

Recurrence-based Vanishing Point Detection

Skanda Bharadwaj Robert T. Collins Yanxi Liu
School of EECS, The Pennsylvania State University, University Park, PA
skanda.bharadwaj@psu.edu, {rcollins, yanxi}@cse.psu.edu

Abstract

Classical approaches to Vanishing Point Detection (VPD) rely solely on the presence of explicit straight lines in images, while recent supervised deep learning approaches need labeled datasets for training. We propose an alternative unsupervised approach: Recurrence-based Vanishing Point Detection (R-VPD) that uses implicit lines discovered from recurring correspondences in addition to explicit lines. Furthermore, we contribute two Recurring-Pattern-for-Vanishing-Point (RPVP) datasets: 1) a Synthetic Image dataset with 3,200 ground truth vanishing points and camera parameters, and 2) a Real-World Image dataset with 1,400 human annotated vanishing points. We compare our method with two classical methods and two state-of-the-art deep learning-based VPD methods. We demonstrate that our unsupervised approach outperforms all the methods on the synthetic images dataset, outperforms the classical methods, and is on par with the supervised learning approaches on real-world images. Code and data can be found here: <http://vision.cse.psu.edu/data/data.shtml>

1. Introduction

Under perspective projection, parallel lines in 3D meet at a point called the *vanishing point* (VP) [32] which may lie inside or outside of the 2D image (Figure 1). Vanishing Point Detection (VPD) plays a vital role in computer vision applications such as structure from motion [16, 23], 3D reconstruction [17, 41, 47], camera calibration [6, 9, 11, 17], scene understanding [14], SLAM [21, 22, 45], facade detection [27] and autonomous driving [20]. VPD algorithms can typically be separated into two categories: classical methods that are based on projective geometry [10, 24, 30, 39, 48] and deep learning methods [5, 7, 25, 28, 46]. Classical VPD methods usually consist of two stages - 1) Line Segment Detection (LSD) [13] and 2) VP estimation using line clustering [3, 31] or voting [4, 15]. Deep learning VPD methods are data-driven approaches using filters learned by training on hundreds of thousands of annotated images [18, 19].

We propose an alternative approach to detect VPs using




Figure 1. Two images containing examples of “things that recur” (Recurring Patterns). Top row: images with ground truth VP (x) indicated. Bottom row: VP prediction results from four VPD methods R-VPD (our method), NeurVPS [46], GPVPD [25] and J-Linkage [42] respectively, where VP predictions that are too far from the ground truth are not shown in the images ((a) 2 misses, (b) 3 misses). Due to a lack of explicit straight lines, automatic vanishing point detection on these images poses challenges to explicit line-based methods.

implicit lines discovered from Recurring Patterns (RPs), defined as “things that recur” [26], which have common and frequent appearances in real-world images (Figure 1). Different from existing methods on VPD, we use groups of corresponding feature points, discovered during the RP detection process, to construct and validate the co-linearity of implicit lines. We hypothesize intersections of these virtual lines in addition to explicit lines and find the VP through a robust, joint, weighted RANSAC procedure. An overview of this approach is shown in Figure 2. Our Recurrence-based VP method (R-VPD) provides additional robustness for VPD on images without explicit lines leading to VPs. Our contributions include:

1. We propose a novel methodology, **Recurrence-based Vanishing Point Detection (R-VPD)**, to detect vanishing points in an unsupervised manner using the implicit lines obtained from feature correspondences of RPs in addition to explicit image lines.

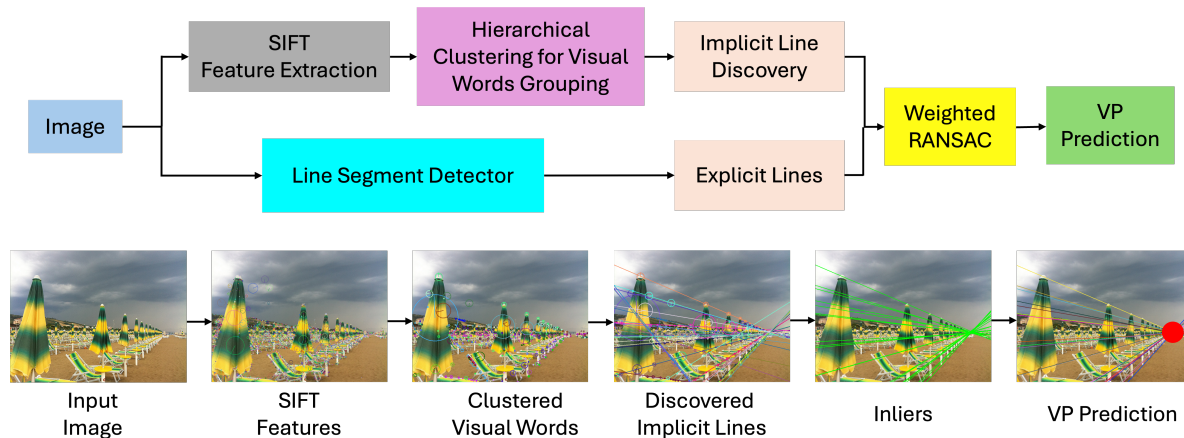


Figure 2. Recurrence-based Vanishing Point Detection: An overview of R-VPD is presented in the above image. SIFT features are extracted from the input image and are clustered hierarchically. Forward selection of the features is performed using geometric constraints based on linearity, angle, and scale. Implicit lines are fitted to the selected feature clusters, and a weighted RANSAC algorithm is employed to find the intersection of both implicit and explicit lines to locate the vanishing point. For this image, there are no explicit long lines.

