The statistics of contour fragments in natural scenes

Josh Merel*, Gasper Tkacik†, Adam Gifford‡, Jason Prentice†, Vijay Balasubramanian†‡
*Biological Basis of Behavior Program & Departments of †Physics and ‡Neuroscience, University of Pennsylvania

1. Introduction

Recent recordings in macaque suggest that single neurons in V4/IT responding to closed contours are sensitive to "contour fragments" with defined orientation, curvature and relative position (Brincat and Connor, Nat.Neurosci.). Motivated by this, we sought to (i) characterize the statistics of contour fragments in natural scenes, and (ii) to generate synthetic images which reflect the measured fragment statistics. We reasoned that if efficient coding extends to higher visual areas, then shape coding in V4/IT should be adapted to the natural statistics of contour fragments. This hypothesis could be tested using synthetic stimuli with controllable contour-fragment statistics.

2. Motivation

Brincat et al [1] parameterize neuronal response to shapes in terms of a small number of key contour features in specific geometric arrangements. Motivated by this experimental characterization of IT cells, we adopt a parameterization of shapes using a small number of contour fragments.

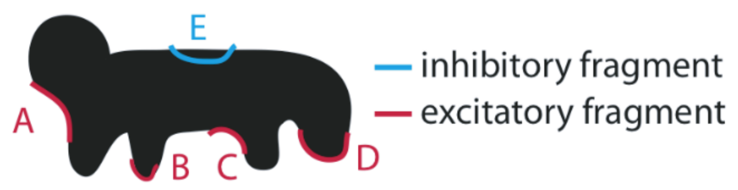

The response to the above shape could be characterized as a a nonlinear sum of the responses to each labelled contour fragment. We begin our image analysis by looking for oriented curved contour fragments in an ensemble of natural images.

3. Set of contour fragments

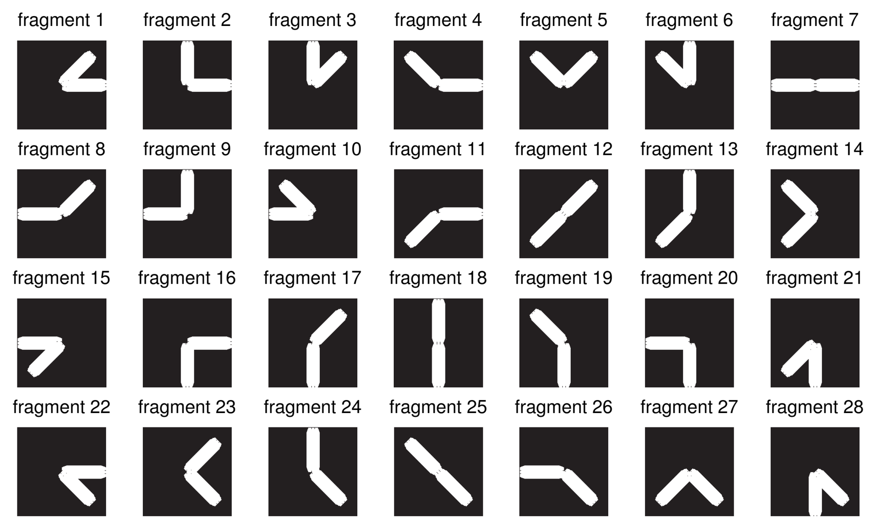

We use a simple basis of contour fragments encompassing a range of opening angles. Though the contours are depicted and formulated as angles, curves similar to a given fragment trigger a response as well. Edges composing the contours depicted here are at 45º angles - this easily extends to basis of opening angles with more finely spaced increments. Each contour fragment is implemented with Gabor filters for each edge, and a thresholding nonlinearity serving as a logical 'AND' to require the presence of both edge components

A note on the images: we have used various image sets throughout this work. including our lab's natural image database [4], an online object images database [3], and frames from video of the BBC production "Planet Earth" found on YouTube. The figures in section 4 use data from "Planet Earth", and sections 5 & 6 use data and images from [3].

4. Image statistics

First order statistics

We measured the frequency with which each fragment is detected in an image enemble (first order statistics). The distribution of relative frequency of occurence of each fragment depends heavily on the nature of the image ensemble.

Second order statistics

We measured the cross-covariance functions of each pair of fragments (second order statistics). We assumed that natural images are translation invariant thereby allowing us to approximate the cross- covariance function with a more efficient computation in the fourier domain involving a convolution-like operation. Ensemble cross-covariances were determined by computing the cross-covariances for each single image and then averaging across images.

