Detecting Vanishing Points using Global Image Context in a Non-Manhattan World

Menghua Zhai, Scott Workman, Nathan Jacobs
{ted, scott, jacobs}@cs.uky.edu
http://cs.uky.edu/~ted

UNIVERSITY OF KENTUCKY

IEEE 2016 Conference on Computer Vision and Pattern Recognition
CVPR 2016

Goal

Goal: Given a single image, detect the horizon line and vanishing points without making a Manhattan-world assumption.

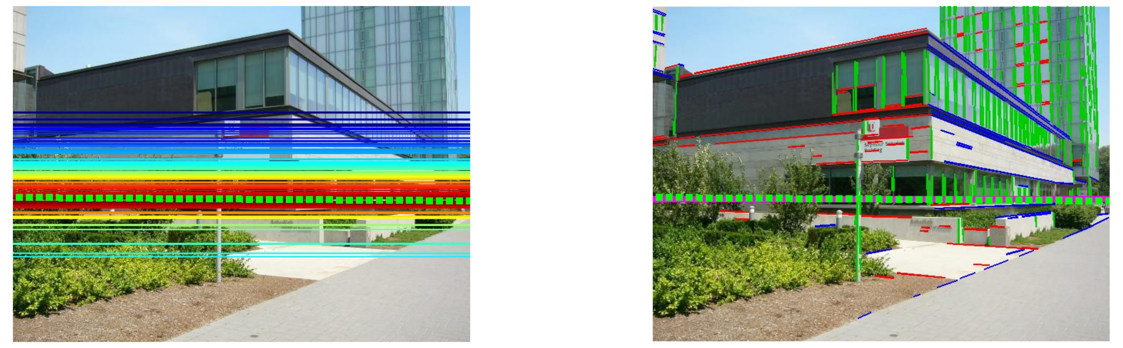

Algorithm Overview

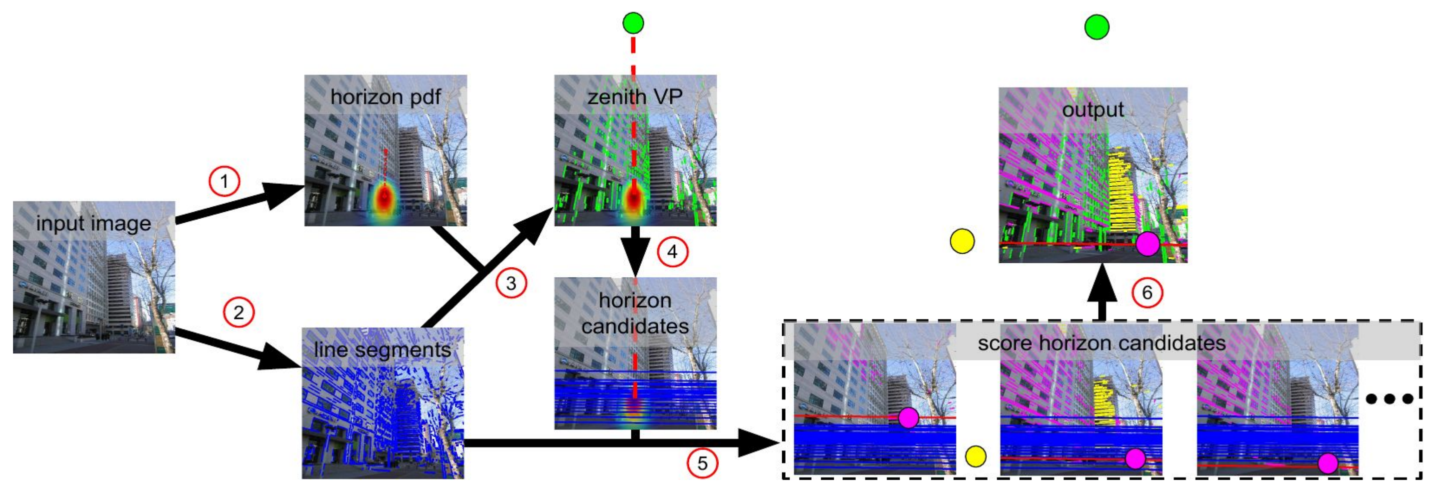

1. Extract global image context (horizon-line prior) using a CNN,
2. Detect line segments using LSD [R. G. Von Gioi et al. 2010],
3. Localize zenith, using RANSAC to fuse prior and segments,
4. Sample horizon-line candidates using prior and the zenith direction,
5. Find vanishing points with many consistent line segments for each horizon-line candidate,
6. Return the horizon line with highest score (i.e., many unique, high quality vanishing points).