2. We provide two benchmark datasets - (a) **Synthetic** images (RPVP-Synthetic) comprised of 3,200 images created in Blender with the ground truth VP and camera parameters (intrinsic and extrinsic matrices) used to create the 2D images, and (b) **Real**-world images (RPVP-Real) with 1,400 annotated real-world images containing RPs, where the ground truth VP is hand-labeled using a *LabelMe* tool [38].
3. We establish the first VPD benchmark on RP-datasets by comparing four VPD algorithms including two classical approaches: a simple baseline Line Segment Detection (LSD) and J-Linkage [42] and two state-of-the-art deep learning approaches: NeurVPS [46] and GPVPD [25].

2. Related Work

Broadly speaking, VPD methods can be separated into two categories: 1) Classical methods; and 2) Deep learning methods. Early classical methods detect VPs as intersections of straight line segments, typically using the Hough transform [29, 37] followed by line clustering [31] or other voting schemes [15]. Barnard [2] introduced a Gaussian sphere representation where lines in the image map to great circles on a sphere, a bounded space more representative of the projective plane than the unbounded image plane and thus better able to support the detection of vanishing points formed by parallel image lines. Motivated by this representation, Collins and Weiss [10] formulated vanishing point calculation as a statistical estimation problem on the unit sphere, estimating a vanishing point location and associated confidence region as the polar axis of an equatorial distribution on the sphere. Schaffalitzky *et al.* [39] presented a proof-of-concept planar grouping approach that exploits

specific geometric configurations, such as equally spaced coplanar lines, to estimate vanishing points and lines. Zhou *et al.* [48] addressed the problem of understanding linear perspective in landscape photography, detecting dominant vanishing points by exploiting global structures in the scene via contour detection. One of the most recent VP estimation methods in the “classical” vein (Li *et al.* [24]) makes a Manhattan world assumption to formulate VP estimation as computing the rotation between the Manhattan frame and the camera frame.

Another category of work leverages groupings of coplanar interest points [8, 34–36, 39] to identify a vanishing line from which the planar surface can be rectified. In particular, Pritts *et al.* [34–36] leverage differing sizes of conjugately translated planar elements to estimate the vanishing line, allowing affine rectification, and with pattern-dependent constraints sometimes improving the rectification to a similarity transformation. Unlike most other works in projective geometry, they are able to estimate and undo the nonlinear warping effects of radial lens distortion.

With recent advancements in deep learning, most research focuses on solving problems using learnable filters [7, 25, 46]. NeurVPS [46] introduced an efficient VP estimation method using a conic convolution operator that extracts and aggregates features along structural lines. Building on NeurVPS and the early work of Barnard [2], Lin *et al.* [25] propose two trainable geometric priors, the Hough transform and Gaussian sphere representation, to reduce network sensitivity to dataset variations. Their method first extracts features with a stacked hourglass network [33], maps them to a Hough space to identify lines, and then uses spherical convolution on the Gaussian sphere to detect VPs. A recent work [43] by Tong *et al.* does not have publicly available code, making direct comparison infeasible.

Most classical and deep-learning VPD methods depend on explicit line segments in the image. Even the state-of-the-art [46] relies on a sufficient number of edges oriented towards the VP. To address VP estimation in images lacking explicit line segments, we propose a novel method that leverages recurring patterns and their properties, which offer strong cues about the 3D world from a 2D image.

3. R-VPD: Recurrence-based Vanishing Point Detection

Recurring Patterns (RPs) are a set of “things that recur” [26]. RPs embody the powerful organizing principle that things that co-occur are not accidental.

A *unit RP* (URP), namely the smallest RP (Figure 3) is composed of two *RP Instances* (RPI), and an RPI has to contain at least one pair of distinct visual words. We use the SIFT feature extractor to generate the pool of candidate features for forming visual words in an image. The initial phase of SIFT involves constructing a difference of Gaussian (DoG) image pyramid. Feature points are identified by searching for local extrema across different scales and spatial positions. Once features are localized, a 128-dimensional custom vector is calculated as the feature descriptor. The 64x64 regions surrounding the key points are subdivided into 16 smaller blocks of size 4x4, and an eight-bin orientation histogram is created, resulting in the final generation of a 128-dimensional vector.

Starting with a URP, a 2x2 matrix is constructed with each row being a visual word and each column being an RP instance. An unsupervised, randomized optimization algorithm called GRASP [26] performs a set of random “moves” that progressively grow, shrink, or replace entries in the matrix, seeking to maximize its size while being constrained by a set of numeric consistency measures. For more information on optimization-based RP detection, we refer readers to [26]. The authors of [44] develop an alternative two-stage unsupervised architecture for RP discovery instead of GRASP, yielding enhanced speed.

3.1. Implicit Recurring Correspondences Discovery

Different from both [26, 44], we do not explicitly generate RPs from a given image but employ in novel ways the corresponding visual words discovered during the optimization process to fit and merge implicit lines with explicit lines for subsequent VP detection. By fitting lines through two or more such correspondences and finding their intersections we can estimate a potential VP. Hence we give the name of our approach *Recurrence-based VPD*.

The basic difference between our method and existing VPD algorithms is the use of implicit lines, such that our method does not rely only on explicit edge in an image. Traditional VP detection methods fail in scenes where there

are not enough salient line segments in the image (Figure 1), while our method takes advantage of implicit lines together with any explicit lines to locate the VP.




Figure 3. A Unit Recurring Pattern (URP) is an RP with at least two RP instances (RPIs) and each RPI has at least two distinct visual words (e.g. clustered SIFT descriptors) that correspond across the RPIs (this figure is adapted from [44]). Here we add the two blue dotted lines passing through pairs of corresponding features, (f_{11}, f_{12}) and (f_{21}, f_{22}) , whose intersections identify the VP.

3.2. Hierarchical clustering

We use hierarchical clustering, a traditional unsupervised method that groups high-dimensional features into a clustering tree, to further group visual words for line fitting. Unlike other methods, it does not require a preset number of classes, making it ideal for this scenario. A feature-matching matrix is first created using Euclidean distances between features. Hierarchical clustering is then performed by iteratively merging the closest subgroups of features until all features form a single group.

