

How good are detection proposals, really?

Jan Hosang

<http://mpi-inf.mpg.de/~jhosang>

Rodrigo Benenson

<http://mpi-inf.mpg.de/~benenson>

Bernt Schiele

<http://mpi-inf.mpg.de/~schiele>

MPI Informatics

Saarbrücken, Germany

Abstract

Current top performing Pascal VOC object detectors employ detection proposals to guide the search for objects thereby avoiding exhaustive sliding window search across images. Despite the popularity of detection proposals, it is unclear which trade-offs are made when using them during object detection. We provide an in depth analysis of ten object proposal methods along with four baselines regarding ground truth annotation recall (on Pascal VOC 2007 and ImageNet 2013), repeatability, and impact on DPM detector performance. Our findings show common weaknesses of existing methods, and provide insights to choose the most adequate method for different settings.

1 Introduction

Object detection is traditionally formulated as a classification problem in the well known “sliding window” paradigm where the classifier is evaluated over an exhaustive list of positions, scales, and aspect ratios. Steadily increasing the sophistication of the core classifier has led to increased detector performance [1, 2, 3].

A typical sliding window detector requires $\sim 10^6$ classifier evaluations per image. One approach to overcome the tension between computational tractability and high detection quality, is the notion of “detection proposals” (sometimes called “objectness” or “selective search”). Under the assumption that all objects of interest share common visual properties that distinguish them from the background, one can train a method that, given an input image, outputs a set of detection window proposals that are likely to contain the objects. If high recall can be reached with $\sim 10^4$ or less windows, significant speed-ups can be achieved, enabling the use of even more sophisticated classifiers.

Besides the potential to improve speed, the use of detection proposals changes the data distribution that the classifier handles, thus also potentially improving detection quality (reduce false positives). It is interesting to note that the current two top performing detection methods for Pascal VOC [4] use detection proposals [5, 6].

Contributions In contrast to previous work that typically performs a partial evaluation when introducing a novel method, this paper aims to revisit existing work on detection




Figure 1: How to evaluate the quality of such detection proposals?

proposals and compare all publicly available methods in a common framework to better understand their benefits and limitations. We attempt to provide the necessary insights to understand when to choose which method.

Our contributions are the following. (1) We review existing detection proposal methods (§2) and compare them in a unified evaluation. Our evaluation aims to provide an unbiased view of existing methods. (2) We introduce the notion of repeatability for detection proposals, discuss why it matters when considering detection proposals for object detection, and compare repeatability of existing methods (§3). (3) We evaluate overlap with annotated ground truth objects on the Pascal VOC 2007 test set, and for the first time, over the larger and more diverse ImageNet 2013 validation set (§4). This latter experiment aims at detecting possible biases towards the Pascal VOC objects categories. (4) Finally we evaluate the influence on detector performance using different proposal methods (§5). (5) All bounding boxes from our experiments and the evaluation scripts will be released with the paper. The results presented summarise more than 500 experiments over different data sets, totalling to more than 2.5 months of CPU computation.

2 Detection proposal methods

Interestingly, the spirit of detection proposals is similar to the idea of interest point detection [23, 28]. Interest points were proposed at a time when computing feature descriptors densely was computationally too expensive and some selection of interest points was important. Feature descriptors computed around such interest points were used successfully for classification, retrieval, and detection. Today however, with the increase of computing power, it is standard practice to use dense feature extraction instead [27]. Thus we may ask: do object proposals help detection quality, or are they just a transition technology until we have sufficient computing power?

Detection proposal methods are based on low-level image features to generate candidate windows. One can reinterpret this process as a discriminative one; given the low-level features the method quickly decides whether a window should be considered for detection or not. In this sense detection proposal methods are related to cascade methods [8, 18, 51], which use a fast (but inaccurate) classifier to discard the vast majority of unpromising proposals. Although traditionally used for class specific detection, cascade methods also apply to sets of categories [26], and in principle can be applied to a very large set of categories (see Bing method below).

Since their introduction [11, 16] multiple works have explored the idea of generating detection proposals. We briefly review all methods we are aware of in chronological order. **gPbUCM** [16] is a leading method for grouping pixels into objects. Given an input image it hierarchically groups segments and uses these segments (or bounding boxes around them) as object detection candidates [16, 17]. This paper evaluates the bounding box proposals since not all methods generate a segmentation of the proposals.