Example cross-covariances for pairs of fragments are presented here. Panels on the right show contour fragments at a displacement corresponding to the arrow superimposed over the cross-covariance. Axes indicate pixel-displacement between the vertices of the fragments with the center corresponding to no displacement.

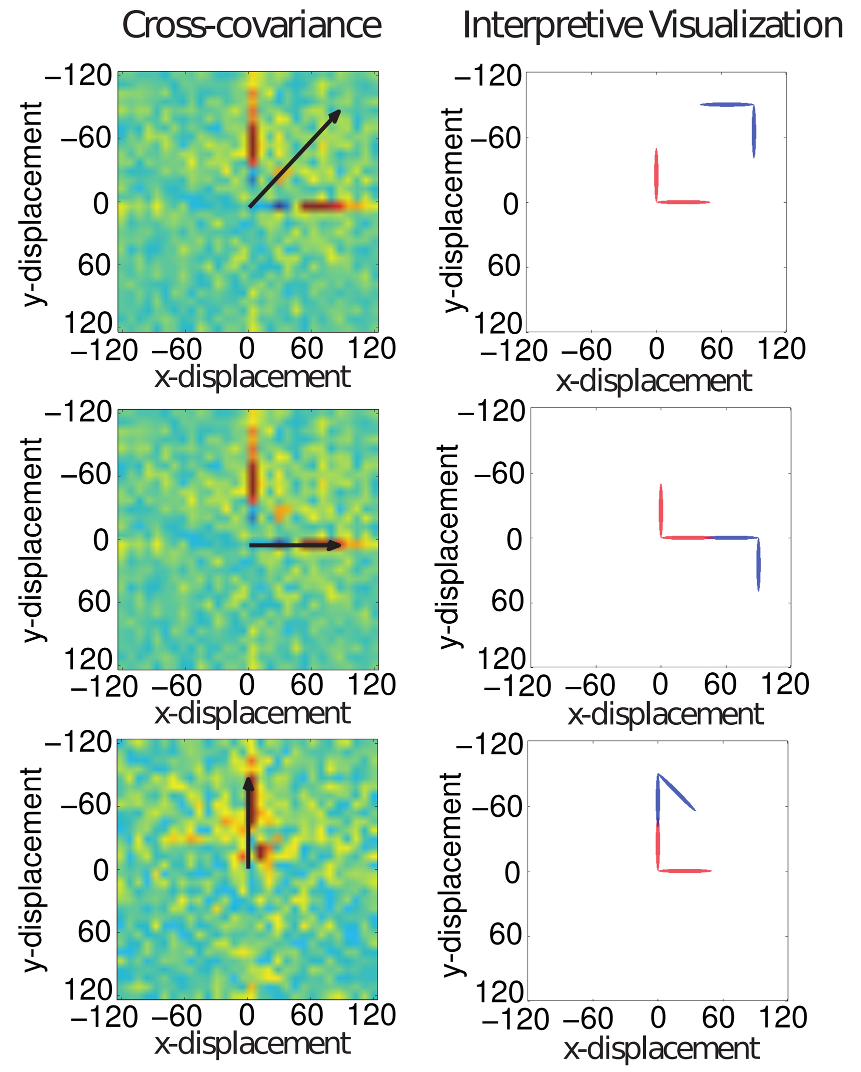

Results & Interpretation

Correlation between fragments is greatest when fragments are collinear (see also [2]), with the peak being most pronounced where edge-detectors overlap. Our statistics, however, potentially afford novel insights beyond those of edge fragment statistics as we may compare the relative correlation between a single fragment and various other fragments with which it is collinear.

Presented below are graphs of correlation as a function of distance (along a collinear segment) for a single fragment and all others with which it is collinear along its horizontal edge.