Specifically, let S be the set of feature groups and let $A, B \in S$. Let N_A and N_B be the number of features within the groups A and B , respectively. The distance between the groups A and B is defined as:

$$Dist_{AB} = \min\{dist_{ij}\}, \forall i \in [1, N_A], \forall j \in [1, N_B] \quad (1)$$

$$dist_{ij} = \|fd_i - fd_j\|_2$$

where $dist_{ij}$ represents the distance between features i and j and fd_k represents the 128-dimensional descriptor of feature k . Then, merging the closest neighboring groups is performed simply as $C = (A \cup B)$, and distances are updated between this new group and existing groups using the equation:

$$Dist_{CP} = \min\{Dist_{AP}, Dist_{BP}\}, \forall P \in \{S \setminus \{A, B\}\}. \quad (2)$$

3.3. Forward Feature Selection using Bottom-up Search

Forward feature selection identifies features within a visual word that are best suited for fitting implicit lines. This is achieved using three scores—Linearity, Angle, and Scale—that assess distinct aspects of geometric consistency:




Figure 4. Sample images from RPVP-Synthetic image dataset showing scenes in which RPs are oriented towards the VP. A relatively lesser number of explicit lines makes this dataset challenging for VPD. The "x" represents the ground truth VP.

Linearity Score (S_L): This score measures how well the (x, y) locations of keypoints align along a dominant orientation. It is calculated as the average perpendicular distance of keypoints from a fitted line. A lower score indicates stronger alignment, making the features more suitable for line fitting.

Angle Score (S_A): This evaluates the consistency of directional changes across an ordered sequence of keypoints. For triplets (A, B, C) , it calculates the absolute difference in angular changes between segments $A \rightarrow B$ and $B \rightarrow C$. These differences are averaged across all triplets, with lower scores reflecting greater uniformity in direction.

Scale Score (S_S): This assesses the consistency of relative size changes between keypoints, influenced by projective transformations. For triplets (A, B, C) , it calculates the absolute difference in scale ratios between $A \rightarrow B$ and $B \rightarrow C$. These differences are averaged, with lower scores indicating more uniform scale progression.

Composite Score (S_C): The overall suitability of a feature set is determined by: $S_C = S_L \times \exp(S_A + S_S)/N^2$, where N is the number of features in the visual word. This score integrates collinearity, angular consistency, and scale variability while normalizing for feature set size, ensuring that only geometrically coherent features are selected.

By combining these metrics, forward feature selection prioritizes features that exhibit strong geometric consistency, enhancing the reliability of vanishing point estimation. A more elaborate explanation of the scores can be found in the Appendix.

3.4. Implicit Line Fitting and Weighted RANSAC

Given a set of corresponding SIFT features, we can fit an orientation vector using the least squares method. The orientation of this vector is determined by the scale of the associated features, with vectors from larger to smaller scale features regarded as pointing toward the vanishing point. A weighted RANSAC algorithm is employed to remove out-

liers. This algorithm entails randomly sampling two lines based on their weights (see below), calculating their intersection, and tallying the votes for this particular prediction among all other lines. These steps are repeated iteratively until the optimal prediction with the highest number of votes is identified. The vanishing point location is ultimately estimated by least squares from all the inliers.

The initial weight w_i for each line l_i is computed by evaluating the angular relationships between l_i and every other line l_j in the dataset. Specifically, the weight w_i is determined as: $w_i = \sum_{\substack{j=1 \\ j \neq i}}^n e^{-\theta_{ij}}$, where θ_{ij} is the acute angle between lines l_i and l_j . This acute angle is calculated as the arccos of the dot product of the direction vectors of the two lines, and if this is an obtuse angle it is subtracted from π to ensure that $0 \leq \theta_{ij} \leq \pi/2$. The exponential decay function emphasizes smaller angles, thereby assigning higher weights to lines that are closer to being parallel, reflecting a higher probability of contributing to the detection of a consistent vanishing point.

During each iteration of the Weighted RANSAC algorithm, an intersection is estimated and lines are classified as inliers based on their distance to the point. The weights are then updated to refine the selection process in subsequent iterations. For inliers ($l_i \in I$), the weight w_i is incremented as $w'_i = w_i \cdot \alpha$ and for outliers ($l_j \in O$), the weight w_j is decremented as $w'_j = w_j \cdot \beta$, where $\alpha = 1.2$ and $\beta = 0.8$ in our algorithm.

We note that the intersection of near-parallel lines is ill-conditioned and the intersection points could be incorrect. To remedy this, we run the weighted RANSAC multiple times by reinitializing the weights and adaptively finding an optimal threshold to find inliers. Finally, the vanishing point associated with the largest set of inliers is calculated using eigenvalue decomposition.

4. Experimental Setup

4.1. Datasets

We evaluate and compare several VPD algorithms on four different datasets. Two of the datasets, RPVP-Synthetic and RPVP-Real, are created/collected and labeled by us. They have the property that each image contains one or more recurring patterns (RPs) but not necessarily explicit straight lines. We also consider the existing TMM17 dataset used in [46,48] containing natural scenes with one dominant vanishing point, some featuring long, straight lines leading to the VP but not necessarily recurring patterns. Specifically, we use the test-image set used by [46], TMM17-Test, to evaluate all the VPD methods compared in this paper. Furthermore, we note there is some overlap between images containing RPs in our RPVP-Real and TMM17. Since NeurVPS-tmm17 [46] is trained on a large portion of TMM17, to avoid data leakage when comparing against

Table 1. RPVP-Synthetic dataset (3200 images) performance measured by AA^{θ° representing the *AUC of angle accuracy* at the threshold angle θ° . The p -values are computed against R-VPD (Ours) to indicate whether the difference in performance is statistically significant. Corresponding AA curves are shown in Figure 6.

