Supplementary Material for

GapFlyt: Active Vision Based Minimalist Structure-less Gap Detection For Quadrotor Flight

Nitin J. Sanket*, Chahat Deep Singh*, Kanishka Ganguly, Cornelia Fermüller, Yiannis Aloimonos

S.I. Robustness of TS^{2}P against different textures

Optical flow algorithms are generally biased to specific textures, and there is a high correlation between highly textured surfaces and good optical flow computation. To demonstrate the robustness of our approach (TS^{2}P) we test the algorithm against ten additional setups. The various scenarios are tabulated in Table I. Each scenario is a combination of different textures from the following set: Bumpy, Leaves, Door, Newspaper, Wall, Low-Texture and Cloth. We now describe each of textures used.

Bumpy texture provides an uneven texture over the original newspaper scenario. These "bumps" are made of crumpled paper. The depth (protrusion) of the bumps are large and are about 25% of the distance to gap from the quadrotor's initial position. This scenario mimics the uneven texture on rock walls around a cave opening, for example.

Leaves texture mimics foliage. In this setup Magnolia and Panicle Hydrangea leaves are glued onto foam-board. The two leaves used are of very different sizes and shapes. The leaves texture are also uneven with depth variation as large as 10% of the distance to the quadrotor's initial position. We use both sides of the leaves. The front-side of the leaves are of a glossy texture with very high reflectance while the back-side are of matte texture with very low reflectance. Also, the leaves look similar to each other. This texture provides similar repeating patterns and large changes in reflectance.

Door texture is a wall with a door. Both these are white with very low texture.

Newspaper texture is the setup similar to the one used in the main paper. Newspaper is glued onto foam-board. This presents an artificial high-texture scenario.

Wall texture is foam-core with a small amount of logos. We consider this as a medium-texture scenario.

Table SI
Comparison of our approach with different setups*

Scenario	Foreground	Background
1	Bumpy	Leaves
2	Bumpy	Door
3	Bumpy	Newspaper
4	Bumpy	Cloth
5	Leaves	Wall
6	Leaves	Newspaper
7	Cloth	Wall
8	Low-Texture	Leaves
9	Low-Texture	Leaves
10	Low-Texture	Leaves

*Refer to Fig. 1 for a visual aid.
Rows in Fig. 1 correspond to rows in this table.

Low-Texture is white foam-core with a few scribbles near the window which are required for tracking. This is artificially crafted to mimic a minimal-textured plain wall with windows.

Cloth texture is created by the usage of wrinkled bed sheets. This scenario mimics hanging curtains, hanging paintings and hanging flags.

A combination of the aforementioned textures creates scenarios which test the bias of the TS^{2}P algorithm. In all the above cases, ${}^0Z_{\mathcal{F}} \sim 2.8$m and ${}^0Z_{\mathcal{B}} \sim 5.6$m and $N = 3$ frames are used for stacking/averaging in all the cases.

Our approach works in most of the scenarios as presented in Fig. 1 since it uses deep learning based optical flow to estimate the position of the gap in the image plane. Even in the low-textured scenarios, the window detection output $\mathcal{O}$ still has at least 75% overlap with the ground truth as mentioned in our paper. TS^{2}P works even with no textures on one of the foreground or background. Though tracking the $\mathcal{F}$ and $\mathcal{B}$ without any textures is not possible.

*Equal Contribution

All authors are associated with University of Maryland Institute for Advanced Computer Studies, College Park. Emails: {nitin, chahat, kganguly, fer, yiannis} @umiacs.umd.edu

