

# Context first

Daniel Heesch, Robby Tan and Maria Petrou

Imperial College London, SW7 2AZ, UK

**Abstract.** We propose a probabilistic model that captures contextual information in the form of typical spatial relationships between regions of an image. We represent a region’s local context as a combination of the identity of neighbouring regions as well as the geometry of the neighbourhood. We subsequently cluster all the neighbourhood configurations with the same label at the focal region to obtain, for each label, a set of configuration *prototypes*. We propose an iterative procedure based on belief propagation to infer the labels of regions of a new image given only the observed spatial relationships between the regions and the hitherto learnt prototypes. We validate our approach on a dataset of hand segmented and labelled images of buildings. Performance compares favourably with that of a boosted, non-contextual classifier.

## 1 Introduction

Object recognition in general scenes remains a formidable task for artificial systems. Part of the difficulty stems from the way humans categorise things: it is first and foremost shared functional or causal characteristics that define most classes of interest, not similarity in appearance [19]. The problem is compounded by the observation that even very similar objects may look very different under different viewing angles and partial occlusion. Amongst the most successful approaches is that of modelling objects in terms of large sets of discriminant keypoints in conjunction with representations that are invariant with respect to several types of transformation (e.g. [14], [15], [16], [22], [8], [25] and [20]). These appearance-based models have in common that the number of classes to be distinguished is very small and that images contain objects only of one class. It is not clear how to scale to the several thousands of categories humans discriminate without effort.

Promising alternatives are hierarchical models in which features are allowed to be shared between different classes thus reducing the computational burden over flat models, e.g. [9]. Another route is to employ information about other objects of the same scene or information about the type of scene. Such contextual information may reduce the space of plausible object hypotheses and suggest a smaller set of dedicated non-contextual classifiers. It is known that the gist of the scene or the relationships between objects can be captured by the low-frequency content of an image [18]. It has also been shown in neuro-physiological studies that low-frequency information is processed relatively early during visual recognition [2]. Combining these two observations suggests that context may play a pivotal role as an early facilitator during visual recognition.

While other authors have explored the use of scene information for object recognition [23] or region information for scene classification [4], [5], [10], our work investigates the question how and to what extent local geometric and topological relationships between objects can be exploited for object classification. Our approach is motivated by the discovery in [1] of cortical ‘context networks that have been implicated in the storage of typical configurations of objects. Spatial relationships can arguably be extracted more easily than specific details of individual regions. Their extraction is also virtually insensitive to photometric variation. It may therefore not be surprising that these spatial networks also exhibit early activation during visual recognition tasks [1].

In contrast with other contextual models, we do not believe that contextual information should only be used to resolve tension between conflicting non-contextual evidence, e.g. [12] and more recently [21] in the context of probabilistic relaxation. Rather, we believe that context on its own can get us a long way and indeed may be the crucial ingredient to make object recognition scalable.

This paper makes three contributions: (i) we propose a fuzzy representation of the local neighbourhood of a region and a method to obtain *typical* neighbourhood configurations or *prototypes* (ii) we propose a way to use these prototypes in the formulation of a random field over image regions; (iii) we provide an optimisation technique based on belief propagation to relax the random field.

The paper is structured as follows. Section 2 describes related work. In Section 3 we formulate the graphical model, formalise the set of spatial relationships used and describe how representative configurations, or prototypes, are obtained from a training set of annotated images. Section 4 explains how inference is performed. Section 5 describes our experiments. Section 6 ends the paper with a discussion.

## 2 Related work

Several contextual models have been formulated that are concerned with dependencies between objects (as opposed to hierarchical dependencies). Amongst the probabilistic models, Markov random fields are the most popular, e.g. [17], [6], [13], [11], [20]. The authors in [11] and [20] define a conditional random field over individual pixels. In [20], contextual information is incorporated by using the joint boosting algorithm [24] for learning potential functions. Neither work explicitly considers spatial relationships, although [11] includes the absolute position of a site in the potential function.

