Beyond Trees: MRF Inference via Outer-Planar Decomposition

Dhruv Batra (CMU), Andrew C. Gallagher (Eastman Kodak Company), Devi Parikh (TTIC), Tsuhan Chen (Cornell University)

Structured Discrete Labelling Problems in Computer Vision

Binary Segmentation

Gender Labelling

Geometric Labelling

Optical Flow

Semantic Image Labelling

Stereo

Name-Face association

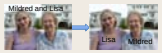

Denoising

Structured Prediction Framework

Graph Structure

$G = (\mathcal{V}, \mathcal{E})$

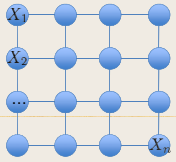

Discrete Random Variables:

$$\mathcal{X} = \{X_1, X_2, \ldots, X_n\}$$

$$X_1 = \{1, 2, \ldots, k\}$$

Pair-wise Energy/Cost Function:

$$E(\mathcal{X}) = \sum_{i \in \mathcal{V}} E_i(X_i) + \sum_{(i,j) \in \mathcal{E}} E_{ij}(X_i, X_j)$$

Node Energies or Local Costs

Edge Energies or Distributed Prior

Map Inference:

$$\mathcal{X}^* = \operatorname*{argmin}_{\mathcal{X} \in \mathcal{L}} E(\mathcal{X})$$

Exact Algorithms for Subclasses:

- Trees & Low Treewidth models [Pearl, AAAI '82]
- Submodular Energies [Hammer '65, Kolmogorov PAMI '04]
- Planar Ising Models / Outer-planar graphs [Schraudolph NIPS '08]

Approximate Algorithms:

- Loopy BP [Pearl, '88]
- Tree-Reweighted MP [Wainwright '05, Kolmogorov '06, Komodakis '07]
- **Outer-Planar Decomposition (OPD)**

Outer-Planarity and Exact Inference in Outer-Planar Graphs

G is Outer-Planar iff

1.Allows a planar embedding AND

2.All nodes lie on an unbounded external face

(i.e., accessible from "outside")

Examples: Trees, lot more

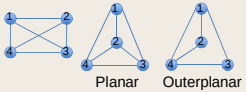

In Computer Vision:

Pixel Grid Graph

Superpixel Graph

Planar = 2 Outerplanar [Goncalves '05]

Exact Inference:

Dynamic Programming

$O(n \cdot L^3)$

Low Treewidth = 2

Reduction to Max-Cut [Hammer '65; Schraudolph '08]

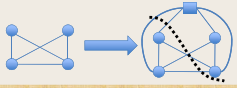

Outer-Planar Decomposition (OPD)

Non-outerplanar graph

OPD

Outer-Planar Sub-Graphs

$$E(\mathcal{X}) = \sum_{i=1}^{p} E^i(\mathcal{X}^i)$$

Accumulator Variables

Messages

BP and TRW are special cases of OPD!

Message-passing

OPD-MP

- Messages: **Min-marginals**

$$\mu_{i:s} = \min_{X_i = s} E(\mathcal{X}) \quad \forall i \in \mathcal{V}, s \in \mathcal{L}$$

- BP is a special case
- To Explore: If optimality properties of BP [Weiss & Freeman '01] carry over to the generalized case

OPD-DD

- Messages: **Map-states**

$$\mathcal{X}^{i*} = \operatorname*{argmin}_{\mathcal{X} \in \mathcal{L}} E^i(\mathcal{X}^i)$$

- TRW –DD is a special case [Komodakis '07, '09]
- Projected subgradient ascent on Lagrangian dual.
- Guaranteed to converge.
- Guaranteed to perform better than TRW-DD!

Synthetic Experiments (subset)

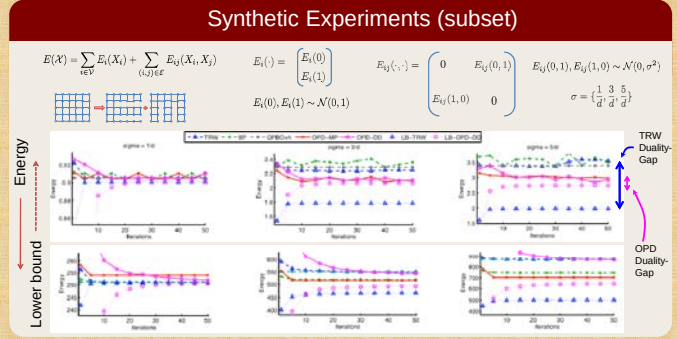

Real Experiments (subset)

Semantic image Labelling

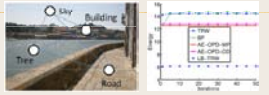

Face-Gender Assignment

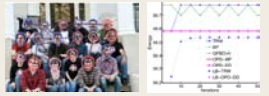

Optical Flow

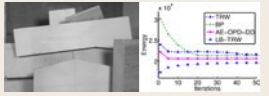

Bigger Pictures

Tighter Outer-Bound on the Marginal Polytope

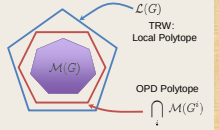

Finding Tightest OPD-Approximation is NP-hard

Greedy Heuristic: