

Minimizing Supervision for Free-space Segmentation

Satoshi Tsutsui*[†]
Indiana University

stsutsui@indiana.edu

Tommi Kerola[†]
Preferred Networks, Inc.

{tommi,shunta}@preferred.jp

David J. Crandall
Indiana University

djcran@indiana.edu

Abstract

Identifying “free-space,” or safely driveable regions in the scene ahead, is a fundamental task for autonomous navigation. While this task can be addressed using semantic segmentation, the manual labor involved in creating pixel-wise annotations to train the segmentation model is very costly. Although weakly supervised segmentation addresses this issue, most methods are not designed for free-space. In this paper, we observe that homogeneous texture and location are two key characteristics of free-space, and develop a novel, practical framework for free-space segmentation with minimal human supervision. Our experiments show that our framework performs better than other weakly supervised methods while using less supervision. Our work demonstrates the potential for performing free-space segmentation without tedious and costly manual annotation, which will be important for adapting autonomous driving systems to different types of vehicles and environments.

1. Introduction

A critical perceptual problem in autonomous vehicle navigation is deciding whether the path ahead is safe and free of potential collisions. While some problems (like traffic sign detection) may just require detecting and recognizing objects, avoiding collisions requires fine-grained, pixel-level understanding of the scene in front of the vehicle, to separate “free-space” [24] – road surfaces that are free of obstacles, in the case of autonomous cars, for example – from other scene content in view.

Free-space segmentation can be addressed by existing fully-supervised semantic segmentation algorithms [33]. But a major challenge is the cost of obtaining pixel-wise ground truth annotations to train these algorithms: human-labeling of a single object in a single image can take approximately 80 seconds [7], while annotating all road-related

objects in a street scene may take over an hour [12]. The high cost of collecting training data may be a substantial barrier for developing autonomous driving systems for new environments that have not yet received commercial attention (e.g. in resource-poor countries, for off-road contexts, for autonomous water vehicles, etc.), and especially for small companies and research groups with limited resources.

In this paper, we develop a framework for free-space segmentation that minimizes human supervision. Our approach is based on two straightforward observations. First, free-space has a strong *location prior*: pixels corresponding to free space are likely to be located at the bottom and center of the image taken by a front-facing camera, since in training data there is always free-space under the vehicle (by definition). Second, a free-space region generally has *homogeneous texture* since road surfaces are typically level and smooth (e.g. concrete or asphalt in an urban street).

To take advantage of these observations, we first group together pixels with low-level homogeneous texture into superpixels. We then select candidate free-space superpixels through a simple clustering algorithm that incorporates both the spatial prior and appearance features (§3.3). The remaining challenge is to create higher-level features for each superpixel that semantically distinguish free-space. We show that features from a CNN pre-trained on ImageNet (§3.1) perform well for free-space when combined with *superpixel alignment*, a novel method that aligns superpixels with CNN feature maps (§3.2). Finally, these results are used as labels to train a supervised segmentation method (§3.4) for performing segmentation on new images.

We note that our framework does not need any image annotations, so collecting annotated data is a simple matter of recording vehicle-centric images while navigating the environment where free-space segmentation is needed, and then running our algorithm. The human effort required is reduced to specifying the location prior and adjusting hyperparameters such as superpixel granularity and the number of clusters. This form of supervision requires little effort because the technique is not very sensitive to the exact values of these parameters, as we empirically demon-

*Part of this work was done as an intern at Preferred Networks, Inc.

[†]The first three authors contributed equally. The order was decided by using the paper ID as a random seed and then calling `np.random.permutation`.

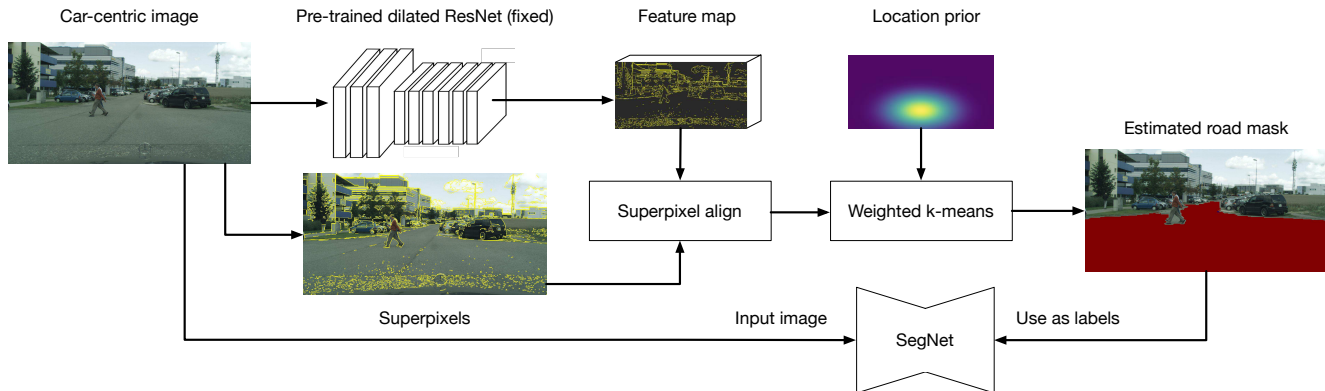


Figure 1. Overview of our method. We extract features from a dilated ResNet pre-trained on ImageNet and leverage a novel superpixel alignment method and location prior clustering to generate masks for training a segmentation CNN. Our method requires no manual annotation of free-space labels. Best viewed in color.

strate with experiments on the well-established, publicly-available Cityscapes dataset [12]. Our quantitative evaluation shows that our framework yields better performance than various baselines, even those that use more supervision than we do (§4).