In [6], each image is assumed to be associated with a bag of words and the precise term-region associations have to be learnt from training data. The Markov random field is specified through single and pair-wise clique potential functions which are learnt on the assumption that they are symmetric. The model therefore does not capture asymmetries in the dependency relationship. The model also does not take into account spatial relationships and thus is indifferent to whether, for example, a blue patch is above (sky) or below (sea) another.

84 In [17], a graphical model is defined over image regions by specifying the 84  
 85 clique functions for all types of single and pair-wise cliques. The potential functions 85  
 86 are weighted sums of basis functions with the parameters being set manually. 86  
 87 Our work has the same objectives as those of [6] and [17]. What sets it 87  
 88 apart is that we allow neighbouring regions to influence each other differently 88  
 89 depending on their relative spatial positions and topological relationships. This 89  
 90 added complexity is best handled by specifying the field in terms of local conditional 90  
 91 probability distributions which are obtained empirically from a training 91  
 92 set. 92

### 93 3 Spatial Context Model 93

94 Let  $S = \{1, \dots, N\}$  index a set of regions in an image. We assume that each 94  
 95 region is associated with a random variable  $x_i$  which takes its value from a 95  
 96 discrete set of class labels. The neighbourhood configuration of the  $i$ th region, 96  
 97  $\mathcal{N}_i$ , comprises the labels and spatial relationships of regions that are within 97  
 98 some radius  $r$  of the focal region. We define the probability with which label  $l$  98  
 99 is assigned to region  $i$  as 99

$$P(x_i = l | \mathcal{N}_i) \equiv \frac{1}{Z} \exp(-\psi(\mathcal{N}_i, R_l)), \quad (1)$$

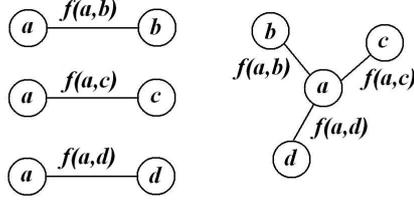
100 where  $Z$  is a normalising constant, and  $\psi(\mathcal{N}_i, R_l)$  is a function that measures the 100  
 101 distance between neighbourhood configuration  $\mathcal{N}_i$  and the set of prototypes with 101  
 102 label  $l$  at their focal region. The next section will describe the set of relations 102  
 103 we use. The subsequent section explains how  $\psi$  is defined and how prototypes 103  
 104  $R_l$  are obtained. 104

#### 105 3.1 Pairwise Relations 105

106 Spatial relationships are often modelled such that only one relationship holds 106  
 107 between any two objects, e.g. [21] and [5] in the context of scene understanding. 107  
 108 The representational convenience of crisp relations comes at the cost of increased 108  
 109 sensitivity to errors with respect to the spatial localisation and geometry of the 109  
 110 input data. We believe that much can be gained by modelling relationships 110  
 111 as fuzzy concepts [3]. A fuzzy relation holds to a variable *degree* determined 111  
 112 by a membership function associated with that relation. We here consider five 112  
 113 relations between region pairs. These are their relative vertical orientation, their 113  
 114 relative horizontal orientation, their containment relation, and the ratio of their 114  
 115 widths and heights, respectively. These are defined as follows. 115

116 **Vertical and Horizontal Relationships.** Let  $p_{cnt_i}$  and  $p_{n_i}$  be points from 116  
 117 each of the two regions. We measure the angle  $\phi_i$  between vector  $(p_{n_i} - p_{cnt_i})$  117  
 118 and the unit vector  $(-1, 1)^T$ . The degree of aboveness (or belowness) of  $p_n$  with 118  
 119 respect to  $p_{cnt}$  is then computed as 119

$$f_{v_i}(p_{n_i}, p_{cnt_i}) = \sin \phi_i \quad (2)$$



**Fig. 1.** The pictorial description of pairwise connections and a configuration, where  $a, b, c, d$  may correspond to wall, sky, roof and door, respectively.  $f(a, b)$  in the diagram is a vector that consist of all components of the relationships between regions  $a$  and  $b$ .

