

Toward Parts-Based Scene Understanding with Pixel-Support Parts-Sparse Pictorial Structures

Jason J. Corso

Computer Science and Engineering
SUNY at Buffalo
jcorso@buffalo.edu

Abstract

Scene understanding remains a significant challenge in the computer vision community. The visual psychophysics literature has demonstrated the importance of interdependence among parts of the scene. Yet, the majority of methods in scene understanding remain local. Pictorial structures have arisen as a fundamental parts-based model for some vision problems, such as articulated object detection. However, the form of classical pictorial structures limits their applicability for global problems, such as semantic pixel labeling. In this paper, we propose an extension of the pictorial structures approach, called pixel-support parts-sparse pictorial structures, or PS3, to overcome this limitation. Our model extends the classical form in two ways: first, it defines parts directly based on pixel-support rather than in a parametric form, and second, it specifies a space of plausible parts-based scene models and permits one to be used for inference on any given image. PS3 makes strides toward unifying object-level and pixel-level modeling of scene elements. In this paper, we implement the first half of our model and rely upon external knowledge to provide an initial graph structure for a given image. Our experimental results on benchmark datasets demonstrate the capability of this new parts-based view of scene modeling.

Keywords: scene understanding, pictorial structures, image segmentation, object recognition, stochastic optimization

1. Introduction

We consider the semantic pixel labeling problem: given a set of semantic classes, such as tree, cow, etc., the task is to associate a label with every pixel. Although hotly studied in recent years, semantic labeling remains a critical challenge in the broader scene understanding community, for obvious reasons like high intraclass variability, occlusion, etc. Early approaches have relied on texture clustering and segmentation, e.g., (Carson et al., 2002). More recently, conditional random fields have become the de facto representation for the problem, e.g., (Shotton et al., 2009). Most such methods learn a strong classifier based on local patches or superpixels and specify some form of a smoothness prior over the field.

Although these methods have demonstrated good success on challenging real world datasets, e.g., (Shotton et al., 2009), their performance remains limited for one key reason: they are intrinsically local making it difficult to incorporate any notion of object and even region semantics. Yet, the visual psychophysics literature has demonstrated the clear importance of modeling at the object- and inter-object relational level for full scene understanding (Biederman, 1981). Although there has been some work on overcoming the challenge of locality (see Related Work), there has been little work on a notion of scene parts nor the inter-relationship among the parts.

In contrast, we present a parts-based approach to full semantic pixel labeling that marries an object-level model of the parts of the scene with a pixel-level representation, rather than a strictly pixel- or region-level model. Our method sits in the broad class of pictorial structures (Fischler and Elschlager, 1973), which has shown notable success at articulated object modeling in recent years, e.g., (Felzenszwalb and Huttenlocher, 2005). However, classical pictorial structures are not well-suited to semantic image labeling: they (1) parameterize object parts and abstract them completely from the pixel-level, (2) require all parts to be present in the scene, and (3) typically adopt simple relational models (linear springs). These three characteristics of the classical models make them unsuitable for image labeling problems.

Our method, called pixel-support parts-sparse pictorial structures, or PS3, overcomes these limitations and takes a step towards a parts-based view of image understanding by proposing a joint global model over image parts—objects in the scene such as trees, cars, the road, etc.— which are each nodes in the pictorial structures graph. It directly ties each part to a set of pixels without any restrictive parameterization, which affords a rich

set of object-level measurements, e.g., global shape. PS3 also defines a space of plausible part-graphs and learns complete relation models between the pairwise elements. At inference, a suitable part-graph is selected (in this paper, manually) and then optimized, which jointly localizes the parts at an object level and performs semantic labeling at the pixel level. Due to the global nature of the PS3 models, de facto energy minimizers are unusable; we use a Metropolis-Hasting Simulated Annealing algorithm to solve the energy minimization problem. We propose a novel principled method of setting the cooling schedule that adapts to each image and requires no manual tweaking.

We have tested our method on the MSRC and the SIFT-Flow benchmarks and demonstrate better performance with respect to maximum likelihood and Markov random field performance in a controlled experimental setting (exact same appearance models). We also compare our methods to existing semantic pixel labeling approaches, but do so with limited significance due to our assumption of being given the parts-graph for a test image. In the remainder of the paper, we present some related papers, then describe classical pictorial structures, our extensions including an appropriate inference algorithm, our experimental results, and conclusion and future work.

2. Related Papers

Several other recent papers have similarly demonstrated the significance of moving beyond local methods. (Hoiem et al., 2008) demonstrate the value of incorporating partial 3D information about the scene during detection. (Li-Jia et al., 2009) take a hierarchical approach to full scene understanding by integrating patch-level, object-level, and textual tags into a generative model. These examples hold strong promise for scene understanding, but are not directly applicable to labeling. One promising method applicable to labeling is (Gould et al., 2008), which proposes a relative location prior for each semantic class and model it with a conditional random field over all of the superpixels in an image. Whereas their approach defines a joint distribution over each of the superpixels in an image, which potentially remains too local, our approach defines it essentially in a layer above the superpixels, affording global coverage and the capability to also model the part-shape.

Another strategy has been to share information among different sub-problems in image understanding. The Cascaded Classification Models approach (Heitz et al., 2008) shares information across object detection and geometrical reasoning. Yang et al. (2010) drive a pixel-labeling process by a

bank of parts-based object detectors; their method demonstrates the power of explicitly modeling objects and their parts within the labeling process.

The Layout Consistent CRF (Winn and Shotton, 2006) uses a parts-based representation of object categories to add robustness to partial occlusion and captures different types of local transitions between labels. Other methods look to hierarchies. (Ladicky et al., 2009) propose an elegant hierarchical extension to the problem that currently performs best on the classic MSRC benchmark (Shotton et al., 2009). However, in principle, it remains local and does not incorporate a notion of scene parts nor inter-relationship among the scene parts; indeed, nor do any of these prior methods.