In summary, we make the following contributions:

- We develop a novel framework for free-space segmentation that does not require any image-level annotations, by taking advantage of the unique characteristics of free-space;
- We propose a novel algorithm for combining CNN feature maps and superpixels, and a clustering method that incorporates prior knowledge about the location of free-space; and
- We show that our approach performs better than other baselines, even those that require more supervision.

2. Related Work

Fully supervised segmentation. Many recent advances in semantic segmentation have been built on fully convolutional networks (FCNs) [31], which extend CNNs designed for image classification by posing semantic segmentation as a dense pixel-wise classification problem. This dense classification requires high resolution feature maps for prediction, so FCNs add upsampling layers into the classification CNNs (which otherwise usually perform downsampling through pooling layers). SegNet [6] improves upon this and introduces an unpooling layer for upsampling, which reflects the pooling indices used in the downsampling phase. We use SegNet here, although our technique is flexible enough to be used with other FCNs as well.

A problem with CNN pooling layers is that they discard spatial information that is critical for image segmentation. One solution is to use dilated (or ‘atrous’) convolutions [46], which allow receptive field expansion without

pooling layers. Dilated convolutions have been incorporated into recent frameworks such as DeepLab [9] and PSP-Net [48]. Although our work does not focus on engineering CNN architectures, this direction inspired our choice of CNN for image feature extraction, since we similarly want to obtain a high resolution feature map. In particular, we use dilated ResNet [47] trained on ImageNet, yielding a higher resolution feature map than the normal ResNet [21].

Weakly supervised segmentation. Since ground-truth segmentation annotations are very costly to obtain, many techniques for segmentation have been proposed that require weaker annotations, such as image tags [13, 26, 34, 36, 39, 45], scribbles [29], bounding boxes [25], or videos of objects [22, 40]. At a high level of abstraction, our work can be viewed as a tag-based weakly supervised method, in that we assume all images have a “tag” of free-space. However, most previous studies mainly focus on foreground objects, so are not directly applicable for free-space, which can be regarded as background [35]. From a technical perspective, some methods propose new CNN architectures [13] or better loss functions [26], while others focus on automatically generating segmentation masks for training available CNNs [25]. We follow the latter approach here of generating segmentation masks for CNNs. We also do not use the approach of gradually refining the segmentation mask [39], because we believe that autonomous vehicles require a high-quality trained CNN even at the stage of initial deployment.

Free-space segmentation. Free-space segmentation is the task of estimating the space through which a vehicle can drive safely without collision. This task is critical for autonomous driving and has traditionally been addressed by geometric modeling [3, 5, 27, 44], handcrafted features [4, 18], or even a patch-based CNN [2]. We use

FCNs in this paper, which Oliveira *et al.* [33] demonstrated to be efficient for road segmentation.

Since pixel-wise ground truth annotations are so expensive to obtain, several papers have investigated weakly supervised free-space segmentation. While an early study [17] trains a probabilistic model, other papers train FCNs [28, 35, 37, 43]. Saleh *et al.* [35] develop a video segmentation algorithm for general background objects including free-space on a road. Tsutsui *et al.* [43] propose distantly supervised monocular image segmentation. However, both methods require additional images to train a saliency or attention extractor. Laddha *et al.* [28] use external maps of the road indexed against the vehicle position according to GPS. Sanberg *et al.* [37] and Guo *et al.* [17] use stereo information for automatically generating segmentation masks. We distinguish our work from these studies in that we only use a collection of monocular vehicle-centric images, which makes our approach even less supervised than most others.

For evaluating free-space segmentation, KITTI [16] and CamVid [8] are older datasets that are not large enough to leverage the power of CNNs. Recently, a larger dataset called Cityscapes [12] was proposed for object segmentation in autonomous driving. We conduct our experiments on Cityscapes, since existing work has found that CNNs trained on Cityscapes perform better than other state-of-the-art methods on KITTI and CamVid [12].

3. Our approach

We now describe our technique for automatically generating annotations suitable for training a free-space segmentation CNN. Our technique relies on two main assumptions about the nature of free-space: (1) that free-space regions tend to have homogeneous texture (e.g., caused by smooth road surfaces), and (2) there are strong priors on the location of free-space within an image taken from a vehicle. The first assumption allows us to use superpixels to group similar pixels. As in previous work [10, 22, 29, 40, 43], we use the Felzenszwalb and Huttenlocher graph-based segmentation algorithm [15] to create the superpixels, since the specific superpixel algorithm is not the focus of this study.

The second assumption allows us to find “seed” superpixels that are very likely to be free-space, based on the fact that free-space is usually near the bottom and center of an image taken by a front-facing in-vehicle camera. A very naive method would be to select superpixels covering predefined locations based on the prior, but this would ignore any semantic or higher level features other than the texture features used for generating superpixels. We thus cluster superpixels based on semantic features and automatically select the cluster likely corresponding to free-space based on the location prior, as described in §3.3. We perform this clustering on multiple images at a time, to be robust against occasional images which do not satisfy the prior assumption.

tion.

An important question is how to extract semantic-level features from each superpixel. We show that the features extracted from CNNs pre-trained on ImageNet are generic enough (§3.1) for our task, and we develop a novel technique called *superpixel alignment* that efficiently aggregates CNN features for the region within a superpixel (§3.2). Finally, superpixel clustering automatically generates a free-space pixel mask, which we then use to train supervised CNNs for segmentation (§3.4).