120 where  $f_{v_i}$  represents the vertical relationship of a point pair. Similarly, we rep- 120  
 121 resent the horizontal relationship between point pairs as 121

$$f_{h_i}(p_{n_i}, p_{cnt_i}) = \cos \phi_i. \quad (3)$$

122 To represent the vertical and horizontal relationship between two regions, we 122  
 123 compute the average over point-wise membership values:  $f_v = \frac{1}{N} \sum_i^N f_{v_i}$  and 123  
 124  $f_h = \frac{1}{N} \sum_i^N f_{h_i}$ . To be computationally efficient, we generate  $(p_{n_i}, p_{cnt_i})$  ran- 124  
 125 domly within the respective regions. 125

126 **Containment Relationships** To measure whether region  $r_n$  includes region 126  
 127  $r_{cnt}$ , we are guided by the following decision rule: 127

$$f_{ct}(r_n, r_c) = \begin{cases} -1 & \text{if } (r_n) \text{ contains } (r_{cnt}) \\ +1 & \text{if } (r_n) \text{ is contained in } (r_{cnt}) \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

128 **Width and Height Relationships** We define these as the ratios between 128  
 129 the widths and heights, respectively, of region  $r_n$  and those of region  $r_{cnt}$ . In 129  
 130 our particular application domain, these relationships are useful to distinguish 130  
 131 between, for example, roofs and chimneys, which are indistinguishable under all 131  
 132 other relations. Formally, the width ratio is 132

$$f_{wr}(r_n, r_c) = \begin{cases} 1 - w_{cnt}/w_n & \text{if } w_n/w_{cnt} \geq 1 \\ w_n/w_{cnt} - 1 & \text{otherwise} \end{cases} \quad (5)$$

133 where  $w$  represents the width of a region with respect to its principal axis. The 133  
 134 height ratio is defined analogously. 134

135 The spatial relationship between two regions can then be modeled as a vector 135  
 136 with values in  $[-1, 1]$  each component of which is the membership value for the 136  
 137 corresponding relation. 137

### 138 3.2 Configurations and Prototypes 138

139 A neighbourhood configuration consists of the labels of the neighbours and their 139  
 140 spatial relationships with respect to the focal region. Formally, it is an ordered 140



**Fig. 2.** Left: Value of the energy function at initialisation (top) and after convergence (bottom). Right: distances between clusterings for random clusters (top) and those obtained from the real data (bottom).

141 set of relationship vectors with each vector being associated with a particular 141  
 142 label. See Figure 1 for an illustration. 142

143 **Prototypes** The purpose of the next step is to identify for each region label 143  
 144 a small set of typical neighbourhood configurations, or prototypes. This is accom- 144  
 145 plished by clustering all those configurations that have the same label at the focal 145  
 146 region. Clustering is based on the pair-wise distances between the configurations' 146  
 147 respective matrix representations as described below. 147

148 Let  $P$  and  $Q$  denote the relation matrices of two configurations  $A$  and  $B$  148  
 149 of size  $M$  and  $N$ , respectively. Let the labels of the regions be represented by 149  
 150 vectors  $p$  and  $q$ , respectively. For each region of configuration  $A$ , we determine 150  
 151 its distance from all those regions of configuration  $B$  that bear the same label. 151  
 152 This distance is computed by applying the  $l_1$  metric to the respective row vectors 152  
 153 of  $P$  and  $Q$ . We consider the closest region as the best match to the region of 153  
 154 configuration  $A$ , add the distance to our overall cost and exclude the matching 154  
 155 region from all subsequent comparisons. If configuration  $B$  does not have any 155  
 156 region of that label, a fixed cost is applied to penalise label discrepancies. This 156  
 157 is repeated for all other configurations of region  $A$ . The overall cost reflects both 157  
 158 differences in the labels as well as differences in the geometry and topology of 158  
 159 regions carrying the same label. 159

160 We employ the  $k$ -medoid algorithm to cluster configurations. Like  $k$ -means, 160  
 161 the algorithm is guaranteed to converge because the sum of the distances be- 161  
 162 tween all points and their respective cluster centroid cannot increase and is also 162  
 163 bounded from below. However, like any gradient descent algorithm, the final so- 163  
 164 lution depends on the initialisation and is thus not guaranteed to be the global 164  
 165 optimum. To assess the stability of the solution, we run the algorithm several 165  
 166 times and compare the energy before and after convergence. As Figure 2 indi- 166  
 167 cates, the final energy remains within narrow bounds and suggests that the final 167  
 168 solutions come close to the global optimum. 168

