

Fast Scene Understanding for Autonomous Driving

Davy Neven, Bert De Brabandere, Stamatios Georgoulis, Marc Proesmans and Luc Van Gool

Abstract—Most approaches for instance-aware semantic labeling traditionally focus on accuracy. Other aspects like runtime and memory footprint are arguably as important for real-time applications such as autonomous driving. Motivated by this observation and inspired by recent works that tackle multiple tasks with a single integrated architecture [13], [20], [22], in this paper we present a real-time efficient implementation based on ENet [18] that solves three autonomous driving related tasks at once: semantic scene segmentation, instance segmentation and monocular depth estimation. Our approach builds upon a branched ENet architecture with a shared encoder but different decoder branches for each of the three tasks. The presented method can run at 21 fps at a resolution of 1024x512 on the Cityscapes dataset without sacrificing accuracy compared to running each task separately.

I. INTRODUCTION

The last years the re-appearance of Convolutional Neural Networks (CNNs), whose origin traces back to the 1970s and 1980s, has led to significant advances in many computer vision tasks, such as image classification [14], object detection [8], semantic scene segmentation [16], instance segmentation [9], and monocular depth estimation [6] to name a few. The majority of these works rely on fine-tuning or slightly altering a CNN architecture, typically the VGG network [19], resulting in task-specific CNNs with long inference times that each require a single GPU to run. Admittedly, this is not enough for autonomous driving applications where many of the aforementioned tasks should run in parallel, in real-time, and on a limited number of GPU devices. Furthermore, as shown in recent works [13], [20], [22] there is merit in combining multiple tasks in a single integrated architecture, as one task might benefit from another leaving smaller space for 'blindspots', which is crucial for self-driving vehicles.

Motivated by these observations, in this paper we focus on street scene understanding and present an efficient implementation that combines the tasks of semantic scene segmentation, instance segmentation, and monocular depth estimation. Unlike state-of-the-art methods, that use networks with huge number of parameters and long inference times (e.g. VGG [19], SegNet [2], FCN [16]), we build upon a real-time architecture, in particular ENet [18] that has proven to offer image processing rates higher than 10 fps on a single GPU device. Specifically, we use a common ENet encoding step for all tasks, but introduce a branched ENet architecture for the decoding step (i.e. one branch for each of the three different tasks). Fig. 1 gives an overview of our approach.

Although we do not introduce a new architecture, in this paper we show how to efficiently combine existing

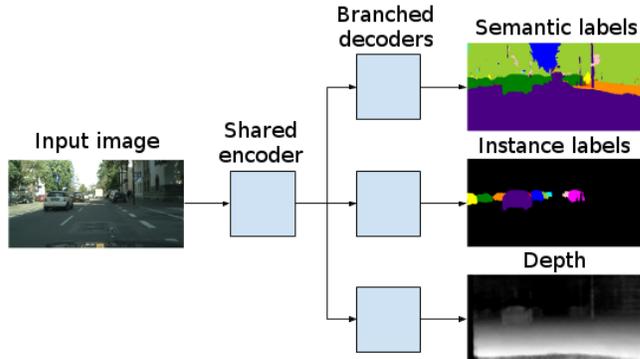


Fig. 1. Overview of our method. From left to right: the input image is passed through the encoding step of an ENet-inspired architecture to create feature maps, which in turn are forwarded to different branches that perform decoding to arrive at the three outputs, i.e. semantic labels, instance labels and depth. For more details regarding the architecture visit Sec. III.

components to build a solid architecture for real-time scene understanding. In Sec. II we describe related work on integrated architectures that tackle multiple tasks. Next, we present the implementation details of our method in Sec. III. Finally, in Sec. IV and V we report results for each of the tasks and provide some insights into the strengths and limitations of the presented approach.

II. RELATED WORK

The amount of research performed in literature on the three main tasks studied in this paper, i.e. semantic scene segmentation, instance segmentation, and monocular depth estimation, is vast. In what follows, we solely focus on related works that have combined one or more of these tasks in a single integrated architecture.