Consistency Measurement

In homogeneous coordinates, lines and points are represented as three dimensional vectors. We define our consistency measurement between a line, l, and a vanishing point, p, as follows:

$$f_c(l, p) = \max(\theta - \Theta_{l,p}, 0),$$

where $\Theta_{l,p} = |\cos^{-1}(l^T p)|$, and θ is some threshold.

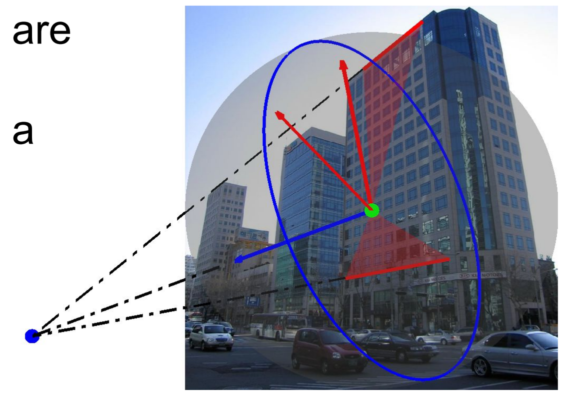

Extract Global Image Context

Definition of Global Image Context: A probability distribution over all possible horizon-line locations on the image.

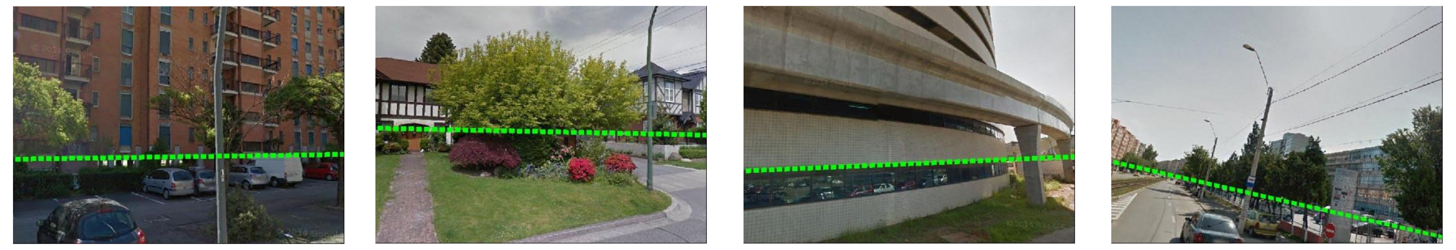

Annotated training examples.

Training Data:
Perspective images "cutout" from Google Street View Panoramas. The horizon line location is computed using the panorama metadata. For training, we extract 10 cutouts from each of the 11,001 panoramas.

Horizon-Line Parameterization:
We use the slope/offset representation and uniformly discretize slope and offset (in homogeneous space) into 500 bins. We assume slope and offset are independent, which results in two 500-class classification problems.

CNN Training:
Replace the last layer of AlexNet with two fully-connected layers, one for slope and one for offset. We initialize our network with weights pre-trained for image classification on the ImageNet dataset.

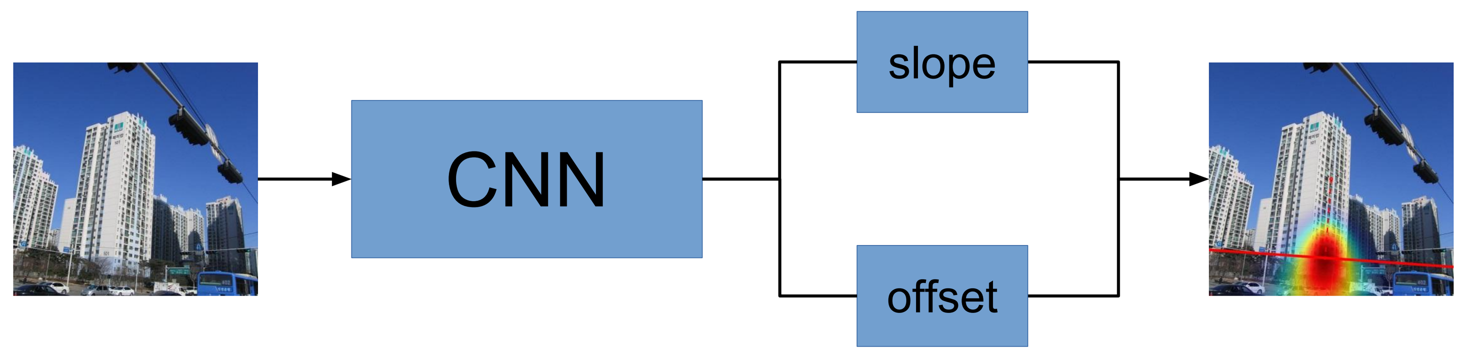

Find Zenith Vanishing Point

1. Filter out line segments unrelated to the zenith vanishing point using the horizon-line slope predicted from the CNN.
2. Find the zenith vanishing point by maximizing the total consistency with the filtered line segments (use RANSAC for robust fitting).