Objectness [11, 1] is a term that refers to a measure of how likely a detection window contains an object (of any category). [1] estimates this score based on a combination of multiple cues such as saliency, colour contrast, edge density, location and size statistics, and how much such windows overlap with superpixel segments.

CPMC [1, 5]: Most methods are built upon some form of hierarchical segmentation. CPMC avoids this by generating a set of overlapping segments. Each proposal segment is the solution of a binary segmentation problem, initialised with diverse seeds. Up to 10^4 segments are

generated per image, which are subsequently ranked by objectness using a trained regressor. **Endres2010** [9, 10]: The approach from [10] mixes a large set of cues. It uses a hierarchical segmentation, a learned regressor to estimate boundaries between surfaces with different orientations, graph cuts with different seeds, and parameters to generate diverse segments (similar to CPMC). It also learns a ranking function for the proposed segments.

SelectiveSearch [29, 30] is a method where no parameters are learned. The authors carefully engineer features and score functions that greedily merge low-level superpixels. The authors obtain state of the art object detections on Pascal VOC and ILSVRC2011.

Rahtu2011 [24] revisits [9] and obtains a significant improvement by proposing new objectness cues and a different method to learn how to combine them.

RandomizedPrim's [22] is similar to *SelectiveSearch* in terms of features to merge low-level superpixels. The key difference is that the weights of the merging function are learned and the merging process is randomised.

Bing [8] is one of the only two methods that are not based on segmentation. A simple linear classifier over edge features is trained and applied in a sliding window manner. Using adequate approximations a very fast class agnostic detector is obtained (1 ms/image on CPU).

MCG [3] is one of the most recent methods combining gPbUCM and CPMC. The authors propose an improved multi-scale hierarchical segmentation (similar to gPbUCM), a new strategy to generate proposals by merging up to 4 segments, and (similar to CPMC) a new ranking procedure to select the final detection proposals.

Rantalankila2014 [25], also recently proposed, combines *SelectiveSearch* and CPMC. Starting from low-level superpixels the authors propose a merging strategy, similar to *SelectiveSearch* but using different features. These merged segments are then used as seeds for a CPMC-like process to generate larger segments.

Rigor [20] is a variant of CPMC obtaining higher quality by using different low-level superpixels and different features for merging segments. They also obtain higher speed by minimizing redundant computations.

EdgeBoxes [33] is similar in spirit to *Bing*, a scoring function is evaluated in a sliding window fashion. This method uses object boundaries estimates (obtained via structured decision forests) as feature for the scoring. Interestingly, the authors propose tuning parameters to optimize recall at a desired overlap threshold (see section 4).

Recently, Kang et al. [21] proposed a “data-driven objectness” approach. Although it showed promising results for their indoor application scenario, the method seems of limited applicability to other datasets, and thus we do not consider it in our evaluation.

The majority of the methods is based on some low-level segmentation: five use [13], two use a variant of the gPbUCM segmentation method, and *Endres2010* uses its own custom low-level segmentation (also used by *Rigor*). Only three methods work without computing low-level segments (CPMC, *Bing*, *EdgeBoxes*). At the time of writing no code is available for gPbUCM and *Rigor*¹, thus we do not consider these in our evaluation. In total seven out of the ten methods evaluated in this paper use some low-level segmentation, five out of these eight use [13].

It should also be noted that since [10], all related work evaluates the quality of detection proposal methods based on the overlap with ground truth annotations of the Pascal VOC dataset (2007 or 2010) [10]. Although a relevant metric, its repeated use opens the doors for over-fitting to this specific dataset. In this paper we also consider other evaluations.

Number of windows The different methods listed above provide different numbers of de-

¹Recently published at CVPR 2014.

tection candidates. Some provide rather few windows ($\sim 10^2$), others provide a large number ($\sim 10^5$). Some methods do provide sorted (scored) windows, others do not. Having more windows increases the chance for high recall, thus for each method in all experiments we do our best effort to generate a similar average number of candidate windows per image. See supplementary material for details on how this is achieved for each method.

2.1 Baselines

Besides the above methods, we also consider a set of baselines that serve as quantitative reference points. All of the above candidate methods and the following baselines are class independent.

