# Associative Hierarchical Random Fields

Ľubor Ladický[1], Chris Russell[1], Pushmeet Kohli, Philip H. S. Torr

**Abstract**—This paper makes two contributions: the first is the proposal of a new model – the associative hierarchical random field (AHRF), and a novel algorithm for its optimisation; the second is the application of this model to the problem of semantic segmentation.

Most methods for semantic segmentation are formulated as a labelling problem for variables that might correspond to either pixels or segments such as super-pixels. It is well known that the generation of super pixel segmentations is not unique. This has motivated many researchers to use multiple super pixel segmentations for problems such as semantic segmentation or single view reconstruction. These super-pixels have not yet been combined in a principled manner, this is a difficult problem, as they may overlap, or be nested in such a way that the segmentations form a segmentation tree. Our new hierarchical random field model allows information from all of the multiple segmentations to contribute to a global energy. MAP inference in this model can be performed efficiently using powerful graph cut based move making algorithms.

Our framework generalises much of the previous work based on pixels or segments, and the resulting labellings can be viewed both as a detailed segmentation at the pixel level, or at the other extreme, as a segment selector that pieces together a solution like a jigsaw, selecting the best segments from different segmentations as pieces. We evaluate its performance on some of the most challenging data sets for object class segmentation, and show that this ability to perform inference using multiple overlapping segmentations leads to state-of-the-art results.

**Index Terms**—Conditional Random Fields, Discrete Energy Minimisation, Object Recognition and Segmentation.

## 1 Introduction

Many problems of image understanding can be formulated as semantic segmentation, or the assignment of a 'semantic class' label – such as *person*, or *road*, to every pixel in the image. Until recently, image labelling problems have been formulated using pairwise Markov random fields. However, these pairwise fields are unable to capture the higher-order statistics of natural images which can be used to enforce the coherence of regions in the image or to encourage particular regions to belong to a certain class. Despite these limitations, the use of pairwise Markov models is prevalent in vision. This can largely be attributed to the pragmatism of computer vision researchers; although such models do not fully capture image statistics, they serve as an effective discriminative model that prevents individual pixels from being mislabelled. Moreover, unlike many optimisation approaches for higher-order models, approximation algorithms exist for pairwise models, that are guaranteed to find a solution whose cost must lie within a fixed bound of the cost of global optima.


- *Ľubor Ladický is with the Department of Engineering Science, University of Oxford, Oxford, U.K.*
- *Chris Russell is with Queen Mary College, University of London, London, U.K.*
- *Pushmeet Kohli is with the Microsoft Research, Cambridge, U.K.*
- *Philip H. S. Torr is with the Department of Computing, Oxford Brookes University, Oxford, U.K.*

1. The authors assert equal contribution and joint first authorship


This work is supported by EPSRC research grants, HMGCC, the IST Programme of the European Community, under the PASCAL2 Network of Excellence, IST-2007-216886. P. H. S. Torr is in receipt of Royal Society Wolfson Research Merit Award.

Many approaches to multi-scale vision have been proposed where either inference is performed in a 'top-down' approach, i.e. an initial scene based estimates is made, followed by the successive labelling of smaller regions that must be consistent with the initial labelling; or they take a bottom-up approach where starting from the labelling of small regions they progressively assign larger regions a labelling consistent with the small regions.

In this work we propose a novel formulation for multi-scale vision, as captured by a hierarchy of segmentations. This allows for the integration of cues defined at any scale, or over any arbitrary region of the image, and provides a generalisation of many of the segmentation methods prevalent in vision. Our approach provides a unification of the 'top-down' and 'bottom-up' approaches to common to many problems of computer vision. To do this, we propose a new model: the associative hierarchical random field (AHRF) and show how it can be solved efficiently using graph-cut based move-making algorithms.

If two AHRFs are added together to produce a new cost, the resulting cost is also an AHRF, and the summand can also be solved effectively. This allows many different potentials and cues to be incorporated within the model, and for inference to remain practical. This flexibility contributes substantially to the success of our method, and allows us to obtain state of the art results on most existing segmentation data-sets. Thus this work contributes both to the general area of random fields, and in its application to computer vision. Next, we shall give some background to the problem of semantic segmentation in computer vision.

### 1.1 Semantic Segmentation

Over the last few years many different methods have been proposed for semantic segmentation *i.e.* the problem of assigning a set of given object labels such as *person*, *car*, or *road* to each pixel of a given image, in a manner consistent with human annotations.

Most methods for semantic segmentation are formulated in terms of pixels (Shotton et al, 2006), other methods used segments (Batra et al, 2008; Galleguillos et al, 2008; Yang et al, 2007), groups of segments (Rabinovich et al, 2007), or the intersections of multiple segmentations (Pantofaru et al, 2008), while some have gone to the extreme of looking at the whole image in order to reason about object segmentation (Larlus and Jurie, 2008).

Each choice of image representation comes with its share of advantages and disadvantages. Pixels might be considered the most obvious choice of segmentation. However, pixels by themselves contain a limited amount of information. The colour and intensity of a lone pixel is often not enough to determine its correct object label. Ren and Malik (2003)'s remark that *'pixels are not natural entities; they are merely a consequence of the discrete representation of images'* captures some of the problems of pixel-based representations.

The last few years have seen a proliferation of super-pixel (Comaniciu and Meer, 2002; Felzenszwalb and Huttenlocher, 2004; Shi and Malik, 2000) based methods, that perform an initial *a priori* segmentation of the image, applied to object segmentation (Batra et al, 2008; Galleguillos et al, 2008; He et al, 2006; Russell et al, 2006; Yang et al, 2007), and elsewhere (Hoiem et al, 2005; Tao et al, 2001). These rely upon an initial partitioning of the image, typically based upon a segmentation of pixels based upon spatial location and colour/texture distribution. This clustering of the image allows the computation of powerful region-based features which are partially invariant to scale (Wang et al, 2005).

Super-pixel based methods work under the assumption that some segments share boundaries with objects in an image. This is not always the case, and this assumption may result in dramatic errors in the labelling (see figure 1). A number of techniques have been proposed to overcome errors in super-pixels. Rabinovich et al (2007) suggested finding the most stable segmentation from a large collection of multiple segmentations in the hope that these would be more consistent with object boundaries. Larlus and Jurie (2008) proposed an approach to the problem driven by object detection. In their algorithm, rectangular regions are detected using a bag-of-words model based upon affine invariant features. These rectangles are refined using graph cuts to extract boundaries in a manner similar to (Rother et al, 2004). Such approaches face difficulties in dealing with cluttered images, in which multiple object classes intersect. Pantofaru et al (2008) observed that although segments may not be consistent with object boundaries, the segmentation map formed by taking the intersections of multiple segmentations often is. They proposed finding the most probable labelling of intersections of segments based upon the features of their parent segments. This scheme effectively reduces the size of super-pixels. It results in more consistent segments but with a loss in the information content and discriminative power associated with each segment.

We shall show that each of these models are AHRFs, and that we are free to combine them additively and solve the resulting AHRF.

### 1.2 Choosing the Correct Segmentation

An earlier approach to dealing with the difficultly of choosing good super-pixels is to delay their choice until much later, and picking super-pixels that are consistent with a 'good' labelling of the image. Gould et al (2009a) proposed an approach in which the choice of super-pixels was integrated with the labelling of the image with object instances. Under their interpretation, super-pixels should physically exist and represent either the entirety of an object or a planar facet if the object class is amorphous and can not be decomposed into individual objects (this includes classes such as *grass*, *building*, or *sky*). Consequently, in their final labelling each pixel belongs to exactly one super-pixel chosen to represent a single instance of an object.

This process of shaping super-pixels to match object outlines is computationally challenging. As discussed in Gould et al (2009b), the optimisation techniques proposed frequently fail to recognise individual instances. Their algorithm is often unable to merge the super-pixels contained within a single instance, even if the super-pixels are correctly labelled by class. The recent work by Kumar and Koller (2010) goes some way to addressing these issues. By using sophisticated LP-relaxations they are able to trade computation time against the quality of the solution found. However, the solutions found still lack the approximation guarantees of our work, as described in section 4.2.

Another method to overcome these issues was proposed by (Kohli et al, 2008b). By formulating the labelling problem as a higher-order RF defined over pixels, they were able to recover from misleading segments which spanned multiple object classes. Further, they were able to encourage individual pixels within a single segment to share the same label by defining higher-order potentials (functions defined over cliques of size greater than 2) that penalised inhomogeneous labellings of segments. Their method can be understood as a relaxation of the hard constraint of previous methods, that the image labelling must follow super-pixel boundaries, to a softer constraint in which a penalty is paid for non-conformance.

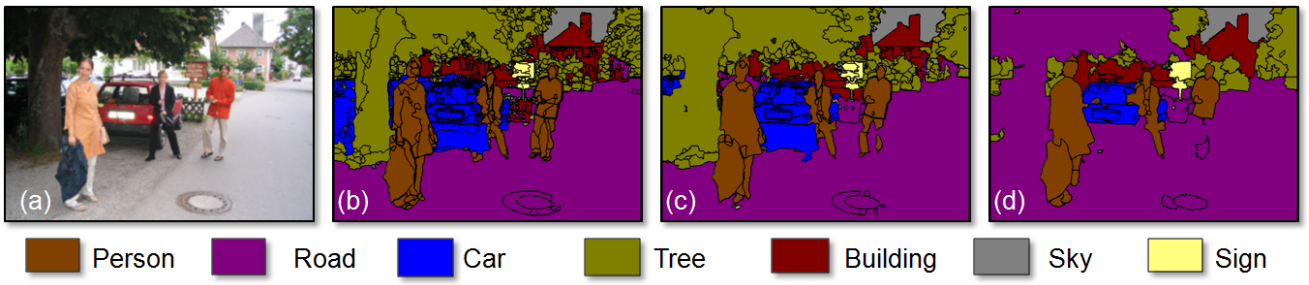



Fig. 1. *Multiple unsupervised image segmentations. (a) Original image. (b)-(d) Unsupervised image segmentations with different size super-pixels. (b), (c) and (d) use three different unsupervised segmentations of the image, in this case mean-shift, with different choices of kernel, to divide the image into segments. Each segment is assigned the label of the dominant object present in it. It can be seen that segmentation (b) is the best for tree, road, and car. However, segmentation (d) is better for the left person and the sign board.*

### 1.3 Overview of our Model

In this paper we propose a novel associative hierarchical RF formulation of semantic segmentation that allows us to combine models defined over different choices of super-pixel, avoiding the need to make a decision of which is most appropriate. It allows for the integration of features derived from different image scales (pixel, segment, and segment union/intersection). We will demonstrate how many of the state-of-the-art methods based on different fixed super-pixels can be seen as special cases of our model.

Inferring the Maximum a Posteriori solution in this framework involves the minimisation of an energy function that contain higher-order potentials defined over several thousand random variables. We show that the solutions of these difficult problems can be efficiently computed using graph-cut based algorithms similar to the pairwise methods of Boykov et al (2001). We further show that our inference techniques over these high-order energies share the bounds of such pairwise methods. The contribution of our work not limited to the problem of inference, and its application of the novel associative hierarchical RF framework to object class segmentation. We propose new sophisticated potentials defined over the different levels of the hierarchy. Evaluating the performance of our framework on some of the most challenging data sets for object class segmentation, we show that it outperforms state-of-the-art methods based on a single choice of scale. We believe this is because:

1) Our methods generalises these previous methods allowing them to be represented as particular parameter choices of our hierarchical model.
2) We go beyond these models by being able to use multiple hierarchies of segmentation simultaneously.
3) In contrast to many previous methods that do not define any sort of cost function, or likelihood, we cleanly formulate the RF energy of our model and show how it can be minimised effectively.

### 1.4 Efficient bounds for inference

The last decade has seen much work on the problem of estimating the maximum a posteriori (MAP) solution of a random field (Szeliski et al, 2006; Kolmogorov, 2006; Komodakis and Paragios, 2008; Kumar and Torr, 2008; Sontag et al, 2008; Wainwright et al, 2005; Weiss and Freeman, 2001). However, most of this research effort has focused on inference over pairwise Random Fields. Of particular interest are the families of associative pairwise potentials (Taskar et al, 2004a), which include truncated convex costs (Kumar et al, 2011), metrics (Boykov et al, 2001), and semi-metrics (Kumar and Koller, 2009). In these associative models, connected variables are assumed to be more likely than not to share the same label. Inference algorithms targeting these associative potentials, often carry bounds which guarantee the cost of the solution found must lie within a bound, specified as a fixed factor of $n$ of the cost of the minimal solution. These worst case bounds are of great interest as they often translate into algorithms with good expected performance.

Although higher-order Random Fields (i.e. those with a clique size greater than two) have been used to obtain impressive results for a number of challenging problems in computer vision (Roth and Black, 2005; Komodakis and Paragios, 2009; Vicente et al, 2009; Kohli et al, 2008a; Lan et al, 2006; Werner, 2009; Potetz and Lee, 2008; Woodford et al, 2008; Rother et al, 2009) there has been little work on deriving models and algorithms for them that carry a worst case bound.

In sections 4.2 – 5, we address the problem of performing graph cut based inference in our new model: the Associative Hierarchical Random Field (AHRF), first published in Ladicky et al (2009). Our approach includes the higher-order Associative Markov Networks (AMNs) (Taskar et al, 2004a) or $P^n$ potentials (Kohli et al, 2007), the Robust $P^n$ (Kohli et al, 2008b) model, and the recent pylon model Lempitsky et al (2011), as special cases, and we derive an approximation bound of 4 (see sections 4.2.1 and 4.3) for the quality of our solutions.

### 1.5 Hierarchical Models and Image Context

The use of image context has been well documented for object recognition and segmentation. It is particularly useful in overcoming ambiguities caused by limited evidence; this often occurs in object recognition where we frequently encounter objects at small scales or low resolution images (Hoiem et al, 2006). Classical Random Field models exploit context in a local manner by encouraging adjacent pixels or segments to take the same label. To encode context at different scales Zhu et al (2008) introduced the hierarchical image model (HIM) built of rectangular regions with parent-child dependencies. This model captures large-distance dependencies and can be solved efficiently using dynamic programming. However, it supports neither multiple hierarchies, nor dependencies between variables at the same level. To encode semantic context and to combine top-down and bottom-up approaches Tu et al (2003) proposed a framework in which they showed that the use of object specific knowledge helps to disambiguate low-level segmentation cues.

Our hierarchical RF model uses a novel formulation that allows context to be incorporated at multiple levels of multiple quantisation, something not previously possible. As we show in section 6 it leads to state of the art segmentation results.

## 2 Random Field Formulations for Semantic Segmentation

Consider an ordered set of variables $\mathbf{X} = [X_1, X_2, \ldots, X_n]$, where each variable $X_i$ takes a label from a set $\mathcal{L}$ corresponding to object classes. We write $\mathbf{x} \in \mathcal{L}^n$ for a labelling of $\mathbf{X}$, and use $x_i$ to refer to the labelling of the variable $X_i$. At times we will refer to the labelling of a subset of variables, corresponding to a *clique* $c$, for this we use the notation $\mathbf{x}_c$. We use $\mathcal{V} = \{1, 2, \ldots, n\}$ to refer to the set of valid vertices (or indexes) of $\mathbf{X}$, and make use of the common mathematical short hand of $\vee$, $\wedge$ to represent *'or'* and *'and'* respectively. We use $\Delta$ as an indicator function, *i.e.*

$$\Delta(\cdot) = \begin{cases} 1 & \text{if } \cdot \text{ is true} \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

We formulate the general problem of inference, as finding $\mathbf{x}$ the minimiser of an arbitrary cost function

$$\arg\min_{\mathbf{x}} E(\mathbf{x}). \tag{2}$$

Finding $\mathbf{x}$ may correspond to finding the maximum a posteriori (MAP) labelling. Labelling problems in vision are typically formulated as a pairwise random field whose energy can be written as a sum of unary and pairwise potentials:

$$E(\mathbf{x}) = \sum_{i \in \mathcal{V}} \psi_i(x_i) + \sum_{i \in \mathcal{V}, j \in \mathcal{N}_i} \psi_{ij}(x_i, x_j), \tag{3}$$

where $N_i$ is the set of neighbours of vertex $i$.

The pairwise random field formulation suffers from a number of problems stemming from its inability to express high-level dependencies between pixels, for example, the use of only pairwise smoothing constraints can result in a *shrinkage bias* (Kohli et al, 2008b). Despite these limitations, it is widely used and highly effective. Shotton et al (2006) applied the pairwise random field to the object class segmentation problem. They defined unary likelihood potentials using the result of a boosted classifier over a region about each pixel, that they called *TextonBoost*, which provided a substantial increase in performance of existing methods at the time.

This random field (RF) model has several distinct interpretations: As a probabilistic model (Besag, 1986; Lafferty et al, 2001), the unary potentials $\psi_i(x_i)$ of the random field can be interpreted as the negative log likelihood of variable $X_i$ taking label $x_i$, while the pairwise potential encodes a smoothness prior over neighbouring variables. Under this formulation the *maximum a posteriori* (MAP) estimate corresponds to the minimal cost labelling of cost (3). These cost can also be interpreted as defining a *structured discriminative classifier* (Tsochantaridis et al, 2005; Nowozin et al, 2010), *i.e.* a classifier whose costs do not characterise the $\log$ marginal distribution of images, but whose minimum cost labelling is likely to correctly label most of the image. These classifiers can be efficiently learnt in a max-margin framework Taskar et al (2004a); Nowozin et al (2010); Alahari et al (2010); Szummer et al (2008).

The primary contribution of this work is in the proposal of a framework that allows a principled contribution of arbitrary cues from many different cues and models. Therefore, we use the more general discriminative interpretation of our work. The learning method discussed in section 7 is discriminative rather than probabilistic.

