Deep Hough-Transform Line Priors Supplementary material

Yancong Lin, Silvia L. Pintea, and Jan C. van Gemert

Computer Vision Lab
Delft University of Technology, Netherlands

1 Exp 1: Qualitative results on the Line-Circle dataset

Figure 1 visualizes detected lines on the Line-Circle dataset from the local-only, global-only and local+global models. Using the global information learned by our HT-IHT block combined with the local information provided by the convolutional layers, we propose a local+global approach that can predict both the direction of the lines and their extent.

2 Exp 3.(a): Qualitative results using Wireframe subsets

Figure 2 visualizes detected wireframes from our HT-LCNN (9.3M) and LCNN (9.7M) [7] trained on various Wireframe subsets [3]. We display the top 100 line segments. In the first example, our HT-LCNN is better than LCNN in detecting wireframes of windows on various subsets. However, our HT-LCNN is not able to ignore the shadow of objects, compared to LCNN, as shown in the last example. In general, HT-LCNN outperforms LCNN when training data is limited.

3 Exp 3.(b): Qualitative comparison with the state-of-the-art on the Wireframe dataset

Figure 3 visualizes detected line segments from different approaches on the Wireframe dataset [3]. We follow [5] to set up thresholds for LSD [4] and WF-Parser [3], and select the top 100 line segments for other methods (HT-HAWP, HT-LCNN, HAWP[6], LCNN [7], AFM [5], MCMLSD [1] and Linelet [2].) Learning-based models predict line segments more precisely than the non-learning methods. In general, our models with HT-IHT block perform competitively with the state-of-the-art.

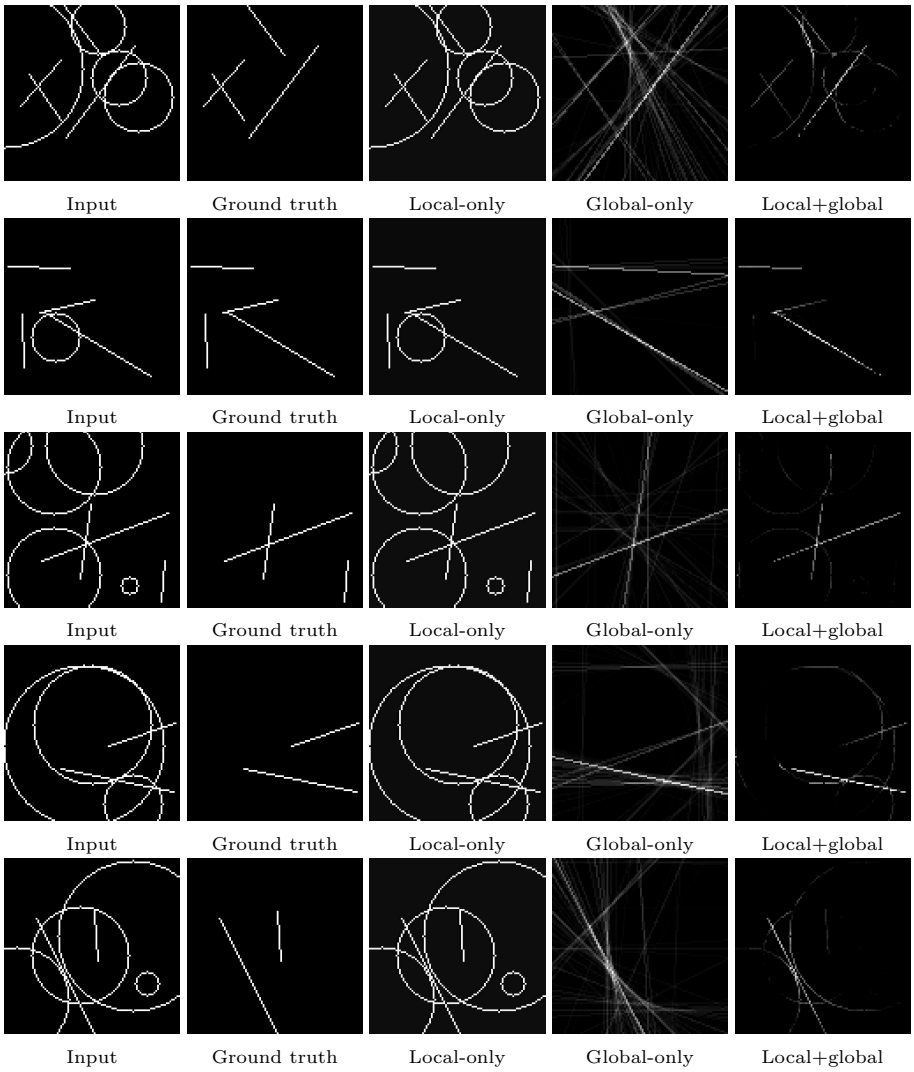

Fig. 1: **Exp 1:** Visualization of detected lines on the toy Line-Circle dataset. The local+global model successfully removed the circle pixels and retains the pixels along the line. Combing local and global information detects not only the direction of the lines but also their extent.