Uniform: To generate detection proposals, we uniformly sample the bounding box parameters centre position, square root area, and log aspect ratio. We estimate the range of these parameters on the Pascal VOC 2007 training set after discarding 0.5% of the smallest and largest values, so that our estimated distribution covers 99% of the data.

Gaussian: We also estimate a multivariate Gaussian distribution for the bounding box parameters centre position, square root area, and log aspect ratio. After calculating mean and covariance on the training set we sample proposals from this distribution.

SlidingWindow places windows on a regular grid as common for sliding window object detectors. The contingent of the requested number of proposals is distributed across different windows sizes. For each window size, we place the windows uniformly across the image. The procedure is inspired by the implementation of Bing [6].

Superpixels: As discussed in the next sections, low-level superpixels have an important influence on the behaviour of the detection proposal methods. As five of the compared methods build on [13], we use it as a baseline: each low-level segment is used as a detection proposal. We should expect this method to have low recall for objects (§4), but high repeatability (§3).

3 Proposal repeatability

Before looking into how well the different object proposals overlap with ground truth annotations of objects, we want to answer a more basic question. Training a detector on detection proposals rather than on all sliding windows modifies the distribution of negative windows that the classifier is trained on. If the proposal method does not consistently propose windows on similar image content without objects or with partial objects, the classifier cannot produce useful scores on negative windows on the test set. We call the property of proposals being placed on similar image content the repeatability of a proposal method. Intuitively the detection proposals should be repeatable on slightly different images with the same image content. To evaluate repeatability we project proposals from one image into another slightly modified image. Pascal does not contain suitable images. An alternative is the dataset of [23], but it only consists of 54 images and



Figure 2: Example of the image perturbations considered. Top to bottom, left to right: original, then blur, illumination, JPEG artefact, rotation, and scale perturbations.

only few of the images contain objects. Instead, we opt to apply synthetic transformations to Pascal images, so we know how to project proposals from one image into another.

3.1 Evaluation protocol

Our evaluation protocol is inspired by [24], which evaluates interest point repeatability. For each image in the Pascal VOC 2007 test set, we generate several perturbed versions. We consider changes in scale, blur, small rotation, illumination, and JPEG compression (see figure 2). The details of the transformations and additional examples of the perturbed images are provided in the supplementary material.

For each pair of reference and perturbed images we compute detection proposals with a given algorithm (requesting 1 000 windows per image). The proposals are projected back from the perturbed into the reference image and then matched to the proposals in the reference image. All proposals whose centre lies outside the image after projection are removed before matching (which can only happen for rotation). Matching is done greedily according to the intersection over union (IoU) criterion. Given the matching, we plot the recall for every IoU threshold and define the area under this “recall versus IoU threshold” curve to be the repeatability. Methods that propose windows at similar locations at high IoU are more repeatable, since the area under the curve is larger.


One issue regarding such proposal matching is that large windows are more likely to match than smaller ones. This effect is important to consider since different methods have quite different distributions of window areas (see figure 3a). To reduce the impact of this effect, we bin the original image windows by area (into 10 groups), and evaluate the area under the recall versus IoU curve per size group. The plots in figure 3 show the unweighted average across size bins.

Figure 3a shows that the different proposal methods exhibit a variety of different distributions of proposal window sizes and in general different from the ground truth annotation distribution. This confirms the need for normalisation based on the window size. Figure 3b shows the recall versus intersection over union (IoU) curve for a blur perturbation. This example shows the effect, that large rectangles have higher repeatability, which is why we use the average across size bins to analyse repeatability in the following.


We omit the slowest two methods, CPMC and Endres2010, from this experiment because it involves running the candidate detector over the Pascal test set ~50 times (once for every perturbation).

3.2 Repeatability results


There are some salient aspects of the result curves in figure 3 that need some explanation. First, not all methods have 100% repeatability when there is no perturbation. This is due to random components in the selection of proposals for several methods (BING being the most notable exception). It is possible to control the random aspect by fixing the random seed, but this would not be effective when perturbations are present and the underlying segmentation changes: a drop in performance could be due to pseudo-random samples or to an actual change in behaviour. Another approach would be to remove the random aspect from all methods, but this would denature the method being evaluated and we would actually evaluate a different algorithm. By keeping the random aspect when no transformation is applied to the image, we can compare values along the change in perturbation.