The reader may wonder why we do not cluster the CNN features directly, given that they capture semantic information. However, because certain parts of free-space is semantically more important than others, direct clustering results would not be smooth and cohesive. This is visually confirmed by comparing the two clustering results shown in Figure 3.

3.1. Features for Clustering

We cluster superpixels based on features extracted from a CNN pre-trained on ImageNet. Such features have been found to capture latent features having rich semantic information for semantic segmentation [36], even though the ImageNet challenge does not include free-space as one of its annotated classes. Much work has found that these features are surprisingly general across vastly different domains including document image analysis [19] and medical image analysis [14]. This is probably because early layers in convolutional neural networks tend to learn low-level features (e.g., edges), while later layers capture increasing amounts of semantic information, with the final layers capturing features suitable for the explicit classification problem (e.g., object types like cars) that the network was trained to solve [36]. We confirmed this tendency by visualizing feature maps of the 26-layer dilated ResNet [47] that is trained for the task of ImageNet classification, and decided to use the last layer feature map, which indeed seemed to capture higher level information. We note that this type of manual inspection has also been performed in previous work [36].

Among other CNN architectures, we intentionally select a dilated network architecture in order to produce higher resolution feature maps, which are important for being able to localize the road with a fine level of granularity.

3.2. Superpixel Alignment

We now wish to extract appearance features for each superpixel. While the dilated ResNet features capture semantic information, they are not well localized for free-space, so we align them to spatially coherent superpixels to create a better representation of the scene. To do this, we propose a new method called *superpixel alignment*, which is inspired by RoIAlign [20]. The technique applies bilinear interpo-

function LOCATIONPRIORKMEANS($k, \{S_i\}_{\forall i}, \mu_{\text{prior}}, \sigma_{\text{prior}}$)

$\forall i, \forall S_{xy} \in S_i : p_{xy} = \text{spatial_coord}(S_{xy})$

$\forall i : w_i = \frac{1}{|S_i|} \sum_{S_{xy} \in S_i} e^{-\|p_{xy} - \mu_{\text{prior}}\|^2 / (2\sigma_{\text{prior}}^2)}$

$\forall i : m_i \leftarrow \begin{cases} 0 & \text{if } w_i > \text{median}_j w_j \\ \text{unif}(1, k-1) & \text{otherwise} \end{cases}$

while not converged **do**

$c_0 \leftarrow \sum_{m_i=0} w_i S_i / \sum_{m_i=0} w_i$

$\forall q > 0 : c_q \leftarrow \sum_{m_i=q} (1 - w_i) S_i / \sum_{m_i=q} w_i$

$\forall i : m_i \leftarrow \text{argmin}_q \|S_i - c_q\|_2^2$

end while

return $\{c_0, \dots, c_{k-1}\}$ \triangleright Cluster centers

end function

Figure 2. Location prior k-means.



Figure 3. Example of using location prior k-means on dilated ResNet (CNN) or superpixel alignment features to find the location of the free-space. Note that due to the design of our algorithm, the free-space will typically be the first cluster (red), which means that we do not need any sophisticated post-processing in order to find the free-space cluster. Best viewed in color.

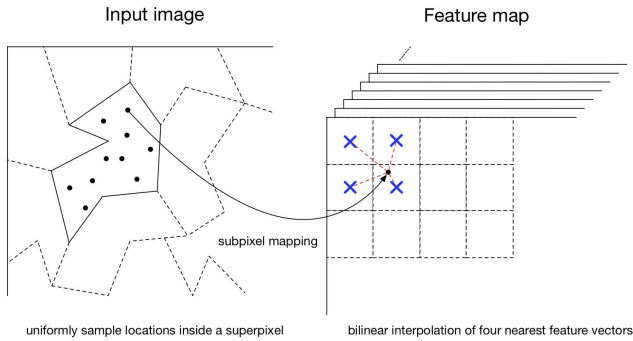


Figure 4. Illustration of superpixel alignment. The feature vector for each superpixel is defined by randomly sampling 10 points inside the superpixel, and taking the average of their bilinearly interpolated feature maps.

lation of the CNN feature maps for a random subset of the pixels inside each superpixel. More precisely, we perform bilinear interpolation [23] of dilated ResNet features F_{mn}^c

at spatial location (m, n) and channel c as

$$S_{xy}^c = \sum_{(m,n) \in \mathcal{N}_{xy}} F_{nm}^c \max(0, \hat{\Delta}_{xm}) \max(0, \hat{\Delta}_{yn}), \quad (1)$$

where $\hat{\Delta}_{xm} = 1 - |x - m|$ and \mathcal{N}_{xy} is the set of neighbors for spatial location (x, y) in superpixel S . We sample 10 locations uniformly at random inside each superpixel, and then use the four nearest neighbors of each selected pixel for the bilinear interpolation. Finally, we aggregate the features inside each superpixel using average pooling. Note that unlike RoIAlign, we assume that each superpixel consists of a homogeneous set of pixels; this avoids the need for computing the bilinear interpolation densely for all pixels by instead using a small randomly sampled set, which we have found works well in practice. To improve the spatial cohesiveness of the feature, we append the centroid of the spatial coordinates of the superpixel to the pooled feature vector. This gives us one image feature for each superpixel. The procedure for superpixel alignment is summarized in Fig. 4.

