

From Neurons to Roads: Detection of Curvilinear Structures in Images

AGATA MOSINSKA, COMPUTER VISION LAB, EPFL



Linear Structures





Motivation

Brain connectivity



Diagnostics



Autonomous driving



Material properties





Challenges

Requires large context







Challenges

Occlusions and noise







Challenges

Variability of appearances, presence of other linear structures









Detection Pipeline





Segmentation



Final Graph

Image



Current Approach – Pixelwise Loss

Encoder-decoder architecture trained using binary cross-entropy



$$egin{aligned} \mathcal{L}_{bce}(\mathbf{x},\mathbf{y},\mathbf{w}) &= -\sum_i \left[(1-\mathbf{y}_i) \cdot \log(1-f_i(\mathbf{x},\mathbf{w}))
ight. \ &+ \mathbf{y}_i \cdot \log f_i(\mathbf{x},\mathbf{w})
ight]. \end{aligned}$$

- \mathbf{x} input patch
- \mathbf{y} ground truth
- f function parametrized by **w**



Topology Loss

Train using direct pixel-wise loss and topology loss that captures higher-level features



$$\mathcal{L}_{top} = \sum_{n=1}^{N} \frac{1}{M_n W_n H_n} \sum_{m=1}^{M} \left(l_n^m(\mathbf{y}) - l_n^m (f(\mathbf{x}, \mathbf{w})) \right)^2$$

 l_n^m – nth channel of mth layer of a pretrained VGG

$$\mathcal{L} = \mathcal{L}_{bce} + \mu \mathcal{L}_{top}$$

Mosinska et al., CVPR 2018



Pixelwise vs. Topology Loss

Binary cross entropy is the same in every case, but topology loss differentiates between mistakes depending on how much they alter the geometry.





Activations

•Different filters are scale- and orientation-selective

•Some respond to long, smooth lines, while the others are activated by small background noise





Relation to Perceptual Loss

•Perceptual loss measures perceptual similarity between the original and generated images

•Unlike in our application there is no defined ground truth for style transfer

•Pre-trained VGG is applied here to natural images rather than binary masks

Content







Style-transferred image





Results





Results





Recursive Refinement

•Use <u>the same network</u> to progressively refine the results keeping the number of parameters constant







Results





Inferring Connectivity





Inferring Connectivity

•Segmentation may still contain some mistakes – gaps or redundant predictions.

•It is not enough just to threshold it.

Should those road segments be connected?



How are those neurons connected?



•We need to infer long-range relationships between the detected segments to recover the connectivity.



Path Classification

•Sample possible paths from the segmentation.

- •For every candidate path extract HOG-like features that describe its curvature and neighbourhood.
- •Train a classifier to assign probability to every path.
- •The result is a weighted graph where weight is the cost of keeping path in the final solution





Path Classification



Image



Segmentation

Sampled paths with overlaid probabilities



Mixed Integer Program

In order to ensure connectivity and other constraints:

- •Every path treated as a binary variable:
 - "1" leave in the solution
 - "0" discard it

•Find minimum cost subgraph subject to constraints:

- Allow/disallow loops
- Conservation of flow
- Connection to the root

Turetken et al., PAMI 2014

From Segmentation to Large-scale Reconstruction







Conclusion

- Investigating larger context and higher-level features of linear structures is essential for effective delineation
- •Topology quality can be approximated by the features extracted using a pretrained VGG.
- •By reasoning about the long-range interactions between linear segments we can recover full connectivity of curvilinear structures.

Thank you for attention Agata Mosinska agata.mosinska@epfl.ch