(a) Histogram of proposal window sizes on Pascal VOC 2007 test.




(b) Example of recall for different candidate window sizes.




(c) Scale change.



(d) JPEG artefacts.



(e) Rotation.



(f) Illumination.

Figure 3: Repeatability results under perturbations. See also supplementary material.

A second important aspect is the large drop of repeatability for most methods, even for very subtle image changes. We have created the `Superpixels` baseline as a means to validate our hypothesis regarding this effect. Most methods use low-level superpixels as a building block for the proposal windows. We suspect that the superpixels themselves are unstable under small perturbations. This is supported by the fact that the `Superpixels` baseline also shows a strong drop (since it is a direct reflection of the superpixels themselves). Since proposals on larger perturbations are all matched against the same proposals in the reference (unperturbed) image, after the very first small perturbation most windows are “lost”, the remaining windows that are correctly matched are progressively lost as the perturbation level increases. Inversely we notice that methods that are not based on superpixels are most robust to small image changes (`Bing` and the baselines that ignore image content being the most noticeable ones).

We now briefly discuss the effects under each perturbation, shown in figure 3:

Scale change 3c: All methods except `Bing` show a drastic drop with small scale changes, and minor relative degradation for larger changes. `Bing`’s drop is less extreme for very small scale changes. For larger changes, however, it presents a non monotonic behaviour because it uses a coarse set of box sizes while searching for detection candidates. Such coarse set impedes `Bing` to detect the same object at two slightly different scales. Logically, our `SlidingWindow` baseline suffers from the same effect.

JPEG artefacts 3d: Similar to scale change, even small JPEG artefacts have a large effect and more aggressive JPEG compression factors show monotonic degradation. Despite using gradient information `Bing` is the method most robust to these kind of changes. It should be noted that JPEG images with 100% quality are still lossy. Only when requesting lossless JPEG the superpixels are preserved because the image content does not change.

Rotation 3e: Here all methods are equally affected by the perturbation. The drop of the Uniform and Gaussian baselines indicate the repeatability loss due to the fact that we are matching rotated bounding boxes.

Illumination change 3f shows a similar trend to JPEG artefacts. We notice here that the gradient information used by Bing is more robust to small changes than the superpixels that are used by other methods. Changes in illumination reduce the difference between neighbour pixels, and in turn, affect how they are merged into larger superpixels.

Blurring exhibits a similar trend to JPEG artefacts, although the drop is stronger for a small σ . The corresponding plot can be found in the supplementary material.

Overall it seems that Bing and EdgeBoxes are more repeatable than other methods, possibly because both use machine components (SVM classifier for scoring and decision forests for features computation, respectively). We also conclude that the sensitivity of superpixels to image perturbations is a major cause for degradation in repeatability of several detection proposal methods.

To better understand how much repeatability influences detector performance, one should train detectors on proposals and compare the detection performance of the detector using different proposal methods. This is, however, a more complicated procedure that involves many new parameters (see e.g. [15, 32]) which is why we leave this for future work.

4 Proposal recall

When using detection proposals it is important to have a good coverage of the true objects in the test image, since missed objects will never be recovered. Thus it is common practice to evaluate the quality of proposals based on ground truth annotations.

4.1 Evaluation protocol

The protocol introduced in [14] has served as a guideline for most other evaluations in the literature. Typically the detection proposal method is trained on a subset of the Pascal VOC categories, and tested on the test set, including unseen categories. The metric of interest we use is the fraction of ground truth annotation covered above a certain intersection over union (IoU) threshold. Note that evaluating (class agnostic) detection proposals is quite different from traditional class-specific detection [19], since most of the metrics (class confusion, background confusion, precision, etc.) do not apply.

Previous work includes comparisons, however the train and test sets vary amongst papers, and the metrics shown tend to favour different methods. We provide an extensive unified evaluation, that shows how a different perspective can offer different ranking of the methods (see figure 4b versus 5b).

We use the methods as provided by the authors. Different methods may be trained on different sets, but we think this is still fair since we care about absolute quality. Some methods do not have training at all (SelectiveSearch), yet provide competitive results. To mitigate this effect we evaluate on the full Pascal VOC 2007 test set, including all categories.