3.3. Superpixel Clustering

Using the features defined in the last section, we can now apply any standard clustering algorithm. An important remaining problem, however, is how to determine which cluster corresponds with free-space. A simple solution would be to select the largest cluster appearing in the bottom half the image, for example, but this would fail in crowded scenes with large numbers of foreground objects on the road.

Instead, inspired by previous work [1, 42], we use prior information about the spatial location of free-space, namely that the road surface should usually be immediately above the visible chassis of the ego-vehicle. To do this, we adapt Lloyd’s algorithm [30] for solving a weighted variant of the k-means clustering problem. We represent the prior as an average of Gaussians

$$w_i = \frac{1}{|S_i|} \sum_{S_{xy} \in S_i} \exp\left(-\frac{\|p_{xy} - \mu_{\text{prior}}\|^2}{2\sigma_{\text{prior}}^2}\right), \quad (2)$$

such that each superpixel S_i has a prior weight w_i that is parameterized by $\mu_{\text{prior}} = [0.75, 0.5]$ and $\sigma_{\text{prior}} = [0.1, 0.1]$ wrt. the image dimensions (estimated empirically) and the spatial coordinates p_{xy} of each pixel inside the superpixel. In practice, we manually adjust the prior parameters empirically with a small number of example images. Subsequently, we initialize half of the pixels to the free-space cluster (which we assume is the first cluster) based on these weights. The first cluster is then encouraged to consist of pixels corresponding to free-space by setting its cluster center to be the spatially weighted average of features assigned to it. The other clusters have a repellent weight assigned to




Figure 5. Batch clustering. (top) Input images. (middle) Results of clustering each image individually. (bottom) Results of clustering the four images together in a batch manner. The left image has a spot around the center of the location prior so only the spot is recognized as free-space when clustered individually, but batch clustering avoids this mistake.

their members to encourage them to spatially spread away from the location prior. Cluster memberships are updated in the same manner as the standard k-means algorithm without taking the weights into account. Our algorithm is summarized in Fig. 2, and an example of the output of the algorithm is shown in Fig. 3.

Although our cluster update breaks the convergence criterion of standard k-means clustering [11], we have found that in practice it usually converges to a stable solution. We note that similar prior information could also be incorporated into other types of clustering algorithms, such as spectral clustering [38].

Batch image clustering. Of course, while the spatial prior assumption on free-space is reasonable in general, it is often violated in individual images (e.g., a vehicle or pedestrian could be located in the center of the location prior, which could cause the algorithm to incorrectly assign the first cluster to consist of features corresponding to non-road locations). We circumvent this issue by clustering superpixels from multiple images at the same time, which we call batch clustering. In Fig. 5, we show an example where only a single spot at the center of the location prior is recognized as free-space, but clustering with three other images prevents this mistake. As our experiments will show, batch clustering is effective for generating higher quality segmentation masks.

3.4. CNN Training from Generated Mask

Once the masks for free-space have been obtained by superpixel clustering (an example can be seen in Fig. 6), we then use these automatically generated masks to train a road segmentation CNN using supervised training.

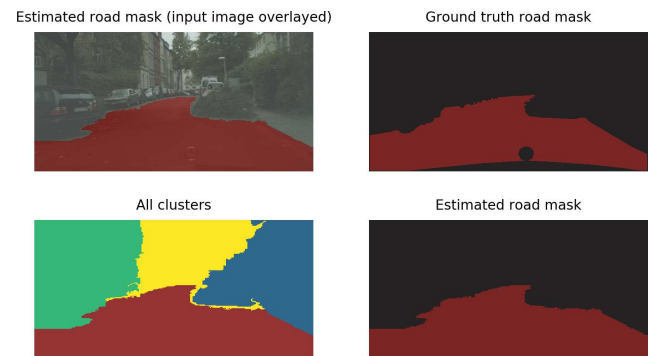


Figure 6. Example of superpixel alignment compared to ground truth.

4. Experiments

Dataset. We conducted a series of experiments on the established Cityscapes [12] dataset to evaluate our proposed method. This dataset is designed for evaluating segmentation algorithms for autonomous driving applications, and includes a set of fine-grained pixel-wise annotations for 19 types of traffic objects. We only use the ‘road’ class and treat it as free-space. We report the intersection over union (IoU) metric, while ignoring void regions not defined in the ground truth.

4.1. Automatic Free-Space Mask Generation.

We first evaluated the quality of the automatically generated masks for free-space, and conducted ablation experiments to study how each part of our technique contributes to the algorithm. Table 1 summarizes the results, in terms of IoU on the Cityscapes dataset. We emphasize that our model has never seen the training set ground truth before.

Technique	IoU
raw CNN features location prior clustering	0.530
+ batch clustering	0.568
+ superpixel overlap [43]	0.620
superpixel align location prior clustering	0.758
+ batch clustering	0.764

Table 1. Ablation study results for automatic road mask generation on the Cityscapes training set.

We compare our proposed superpixel alignment road prior clustering method with directly clustering the CNN features from the dilated ResNet. As can be seen, superpixel alignment achieves higher IoU than the raw CNN features. We can also see that batch clustering improves both methods, and helps superpixel alignment to achieve higher IoU. For the sake of comparison, we also compare with previous work that combines superpixels and a saliency map [43]. We treat the free-space cluster from the raw CNN features as saliency, and use these for selecting superpixels. As can be seen, this technique improves the performance of the raw CNN features, but is still unable to beat superpixel alignment.

