FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics

Diplomarbeitspräsentation

Masterstudium:
Computergraphik/Digitale Bildverarbeitung

Extracting Vanishing Points across Multiple Views

Michael Hornáček, BA

Technische Universität Wien
Institut für Computergraphik und Algorithmen
Arbeitsbereich: Computergraphik
Betreuer: Univ.-Prof. Dipl.-Ing. Dr. Werner Purgathofer
Mitwirkung: Dipl.-Ing. Dr. Robert F. Tobler (VRVis)

Motivation and Objective

The realization that we see lines we know to be parallel in space as lines that appear to converge in a corresponding vanishing point has led to techniques employed by artists to render a credible impression of perspective. More recently, it has also led to techniques for recovering information embedded in images—or, indeed, in paintings that feature correct perspective—concerning the geometry of their underlying scene.

In this master's thesis, we explore the extraction of vanishing points in the aim of facilitating the reconstruction of typical urban scenes. In departure from most vanishing point extraction methods, ours extracts a constellation of vanishing points across multiple views rather than in a single image alone. By making use of a strong orthogonality criterion per view, optimal segment intersection estimation and a novel tripod fitting technique, our approach allows for the extraction of results that correspond closely to the dominant three pairwise-orthogonal line orientations of a typical urban scene. These three orientations serve at the same time as normal vectors corresponding closely to the scene's dominant three pairwise-orthogonal plane orientations. Accordingly, ours can fairly be described as a material refinement of the approach proposed in Sinha et al. [1] as part of their reconstruction pipeline.

The Geometry of Vanishing Points

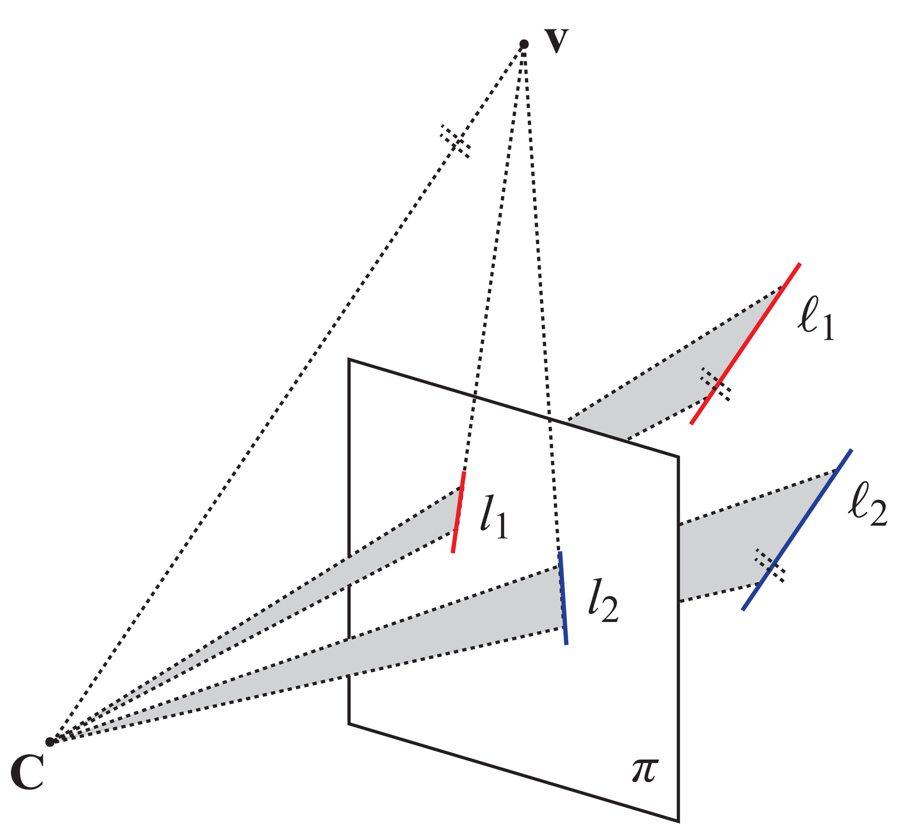

Figure: The projections $l_1, l_2 \subset \pi$ of two lines ℓ_1, ℓ_2 in space converge in a corresponding vanishing point $\mathbf{v}$ in the image plane π. Note that under known camera geometry, the lines ℓ_1, ℓ_2 in space thus have the same orientation as the ray extending from the camera center $\mathbf{C}$ through $\mathbf{v}$. We call that ray the back-projection of $\mathbf{v}$ with respect to the given camera.