169 A similar energy upon convergence does not imply, however, that the cluster- 169  
 170 ings are the same or similar. Let  $\mathcal{A} = \{A_1, A_2, \dots, A_m\}$  and  $\mathcal{B} = \{B_1, B_2, \dots, B_n\}$  170  
 171 denote two clusterings. The members of each clustering are themselves sets of 171

172 indices denoting a particular configuration. To assess the quality of the result of 172  
 173 the  $k$ -medoid algorithm, we run it several times on the real data and compute 173  
 174 a distance measure for all pairs of clusterings thus obtained. We then generate 174  
 175 random clusterings consisting of the same number of clusters and the same cluster 175  
 176 size distribution as the real clustering, and measure their pair-wise distances. 176  
 177 We compute the distance between two clusterings  $\mathcal{A}$  and  $\mathcal{B}$  as 177

$$\sum_{i=1}^m d(A_i, \mathcal{B}) = \sum_{i=1}^m \left( |A_i| - \max_j |A_i \cap B_j| \right). \quad (6)$$

178 The distances between the random clusterings and the true clusterings are plot- 178  
 179 ted in Figure 2 on the right. The plot demonstrates that the set of clusterings 179  
 180 obtained by running  $k$ -medoid repeatedly on the same distance matrix are more 180  
 181 similar to each other than a set of random clusterings. We take this as circum- 181  
 182 stantial evidence that  $k$ -medoid does capture intrinsic structure in the space of 182  
 183 configurations. 183

184 Cluster centroids are those configurations for which the sum of the distances 184  
 185 to all other members of the respective cluster is minimal. Prototypes correspond 185  
 186 to the cluster centroids and thus are themselves configurations. 186

## 187 4 Inference 187

188 Given a set of regions in an image, we intend to label them using the prototypes 188  
 189 we have generated. To arrive at correct labellings, we define a cost function that 189  
 190 is based on the distance between the observed configurations and their closest 190  
 191 prototype. Formally, we define the potential function as 191

$$\psi(\mathcal{N}_i, R_l) = \min_{R \in R_l} d(\mathcal{N}_i, R) \quad (7)$$

192 where  $d(\mathcal{N}_i, R)$  is the distance between a configuration and a prototype  $R$  as de- 192  
 193 fined in Section 3.2. We intend to obtain the closest distance of all configurations 193  
 194 from the corresponding prototypes, that is we want to minimise 194

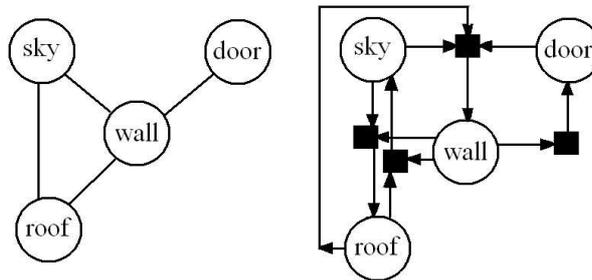
$$E(\mathbf{x}) \equiv E(x_1 = l_1, \dots, x_N = l_N) = \sum_{i \in S} \psi(\mathcal{N}_i, R_l) \quad (8)$$

195 In order to apply the technique of belief propagation, we change the undi- 195  
 196 rected configurations into directed ones and generate a factor graph. Figure 3 196  
 197 shows an example of the transformation. We may subsequently use the following 197  
 198 equations to optimise the cost function. 198

$$b(x_i) = \prod_{c \in N_i} m_{c \rightarrow i}(x_i) \quad (9)$$

$$m_{c \rightarrow i}(x_i) = \sum_{\{x_c\} - x_i} \psi \prod_{i' \in N_{c-i}} n_{i' \rightarrow c}(x_{i'}) \quad (10)$$

$$n_{i \rightarrow c}(x_i) = \prod_{c' \in N_{i-c}} m_{c' \rightarrow i}(x_i) \quad (11)$$