Parameter sensitivity. Although our method does not use any annotations, it does rely on some manually selected parameters. In practice, we chose these values by visually investigating a small (~ 10) number of images. To measure the sensitivity of our method to these values, we changed each of three key parameters and compared the final road IoU on the training set: number of clusters (default 4), batch size (default 30), and superpixel granularity scale (default 300). Results are shown in Figure 7. While performance did vary with differing parameter values, of course, we found that the final IoU metric differed by only a few percent across even relatively extreme parameter settings.

In more detail, Figure 7(a) shows the sensitivity for the number of clusters. We see that having too few clusters makes it difficult to separate road from other parts of the image, while having too many also has diminishing returns as the free-space is eventually split into multiple clusters. The effect of varying the batch size is shown in Figure 7(b); increasing the batch size improves the IoU. Finally, Figure 7(c) shows that smaller superpixels tend to work slightly better, presumably since they avoid under-segmentation which can lead to false positives (e.g., due to merging free-space with building walls). These results suggest that our method is relatively robust to the choice of parameter values.

4.2. Training a CNN from the Generated Mask

We next tested our algorithm in the context that it was designed for: automatically generating pixel-level annotations

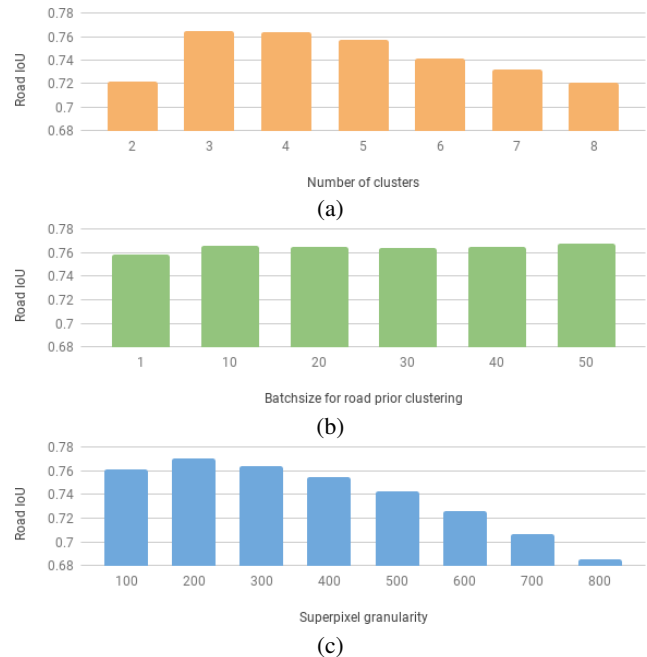


Figure 7. Parameter sensitivity with respect to (a) number of clusters, (b) batch size, and (c) superpixel granularity.

for training a supervised free-space segmentation model. In particular, we used our automatically generated pixel-level annotations from the previous section to train SegNet [6], although we note that our method is agnostic to the choice of model so any CNN could be used instead.

Experimental setup. We trained SegNet Basic with our generated masks as labels for the Cityscapes training images using the Chainer framework [32, 41]. We used the validation set to evaluate our method against several baselines, since the test set annotations are not publicly available (and the evaluation server restricts the number of submissions to avoid overfitting to the test dataset). We emphasize that no hyper-parameters were tuned based on the validation set; we treated it as if it were the test set.

Baselines. We compared our technique with six baseline methods, as summarized in Table 2. The first two baselines serve as simple indications of how well trivial solutions work on this task: **Largest superpixel** uses just the single largest superpixel as the free-space annotation mask, and **Bottom half** blindly uses the bottom half of the image as the free-space mask. In contrast, **Ground truth as road cluster** uses the ground truth mask as the clustering results and combines them with the superpixels in a similar manner as in previous work [43]. **Distant supervision** is the technique of Tsutsui *et al.* [43], which shares a similar motivation with us and uses external images (which they call a ‘distant supervisor’) to perform road segmentation in

Method	Required annotation	IoU
largest superpixel	none	0.659
bottom half	none	0.720
ground truth as road cluster	-	0.824
distant supervision [43]	additional images	0.800
video segmentation [35]	additional images	0.759
fully supervised (from [43])	pixel-wise	0.853
ours (generated masks)	none	0.761
ours (CNN training)	none	0.835

Table 2. Results evaluated on the Cityscapes validation set.

a weakly supervised manner. **Video segmentation** is the technique of Saleh *et al.* [35], which was originally proposed for general background segmentation and not only uses external images but also videos. Finally, **Fully supervised** trains the SegNet model from ground truth annotations.

Evaluation. We computed the IoU of SegNet trained on the output of our weakly-supervised algorithm, and obtained an IoU of 0.835 on the Cityscapes validation set. This is much higher than the trivial baselines **largest superpixel** and **bottom half**, which yielded IoUs of 0.659 and 0.720, respectively. The relatively high IoU of the bottom half baseline might make this task seem easy, but we emphasize that our method has a much lower false positive rate, which is crucial for employing the method in a practical system to avoid collisions. In particular, **Bottom half** gives precision 0.754, while our generated masks and trained CNN have a precision of 0.867 and 0.898, respectively, thus showing that our method is less prone to fatal false positives, such as a car being mistaken for road. In contrast, false negatives are less important in our application, since they may just mean that the car is unable to drive to a certain point, still preserving safe behavior.

Our technique also outperforms **distant supervision** and **video segmentation**, even though they require more annotations. Of course, our technique also imposes more assumptions, since those approaches were designed for general video segmentation and our cues are customized to free-space, but nonetheless we believe it should be notable because we do not use any motion cues with video. **Ground truth as road cluster**, which can be viewed as an upper bound on the performance of any technique using superpixels (e.g., [10, 22, 29, 40]), yields an IoU of 0.824.