Sample Horizon Line Candidates

- Fix the horizon line slope with the zenith vanishing point found from the previous step,
- Sample horizon line candidates using the probability of the horizon line offset.

Horizon-line candidates sampled from the offset prior (right) are more dense around the ground truth than uniform sampling (left).

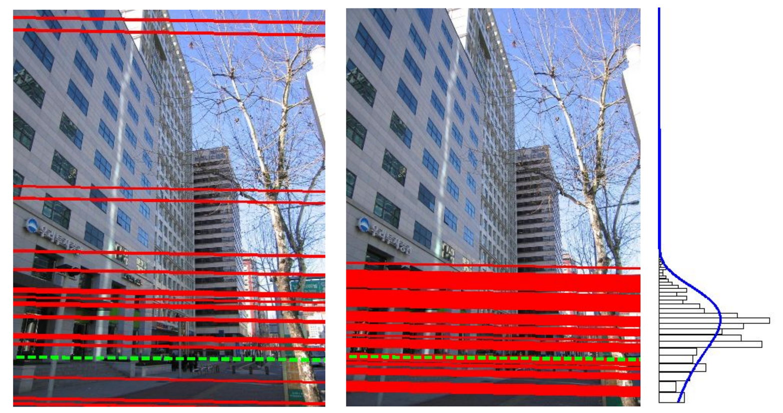

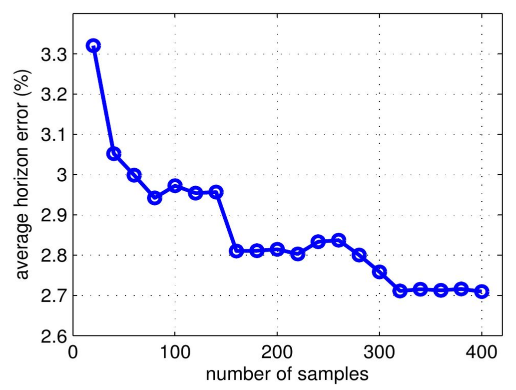

Sampling more horizon line candidates leads to lower horizon line prediction error, but it also takes longer to compute. In our experiments, we sample 300 horizon-line candidates.

Score Horizon-Line Candidates

Intuition: A good horizon line will have multiple distinct vanishing points, each with many consistent line segments.

Challenge: Many potential, non-distinct vanishing points on each horizon-line candidate.

Solution: We propose a discrete-continuous optimization process:

- Find an initial subset of candidate vanishing points (reduces to maximum weighted independent set with a ring-like structure). Solve using dynamic programming.
- Locally optimize the selected vanishing points along the horizon-line candidate using an EM-like approach.

$$g(\mathcal{P}|\mathbf{h}, \mathcal{L}) = -\sum_{\mathbf{p}_i \in \mathcal{P}} \sum_{\mathbf{l}_j \in \mathcal{L}} f_c(\mathbf{p}_i, \mathbf{l}_j)$$

subject to:

$$\Theta_{\mathbf{p}_i, \mathbf{p}_j} > \theta_{dist} \text{ and } \langle \mathbf{p}_i, \mathbf{h} \rangle = 0, \ \forall (i, j) \ .$$

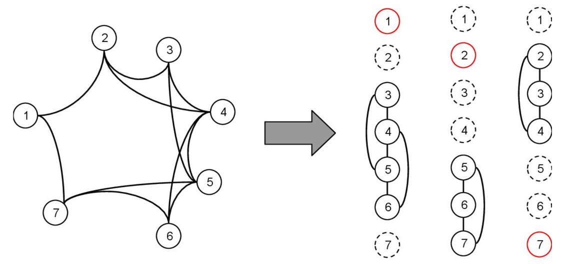

Results

Our method outperforms state-of-the-art methods on two popular benchmark datasets, and a new, more challenging dataset with ~2000 test images.
All our results were obtained using the **same set of parameters**!

