

Contents

1	Introduction	1
1.1	List of Contributions	5
1.2	Thesis Outline	5
1.3	Publications	6
2	Information Theoretic Clustering	9
2.1	Related Work	10
2.2	Clustering using Non-Parametric Entropy Estimates	12
2.2.1	Minimum Spanning Tree Based Entropy Estimation	13
2.2.2	Finding Euclidean Minimum Spanning Tree Clusterings	14
2.2.3	Estimating Intrinsic Dimensionality	17
2.3	Experiments	18
2.3.1	Experimental Setup	19
2.3.2	Qualitative Results	19
2.3.3	Quantitative Results	20
2.4	Summary	21
3	Weakly Supervised Object Segmentation	23
3.1	Related Work	24
3.1.1	Object Segment Proposals	24
3.1.2	Multi-Instance Methods	25
3.1.3	Semantic Scene Segmentation using Weak Annotation	26
3.2	Multi-Instance Kernels for Image Segmentation	27
3.2.1	Constraint Parametric Min-Cuts (CPMC)	27
3.2.2	Multi-Instance Learning using MI-Kernels	27
3.2.3	Segment Features	28
3.2.4	Combining Segments	28

3.3	Experiments	29
3.3.1	Instance-Level Predictions using MI-Kernel	29
3.3.2	Partially Supervised Image Segmentation on Graz-02	30
3.4	Summary	32
4	Learning Conditional Random Fields	33
4.1	Basic Concepts in Structured Prediction	34
4.1.1	Factor Graphs and the Relation to Graphical Models	34
4.2	Learning Max-Margin Structured Prediction	37
4.2.1	Stochastic Subgradient Descent	38
4.2.2	The n -Slack Cutting Plane Method	39
4.2.3	The 1-Slack Cutting Plane Method	41
4.2.4	The BCFW Algorithm	42
4.3	Conditional Random Fields for Semantic Segmentation	45
4.3.1	Fundamentals of Conditional Random Fields	45
4.3.2	Data, Features and Superpixels	47
4.3.3	Previous Work	48
4.4	Casting Structured Prediction into Software	49
4.4.1	Library Structure and Content	49
4.4.2	Project Goals	51
4.4.3	Usage Example: Semantic Image Segmentation	52
4.4.4	Experiments	52
4.5	Summary	54
5	Empirical Comparison of Learning Algorithms	55
5.1	Datasets and Models	55
5.1.1	Multi-Class Classification (MNIST)	55
5.1.2	Sequence Labeling (OCR)	56
5.1.3	Multi-Label Classification	57
5.1.4	2D Grid CRF (Snakes)	58
5.1.5	Superpixel CRFs for Semantic Segmentation	59
5.2	Experiments	60
5.2.1	Experiments using Exact Inference	61
5.2.2	Experiments using Approximate Inference	66
5.3	Summary	70

6	Learning Loopy CRF Exactly	73
6.1	Introduction	73
6.2	Related Work	74
6.3	Learning SSVMs with Approximate Inference	75
6.3.1	Bounding the Objective	75
6.4	Efficient Exact Cutting Plane Training of SSVMs	77
6.4.1	Combining Inference Procedures	77
6.4.2	Dynamic Constraint Selection	78
6.5	Experiments	79
6.5.1	Inference Algorithms	79
6.5.2	Semantic Image Segmentation	79
6.5.3	Caching	80
6.5.4	Implementation Details	83
6.6	Summary	84
7	Learning Depth-Sensitive Conditional Random Fields	85
7.1	Related Work	86
7.2	Learning Depth-Sensitive Conditional Random Fields	88
7.2.1	Low Level Segmentation	88
7.2.2	Unary Image Features	90
7.2.3	Pairwise Depth-Sensitive Features	90
7.3	Experiments	92
7.4	Summary and Discussion	95
8	Conclusion	97
8.1	Future Directions	98
9	Bibliography	101