| Method | $AA^{2^\circ}(\uparrow)$ | $AA^{5^\circ}(\uparrow)$ | $AA^{10^\circ}(\uparrow)$ | median(\downarrow) | p -value |
|----------------------|--------------------------|--------------------------|---------------------------|--------------------------------|------------|
| LSD | 0.14 | 0.45 | 1.12 | 45.47 $^\circ$ | 1.80e-16 |
| J-Linkage [42] | 0.42 | 1.46 | 3.43 | 29.41 $^\circ$ | 5.78e-12 |
| GPVPD-nyu [25] | 0.04 | 0.26 | 0.71 | 88.38 $^\circ$ | 1.73e-10 |
| GPVPD-su3 [25] | 0.002 | 0.02 | 0.15 | 73.88 $^\circ$ | 1.03e-18 |
| GPVPD-scannet [25] | 0.004 | 0.04 | 0.22 | 71.01 $^\circ$ | 1.56e-18 |
| NeurVPS-su3 [46] | 0.17 | 0.72 | 1.88 | 91.83 $^\circ$ | 9.37e-17 |
| NeurVPS-scannet [46] | 0.19 | 0.74 | 1.80 | 96.94 $^\circ$ | 3.06e-16 |
| NeurVPS-tmm17 [46] | 0.32 | 1.26 | 3.28 | 12.77 $^\circ$ | 1.15e-07 |
| R-VPD | 0.97 | 3.49 | 7.44 | 0.60$^\circ$ | - |

Table 2. RPVP-Real dataset (1400 images) performance measured by AA^{θ° and p -value. Corresponding AA curves are shown in Figure 7.

| Method | $AA^{2^\circ}(\uparrow)$ | $AA^{5^\circ}(\uparrow)$ | $AA^{10^\circ}(\uparrow)$ | median(\downarrow) | p -value |
|----------------------|--------------------------|--------------------------|---------------------------|--------------------------------|------------|
| LSD | 0.36 | 1.29 | 2.96 | 42.95 $^\circ$ | 4.37e-13 |
| J-Linkage [42] | 0.68 | 2.58 | 6.10 | 1.47 $^\circ$ | 5.27e-04 |
| GPVPD-nyu [25] | 0.19 | 1.05 | 3.06 | 71.53 $^\circ$ | 3.12e-11 |
| GPVPD-su3 [25] | 0.05 | 0.36 | 1.38 | 22.04 $^\circ$ | 4.82e-14 |
| GPVPD-scannet [25] | 0.04 | 0.31 | 1.19 | 26.75 $^\circ$ | 6.42e-15 |
| NeurVPS-su3 [46] | 0.04 | 0.17 | 0.52 | 88.70 $^\circ$ | 6.14e-18 |
| NeurVPS-scannet [46] | 0.23 | 1.12 | 2.92 | 32.35 $^\circ$ | 3.71e-12 |
| R-VPD | 1.03 | 3.75 | 8.63 | 0.74$^\circ$ | - |

that algorithm on RPVP-Real we form a subset RPVP-Real-Exclusive of images in RPVP-Real that are not found in TMM17.

4.1.1 RPVP-Synthetic: Recurring Pattern-based Vanishing Point Dataset with Synthetic Images

A custom dataset is created using *Blender*, an open-source 3D computer graphics software tool [1]. We generate 3,200 images from 16 different synthetic 3D scenes (shown in Figure 4), each contributing 200 images. In this dataset, objects are placed along a straight line and oriented toward a vanishing point, with each scene featuring varying objects, backgrounds, scales, and camera orientations. Ground truth annotations including the (x, y) coordinates of the vanishing point and the camera parameters are provided for each image. More information on scene creation and frame extraction is provided in the Appendix. It is worth noting that a majority of images from this dataset do not contain explicit straight lines oriented towards the vanishing point.

4.1.2 RPVP-Real: Recurring Pattern-based Vanishing Point Dataset with Real-world Images

RPVP-Real contains 1400 “real world” images with VPs that are produced from recurring patterns under perspective

projection. Explicit straight lines may or may not be available. The VPs on these images are manually labeled using the *LabelMe* tool [38] and ground truth annotations are provided as text files for each image with the (x, y) coordinates of the vanishing point. The ground truth vanishing point is then calculated as the intersection of these annotated lines.

4.1.3 TMM17-Test: Natural Scenes Dataset

The TMM17 dataset [46, 48] consists of 2275 images of natural scenes with one dominant vanishing point. As illustrated in Figure 5 TMM17 overlaps with over 70% of RPVP-Real images (images with recurring patterns). The one split [46] used 2000 images for training and 275 images as the test image set. To avoid any data leaking, we choose TMM17-Test, as a validation image set for all VPD methods. It is worth noting that TMM17-Test contains both images with and without RPs (Figure 5). Even though R-VPD is designed primarily to deal with images with RPs, we use TMM17-Test to test the robustness of R-VPD on real world scenes beyond those with recurring patterns.

4.1.4 RPVP-Real-Exclusive

The RPVP-Real dataset has an overlap of about 70.3% with the TMM17 dataset (Sec 4.1.3) used to train NeurVPS-tmm17 [46]. To avoid potential data leaking when compar-

ing with NeurVPS, we form an additional dataset which is the subset (29.7%) of RPVP-Real that has no overlap with TMM17, which we call RPVP-Real Exclusive dataset that consists of 416 images (Figure 5).

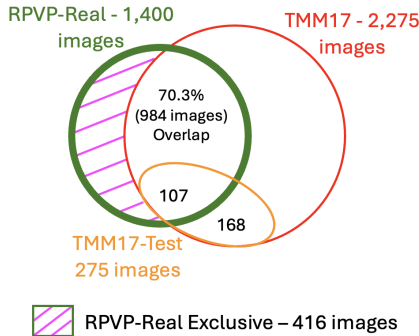


Figure 5. The relationship between RPVP-Real, TMM17, RPVP-Real Exclusive and TMM17-Test. TMM17/TMM17-Test contains natural images with/without RPs while RPVP-Real/RPVP-Real Exclusive contains images with RPs only.

