Pattern Recognition 48 (2015) 3993–4011

Contents lists available at ScienceDirect

Pattern Recognition

journal homepage: www.elsevier.com/locate/pr

# A generalisable framework for saliency-based line segment detection

Mark Brown*, David Windridge, Jean-Yves Guillemaut

*Centre for Vision, Speech and Signal Processing, University of Surrey, Guildford, Surrey GU2 7XH, UK*

ARTICLE INFO

*Article history:*
Received 26 January 2015
Received in revised form
14 May 2015
Accepted 25 June 2015
Available online 6 July 2015

*Keywords:*
Line detection
Feature detection
Feature matching
Registration
Saliency

ABSTRACT

Here we present a novel, information-theoretic salient line segment detector. Existing line detectors typically only use the image gradient to search for potential lines. Consequently, many lines are found, particularly in repetitive scenes. In contrast, our approach detects lines that define regions of significant divergence between pixel intensity or colour statistics. This results in a novel detector that naturally avoids the repetitive parts of a scene while detecting the strong, discriminative lines present. We furthermore use our approach as a *saliency filter* on existing line detectors to more efficiently detect salient line segments. The approach is highly generalisable, depending only on image statistics rather than image gradient; and this is demonstrated by an extension to depth imagery. Our work is evaluated against a number of other line detectors and a quantitative evaluation demonstrates a significant improvement over existing line detectors for a range of image transformations.


## 1. Introduction

Line segments are an important low-level feature, particularly where man-made structures are present. In many situations they may be used in a similar manner to points, e.g. pose estimation [5], stereo matching [9], or structure from motion [8]. This may often be helped by using the *duality* between lines and points, resulting in similar registration approaches for the two types of feature [26]. Further, there are tasks especially suited to lines, e.g. vanishing point estimation for camera calibration [10], image resizing [17], or structural graph matching [19].

Existing line detection methods either first use a derivative-based edge detector and detect lines from the edges (e.g. [4] or via the Hough Transform [6]), or they directly group pixels in the image into line segments based on the magnitude and direction of their derivative [49,14]. However, these all act locally on the image, detecting a large number of lines, particularly in repetitive scenes. This limitation is illustrated[1] in Fig. 1: state of the art line detection detects all lines regardless of their significance, whereas, ideally, the non-repetitive lines denoting the geometrically significant edges would be preferentially detected.

To address this, we propose to detect only the *salient* line segments, an area that, to the best of the authors' knowledge, has not been addressed in the literature. Instead, saliency detection commonly refers to the computation of a *saliency map* (e.g. [31]), with some work addressing salient edge detection [28] and salient point detection [32]. In detecting only the salient line segments, we propose an approach that is fundamentally different from existing methods for line segment detection in that it is not derivative-based: instead, it seeks informational contrast between regions and thereby favours non-repetitive edges. The information is expressed in terms of distributions of pixel intensities taken from rectangles of a variable width, meaning our approach operates over a larger scale than other detectors and so naturally avoids repetitive parts of a scene.

We measure the contrast between the two distributions on either side of the line using the information-theoretic Jensen–Shannon Divergence (JSD). This measure has been used elsewhere for edge detection [39], unlabeled point-set registration [50], and DNA segmentation [25]. It has many interpretations, e.g. it may be expressed in terms of other information-theoretic quantities such as the Kullback–Liebler Divergence and Mutual Information, having further interpretations in both statistical physics and mathematical statistics [25], and is the square of a metric.

Our measure of line saliency may further be used as a *saliency filter* on existing line detectors. This allows it to cull the non-salient line segments computed by other detectors and localise the position of salient lines under our saliency measure. It furthermore increases the speed of salient line detection by orders of magnitude over the naive approach of determining the saliency measure of every possible line segment on the image.

This distribution-based approach to line detection we propose is highly generalisable, being applicable to any situation where informational contrast can be found. As such, we implement an extension for line detection in depth images, whereby lines that

* Corresponding author. Tel.: +44 1483 686046.
*E-mail addresses:* m.r.brown@surrey.ac.uk (M. Brown), d.windridge@mdx.ac.uk (D. Windridge), j.guillemaut@surrey.ac.uk (J.-Y. Guillemaut).

[1] Image by Gila Brand, http://commons.wikimedia.org/wiki/File:Washington_DC_windows.jpg. Licensed under CC BY 2.5. Grayscale of original.

http://dx.doi.org/10.1016/j.patcog.2015.06.015
0031-3203/

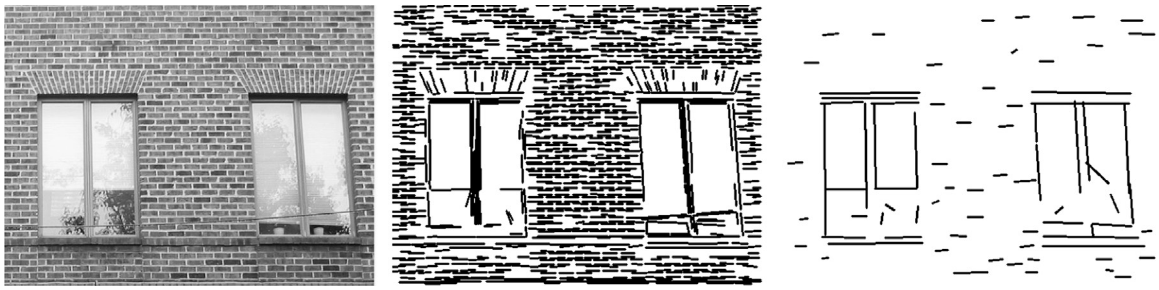

**Fig. 1.** Left: Input image. Centre: LSD algorithm, [49], returning 1026 lines. Right: Our proposed approach, returning 75 lines, indicative of the broad structure of the scene.

jointly delineate changes in surface orientation or texture are detected. These are reprojected, allowing for 3D salient line detection and hence potential multi-modal applications.