Implementation

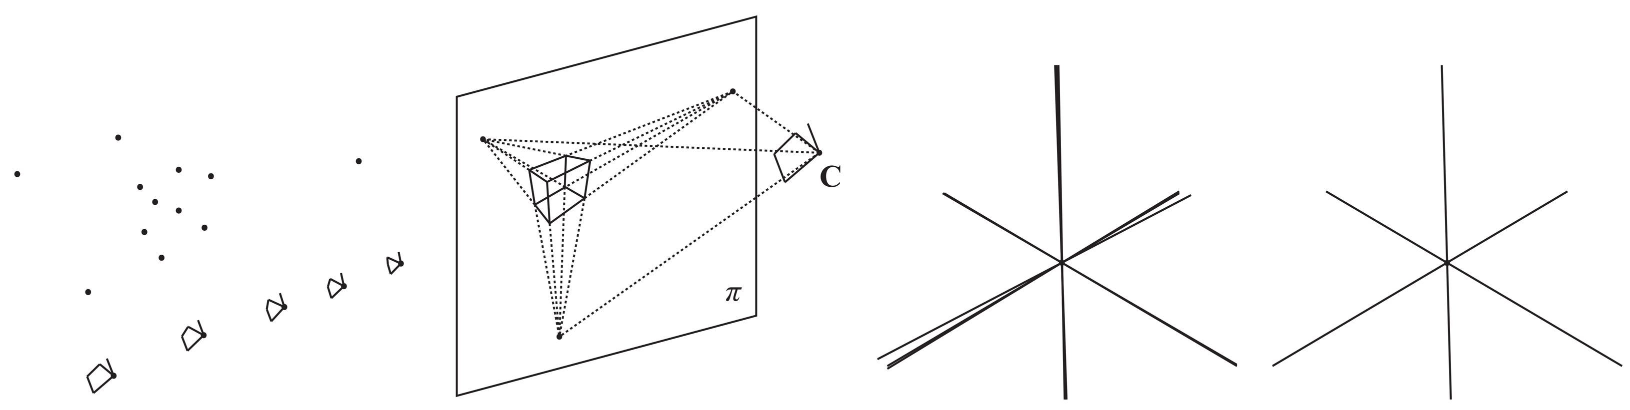

(a) Recovery of camera geometry for the k available views.

(b) Extraction of a constellation of two or three candidate vanishing points in a single view, constrained to correspond closely to pairwise-orthogonal scene orientations.

(c) Antipodal unit direction vectors corresponding to the orientations computed from candidate vanishing points extracted across the k views.

(d) The pairwise-orthogonal orientations corresponding to the best-fit tripod.

Figure: The processing pipeline of our approach.