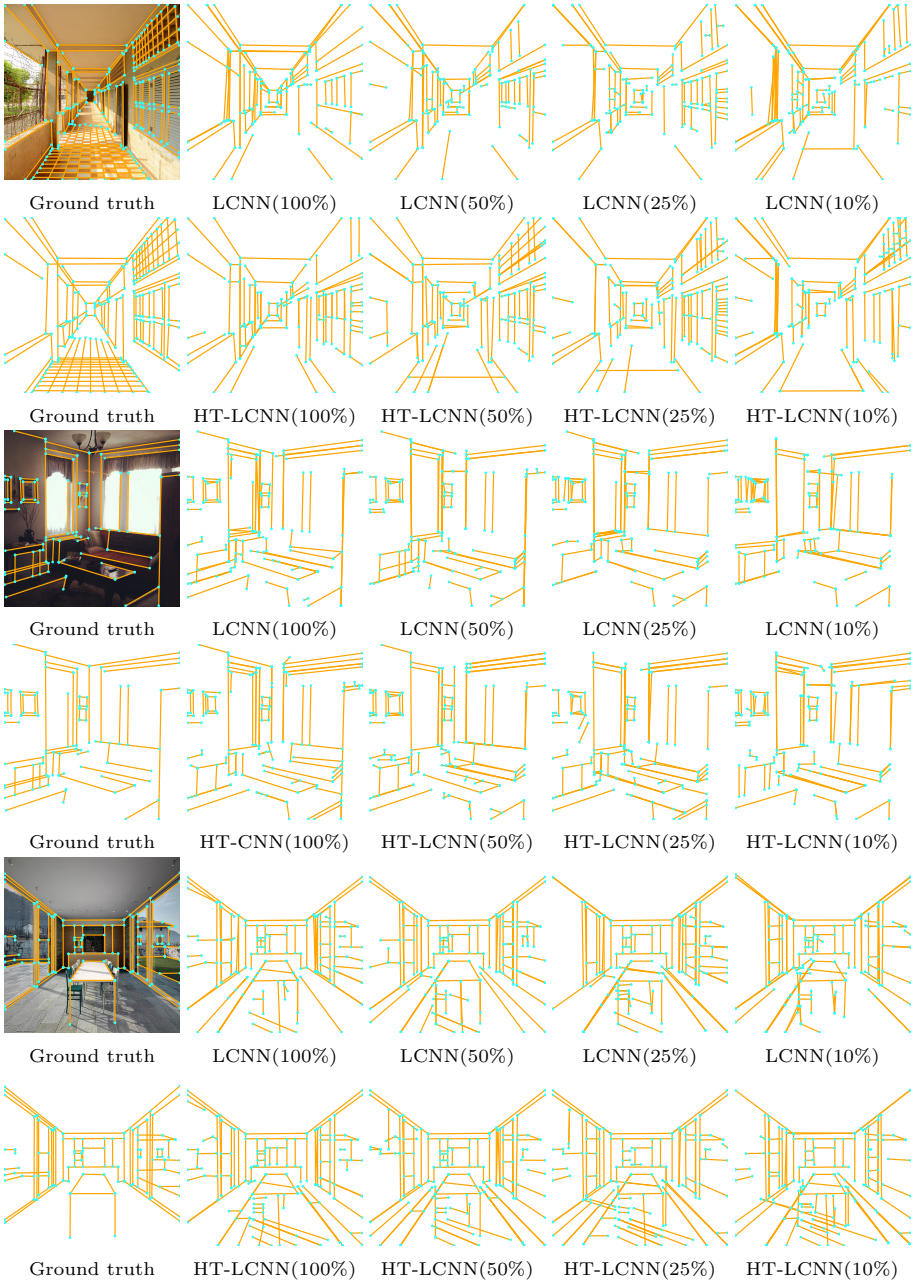

Fig. 2: **Exp 3.(a):** Visualization of detected wireframes from HT-LCNN (9.3M) and LCNN (9.7M) [7] trained on various Wireframe subsets [3]. Our HT-LCNN can more precisely detect the wireframes of the windows than LCNN, as shown in the first example. However, our HT-LCNN generates more false-positive predictions from the shadow of objects, when compared to LCNN, as shown in the last example.