### 2.1 The Robust $P^N$ model

The pairwise random field formulation of (Shotton et al, 2006) was extended by (Kohli et al, 2008b) with the incorporation of robust higher-order potentials defined over segments. Their formulation was based upon the observation that pixels lying within the same segment are more likely to take the same label. The energy of the higher-order random field proposed by (Kohli et al, 2008b) was of the form:

$$E(\mathbf{x}) = \sum_{i \in \mathcal{V}} \psi_i(x_i) + \sum_{i \in \mathcal{V}, j \in \mathcal{N}_l} \psi_{ij}(x_i, x_j) + \sum_{c \in \mathcal{S}} \psi_c^h(\mathbf{x}_c), \tag{4}$$

where $\mathcal{S}$ is a set of cliques (or segments), given by one or more super-pixel algorithms as shown in figure 1, and $\psi_c^h$ are higher-order potentials defined over the cliques. The higher order potentials took the form of

a Robust $P^N$ model defined as:

$$\psi_c^h(\mathbf{x}_c) = \min_{l \in \mathcal{L}} \left( \gamma_c^{\max}, \gamma_c^l + \sum_{i \in c} w_i k_c^l \Delta(x_i \neq l) \right), \quad (5)$$

where $w_i$ is the weight of the variable $x_i$, and the variables $\gamma$ satisfy

$$\gamma_c^l \leq \gamma_c^{\max}, \forall l \in \mathcal{L}. \quad (6)$$

The potential has a cost of $\gamma_c^l$ if all pixels in the segment take the label $l$. Each pixel not taking the label $l$ is penalised with an additional cost of $w_i k_c^l$, and the maximum cost of the potential is truncated to $\gamma_c^{max}$. This framework enabled the integration of multiple segmentations of the image space in a principled manner. segment to belong to a particular label.

### 2.2 Hierarchical Random Fields

We now show how the higher-order $P^N$ potentials of (5) are equivalent to the minimisation of a pairwise graph defined over the same clique $\mathbf{x}_c$ and a single auxiliary variable $x_c^{(1)}$ takes values from an extended label set

$$\mathcal{L}^E = \mathcal{L} \cup \{L_F\}. \quad (7)$$

In a minimal cost solution, $x_c^{(1)}$ taking a value $l$ will correspond to the clique $\mathbf{x}_c$ having a dominant label $l$, *i.e.* the majority of pixels within $c$ must take label $l$. If it takes the free label $L_F$, it means that there is no dominant label in the clique, and that segment is unassigned. The cost function over $\mathbf{x}_c \cup \{x_c^{(1)}\}$ takes the form:

$$\psi_c^p(\mathbf{x}_c, x_c^{(1)}) = \phi_c(x_c^{(1)}) + \sum_{i \in c} \phi_c(x_c^{(1)}, x_i). \quad (8)$$

where the unary potential over $x_c^{(1)}$, $\phi_c(x_c^{(1)})$ associates the cost $\gamma_c^l$ with $x_c^{(1)}$ taking a label in $\mathcal{L}$, and $\gamma_c^{\max}$ with $x_c^{(1)}$ taking the *free* label $L_F$. The pairwise potentials $\phi_c(x_c^{(1)}, x_i)$ are defined as:

$$\phi_c(x_c^{(1)}, x_i) = \begin{cases} 0 & \text{if } x_c^{(1)} = L_F \vee x_c^{(1)} = x_i \\ w_i k_c^{x_c^{(1)}} & \text{otherwise.} \end{cases} \quad (9)$$

Then:

$$\psi_c^h(\mathbf{x}_c) = \min_{x_c^{(1)}} \psi_c^p(\mathbf{x}_c, x_c^{(1)}). \quad (10)$$

By ensuring that the pairwise edges between the auxiliary variable and its children satisfy the constraint:

$$\sum_{i \in c} w_i k_c^l \geq 2(\phi_c(L_F) - \phi_c(l)), \forall l \in \mathcal{L}, \quad (11)$$

(see section 4.4.1) we can guarantee that the labels of these auxiliary variables carry a clear semantic meaning. If this constraint is satisfied an auxiliary variable may take state $l \in \mathcal{L}$ in a minimal cost labelling, if and only if, the weighted majority of its child nodes take state $l$. State $L_F$ indicates a heterogeneous labelling of a segment in which no label holds a significant majority.

We now extend the model to include pairwise dependencies between auxiliary variables:

$$E(\mathbf{x}) = \sum_{i \in \mathcal{V}} \psi_i(x_i) + \sum_{i \in \mathcal{V}, j \in \mathcal{N}_i} \psi_{ij}(x_i, x_j) + \min_{\mathbf{x}^{(1)}} \left( \sum_{c \in \mathcal{S}} \psi_c^p(\mathbf{x}_c, x_c^{(1)}) + \sum_{c,d \in \mathcal{S}} \psi_{cd}(x_c^{(1)}, x_d^{(1)}) \right). \quad (12)$$

These pairwise terms can be understood as encouraging consistency between neighbouring cliques. This framework can be further generalised to a hierarchical model where the connection between layers takes the form of (8) and the weights for each child node in $\phi_c(\cdot)$ are proportional to the sum of the weights in the *"base layer"* belonging to the clique $c$.

The energy of our new hierarchical model is of the form:

$$E(\mathbf{x}) = \sum_{i \in \mathcal{V}} \psi_i(x_i) + \sum_{i \in \mathcal{V}, j \in \mathcal{N}_i} \psi_{ij}(x_i, x_j) + \min_{\mathbf{x}^{(1)}} E^{(1)}(\mathbf{x}, \mathbf{x}^{(1)}), \quad (13)$$

where $E^{(1)}(\mathbf{x}, \mathbf{x}^{(1)})$ is recursively defined as:

$$E^{(\mathrm{n})}(\mathbf{x}^{(n-1)}, \mathbf{x}^{(\mathrm{n})}) = \sum_{c \in \mathcal{S}^{(\mathrm{n})}} \psi_c^p(\mathbf{x}_c^{(\mathrm{n-1})}, x_c^{(\mathrm{n})}) + \sum_{c,d \in \mathcal{S}^{(\mathrm{n})}} \psi_{cd}(x_c^{(\mathrm{n})}, x_d^{(\mathrm{n})}) + \min_{\mathbf{x}^{(n+1)}} E^{(n+1)}(\mathbf{x}^{(\mathrm{n})}, \mathbf{x}^{(n+1)}). \quad (14)$$

Where $\mathbf{x}^{(0)} = \mathbf{x}$ refers to the state of the base level, and $\mathbf{x}^{(\mathrm{n})}$ for $n \geq 1$ the state of auxiliary variables.

The inter-layer potential between between two layers of auxiliary variables takes the form of weighted Robust $P^N$:

$$\phi_c(\mathbf{x}_c^{(\mathrm{n-1})}, x_c^{(\mathrm{n})}) = \begin{cases} 0 & \text{if } x_c^{(\mathrm{n})} = L_F \wedge x_c^{(\mathrm{n})} = \mathbf{x}_c^{(\mathrm{n-1})} \\ w_{\mathbf{x}_c^{(\mathrm{n-1})}} k_c^{\mathbf{x}_c^{(\mathrm{n-1})}} & \text{otherwise,} \end{cases} \quad (15)$$

where the weights are summed up over the base layer as:

$$w_{\mathbf{x}_c^{(\mathrm{n-1})}} = \sum_{i \in \mathbf{x}_c^{(\mathrm{n-1})}} w_i. \quad (16)$$

## 3 Relation to Previous Models

In this section, we draw comparisons with the current state-of-the-art models for object segmentation (Galleguillos et al, 2008; Pantofaru et al, 2008; Rabinovich et al, 2007; Yang et al, 2007) and show that at certain choices of the parameters of our model, these methods fall out as special cases (illustrated in figure 2). Thus, our method not only generalises the standard pairwise RF formulations over pixels, but also the previous work based on super-pixels and (as we shall see) provides a global optimisation framework allowing us to combine features at different quantisation levels.

We will now show that our model is not only a generalisation of RFs over pixels, but also of two

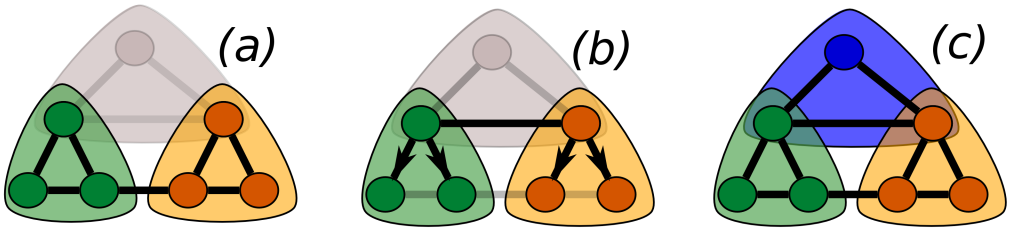



Fig. 2. Existing models as special cases of our hierarchical model. *The lowest layer of the image represents the pixel layer, the middle layer potentials defined over super-pixels or segments, and the third layer represents our hierarchical terms.* (a) *shows the relationships permitted in a pixel-based* RF *with Robust* $P^N$ *potentials.* (b) *shows relationships contained within a super-pixel-based* RF *(the directed edges indicate the one way dependence between the labellings of pixels and super-pixels).* (c) *Our hierarchical* RF. *See section 3.*

classes of preexisting model: *(i)* RFs based upon disjoint segments (Batra et al, 2008; Galleguillos et al, 2008; Yang et al, 2007) (see figure 2(b)), and *(ii)* RFs based upon the intersection of segments (Pantofaru et al, 2008).

### 3.1 Equivalence to RFs based on Segments

Consider a hierarchy composed of only pixels, and cliques corresponding to one super-pixel based segmentation of the image. In this case, all the segments are disjoint (non-overlapping). In this case, our model becomes equivalent to the pairwise RF models defined over segments (Batra et al, 2008; Galleguillos et al, 2008; Rabinovich et al, 2007; Yang et al, 2007).

To ensure that no segment takes the label $L_F$, we assign a high value to $\gamma_c^{\max} \rightarrow \infty, \forall c \in \mathcal{S}^{(1)}$. In this case, the optimal labelling will always be *segment consistent* (i.e. all pixels within the segment will take the same label) and the potential $\psi_c^p(\mathbf{x}_c, x_c^{(1)})$ can now be considered as a unary potential over the auxiliary (segment) variable $x_c^{(1)}$. This allows us to rewrite (12) as:

$$E(\mathbf{x}^{(1)}) = \sum_{c \in \mathcal{S}^{(1)}} \psi_c(x_c^{(1)}) + \sum_{c,d \in \mathcal{S}^{(1)}} \psi_{cd}(x_c^{(1)}, x_d^{(1)}) \quad (17)$$

which is exactly the same as the cost associated with the pairwise RF defined over segments with $\psi_c(x_c^{(1)} = l) = \gamma_c^l$ as the unary cost and $\psi_{cd}(\cdot)$ as the pairwise cost for each segment.

### 3.2 Equivalence to Models of Segment Intersections

We now consider the case with multiple overlapping segmentations and unary and pairwise potentials defined only upon these segments; this is analogous to the construct of Pantofaru et al (2008). If we set:

$$w_i k_c^l = \gamma_c^{\max}, \forall i \in \mathcal{V}, l \in \mathcal{L}, c \in \mathcal{S}, \quad (18)$$

then:

$$x_c^{(1)} \neq L_F \text{ only if } x_i = x_c^{(1)}, \forall i \in c. \quad (19)$$

In this case, only the potentials $\sum_{c \ni i} \psi_c^p(\mathbf{x}_c, x_c^{(1)})$ act on $x_i$.

Consider a pair of pixels $i, j$ that lie in the same intersection of segments *i.e.* :

$$\{c \in S : c \ni i\} = \{c \in S : c \ni j\}. \quad (20)$$

Then, in a minimal labelling, either $\exists x_c^{(1)} = x_i$, and hence $x_j = x_c^{(1)} = x_i$, or $\forall c \ni i : x_c^{(1)} = L_F$. In the second degenerate case there are no constraints acting on $x_i$ or $x_j$, and a minimal cost labelling can be chosen such that $x_i = x_j$.

Consequently, there is always a minimal cost labelling consistent with respect to the intersection of segments, in this sense our model is equivalent to that proposed in Pantofaru et al (2008).

### 3.3 Equivalence to tree structured associative models

Tree structured hierarchies such as (Zhu and Yuille, 2005; Lim et al, 2009; Nowozin et al, 2010; Reynolds and Murphy, 2007) have been proposed for semantic segmentation. The structure of these models is clearly a strict subset of ours, as it does not support pairwise connections between variables in the same level, and each variable may only be attached to one variable in the layer above. In the restricted case, in which the label space and edge costs between parent and child are of the same form as those we consider, these models can also be contained in our approach, although this need not hold in general.

### 3.4 The Relationship with Directed Models

A hierarchical, two-layer, directed model was proposed in (Kumar and Hebert, 2005). This is a hybrid model relatively similar to ours, with unary and pairwise potentials defined over both super-pixels and pixels and pairwise connections between the layers, enforcing consistency. It principally differs from ours in the use of directed edges between layers. These directed edges mean that max-marginals can be computed in a piecewise manner, and propagated from one layer to the other. This makes it more suitable for inference with message passing algorithms than our framework (see section 8.1).

This directed approach does not propagate information through-out the structure. In order to arrive at a

consistent hypothesis, that takes account conflicting clues from all levels of the hierarchy, there are two desirable criteria for the propagation of information.

1) We wish for information to be transmitted from the pixel to the segment level and from there back to the pixel level. That is, the labelling of one pixel should affect the label of the segment potential and, from this, the label of other pixels in the same segment.
2) Information should also be transmitted from a segment to the pixel level and back to the segment level. This means that if two segments overlap, the optimal label of one segment should indirectly depend on the labelling of the other.

At most one of the above-mentioned conditions can hold if the connections between layers of the hierarchical model form a directed acyclic graph (DAG). This is the case in the model of Kumar and Hebert (2005), and in *Deep Belief Nets* (Bengio et al, 2007; Hinton and Osindero, 2006), that only satisfy the first criteria. In simple hierarchies such as those proposed by Kumar and Herbert, in which each pixel belongs to only one segment, the second criteria is less important. However, the integration of multiple segmentations into a single coherent hypothesis depends upon the transmission of information in both directions.

### 3.5 Equivalence to the Pylon Model

The recent work by Lempitsky et al (2011) proposed a new approach to segmentation in which, given a tree-structured hierarchy of super-pixels $\mathcal{S}$, they seek a labelling $\mathbf{x} \in \mathcal{L} \cup \{L_F\}$ that minimise a strictly positive cost of the form:

$$E'(\mathbf{x}) = \sum_{c \in \mathcal{S}} \psi'_c(x_c) + \sum_{i \in \mathcal{V}, j \in \mathcal{N}_i} \psi_{i,j}(x_i, x_j) \qquad (21)$$

here $\mathcal{V}$ denotes the set of pixels in the image. This is a restricted case of our work, in that the pairwise costs $\psi'$ are only defined over pixels, and not super-pixels, while unary potentials are not defined over individual pixels. Unlike our work, they impose the constraint that

$$\psi'_c(L_F) = 0 \ \ \forall c \in \mathcal{S}, \qquad (22)$$

and that for every pixel $i$ in the image $\mathcal{V}$ there exists at least one segment $c \in \mathcal{S}$ such that:

$$x_c \neq L_F \wedge i \in c \qquad (23)$$

or equivalently, every pixel belongs to at least one 'active' segment not taking label $L_F$. In this formulation, assigning the label $L_F$ to a segment $c$ associates no cost with it, and is analogous to discarding it. In a minimum cost labelling, only one super-pixel containing any pixel will take a state other than $L_F$, and this objective function is similar to those proposed by Gould et al (2009a). Note that every pixel $i$ always has a label $x_i \neq L_F$ associated with it, and this label is shared by any super-pixel that contains it, and does not take label $L_F$.

In Lempitsky et al (2011), the authors claim that inference of their model can be performed using the bounded inference techniques we discuss in this paper.

We will now show how this is the case by constructing an equivalent cost function of the form described in section 2.2.

We replace the implicit constraint of Lempitsky et al (2011), that if a super-pixel $c \in \mathcal{S}$ takes label $l \in \mathcal{L}$, all its children *i.e.* $j \in \mathcal{S} : j \subset c$ take the label $L_F$, with the constraint that:

$$x_j = x_c = l \ \forall j \subset c \ \forall x_c \neq L_F, \qquad (24)$$

*i.e.* that a child's labelling must agree with its parent. This can be enforced by setting the weights $w_i k_c$ of equation (9) to be arbitrarily high. We take the same tree of super-pixels $S$ as our hierarchy, and writing $C(c)$ for the children of $c$, we set unary potentials to be:

$$\psi_c(x_c) = \psi'_c(x_c) - \sum_{j \in C(c)} \psi'_j(x_c) \ \forall c, x_c \in \mathcal{L}, \qquad (25)$$

and

$$\psi_c(L_F) = 0. \qquad (26)$$

Taking the pairwise potentials $\psi_{i,j}$ to be the same as the pylon model, we consider the cost:

$$E(\mathbf{x}) = \sum_{c \in \mathcal{S}} \psi_c(x_c) + \sum_{i \in \mathcal{I}, j \in \mathcal{N}_i} \psi_{i,j}(x_i, x_j) \qquad (27)$$

Now, given any labelling $\mathbf{x}$, we can create a new label $\mathbf{x}'$ by the following transform:

$$x'_{C(i)} = \begin{cases} L_F & \text{if } x_{C(i)} = L_F \wedge x_{C(i)} = x_i \\ x_{C(i)} & \text{otherwise.} \end{cases} \qquad (28)$$

Then $E(\mathbf{x}) = E'(\mathbf{x}')$[2], and any minimal $\mathbf{x}$ labelling of $E(\mathbf{x})$, must induce a minimal labelling $\mathbf{x}'$of $E(\mathbf{x}')$, and every pylon model is equivalent to an AHRF.