4.2. Evaluation Metrics

We compare our VP detection method with two classical methods: a baseline Line Segment Detection (LSD) method and a more sophisticated J-Linkage [42] approach. We also compare two state-of-the-art deep learning methods, NeurVPS [46] and GPVPD [25]. For NeurVPS, we consider three networks: 1) NeurVPS-tmm17, 2) NeurVPS-su3 and 3) NeurVPS-scanet, where NeurVPS is trained on TMM17 [48], SU3 [47], or ScanNet [12] dataset, respectively. Similarly, for GPVPD, we compare with three variants of their model since they claim in [25] that geometric priors make the dataset variations vanish. The three variants are 1) GPVPD-nyu, 2) GPVPD-su3, and 3) GPVPD-scanet, where GPVPD is trained on NYU [40], SU3 [47], or ScanNet [12] dataset, respectively. With our RPVP datasets we emphasize different algorithms failing to find a VP in the absence of explicit line segments.




Figure 6. RPVP-Synthetic dataset (3200 images): Success rate curves for AA @ 10°. Corresponding AUC is given in Table 1.

We use a vector orientation-based method for comparing a detected VP and the ground truth. In this method,




Figure 7. RPVP-Real dataset (1400 images): Success rate curves for AA @ 10°. Corresponding AUC and are given in Table 2.

the VP is represented as a unit vector in the direction $(vp_x - x_0, vp_y - y_0, f)$, where (x_0, y_0) is the image center, and f is a nominal value that would be the camera focal length if known, but otherwise is chosen to be $(\text{image width} + \text{image height})/4$. Distance in this case is the angle between the detected and ground truth VP unit vectors. We call this metric *angle accuracy* (AA) (also used in [46]). With this metric, we can plot the *success rate curves* for AA for each method by varying the angle thresholds. The success rate of VP detection is computed as the ratio of the number of acceptable VPs against ground truth over all detected VPs.

Along with success rate curves we also provide the *AUC* (area under the curve) to evaluate the performance of different algorithms. For *AUC* we evaluate all algorithms at thresholds $AA = [2, 5, 10]^\circ$. In our *AUC* tables, AA^{θ° represents the *AUC* of *angle accuracy* at threshold angle θ° .

5. Results and Discussions

5.1. P-values

Table 1 provides *AUC* (Figure 6) for all the methods on the RPVP-Synthetic dataset at thresholds $AA = [2, 5, 10]^\circ$ respectively. The last column of the table reports *p*-values, a measure that indicates whether a difference in numerical performance between our method and each of the other methods is statistically significant, and is based on comparing the entire *AUC* curve (all the possible thresholds) between our method and the other methods. Although not commonly used in computer vision (but widely used in scientific domains such as biomedical research), *p*-values indicate whether the numerical results of one method are *statistically significantly different* from another, that is, whether the difference in performance is likely due to a real effect or just random chance. Typically one chooses a threshold of 0.05 or 0.001 to classify *p*-values into significant (lower than threshold) or not significant (higher than threshold).

Table 2 and Figure 7 provide results on the RPVP-Real dataset. NeurVPS-tmm17 is excluded from this comparison due to the overlap between RPVP-Real and TMM17, however a direct and fair comparison between our method and

and NeurVPS-tmm17 will be presented below.

Both quantitative results on RPVP-Synthetic and RPVP-Real show that our method R-VPD outperforms the other compared methods *significantly* ($p\text{-value} \ll 0.001$). This suggests that classical and current learning-based VPD methods struggle when explicit straight-line segments are absent.

5.2. RPVP-Real-Exclusive: RP Image Exclusive

To avoid any potential data leaking issues, we restrict our evaluation of NeurVPS-tmm17 to the RPVP-Real Exclusive dataset that is non-overlapping with TMM17 (Figure 5). Here, our R-VPD method performs slightly better than NeurVPS-tmm17 (Table 3, Figure 8) with $p\text{-value} = 0.58$. Thus, we can state that our method is on par with the state of the art deep learning method such as NeurVPS-tmm17 on VPD.

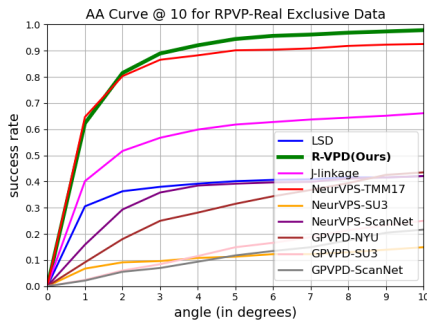


Figure 8. Success Rate curves for the different methods on RPVP-Real Exclusive dataset (416 images). Corresponding AUC is given in Table 3.




Figure 9. Success rate curves for $AA @ 10^\circ$ for TMM17-Test set (275 images). Corresponding AUC is given in Table 4.

