

Evolution of a local Boundary Detector for Natural Images via Genetic Programming and Texture Cues

Ilan Kadar
Dept. of Computer Science
Ben-Gurion University
Beer-Sheva, Israel
ilankad@cs.bgu.ac.il

Ohad Ben-Shahar
Dept. of Computer Science
Ben-Gurion University
Beer-Sheva, Israel
ben-shahar@cs.bgu.ac.il

Moshe Sipper
Dept. of Computer Science
Ben-Gurion University
Beer-Sheva, Israel
sipper@cs.bgu.ac.il

ABSTRACT

Boundary detection constitutes a crucial step in many computer vision tasks. We present a learning approach for automatically constructing high-performance local boundary detectors for natural images via genetic programming (GP). Our GP system is unique in that it combines filter kernels that were inspired by models of processing in the early stages of the primate visual system, but makes no assumptions about what constitutes a boundary, thus avoiding the need to make *ad hoc* intuitive definitions. By testing our evolved boundary detectors on a highly challenging benchmark set of natural images with associated human-marked boundaries, we show performance that outperforms most existing approaches.

Categories and Subject Descriptors

I.4.6 [Computing Methodologies]: Image Processing and Computer Vision—*Segmentation*

General Terms

Algorithms, Measurement, Performance, Design

Keywords

Boundary Detection, Computer Vision, Machine Learning, Evolutionary Algorithms

1. INTRODUCTION

Boundary detection in images is a fundamental problem in computer vision. The performance of many high-level computer vision tasks, such as segmentation and object recognition, is highly dependent upon the boundary map of an image. A boundary is a contour in the image plane that represents a change in the pixel’s “ownership” from one object or surface to another. In general, there are different types of boundaries: for example, those formed between two regions with an abrupt change in the image brightness, and those formed between two regions with a change in the texture. Clearly, boundaries in natural images are marked by change both in brightness and texture. There are

some attempts in computer vision to address both brightness and texture cues using complex and computationally intensive schemes [5]. In contrast, humans have an outstanding ability to detect boundaries pre-attentively, and hence very fast. Correspondingly, evidence from behavioral science and neuroscience strongly suggests that this process occurs in early stages of visual processing. This paper presents an approach that aims to use genetic programming (GP) as a learning framework for evolving detectors. The detectors are evaluated against human-marked boundary maps in order to accurately detect and localize boundaries in grayscale natural images. The evolving programs use both linear and non-linear operators to combine multiple cues from the early stages in the visual cortex. The presented learning framework was developed based on insights from our recent work [3] with the critical improvement of incorporating texture cues. Our results show that this approach is highly effective at automatically generating boundary detectors. By testing the evolutionary algorithm on a highly challenging benchmark set of natural images with associated human-marked boundaries, we show performance to be quantitatively human competitive [4].

2. METHOD

We present a visual learning approach to automatically construct a boundary detector using *GP* and *texture* cues. Each individual in the GP population represents a candidate boundary cue, which is then combined with a *texture gradient* cue into a single detector on a trained *logistic regression classifier*. *Fitness* is defined as the *F-measure*, computed for a set $T = \{I_i\}$ of n images taken from the training set of the Berkeley data set. The *terminal set* is image independent, such that the terminals for image I_i , given as an array of matrices, are the convolution of the image I_i with filter kernels tuned to various orientations. These filter kernels are inspired by models of processing in the early stages of the primate visual system, which model both odd- and even-symmetric simple cell receptive fields at various orientations (see Figure 1). The *function set* contains both unary and binary functions. The input and output of all functions are arrays of length N of data matrices with the same size as images in T (see Table 1). In order to measure the degree to which texture varies at a location (x, y) in direction θ , we used a *texture gradient* approach inspired by Martin and Malik [5].

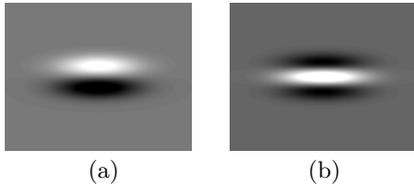


Figure 1: (a) Odd-symmetric simple cell with preferred orientation 0° . (b) Even-symmetric simple cell with preferred orientation 0° .