The contributions of this paper are as follows: firstly, a distribution-based salient line segment detector is formulated and implemented: the first known method for salient line segment detection. Secondly, the notion of *saliency-based filtering* is applied to existing line detectors for efficient salient line detection. Thirdly, an extension to depth imagery is implemented, allowing for the detection of salient lines in 3D structures. An evaluation shows that, when considering that we detect only a small number of lines, our approaches significantly outperform the others in terms of repeatability and homography estimation. It demonstrates that they are representative of the underlying aspects of the scene, with potential use for problems that benefit from fewer, but more reliable, features e.g. [20].

The structure of the paper is as follows: in Section 2, we review related work in line detection, edge detection, and line detection in depth imagery. In Section 3 the methodology is described for both salient line detection and saliency filtering, with the extension to depth imagery (and subsequently 3D by reprojection) described in Section 4. In Section 5 a range of qualitative and quantitative results are given, and in Section 6 our conclusions and ideas for future work are presented.

## 2. Related work

Since we are unaware of any research into salient line detection (or any line detection method that does not act locally on the derivative of the image) we firstly review line segment detection, before reviewing relevant edge detection methods. Finally, line detection in other modalities (depth images, 3D data) is reviewed.

### 2.1. Line detection

Most early methods of line detection relied upon the Hough Transform (HT) [6] to determine a set of lines from a set of edges (typically extracted from the image by the Canny edge detector [15]). The HT exhaustively searches the space of all possible infinitely long lines, determining how many edge pixels are consistent with each line; lines with a suitably large number of edge pixels lying on them are returned as the output of the algorithm. In its naive form there are many drawbacks, for example it only depends on the magnitude of the gradient and not the orientation, and leaves a problem of how to accurately determine the endpoints of the lines. However, there are many variants of the Hough Transform [30] that seek to solve some of these problems.

Regardless of the approach to line detection, early methods particularly suffered from the problem of setting meaningful thresholds. This was addressed by the Progressive Probabilistic Hough Transform (PPHT) [41] by Matas et al. where it is achieved in a probabilistic manner: the threshold is expressed in terms of the probability of the line occurring by chance. The idea was extended by Desolneux et al. [21] who exhaustively search every line segment on the image and define an *a contrario* model to control the number of false detections. The latter part is a straightforward extension: if there are $N$ possible line segments on an image and $p$ is the probability of that line segment occurring by chance, then accepting the line if $p < \epsilon/N$ guarantees, on average, $\epsilon$ false detections per image.

However, Grompone von Gioi et al. [48] note that this model, in its current form, is too simple. Given an array of line segments, the model tends to interpret it as one long line, leading to unsatisfactory results. This is not a fault of the *a contrario* model, but rather that it is applied to each line individually. If instead it is applied to groups of lines at a time it will segment a line into its components in the correct manner, known as 'multi-segment analysis'. However, this adds another layer of complexity, becoming $O(N^5)$ for an $N \times N$ image.

Grompone von Gioi et al. subsequently implemented a linear-time Line Segment Detector (LSD) [49]. It is based on both the *a contrario* model and an earlier line detection algorithm by Burns et al. [14]. It is a spatially based approach, starting from small line segments and growing them. Furthermore, each segment has its own *line support region*, constructed by grouping nearby pixels that have a similar gradient, thus detecting lines of variable width.

The *a contrario* model has also been implemented in the EDLines detector by Akinlar and Topal [4]. The approach performs similarly to LSD but up to ten times faster due to its very fast edge detection algorithm that simultaneously detects edges and groups them into connected chains of pixels. Less processing time is required for subsequent line detection, resulting in a real-time line segment detector.

All line detection methods reviewed above are unable to detect lines based on their significance or surroundings. Consequently, they tend to return a large number of lines which does not capture the general structure of the scene.

### 2.2. Edge detection in images

Similarly to approaches to line detection, many approaches to edge detection act locally on the image. One of the earliest algorithms, the Canny edge detector [15], convolves the image with a Gaussian filter before computing the magnitude of the gradient at each pixel. Variants have been proposed in particular for the convolution stage; notably Liu and Feng [38] use an anisotropic Gaussian filter that only operates perpendicularly to an edge. It is combined with a multi-pixel search to detect longer edges than other approaches, culminating in the detection of short edge-line segments. Their results indicate superior performance compared to existing edge detectors in the presence of different levels of Gaussian noise. However, both approaches are fundamentally derivative based, acting locally on the image regardless of the structure of the scene.