### 3.6 Co-occurrence, Harmony and Stack Potentials

Hierarchical models allow the modelling of contextual relationships. The use of context has been well documented by various researchers, for example, the harmony potentials of Boix et al (2011), the contextual potentials of Lim et al (2009) or the stack parameters of Munoz et al (2010). The above works proposed the use of potentials that penalise the labelling of a subset of the image if they do not naturally occur together in an image. For example, the combination of *cow, aeroplane* and *sky* would be penalised, while *cow*, *grass*, and *sky* would not. This is in contrast to our potentials which penalises *inhomogeneous* labellings of a region.

2. Proof follows by induction over the hierarchy from fine to coarse.

Although these works do not formally characterise the potentials they can solve, they all appear to be *monotonic increasing i.e.* their cost never decreases when a new label is added to a subset of an image. As such, these potentials can be understood as *local co-occurrence costs* (or co-occurrence costs defined over a subset of the image) and integrated in our framework using the work (Ladicky et al, 2010).

This use of *local co-occurrence* rather than $P^n$ like potentials, is useful in tree-structured hierarchies that can not make use of pairwise potentials to encourage neighbouring regions to take the same label. In our framework, a region containing a large amount of both *grass* and *cow* would be marked as inhomogeneous and take label $L_F$, while using *local co-occurrence* the region could be considered as a *cow-field*, and further regions could be merged with it, without modification, only if they take label *cow* or *grass*. This is particularly desirable if no pairwise smoothing exists between adjacent regions in the image. While, in common with Boix et al (2011), we have found *global co-occurrence potentials* useful, we have been unable to find a use for *local co-occurrence*, in our framework, that could not be better handled by pairwise potentials (see Galleguillos et al (2008) for an example of how pairwise potentials can be used to capture local co-occurrence).

### 3.7 Robustness to Misleading Segmentations

The quantisation of image space obtained using unsupervised segmentation algorithms may be misleading since segments may contain multiple object classes. Assigning the same label to all pixels of such segments will result in an incorrect labelling. This problem can be overcome by using segment quality measures proposed by (Rabinovich et al, 2007; Ren and Malik, 2003) which can be used to distinguish the *good* segments from *misleading* ones. These measures can be seamlessly integrated in our hierarchical framework by modulating the strength of the potentials defined over segments. Formally, this is achieved by weighting the potentials $\psi_c^h(\mathbf{x}_c, x_c^{(1)})$ according to a quality sensitive measure $Q(c)$ for any segment $c$, as in Kohli et al (2008b).

## 4 Inference for Hierarchical RFs

The preceding sections introduced the AHRF model and explained how it can be used for semantic segmentation. However, a model is of little use without an efficient and practical method for optimisation. As the associative hierarchical random field is a pairwise model over $|\mathcal{L}|+1$ variables, any method designed for solving general pairwise models could be used. In this section we analyse the suitability of various inference methods for AHRF and propose a novel move making algorithm, that outperforms general pairwise methods for this task.

### 4.1 Inference in Pairwise Random Fields

Although the problem of MAP inference is NP-hard for most associative pairwise functions defined over more than two labels, in real world problems many conventional algorithms provide near optimal solutions over grid connected random fields (Szeliski et al, 2006). However, the dense structure of hierarchical random fields makes traditional message passing algorithms such as loopy belief propagation (Weiss and Freeman, 2001) and tree-reweighted message passing (Kolmogorov, 2006) converge slowly to high cost solutions (Kolmogorov and Rother, 2006). The difficulties faced by these message passing algorithms can be attributed to the presence of frustrated cycles (Sontag et al, 2008; Werner, 2009) that can be eliminated via the use of cycle inequalities, but only by significantly increasing run time.

Graph cut based move making algorithms do not suffer from this problem and have been successfully used for minimising pairwise functions defined over densely connected networks encountered in vision.

Examples of move making algorithms include $\alpha$-expansion which can only be applied to metrics, $\alpha\beta$ swap which can be applied to semi-metrics (Boykov et al, 2001), and range moves (Veksler, 2007; Kumar et al, 2011) for truncated convex potentials. These moves differ in the size of the space searched for the optimal move. While expansion and swap search a space of size at most $2^n$ while minimising a function of $n$ variables, range moves explores a much larger space of $K^n$ where $K > 2$ (see Kumar et al (2011); Veksler (2007) for more details). Of these move making approaches, only $\alpha\beta$ swap can be directly applied to associative hierarchical random fields – other methods require that the inter-layer cost either form a metric, or are truncated convex with respect to some ordering of the labels.

Move-making algorithms start from an arbitrary initial solution of the problem and proceed by making a series of changes each of which leads to a solution of the same or lower energy (Boykov et al, 2001). At each step, the algorithms project a set of candidate moves into a Boolean space, along with their energy function. If the resulting projected energy function (also called the *move energy*) is both submodular and pairwise, it can be exactly minimised in polynomial time by solving an equivalent st-mincut problem. These optima can then be mapped back into the original space, returning the optimal move within the move set. The move algorithms run this procedure until convergence, iteratively picking the best candidate as different choices of range are cycled through.

#### 4.1.1 Minimising Higher-Order Functions

A number of researchers have worked on the problem of MAP inference in higher-order RFs. Lan et al (2006) proposed approximation methods for BP to make

| $\mathbf{x}$ | ... | $\beta$ | $\gamma$ | $\alpha$ | $\gamma$ | $\gamma$ | ... |
|---|---|---|---|---|---|---|---|
| Proposed Moves | ... | $\alpha$ | $\alpha$ | $\alpha$ | $\alpha$ | $\alpha$ | ... |
| Move choice $\mathbf{t}$ | ... | 0 | 1 | 1 | 1 | 0 | ... |
| $\mathbf{x}'$ | ... | $\alpha$ | $\gamma$ | $\alpha$ | $\gamma$ | $\alpha$ | ... |

Fig. 3. *An illustration of move encoding in $\alpha$-expansion. Starting from an initial model $\mathbf{x}$ a new move $\mathbf{t}$ is proposed which causes two variables to change their label to $\alpha$. This results in the new labelling $\mathbf{x}'$.*

efficient inference possible in higher-order MRFs. This was followed by the recent works of Potetz and Lee (2008); Tarlow et al (2008, 2010) in which they showed how belief propagation can be efficiently performed in random fields containing moderately large cliques. However, as these methods were based on BP, they were quite slow and took minutes or hours to converge, and lack bounds.

To perform inference in the $P^n$ models, Kohli et al (2007, 2008b), first showed that certain projection of the higher-order $P^n$ model can be transformed into submodular pairwise functions containing auxiliary variables. This was used to formulate higher-order expansion and swap move making algorithms.

### 4.2 Pairwise Inference with $\alpha$-Expansion

The algorithm $\alpha$-expansion makes a sequence of progressive decisions to replace the labels of a subset of the RF with label $\alpha$. At each step in the sequence, for a given choice of $\alpha$, the optimal subset which results in the lowest possible cost is found via graph-cuts (see figure 3). The steps proposed by $\alpha$-expansion, can be can be encoded as a *transformation vector* of binary variables $\mathbf{t} = \{t_i, \forall i \in \mathcal{V}\}$. Each component $t_i$ of $\mathbf{t}$ encodes a partial decision, about what the state of the variable $x_i$ should change to. In $\alpha$-expansion, $t_i = 1$ encodes the decision that $x_i$ should remain constant, while $t_i = 0$ indicates that it should transition to a new label $\alpha$. See figure 3 for an illustration.

The use of these transformation vectors makes the problem of finding the optimal move equivalent to minimising a pseudo-Boolean cost function from $2^n \rightarrow \mathbb{R}$. Consequently, it follows from Kolmogorov and Zabih (2004), that if these cost functions can be shown to be pairwise submodular, the optimal move can be efficiently found using graph-cut.

The inter-layer pairwise potentials associated with our hierarchy (see (8)) are not positive-definite, nor symmetric, so we can not directly solve these problems with the standard form of $\alpha$-expansion.

This requires a transform of all functions in the pairwise representation, so that they can be representable as a metric (Boykov et al, 2001). This transformation is non-standard and should be considered a contribution of this work.

We alter the form of the potentials in two ways. First, we assume that all variables in the hierarchy take values from the same label set $\mathcal{L}^E$ defined in equation (7). Where this is not true, for example, original variables $\mathbf{x}^{(0)}$ at the base of the hierarchy never take label $L_F$, we augment the label set with the label $L_F$ and associate a prohibitively large unary cost with it. Secondly, we make the inter-layer pairwise potentials *symmetric*[3] by performing a local reparameterisation operation.

*Lemma 1:* The inter-layer pairwise functions

$$\phi_{ic}^{(n)}(x_i^{(n-1)}, x_c^{(n)}) = \begin{cases} 0 & \text{if } x_c^{(n)} = L_F \text{ or } x_c^{(n)} = x_i^{(n-1)} \\ k_c^i & \text{if } x_c^{(n)} = l \in \mathcal{L} \text{ and } x_i^{(n-1)} \neq l \end{cases} \tag{29}$$

of (13) can be written as:

$$\begin{aligned} \phi_{ic}^{(n)}(x_i^{(n-1)}, x_c^{(n)}) &= \psi_i^{(n-1)}(x_i^{(n-1)}) + \psi_c^{(n)}(x_c^{(n)}) \\ &+ \Phi_{ic}^{(n)}(x_i^{(n-1)}, x_c^{(n)}), \end{aligned} \tag{30}$$

$$\text{where} \quad \Phi_{ic}^{(n)}(x_i^{(n-1)}, x_c^{(n)}) = \begin{cases} 0 & \text{if } x_i^{(n-1)} = x_c^{(n)} \\ k_c^i/2 & \text{if } x_i^{(n-1)} = L_F \\ & \text{or } x_c^{(n)} = L_F \\ & \text{and } x_i^{(n-1)} \neq x_c^{(n)} \\ k_c^i & \text{otherwise,} \end{cases} \tag{31}$$

and

$$\psi_c^{(n)}(x_c^{(n)}) = \begin{cases} 0 & \text{if } x_c^{(n)} \in \mathcal{L} \\ -k_c^i/2 & \text{otherwise,} \end{cases} \tag{32}$$

$$\psi_i^{(n-1)}(x_i^{(n-1)}) = \begin{cases} 0 & \text{if } x_i^{(n-1)} \in \mathcal{L} \\ k_c^i/2 & \text{otherwise.} \end{cases} \tag{33}$$

*Proof:* Consider a clique containing only one variable, the general case will follow by induction. Note that if no variables take state $L_F$ the costs are invariant to reparameterisation. This leaves three cases:

$$\begin{array}{|r|} \hline \mathbf{x_c^{(n)}} = \mathbf{L_F}, \mathbf{x_i^{(n-1)}} \in \mathcal{L} \\ \psi_c(x_c^{(n)}) + \psi_{ic}(x_c^{(n)}, x_i^{(n-1)}) = -k/2 + k/2 = 0 \\ \hline \mathbf{x_c^{(n)}} \in \mathcal{L}, \mathbf{x_i^{(n-1)}} = \mathbf{L_F} \\ \psi_i(x_i^{(n-1)}) + \psi_{ic}(x_i^{(n-1)}, x_c^{(n)}) = k/2 + k/2 = k \\ \hline \mathbf{x_c^{(n)}} = \mathbf{L_F}, \mathbf{x_i^{(n-1)}} = \mathbf{L_F} \\ \psi_i(x_i^{(n-1)}) + \psi_{ic}(x_i^{(n-1)}, x_c^{(n)}) + \psi_c(x_c^{(n)}) = \frac{k-k}{2} = 0 \\ \hline \end{array} \tag{34}$$

□

#### 4.2.1 Bounded Higher-order Inference

We now prove bounds for $\alpha$-expansion over an AHN.

- By inspection, the pairwise function of lemma 1, is positive definite, symmetric, and satisfies the triangle inequality

$$\psi_{a,b}(x, z) \leq \psi_{a,b}(x, y) + \psi_{a,b}(y, z) \forall x, y, z \in \mathcal{L} \cup \{L_F\}. \tag{35}$$

  Hence it is a metric, and the algorithms $\alpha\,\beta$ swap and $\alpha$-expansion can be used to minimise it.

3. A pairwise potential $\psi_{i,j}(x_i, x_j)$ is said to be *symmetric* if $\psi_{i,j}(x_i, x_j) = \psi_{i,j}(x_j, x_i)$ for all choices of $x_i$ and $x_j$.

- By the work of Boykov et al (2001), the $\alpha$-expansion algorithm is guaranteed to find a solution within a factor of

$$k = 2\max\left(2, \max_{E\in\mathcal{E}^{(1)}} \frac{\max_{x_i,x_j\in\mathcal{L}} \psi_E(x_i,x_j)}{\min_{x_i,x_j\in\mathcal{L}} \psi_E(x_i,x_j)}\right) \tag{36}$$

of the global optima, where $\mathcal{E}^{(1)}$ is the set of edges in the base layer. This has a value of 4 if the potentials defined over the base layer of hierarchy take the form of a Potts model).

We now prove that these bounds transfer to $\alpha$-expansion over an higher-order function $E(\mathbf{x}^{(0)})$ that can be written as a pairwise function with minimised auxiliary variables i.e.

$$E(\mathbf{x}^{(0)}) = E'(\mathbf{x}^{(0)}) + \min_{\mathbf{x}^{(a)}} E^a(\mathbf{x}^{(0)}, \mathbf{x}^{(a)}). \tag{37}$$

Intuitively, the bounds must transfer to the higher-order form for two reasons: the higher-order and pairwise forms share global minima of the same cost; and for any choice of $\mathbf{x}^{(0)}$, $\mathbf{x}^{(a)}$ the value of the high-order function (formed by minimising over $\mathbf{x}^{(a)}$) is always be less than or equal to the equivalent pairwise cost. As such, the transformation into a higher-order function must tighten rather than loosen the bound.

More formally, we perform inference of the pairwise form of the higher-order function, and assume the cost of the solution $(\mathbf{x}', \mathbf{x}^+)$ found via $\alpha$-expansion lies within a factor 4 of the global minima $\min_{\mathbf{x},\mathbf{x}^a} E'(\mathbf{x}) + E^a(\mathbf{x}, \mathbf{x}^a)$ then,

$$E'(\mathbf{x}') + \min_{\mathbf{x}^{(a)}} E^a(\mathbf{x}', \mathbf{x}^{(a)}) \leq E'(\mathbf{x}') + E^a(\mathbf{x}', \mathbf{x}^+) \tag{38}$$

$$\leq k \min_{\mathbf{x}^{(0)},\mathbf{x}^a} \left(E'(\mathbf{x}^{(0)}) + E^a(\mathbf{x}^{(0)}, \mathbf{x}^{(a)})\right) \tag{39}$$

and

$$E(\mathbf{x}') \leq k \min_{\mathbf{x}^{(0)}} E(\mathbf{x}^{(0)}). \tag{40}$$

Consequently, the bound of 4 is preserved in the transformation that maps the pairwise energy back to its higher-order form.

### 4.3 Range-move $\alpha$-Expansion

Let us consider a generalisation of the swap and expansion moves proposed in Boykov et al (2001). In a standard swap move, the set of all moves considered is those in which a subset of the variables currently taking label $\alpha$ or $\beta$ change labels to either $\beta$ or $\alpha$. In our range-swap the moves considered allow any variables taking labels $\alpha$, $L_F$ or $\beta$ to change their state to any of $\alpha$, $L_F$ or $\beta$. Similarly, while a normal $\alpha$ expansion move allows any variable to change to some state $\alpha$, our range expansion allows any variable to change to states $\alpha$ or $L_F$.

This approach always considers a set of moves that completely contains the moves considered by $\alpha$-expansion. As such any local optima of range-move expansion is a local optima of $\alpha$-expansion, and the bound on the cost of the solution found must be at most that of $\alpha$-expansion, *i.e.* 4.

These moves can be seen as a variant on the ordered range moves proposed in Kumar et al (2011). However, Kumar et al (2011) required that there exists an ordering of the labels $\{l_1, l_2, \ldots, l_n\}$ such that the cost function is convex over the range $\{l_i, l_{i+1} \ldots l_{i+j}\}$ for some $j \geq 2$, and our range moves require no such ordering.

### 4.4 Transformational Optimality

Consider an energy function defined over the variables $\mathbf{x} = \{\mathbf{x}^{(h)}, h \in \{1, 2, \ldots, H\}\}$ of a hierarchy with $H$ levels. We call a move making algorithm *transformationally optimal* if and only if any move $(\mathbf{x}^*, \mathbf{x}^a)$ proposed by the algorithm satisfies the property:

$$E^a(\mathbf{x}^*, \mathbf{x}^a) = \min_{\mathbf{x}'} E^a(\mathbf{x}^*, \mathbf{x}') \tag{41}$$

*i.e.* $\mathbf{x}^a$ is a minimiser of $E^a(\mathbf{x}^*, \cdot)$. Inserting this into equation (37) we have:

$$E(\mathbf{x}^*) = E'(\mathbf{x}^*) + E^a(\mathbf{x}^*, \mathbf{x}^a). \tag{42}$$

This implies that the partial move $\mathbf{x}^*$ proposed by a transformationally optimal algorithm over

$$E'(\mathbf{x}^{(1)}) + E^a(\mathbf{x}^{(1)}, \mathbf{x}^a) \tag{43}$$

must function as a move that directly minimise the higher-order cost of equation (37). Experimentally, our transformationally optimal algorithms converge faster, and to better solutions than standard approaches, such as $\alpha$-expansion. Moreover, unlike standard approaches, our transformationally optimal algorithms are guaranteed find the exact solution for binary AHNs.

We now show that when applied to hierarchical random fields, subject to two natural requirements, namely hierarchical consistency, and a restricted choice of pairwise potentials, *range* moves are transformationally optimal.