Table 1: Function set

$M_{\forall i \in 1 \dots N} = \text{Add}(M_{1_i}, M_{2_i})$	Matrix addition
$M_{\forall i \in 1 \dots N} = \text{Sub}(M_{1_i}, M_{2_i})$	Matrix subtraction
$M_{\forall i \in 1 \dots N} = \text{Mul}(M_{1_i}, M_{2_i})$	Matrix multiplication
$M_{\forall i \in 1 \dots N} = \text{Max}(M_{1_i}, M_{2_i})$	Largest elements taken from M_{1_i} or M_{2_i}
$M_{\forall i \in 1 \dots N} = \text{Min}(M_{1_i}, M_{2_i})$	Smallest elements taken from M_{1_i} or M_{2_i}
$M_{\forall i \in 1 \dots N} = \text{Pow2}(M_{1_i})$	Elements are 2 raised to the power M_{1_i}
$M_{\forall i \in 1 \dots N} = \text{Sqrt}(M_{1_i})$	Square root of each element of M_{1_i}
$M_{\forall i \in 1 \dots N} = M_{(i+N/2)\%N}$	R90 (90° Rotation)
$M_{\forall i \in 1 \dots N} = M_{1_i} * K_{\text{odd}_i}$	M_{1_i} convolved with oriented odd symmetric kernel

3. RESULTS

The output of each detector for a given image I_i is a soft boundary map, which provides the probability of a boundary at each image location. The fittest individual generated with our approach is presented in Figure 2. The presented

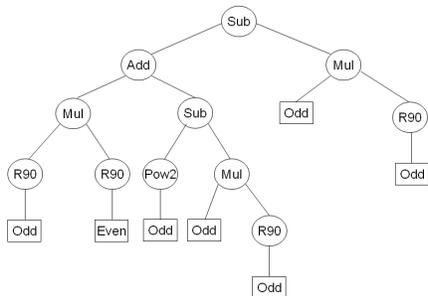


Figure 2: The fittest individual generated with our approach, with performance F -measure = 0.62.

detector was tested on the Berkeley test set of 100 images [6], and the overall performance was computed using the Berkeley benchmark algorithm [1], which computes a score based on the F -measure. Table 2 shows the obtained score of our approach compared with other existing approaches (Note that like us, most of these approaches use training). One of the obtained soft boundary maps is shown in Figure 3.

Analyzing the top individuals might provide additional insights that could improve the selection of function and



Figure 3: (a) Sample test image (b) Extracted soft boundary map, using our best evolved detector, with $F = 0.81$.

Table 2: Performance summary table of local boundary detectors

Method	Performance
Brightness and texture gradient	0.63
GP/TG Detector	0.62
OE/TG Detector	0.61
Learning of the brightness distribution (brightness Gradient)	0.60
GM/TG	0.58
Texture gradient (TG)	0.58
Multiscale gradient magnitude	0.58
Second moment matrix	0.57
Gradient magnitude (GM)	0.56
Segmentation induced by scale invariance	0.48

terminal sets. Similarly, adding mid- and high-level cues that have been shown [2] to improve overall performance on the Berkeley benchmark [1] might also contribute to the evolutionary process. These, and the informed use of additional computational resources to improve performance, all constitute our short-term future work.

4. REFERENCES

- [1] The Berkeley segmentation dataset and benchmark www.cs.berkeley.edu/projects/vision/bsds/.
- [2] P. Dollar, Z. Tu, and S. Belongie. Supervised learning of edges and object boundaries. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2006.
- [3] I. Kadar, O. Ben-Shahar, and M. Sipper. Evolving boundary detectors for natural images via genetic programming. In *Proc. of the 19th International Conference on Pattern Recognition*, Tampa, Florida, 2008.
- [4] J. R. Koza, M. A. Keane, M. J. Streeter, W. Mydlowec, J. Yu, and G. Lanza. *Genetic Programming IV: Routine Human-Competitive Machine Intelligence*. Kluwer Academic Publishers, 2003.
- [5] D. Martin, C. Fowlkes, and J. Malik. Learning to detect natural image boundaries using local brightness, color and texture cues. *PAMI04*, 2004.
- [6] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its applications to evaluating segmentations algorithms and measuring ecological statistics. In *Proc. of international Conference on Computer Vision*, 2001.