3. Classical Pictorial Structures

Pictorial structures (PS) are a parts-based representation for objects in images. Classical PS models (Fischler and Elschlager, 1973) are in the class of undirected Markov graphical models. Concretely, pictorial structures represent an object as a graph $G = (V, E)$ in which each vertex v_i , $i = 1, \dots, n$ is a *part* in the n -part model and the edges $e_{ij} \in E$ depict those parts that are connected. A *configuration* $L = \{l_1, \dots, l_n\}$ specifies a complete instance of the model, with each l_i specifying the parametric description of each part v_i . For example, in human pose estimation, each l_i can specify the location, scale, and in-plane rotation of each body part.

The best configuration for a given image I is specified as the one minimizing the following energy:

$$L^* = \arg \min_L \left(\sum_{i=1}^n m_i(l_i|\theta) + \sum_{e_{ij}} d_{ij}(l_i, l_j|\theta) \right), \quad (1)$$

where the m_i and d_{ij} potentials specify the unary and binary potentials, respectively, for parts l_i and l_j , and θ specify model parameters. The specific form of these potential functions is arbitrary, but they are most commonly Gaussian functions, which gives rise to a spring model interpretation.

Such parts-based models have found success in the computer vision community for object recognition problems. Firstly, pictorial structures are a general framework for parts-based modeling. Secondly, although the optimization is, in general, NP-hard, under certain conditions, such as a tree-structured graph (Felzenszwalb and Huttenlocher, 2005), the global optimum

can be reached efficiently. Thirdly, pictorial structures have a clear statistical interpretation in the form of a Gibbs distribution:

$$P(L|I, \theta) = \frac{1}{Z(\theta)} \exp[-H(L|I, \theta)] \quad , \quad (2)$$

where $Z(\cdot)$ is the partition, or normalizing, function and $H(\cdot)$ is the energy function defined in (1). This statistical view permits principled estimation of the model parameters and globally convergent inference algorithms even in the case of general potentials.

However, classical pictorial structures have significant limitations when applied to more general problems in which (1) some parts may be missing, (2) a distribution over structures is present rather than a single one, and (3) a precise segmentation of each part is required rather than strictly its parametric description. One such problem is semantic pixel labeling. In most images, only a few of the classes are present: e.g., four to five for the 21 class MSRC data set (Shotton et al., 2009). Furthermore, the standard parametric descriptions of the parts l_i do not readily map to pixel labels.

4. The PS3 Model for Semantic Labeling

We begin with a concrete problem definition for semantic scene labeling. Let Λ be the pixel lattice and define the basic elements $\lambda \subset \Lambda$ to be either individual pixels, patches, or superpixels, such that $\bigcup \lambda = \Lambda$ and $\lambda_1 \cap \lambda_2 = \emptyset$. Let \mathcal{Z} specify the set of semantic class labels, e.g., car, tree, etc., and denote z_λ as the label for element λ . In the maximum a posteriori view, the labeling problem is to associate the best label with each element

$$\{z_\lambda\}^* = \arg \max_{\{z_\lambda\}} P(\{z_\lambda\}|I, \theta) \quad , \quad (3)$$

but we do not directly model the problem at the pixel level. Rather, we model it at the object level l_i as we now explain.

Parts with Direct Pixel Support. We take a nonparametric approach and directly represent the part l_i based on its pixel support. Each part l_i comprises a set of basic elements $\{\lambda^{(1)}, \lambda^{(2)}, \dots\}$, and induces a binary map, $B_i: \Lambda \mapsto \{0, 1\}$. A configuration L jointly represents a high-level description of the elements in a scene, and also a direct semantic labeling of each pixel in the image. Furthermore, rich, pixel-level descriptions of part-appearance and part-shape are now plausible. However, it adds significant complexity

into the estimation problem: fast inference based on max-product message passing (Felzenszwalb and Huttenlocher, 2005) is no longer a viable option as the parts have a more complex interdependent relation among their supports. **Parts-Sparse Pictorial Structures.** Classically, pictorial structures models are defined by a fixed set of n parts, and all are expected in the image. In scene labeling, however, most images contain a small subset of the possible labels \mathcal{Z} . We consider the space Ω containing all plausible pictorial structures for scene labels. Ω is large, but finite: for an image of size \mathbf{w} , the upper bound on nodes in a PS3 model is $|\mathbf{w}|$, but the typical number is quite smaller, e.g., around four to five for the MSRC dataset. Each node can be of one class type from \mathcal{Z} . Whereas classical pictorial structures model the parameters θ for a specific structure, in PS3, we model θ in the unary and binary terms at an individual and pairwise level, independent of the structure. Then, for any plausible layout of parts, we can immediately index into their respective parameters and use them for inference.

In this paper, we do not define an explicit form on how Ω is distributed. Rather, we enumerate a plausible set of structures and tie one to each image, but in the general case, PS3 samples from Ω . In spirit, this notion of parts-sparse pictorial structures has appeared in (Hess et al., 2007). Their mixture-of-parts pictorial structures model has similarly relaxed the assumption that the full set of parts needs to appear. One can indeed use this mixture distribution on Ω . See Section 6 for more comparative discussion.

Standard Form of PS3. The terms of the energy function underlying PS3 operate on functions of the parts $\phi(\cdot)$ and $\psi(\cdot)$ rather than the parts directly. These functions are arbitrary and depend on how the potentials will be modeled (we specify exact definitions in the next section). The standard form of a PS3 from Ω is

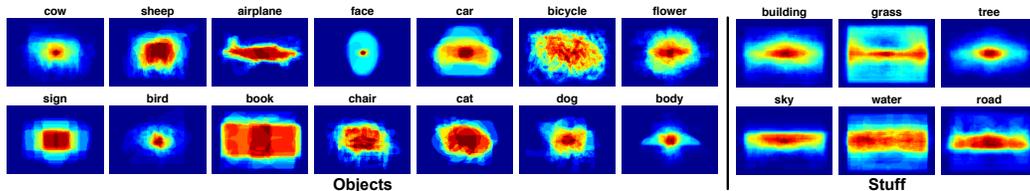
$$H(L|I, \theta) = \left(\sum_{i=1}^n m(\phi(l_i)|\theta) + \sum_{e_{ij}} d(\psi(l_i), \psi(l_j)|\theta) \right). \quad (4)$$