Table 3. Performance of NeurVPS and R-VPD on RPVP-Real Exclusive dataset (416 images). Corresponding AA curves are shown in Figure 8.

| Method | $AA^{5^\circ} (\uparrow)$ | $AA^{10^\circ} (\uparrow)$ | median (\downarrow) | $p\text{-value}$ |
|--------------------|---------------------------|----------------------------|--------------------------------|------------------|
| NeurVPS [46] | 3.64 | 8.21 | 0.72° | 0.58 |
| R-VPD(Ours) | 3.71 | 8.54 | 0.71° | - |

Table 4. Performance of all VPD methods on TMM17-Test set (275 images, Fig. 5). Corresponding AA curves are shown in Fig 9.

| Method | $AA^{5^\circ} (\uparrow)$ | $AA^{10^\circ} (\uparrow)$ | median (\downarrow) | $p\text{-value}$ |
|--------------------|---------------------------|----------------------------|--------------------------------|------------------|
| LSD | 1.07 | 2.56 | 71.30° | $8.41e-10$ |
| J-Linkage [42] | 2.23 | 5.59 | 2.42° | 0.52 |
| GPVPD-su3 [25] | 0.53 | 1.76 | 31.83° | $8.16e-10$ |
| GPVPD-nyu | 0.55 | 1.96 | 52.80° | $7.43e-09$ |
| GPVPD-scannet | 0.45 | 1.60 | 35.05° | $6.81e-11$ |
| NeurVPS [46] | 3.49 | 8.14 | 0.86° | 0.003 |
| R-VPD(Ours) | 2.42 | 5.95 | 1.85° | - |

5.3. TMM17-Test: Going beyond RP images

To evaluate on images that do not necessarily have recurring patterns in them, we perform an evaluation on the same test set used in [46], TMM17-Test. Figure 9 and Table 4 provide AUC for all the methods on TMM17-Test. R-VPD performs second to the best on TMM17-Test without any training, and with a $p\text{-value}$ of $p = 0.003$ (> 0.001) the best performer NeurVPS-tmm17 does not significantly outperform our approach. It is worth noting that of the 275 images in the TMM17-Test set, a considerable portion has no RP on the images (Figure 5).

5.4. Qualitative Evaluation and Speed

Figures 10 and 11 present results on some representative images, showing the effectiveness of our R-VPD method compared with ground truth and all the SOTA methods. Classical LSD-based algorithms depend on explicit line segments in images [44] while learning-based approaches are learning edge filters oriented locally towards the VP. For this reason, real-world images that do not have explicit straight lines or edges oriented towards the VP can lead these VPD methods to make erroneous predictions of the VP (Figure 10).

Processing time comparisons are shown in Table 5. NeurVPS and GPVPD achieve their speed by using a GPU and require up to 400 MB of memory, which can be limiting in memory-constrained settings. J-Link appears to be faster here due to fewer lines to process.

The complexity of our algorithm depends on the number n of SIFT features detected and is dominated by the calculation of the distance matrix between these features. Therefore our algorithm's time and space complexity is $O(n^2)$.

Table 5. Processing time per image comparison (median times).

| Method | Hardware | RPVP-Real | RPVP-Synth |
|---------|----------|-----------|------------|
| NeurVPS | GPU | 0.86s | 0.52s |
| GPVPD | | 0.52s | 0.52s |
| J-Link | CPU | 18.67s | 0.75s |
| R-VPD | | 10.73s | 1.31s |

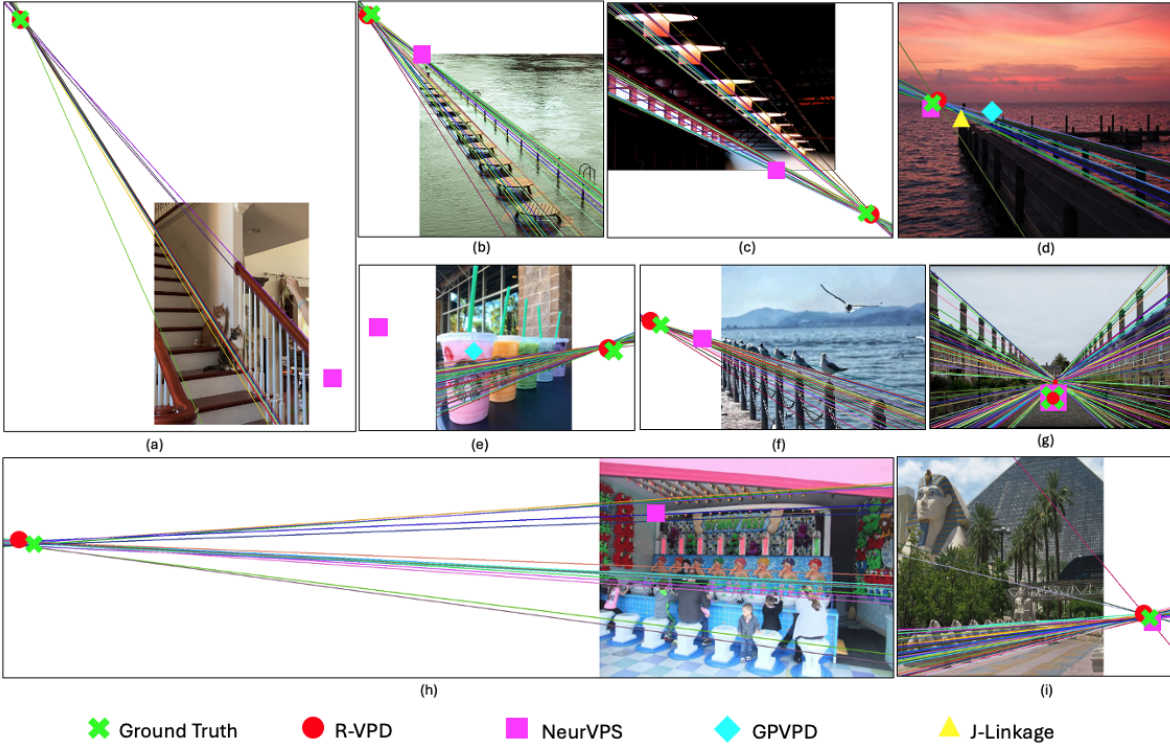


Figure 10. Sample results for VP detection on real-world images using our method and a comparison with other approaches. The absence of a specific symbol in any image indicates the failure of the corresponding method to detect the vanishing point in that image.




Figure 11. One application using a single discovered VP in one-point perspective images is to perform rectification of the original image into an affine (parallelism preserving) frontal view.

