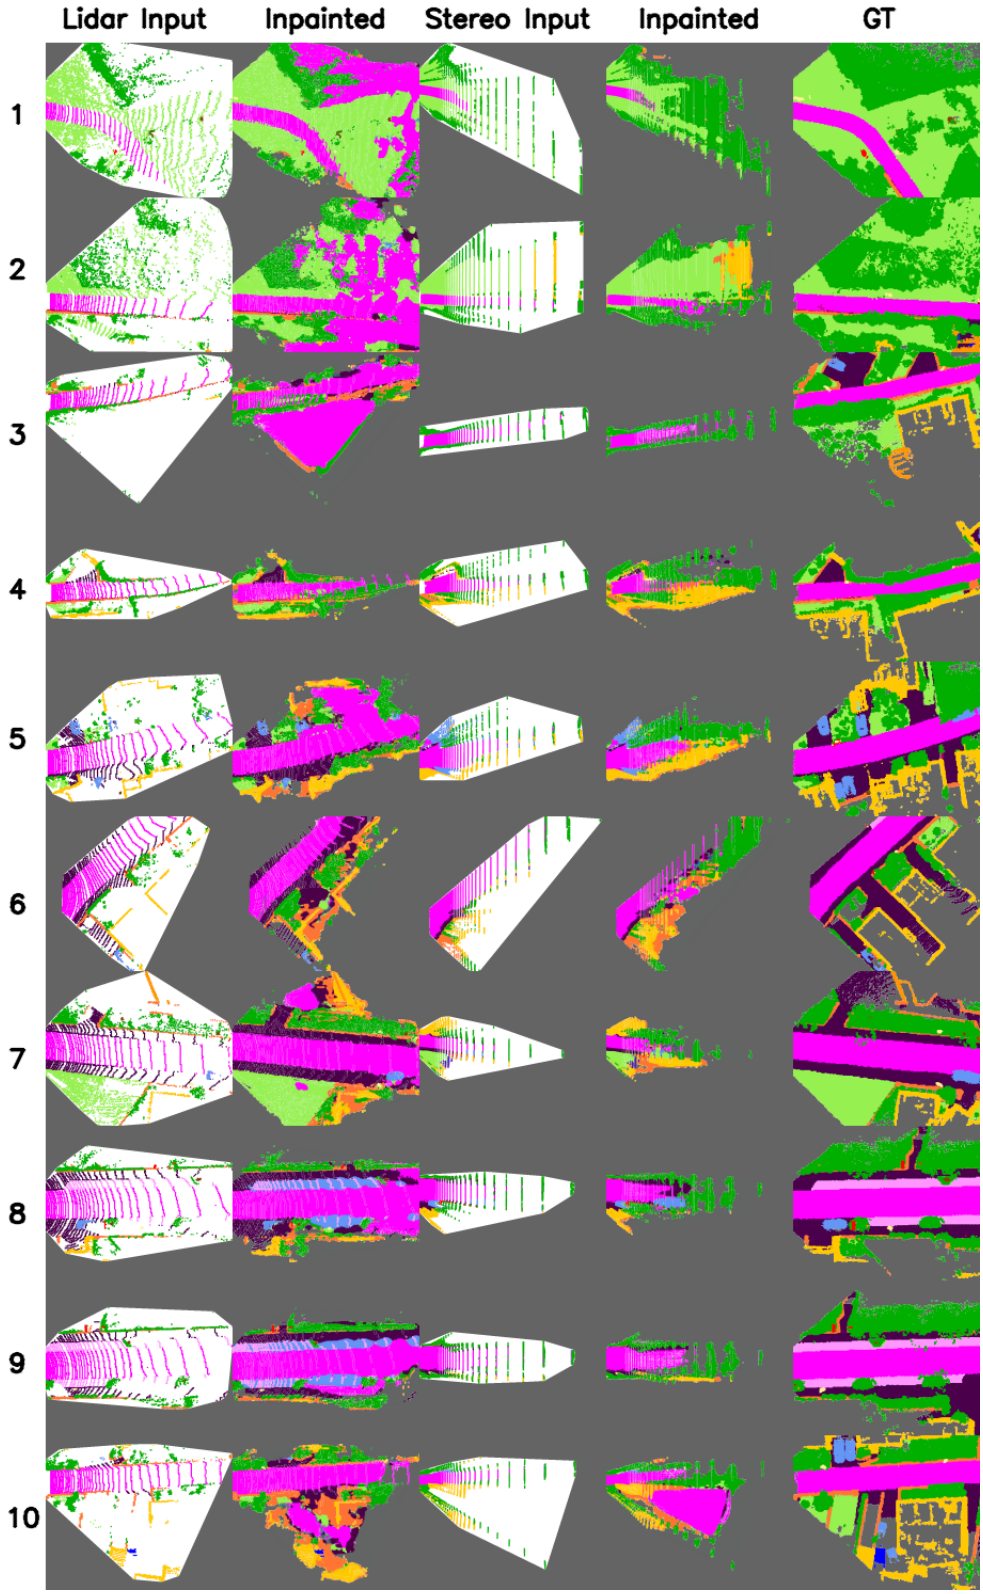

Figure 1: Qualitative evaluation of the inpainting network performance. Columns from left to right: (1) lidar input, (2) inpainted image, (3) stereo input, (4) inpainted image, (5) ground truth (GT).