Eigen and Fergus [5] addressed the tasks of depth prediction, surface normal estimation, and semantic labeling using a multiscale convolutional network architecture that progressively refines predictions from a sequence of scales. Uhrig *et al.* [22] presented a method that leverages a FCN network to predict semantic labels, depth, and an instance-based encoding using each pixel's direction towards its corresponding instance center and consequently applying low-level computer vision techniques. Kokkinos [13] went one step further from the previous approaches, and introduced a CNN, namely UberNet, that jointly handles low-, mid-, and high-level vision tasks in a unified architecture. His universal network tackles boundary detection, normal estimation, saliency estimation, semantic segmentation, human part segmentation, semantic boundary detection, region proposal

generation, and object detection. Despite obtaining competitive performance while jointly addressing many different tasks, all these approaches suffer from poor inference times making them unsuitable for real-time autonomous driving applications with high frame-rate demands.

Recently, Teichmann *et al.* [20] argued that improving the computational times is more important than improving performance, especially for the case of self-driving vehicles. They presented an approach to joint classification, detection, and semantic segmentation via a unified architecture where the encoder is shared amongst the three tasks, marginally reaching a computational time of 10 fps on the KITTI dataset. Our approach also focuses on further improving the computational times but addresses different tasks, in particular semantic scene segmentation, instance segmentation, and monocular depth estimation, and achieves a computational time of 21 fps on the Cityscapes dataset. To our knowledge this is the first system to estimate depth, semantic and instance segmentation at these frame-rates.

III. METHOD

In order to predict depth, semantic and instance segmentation in real-time, we modify the ENet architecture into a multi-branched network, having three output branches, one for each task (see Fig. 1). The original network, as described in [18], consists of an encoding step that has three stages (stage 1, 2, 3) and a decoding step that has two stages (stage 4, 5). Since the ENet decoding step is merely for upscaling and finetuning the output of the encoding step, sharing the full encoder (stages 1, 2, 3) between all branches would lead to poor results. Instead, our multi-branch network is constructed as follows: our shared "encoder" consists of stages 1 and 2 of the original Enet network, before continuing to each branch that combines stage 3 of the original ENet encoder with stages 4 and 5 of the original ENet decoder. In what follows, we dive into the details of the individual branches that each performs one task.

Semantic segmentation The semantic segmentation branch is trained using the standard pixel-wise cross-entropy loss. The classes are weighted using the method described in [18] and trained until convergence. The semantic segmentation is used for free space detection as well for classifying the objects found by the instance segmentation branch.

Instance segmentation In order to perform instance segmentation using a typical feed-forward network without having to resort to slower detect-and-segment approaches, we use a recently introduced discriminative loss function [4] suited for real-time instance segmentation that can be plugged into an off-the-shelf network. The intuition behind the proposed loss function is that pixel embeddings (i.e. the network's output for each pixel) with the same label (i.e. same instance) should end up close together, while embeddings with a different label (i.e. different instance) should end up far apart.

Inspired by Weinberger *et al.* [23] and other distance metric learning approaches, the authors propose a loss function with two competing terms to achieve this objective: a

variance term pulling pixel embeddings towards the mean embedding of their cluster, and a *distance term* pushing the clusters away from each other. To relax the constraints on the network, the variance and distance terms are hinged: embeddings within a distance of δ_v from their cluster centers are no longer attracted to it and cluster centers further apart than $2\delta_v$ are no longer repulsed. A small regularization pull-force that draws all clusters towards the origin keeps the activations bounded. These three terms can be written as follows, with C the number of clusters in ground truth, N_c the number of elements in cluster c , x_i an embedding, μ_c the mean embedding of cluster c , $\|\cdot\|$ the L2 distance, and $[x]_+ = \max(0, x)$ denotes the hinge:

$$\begin{cases} L_{var} = \frac{1}{C} \sum_{c=1}^C \frac{1}{N_c} \sum_{i=1}^{N_c} [\|\mu_c - x_i\| - \delta_v]_+^2 \\ L_{dist} = \frac{1}{C(C-1)} \sum_{c_A=1}^C \sum_{c_B=1, c_A \neq c_B}^C [2\delta_d - \|\mu_{c_A} - \mu_{c_B}\|]_+^2 \\ L_{reg} = \frac{1}{C} \sum_{c=1}^C \|\mu_c\| \end{cases} \quad (1)$$

The final loss can then be written as the sum of the above terms: $L_{inst} = L_{var} + L_{dist} + L_{reg}$. When the loss has converged, all pixel embeddings are within a distance of δ_v from their cluster center and all cluster centers are at least $2\delta_d$ apart. By setting $\delta_d > 2\delta_v$, each embedding is closer to all embeddings of its *own* cluster than to any embedding of a *different* cluster. During inference we can then threshold with bandwidth $b = \delta_v$ around any embedding to select all embeddings belonging to the same cluster. Since the loss on the test set will not be zero, we apply a GPU accelerated variant of the mean-shift algorithm [7] to shift to a center pixel around which we threshold, avoiding outliers.