6. Conclusion

We have proposed an alternative unsupervised vanishing point detection algorithm, R-VPD, that uses feature correspondences in recurring patterns (RP) to form implicit lines for VPD. We contribute two RP-based vanishing point datasets: RPVP-Synthetic has 3200 synthetic images with VPs and camera parameters, and RPVP-Real has 1400 real-world images with vanishing point annotations. Through a benchmark evaluation of four VPD algorithms - a baseline LSD, J-Linkage [42], NeurVPS [46], and GPVPD [25] against R-VPD, our quantitative results show that R-VPD

outperforms statistically significantly all four approaches on RPVP-Synthetic (Table 1), and on RPVP-Real (Table 2; excluding NeurVPS-tmm17). For fair comparison with NeurVPS-tmm17 we use the RPVP-Real Exclusive dataset, where R-VPD performs better than this SOTA deep learning method (Table 3) but not statistically significantly so, thus “on par” on this real-world RP image set. Finally, when extending to images with and without RPs (for which R-VPD is not designed for), R-VPD performs second to the best with a statistically insignificant difference (with p -value = 0.001 used as the threshold, Table 4).

R-VPD has limitations on RP images with strong perspective and on non-RP images. In this work we focused on natural scenes with a single vanishing point (Figure 11), not Manhattan-world scenes with multiple vanishing points such as those found in indoor or urban environments. We aim to extend current R-VPD to go beyond images with single VPs and to be more robust on general indoor/outdoor scenes.

Acknowledgments. This research is funded in part by NSF Award #1909315 and in part by NSF Award #1248076. We thank all the PSU LPAC members, especially Keaton Yukio Kraiger, for their valuable input.

References

- [1] Blender - a 3d modelling and rendering package. <https://www.blender.org>. 5
- [2] Stephen T. Barnard. Interpreting perspective images. *Artificial Intelligence*, 21(4):435–462, 1983. 2
- [3] Jean-Charles Bazin, Yongduek Seo, Cédric Démonceaux, Pascal Vasseur, Katsushi Ikeuchi, Inso Kweon, and Marc Pollefeys. Globally optimal line clustering and vanishing point estimation in manhattan world. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 638–645. IEEE, 2012. 1
- [4] Jean-Charles Bazin, Yongduek Seo, and Marc Pollefeys. Globally optimal consensus set maximization through rotation search. In *Computer Vision—ACCV 2012: 11th Asian Conference on Computer Vision, Daejeon, Korea, November 5–9, 2012, Revised Selected Papers, Part II 11*, pages 539–551. Springer, 2013. 1
- [5] Ali Borji. Vanishing point detection with convolutional neural networks. *arXiv preprint arXiv:1609.00967*, 2016. 1
- [6] Bruno Caprile and Vincent Torre. Using vanishing points for camera calibration. *International journal of computer vision*, 4(2):127–139, 1990. 1
- [7] Chin-Kai Chang, Jiaping Zhao, and Laurent Itti. Deepvp: Deep learning for vanishing point detection on 1 million street view images. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4496–4503. IEEE, 2018. 1, 2
- [8] Ondřej Chum and Jiří Matas. Planar affine rectification from change of scale. In Ron Kimmel, Reinhard Klette, and Akihiro Sugimoto, editors, *Computer Vision – ACCV 2010*, pages 347–360. Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. 2
- [9] Roberto Cipolla, Tom Drummond, and Duncan P Robertson. Camera calibration from vanishing points in image of architectural scenes. In *BMVC*, volume 99, pages 382–391, 1999. 1
- [10] Robert T Collins and Richard S Weiss. Vanishing point calculation as a statistical inference on the unit sphere. In *ICCV*, volume 90, pages 400–403, 1990. 1, 2
- [11] James M Coughlan and Alan L Yuille. Manhattan world: Compass direction from a single image by bayesian inference. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, pages 941–947. IEEE, 1999. 1
- [12] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 6
- [13] Patrick Denis, James H Elder, and Francisco J Estrada. Efficient edge-based methods for estimating manhattan frames in urban imagery. In *Computer Vision—ECCV 2008: 10th European Conference on Computer Vision, Marseille, France, October 12–18, 2008, Proceedings, Part II 10*, pages 197–210. Springer, 2008. 1
- [14] Alex Flint, David Murray, and Ian Reid. Manhattan scene understanding using monocular, stereo, and 3d features. In *2011 International Conference on Computer Vision*, pages 2228–2235. IEEE, 2011. 1
- [15] Paolo Gamba, Alessandro Mecocci, and U Salvatore. Vanishing point detection by a voting scheme. In *Proceedings of 3rd IEEE International Conference on Image Processing*, volume 2, pages 301–304. IEEE, 1996. 1, 2
- [16] Yuan Gao and Alan L Yuille. Exploiting symmetry and/or manhattan properties for 3d object structure estimation from single and multiple images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7408–7417, 2017. 1
- [17] Erwan Guillou, Daniel Meneveaux, Eric Maisel, and Kadi Bouatouch. Using vanishing points for camera calibration and coarse 3d reconstruction from a single image. *The Visual Computer*, 16:396–410, 2000. 1
- [18] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. 1
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 1
- [20] Seokju Lee, Junsik Kim, Jae Shin Yoon, Seunghak Shin, Oleksandr Bailo, Namil Kim, Tae-Hee Lee, Hyun Seok Hong, Seung-Hoon Han, and In So Kweon. Vpnet: Vanishing point guided network for lane and road marking detection and recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 1947–1955, 2017. 1
- [21] Haoang Li, Yazhou Xing, Ji Zhao, Jean-Charles Bazin, Zhe Liu, and Yun-Hui Liu. Leveraging structural regularity of atlanta world for monocular slam. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2412–2418. IEEE, 2019. 1
- [22] Haoang Li, Jian Yao, Jean-Charles Bazin, Xiaohu Lu, Yazhou Xing, and Kang Liu. A monocular slam system leveraging structural regularity in manhattan world. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2518–2525. IEEE, 2018. 1
- [23] Haoang Li, Ji Zhao, Jean-Charles Bazin, Wen Chen, Kai Chen, and Yun-Hui Liu. Line-based absolute and relative camera pose estimation in structured environments. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6914–6920. IEEE, 2019. 1
- [24] Haoang Li, Ji Zhao, Jean-Charles Bazin, and Yun-Hui Liu. Quasi-globally optimal and near/true real-time vanishing point estimation in manhattan world. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3):1503–1518, 2020. 1, 2
- [25] Yancong Lin, Ruben Wiersma, Silvia L Pinteá, Klaus Hildebrandt, Elmar Eisemann, and Jan C van Gemert. Deep vanishing point detection: Geometric priors make dataset variations vanish. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6103–6113, 2022. 1, 2, 5, 6, 7, 8