**Fig. 3.** The pictorial description of the transformation from an undirected graph into a directed graph.



**Fig. 4.** Examples of manually segmented images

199 where  $b(x_i)$  is the belief that  $x_i$  has a particular label.  $m_{c \rightarrow i}(x_i)$  is the message 199  
 200 from the neighbouring nodes of  $i$  with respect to the label of  $x_i$ . In Figure 3, 200  
 201  $m_{c \rightarrow i}(x_i)$  is represented by the black boxes that connect the regions or nodes. 201  
 202 Since the configurations in the factor graph are directed configurations, we may 202  
 203 simplify the equations thus 203

$$b(x_i) = m_{c \rightarrow i}(x_i) \quad (12)$$

$$m_{c \rightarrow i}(x_i) = \sum_{\{x_c\} - x_i} \psi \prod_{i' \in N_{c \rightarrow i}} n_{i' \rightarrow c}(x_{i'}) \quad (13)$$

$$n_{i \rightarrow c}(x_i) = m_{c' \rightarrow i}(x_i) \quad (14)$$

## 204 5 Evaluation 204

205 The image collection used for training and testing consists of photographs depicting 205  
 206 buildings from different cities of several countries, mostly taken from 206  
 207 street-level. Each image was manually segmented and labeled with one of the 207  
 208 following nine classes: ‘window’, ‘chimney’, ‘roof’, ‘door’, ‘wall’, ‘stairs’, ‘pipe’, 208  
 209 ‘sky’, and ‘vegetation’. Figure 4 shows two examples. The training set contains 209  
 210 197 images with a total of 3,675 regions. The test set comprises 80 images with 210  
 211 a total of 1,372 regions. Images were randomly assigned to one of the two sets. 211

212 We shall note that we do not address the challenge of segmenting regions 212  
 213 automatically. Instead, we assume that for learning and testing regions have been 213  
 214 manually segmented. The recent work of [5] indicates that one may profitably 214  
 215 start with an automated segmentation when the regions to be separated are 215  
 216 visually distinct (e.g. material types like water, sky, sand). In our case, however, 216  
 217 a general-purpose segmentation routine is unlikely to achieve sufficient accuracy 217  
 218 for our contextual model. 218

219 Table 1 shows the confusion matrix for the proposed method. We compare 219  
 220 performance with a non-contextual AdaBoost classifier which is trained to find 220  
 221 the optimal combination of unary attributes pertaining to the shape of regions 221  
 222 including the ratio of principal axes, compactness, variances, and elliptical vari- 222  
 223 ances. The confusion matrix for the labelling based on the adaboost algorithm 223  
 224 is shown in Table 2. The total error is 576 (42%) compared to 473 (34%) for the 224  
 225 contextual labelling. 225

	Win	Chi	Roo	Doo	Wal	Dor	Sta	Pip	Sky	Veg
Window	487	11	9	77	2	11	1	1	23	17
Chimney	10	70	1	0	0	24	0	0	2	0
Roof	18	0	31	0	0	2	1	0	31	8
Door	27	3	0	84	0	9	1	0	1	10
Wall	30	8	2	5	45	3	1	5	10	17
Dormer	7	1	0	1	0	2	0	0	3	1
Stairs	0	0	0	0	0	0	19	0	0	1
Pipes	12	16	0	2	3	0	0	48	1	1
Sky	11	2	1	0	0	0	0	1	79	0
Vegetation	7	2	2	3	1	1	6	5	1	35

**Table 1.** Confusion matrix for contextual classification (rows: true labels; columns: hypothesised labels).

	Win	Chi	Roo	Doo	Wal	Dor	Sta	Pip	Sky	Veg
Window	435	56	9	49	4	78	4	0	1	4
Chimney	71	21	2	7	0	5	0	0	1	0
Roof	5	1	45	0	0	2	0	1	37	0
Door	71	3	0	53	3	5	0	0	0	0
Wall	2	9	3	1	83	10	1	2	14	1
Dormer	7	3	3	0	0	2	0	0	0	0
Stairs	15	1	2	0	0	2	0	0	0	0
Pipes	0	0	0	0	0	4	0	79	0	0
Sky	7	1	6	0	5	1	0	0	72	2
Vegetation	10	0	10	3	19	4	0	3	6	8