#### 4.4.1 Hierarchical Consistency

To guarantee transformational optimality, for $\alpha$-expansion and $\alpha\beta$-swap, we must constrain the set of higher order potentials further. To do this we will introduce the notion of hierarchical consistency, which simply says auxiliary variables should agree with the state of their child variables in a minimal cost labelling. This has two important properties: first it restricts the space of possible labels taken by an auxiliary variable, allowing us to guarantee transformational optimality; more importantly it will enforce agreement between the different layers of the hierarchy and guarantee a consistent labelling across the entire hierarchy.

Consider a clique $c$ with an associated auxiliary variable $x_c^{(i)}$. Let $\mathbf{x}_l$ be a labelling such that $x_c^{(i)} = l \in \mathcal{L}$

and $\mathbf{x}_{L_F}$ be a labelling that differs from it only in that the variable $x_c^{(i)}$ takes label $L_F$. We say a clique potential is *hierarchically consistent* only if it satisfies the constraint:

$$E(\mathbf{x}_l) \geq E(\mathbf{x}_{L_F}) \implies \frac{\sum_{i \in c} k_c^i \Delta(x_i = l)}{\sum_{i \in c} k_c^i} > 0.5. \quad (44)$$

The property of hierarchical consistency is also required in computer vision for the cost associated with the hierarchy to remain meaningful. The labelling of an auxiliary variable within the hierarchy should be reflected in the state of the clique associated with it. If an energy is not hierarchically consistent, it is possible that the optimal labelling of regions of the hierarchy will not reflect the labelling of the base layer.

To understand why this consistency is important, we consider a case where this is violated. Consider a simple energy function consisting of a base layer of 10 pixels $\mathbf{x}^{(0)}$ and only one clique, with associated auxiliary variable $x_c$, defined over the base layer. We assume that unary potentials defined over individual pixels have a preference for the class *cow* while the higher-order potential defined over the clique expresses a preference for class *sheep*. More formally we set:

$$\psi_i(x_i) = \begin{cases} 2 & \text{if } x_i = \textit{sheep} \\ 0 & \text{if } x_i = \textit{cow} \end{cases} \quad \forall x_i \in \mathbf{x}^{(0)} \quad (45)$$

$$\phi_c(x_c) = \begin{cases} 0 & \text{if } x_c = \textit{sheep} \\ 20 & \text{if } x_c = \textit{cow} \\ 20 & \text{if } x_c = L_F \end{cases} \quad (46)$$

And we define the pairwise terms between the clique variables as

$$\phi_{c,i}(x_c, x_i) = \begin{cases} 1 & \text{if } x_c \neq L_F \wedge x_c \neq x_i \\ 0 & \text{otherwise.} \end{cases} \quad (47)$$

For simplicity, we set all pairwise terms within the base layer to 0, and disregard them. Then a minimal labelling of the solution occurs when, $x_i = \textit{cow}$ $\forall x_i \in \mathbf{x}^{(0)}$ and $x_c = \textit{sheep}$. This labelling is incoherent, insomuch as we believe at the base scale that a region is *cow*, and at a coarser scale that the same region is *sheep*. Our requirement of *hierarchical consistency* prohibits such solutions by insisting that the minimal cost labelling of auxiliary variables correspond to either the dominant label in base layer, or to the label $L_F$.

The constraint (44) is enforced by construction, weighting the relative magnitude of $\psi_i(l)$ and $\psi_{i,j}(b_j, x_c^{(i)})$ to guarantee that:

$$\psi_i(l) + \sum_{j \in N_i/c} \max_{b_j \in \mathcal{L} \cup \{L_F\}} \psi_{i,j}(b_j, x_c^{(i)}) < 0.5 \sum_{i \in c} k_i \forall l \in \mathcal{L}. \quad (48)$$

If this holds, in the degenerate case where there are only two levels in the hierarchy, and no pairwise connections between the auxiliary variables, our AHRF is exactly equivalent to the $P^n$ model.

At most one $l \in \mathcal{L}$ at a time can satisfy (44), assuming the hierarchy is consistent. Given a labelling for the base layer of the hierarchy $\mathbf{x}^{(1)}$, an optimal labelling for an auxiliary variable in $\mathbf{x}^{(2)}$ associated with some clique must be one of two labels: $L_F$ and some $l \in \mathcal{L}$. By induction, the choice of labelling of any clique in $\mathbf{x}^{(j)} : j \geq 2$ must also be a decision between at most two labels: $L_F$ and some $l \in \mathcal{L}$.

#### 4.4.2 Transformational Optimality under Swap range moves

Transformational optimality for $\alpha\beta$-swap based range moves for AHRFs further requiresthat there are no pairwise connections between variables in the same level of the hierarchy, except in the base layer.

From (14), if an auxiliary variable $x_c$ may take label $\gamma$ or $L_F$, and one of its children $x_i | i \in c$ take label $\delta$ or $L_F$, the cost associated with assigning label $\gamma$ or $L_F$ to $x_c$ is independent of the label of $x_i$ with respect to a given move. This means that under a swap move, a clique currently taking label $\delta \notin \{\alpha, \beta\}$ will continue to do so. This follows from (5) as the cost associated with taking label $\delta$ is only dependent upon the weighted average of child variables taking state $\delta$, and this remains constant. Hence the only clique variables that may have a new optimal labelling under the swap are those currently taking state $\alpha, L_F$ or $\beta$, and these can only transform to one of the states $\alpha, L_F$ or $\beta$. As the range moves map exactly this set of transformations, the move proposed must be transformationally optimal, and consequently the best possible $\alpha\beta$ swap over the energy (13).

#### 4.4.3 Transformational Optimality under Expansion Moves

Using range-expansion moves, we can maintain transformational optimality while incorporating pairwise connections into the hierarchy — provided condition (44) holds, and the energy can be exactly represented in our submodular moves.

In order for this to be the case, the pairwise connections must be both convex over any range $\alpha, L_F, \beta$ and a metric. The only potentials that satisfy this are linear over the orderings $\alpha, L_F, \beta$ $\forall \alpha, \beta$. Hence all pairwise connections must be of the form:

$$\psi_{i,j}(x_i, x_j) = \begin{cases} 0 & \text{if } x_i = x_j \\ K/2 & \text{if } (x_i = L_F \vee x_j = L_F) \wedge x_i \neq x_j \\ K & \text{otherwise.} \end{cases} \quad (49)$$

where $K \in \mathbb{R}_0^+$.

A similar argument to that of the optimality of $\alpha\beta$ swap can be made for $\alpha$-expansion. As the label $\alpha$ is 'pushed' out across the base layer, the optimal labelling of some $x^{(n)}$ where $n \geq 2$ must either remain constant or transition to one of the labels

$L_F$ or $\alpha$. Again, the range moves map exactly this set of transforms and the suggested move is both transformationally optimal, and the best expansion of label $\alpha$ over the higher-order energy of (13).

As the pairwise cost we use in the hierarchy are metric, equation (36) gives us a bound of 4, when they are used in conjecture with $\alpha$-expansion, or range-expansion.

## 5 Graph Constructions for Range-move Expansion

The graph constructions for swap and expansion range moves can be built using graph constructions of Ishikawa (2003) over the range $(\alpha, L_F, \beta)$ or $(\alpha, L_F, \text{old label})$ respectively. Compared to the construct we will present, use of an Ishikawa construction has several disadvantages: it is slow and memory inefficient, requiring approximately 3 times the number of binary variables we do; more importantly, use of the Ishikawa construct in this way would be a non-standard encoding, and mean that these graphs can not be directly combined with other potentials such as Delong et al (2010); Ladicky et al (2010). In this section, we show an alternative graph constructs that uses a single binary variable to encode the state change of each base layer variableand two binary variables for the state change of variables in all other layers.

The move energy will be encoded using one binary variable $t_i$ for each variable $x_i$ in the base layer. This captures the two possible states $\{\alpha, x_i\}$ of the base layer variables after the move. It will use two binary variables $a_c^{(\mathrm{n})}, b_c^{(\mathrm{n})}$ for each variable $x_c^{(\mathrm{n})}$ in the auxiliary layer to encoding their three possible states $\{\alpha, L_F, x_c^{(\mathrm{n})}\}$ under a move, where $x_i$ and $x_c^{(\mathrm{n})}$ are the states of the corresponding variables before the move.

The transformation vector (see figure 3) for the base layer variables is encoded the same way as standard $\alpha$-expansion:

$$T_\alpha(x_i, t_i) = \begin{cases} \alpha & \text{if} \quad t_i = 0 \\ x_i & \text{if} \quad t_i = 1. \end{cases} \tag{50}$$

While the transformation vector for the auxiliary variables is encoded as:

$$T_\alpha(\mathbf{x}_c^{(\mathrm{n})}, a_c^{(\mathrm{n})}, b_c^{(\mathrm{n})}) = \begin{cases} \alpha & \text{if} \quad a_c^{(\mathrm{n})} = 0 \wedge b_c^{(\mathrm{n})} = 0 \\ x_c^{(\mathrm{n})} & \text{if} \quad a_c^{(\mathrm{n})} = 1 \wedge b_c^{(\mathrm{n})} = 1 \\ L_F & \text{if} \quad a_c^{(\mathrm{n})} = 1 \wedge b_c^{(\mathrm{n})} = 0. \end{cases} \tag{51}$$

To prohibit the combination $a_c^{(\mathrm{n})} = 0$ and $b_c^{(\mathrm{n})} = 1$, we add an edge $K(1-a_c^{(\mathrm{n})})b_c^{(\mathrm{n})}$ with sufficiently large $K \to \infty$. The energy is additive, thus we can find equivalent graph constructions for each term separately.

### 5.1 Graph Construction for the Inter-layer Potential

Let us first assume none of the variables currently takes a label $\alpha$ or $L_F$ and consider the inter-layer term

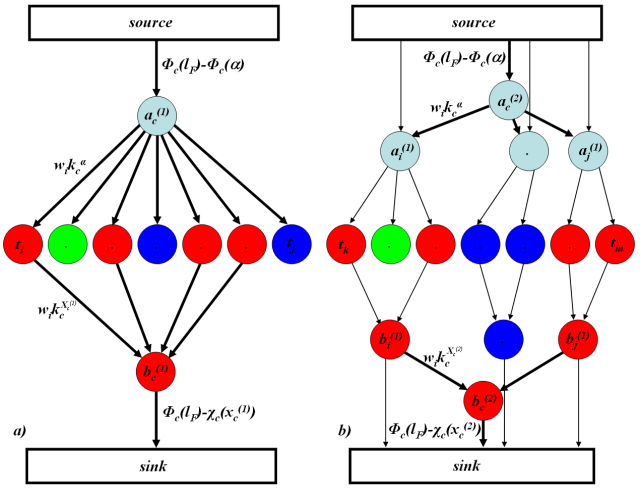



Fig. 4. *A graph construction for the $\alpha$-expansion move of the inter-layer connection between a) base layer and the first auxiliary layer, b) between two auxiliary levels. The colour of variables $t_i$ and $b_c^{(\mathrm{n})}$ corresponds to the label before the move. Each variable $a_c^{(\mathrm{n})}$ is connected to each of the variables $t_i$ respectively $a_i^{(\mathrm{n-1})}$ in the clique of the previous level, each variable $b_c^{(\mathrm{n})}$ is connected to each of the variables $t_i$ respectively $b_i^{(\mathrm{n-1})}$ in the clique of the previous level. Edges modelling corresponding inter-layer connection are bold.*

between the base layer $\mathbf{x}_c$ and the first auxiliary layer:

$$\psi_c^p(\mathbf{x}_c, x_c^{(1)}) = \phi_c(x_c^{(1)}) + \sum_{i \in c} \phi_c(x_c^{(1)}, x_i), \tag{52}$$

where

$$\phi_c(x_c^{(1)}, x_i) = \begin{cases} 0 & \text{if } x_c^{(1)} = L_F \vee x_c^{(1)} = x_i \\ w_i k_c^{x_c^{(1)}} & \text{otherwise.} \end{cases} \tag{53}$$

Writing

$$K_{i,c}^{(\mathrm{n})} = w_i k_c^{x_c^{(\mathrm{n})}} \Delta(x_i = x_c^{(\mathrm{n})}) \tag{54}$$

and

$$\chi_c(x_c^{(\mathrm{n})}) = \phi(x_c^{(\mathrm{n})}) + \sum_{i \in c} w_i k_c^{x_c^{(\mathrm{n})}} \Delta(x_i \neq x_c^{(\mathrm{n})}) \tag{55}$$

The move energy of this potential is:
$\psi_c^p(\mathbf{t}_c, a_c^{(1)}, b_c^{(1)}) =$

$$\begin{cases} \phi_c(\alpha) + \sum_{i \in c} w_i k_c^\alpha t_i & \text{if } a_c^{(1)} = 0 \wedge b_c^{(1)} = 0 \\ \chi_c(x_c^{(1)}) + \sum_{i \in c} K_{i,c}^{(1)}(1 - t_i) & \text{if } a_c^{(1)} = 1 \wedge b_c^{(1)} = 1 \\ \phi_c(L_F) & \text{if } a_c^{(1)} = 1 \wedge b_c^{(1)} = 0, \end{cases} \tag{56}$$

The move energy can be transformed into:

$$\begin{aligned} \psi_c^p(\mathbf{t}_c, a_c^{(1)}, b_c^{(1)}) =& \phi_c(\alpha) + \chi_c(x_c^{(1)}) - \phi_c(L_F) \\ &+ \sum_{i \in c} w_i k_c^\alpha t_i (1 - a_c^{(1)}) \\ &+ (\phi_c(L_F) - \phi_c(\alpha))\, a_c^{(1)} \\ &+ \sum_{i \in c} w_i k_c^{x_c^{(1)}} \Delta(x_i = x_c^{(1)})(1 - t_i) b_c^{(1)} \\ &+ (\phi_c(L_F) - \chi_c(x_c^{(1)}))\,(1 - b_c^{(1)}). \end{aligned} \tag{57}$$

The equivalence can be shown by checking the value of the transformed move energy for each combination of $a_c^{(1)}$ and $b_c^{(1)}$. The move energy is pairwise submodular and thus represents our inter-layer potential. The graph is equivalent to the Robust-$P^N$ graph construction in (Kohli et al, 2008b).

For the inter-layer potential between two auxiliary layers $\mathbf{x}^{(n)}$ and $\mathbf{x}^{(n-1)}$ where $n > 1$, the pairwise cost becomes:

$$\phi_c(x_c^{(n)}, x_d^{(n-1)}) = \begin{cases} 0 & \text{if } x_c^{(n)} = L_F \\ & \vee x_c^{(n)} = x_d^{(n-1)} \\ w_d k_c^{x_c^{(n)}} & \text{otherwise.} \end{cases} \quad (58)$$

The condition $x_c^{(n)} = x_d^{(n-1)}$ is satisfied if both auxiliary variables satisfy $a_c^{(n)} = a_d^{(n-1)}$ and $b_c^{(n)} = b_d^{(n-1)}$. A label of a child is not consistent with a label $\alpha$ if $a_i^{(n-1)} = 1$, a label of a child is not consistent with an old label if $b_i^{(n-1)} = 0$. Thus, the move energy of this potential is:
$\psi_c^p(\mathbf{a}^{(n-1)}, \mathbf{b}^{(n-1)}, a_c^{(n)}, b_c^{(n)}) =$

$$\begin{cases} \phi_c(\alpha) + \sum_{i \in c} w_i k_c^{\alpha} a_i^{(n-1)} & \text{if } a_c^{(n)} = 0 \wedge b_c^{(n)} = 0 \\ \chi_c(x_c^{(n)}) + \sum_{i \in c} K_{i,c}^{(n)} (1 - b_i^{(n-1)}) & \text{if } a_c^{(n)} = 1 \wedge b_c^{(n)} = 1 \\ \phi_c(L_F) & \text{if } a_c^{(n)} = 1 \wedge b_c^{(n)} = 0, \end{cases}$$

As with the previous case, the move energy can transformed into:
$\psi_c^p(\mathbf{a}^{(n-1)}, \mathbf{b}^{(n-1)}, a_c^{(n)}, b_c^{(n)}) =$

$$\begin{aligned} & \phi_c(\alpha) + \chi_c(x_c^{(n)}) - \phi_c(L_F) \\ & + \sum_{i \in c} w_i k_c^{\alpha} a_i^{(n-1)} (1 - a_c^{(n)}) + (\phi_c(L_F) - \phi_c(\alpha)) a_c^{(n)} \\ & + \sum_{i \in c} w_i k_c^{x_c^{(n)}} \Delta(x_i^{(n-1)} = x_c^{(n)}) (1 - b_i^{(n-1)}) b_c^{(n)} \\ & + (\phi_c(L_F) - \chi_c(x_c^{(n)}))(1 - b_c^{(n)}). \end{aligned} \quad (59)$$

The graph constructions for both cases of inter-layer connection are given in figure 4.