- [26] Jingchen Liu and Yanxi Liu. Grasp recurring patterns from a single view. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2003–2010, 2013. [1](#), [3](#)
- [27] Jingchen Liu and Yanxi Liu. Local regularity-driven city-scale facade detection from aerial images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3778–3785, 2014. [1](#)
- [28] Shichen Liu, Yichao Zhou, and Yajie Zhao. Vapid: A rapid vanishing point detector via learned optimizers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12859–12868, 2021. [1](#)
- [29] Evelynne Lutton, Henri Maitre, and Jaime Lopez-Krahe. Contribution to the determination of vanishing points using hough transform. *IEEE transactions on pattern analysis and machine intelligence*, 16(4):430–438, 1994. [2](#)
- [30] Michael J Magee and Jake K Aggarwal. Determining vanishing points from perspective images. *Computer Vision, Graphics, and Image Processing*, 26(2):256–267, 1984. [1](#)
- [31] Gerard F McLean and D Kotturi. Vanishing point detection by line clustering. *IEEE Transactions on pattern analysis and machine intelligence*, 17(11):1090–1095, 1995. [1](#), [2](#)
- [32] Joseph L. Mundy and Andrew Zisserman, editors. *Geometric Invariance in Computer Vision*. MIT Press, Cambridge, MA, USA, 1992. [1](#)
- [33] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hour-glass networks for human pose estimation. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 483–499. Springer, 2016. [2](#)
- [34] James Pritts, Ondrej Chum, and Jiri Matas. Rectification, and segmentation of coplanar repeated patterns. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2973–2980, 2014. [2](#)
- [35] James Pritts, Zuzana Kukelova, Viktor Larsson, Yaroslava Lochman, and Ondrej Chum. Minimal solvers for rectifying from radially-distorted scales and change of scales. *Int. J. Comput. Vis.*, 128(4):950–968, 2020. [2](#)
- [36] J. Pritts, Z. Kukelova, V. Larsson, Y. Lochman, and O. Chum. Minimal solvers for rectifying from radially-distorted conjugate translations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):3931–3948, Nov 2021. [2](#)
- [37] Long Quan and Roger Mohr. Determining perspective structures using hierarchical hough transform. *Pattern Recognition Letters*, 9(4):279–286, 1989. [2](#)
- [38] Bryan C. Russell, Antonio Torralba, Kevin P. Murphy, and William T. Freeman. LabelMe: A database and web-based tool for image annotation. *International journal of computer vision*, 77(1):157–173, 2008. [2](#), [5](#)
- [39] Frederik Schaffalitzky and Andrew Zisserman. Planar grouping for automatic detection of vanishing lines and points. *Image and Vision Computing*, 18(9):647–658, 2000. [1](#), [2](#)
- [40] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part V 12*, pages 746–760. Springer, 2012. [6](#)
- [41] Sudipta Sinha, Drew Steedly, and Rick Szeliski. Piecewise planar stereo for image-based rendering. In *2009 International Conference on Computer Vision*, pages 1881–1888, 2009. [1](#)
- [42] Roberto Toldo and Andrea Fusiello. Robust multiple structures estimation with j-linkage. In *European Conference on Computer Vision*, 2008. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [43] Xin Tong, Xianghua Ying, Yongjie Shi, Ruibin Wang, and Jinfa Yang. Transformer based line segment classifier with image context for real-time vanishing point detection in manhattan world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6093–6102, 2022. [2](#)
- [44] Shimian Zhang, Skanda Bharadwaj, Keaton Kraiger, Yashasvi Asthana, Hong Zhang, Robert Collins, and Yanxi Liu. Novel 3d scene understanding applications from recurrence in a single image. *arXiv preprint arXiv:2210.07991*, 2022. [3](#), [7](#)
- [45] Huizhong Zhou, Danping Zou, Ling Pei, Rendong Ying, Peilin Liu, and Wenxian Yu. Structslam: Visual slam with building structure lines. *IEEE Transactions on Vehicular Technology*, 64(4):1364–1375, 2015. [1](#)
- [46] Yichao Zhou, Haozhi Qi, Jingwei Huang, and Yi Ma. Neurvps: Neural vanishing point scanning via conic convolution. *Advances in Neural Information Processing Systems*, 32, 2019. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [47] Yichao Zhou, Haozhi Qi, Yuexiang Zhai, Qi Sun, Zhili Chen, Li-Yi Wei, and Yi Ma. Learning to reconstruct 3d manhattan wireframes from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7698–7707, 2019. [1](#), [6](#)
- [48] Zihan Zhou, Farshid Farhat, and James Z Wang. Detecting dominant vanishing points in natural scenes with application to composition-sensitive image retrieval. *IEEE Transactions on Multimedia*, 19(12):2651–2665, 2017. [1](#), [2](#), [4](#), [5](#), [6](#)