Of course, **fully supervised** somewhat outperforms our results (0.853¹ vs 0.835). Nonetheless, it is impressive that

¹This is a bit worse than the original SegNet [6] because we use their simplest model, SegNet Basic, and only train with binary classes while the original one used more classes.

Method	Required annotation	IoU
video segmentation [35]	additional images	0.785
ours (CNN training)	none	0.857

Table 3. Results evaluated on the Cityscapes test set.

our technique achieves 98% of the IoU of the fully supervised model, without requiring the tedious pixel-wise annotations for each image. This indicates that our proposed method is able to perform proper free-space segmentation while using no manual annotations for training the CNN.

Some sample results on the validation set can be seen in Figure 8. We see that while our method typically follows the shape of the true road and avoids labeling cars as road, it has some trouble with *e.g.* pedestrian legs and some parts of the sidewalk being labeled as road. The last row also shows an example of a false positive, where a car in front of the ego-vehicle is not able to be separated from the estimated free-space. In future work, it would be interesting to investigate more powerful (albeit more computationally heavy) CNN architectures that might help mitigate these problems [48].

Finally, we also evaluated our best model on the Cityscapes test set evaluation server. Results are shown in Table 3, where we can see that consistent with the validation set results, our method is able to gain better performance than the general video segmentation approach [35].

5. Conclusion

In this paper, we developed a new framework for minimizing human supervision for free-space segmentation, using assumptions of the the characteristics of free-space. Our method extracts free-space by performing clustering of superpixel features, which are created by a novel superpixel alignment method that bases features on the last layer of an ImageNet-pretrained CNN. We use a location prior to select the cluster corresponding to free-space and then perform training of a free-space segmentation CNN. Unlike previous work, our method needs no annotations, and experimental results demonstrate superior performance compared to other methods, even ones that use more information.

As future work, we plan to automatically generate the location prior conditioned on the input image to better handle segmentation of distant free-space, which is a weakness of the current model. Extending the model to other application domains with high cost of collecting training data, such as robots moving on a house floor or autonomous water vehicles, is another interesting direction.

Acknowledgments We would like to thank the authors of the video segmentation paper [35] for sharing their results on the validation set. We would also like to thank Masaki Saito and Richard Calland for helpful discussions.