### 5.1.1 Graph Construction for the Pairwise Potentials of the Auxiliary Variables

We assume that the pairwise potentials considered are in the same form as equation (49). There are two case to consider: In the first case, $x_c^{(n)} = x_d^{(n)}$, and the move energy of the pairwise potentials between auxiliary variables is:
$\psi_{cd}^p(a_c^{(n)}, b_c^{(n)}, a_d^{(n)}, b_d^{(n)}) =$

$$\begin{cases} 0 & \text{if } a_c^{(n)} = a_d^{(n)} \wedge b_c^{(n)} = b_d^{(n)} \\ \frac{K}{2} & \text{if } (a^{(n)} \neq a_d^{(n)} \wedge b_c^{(n)} = b_d^{(n)}) \\ & \vee (a^{(n)} = a_d^{(n)} \wedge b_c^{(n)} \neq b_d^{(n)}) \\ K & \text{if } a_c^{(n)} \neq a_d^{(n)} \wedge b_c^{(n)} \neq b_d^{(n)}. \end{cases} \quad (60)$$

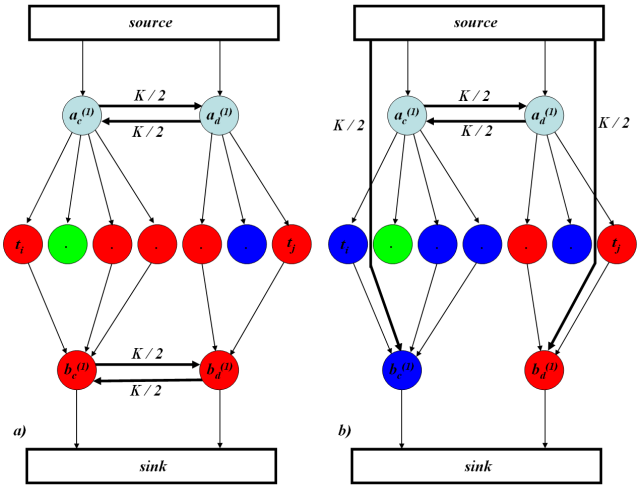



Fig. 5. *A graph construction for the $\alpha$-expansion move of the pairwise potential on the auxiliary level if the label before the move was a) the same, b) different. The colours of variables $t_i$ and $b_c^{(n)}$ correspond to the label before the move. Edges modelling corresponding pairwise potentials are shown in bold.*

This move energy can be transformed into a pairwise submodular one as:

$$\begin{aligned} \psi_{cd}^p(a_c^{(n)}, b_c^{(n)}, a_d^{(n)}, b_d^{(n)}) = & \frac{K}{2} a_c^{(n)} (1 - a_d^{(n)}) + \frac{K}{2} (1 - a_c^{(n)}) a_d^{(n)} \\ & + \frac{K}{2} b_c^{(n)} (1 - b_d^{(n)}) + \frac{K}{2} (1 - b_c^{(n)}) b_d^{(n)}. \end{aligned} \quad (61)$$

The equivalence can be seen by checking all possible combinations of $a_c^{(n)}$, $b_c^{(n)}$, $a_d^{(n)}$ and $b_d^{(n)}$.

In the case where $x_c^{(n)} \neq x_d^{(n)}$ the move energy of the pairwise potential between auxiliary variables is:
$\psi_{cd}^p(a_c^{(n)}, b_c^{(n)}, a_d^{(n)}, b_d^{(n)}) =$

$$\begin{cases} 0 & \text{if } a_c^{(n)} = a_c^{(n)} \wedge b_c^{(n)} = b_d^{(n)} = 0 \\ K & \text{if } a_c^{(n)} = a_d^{(n)} \wedge b_c^{(n)} = b_d^{(n)} = 1 \\ \frac{K}{2} & \text{if } (a^{(n)} \neq a_d^{(n)} \wedge b_c^{(n)} = b_d^{(n)}) \\ & \vee (a^{(n)} = a_d^{(n)} \wedge b_c^{(n)} \neq b_d^{(n)}) \\ K & \text{if } a_c^{(n)} \neq a_d^{(n)} \wedge b_c^{(n)} \neq b_d^{(n)}, \end{cases} \quad (62)$$

and the equivalent pairwise submodular move energy is:

$$\begin{aligned} \psi_{cd}^p(a_c^{(n)}, b_c^{(n)}, a_d^{(n)}, b_d^{(n)}) = & \frac{K}{2} a_c^{(n)} (1 - a_d^{(n)}) \\ & + \frac{K}{2} (1 - a_c^{(n)}) a_d^{(n)} + \frac{K}{2} b_c^{(n)} + \frac{K}{2} b_d^{(n)}. \end{aligned} \quad (63)$$

This equivalence holds for the $3 \times 3$ allowed configurations of $a_c^{(n)}$, $b_c^{(n)}$, $a_d^{(n)}$ and $b_d^{(n)}$. Graph constructions for both cases $x_c^{(n)} = x_d^{(n)}$ and $x_c^{(n)} \neq x_d^{(n)}$ are given in figure 5.

All the previous constructions were made under the assumption that none of the variables already takes the label $\alpha$ or $L_F$. If a variable in the base layer already takes the label $\alpha$, the problem is equivalent to changing each $t_i$ to 0 in all pairwise submodular expressions. If the variable in the auxiliary layer already takes the the label $\alpha$, both $a_c^{(n)}$ and $b_c^{(n)}$ have

to be changed to $0$ in all derived expressions. In the case that the auxiliary variable takes the label $L_F$, the variable can take only label $\alpha$ and label $L_F$ after the move and thus $b_c^{(\mathrm{n})}$ has to be changed to $0$. Setting the label of any variable to $0$ is equivalent to tying it to the sink or equivalently changing each incoming edge to this variable to the edge going to the sink. Setting the label of any variable to $1$ is equivalent to tying it to the source or equivalently changing each outgoing edge of this variable to the edge going to the source. The infinite edge between $a_c^{(\mathrm{n})}$ and $b_c^{(\mathrm{n})}$ is not necessary if the hierarchy is *hierarchically consistent*, see 4.4.1.

# 6 Hierarchical RFs for Semantic Segmentation

Having described the definition and intuition behind the AHRF framework, in this section we describe the set of potentials we use in the object-class segmentation problem. This set includes unary potentials for both pixels and segments, pairwise potentials between pixels and between segments and connective potentials between pixels and their containing segments.

In the previous sections we decomposed the energy (14) into a set of potentials $\psi_c(\mathbf{x}_c)$. In this section we will decompose them further, writing $\psi_c(\mathbf{x}_c) = \lambda_c \xi_c(\mathbf{x}_c)$, where $\xi_c$ is a feature based potential over $c$ and $\lambda_c$ its weight. Initially we will discuss the learning of potentials $\xi_c(\mathbf{x}_c)$, and later discuss the learning of the weights $\lambda_c$.

For our application we used potentials defined over a three-level hierarchy. We refer to elements of each layer as pixels, segments and super-segments respectively. Unsupervised segments are initially found using multiple applications of a fine scale mean-shift algorithm (Comaniciu and Meer, 2002). "Super-segments" are based upon a coarse mean-shift segmentation, performed over the result of the previous segmentations.

## 6.1 Features

Several well-engineered features were experimentally found to be more discriminative then the raw RGB values of pixels. In our application we use textons (Malik et al, 2001), local binary patterns (Ojala et al, 1994), multi-scale (Bosch et al, 2007) dense SIFT (Lowe, 2004) and opponent SIFT (van de Sande et al, 2008). Textons (Malik et al, 2001) are defined as a clustered $16$-dimensional response to $16$ different filters - Gaussian, Gaussian derivative and Laplacian filters at different scales. Local binary pattern (Ojala et al, 1994) is a $8$-dimensional binary feature consisting of $8$ comparisons of the intensity value of the centre pixel with its neighbours. The SIFT (Lowe, 2004) feature contains the histograms of gradients of $4 \times 4$ cells quantised into $8$ bins. The resulting $128$ dimensional vector is normalised to $1$. Opponent SIFT (van de Sande et al, 2008) is a variant of coloured SIFT and is built of separate histograms of gradients for $3$ channels in the transformed colour space. All features except local binary patterns are quantised to $150$ clusters using standard $K$-means clustering.

## 6.2 Unary Potentials from Pixel-wise Features

Unary potentials from pixel-wise features are derived from *TextonBoost* (Shotton et al, 2006), and allow us to perform texture based segmentation, at the pixel level, within the same framework. The features used for constructing these potentials are computed on every pixel of the image, and are also called *dense* features. TextonBoost estimates the probability of a pixel taking a certain label by boosting weak classifiers based on a set of shape filter responses. The shape filters are defined by a texton $t$ and rectangular region $r$. Their response $v_{[t,r]}(i)$ for a given point $i$ is the number of textons $t$ in the region $r$ placed relative to the point $i$. Corresponding weak classifiers are decision stumps, which split on a shape filter response and one of a set of thresholds. The most discriminative weak classifiers are found using multi-class Gentle Ada-Boost (Torralba et al, 2004).

We observed that textons were unable to discriminate between some classes of similar textures. This motivated us to extend the *TextonBoost* framework by boosting classifiers defined on multiple dense features (such as colour, textons, histograms of oriented gradients (HOG) (Dalal and Triggs, 2005), and pixel location) together. Generalised shape filters are defined by feature type $f$, feature cluster $t$ and rectangular region $r$. Their response $v^f_{[t,r]}(i)$ for given point $i$ is the number of features of type $f$ belonging to cluster $t$ in the region $r$ placed relative to the point $i$. The pool of weak classifiers contains decision stumps based on the generalised shape filters against a set of thresholds $\theta$. See (Torralba et al, 2004; Shotton et al, 2006) for further details of the procedure. Our results show that the boosting of multiple features together results in a significant improvement of the performance (note the improvement from the 72% of (Shotton et al, 2006) to 81% of our similar pixel-based RF in figure 1). Further improvements were achieved using exponentially instead of linearly growing thresholds and Gaussian instead of uniform distribution of rectangles around the point. The potential is incorporated into the framework in the standard way as a negative log-likelihood:

$$\phi_i(x_i = l) = -\log \frac{e^{H_l(i)}}{\sum_{l' \in \mathcal{L}} e^{H_{l'}(i)}} = -H_l(i) + K_i, \quad (64)$$

where $H_l(i)$ is the Ada-Boost classifier response for a label $l$ and a pixel $i$ and $K_i = \log \sum_{l' \in \mathcal{L}} e^{H_{l'}(i)}$ a normalising constant.

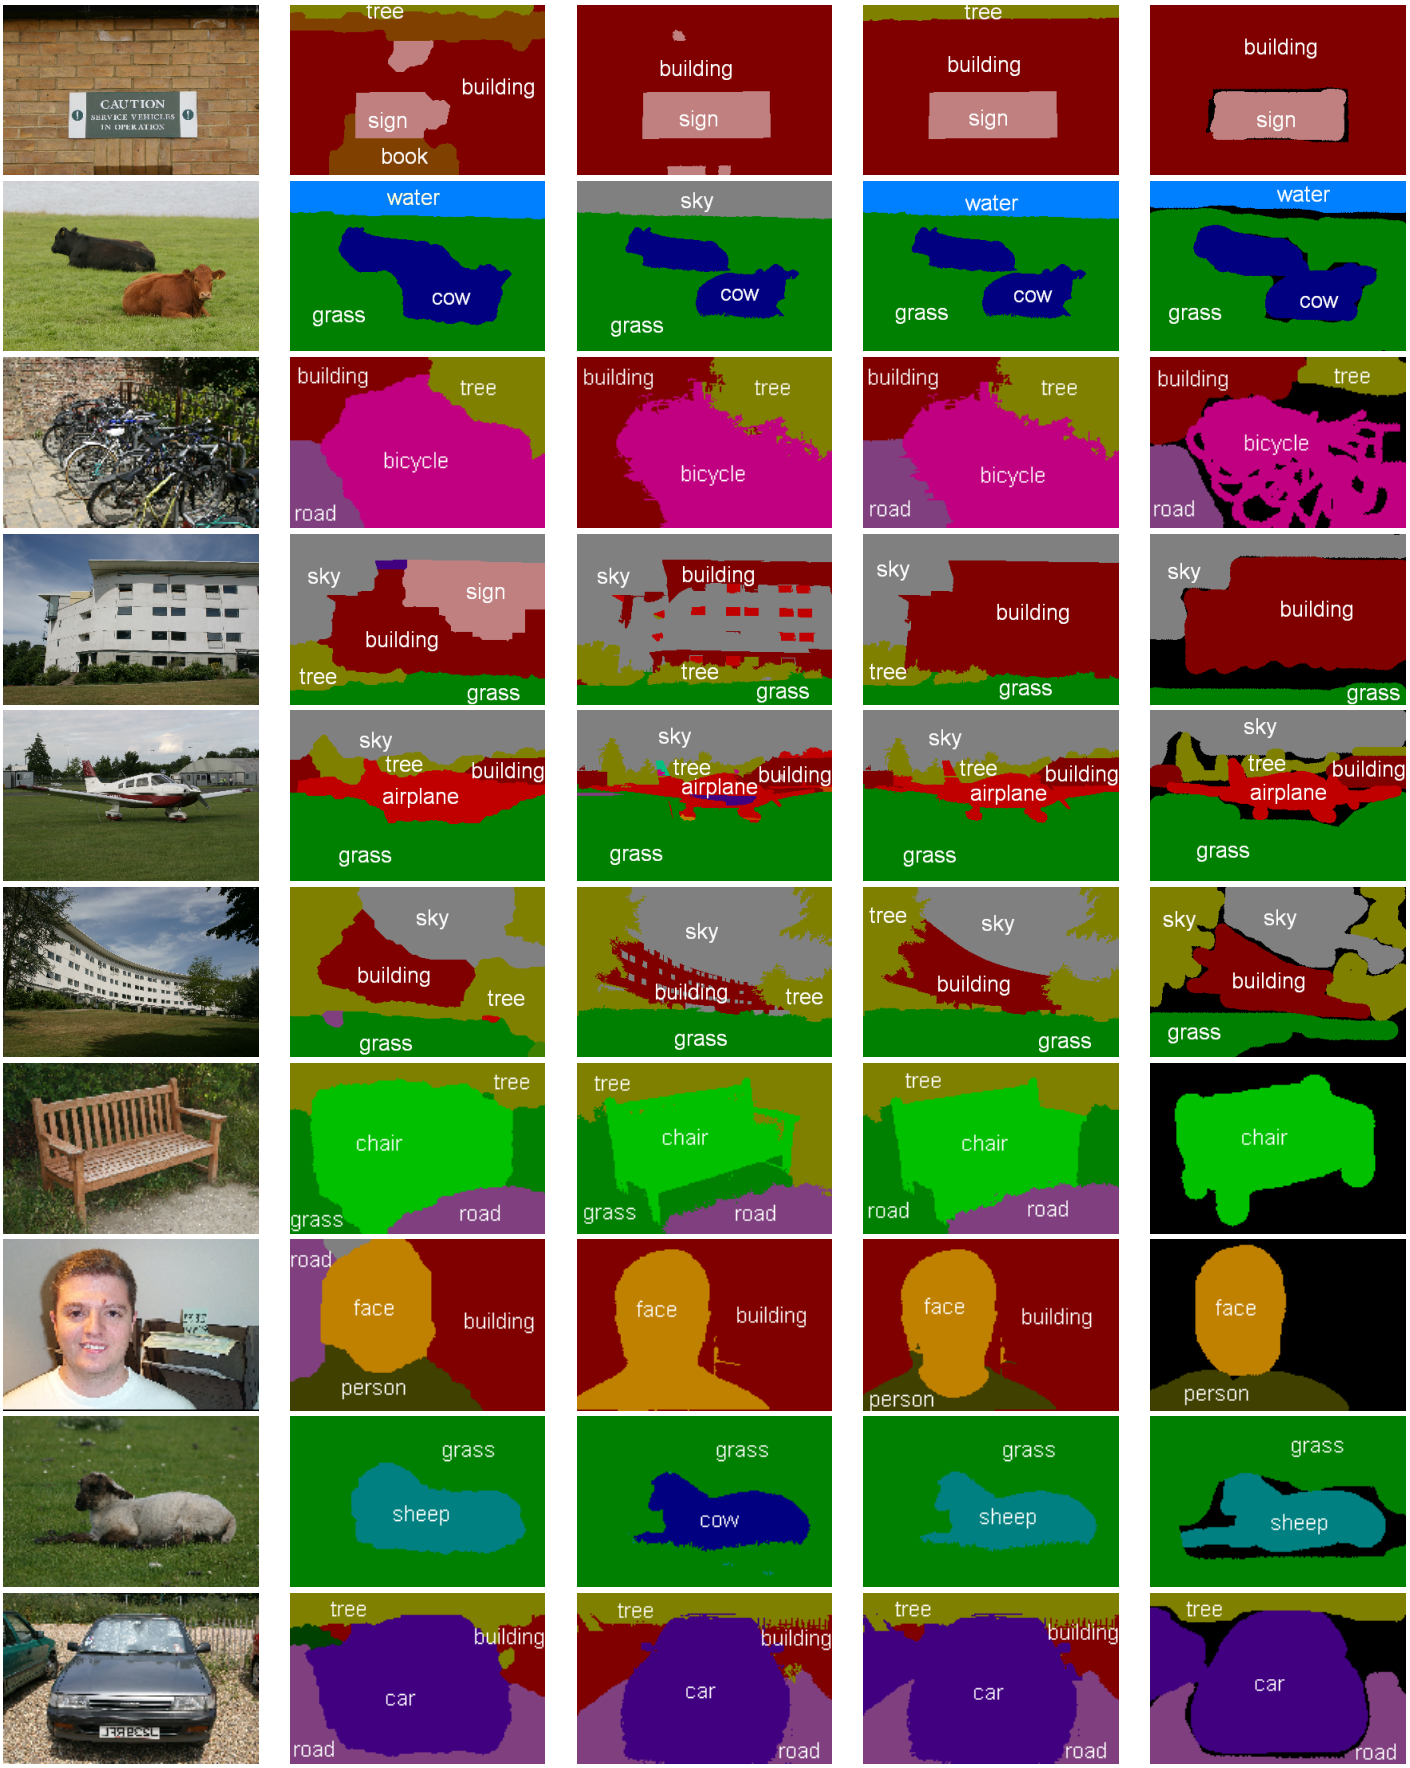
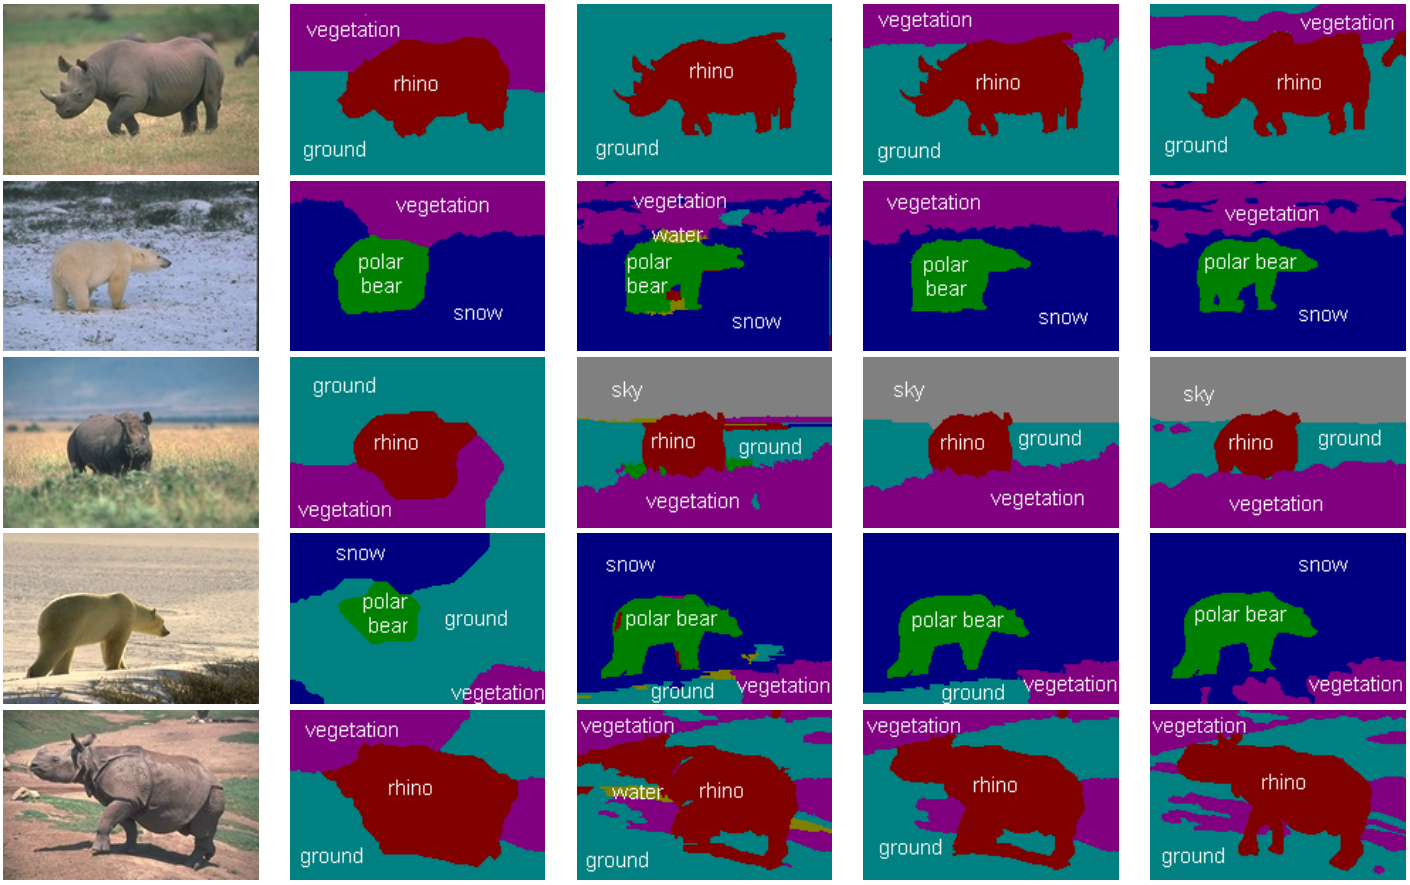
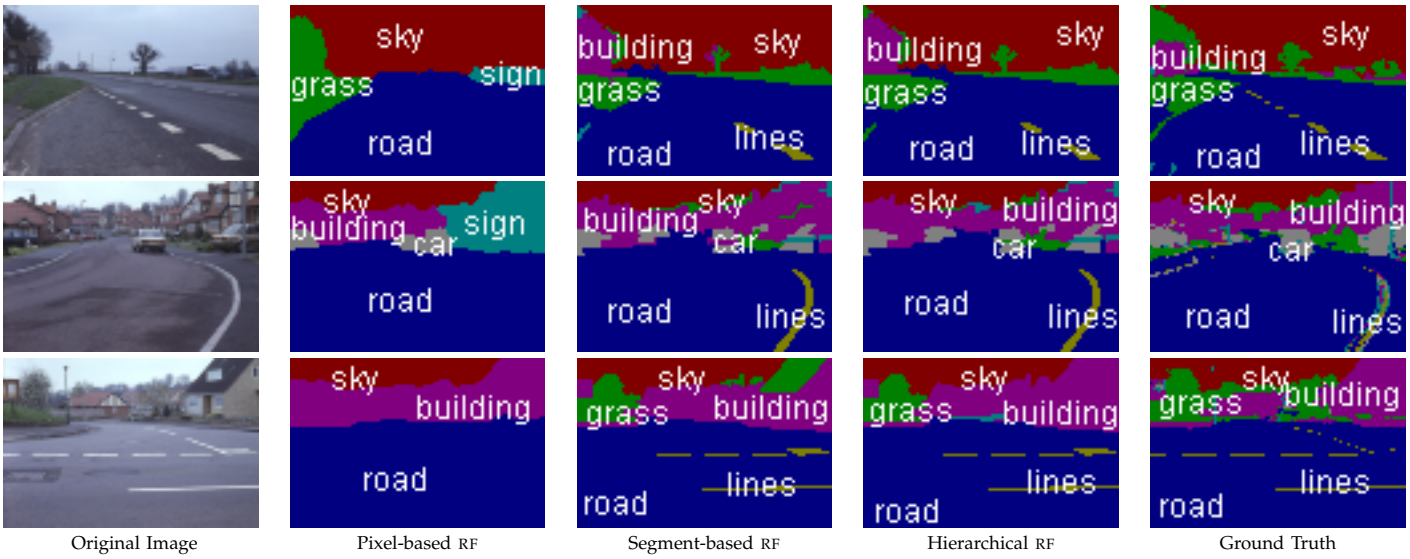