**Table 2.** Confusion matrix for boosting on shape descriptors (rows: true labels; columns: hypothesised labels).

## 226 6 Discussion 226

227 In most computer vision applications that take into account contextual informa- 227  
 228 tion, regions are initially assigned labels on the basis of their unitary attributes, 228  
 229 and the label assignment is subsequently refined through the use of contextual 229  
 230 information (e.g. [7]). Inspired by neuro-physiological findings about human vi- 230  
 231 sual processing, we here advocate the view that the order in which information 231  
 232 is processed ought to be reversed. When humans view a scene, they first view it 232  
 233 as a whole before focussing on particular details that merit further interpreta- 233  
 234 tion. Anecdotal evidence supporting this idea is the preference people show for 234  
 235 seats in trains that look forward over seats that look backward in relation to the 235  
 236 travelling direction. When looking backwards, one sees first the detail and then 236  
 237 the context of the object; when looking forward, the global picture is captured 237  
 238 first. 238

239 We proposed a probabilistic model in which a region’s local neighbourhood 239  
 240 is represented in the form of fuzzy relationship matrices. Because of the inability 240  
 241 to define cliques in directed graphical models, the joint label probability over all 241  
 242 regions cannot be expressed as a Gibbs distribution. Instead, we defined local 242  
 243 conditional probabilities of a region’s label that depend on the neighbourhood 243  
 244 through a potential function that takes into consideration differences in the ge- 244  
 245 ometry of and the labels present in the neighbourhood. Central to our approach 245  
 246 is the idea of typical neighbourhood configurations or prototypes which are ob- 246  
 247 tained from a training set through clustering. Every label is associated with a 247  
 248 small number of prototypes and inference aims to find a labelling of all regions 248  
 249 such that the observed configurations are close the labels’ closest prototypes. 249  
 250 Comparison with a non-contextual Adaboost classifier trained on a variety of 250  
 251 shape features support the view that contextual information can provide power- 251  
 252 ful information for labelling structured scenes. 252

253 **Acknowledgements:** This work was supported by the EU project eTRIMS. 253

## 254 References 254

- 255 1. M Bar and E Aminoff. Cortical analysis of visual context. *Neuron*, 38:347–358, 255  
 256 2003. 256
- 257 2. M Bar, K Kassam, A Ghuman, J Boshyan, A Schmidt, A Dale, M Hämäläinen, 257  
 258 K Marinkovic, D Schacter, B Rosen, and E Halgren. Top-down facilitation of visual 258  
 259 recognition. *Proc National Academy of Sciences*, 103(2):449–454, 2006. 259
- 260 3. I Bloch. *Fuzzy Representations of Spatial Relations for Spatial Reasoning*. John 260  
 261 Wiley & Sons, 2007. 261
- 262 4. M Boutell, C Brown, and J Luo. Learning spatial configuration models using 262  
 263 modified Dirichlet priors. In *Workshop on Statistical Relational Learning*, 2004. 263
- 264 5. M Boutell, J Luo, and C Brown. Factor graphs for region-based whole-scene classi- 264  
 265 fication. In *Proc IEEE Conf Computer Vision and Pattern Recognition, Semantic 265  
 266 Learning Workshop*, 2006. 266
- 267 6. P Carbonetto, N de Freitas, and K Barnard. A statistical model for general contex- 267  
 268 tual object recognition. In *Proc European Conf Computer Vision*, pages 350–362, 268  
 269 2004. 269