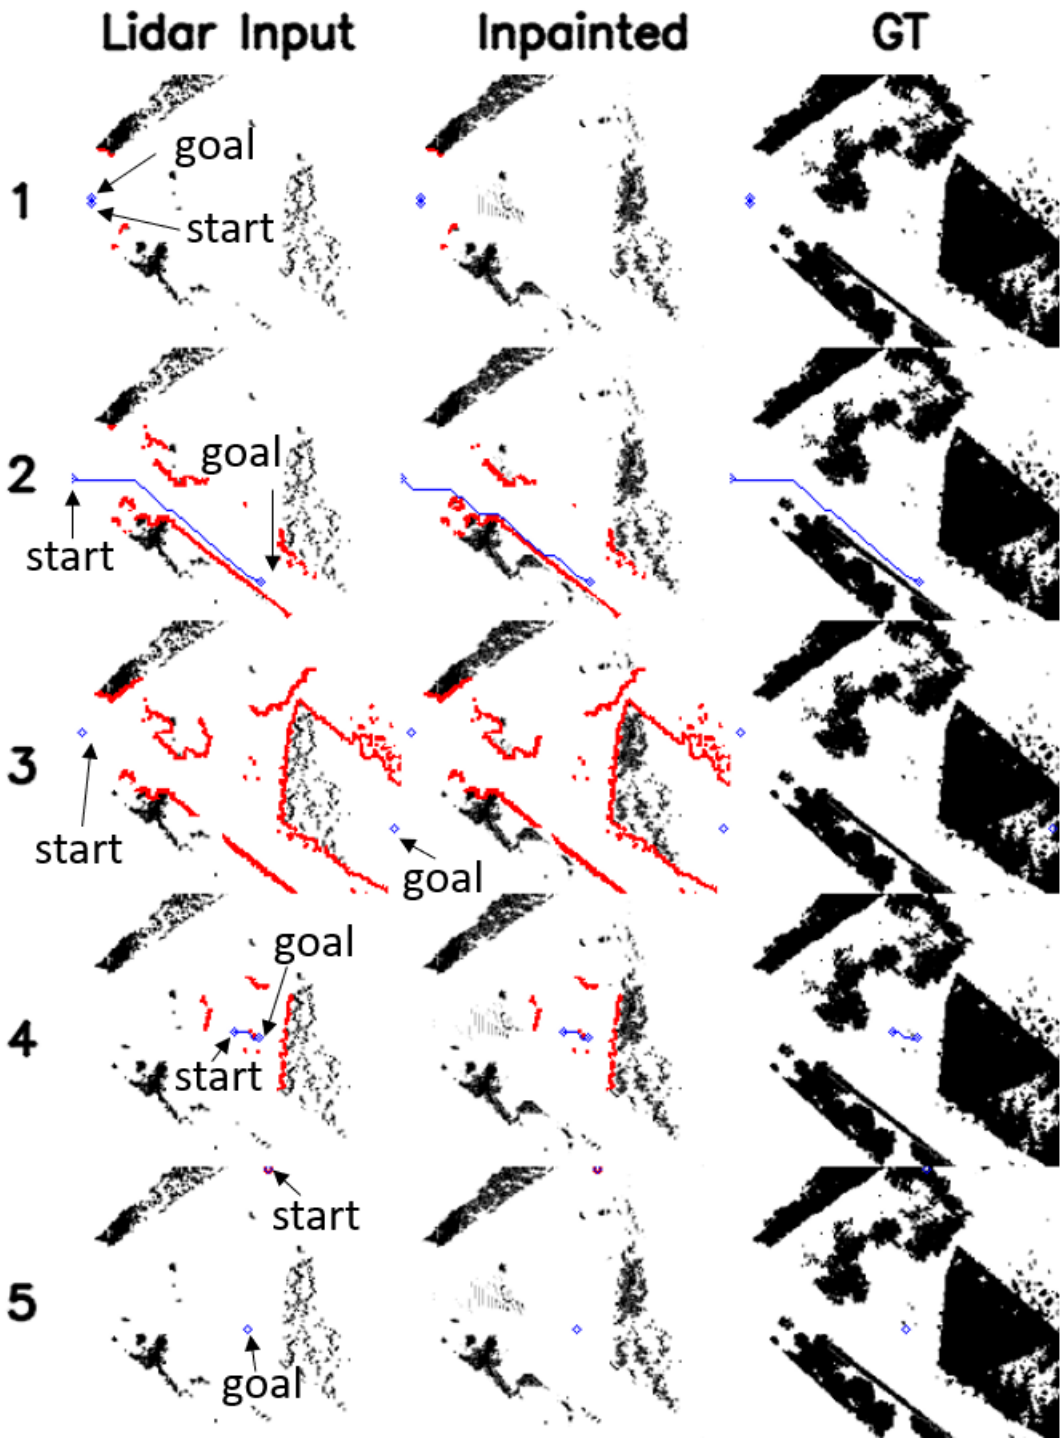

Figure 2: Randomly chosen examples of path planner performance for image 45. Columns from left to right: (1) lidar input, (2) inpainted image, (3) ground truth (GT). The