The Pascal VOC 2007 test set has only 20 categories, present in a diverse set of ~ 5000 unconstrained images, yet detection proposal methods attempt to predict “any” object. On principle, methods which are particularly good at detecting non-Pascal categories will be penalised compared to the ones that are good at these 20 categories but nothing else. To




Figure 4: Recall versus IoU threshold on the Pascal VOC 2007 test set. Numbers next to label indicate area under the curve and average number of windows per image, respectively.




Figure 5: Recall versus number of proposed windows on the Pascal VOC 2007 test set.

to investigate this dataset bias, we also evaluate methods on the larger ImageNet [14] 2013 validation set, which contains annotations for 200 categories over ~ 20000 images. It should be noted that these 200 categories are *not* fine grained versions of the Pascal ones. It includes additional types of animals (e.g. crustaceans), food items (e.g. hot-dogs), house hold items (e.g. diapers), and other diverse object categories.

4.2 Recall results

Results in figure 4 present a consistent trend across the different metrics. MCG, EdgeBoxes, and SelectiveSearch seem to be the best methods across different number of proposals.

We show EdgeBoxes when tuned for IoU at 0.7, which results in a clear bump in recall. When tuned for IoU = 0.9, for 10^3 proposals, EdgeBoxes is below MCG, and when tuned for IoU = 0.5 it reaches about 93% recall. MCG shows the most consistent results when varying the number of proposals. SelectiveSearch is surprisingly effective despite being a fully hand-crafted method (no machine learning involved). When considering less than 10^3 proposals, MCG, Endres2010, and CPMC provide strong results.




Figure 6: Recall versus IoU threshold, for 1000 proposals per image on the ImageNet 2013 validation set. See figure 4 and supplementary material.

Overall, we can group the methods into two groups: well localized methods that gradually lose recall as the IoU threshold increases and methods that only provide approximate bounding box locations, so their recall drops dramatically. All baseline methods, as well as Bing, Rahtu2011, and EdgeBoxes fall into the latter category. Bing in particular, while providing high repeatability, only provides high recall at $\text{IoU} = 0.5$ and drops dramatically when requiring higher overlap.

When inspecting figure 4 from left to right, one notices that with only few proposal windows our baselines provide low recall compared to non-baseline methods (figure 4a). However as the number of proposals increases baselines such as Gaussian or Sliding-Window become competitive (figure 4b, $\text{IoU} > 0.7$). Eventually for enough windows even simple random baselines such as Uniform become competitive (figure 4c). In relative gain, detection proposal methods have most to offer for low numbers of windows.

Figure 6 presents the results over the ImageNet validation dataset. Compared to Pascal VOC 2007 test set, it includes $10\times$ ground truth classes and $4\times$ images. Somewhat surprisingly the results are almost identical to the ones in 4b. A more detailed inspection reveals that, by design, the statistics of ImageNet match the ones of Pascal VOC. In particular the typical image size, and the mean number of object annotation per image (three) is identical in both datasets. This explains why the recall behaviour is similar, and why Selective-Search, EdgeBoxes, and MCG still perform well despite being crafted for Pascal VOC.

Although the curves look very similar, ImageNet covers 200 object categories, many of them unrelated to Pascal VOC. The good news is that having almost identical results in both datasets confirms that these methods do transfer adequately amongst object classes, and thus can be considered as “true objectness” measures. In other words, the experiments of figure 6 indicate that there is no visible over-fitting of the candidate detection methods towards the 20 Pascal VOC categories.

For researchers looking to benchmark their detection proposal method, figure 5a serves as good overall summary. It puts aside the IoU threshold, because detection proposal users can typically estimate their computation constrains (maximum number of proposals) better than the needed detection proposals localization (IoU threshold). If a specific IoU threshold is desired, then summary figure such as 5b or 5c are suitable.

5 Using the detection proposals

This section analyses the detection proposals in the context of a DPM object detector, namely the LM-LLDA variant [14]. We use a pre-trained model and different proposals to filter its detections at test time. This experiment does not speed-up the detection, but enable evaluating the effect of proposals on the detection quality. A priori detections might get worse (by losing recall), but can also improve if the detection proposal method discards windows that would otherwise be false positives.