4.1. Model Details for Semantic Image Labeling

Unary Term. The unary term will capture the appearance, shape and the location of the parts:

$$\begin{aligned} m(\phi(l_i)|\theta) = & \alpha_A m_A(l_i|\theta) + && \leftarrow \text{appearance,} \\ & \alpha_S m_S(l_i|\theta) + && \leftarrow \text{shape,} \\ & \alpha_L m_L(\mu_i|\theta) && \leftarrow \text{location.} \end{aligned} \quad (5)$$

Figure 1: Visual rendering of some of the shape models on MSRC data. Each image shows the map B_i that is centered around the part centroid and normalized to a unit-box coordinate system. The images are rendered using the *jet* colormap in Matlab. The figure has been broken into objects (left) and stuff (right) to emphasize the disparity in expressiveness between the shape maps for the two part types.



The α_i are coefficients on each term. The ϕ function maps the pixel support part l_i to the pair (l_i, μ_i) , where μ_i is the centroid of l_i : $\mu_i \doteq \frac{1}{|l_i|} \sum_{\lambda \in l_i} \sum_{\mathbf{x} \in \lambda} \mathbf{x}$. The terms of (5) are described next.

Appearance. We model appearance in four-dimensions: Lab color-space and texton space. The texton maps use a 61-channel filter bank (Varma and Zisserman, 2005) of combined Leung-Malik (Leung and Malik, 2001) and Schmid filters (Schmid, 2001), followed by a k-Means process with 64 cluster centers. No experimentation was performed to optimize the texton codebook filter and size choice. We compute histograms for the foreground and class-specific background (computed in a narrowband of pixels around the boundary of the part). The appearance potential is computed as the ratio of the cross-element (foreground to background) histogram intersection to the internal element (e.g., foreground model to foreground evidence) intersection.

Shape. The capability to model global part shape is a key unique feature of the PS3 model. We model it nonparametrically using a kernel density estimator and quantize the density as a discrete map of 201×201 normalized pixels; call this map S_{z_i} for class z and part l_i . Recall, a part l_i induces a binary membership map B_i at the (normalized) pixel level. Finally, the shape potential is defined as the mean shape probability over the part’s constituent pixels:

$$m_S(l_i|\theta) = -\log \left(\frac{1}{|\mathbf{w}|} \|B_i \odot S_{z_i}\|_F + \|(1 - B_i) \odot (1 - S_{z_i})\|_F \right) \quad (6)$$

where \odot is the element-wise product and the Frobenius norm is used.

Example shape densities are shown in Figure 1. There is a clear distinction in the expressiveness of this shape model between the objects (e.g., airplane, face) and the stuff (e.g., sky, road). The maps for the stuff classes tend to be diffuse and indiscriminate, whereas the maps for the object classes are mostly recognizable. Section 6 shows that this object-level modeling significantly aids in labeling of the object-type classes.

Location. We model the part location with a Gaussian distribution on its centroid μ_i and hence use the Mahalanobis distance for the location potential: $m_L(\mu_i|\theta) = (\mu_i - \nu_z)^T \Sigma^{-1} (\mu_i - \nu_z)$, where ν_z and Σ_z are the mean centroid and covariance matrix for class z respectively.

Binary Term. The binary term will capture the relative distance and angle of pairwise connected parts:

$$d(\psi(l_i), \psi(l_j)|\theta) = \alpha_D d_D(\mu_i, \mu_j|\theta) + \leftarrow \text{distance,} \\ \alpha_R d_R(\mu_i, \mu_j|\theta) \leftarrow \text{angle.} \quad (7)$$

The α are again coefficients on each term. The ψ function maps the pixel support part l_i to the μ_i , is the centroid of l_i . More sophisticated ϕ and ψ functions are plausible with the PS3 framework, but we do not explore them in this paper.

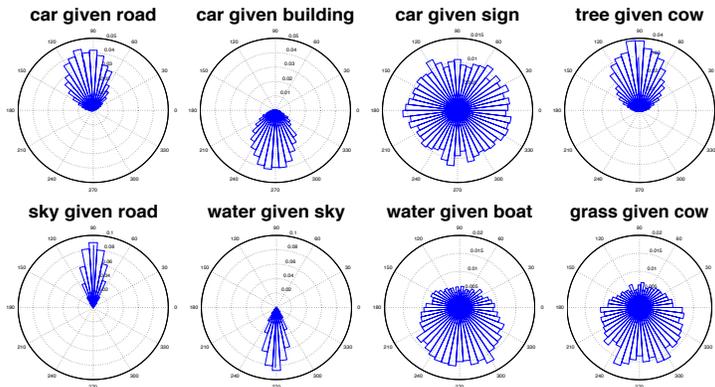
Distance. The relative part distance is captured simply by the distance between the parts (classical pictorial structures). For parts l_i and l_j , we evaluate the distance $v_{ij} = \|\mu_i - \mu_j\|_2$ and model it by a Gaussian parameterized by $(\nu_{ij}, \sigma_{ij}^2)$. The distance potential is $d_D(\mu_i, \mu_j|\theta) = (v_{ij} - \nu_{ij})^2 / \sigma_{ij}^2$.

Angle. We model the relative angle between the two parts by a von Mises distribution, which is a symmetric and unimodal distribution on the circle (Berens, 2009). Let r_{ij} denote the angle relating part i with respect to part j . The angle potential is the negative log, $d_R(\mu_i, \mu_j|\theta) = -\log P(r_{ij}|\theta, z_i, z_j)$, of the von Mises density. Some examples are presented in Figure 2. These examples suggest this angle potential is jointly useful for the objects and the shape, especially how they inter-relate. For example, consider the rightmost plots of tree-given-cow: in the MSRC dataset, cows appear in (or on) pasture nearly always and there are often trees on the horizon.