Depth estimation from a single image The standard loss used in most regression problems, like monocular depth estimation, is the L2 loss. It minimizes the difference between predicted D and ground truth D^* depth: $L_2(D, D^*) = \frac{1}{n} \sum_i d_i^2$, with $d = D - D^*$. Recently, Eigen and Fergus [5] added two more terms to the typical L2 loss for the depth estimation task; one for scale invariance ($-\frac{1}{2n^2} (\sum_i d_i)^2$), and another for similarity in local structure ($\frac{1}{n} \sum_i [(\nabla_x d_i)^2 + (\nabla_y d_i)^2]$, with ∇_x and ∇_y denoting the horizontal and vertical image gradients). Instead, the depth estimation branch uses the reverse Huber loss (berHu) [17],

$$L_{dep} = \begin{cases} |d| & |d| \leq c \\ \frac{d^2 + c^2}{2c} & |d| > c, \end{cases} \quad (2)$$

that shows a good balance between penalizing high residuals that usually account for the mean depth and low residuals that explain the smaller depth details. We have experimentally found that this choice yields a better final error than using the L2 loss, even with the added terms. Notice that, the reverse Huber loss formulation above is continuous and first order differentiable at point c , which is set to $c = \frac{1}{5} \max_i(d_i)$ as in [15]. We use the SGM-calculated disparity depth maps of the Cityscapes dataset as ground truth for this task.

TABLE I

SEMANTIC SEGMENTATION RESULTS ON THE CITYSCAPES BENCHMARK.

	IoU class	IoU category
Segnet [2]	56.1	79.8
ENet [18]	58.3	80.4
SQ [21]	59.8	84.3
Ours	59.3	80.4

TABLE II

INSTANCE SEGMENTATION RESULTS ON THE CITYSCAPES BENCHMARK.

	AP	AP0.5	AP100m	AP50m
InstanceCut [12]	23.7	44.8	38.9	42.5
PPLoss	24.4	43.2	40.0	44
Pixelwise DIN [1]	25.7	45.7	41.1	44.2
DWT [3]	31.5	48.5	50.8	53.5
Shape-aware [10]	35.7	54.7	58.2	63.1
SGN	39.4	59.7	60.1	63.2
Mask R-CNN [11]	46.9	68.3	65.5	67.4
Ours	21.0	38.6	34.8	38.7

Training To train our multi-task network, the three losses described above are summed and equally weighted. Although different weights can also be used for each task we found that using equal weights already leads to good performance. We start from a pretrained encoder, trained for Cityscapes segmentation, and continue training the three tasks together. We train with a batch size of 10 at a resolution of 1024x512 and use Adam with a learning rate of 5e-4. Note that, we keep the parameters of the batch norm layers fixed.

IV. RESULTS

Semantic and instance segmentation We report Cityscapes semantic segmentation results in Tab. I and instance segmentation results on the car class in Tab. II. We notice that by jointly training our network for 3 different tasks, we match and even slightly outperform standard ENet for semantic scene segmentation. This justifies our hypothesis that training with multiple tasks at once can increase the performance of each individual task.

As expected, our result for instance segmentation lacks behind the other methods on the Cityscapes benchmark, since they are all optimized for accuracy and are far from real-time. They either rely on a big network or use highly accurate pre-generated semantic segmentation labels, which explains their significantly higher performance, compared to our result. Nevertheless, this work can serve as a baseline for methods that also focus on speed.

Depth In Fig. 2 we plot for each car in the dataset its ground truth depth versus its predicted depth, which is calculated as an average over the predicted depth map masked out with the ground truth instance mask. The expected trend of nearby cars being predicted more accurately than far-away cars is clearly visible. Some of the extreme outliers are caused by cars that are mostly occluded and thus only consist of a few pixels. These extreme cases can in principle be detected and filtered out using the instance mask. We encourage others to include similar plots in their work on car depth estimation, as it is more informative than a single

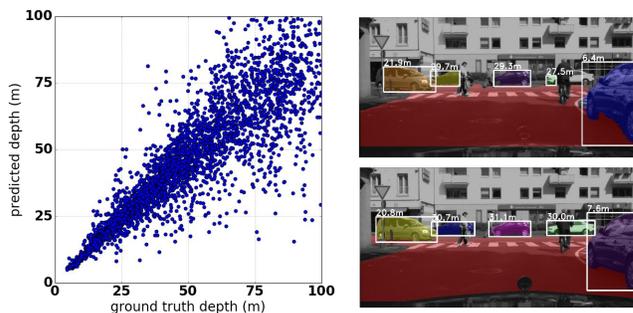


Fig. 2. Left: ground truth and predicted depth for each car in the dataset. The expected trend of nearby cars being predicted more accurately than far-away cars is clearly visible. Right: a qualitative result (top: prediction, bottom: ground truth).