Results

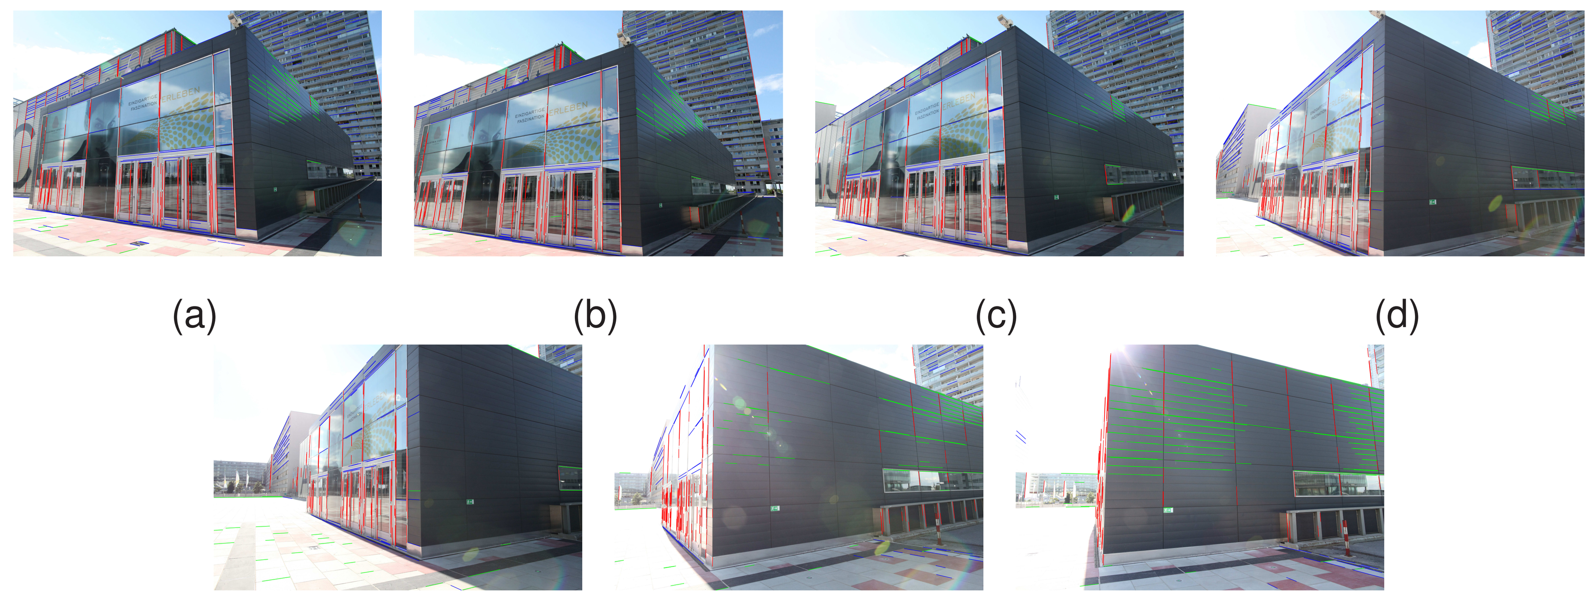

(a) (b) (c) (d)

(e) (f) (g)

Figure: The acv data set with an approximation of its dominant three pairwise-orthogonal scene orientations extracted via our approach, with the inlier segments of their respective corresponding vanishing points shown per view.

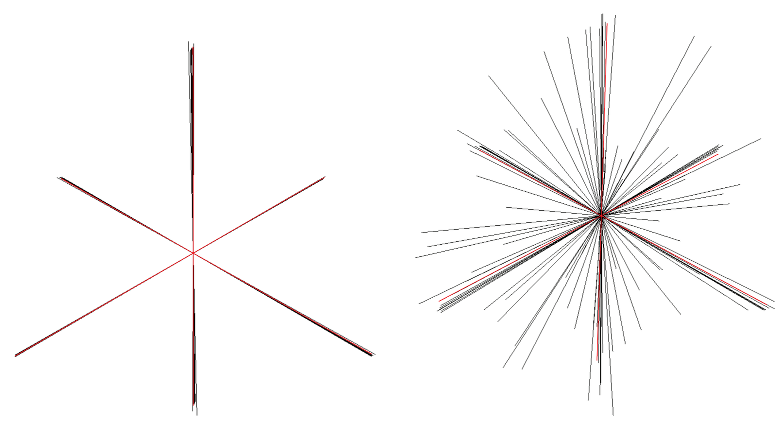

(a) The best-fit tripod w.r.t. the antipodal directions obtained per view via our approach.

(b) The best-fit tripod w.r.t. the antipodal directions obtained per view via the approach of Sinha et al.

Figure: Antipodal unit direction vectors extracted across all views of the given data set, with the corresponding best-fit tripod indicated in red. The tripod for our approach corresponds to the results given in the figure immediately above.