4.2. Learning the PS3 Model

For a training set of images $\{I_1, \dots, I_N\}$ and corresponding configurations $\{L_1, \dots, L_N\}$, which are essentially pixel-wise image labelings, learning the parameters is cast as a maximum likelihood (MLE) problem. (Felzenszwalb

Figure 2: Samples of the angle distribution.



and Huttenlocher, 2005) show that the parameters θ on the unary potentials can be learned independently, and this holds for our pixel-support parts. In our case, we do not seek to learn a tree-structured graph. Instead, we define a deterministic mapping from a label image, or configuration, L_i to a graph G in the following manner: for each connected component in L_i create a part in the graph. Two parts are adjacent in the graph if any pixel in their respective connected components are adjacent in the label image L_i . It is our assumption that this general structure adds necessary descriptiveness for the labeling problem. Finally, for each pair of adjacent parts, we learn the parameters on the binary potentials via MLE.

The last part of learning is to estimate the five α weights on the various potentials. It is widely known that estimating these weights is a significant problem as it requires estimation of the full partition function, which is intractable (Winkler, 2006). However, in our case, the problem is compounded, even some standard approximations like pseudo-likelihood are intractable because of the pixel-support nature of the parts. Because of these complexities, we simply set the coefficients such that the relative scale of each potential is normalized and finally we ensure the weights specify a convex combination over the potentials.

5. Inference with Data-Adaptive MCMC

Inference with the PS3 model has two main components: (1) determine the structure of the PS3 for the image at hand, and (2) determine the optimal configuration L^* given a structure. In this paper, we study the latter and

leave the former for future work. Although this limits the generality of the contributions proposed herein to cases in which a suitable structure for the PS3 could be determined or given, we show that even the determination of the optimal configuration alone is a significant problem. Furthermore, direct extensions of the proposed methods present viable options for handling the structure inference, as we will discuss. The configuration inference problem is posed as

$$L^* = \arg \max_L \exp[-H(L|I, \theta)] . \quad (8)$$

The corresponding energy minimization problem is $\arg \min_L H(L|I, \theta)$. In general, this problem is NP-hard, but seems similar to the standard form for which our community has hotly studied approximate solutions over the past decade (Szeliski et al., 2006). However, as noted by (Felzenszwalb and Huttenlocher, 2005), the structure of the graph and the space of possible solutions differ substantially; in other words, the minimization problem cannot be cast as a local labeling problem, as we know explain in more detail.

Consider the variables in question, $L = \{l_1, \dots, l_n\}$; we adopt a superpixel specification of the elements $\{\lambda\}$, computed via (Felzenszwalb and Huttenlocher, 2004), in our implementation. We already know the class z_i of each part and each l_i has a complex interrelationship to the other parts via its pixel support. For example, taking one element λ away from l_i and moving it to l_j has part-global effects on both l_i and l_j in terms of appearance and shape, which differs quite drastically from these prior methods. One could consider defining the PS3 inference as a labeling problem over the elements $\{\lambda\}$ with each part l_i being a labeling index and associating a label variable, say ξ_j , with each element λ_j . However, inference would remain outside of the scope of these methods, again because a local change of one label variables ξ_j would have a far-reaching affect on many other (potentially all) elements $\{\lambda_k : \xi_j \equiv \xi_k\}$.

In addition, classical pictorial structures use parametric representations of l_i , such as part-centroid, and for the typical spring-model case, define a Mahalonobis distance to capture the ideal relative location between parts. Casting our nonparametric form l_i into this framework would yield an intractable high-dimension problem: even though we rely on parametric functions of l_i for our binary potentials (7) no convenient form of the ideal location is possible since the l_i are tied directly to the pixel support.

MCMC Sampler. We hence adopt a Metropolis-Hastings (MH) approach

to handle inference (Andrieu et al., 2003), as it is straightforward to implement and known to converge to the underlying invariant distribution $P(L|I, \theta)$ as long as the Markov chain satisfies detailed balance (Winkler, 2006). For space, we do not go into the details of the MCMC sampler other than to note that the proposal distribution is specified to (1) uniformly sample from each of the parts and (2) sample elements in proportion to how well their appearance fits the proposed change.

Data-Adaptive Simulated Annealing. We embed the MH sampler into a simulated annealing process (Geman and Geman, 1984) because we seek the maximum a posteriori samples. Simulated annealing adds a temperature parameter T into the distribution, $P^{(T)}(L|I, \theta) = \frac{1}{Z} \exp[-\frac{1}{T}H(L|I, \theta)]$, such that as $T \rightarrow 0$ the $P^{(T)}(L|I, \theta)$ distribution approaches the modes of $P(L|I, \theta)$. However, the theoretical guarantee exists on fairly restrictive bounds on the *cooling schedule*, the sequence of temperatures as the process is cooled (Andrieu et al., 2003). Furthermore, it is not well understood how to set the cooling schedule in practice, especially for very high-dimensional sample spaces, such as the one at hand. The challenge is that one proposed move L' will change the density quite little resulting in acceptance probabilities near uniform unless the cooling schedule is set just right.

To resolve this issue, we propose an principled approach to set the cooling schedule that adapts to each image at hand and requires no manual tweaking. The basic idea is to directly estimate a map from desired acceptance probabilities to the required temperatures. Denote γ as shorthand for the acceptance probability. Disregarding the proposal distribution, consider γ written directly in terms of the amount ρ of energy the proposed move would make:

$$\gamma = \frac{\exp[-H(L')/T]}{\exp[-H(L^t)/T]} = \frac{\exp[-(H(L^t) + \rho)/T]}{\exp[-H(L^t)/T]} \quad (9)$$

For a specific desired γ value and known ρ , we can solve (9) for $T = -\frac{\rho}{\ln \gamma}$ making it possible to adapt the simulated annealing cooling schedule to each image in a principled manner, rather than manually tuning parameters by hand. Before beginning the annealing process, we sample $P(L|I, \theta)$ to estimate the ρ for the image. Assuming a linear gradient of desired acceptance ratios, the only part that needs to be manually set is the acceptance probability range, γ_1, γ_2 , which we set to 0.9 and 0.1 respectively to cover most of the range of acceptance probabilities but never making them guaranteed or impossible.