TABLE III

SINGLE IMAGE DEPTH RESULTS ON THE CITYSCAPES VALIDATION SET. SOME CAUTION NECESSARY WHEN COMPARING WITH [22] (SECT. IV).

	MAE	RMSE	ARD
Uhrig <i>et al.</i> [22] (val)	7.7m	24.8m	11.3%
Ours (val, < 100m)	7.5m	11.9m	13.8%
Ours (val, < 50m)	3.5m	5.6m	11.0%
Ours (val, < 25m)	1.5m	2.5m	8.9%

summary number.

Nevertheless, we follow [22] and report three metrics in Tab. III: mean absolute error (MAE), root mean squared error (RMSE) and absolute relative error (ARD). Note that we calculate the depth of each car by average pooling the predicted depth map with the *ground truth* instance masks. This is unlike [22], who calculate the depth with the *predicted* instance masks, and report the metrics only over predicted cars that match with ground truth cars. This means that the metric they report does not take the depth estimation of undetected smaller or badly visible cars into account, leading to a number that is dependent on the instance segmentation performance. By reporting the numbers over the *ground truth* car masks we avoid this entanglement, but some caution is necessary when comparing the numbers. We provide the numbers at different maximum depths of 100m, 50m and 25m.

Multi-task network and speed In Tab. IV we provide a comparison between training the tasks separately (each running on an ENet of their own), versus training them together with a shared encoder as explained in the previous section. The benefits of training the three tasks together in a single multi-task network are clear: the speed almost doubles and the memory usage decreases drastically. This makes our approach suitable for real-time autonomous driving applications that require a low memory footprint. Important to note is that the accuracy of the individual tasks does not decrease when training together: in fact we even notice a slight performance increase. This suggests that the shared encoder can effectively learn to exploit the common structure of the three related semantic tasks.