path is in blue pixels. Red pixels are obstacles the robot observes online.

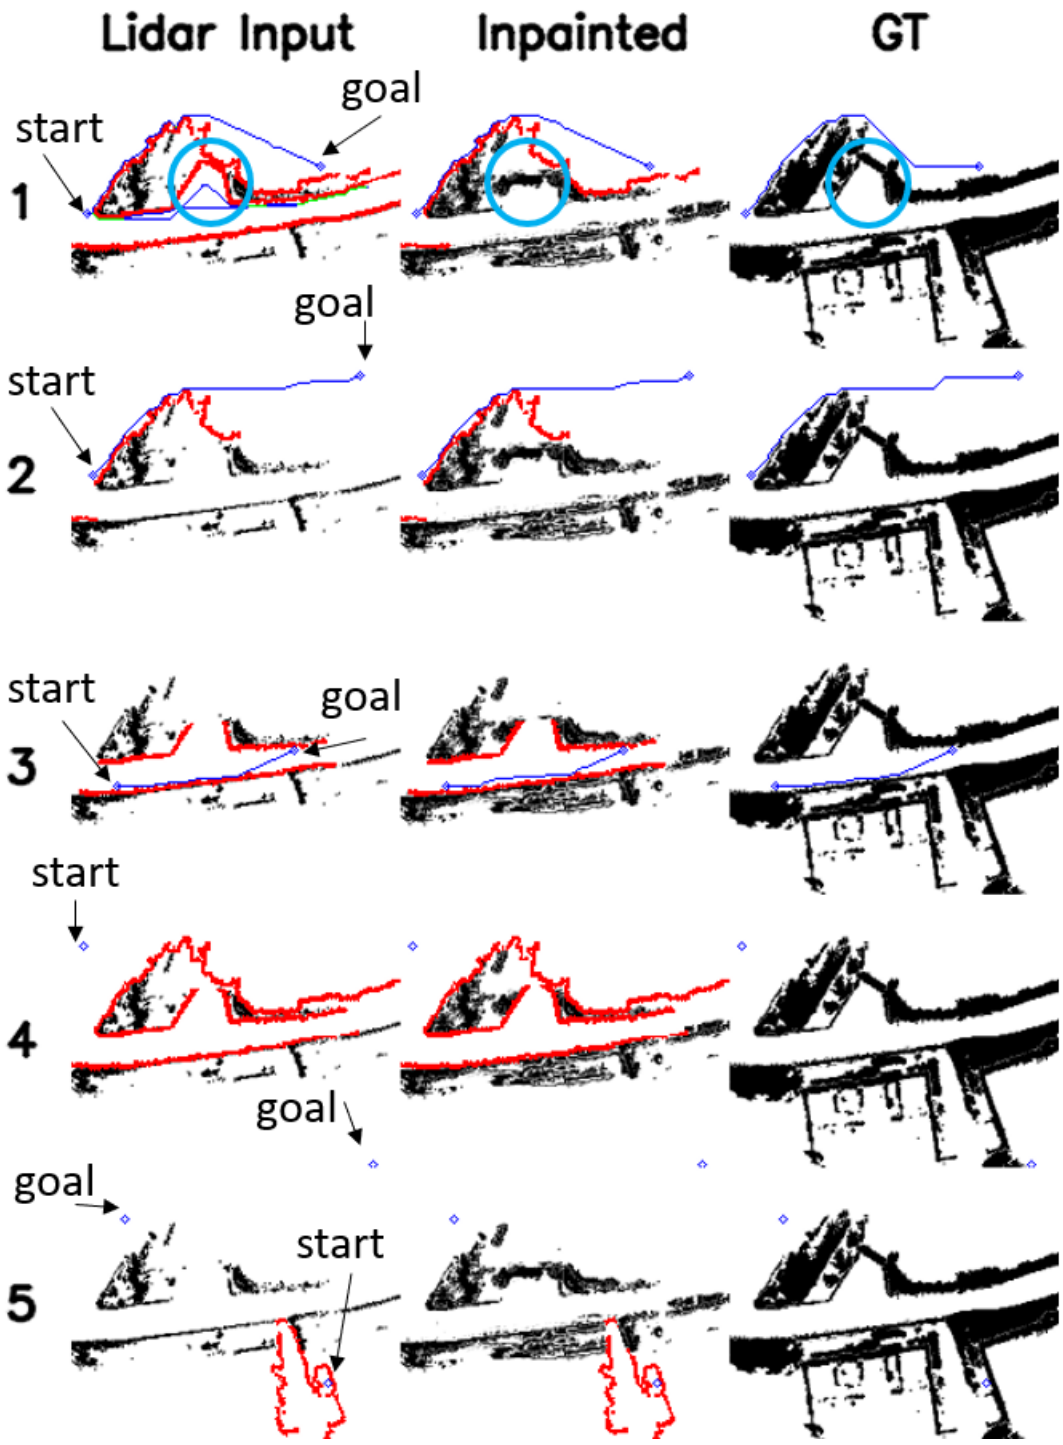

Figure 3: Randomly chosen examples of path planner performance for image 150. Columns from left to right: (1) lidar input, (2) inpainted image, (3) ground truth (GT). The path is in blue pixels. Red pixels are obstacles the robot observes online. Green pixels show where the path overlaps itself.