6. Results and Discussion

We use two pixel-labeling benchmark datasets for our experimental analysis: MSRC (Shotton et al., 2009) and SIFT-Flow (Liu et al., 2009). In brief, MSRC is a 21-class 596-image dataset and SIFT-Flow is a 33-class 2688-image dataset, both of typical natural photos. The gold standard for these data is set by manual human labeling and most images have a large percentage of pixels actually labeled in the gold standard. In both cases, we use the authors’ training-testing splits for learning and evaluation; we note the split is 55% training for MSRC and 90% training for SIFT-Flow. Finally, in the authors’ split for SIFT-Flow, three classes (cow, desert, and moon) do not appear and are hence dropped, yielding a 30-class dataset in actuality.

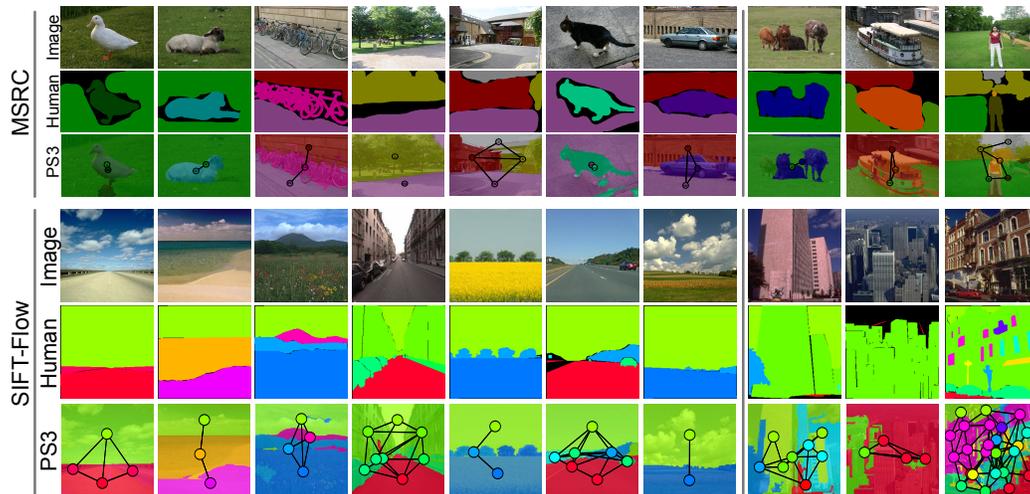
Figure 3: Quantitative results on the two data sets, comparing the MLE classifier, the MRF model over the superpixel basic elements, and our proposed PS3. The table shows % pixel accuracy $N_{ii}/\sum_j N_{ij}$ for different object classes. “Global” refers to the overall error $\frac{\sum_{i \in \mathcal{Z}} N_{ii}}{\sum_{i,j \in \mathcal{Z}} N_{ij}}$, while “average” is $\sum_{i \in \mathcal{Z}} \frac{N_{ii}}{|\mathcal{Z}| \sum_{j \in \mathcal{Z}} N_{ij}}$. N_{ij} refers to the number of pixels of label i labelled j . Note the color in the table columns is intended to serve as the legend for Figure 4.

		MSRC																															
		Global	Average	Building	Grass	Tree	Cow	Sheep	Sky	Aeroplane	Water	Face	Car	Bicycle	Flower	Sign	Bird	Book	Chair	Road	Cat	Dog	Body	Boat									
SIFT-Flow	MLE	79	74	41	96	73	72	73	93	44	79	84	58	72	98	82	78	94	90	89	81	83	49	39									
	MRF	79	75	42	96	74	72	72	93	45	79	85	58	74	98	82	77	94	91	90	81	83	49	37									
	PS3	81	78	74	90	72	72	69	93	65	71	70	70	78	99	83	79	99	90	77	84	88	59	56									
		SIFT-Flow																															
		Global	Average	Awning	Balcony	Bird	Boat	Bridge	Building	Bus	Car	Crosswalk	Door	Fence	Field	Grass	Mountain	Person	Plant	Pole	River	Road	Rock	Sand	Sea	Sidewalk	Sign	Sky	Staircase	Streetlight	Sun	Tree	Window
SIFT-Flow	MLE	56	42	12	51	0	19	35	32	33	34	93	23	46	54	77	50	39	40	21	31	62	44	67	42	44	16	84	26	2.8	100	58	26
	MRF	57	43	12	65	0	18	54	34	26	39	94	23	47	55	79	51	37	44	14	29	63	45	67	45	43	15	84	25	1.8	100	58	24
	PS3	66	55	31	81	2	35	33	41	69	59	39	62	33	73	58	57	60	65	60	67	69	67	78	61	43	43	90	56	34	35	68	78

6.1. Comparisons to Baselines

Our primary evaluation goal is to determine and quantify the benefit gained through the global parts-based structure in a controlled setting. Hence we make a quantitative comparison of our method against an MLE classifier and an MRF, with results shown in Figure 3. We do not make a comparison to other pictorial structures papers as, to the best of our knowledge, no existing pictorial structures method can be directly applied to the pixel labeling problem. To establish a rigorously controlled experimental setting, in all cases, we use the same appearance models and assume full knowledge of the graph structure (for the MLE and MRF methods, this means the

Figure 4: Visual results on the two data sets. Each column shows an example in three rows: (1) original image, (2) human gold standard, and (3) our PS3 result overlaid upon the image. We have also rendered the graph structure on top of the image. The color legend is given in Figure 3. The results on the right side of the figure show some of the worst examples of our performance.



subset of possible class labels) for each test image. In the MLE case, the basic elements are assumed independent, and in the MRF case, a local Potts smoothness model is used between the basic elements. In all three cases, the only aspect that is varied is how the basic elements (superpixels) are related to the overall model; i.e., we are able to evaluate exactly how much the PS3 model improves upon local MLE and MRF models. We note that our proposed PS3 model is also an MRF, but over the global scene parts and not over the superpixel basic elements. Our quantitative scores hence need to be interpreted as relative among the three methods we have displayed. In this setting, it is clearly demonstrated that the proposed model outperforms both the superpixel-independent MLE classifier and the locally connected MRF model.