Fig. 6. *Qualitative results on the MSRC-21 data set comparing non-hierarchical (i.e. pairwise models) approaches defined over pixels (similar to TextonBoost (Shotton et al, 2006)) or segments (similar to (Yang et al, 2007; Pantofaru et al, 2008; Russell et al, 2006) described in section 3) against our hierarchical model. Regions marked black in the hand-labelled ground truth image are unlabelled.*




Original Image | Pixel-based RF | Segment-based RF | Hierarchical RF | Ground Truth

Fig. 7. *Qualitative results on the Corel data set comparing approaches defined over pixels or segments against the hierarchical model.*




Original Image | Pixel-based RF | Segment-based RF | Hierarchical RF | Ground Truth

Fig. 8. *Qualitative results on the Sowerby data set comparing approaches defined over pixels or segments against the hierarchical model.*

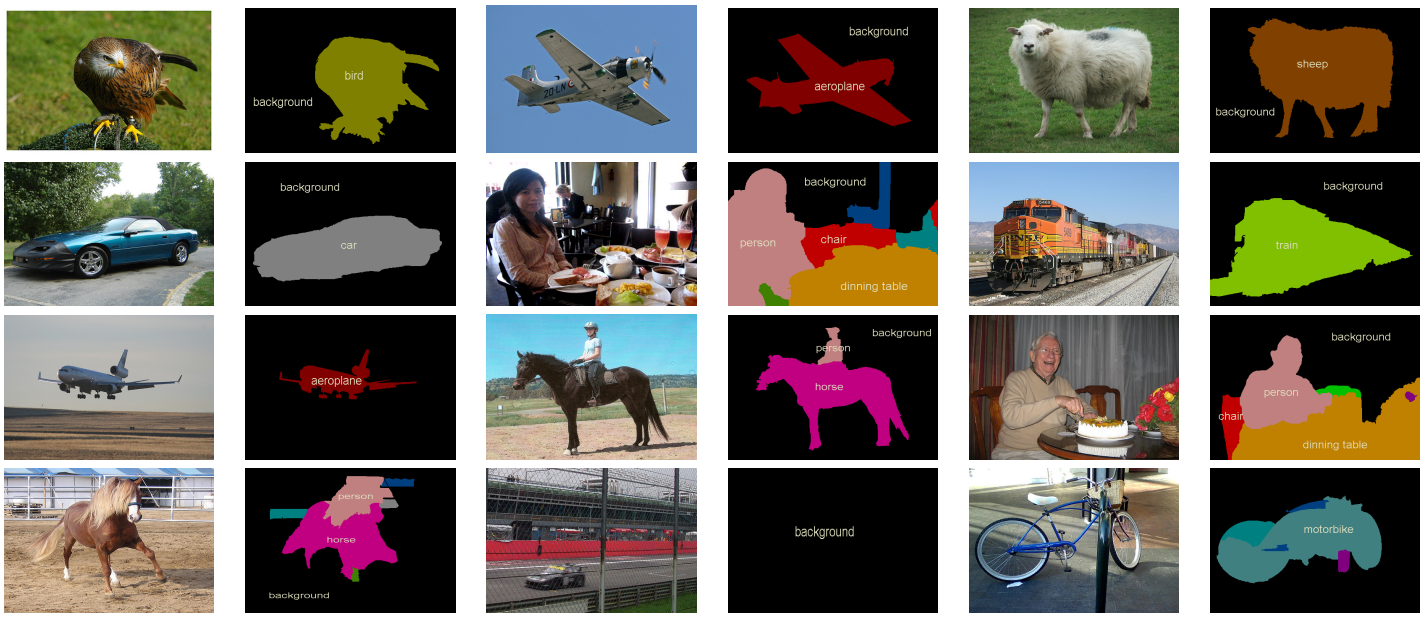
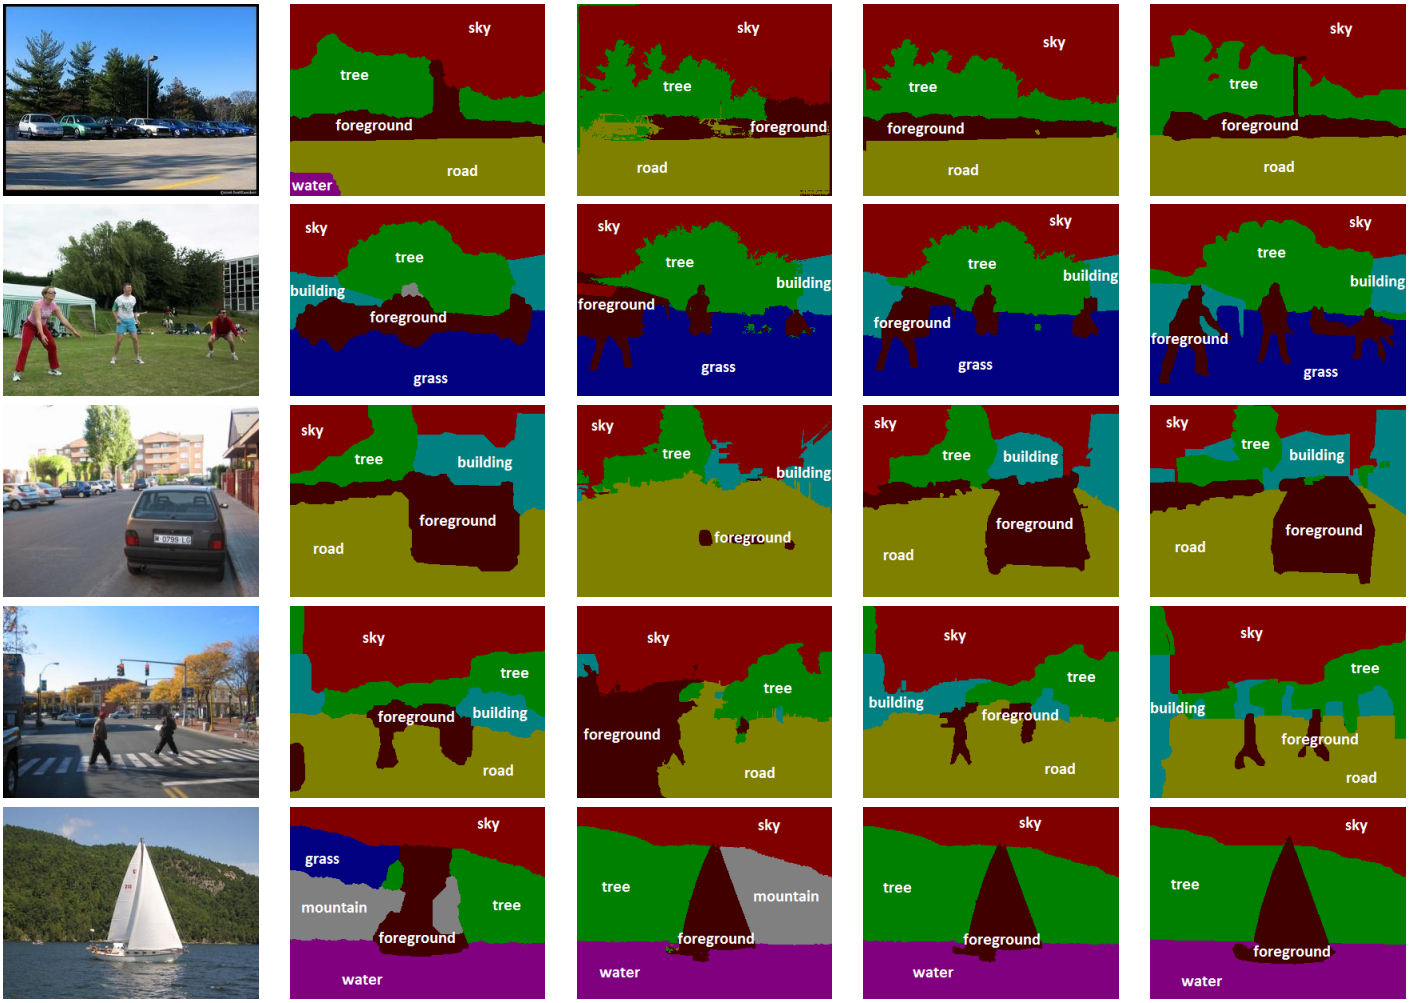



Fig. 9. *Qualitative results on the VOC-2008 data set.* **Successful segmentations** *(top 3 rows) and* **standard failure cases** *(bottom) - from left to right, context error, detection failure and misclassification.*




Original Image | Pixel-based RF | Segment-based RF | Hierarchical RF | Ground Truth

Fig. 10. *Qualitative results on the Stanford data set.*

| | Global | Average | Building | Grass | Tree | Cow | Sheep | Sky | Aeroplane | Water | Face | Car | Bicycle | Flower | Sign | Bird | Book | Chair | Road | Cat | Dog | Body | Boat |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| (Shotton et al, 2008) | 72 | 67 | 49 | 88 | 79 | **97** | **97** | 78 | 82 | 54 | 87 | 74 | 72 | 74 | 36 | 24 | 93 | 51 | 78 | **75** | 35 | 66 | 18 |
| (Shotton et al, 2006) | 72 | 58 | 62 | **98** | 86 | 58 | 50 | 83 | 60 | 53 | 74 | 63 | 75 | 63 | 35 | 19 | 92 | 15 | 86 | 54 | 19 | 62 | 07 |
| (Batra et al, 2008) | 70 | 55 | 68 | 94 | 84 | 37 | 55 | 68 | 52 | 71 | 47 | 52 | 85 | 69 | 54 | 05 | 85 | 21 | 66 | 16 | 49 | 44 | 32 |
| (Yang et al, 2007) | 75 | 62 | 63 | **98** | **89** | 66 | 54 | 86 | 63 | 71 | 83 | 71 | 79 | 71 | 38 | 23 | 88 | 23 | 88 | 33 | 34 | 43 | **32** |
| Pixel-based RF | 84 | 76 | 73 | 93 | 84 | 77 | 84 | 96 | **85** | 91 | 90 | **86** | 91 | 95 | **91** | **41** | 92 | **53** | 87 | 65 | **77** | **70** | 17 |
| Segment-based RF | 81 | 66 | 80 | **98** | 83 | 64 | 81 | **99** | 59 | 89 | 85 | 68 | 68 | **98** | 76 | 26 | 85 | 39 | 84 | 30 | 49 | 50 | 07 |
| Hierarchical RF | **87** | **78** | **81** | 96 | **89** | 74 | 84 | **99** | 84 | **92** | **90** | **86** | **92** | **98** | **91** | 35 | **95** | **53** | **90** | 62 | **77** | **70** | 12 |

TABLE 1

*Quantitative results on the MSRC data set. The table shows % pixel recall measure $N_{ii}/\sum_j N_{ij}$ for different object classes. 'Global' refers to the overall error $\frac{\sum_{i\in\mathcal{L}} N_{ii}}{\sum_{i,j\in\mathcal{L}} N_{ij}}$, while 'average' is $\sum_{i\in\mathcal{L}} \frac{N_{ii}}{|\mathcal{L}|\sum_{j\in\mathcal{L}} N_{ij}}$. $N_{ij}$ refers to the number of pixels of label $i$ labelled $j$. The comparison suggests that the incorporation of the classifiers at different scales leads to a significant improvement of the performance.*

| | Global | Average | Rhino/Hippo | Polar Bear | Water | Snow | Grass | Ground | Sky |
|---|---|---|---|---|---|---|---|---|---|
| (Batra et al, 2008) | 83 | **85** | 87 | **92** | 82 | **91** | 66 | **83** | **94** |
| Pixel-based RF | 76 | 72 | 80 | 85 | 88 | 83 | 75 | 57 | 35 |
| Segment-based RF | 80 | 78 | **92** | 65 | 91 | 84 | 81 | 67 | 73 |
| Hierarchical RF | **84** | **85** | **92** | 82 | **94** | 88 | **83** | 77 | 76 |