Implementation details We take the raw detections of an LM-LLDA model before non-maximum suppression and filter them with the detection proposals of each method. We keep all detections that overlap more than 0.8 IoU with a candidate; the surviving detections are

Method	mAP
LM-LLDA	33.5/34.4
Objectness	25.0/25.4
CPMC	29.9/30.7
Endres2010	31.2/31.6
Sel.Search	31.7/32.3
Rahtu2011	29.6/30.4
Rand.Prim	30.5/30.9
Bing	21.8/22.4
MCG	32.4/32.7
Ranta.2014	30.7/31.3
EdgeBoxes	31.8/32.2
Uniform	16.6/16.9
Gaussian	27.3/28.0
Slid.Window	20.7/21.5
Superpixels	11.2/11.3

Table 1: Detection results on Pascal 2007 (left/right mAP is before/after bounding box regression).

then non-maxima suppressed. As final step we do bounding box regression, as common for DPM models [12]. Note that with this procedure the detector is not only evaluated inside each proposal window, but also around it.

Results Table 1 shows that using 1000 detection proposals decreases detection quality compared to sliding window. When we compare these results with figure 4b, we see that methods with high area under the curve (above 69%), also have high mAP. Although Rahtu and the Gaussian baseline have a similar area under the curve as Objectness and Bing, the latter have lower mAP than the former. We attribute this to higher recall in the high precision area ($\text{IoU} \geq 0.7$). These results show that better localisation of proposals leads to increased detection quality and support the analysis in §4.

The per class recall versus precision plots (in supplementary material) can be grouped into three cases: 1) the best proposal methods do not gain or lose much; 2) proposals sometimes clearly hurt performance (bicycle, bottle, car, chair, motorbike, person, potted plant); 3) proposals improve performance (aeroplane, cat, dining table, dog, sofa). In the case of (2) we see reduced recall but also reduced precision, probably because bad localization decreases scores of strong detections.

6 Conclusion

We evaluated ten detection proposal methods, among which `SelectiveSearch` and `EdgeBoxes` strike by their consistently good performance both in ground truth recall, reasonable repeatability, and tolerable evaluation speed (see table 2). If fast proposals are required properly tuned `EdgeBoxes` seem to provide the best compromise in speed versus quality. When requiring less than 10^3 proposals, `MCG` is the method of choice for high recall if speed is not a concern. `MCG` also obtains top results regarding pure detection quality (table 1).

The extensive evaluation enables researchers to make more informed decisions when considering to use detection proposals. We hope that the provided code and data will encourage authors of detection proposal methods to provide more complete comparisons to related work. Our study also has revealed that most methods suffer from low repeatability due to unstable superpixels, even for the slightest of image perturbation. We foresee room for improvement by using more robust superpixel (or boundary estimation) methods. Finally, our ImageNet experiments have validated that most methods do indeed generalise to different unseen categories (beyond Pascal VOC), and as such can be considered true “objectness” methods.

In future work we plan to study the issue of missing recall in more detail and do further experiments regarding the relation between detection proposals and object detection.

Method		Time	Repeatability	Recall	Detection
<code>Objectness</code> [10]	O	3	.	*	.
<code>CPMC</code> [10]	C	250	-	**	*
<code>Endres2010</code> [10]	E	100	-	**	**
<code>Sel.Search</code> [10]	SS	10	**	***	**
<code>Rahtu2011</code> [10]	R1	3	.	.	*
<code>Rand.Prim</code> [10]	RP	1	*	*	*
<code>Bing</code> [10]	B	0.2	***	*	.
<code>MCG</code> [10]	M	30	*	***	**
<code>Ranta.2014</code> [10]	R4	10	**	.	*
<code>EdgeBoxes</code> [10]	EB	0.3	**	***	**
<code>Uniform</code>	U	0	.	.	.
<code>Gaussian</code>	G	0	.	.	*
<code>SlidingWindow</code>	SW	0	***	.	.
<code>Superpixels</code>	SP	1	*	.	.

Table 2: Overview of detection proposal methods. Time is in seconds. Repeatability, quality, and detection rankings are provided as rough qualitative overview; “-” indicates no data, “.”, “*”, “**”, “***” indicate progressively better results. See paper’s text for details and quantitative evaluations.