Our proposed method performs best in global and average per-class labeling accuracy over the pixel-independent MLE and local-MRF methods on both datasets. On MSRC we see a gain of 2% in global and 3% in average accuracy. These are not significant numbers, overall, but we note the significant improvement in two subsets of the classes. First, in classes with

high intraclass variance, such as building, we see a 30+% increase. Second, in classes with strong global object shape, such as airplane, we see a 20% increase. These exhibit the merits the global modeling of PS3 brings to the problem. The reason why the overall gain is not too much is that the dominant classes, such as sky, grass, and so on, have a strong visual character in the relatively small MSRC data set that is already easily modeled at the local level.

We find a different case in the SIFT-Flow dataset, which is much larger and contains more intra-class variance even for these dominant classes. In the SIFT-Flow cases, a larger increase of 9% in global and 12% in average accuracy is observed. We note the marked improvement in some of the categories corresponding to *things*, such as airplane, car, door, and person. We explain this improvement as being due to the added modeling richness in the parts-based representation: *things* in the image benefit from the rich global description through the shape and part-context. We also note the few categories in SIFT-Flow where PS3 was outperformed by the MLE and MRF methods (bridge, crosswalk, fence, and sun). In these cases, the typical object foreground is *sparse* and the global part description is insufficient to accommodate the weak local cues, which get overtaken by the other nearby classes. Examples of this phenomena (as well as good cases) are given in Figure 4.

6.2. Comparisons to State of the Art

We also make a quantitative (Figure 5) against a range of papers from the state of the art, TextonBoost (Shotton et al., 2009), Mean-Shift Patches (Yang et al., 2007), Graph-Shifts (Corso et al., 2008), TextonForests (Shotton et al., 2008), and Hierarchical CRF (H-CRF) (Ladicky et al., 2009). Nearly all of these papers can be classes within the “local” labeling realm. The state-of-the-art H-CRF approach in (Ladicky et al., 2009) makes a clever extension to define a hierarchy of random fields that has shown great potential to overcome the limitations of purely local labeling methods. However, it still defines the labeling problem based directly on local interactions of label variables rather than on object level interactions, as we do in PS3. None of the existing pictorial structures papers we are aware of can be directly applied to semantic image labeling and are hence not compared here.

Our proposed PS3 method performs best in average per-class labeling accuracy (78%) and shows marked improvement in numerous classes, such as flower, bird, chair, etc. We make careful note that although the table directly

Figure 5: Quantitative results on the MSRC data set, in the same format as (Ladicky et al., 2009) for easy comparison. The table shows % pixel accuracy $N_{ii}/\sum_j N_{ij}$ for different object classes. “Global” refers to the overall error $\frac{\sum_{i \in \mathcal{Z}} N_{ii}}{\sum_{i,j \in \mathcal{Z}} N_{ij}}$, while “average” is $\sum_{i \in \mathcal{Z}} \frac{N_{ii}}{|\mathcal{Z}| \sum_{j \in \mathcal{Z}} N_{ij}}$. N_{ij} refers to the number of pixels of label i labelled j .

	Global	Average	Building	Grass	Tree	Cow	Sheep	Sky	Aeroplane	Water	Face	Car	Bicycle	Flower	Sign	Bird	Book	Chair	Road	Cat	Dog	Body	Boat
TextonBoost	72	58	62	98	86	58	50	83	60	53	74	63	75	63	35	19	92	15	86	54	19	62	7
Mean-Shift Patches	75	62	63	98	89	66	54	86	63	71	83	71	79	71	38	23	88	23	88	33	34	43	32
Graph-Shifts	77	66	72	93	85	66	70	93	86	69	85	66	85	58	64	25	91	44	74	41	30	51	43
TextonForests	72	67	49	88	79	97	97	78	82	54	87	74	72	74	36	24	93	51	78	75	35	66	18
Hierarchical CRF	86	75	80	96	86	74	87	99	74	87	86	87	82	97	95	30	86	31	95	51	69	66	9
PS3	81	78	74	90	72	72	69	93	65	71	70	70	78	99	83	79	99	90	77	84	88	59	56

compares PS3 to the other literature, we assume the graph structure for each testing image is known. Our findings hence need to be taken with caution: the assumption on knowing the graph structure that we have made limits the significance and comparability of our proposed method against these others. Notwithstanding this point, we do feel it is important to demonstrate the comparative performance against the state of the art. Furthermore, we note that our unary potentials are comparatively simpler (i.e., color and texton histograms) to those in many of the other methods. Finally, knowing the appropriate graph for the image does not immediately solve the problem: e.g., as Figure 3 points out, on MSRC, an MLE assignment of superpixels to elements in the graph yields global accuracy of 74% and average accuracy of 70%—i.e., the PS3 model is indeed adding power to the problem.

Separating Objects from Stuff. The respective merits of the two top performing approaches, namely H-CRF and ours, PS3, become immediately evident when inspecting how the methods compare on various classes, as we discuss next. As we mentioned earlier, one can group the parts roughly into two types: objects (cow, sheep, aeroplane, face, car, bicycle, flower, sign, bird, chair, road, cat, dog, body, and boat) and stuff (building, grass, tree, sky, water and road)

The objects tend to be small and articulated and have high location variance, whereas the stuff tends to be relatively stable in terms of location and appearance distribution. As we have shown in Section 4.1, the shape distributions of the stuff are uninformative, but for the objects, they are quite informative. We have claimed that a key merit of our method is that it allow the modeler to emphasize global object shape and relationship to

Figure 6: Quantitative results when we group objects and stuff on the MSRC data set. The scores in this table are all average accuracy derived from the per-class accuracy scores from Figure 5; no further processing was performed. See text for discussion and the list of classes for the two groups.