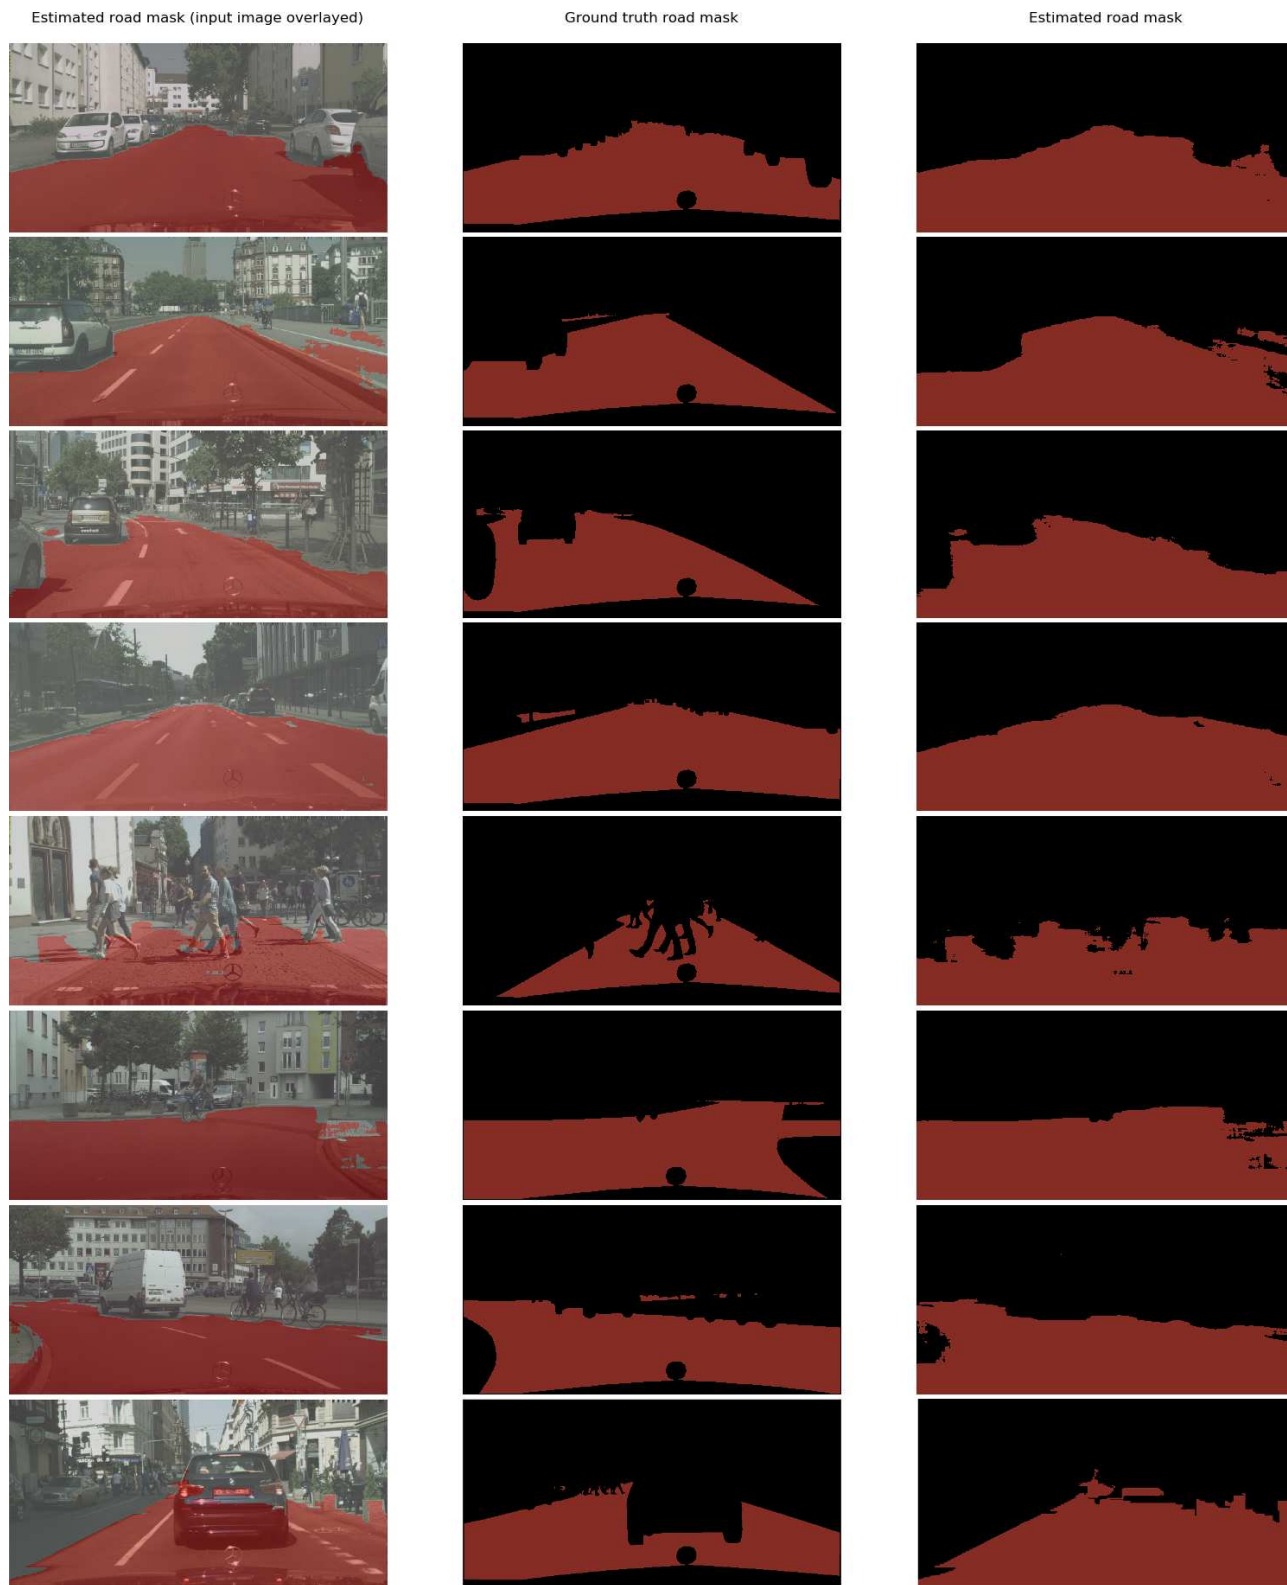


Figure 8. Sample results on the Cityscapes validation set, comparing our method with ground truth. The last row shows a failure case with a false positive.