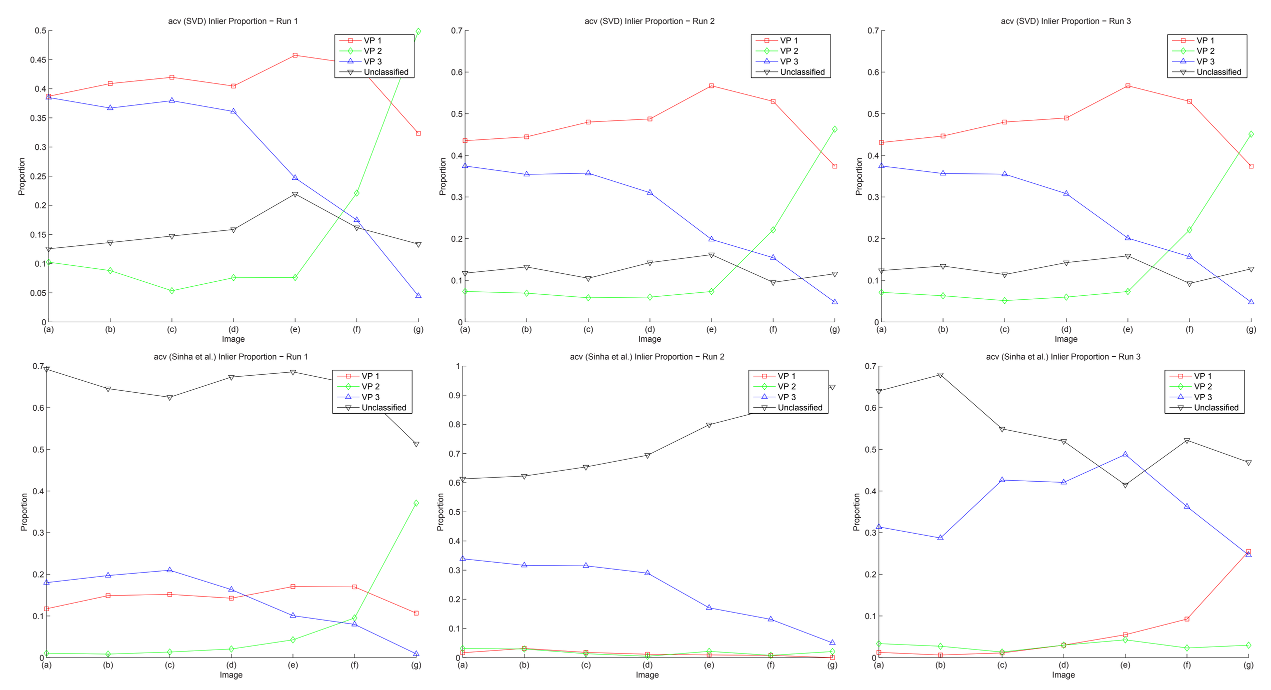

Figure: Inlier proportions for the acv data set across three runs. The top row corresponds to the results obtained using our approach; the bottom, to our tripod fitting with respect to the antipodal directions obtained via the approach of Sinha et al. In both cases, run 1 refers to selfsame respective run that gave rise to the corresponding tripod in the figure immediately above.

Conclusion

Our approach presents a material refinement of the multiple-view vanishing point extraction technique proposed in Sinha et al. [1]. We found in our experiments—in which we considered two data sets in addition to the acv data set—that our method consistently outperformed the fundamental approach of Sinha et al. Our method yielded results that were comparatively more stable across runs and that in each case corresponded closely to the respective dominant three pairwise-orthogonal orientations of each of the three urban scenes considered.

References

[1] S. N. Sinha, D. Steedly, R. Szeliski, M. Agrawala, and M. Pollefeys. Interactive 3D Architectural Modeling from Unordered Photo Collections. ACM Transactions on Graphics, 27(5):1, December 2008.

Kontakt: hornacek@vrvis.at / mho@ipf.tuwien.ac.at