	Objects	Stuff
TextonBoost	50	78
Mean-Shift Patches	53	83
Graph-Shifts	60	81
TextonForests	65	71
Hierarchical CRF	68	91
PS3	77	80

the scene in general. This claim is clearly substantiated when looking at the comparative average accuracy of the objects to the stuff in Figures 5 and 6. We explain it via the components of the PS3 model as follows: our method performs at about the average performance for the stuff classes, which are comparatively easier to infer using location and appearance. Subsequently, these stuff classes are grounded and drive the object classes during inference allowing them to utilize the objects’ richer shape and angle potentials.

6.3. Methodological and Computational Comparisons

Parts-based methods have been the topic of heavy inquiry since their inception (Fischler and Elschlager, 1973). In this section, we discuss the proposed PS3 model against four relevant examples in the literature on the basis of methodological and computational comparisons; the computational comparison is made at the end of the section.

Comparison to DDMCMC (Tu and Zhu, 2002). The seminal DDMCMC work laid the groundwork for our approach to inference in this paper, but the underlying problem and model are quite different. Firstly, the DDMCMC work is an approach to the low-level image segmentation problem. No notion of object class or global object is incorporated into that work, which, as our quantitative results demonstrated is a significant merit of our proposed approach. Secondly, it is primarily seeking samples of image segmentations under all plausible partitions of the image. We have restricted ours to the set of superpixels, but we can plausibly relax our assumption. Lastly, their work did not seem to seek the modes, whereas we propose a data-adaptive method for mode seeking in the MCMC framework.

Comparison to the constellation and star models (Fergus et al., 2007, 2005). This class of methods semi-supervisedly learn translation and

scale-invariant models for the parts of various object categories under specific topological arrangements based on available bottom-up interest point detections. The learning process is expectation-maximization based and automatically selects the parts comprising the model during learning; in the constellation model (fully connected) learning is $O(N^P)$ for N point detections and P model parts and for the star (root-connected) learning is $O(N^2P)$. Detection can be improved to $O(NP)$ using distance transforms. These models, however, are inherently incomparable to the proposed PS3 model because they learn an object-specific model with a fixed number of parts; in the semantic pixel labeling problem studied herein, we focus on scene models in which only a few parts of those possible will appear in the image. Despite its success on specific objects, directly learning a constellation model for semantic pixel labeling is hence not plausible.

Comparison to Deformable Parts Models (DPM) (Felzenszwalb et al., 2010). The state of the art in parts-based models of objects is the DPM, per its broad success on challenges like the PASCAL VOC (Everingham et al., 2010). These models use a latent SVM learning procedure to learn a two-level hierarchy of object-specific part-filters capable of recognizing articulated objects. The natural successors of the earlier constellation and star models, these methods bring more power by maximizing over part-parameters during inference (such as spatial location of the part) but also substantially more cost. Nevertheless, these methods are similarly not designed to deal with the various configurations of different part-types in a single model, as we expect in the semantic pixel labeling case.

Comparison to Mixture-of-Parts Pictorial Structures (MoPPS) (Hess et al., 2007). As far as we know, MoPPS is the first and only other pictorial structures extension to permit part subsets. Like our method, they permit a space of plausible pictorial structures. Then, the MoPPS method carefully specifies a mixture distribution over parts, a set of legal part configurations and a mechanism for returning a pictorial structure given a part subset. In the spirit of sparse-parts pictorial structures, PS3 is similar to MoPPS. But MoPPS remains restricted to the object modeling case. For a given parts subset, the MoPPS structure is classical (Gaussian spring model) whereas our part potential incorporate a more rich set of relations.

Computational Comparisons. The core challenge in PS3 is that we need to address the space Ω of all part-models for scenes, in which some or many may be present in any one instance of PS3. Because of this complexity, we have made an assumption of stationarity: for P parts, we learn the unary

terms separately $O(P)$ and binary pairwise terms $O(P^2)$. These learning costs are comparatively low because we are not selecting optimal parts for a specific object/scene (our models are hence “weaker” in a sense), but this is necessary for the broader problem at hand.

Inference, however, is more expensive than learning in PS3 because we are optimizing over the combinatoric space of associating elements with each graph part and no dynamic programming or distance transforms methods are available. The Metropolis-Hasting stochastic optimizer is well-known to be a costly process; if we fix the number of sweeps to a constant number (say 10, which we have found to work satisfactorily in practice), we can approximate the complexity to $O(N^2P^2)$ for N superpixels and P parts in the graph structure for a given image. The N^2 term results from each superpixel getting visited once during each sweep and possibly affecting a change globally over the graph, and the P^2 term comes from the binary interaction term as we test the membership of different superpixels. We find it converges on average in about a minute of processing the images in this paper (running sequentially on a standard 2GHz PC in Matlab). Comparatively, in the case the graph structure is known (which is also an assumption made in standard parts-based methods) the PS3 method has a similar complexity to the other parts-based methods, with an extra P term that is typically much smaller than N . Typical running times also reflect this similarity across the methods.

7. Conclusion

We have presented the pixel-support, parts-sparse pictorial structures, or PS3 model. PS3 makes a step in scene labeling by moving beyond the de facto local and region based approaches to full semantic scene labeling and into a rich object-level approach that remains directly tied to the pixel level. As such, PS3 unifies parts-based object models and scene-based labeling models in one common methodology. Our experimental comparisons demonstrate the merits in moving beyond the local methods in a number of settings on benchmark data sets (MSRC and Sift Flow).

PS3 has, perhaps, opened more problems than it has solved, however. For example, we have assumed that the graph for an image is known during inference. For general applicability, this assumption needs to be relaxed. Extensions of the proposed MCMC methods into jump-diffusion dynamics (Tu and Zhu, 2002) are plausible, but some approximations or other methods to *marginalize* the full sample-space are also plausible. Probabilistic ontologies

and Markov logic present two potential avenues for this problem. Similarly, we have demonstrated that the parameter estimation problem in the PS3 is more complex given the global-local part-pixel dependency. We are not aware of a principled tractable method for estimating these parameters. Finally, we have observed a big disparity in the respective strength of our various model terms for object- and stuff-type classes, but we have not incorporated this distinction into the model itself.