TABLE 2

*Quantitative results on the Corel data set. Segment-based method tend to outperform pixel-based ones. Due to the insufficient amount of data the performance largely depends on the random split of the data. The same error measure as for the MSRC data set has been used. Combining classifiers at different scales led to an improvement of the performance.*

| | Global | Average | Sky | Grass | Road Line | Road | Building | Sign | Car |
|---|---|---|---|---|---|---|---|---|---|
| Pixel-based RF | 83 | 47 | 92 | 83 | 00 | 89 | 28 | 07 | 33 |
| Segment-based RF | 89 | **60** | **94** | 87 | **47** | 94 | **61** | **10** | 35 |
| Hierarchical RF | **91** | **64** | **97** | **96** | 45 | **98** | 59 | 09 | **43** |

TABLE 3

*Quantitative results on the Sowerby data set. The segment-based method tend to outperform pixel-based ones. Context-based pixel method could not capture small objects due to the insufficient size of the images. The same error measure as for the MSRC data set has been used. As with other data sets, the hierarchical* RF *outperformed both approaches over single scale.*

| | Average | Background | Aeroplane | Bicycle | Bird | Boat | Bottle | Bus | Car | Cat | Chair | Cow | Dining table | Dog | Horse | Motor bike | Person | Potted plant | Sheep | Sofa | Train | TV/monitor |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| XRCE | **25.4** | 75.9 | 25.8 | 15.7 | 19.2 | **21.6** | 17.2 | **27.3** | **25.5** | **24.2** | 7.9 | **25.4** | 9.9 | **17.8** | **23.3** | **34.0** | 28.8 | **23.2** | **32.1** | **14.9** | **25.9** | **37.3** |
| UIUC / CMU | 19.5 | **79.3** | 31.9 | **21.0** | 8.3 | 6.5 | **34.3** | 15.8 | 22.7 | 10.4 | 1.2 | 6.8 | 8.0 | 10.2 | 22.7 | 24.9 | 27.7 | 15.9 | 4.3 | 5.5 | 19.0 | 32.1 |
| MPI | 12.9 | 75.4 | 19.1 | 7.7 | 6.1 | 9.4 | 3.8 | 11.0 | 12.1 | 5.6 | 0.7 | 3.7 | 15.9 | 3.6 | 12.2 | 16.1 | 15.9 | 0.6 | 19.7 | 5.9 | 14.7 | 12.5 |
| Hierarchical RF | 20.1 | 75.0 | **36.9** | 4.8 | **22.2** | 11.2 | 13.7 | 13.8 | 20.4 | 10.0 | **8.7** | 3.6 | **28.3** | 6.6 | 17.1 | 22.6 | **30.6** | 13.5 | 26.8 | 12.1 | 20.1 | 24.8 |

TABLE 4

*Quantitative analysis of VOC2008 results (Everingham et al, 2008) based upon performance the intersection vs. union criteria $\left(\frac{\sum_{i\in\mathcal{L}} N_{ii}}{|\mathcal{L}|(-N_{ii}+\sum_{j\in\mathcal{L}} N_{ij}+N_{ji})}\right)$. All other methods used classification and detection priors trained over a much larger data set that included unsegmented images. The reported results are from the actual challenge.*

| | Global | Average | Sky | Tree | Road | Grass | Water | Building | Mountain | Foreground |
|---|---|---|---|---|---|---|---|---|---|---|
| Pixel-based RF | 77.9 | 67.8 | 90.9 | 71.3 | 85.9 | 82.7 | 70.0 | 76.9 | 8.9 | **63.7** |
| Segment-based RF | 77.3 | 68.1 | 94.3 | 65.1 | 89.9 | **88.4** | 70.6 | 77.9 | **17.3** | 49.0 |
| Hierarchical RF | **80.9** | **70.4** | **94.8** | **71.6** | **90.6** | 88.0 | **73.5** | **82.2** | 10.2 | 59.9 |

TABLE 5

*Quantitative results on the Stanford data set. The table shows % pixel recall measure $N_{ii}/\sum_j N_{ij}$ for different object classes. 'Global' refers to the overall score $\frac{\sum_{i\in\mathcal{L}} N_{ii}}{\sum_{i,j\in\mathcal{L}} N_{ij}}$, while 'average' is $\sum_{i\in\mathcal{L}} \frac{N_{ii}}{|\mathcal{L}|\sum_{j\in\mathcal{L}} N_{ij}}$. $N_{ij}$ refers to the number of pixels of label $i$ labelled $j$. The comparison suggests that the incorporation of the classifiers at different scales leads to a significant improvement of the performance. Other papers typically report only the global overall score. For comparison, the pixel CRF approach of (Gould et al, 2009a) gets 74.3%, their region-based method (Gould et al, 2009a) 76.4%, the follow-up paper of (Kumar and Koller, 2010) 79.4%, the segment-based method of (Tighe and Lazebnik, 2010) 77.5% and the image parsing using recursive neural networks of (Socher et al, 2011) 78.1%.*

### 6.3 Histogram-based Segment Unary Potentials

We now explain the unary potential defined over segments and super-segments. For many classification and recognition problems, the distributions of pixel-wise feature responses are more discriminative than any feature alone. For instance, the sky can be either 'black' (night) or 'blue' (day), but is never 'half-black' and 'half-blue'. This consistency in the colour of object instances can be used as a region based feature for improving object segmentation results. The unary potentials of auxiliary segment variables are learnt using multi-class Gentle Ada-Boost (Torralba et al, 2004) over the normalised histograms of multiple clustered pixel-wise features. The pool of week classifiers are decision stumps that return 1 if more that $\theta$ % of a segment belongs to cluster $t$ of feature $f$, and 0 otherwise. The selection and learning procedure is identical to (Torralba et al, 2004).

The segment potential is incorporated into the energy as:

$$\phi_c(x^{(1)} = l) = \lambda_s|c|\min(-H_l(c) + K_c, \alpha^h), \quad (65)$$

$$\phi_c(x^{(1)} = L_F) = \lambda_s|c|\alpha^h, \quad (66)$$

where $H_l(c)$ is the response given by the Ada-boost classifier to clique $c$ taking label $l$, $\alpha^h$ a truncation threshold and $K = \log\sum_{l'\in\mathcal{L}} e^{H_{l'}(c)}$ a normalising constant.

For our experiments, the cost of pixel labels differing from an associated segment label was set to $k_c^l = (\phi_c(x^{(1)} = L_F) - \phi_c(x^{(1)} = l))/0.1|c|$. This means that up to 10% of the pixels can take a label different to the segment label without the segment variable changing its state to $L_F$.

### 6.4 Pairwise Potentials

The pairwise terms on the pixel level $\psi_{ij}(\cdot)$ take the form of the classical contrast sensitive potentials.

$$\xi^p(x_i, x_j) = \begin{cases} 0 & \text{if } x_i = x_j, \\ g(i,j) & \text{otherwise,} \end{cases} \quad (67)$$

where the function $g(i,j)$ is an edge feature based on the difference in the intensity of colours of neighbouring pixels (Boykov and Jolly, 2001). It is typically defined as:

$$g(i,j) = \theta_p + \theta_v \exp(-\theta_\beta||I_i - I_j||^2), \quad (68)$$

where $I_i$ and $I_j$ are the colour vectors of pixel $i$ and $j$ respectively. These encourage neighbouring pixels in the image (having a similar colour) to take the same label. We refer the reader to (Boykov and Jolly, 2001; Rother et al, 2004; Shotton et al, 2006) for details.

To encourage neighbouring segments with similar texture to take the same label, we used pairwise potentials based on the squared Euclidean distance of normalised histograms of colour between corresponding auxiliary variables:

$$\xi_{cd}^p(x_c^{(1)}, x_d^{(1)}) = \begin{cases} 0 & \text{if } x_c^{(1)} = x_d^{(1)}, \\ g(c,d)/2 & \text{if } (x_c^{(1)} = L_F \wedge x_d^{(1)} \neq L_F) \\ & \vee (x_c^{(1)} \neq L_F \wedge x_d^{(1)} = L_F), \\ g(c,d) & \text{otherwise,} \end{cases} \quad (69)$$

where $g(c,d) = |\mathbf{h}(\mathbf{x}_c^{(1)}) - \mathbf{h}(\mathbf{x}_d^{(1)})|_2^2$ and $\mathbf{h}(\cdot)$ is the normalised histogram of colours of given segment.

## 7 Learning Weights for Hierarchical RFs

Having learnt potentials $\xi_c(\mathbf{x}_c)$ as described earlier, the problem remains of how to assign appropriate weights $\lambda_c$. This weighting, and the training of RF parameters in general is not an easy problem and there is a wide body of literature dealing with it (Blake et al, 2004; Hinton and Osindero, 2006; He et al, 2006; Taskar et al, 2004b). The approach we take to learn these weights uses a coarse to fine, layer-based, local search scheme over a validation set.

We first introduce additional notation: $\mathcal{V}^{(i)}$ refers to variables contained in the $i^{\text{th}}$ layer of the hierarchy, while $\mathbf{x}^{(i)}$ is the labelling of $\mathcal{V}^{(i)}$ associated with a MAP estimate over the truncated hierarchical RF consisting

**Algorithm 1** *Weight Learning Scheme.*

**for** $i$ *from* $n$ *down to* $1$ **do**
  $s_1, s_2, s_h, d_1, d_2, d_h = 1$
  **while** $s_1, s_2$ or $s_h \geq \Theta$ **do**
    **for** $t \in \{1, 2, h\}$ **do**
      $\lambda_t'^{(i)} \leftarrow \lambda_t^{(i)} + d_t s_t$
      Perform MAP estimate of $\mathbf{x}_i$ using $\lambda_t'$ instead of $\lambda_t$
      **if** $C(\mathbf{x}_i)$ *has decreased* **then**
        $\lambda_t \leftarrow \lambda_t'$
      **else**
        $s_t \leftarrow s_t/2$, $d_t \leftarrow -d_t$
      **end if**
    **end for**
  **end while**
**end for**

of the random variables $\mathbf{v}' = \{v \in \mathcal{V}^{(k)} : k \geq i\}$. Given the validation data we can determine a dominant label $L_c$ for each segment $c$, such that $L_F = l$ when $\sum_{i \in l} \Delta(x_i = l) = 0.5|c|$, and if there is no such dominant label, we set $L_c = L_F$.

At every level of the hierarchy, the label of a clique $x_c^{(i)}$ must correspond to the dominant label of this clique in the ground truth (or $L_F$) for its pixels to be correctly labelled. Based on this observation, we propose a simple heuristic which we optimise for each layer.

At each layer, we seek to minimise the discrepancy between the dominant ground truth label of a clique $l_c$, and the value $x_c^{(i)}$ of the MAP estimate. Formally, we choose parameters $\lambda$ to minimise

$$C(\mathbf{x}^{(i)}) = \sum_{c \in \mathcal{V}^{(i)}} \Delta(x_c^{(i)} \neq l_c \wedge l_c \neq L_F). \qquad (70)$$

We optimise (70) layer by layer. The full method is given in algorithm 1, where we use $\lambda_1^{(i)}$ to refer to the weighting of unary potentials in the $i^{\text{th}}$ layer, $\lambda_2^{(i)}$ the weight of the pairwise terms and $\lambda_h^{(i+1)}$ a scalar modifier of all terms in the $(i+1)^{\text{th}}$ layer or greater. $\Theta$ is an arbitrary constant that controls the precision of the final assignment of $\lambda$.

An alternative and elegant approach to this is that of (Finley and Joachims, 2008) which we intend to investigate in future work.

## 8 Experiments

We evaluated the performance of our framework on four data sets: Corel, Sowerby, PASCAL VOC 2008 (Everingham et al, 2008) and MSRC-21 (Shotton et al, 2006)

**MSRC-21** The MSRC segmentation data set contains $591$ images of resolution $320 \times 213$ pixels, accompanied with a hand labelled object segmentation of $21$ object classes. Pixels on the boundaries of objects are not labelled in these segmentations. The division into training, validation and test sets occupied 45%, 10% and 45% of the images. Methods are typically compared using global criteria or average-per-class recall criteria (see table 1 for details). For these experiments, the hierarchy was composed of 3 pairs of nested segmentations. The parameters of the mean-shift kernels were chosen as $(6, 5), (12, 10)$; $(6, 7.5), (12, 15)$; and $(6, 9), (12, 18)$. The first value refers to the planar distance between points, and the second refers to the Euclidean distance in the LUV colour space. A quantitative comparison of performance with other methods is given in table 1. Qualitative results are given in figure 6.

**Corel** The Corel segmentation data set contains $100$ images of resolution $180 \times 120$ pixels of natural scenarios, with a hand labelled object segmentation of $7$ object classes. The division into training and test sets occupied 50% and 50% the images. The same parameters as for MSRC data set have been used due to an insufficient amount of data. Unlike in MSRC data-set, segment-based methods performed better than pixel-based (see table 2 for more details). Qualitative results are given in figure 7.

**Sowerby** The Sowerby segmentation data set contains $106$ images of resolution $96 \times 64$ pixels of road scenes, with a hand labelled object segmentation of $7$ object classes. The division into training and test sets occupied 50% and 50% the images. As with the Corel data set, there was insufficient training data to tune parameters on a validation set, and the parameters tuned for the MSRC data-set were used. Segment-based methods perform better than pixel-based (see table 3 for more details). Small classes performed very badly due to their highly variable appearance and insufficient training and test data. Qualitative results are given in figure 8.

**Stanford** The Stanford segmentation data set contains $715$ images of resolution $320 \times 240$ pixels of road scenes, with a hand labelled object segmentation of $7$ background classes one one foreground class. The set of images are split into two equally sized training and testing sets. The same parameters as for the MSRC data set were used. Quantitative comparison of performance is given in table 5. Qualitative results are given in figure 10.

**PASCAL VOC 2008** This data set was used for the PASCAL Visual Object Category segmentation contest 2008. It is especially challenging given the presence of significant background clutter, illumination effects and occlusions. It contains $511$ training, $512$ validation and $512$ segmented test images of $20$ foreground and $1$ background classes. The organisers also provided $10,057$ images for which only the bounding boxes of the objects present in the image are marked. We did not use these additional images for training our framework. For this data set we used a two-level hierarchy. The methods are evaluated using intersec-

| Method | Best | $E(\text{meth}) - E(\min)$ | $\frac{E(\text{meth})}{E(\min)}$ | Time |
|---|---|---|---|---|
| Range-exp | 265 | 75 | 1.000 | 6.1s |
| Range-swap | 137 | 9034 | 1.059 | 20s |
| $\alpha$-expansion | 109 | 256 | 1.002 | 6.3s |
| $\alpha\beta$ swap | 42 | 9922 | 1.060 | 42s |
| TRW-S | 12 | 38549 | 1.239 | 500s |
| BP | 6 | 13456 | 1.081 | 120s |
| ICM | 5 | 45955 | 1.274 | 25s |

Fig. 12. *Comparison of methods on 295 testing images. From left to right the columns show the number of times they achieved the best energy (including ties), the average difference ($E(\text{method}) - E(\min)$), the average ratio ($E(\text{method})/E(\min)$) and the average time taken. All three inference methods proposed in this paper: the $\alpha$-expansion of section 4.2, and the transformationally optimal range expansion and swap (section 5) significantly outperformed existing inference methods both in speed and accuracy. See table 11 for individual examples.*

tion vs. union criteria (Everingham et al, 2008) that penalises the performance of classes $i$ and $j$ given a mislabelling of $i$ as $j$ (see table 4). It should be emphases that this is not the same as the percentage of pixels correctly labelled. Quantitative comparison of performance with other methods is given in 4. Qualitative results are given in table 9. The only comparable methods used classification and detection priors trained over a much larger set of images. The reported results we show are from the actual challenge, and not the performance of the current implementation.

The hierarchical RF significantly outperformed RF approaches at single scale (pixels, segments) on all data sets. Experimentally, the approach was robust to the choice of the parameters and typically the same parameters performed well on all data sets. This suggests that the improvement of the performance comes from the incorporation of the different discriminative cues across multiple scales.

### 8.1 Comparison of Inference Methods

We evaluate $\alpha$-expansion, $\alpha\beta$ swap, TRW-S, Belief Propagation, Iterated Conditional Modes, and both the expansion and swap based variants of our unordered range moves on the problem of object class segmentation over the MSRC data-set (Shotton et al, 2006), in which each pixel within an image must be assigned a label representing its class, such as grass, water, boat or cow. We express the problem as a three layer hierarchy. Each pixel is represented by a random variable of the base layer. The second layer is formed by performing multiple unsupervised segmentations over the image, and associating one auxiliary variable with each segment - note that this use of several hierarchies results in overlapping segments. The children of each of these variables in $x^{(2)}$ are the variables contained within the segment, and pairwise connections are formed between adjacent segments. Further details are given in section 8 (MSRC).

We tested each algorithm on 295 test images, with an average of 70,000 pixels/variables in the base layer and up to 30,000 variables in a clique, and ran them either until convergence, or for a maximum of 500 iterations. In the table in figure 12 we compare the final energies obtained by each algorithm, showing the number of times they achieved an energy lower than or equal to all other methods, the average difference $E(\text{method}) - E(\min)$ and average ratio $E(\text{method})/E(\min)$. Empirically, the message passing algorithms TRW-S and BP appear ill-suited to inference over these dense hierarchical random fields. In comparison to the graph cut based move making algorithms, they had higher resulting energy, higher memory usage, and exhibited slower convergence.

While it may appear unreasonable to test message passing approaches on hierarchical energies when higher-order formulations such as (Komodakis and Paragios, 2009; Potetz and Lee, 2008) exist, we note that for the simplest hierarchy that contains only one additional layer of nodes and no pairwise connections in this second layer, higher-order and hierarchical message-passing approaches will be equivalent, as inference over the trees that represent higher-order potentials is exact. Similar relative performance by message passing schemes was observed in these cases.