References

- [1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2282, 2012. 4
- [2] J. Alvarez, T. Gevers, Y. LeCun, and A. Lopez. Road scene segmentation from a single image. In *ECCV*, 2012. 2
- [3] J. M. Alvarez, T. Gevers, and A. M. Lopez. 3d Scene Priors for Road Detection. In *CVPR*, 2010. 2
- [4] J. M. Alvarez and A. M. Lopez. Road detection based on illuminant invariance. *IEEE Trans. Intell. Transp. Syst.*, 2011. 2
- [5] H. Badino, U. Franke, and R. Mester. Free space computation using stochastic occupancy grids and dynamic programming. In *ICCV Workshop on Dynamical Vision*, 2007. 2
- [6] V. Badrinarayanan, A. Kendall, and R. Cipolla. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2017. 2, 6, 7
- [7] A. Bearman, O. Russakovsky, V. Ferrari, and L. Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *ECCV*, 2016. 1
- [8] G. J. Brostow, J. Fauqueur, and R. Cipolla. Semantic object classes in video: A high-definition ground truth database. *Pattern Recognition Letters*, 30(2):88–97, 2009. 3
- [9] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2016. 2
- [10] Y.-H. Chen, W.-Y. Chen, Y.-T. Chen, B.-C. Tsai, Y.-C. F. Wang, and M. Sun. No more discrimination: Cross city adaptation of road scene segmenters. In *ICCV*, 2017. 3, 7
- [11] D. M. Christopher, R. Prabhakar, and S. Hinrich. Introduction to information retrieval. *An Introduction To Information Retrieval*, 151:177, 2008. 5
- [12] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 1, 2, 3, 5
- [13] T. Durand, T. Mordan, N. Thome, and M. Cord. WILDCAT: Weakly Supervised Learning of Deep ConvNets for Image Classification, Pointwise Localization and Segmentation. In *CVPR*, 2017. 2
- [14] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115–118, 2017. 3
- [15] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *International Journal of Computer Vision (IJCV)*, 59(2):167–181, 2004. 3
- [16] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. *CVPR*, 2012. 3
- [17] C. Guo, S. Mita, and D. McAllester. Robust Road Detection and Tracking in Challenging Scenarios Based on Markov Random Fields With Unsupervised Learning. *IEEE Trans. Intell. Transp. Syst.*, 13(3):1338–1354, 2012. 3
- [18] S. Hänisch, R. H. Evangelio, H. H. Tadjine, and M. Pätzold. Free-Space Detection with Fish-Eye Cameras. In *IEEE Intelligent Vehicles Symposium (IV)*, 2017. 2
- [19] A. W. Harley, A. Ufkes, and K. G. Derpanis. Evaluation of deep convolutional nets for document image classification and retrieval. In *ICDAR*, 2015. 3
- [20] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *ICCV*, 2017. 3
- [21] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2
- [22] S. Hong, D. Yeo, S. Kwak, H. Lee, and B. Han. Weakly supervised semantic segmentation using web-crawled videos. In *CVPR*, 2017. 2, 3, 7
- [23] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *NIPS*, 2015. 4
- [24] J. Janaia, F. Guney, A. Behl, and A. Geiger. Computer vision for autonomous vehicles: Problems, datasets and state-of-the-art. *arXiv 1704.05519*, 2017. 1
- [25] A. Khoreva, R. Benenson, J. Hosang, M. Hein, and B. Schiele. Simple Does It: Weakly Supervised Instance and Semantic Segmentation. In *CVPR*, 2017. 2
- [26] A. Kolesnikov and C. H. Lampert. Seed, Expand and Constrain: Three Principles for Weakly-Supervised Image Segmentation. In *ECCV*, 2016. 2
- [27] H. Kong, J.-Y. Audibert, and J. Ponce. Vanishing point detection for road detection. In *CVPR*, 2009. 2
- [28] A. Laddha, M. K. Kocamaz, L. E. Navarro-Serment, and M. Hebert. Map-supervised road detection. In *IEEE Intelligent Vehicles Symposium (IV)*, pages 118–123. IEEE, 2016. 3
- [29] D. Lin, J. Dai, J. Jia, K. He, and J. Sun. ScribbleSup: Scribble-Supervised Convolutional Networks for Semantic Segmentation. In *CVPR*, 2016. 2, 3, 7
- [30] S. Lloyd. Least squares quantization in pcm. *IEEE Trans. Inf. Theory*, 28(2):129–137, 1982. 4
- [31] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 2
- [32] Y. Niitani, T. Ogawa, S. Saito, and M. Saito. Chainercv: a library for deep learning in computer vision. In *ACM Multimedia*, 2017. 6
- [33] G. L. Oliveira, W. Burgard, and T. Brox. Efficient deep models for monocular road segmentation. In *IROS*, 2016. 1, 3
- [34] P. O. Pinheiro and R. Collobert. From image-level to pixel-level labeling with convolutional networks. In *CVPR*, 2015. 2
- [35] F. Sadat Saleh, M. Sadegh Aliakbarian, M. Salzmann, L. Petersson, and J. M. Alvarez. Bringing background into the foreground: Making all classes equal in weakly-supervised video semantic segmentation. In *ICCV*, 2017. 2, 3, 7
- [36] F. Saleh, M. S. A. Akbarian, M. Salzmann, L. Petersson, J. M. Alvarez, and S. Gould. Incorporating network built-in priors in weakly-supervised semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2017. 2, 3
- [37] W. P. Sanberg, G. Dubbelman, and P. H. de With. Free-space detection with self-supervised and online trained fully

- convolutional networks. In *IS&T Electronic Imaging - Autonomous Vehicles and Machines (EI-AVM)*, pages 54–61, 2017. 3
- [38] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):888–905, 2000. 5
- [39] W. Shimoda and K. Yanai. Distinct class-specific saliency maps for weakly supervised semantic segmentation. In *ECCV*, 2016. 2
- [40] P. Tokmakov, K. Alahari, and C. Schmid. Weakly-supervised semantic segmentation using motion cues. In *ECCV*, 2016. 2, 3, 7
- [41] S. Tokui, K. Oono, S. Hido, and J. Clayton. Chainer: a next-generation open source framework for deep learning. In *NIPS workshop on Machine Learning Systems*, 2015. 6
- [42] G. C. Tseng. Penalized and weighted k-means for clustering with scattered objects and prior information in high-throughput biological data. *Bioinformatics*, 23(17):2247–2255, 2007. 4
- [43] S. Tsutsui, T. Kerola, and S. Saito. Distantly supervised road segmentation. In *ICCV Workshop on Computer Vision for Road Scene Understanding and Autonomous Driving.*, 2017. 3, 6, 7
- [44] A. Wedel, H. Badino, C. Rabe, H. Loose, U. Franke, and D. Cremers. B-spline modeling of road surfaces with an application to free-space estimation. *IEEE Trans. Intell. Transp. Syst.*, 2009. 2
- [45] Y. Wei, J. Feng, X. Liang, M.-M. Cheng, Y. Zhao, and S. Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *CVPR*, 2017. 2
- [46] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. In *ICLR*, 2015. 2
- [47] F. Yu, V. Koltun, and T. Funkhouser. Dilated residual networks. In *CVPR*, 2017. 2, 3
- [48] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid Scene Parsing Network. In *CVPR*, 2017. 2, 7