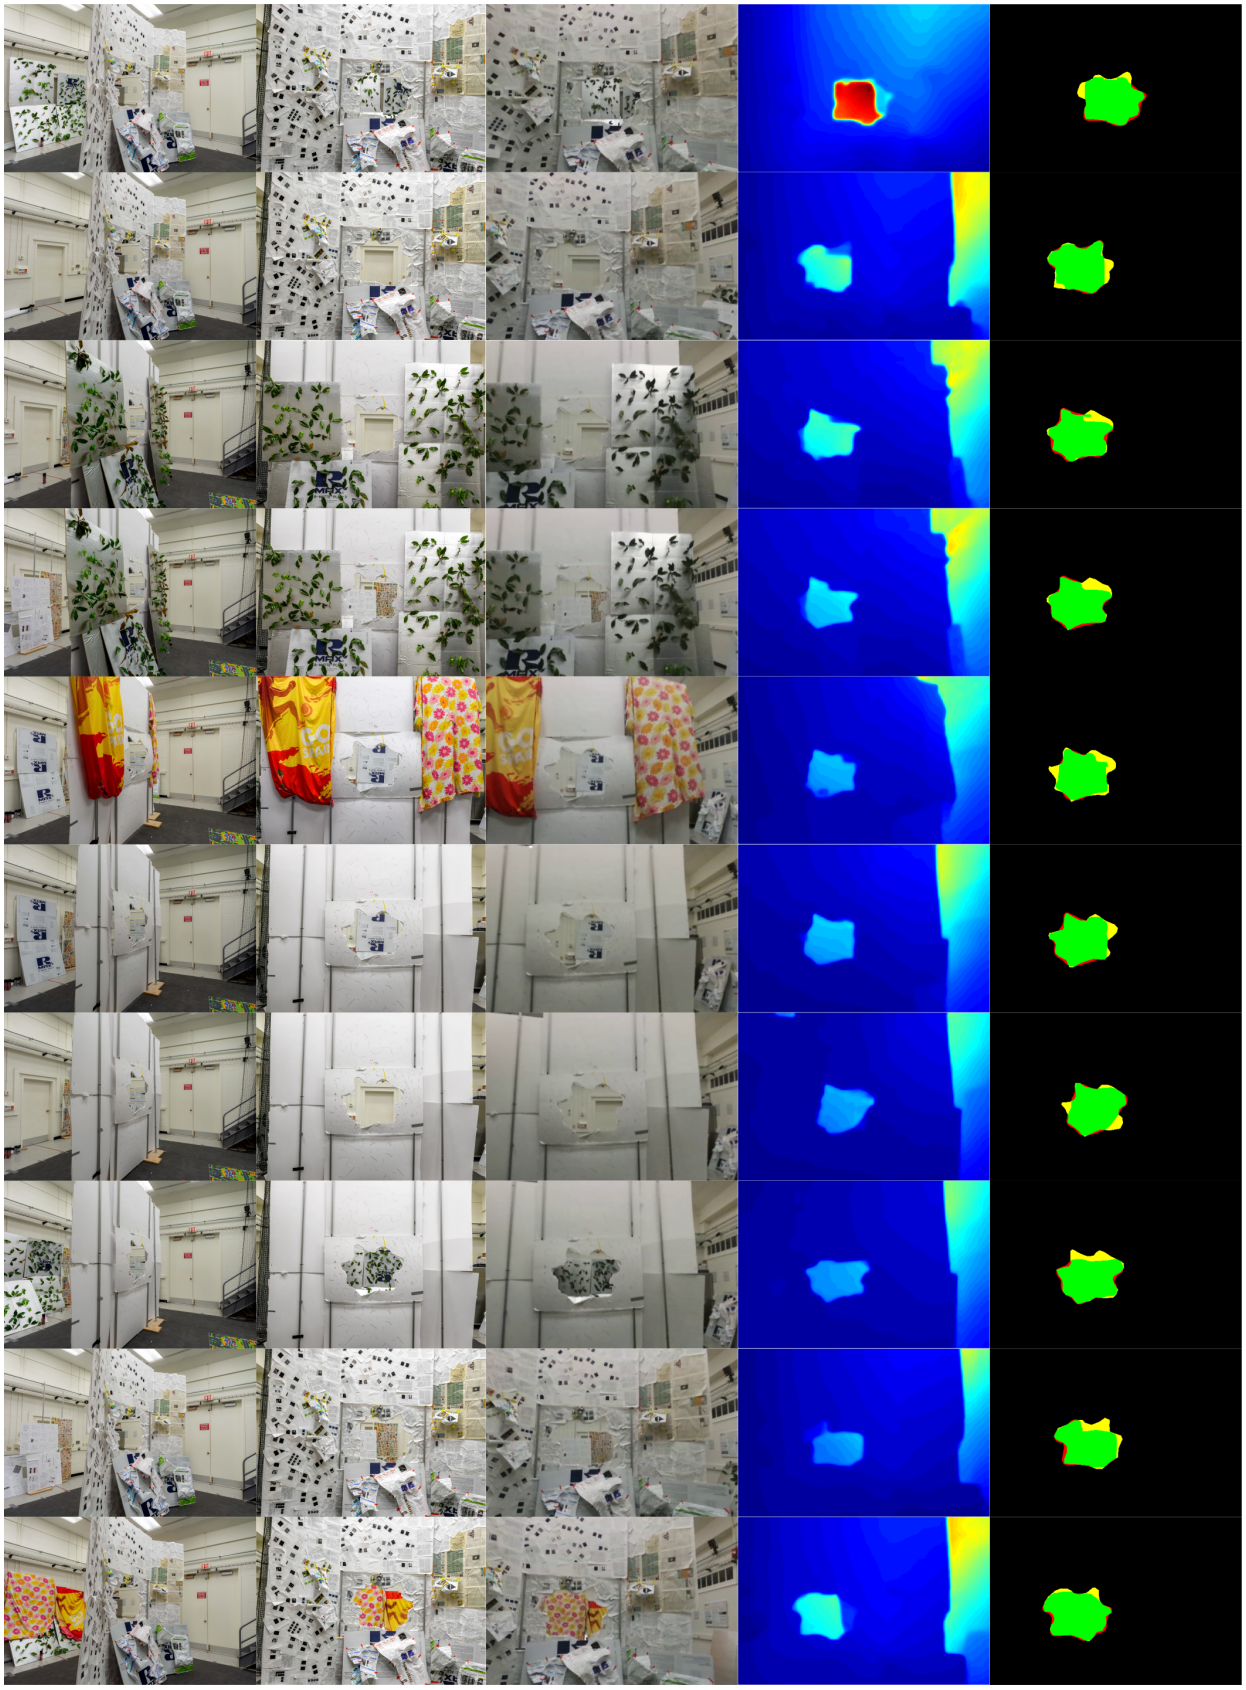

Figure S1. Left to right columnwise: Side view of the setup, Front view of the setup, sample image frame used, Ξ output, Detection output. *(green: $\mathcal{G} \cap \mathcal{O}$, yellow: false negative $\mathcal{G} \cap \mathcal{O}'$, red: false positive $\mathcal{G}' \cap \mathcal{O}$).* Rowwise: Cases are in the order of Table I.