In all tested images both $\alpha$-expansion variants outperformed TRW-S, BP and ICM. These later methods only obtained minimal cost labellings in images in which the optimal solution found contained only one label i.e. they were entirely labelled as grass or water. The comparison also shows that unordered range move variants usually outperform vanilla move making algorithms. The higher number of minimal labellings found by the range-move variant of $\alpha\beta$ swap in comparison to those of vanilla $\alpha$-expansion can be explained by the large number of images in which two labels strongly dominate – unlike standard $\alpha$-expansion both range move algorithms are guaranteed to find a global optimum of such a two label sub-problem. The typical behaviour of all methods alongside the lower bound of TRW-S can be seen in figure 12 and further, alongside qualitative results, in figures 11.

## 9 Conclusion

This work is a generalisation of many previous super-pixel based methods within a principled RF framework. Our approach enabled the integration of features and contextual priors defined over multiple super-pixels in one optimisation framework that supports efficient MAP estimation using graph cut based move making algorithms. To do this, we have examined the use of auxiliary variables in RFs which have been relatively neglected in computer vision over the past twenty years.

In doing so, we have shown that higher-order RFs are intimately related to pairwise models. This observation has allowed us to characterise some of

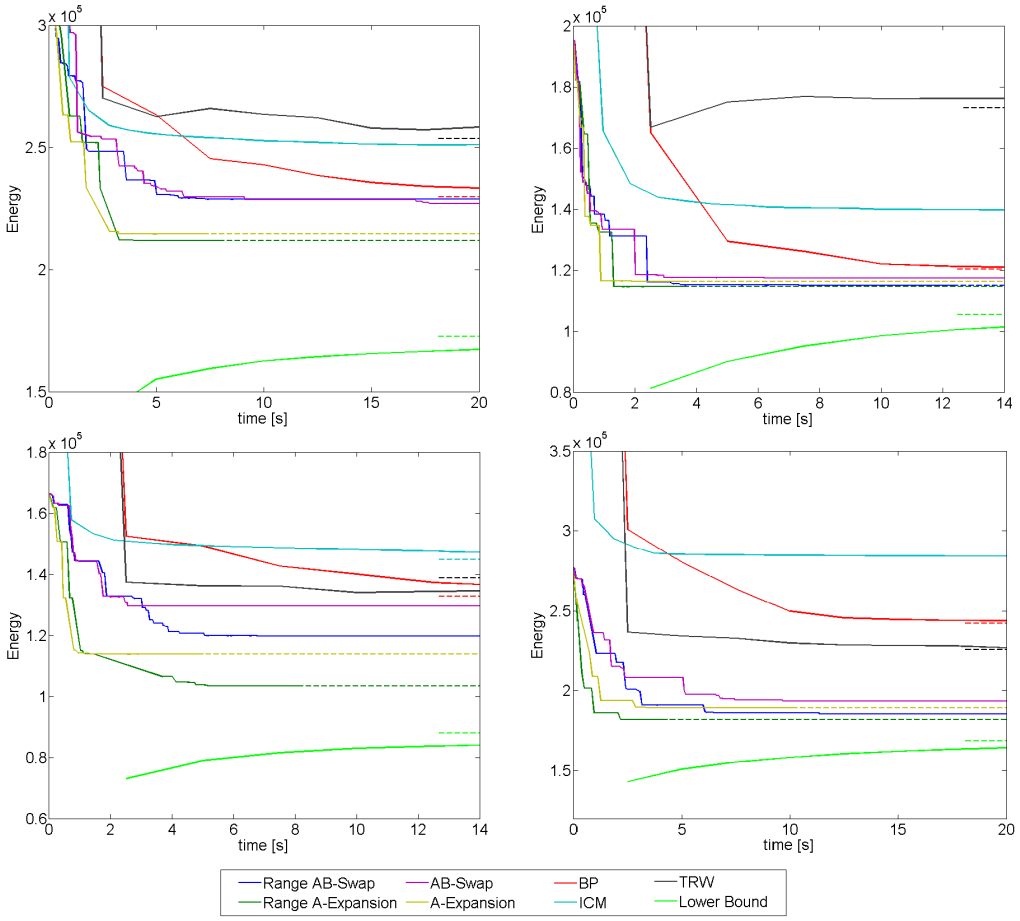



Fig. 11. **Best Viewed in Colour.** *This figure shows additional quantitative results taken from the* MSRC *data set (Shotton et al, 2006). Dashed lines indicate the final converged solution. The slow convergence and poor solutions found by* TRW *and* BP *are to be expected given the large number of cycles present in the graph. Of the remaining move making schemes, the relatively weak performance of $\alpha\beta$-swap and* ICM *is in line with the restricted space of moves available to them. While the three methods derived in this paper significantly outperform all other approaches, range $\alpha$ expansion reliably dominates.*

the higher-order potentials which can be solved using modified expansion and swap algorithms, and derive a bound of 4 for such approaches. We also gave a new transformationally optimal family of algorithms for performing efficient inference in higher-order RF that inherits such bounds.

We demonstrated the usefulness of our algorithms on the problem of object class segmentation where they have been shown to outperform state of the art approaches over challenging data sets both in speed and accuracy. We believe that similar improvements can be achieved for many other higher-order labelling problems both in computer vision and machine learning in general.

The flexibility and generality of our framework allowed us to propose and use novel pixel and segment based potential functions and achieve state-of-the-art results on some of the most challenging data sets for object class segmentation. We believe that use of the hierarchical random fields will yield similar improvements for other labelling problems. The source code is publicly available[4].

## REFERENCES

Alahari K, Russell C, Torr PHS (2010) Efficient piecewise learning for conditional random fields. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition

Batra D, Sukthankar R, Tsuhan C (2008) Learning class-specific affinities for image labelling. In: Conference on Computer Vision and Pattern Recognition

Bengio Y, Lamblin P, Popovici P, Larochelle H (2007) Greedy layer-wise training of deep networks. Advances in Neural Information Processing Systems 19

Besag J (1986) On the statisical analysis of dirty pictures. Journal of the Royal Statistical Society B-48:259–302

Blake A, Rother C, Brown M, Perez P, Torr P (2004) Interactive image segmentation using an adaptive

4. http://cms.brookes.ac.uk/staff/PhilipTorr/ale.htm

GMMRF model. In: European Conference on Computer Vision

Boix X, Cardinal G, van de Weijer J, Bagdanov AD, Serrat J, Gonzalez J (2011) Harmony potentials: Fusing global and local scale for semantic image segmentation. International Journal on Computer Vision URL http://cat.uab.cat/Publications/2011/BGVBSG11

Bosch A, Zisserman A, Munoz X (2007) Representing shape with a spatial pyramid kernel. In: International Conference on Image and Video Retrieval

Boykov Y, Jolly M (2001) Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images. In: International Conference on Computer Vision

Boykov Y, Veksler O, Zabih R (2001) Fast approximate energy minimization via graph cuts. Transactions on Pattern Analysis and Machine Intelligence 23:1222–1239

Comaniciu D, Meer P (2002) Mean shift: A robust approach toward feature space analysis. Transactions on Pattern Analysis and Machine Intelligence

Dalal N, Triggs B (2005) Histograms of oriented gradients for human detection. In: Conference on Computer Vision and Pattern Recognition

Delong A, Osokin A, Isack H, Boykov Y (2010) Fast approximate energy minimization with label costs. Conference on Computer Vision and Pattern Recognition

Everingham M, Van Gool L, Williams CKI, Winn J, Zisserman A (2008) The PASCAL Visual Object Classes Challenge (VOC2008) Results. http://www.pascal-network.org/challenges/VOC/voc2008/workshop/index.html

Felzenszwalb PF, Huttenlocher DP (2004) Efficient graph-based image segmentation. International Journal of Computer Vision 59:167–181

Finley T, Joachims T (2008) Training structural SVMs when exact inference is intractable. In: International Conference on Machine Learning

Galleguillos C, Rabinovich A, Belongie S (2008) Object categorization using co-occurrence, location and appearance. In: Conference on Computer Vision and Pattern Recognition

Gould S, Fulton R, Koller D (2009a) Decomposing a scene into geometric and semantically consistent regions. In: International Conference on Computer Vision

Gould S, Gao T, Koller D (2009b) Region-based segmentation and object detection. In: NIPS

He X, Zemel R, Ray D (2006) Learning and incorporating top-down cues in image segmentation. In: European Conference on Computer Vision

Hinton GE, Osindero S (2006) A fast learning algorithm for deep belief nets. Neural Computation 18:1527–1554

Hoiem D, Efros A, Hebert M (2005) Geometric context from a single image. In: International Conference on Computer Vision

Hoiem D, Efros AA, Hebert M (2006) Putting objects in perspective. In: Conference on Computer Vision and Pattern Recognition

Ishikawa H (2003) Exact optimization for Markov random fields with convex priors. Transactions on Pattern Analysis and Machine Intelligence 25:1333–1336

Kohli P, Kumar M, Torr PHS (2007) $P^3$ and beyond: Solving energies with higher order cliques. In: Conference on Computer Vision and Pattern Recognition

Kohli P, Ladicky L, Torr PHS (2008a) Graph cuts for minimizing robust higher order potentials. Tech. rep., Technical report, Oxford Brookes University, UK

Kohli P, Ladicky L, Torr PHS (2008b) Robust higher order potentials for enforcing label consistency. In: Conference on Computer Vision and Pattern Recognition

Kolmogorov V (2006) Convergent tree-reweighted message passing for energy minimization. Transactions on Pattern Analysis and Machine Intelligence 28:1568–1583

Kolmogorov V, Rother C (2006) Comparison of energy minimization algorithms for highly connected graphs. In: European Conference on Computer Vision

Kolmogorov V, Zabih R (2004) What energy functions can be minimized via graph cuts?. Transactions on Pattern Analysis and Machine Intelligence 26:65–81

Komodakis N, Paragios N (2008) Beyond loose lp-relaxations: Optimizing mrfs by repairing cycles. In: European Conference on Computer Vision

Komodakis N, Paragios N (2009) Beyond pairwise energies: Efficient optimization for higher-order mrfs. In: Conference on Computer Vision and Pattern Recognition, pp 2985–2992

Kumar M, Torr PHS (2008) Efficiently solving convex relaxations for map estimation. In: International Conference on Machine Learning

Kumar MP, Koller D (2009) MAP estimation of semimetric MRFs via hierarchical graph cuts. In: Proceedings of the Conference on Uncertainity in Artificial Intelligence

Kumar MP, Koller D (2010) Efficiently selecting regions for scene understanding. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition

Kumar MP, Veksler O, Torr PHS (2011) Improved moves for truncated convex models. Journal of Machine Learning Research 12:31–67, URL http://dblp.uni-trier.de/db/journals/jmlr/jmlr12.html#KumarVT11

Kumar S, Hebert M (2005) A hierarchical field framework for unified context-based classification. In: International Conference on Computer Vision

Ladicky L, Russell C, Kohli P, Torr PHS (2009) Associative hierarchical CRFs for object class image segmentation. In: International Conference on Computer Vision

Ladicky L, Russell C, Kohli P, Torr PHS (2010) Graph cut based inference with co-occurrence statistics. European Conference on Computer Vision

Lafferty J, McCallum A, Pereira F (2001) Conditional random fields: Probabilistic models for segmenting and labelling sequence data. In: International Conference on Machine Learning

Lan X, Roth S, Huttenlocher D, Black M (2006) Efficient belief propagation with learned higher-order markov random fields. In: European Conference on Computer Vision, pp 269–282

Larlus D, Jurie F (2008) Combining appearance models and Markov random fields for category level object segmentation. In: Conference on Computer Vision and Pattern Recognition

Lempitsky V, Vedaldi A, Zisserman A (2011) A pylon model for semantic segmentation. In: Proceedings of Advances in Neural Information Processing Systems (NIPS)

Lim JJ, Arbeláez P, Gu C, Malik J (2009) Context by region ancestry. In: International Conference on Computer Vision

Lowe DG (2004) Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision 60:91–110

Malik J, Belongie S, Leung T, Shi J (2001) Contour and texture analysis for image segmentation. International Journal of Computer Vision 43:7–27

Munoz D, Bagnell JA, Hebert M (2010) Stacked hierarchical labeling. In: European Conference on Computer Vision (ECCV)

Nowozin S, Gehler PV, Lampert CH (2010) On parameter learning in crf-based approaches to object class image segmentation. In: European Conference on Computer Vision

Ojala T, Pietikainen M, Harwood D (1994) Performance evaluation of texture measures with classification based on Kullback discrimination of distributions. In: Conference on Computer Vision and Image Processing

Pantofaru C, Schmid C, Hebert M (2008) Object recognition by integrating multiple image segmentations. In: European Conference on Computer Vision

Potetz B, Lee TS (2008) Efficient belief propegation for higher order cliques using linear constraint nodes. Computer Vision and Image Understanding

Rabinovich A, Vedaldi A, Galleguillos C, Wiewiora E, Belongie S (2007) Objects in context. In: International Conference on Computer Vision

Ren X, Malik J (2003) Learning a classification model for segmentation. In: International Conference on Computer Vision

Reynolds J, Murphy K (2007) Figure-ground segmentation using a hierarchical conditional random field. In: CRV '07: Proceedings of the Fourth Canadian Conference on Computer and Robot Vision, IEEE Computer Society, Washington, DC, USA, pp 175–182

Roth S, Black M (2005) Fields of experts: A framework for learning image priors. In: Conference on Computer Vision and Pattern Recognition, pp 860–867

Rother C, Kolmogorov V, Blake A (2004) Grabcut: interactive foreground extraction using iterated graph cuts. In: SIGGRAPH

Rother C, Kohli P, Feng W, Jia J (2009) Minimizing sparse higher order energy functions of discrete variables. In: Conference on Computer Vision and Pattern Recognition, pp 1382–1389

Russell B, Freeman W, Efros A, Sivic J, Zisserman A (2006) Using multiple segmentations to discover objects and their extent in image collections. In: Conference on Computer Vision and Pattern Recognition

van de Sande KEA, Gevers T, Snoek CGM (2008) Evaluation of color descriptors for object and scene recognition. In: Conference on Computer Vision and Pattern Recognition

Shi J, Malik J (2000) Normalized cuts and image segmentation. Transactions on Pattern Analysis and Machine Intelligence

Shotton J, Winn J, Rother C, Criminisi A (2006) *TextonBoost*: Joint appearance, shape and context modeling for multi-class object recognition and segmentation. In: European Conference on Computer Vision

Shotton J, Johnson M, Cipolla R (2008) Semantic texton forests for image categorization and segmentation. In: Conference on Computer Vision and Pattern Recognition

Socher R, yu Lin CC, Ng AY, Manning CD (2011) Parsing natural scenes and natural language with recursive neural networks

Sontag D, Meltzer T, Globerson A, Jaakkola T, Weiss Y (2008) Tightening lp relaxations for map using message passing. In: Uncertainty in Artificial Intelligence

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: European Conference on Computer Vision

Szummer M, Kohli P, Hoiem D (2008) Learning crfs using graph cuts. In: Proceedings of the 10th European Conference on Computer Vision: Part II, Springer-Verlag, Berlin, Heidelberg, ECCV '08, pp 582–595, DOI 10.1007/978-3-540-88688-4_43, URL http://dx.doi.org/10.1007/978-3-540-88688-4_43

Tao H, Sawhney H, Kumar R (2001) A global matching framework for stereo computation. In: International Conference on Computer Vision

Tarlow D, Zemel R, Frey B (2008) Flexible priors for exemplar-based clustering. In: Conference on Uncertainty in Artificial Intelligence

Tarlow D, Givoni I, Zemel R (2010) Hop-map: Efficient message passing with high order potentials. In: Artificial Intelligence and Statistics

Taskar B, Chatalbashev V, Koller D (2004a) Learning associative markov networks. In: Proc. ICML, ACM Press, p 102

Taskar B, Chatalbashev V, Koller D (2004b) Learning associative Markov networks. In: International Conference on Machine Learning

Tighe J, Lazebnik S (2010) Superparsing: Scalable nonparametric image parsing with superpixels. In: European Conference on Computer Vision

Torralba A, Murphy K, Freeman W (2004) Sharing features: efficient boosting procedures for multiclass object detection. In: Conference on Computer Vision and Pattern Recognition

Tsochantaridis I, Joachims T, Hofmann T, Altun Y (2005) Large margin methods for structured and interdependent output variables. The Journal of Machine Learning Research 6:1453–1484

Tu Z, Chen X, Yuille AL, Zhu SC (2003) Image parsing: Unifying segmentation, detection, and recognition. International Conference on Computer Vision

Veksler O (2007) Graph cut based optimization for MRFs with truncated convex priors. In: Conference on Computer Vision and Pattern Recognition

Vicente S, Kolmogorov V, Rother C (2009) Joint optimization of segmentation and appearance models. In: International Conference on Computer Vision, IEEE

Wainwright M, Jaakkola T, Willsky A (2005) Map estimation via agreement on trees: message-passing and linear programming. Transactions on Information Theory pp 3697–3717

Wang J, Bhat P, Colburn A, Agrawala M, Cohen M (2005) Interactive video cutout. ACM Transactions on Graphics 24:585–594

Weiss Y, Freeman W (2001) On the optimality of solutions of the max-product belief-propagation algorithm in arbitrary graphs. Transactions on Information Theory 47:736–744

Werner T (2009) High-arity interactions, polyhedral relaxations, and cutting plane algorithm for soft constraint optimisation (map-mrf). In: Conference on Computer Vision and Pattern Recognition

Woodford O, Torr PHS, Reid I, Fitzgibbon A (2008) Global stereo reconstruction under second order smoothness priors. In: Conference on Computer Vision and Pattern Recognition

Yang L, Meer P, Foran DJ (2007) Multiple class segmentation using a unified framework over mean-shift patches. In: Conference on Computer Vision and Pattern Recognition

Zhu L, Yuille AL (2005) A hierarchical compositional system for rapid object detection. In: NIPS

Zhu L, Chen Y, Lin Y, Lin C, Yuille AL (2008) Recursive segmentation and recognition templates for 2D parsing. In: Advances in Neural Information Processing Systems