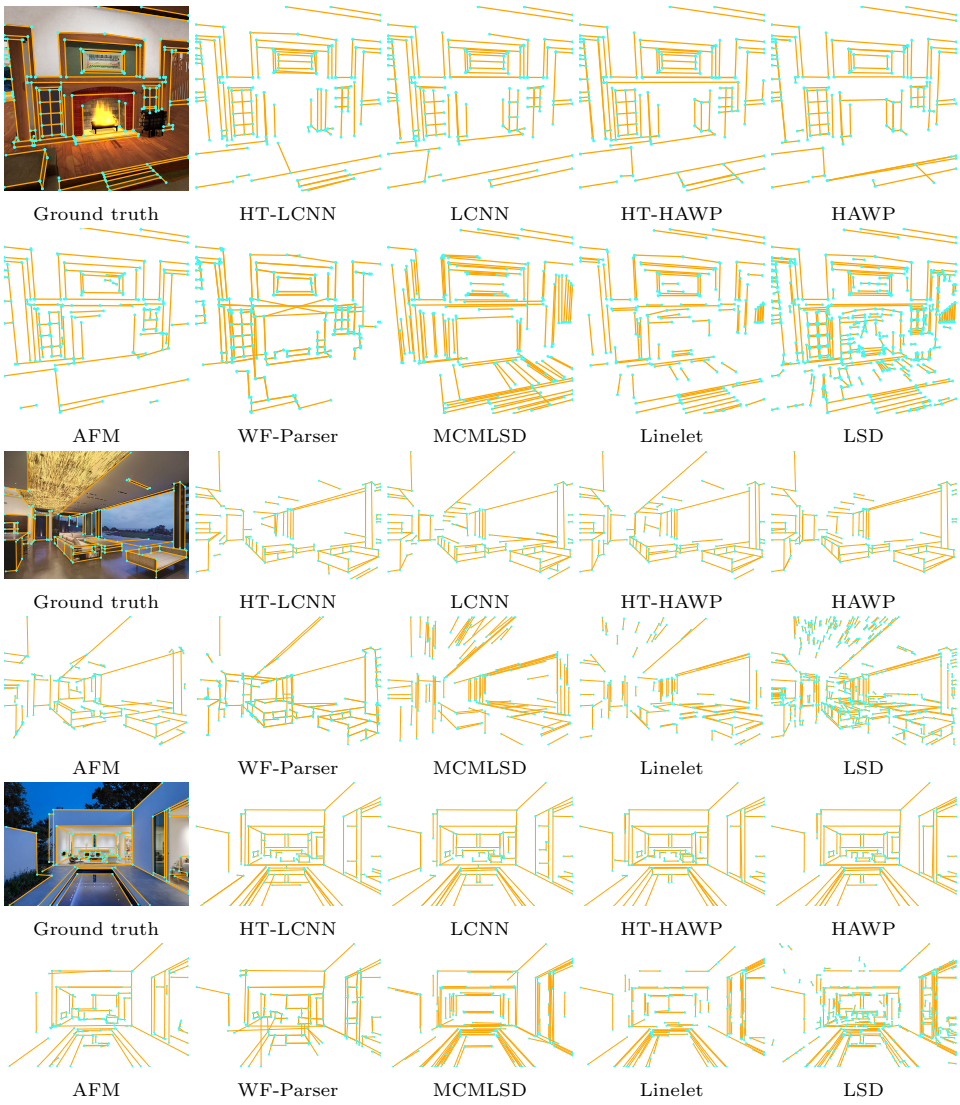

Fig. 3: **Exp 3.(b):** Visualization of detected line segments on the Wireframe dataset [3]. We show predictions from our HT-HAWP, HT-LCNN, and seven other leading methods: HAWP[6], LCNN [7], AFM [5], WF-Parser [3], MCMLSD [1], Linelet [2] and LSD [4]). (Continued on the next page.)

Fig. 3: **Exp 3.(b):** Visualization of detected wireframes (line segments) on the Wireframe dataset [3]. We show predictions from our HT-HAWP, HT-LCNN and seven other leading methods (HAWP[6], LCNN [7], AFM [5], WF-Parser [3], MCMLSD [1], Linelet [2] and LSD [4]). In general, learning-based methods are able to detect line segments more precisely, while MCMLSD, Linelet and LSD generate more false-positive predictions. The HT-LCNN and HT-HAWP predictions preserve both global structures and local details, and show competitive performance with the leading methods.

References

1. Almazan, E.J., Tal, R., Qian, Y., Elder, J.H.: Mcmlsd: A dynamic programming approach to line segment detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2031–2039 (2017)
2. Cho, N.G., Yuille, A., Lee, S.W.: A novel linelet-based representation for line segment detection. IEEE Transactions on Pattern Analysis and Machine Intelligence **40**(5), 1195–1208 (2017)
3. Huang, K., Wang, Y., Zhou, Z., Ding, T., Gao, S., Ma, Y.: Learning to parse wireframes in images of man-made environments. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 626–635 (2018)
4. Von Gioi, R.G., Jakubowicz, J., Morel, J.M., Randall, G.: Lsd: A fast line segment detector with a false detection control. IEEE Transactions on Pattern Analysis and Machine Intelligence **32**(4), 722–732 (2008)
5. Xue, N., Bai, S., Wang, F., Xia, G.S., Wu, T., Zhang, L.: Learning attraction field representation for robust line segment detection. In: The IEEE Conference on Computer Vision and Pattern Recognition (June 2019)
6. Xue, N., Wu, T., Bai, S., Wang, F., Xia, G.S., Zhang, L., Torr, P.H.: Holistically-attracted wireframe parsing. In: Conference on Computer Vision and Pattern Recognition (2020)
7. Zhou, Y., Qi, H., Ma, Y.: End-to-end wireframe parsing. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 962–971 (2019)