- 270 7. W J Christmas, J Kittler, and M Petrou. Structural matching in computer vi- 270  
 271 sion using probabilistic relaxation. *IEEE Transactions on Pattern Analysis and* 271  
 272 *Machine Intelligence*, 17(8):749–764, 1995. 272
- 273 8. G Csurka, C Bray, C Dance, and L Fan. Visual categorization with bags of key- 273  
 274 points. In *Proc European Conf Computer Vision*, 2004. 274
- 275 9. S Fidler, M Boben, and A Leonardis. Similarity-based cross-layered hierarchical 275  
 276 representation for object categorization. In *Proc Int'l Conf Computer Vision and* 276  
 277 *Pattern Recognition*, 2008. 277
- 278 10. D Gökalp and S Aksoy. Scene classification using bag-of-regions representations. 278  
 279 In *Proc IEEE Conf Computer Vision and Pattern Recognition, Beyond Patches* 279  
 280 *Workshop*, 2007. 280
- 281 11. X He, R Zemel, and D Ray. Learning and incorporating top-down cues in image 281  
 282 segmentation. In *Proc European Conf Computer Vision*, 2006. 282
- 283 12. J Kittler and S Hancock. Combining evidence in probabilistic relaxation. *Journal* 283  
 284 *of Pattern Recognition and Artificial Intelligence*, 3:29–51, 1989. 284
- 285 13. S Kumar and H Hebert. Discriminative random fields: a discriminative framework 285  
 286 for contextual interaction in classification. In *Proc Int'l Conf Computer Vision*, 286  
 287 2003. 287
- 288 14. D Lowe. Object recognition from local scale-invariant features. In *Proc Int'l Conf* 288  
 289 *Computer Vision*, pages 1150–1157, 1999. 289
- 290 15. D Lowe. Distinctive image features from scale-invariant keypoints. *Int'l Journal* 290  
 291 *of Computer Vision*, 60:91–110, 2004. 291
- 292 16. J Matas, O Chum, M Urban, and T Pajdla. Robust wide baseline stereo from 292  
 293 maximally stable extremal regions. In *British Machine Vision Conference*, pages 293  
 294 384–393, 2002. 294
- 295 17. J Modestino and J Zhang. A Markov random field model-based approach to image 295  
 296 interpretation. *IEEE Trans Pattern Analysis and Machine Intelligence*, 14(6):606– 296  
 297 615, 1992. 297
- 298 18. A Oliva and A Torralba. Modelling the shape of the scene: a holistic representation 298  
 299 of the spatial envelope. *Int'l Journal Computer Vision*, 42(3):145–175, 2001. 299
- 300 19. M Petrou. Learning in Computer Vision: some thoughts. *Progress in Pattern Recog-* 300  
 301 *niton, Image Analysis and Applications. The 12th Iberoamerican Congress on Pat-* 301  
 302 *tern Recognition, CIARP 2007*, Vina del Mar-Valparaiso, November, L Rueda, D 302  
 303 Mery and J Kittler (eds), LNCS 4756, Springer, pages 1–12, 2007. 303
- 304 20. J Shotton, J Winn, C Rother, and A Criminisi. Textonboost: Joint appearance, 304  
 305 shape and context modeling for multi-class object recognition and segmentation. 305  
 306 In *Proc European Conf Computer Vision*, 2006. 306
- 307 21. A Singhal, J Luo, and W Zhu. Probabilistic spatial context models for scene content 307  
 308 understanding. In *Proc IEEE Conf on Computer Vision and Pattern Recognition*, 308  
 309 2003. 309
- 310 22. J Sivic and A Zisserman. Video google: a text retrieval approach to object matching 310  
 311 in videos. In *Proc Int'l Conf Computer Vision*, pages 1–8, 2003. 311
- 312 23. A Torralba. Contextual priming for object detection. *Int'l Journal of Computer* 312  
 313 *Vision*, 52(2):169–191, 2003. 313
- 314 24. A Torralba, K Murphy, and W Freeman. Sharing features: efficient boosting pro- 314  
 315 cedures for multiclass object detection. In *Proc IEEE Conf Computer Vision and* 315  
 316 *Pattern Recognition*, pages 762–769, 2004. 316
- 317 25. J Winn, A Criminisi, and T Minka. Object categorization by learned universal 317  
 318 visual dictionary. In *Proc Int'l Conf Computer Vision*, pages 1800–1807, 2005. 318