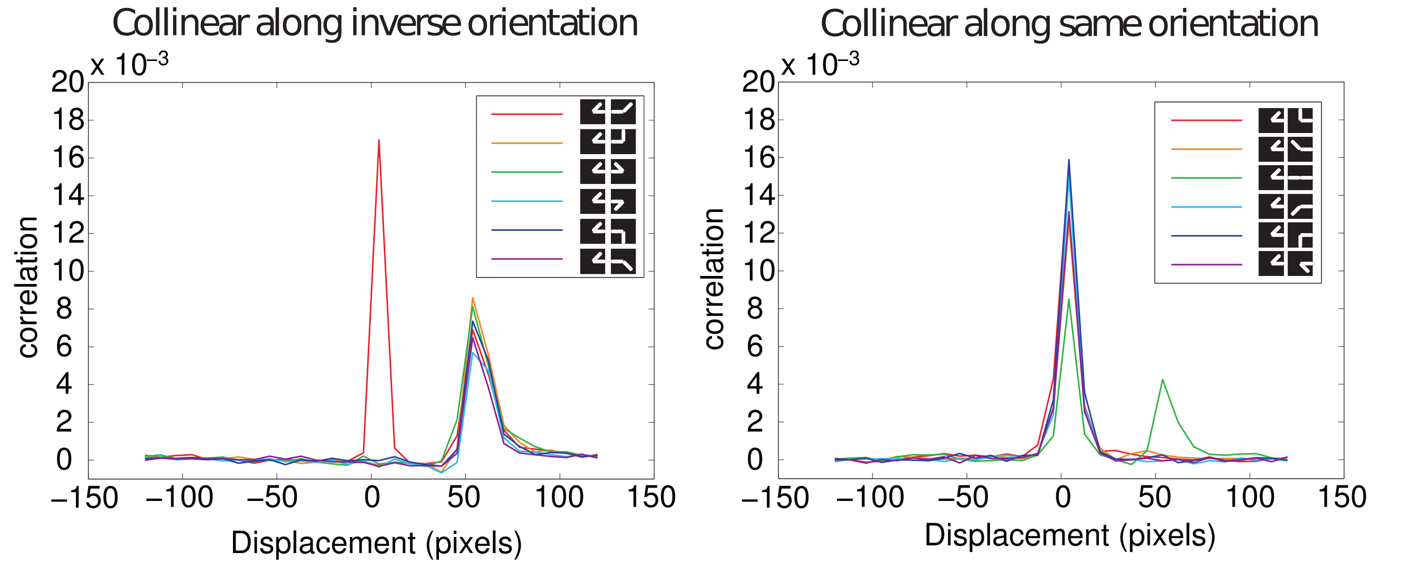

5. Generative model of shape "images"

First and second order statistics (i.e. fragment frequencies and joint probabilities which we approximate via the cross-covariances) were used as constraints in a maximum entropy Ising-like model to generate a probability distribution of shape fragments for four spatial locations, arranged in a square configuration. The basic model can be described by the following equation:

$$P(\{\sigma\}) = \frac{1}{Z(h,J)} exp(-\sum_{i=1}^{N} h_i^\alpha \sigma_i^\alpha - \sum_{i<j}^{N} J_{ij}^{\alpha\beta} \sigma_i^\alpha \sigma_j^\beta)$$

σ_i^α is 1 if fragment α is present at site i

Maximum entropy distribution determined for 4 locations and 24 contour fragments; each 4-fragment pattern is enumerated along the abscissa and its probability is given along the ordinate. The configurations that represent relative peaks in the distribution are drawn connected to their respective peaks by blue arrows. The major peaks show a high prevalence of right-angle shape fragments (reflecting the first order statistics)

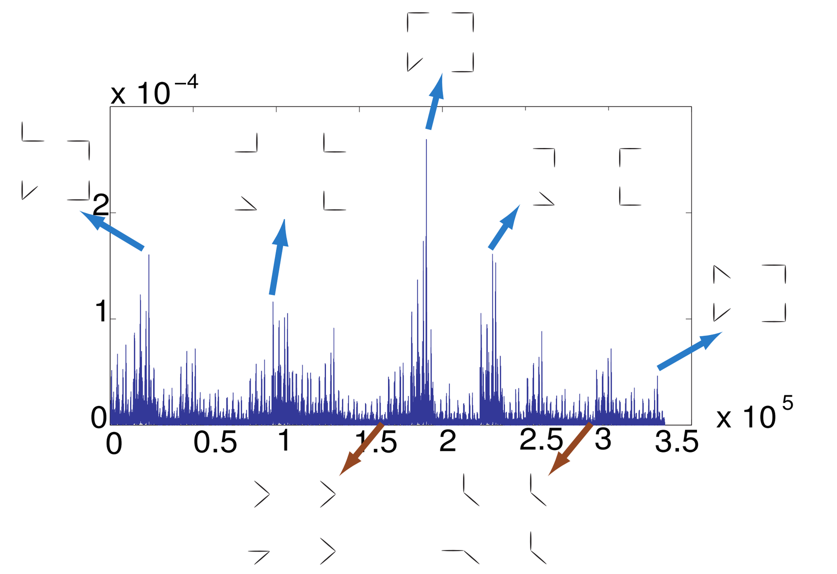

as well as potential collinearity and closedness. Compare this to the two least probable configurations, identified by red arrows, which show little collinearity and shape fragments with low fequencies of occurrence.