References

- [1] B. Alexe, T. Deselaers, and V. Ferrari. What is an object? In *CVPR*, June 2010.
- [2] B. Alexe, T. Deselaers, and V. Ferrari. Measuring the objectness of image windows. *PAMI*, 2012.
- [3] Pablo Arbelaez, J. Pont-Tuset, Jon Barron, F. Marqués, and Jitendra Malik. Multiscale combinatorial grouping. In *CVPR*, 2014.
- [4] J. Carreira and C. Sminchisescu. Constrained parametric min-cuts for automatic object segmentation. In *CVPR*, 2010.
- [5] J. Carreira and C. Sminchisescu. Cpmc: Automatic object segmentation using constrained parametric min-cuts. *PAMI*, 2012.
- [6] Ming-Ming Cheng, Ziming Zhang, Wen-Yan Lin, and Philip H. S. Torr. BING: Binarized normed gradients for objectness estimation at 300fps. In *CVPR*, 2014.
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009.
- [8] P. Dollár, R. Appel, and W. Kienzle. Crosstalk cascades for frame-rate pedestrian detection. In *ECCV*, 2012.
- [9] Ian Endres and Derek Hoiem. Category independent object proposals. In *ECCV*, 2010.
- [10] Ian Endres and Derek Hoiem. Category-independent object proposals with diverse ranking. In *PAMI*, 2014.
- [11] Mark Everingham, S. M. Ali Eslami, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge – a retrospective. *IJCV*, 2014.
- [12] P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *PAMI*, 2010.
- [13] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *IJCV*, 59(2):167–181, 2004.
- [14] R. Girshick and J. Malik. Training deformable part models with decorrelated features. In *ICCV*, 2013.
- [15] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [16] C. Gu, J. Lim, P. Arbelaez, and J. Malik. Recognition using regions. In *CVPR*, 2009.
- [17] C. Gu, P. Arbelaez, Y. Lin, K. Yu, and J. Malik. Multi-component models for object detection. In *ECCV*, 2012.
- [18] H. Harzallah, F. Jurie, and C. Schmid. Combining efficient object localization and image classification. In *ICCV*, 2009.

- [19] Derek Hoiem, Yodsawalai Chodpathumwan, and Qieyun Dai. Diagnosing Error in Object Detectors. In *ECCV*, 2012.
- [20] Ahmad Humayun, Fuxin Li, and James M. Rehg. Rigor: Recycling inference in graph cuts for generating object regions. In *CVPR*, 2014.
- [21] Hongwen Kang, Alyosha Efros, Takeo Kanade, and Martial Hebert. Data-driven objectness. *PAMI*, 2014.
- [22] Santiago Manén, Matthieu Guillaumin, and Luc Van Gool. Prime object proposals with randomized prim’s algorithm. In *ICCV*, 2013.
- [23] Krystian Mikolajczyk, Tinne Tuytelaars, Cordelia Schmid, Andrew Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool. A comparison of affine region detectors. *IJCV*, 2005.
- [24] E. Rahtu, J. Kannala, and M. Blaschko. Learning a category independent object detection cascade. In *ICCV*, 2011.
- [25] P. Rantalankila, J. Kannala, and E. Rahtu. Generating object segmentation proposals using global and local search. In *CVPR*, 2014.
- [26] Antonio Torralba, Kevin P. Murphy, and William T. Freeman. Sharing visual features for multiclass and multiview object detection. *PAMI*, 2007.
- [27] Tinne Tuytelaars. Dense interest points. In *CVPR*, 2010.
- [28] Tinne Tuytelaars and Krystian Mikolajczyk. Local invariant feature detectors: a survey. *Foundations and Trends in Computer Graphics and Vision*, 2008.
- [29] J. Uijlings, K. van de Sande, T. Gevers, and A. Smeulders. Selective search for object recognition. *IJCV*, 2013.
- [30] Koen van de Sande, Jasper Uijlings, Theo Gevers, and Arnold Smeulders. Segmentation as selective search for object recognition. In *ICCV*, 2011.
- [31] P. Viola and M. Jones. Robust real-time face detection. In *IJCV*, 2004.
- [32] Xiaoyu Wang, Ming Yang, Shenghuo Zhu, and Yuanqing Lin. Regionlets for generic object detection. In *ICCV*, 2013.
- [33] C. Lawrence Zitnick and Piotr Dollár. Edge boxes: Locating object proposals from edges. In *ECCV*, 2014.