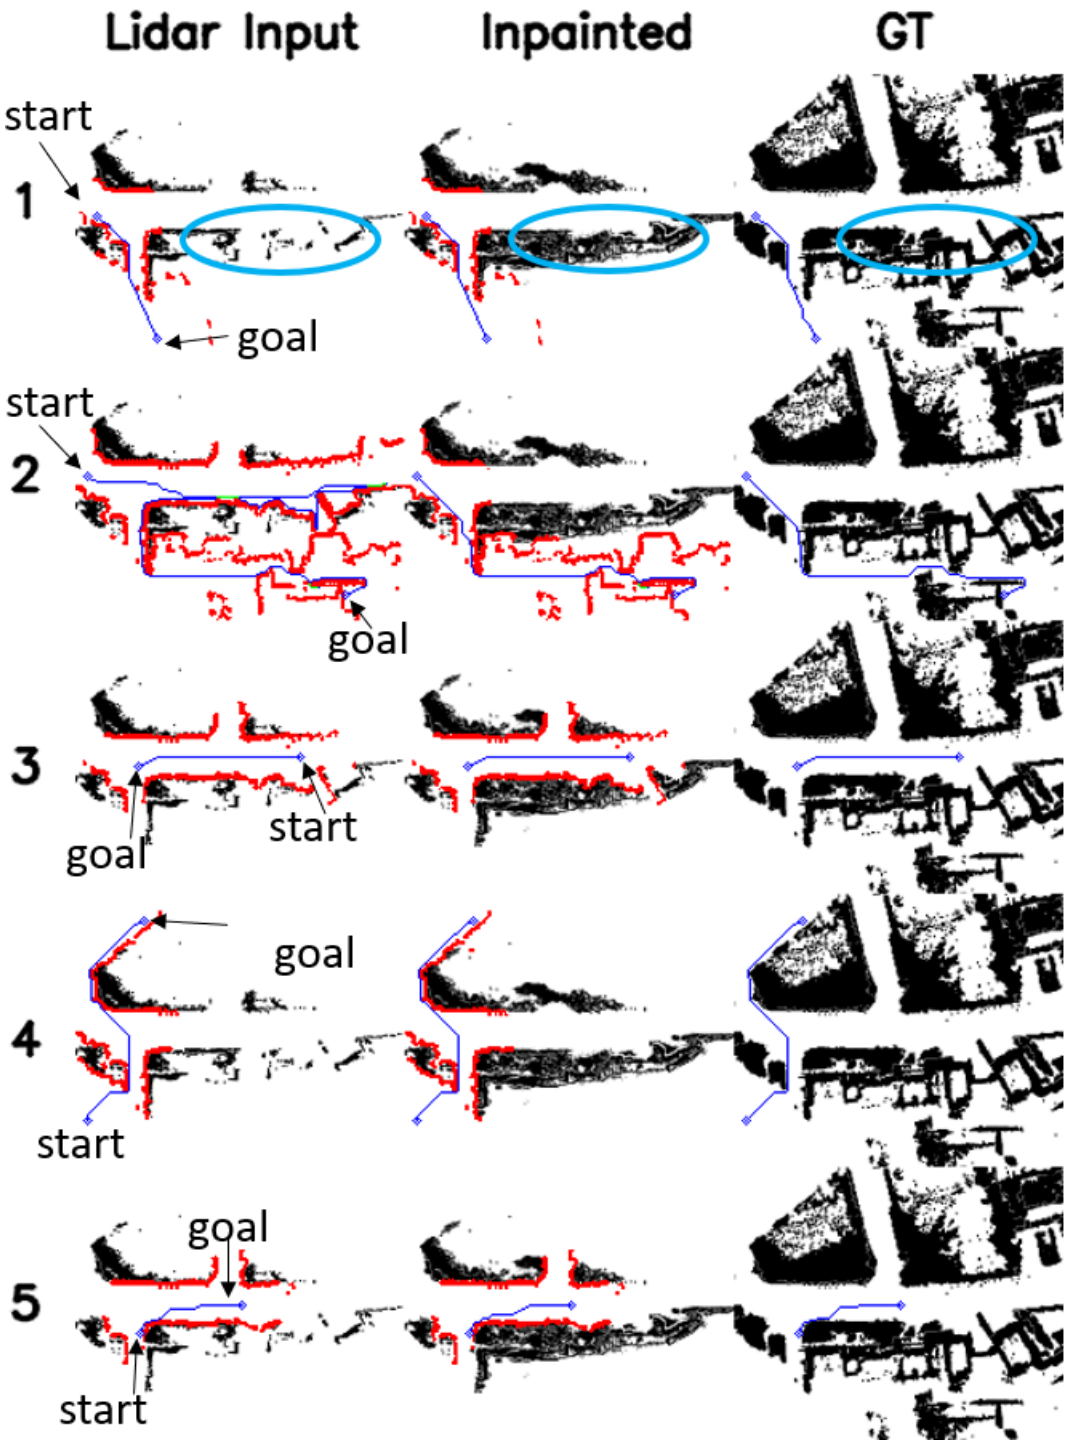

Figure 4: Randomly chosen examples of path planner performance for image 285. Columns from left to right: (1) lidar input, (2) inpainted image, (3) ground truth (GT). The path is in blue pixels. Red pixels are obstacles the robot observes online. Green pixels show where the path overlaps itself.