Acknowledgements

We are grateful for the support in part provided through the following grants: NSF CAREER IIS-0845282, ARO YIP W911NF-11-1-0090, DARPA Minds Eye W911NF-10-2-0062, DARPA CSSG HR0011-09-1-0022. Findings are those of the authors and do not reflect the views of the funding agencies.

References

- Andrieu, C., Freitas, N., Doucet, A., Jordan, M., 2003. Introduction to MCMC for Machine Learning. *Machine Learning* 50, 5–43.
- Berens, P., 2009. Circstat: A matlab toolbox for circular statistics. *Journal of Statistical Software* 31 (10).
- Biederman, I., 1981. On the Semantics of a Glance at a Scene. In: Kubovy, M., Pomerantz, K. R. (Eds.), *Perceptual Organization*. Lawrence Erlbaum Publisher, pp. 213–263.
- Carson, C., Belongie, S., Greenspan, H., Malik, J., 2002. Blobworld: Image Segmentation using Expectation-Maximization and Its Application to Image Querying. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (8), 1026–1038.
- Corso, J. J., Yuille, A., Tu, Z., 2008. Graph-Shifts: Natural Image Labeling by Dynamic Hierarchical Computing. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*.
- Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., Zisserman, A., 2010. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision* 88 (2), 303–338.

- Felzenszwalb, P. F., Girshick, R. B., McAllester, D., Ramanan, D., 2010. Object detection with discriminatively trained part based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 1627–1645. URL <http://people.cs.uchicago.edu/~pff/latent/>
- Felzenszwalb, P. F., Huttenlocher, D. P., 2004. Efficient Graph-Based Image Segmentation. *International Journal of Computer Vision* 59 (2), 167–181.
- Felzenszwalb, P. F., Huttenlocher, D. P., 2005. Pictorial structures for object recognition. *International Journal of Computer Vision* 61 (1), 55–79.
- Fergus, R., Perona, P., Zisserman, A., 2005. A sparse object category model for efficient learning and exhaustive recognition. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*.
- Fergus, R., Perona, P., Zisserman, A., 2007. Weakly supervised scale-invariant learning of models for visual recognition. *International Journal of Computer Vision* 73 (3), 273–303.
- Fischler, M. A., Elschlager, R. A., 1973. The representation and matching of pictorial structures. *IEEE Transactions on Computers* C-22 (1), 67–92.
- Geman, S., Geman, D., 1984. Stochastic Relaxation, Gibbs Distributions, and Bayesian Restoration of Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6, 721–741.
- Gould, S., Rodgers, J., Cohen, D., Elidan, G., Koller, D., 2008. Multi-class segmentation with relative location prior. *International Journal of Computer Vision* 80 (3), 300–316.
- Heitz, G., Gould, S., Saxena, A., Koller, D., 2008. Cascaded classification models: Combining models for holistic scene understanding. In: *Proceedings of Neural Information Processing Systems (NIPS)*.
- Hess, R., Fern, A., Mortensen, E. N., 2007. Mixture-of-parts pictorial structures for objects with variable part sets. In: *Proceedings of International Conference on Computer Vision*.
- Hoiem, D., Efros, A. A., Hebert, M., 2008. Putting objects in perspective. *International Journal of Computer Vision* 80 (1), 3–15.

- Ladicky, L., Russell, C., Kohli, P., Torr, P. H. S., 2009. Associative hierarchical crfs for object class image segmentation. In: Proceedings of International Conference on Computer Vision.
- Leung, T., Malik, J., 2001. Representing and recognizing the visual appearance of materials using three-dimensional textons. *International Journal of Computer Vision* 43 (1), 29–44.
- Li-Jia, L., Socher, R., Fei-Fei, L., 2009. Towards total scene understanding: Classification, annotation and segmentation in an automatic framework. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition.
- Liu, C., Yuen, J., Torralba, A., 2009. Nonparametric scene parsing: Label transfer via dense scene alignment. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition.
- Schmid, C., 2001. Constructing Models for Content-Based Image Retrieval. In: IEEE Conference on Computer Vision and Pattern Recognition. Vol. 2. pp. 96–101.
- Shotton, J., Johnson, M., Cipolla, R., 2008. Semantic texton forests for image categorization and segmentation. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition.
- Shotton, J., Winn, J., Rother, C., Criminisi, A., 2009. TextonBoost for image understanding: Multi-class object recognition and segmentation by jointly modeling texture, layout, and context. *International Journal of Computer Vision* 81 (2), 2–23.
- Szeliski, R., Zabih, R., Scharstein, D., Veksler, O., Kolmogorov, V., Agarwala, A., Tappen, M., Rother, C., 2006. A Comparative Study of Energy Minimization Methods for Markov Random Fields. In: Ninth European Conference on Computer Vision (ECCV 2006). Vol. 2. pp. 16–29.
- Tu, Z., Zhu, S. C., 2002. Image Segmentation by Data-Driven Markov Chain Monte Carlo. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (5), 657–673.

- Varma, M., Zisserman, A., 2005. A statistical approach to texture classification from single images. *International Journal of Computer Vision* 62 (1–2), 61–81.
- Winkler, G., 2006. *Image Analysis, Random Fields, and Markov Chain Monte Carlo Methods*, 2nd Edition. Springer.
- Winn, J., Shotton, J., 2006. The layout consistent random field for recognizing and segmenting partially occluded objects. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*.
- Yang, L., Meer, P., Foran, D. J., 2007. Multiple class segmentation using a unified framework over mean-shift patches. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*.
- Yang, Y., Hallman, S., Ramanan, D., Fowlkes, C., 2010. Layered object detection for multi-class segmentation. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*.