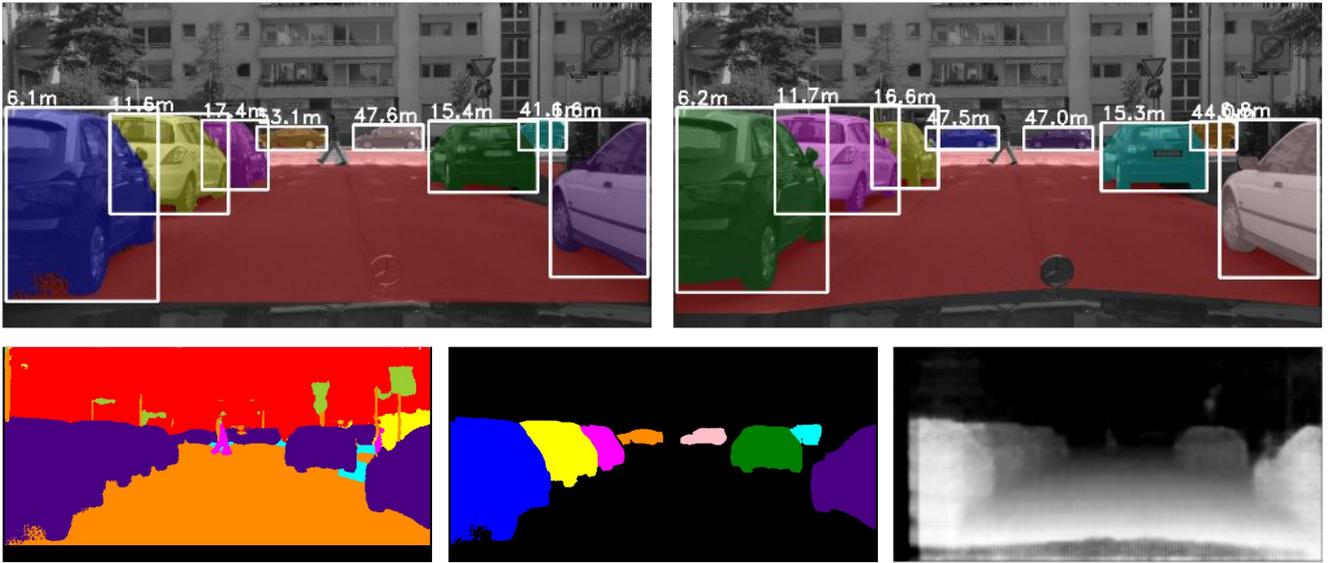


Fig. 3. Another qualitative result on the validation set. Top: predicted and ground truth summary picture. Bottom: predicted semantic, instance and depth maps.

TABLE IV

SEMANTIC SEGMENTATION (IOU_C), INSTANCE SEGMENTATION (AP), DEPTH (MAE_{100M} ON THE VAL SET), MEMORY, AND SPEED (FORWARD PASS ON A PASCAL TITANX) AT TEST TIME WHEN TRAINED SEPARATELY VERSUS TOGETHER.

	semantic	instance	depth	mem	speed
Trained separately	58.3%	0.20%	todo	2.6 GB	12 fps
Trained together	59.3%	0.21%	7.5m	1.2 GB	21 fps

V. CONCLUSION

Overall, our system is fast but lags behind the state-of-art in terms of segmentation accuracy. Nevertheless, we believe that it can serve as a low-complexity baseline for other multi-task approaches that focus on speed, and as a starting point for further exploration of the speed-accuracy trade-off in scene understanding. Furthermore, we observe that jointly training tasks can potentially lead to increased performance.

REFERENCES

- [1] Anurag Arnab and Philip H. S. Torr. Pixelwise instance segmentation with a dynamically instantiated network. In *CVPR*, 2017.
- [2] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *arXiv preprint arXiv:1511.00561*, 2015.
- [3] Min Bai and Raquel Urtasun. Deep watershed transform for instance segmentation. *arXiv preprint arXiv:1611.08303*, 2016.
- [4] Bert De Brabandere, Davy Neven, and Luc Van Gool. Semantic instance segmentation with a discriminative loss function. *manuscript (under submission)*, 2017.
- [5] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, 2015.
- [6] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *NIPS*, 2014.
- [7] Keinosuke Fukunaga and Larry Hostetler. The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Transactions on information theory*, 1975.
- [8] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [9] Bharath Hariharan, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Simultaneous detection and segmentation. In *ECCV*, 2014.
- [10] Zeeshan Hayder, Xuming He, and Mathieu Salzmann. Shape-aware instance segmentation. *arXiv preprint arXiv:1612.03129*, 2016.
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. *arXiv preprint arXiv:1703.06870*, 2017.
- [12] Alexander Kirillov, Evgeny Levinkov, Bjoern Andres, Bogdan Savchynskyy, and Carsten Rother. Instancecut: from edges to instances with multicut. *arXiv preprint arXiv:1611.08272*, 2016.
- [13] Iasonas Kokkinos. Ubernet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. *arXiv preprint arXiv:1609.02132*, 2016.
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [15] Iro Laina, Christian Rupprecht, Vasileios Belagiannis, Federico Tombari, and Nassir Navab. Deeper depth prediction with fully convolutional residual networks. In *3DV*, 2016.
- [16] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [17] Art B Owen. A robust hybrid of lasso and ridge regression. *Contemporary Mathematics*, 2007.
- [18] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. Enet: A deep neural network architecture for real-time semantic segmentation. *arXiv preprint arXiv:1606.02147*, 2016.
- [19] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [20] Marvin Teichmann, Michael Weber, Marius Zoellner, Roberto Cipolla, and Raquel Urtasun. Multinet: Real-time joint semantic reasoning for autonomous driving. *arXiv preprint arXiv:1612.07695*, 2016.
- [21] Michael Treml, José Arjona-Medina, Thomas Unterthiner, Rupesh Durgesh, Felix Friedmann, Peter Schuberth, Andreas Mayr, Martin Heusel, Markus Hofmarcher, Michael Widrich, et al. Speeding up semantic segmentation for autonomous driving.
- [22] Jonas Uhrig, Marius Cordts, Uwe Franke, and T. Brox. Pixel-level encoding and depth layering for instance-level semantic labeling. *GCPR*, 2016.
- [23] Kilian Q Weinberger and Lawrence K Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, 2009.