6. Fragment detector visualization

For a single image taken from [3], we demonstrate our contour fragment detection procedure as well as different parametric randomizations of the detected fragment set.

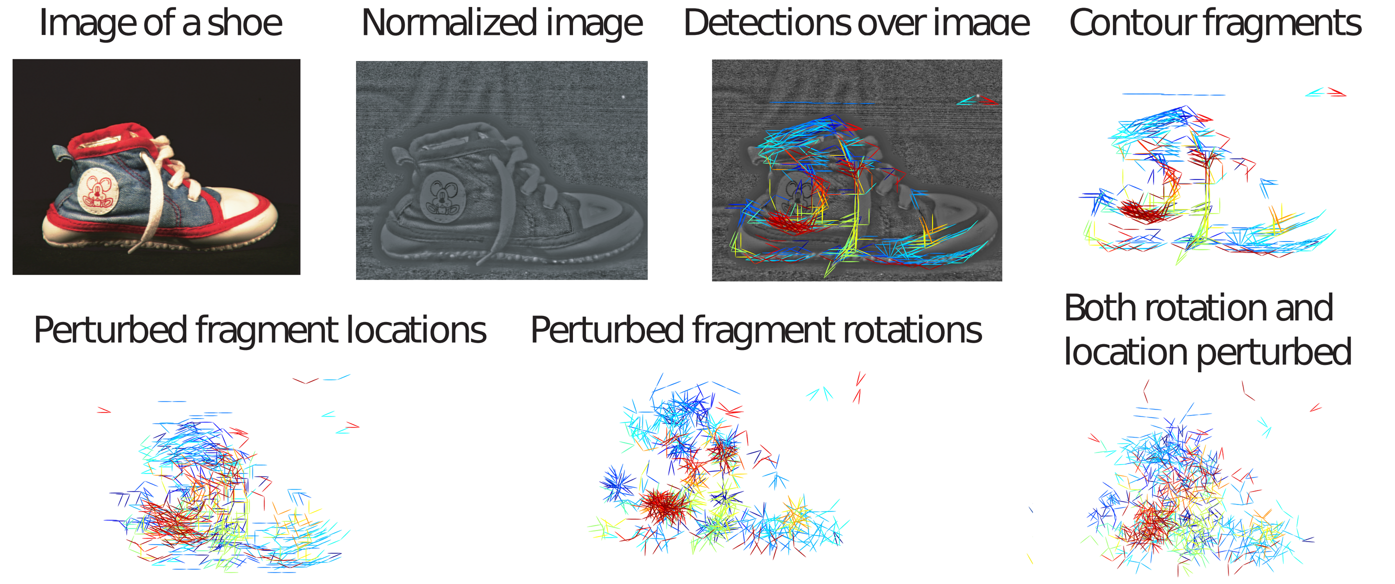

7. Key Points and Future Goals

(1) Contour fragments statistics can be measured using our procedure, and they seem to contain novel information beyond single edge statistics. (2) A generative model which takes into account only first and second order contour statistics may reproduce some key features of shapes - namely collinearity and closedness. (3) Leveraging shape-related statistics measured in natural scenes, we hope to generate theoretical predictions for an optimal representation of shape in IT. (4) We are developing synthetic stimuli based on contour fragment statistics to probe neural mechanisms of shape representation. (5) We intend to extend our procedure to measure 3-point contour fragment statistics more directly related to shape.

References: [1]Brincat, SL, & Connor, CE. 2004. Underlying principles of visual shape selectivity in posterior inferotemporal cortex. Nature Neuroscience, 7, 880–886. [2]Geisler, Wilson S., & Perry, Jeffrey S. 2009. Contour statistics in natural images: Grouping across occlusions. Visual Neuroscience, 26(01), 109–121. [3]Geusebroek, Jan-Mark, Burghouts, Gertjan J, & Smeulders, Arnold W M. 2005. The Amsterdam Library of Object Images. International Journal of Computer Vision, 61(1), 103–112. [4]Tkacik, G., Garrigan, P., Ratliff, C.P., Klein, J., Sterling, P., Brainard, D., & Balasubramanian, V. A calibrated natural image database. Submitted [Preprint].