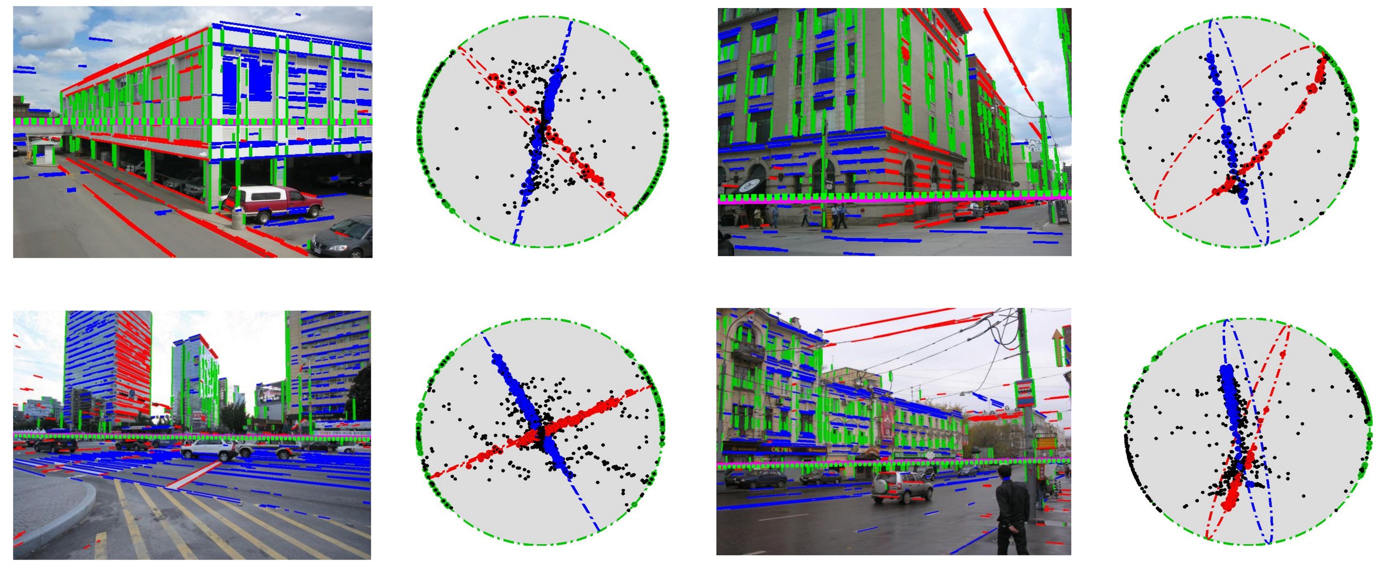

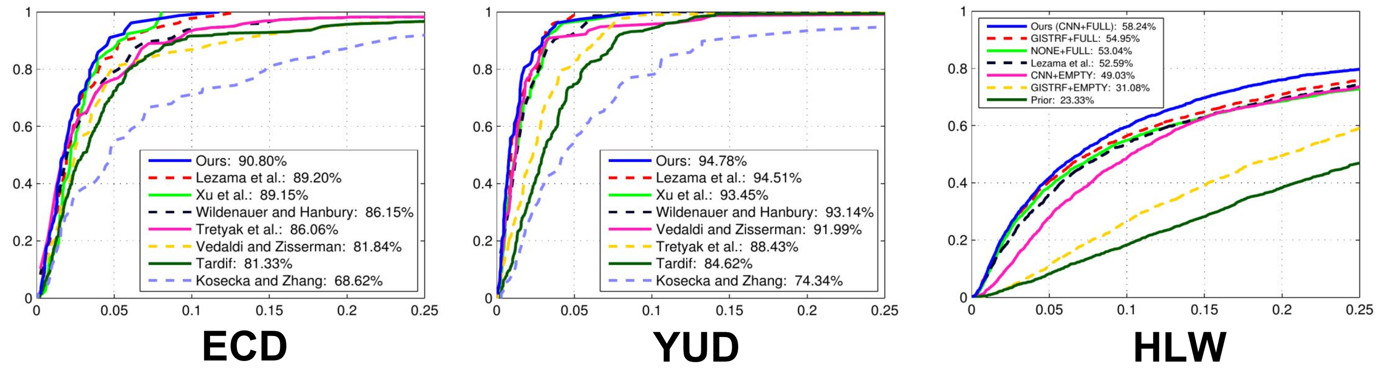

Conclusion

Main innovations:
- Use of global image context to sample possible horizon lines,
- A novel horizon-first framework for vanishing point detection,
- A novel discrete-continuous procedure to score each horizon line by choosing the optimal vanishing points for the line.

Our method is both more accurate and more efficient than the previous state-of-the-art algorithms.

Acknowledgement: We gratefully acknowledge the support of DARPA (contract CSSG D11AP00255). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.