

Layered Interpretation of Street View Images

Ming-Yu Liu*, Shuoxin Lin†, Srikumar Ramalingam*, and Oncel Tuzel*

*Mitsubishi Electric Research Labs (MERL), Cambridge, Massachusetts, USA

Email: {mliu,ramalingam,oncel}@merl.com

†University of Maryland College Park, Maryland, USA

Email: slin07@umd.edu

Abstract—We propose a layered street view model to encode both depth and semantic information on street view images for autonomous driving. Recently, stixels, stix-mantics, and tiered scene labeling methods have been proposed to model street view images. We propose a 4-layer street view model, a compact representation over the recently proposed stix-mantics model. Our layers encode semantic classes like ground, pedestrians, vehicles, buildings, and sky in addition to the depths. The only input to our algorithm is a pair of stereo images. We use a deep neural network to extract the appearance features for semantic classes. We use a simple and an efficient inference algorithm to jointly estimate both semantic classes and layered depth values. Our method outperforms other competing approaches in Daimler urban scene segmentation dataset. Our algorithm is massively parallelizable, allowing a GPU implementation with a processing speed about 9 fps.

I. INTRODUCTION

Consider a typical road scene as shown in Figure 1 while driving a car. We first observe the immediate road, nearby obstacles (pedestrians, cars, cyclists, etc), followed by adjacent buildings and sky. These scene entities, or objects, can be layered in a typical road scene based on their locations. Understanding such a scene would require us to know the type of objects and spatial locations in the 3D world. Most conventional approaches look at this as two different problems: 3D reconstruction and object class segmentation. Recently, both these problems have been merged and solved as a single optimization problem. Along this avenue, several challenges exist. Prior segmentation algorithms focus on classifying each pixel individually into different semantic object classes. Such approaches are computationally expensive and may not respect the layered constraint that is preserved in most road scenes. In this paper, we jointly infer the semantic labels and depths of road scenes using the layered structure of street scenes.

In Figure 1, we use four layers to represent a street view image. The layers are ordered from the bottom of the image and they are associated with semantic classes. The first layer consists of only the ground. The second layer can have dynamic objects—vehicles, pedestrians, and cyclists. The third layer can only have buildings. The fourth layer can only contain sky. Each of these layers are supposed to model planar objects standing upright with respect to the ground at various distances from the camera. The transition between the layers happen at places where there is depth variation. In most road scenes, four layers are sufficient to model important object classes along with their layered depths. Our approach can fail




Fig. 1. We propose a layered interpretation algorithm for street scene understanding. We use four layers to represent every street scene image. We encode the semantic classes of several important object classes in the region between these layers. These classes include ground, pedestrians, cars, building, and sky. In addition to semantic classes, we also model and compute the layer-aware depth of the scene. In other words, the depth increases for pixels as we move to the top of the image along an image column.

in challenging scenarios with bridges or tunnels. However, these cases can be discovered by survey vehicles in an offline process. Autonomous vehicles can be alerted as we encounter these regions using GPS and we can use additional layers to correctly interpret such challenging regions.

We focus on obtaining layer-aware semantic labels and depths jointly from street-view images. We use a stereo camera setup and compute the disparity cost volume for depth cues. For obtaining semantic cues, we use a deep neural network, which extracts deep features from intensity images. The depth and semantic cues are formulated in an energy function that respects the layered street scene constraint. We propose an inference algorithm based on dynamic programming to efficiently minimize this energy function. Our inference algorithm is massively parallelizable. We develop a parallel implementation and achieve a 8.8 fps processing speed on GPU. Our method outperforms the competing algorithms that do not enforce the layered constraint. Our inference algorithm is general and can work on the data from other modalities including LIDAR and Radar sensors.

A. Related work

Our work is related to the general area of semantic segmentation and scene understanding, such as Carreira et al. [4], Fulkerson et al. [11], Krahenbuhl and Koltun [18], Shotton et al. [28], Ess et al. [6], Geiger et al. [12], Wojek et al. [29], Scharwächter et al. [24]. While earlier approaches were based on hand-designed features, it has been shown recently that using deep neural networks for feature learning leads to

better performance on this task Farabet et al. [8], Girshick et al. [13], Pinheiro and Collobert [22], Sharma et al. [25].

The problem of jointly solving both semantic segmentation and depth estimation from stereo camera was addressed in Ladicky et al. [20] as a unified energy minimization framework. Our work focuses on semantic labeling using ordering constraint on road scenes and using fewer classes applicable to road scenes. In Hoiem et al. [16], a typical road scene is classified into ground, vertical objects and sky to estimate the geometric layout from a single image. Objects like pedestrians and cars are segmented as vertical objects. This would be an under-representation for road scene understanding. Felzenszwalb and Veksler [9] modeled the scene using two horizontal curves that divide the image into three regions: top, middle, and bottom.

One popular model for road scene is the stixel world that simplifies the world using a ground plane and a set of vertical sticks on the ground representing obstacles [2]. Stixels are compact and efficient representation for upright objects on the ground. The stixel representation can simply be seen as the computation of two curves. The first curve runs on the ground plane enclosing the free space that can be immediately reached without collision, and the second curve encodes the vertical objects boundary. In order to compute the stixel world, either depth map from semi-global stereo matching algorithm (SGM) [15] or cost volume [3] can be used. As with SGM, dynamic programming (DP) enables fast implementation for the computation of the stixels. Recently, [?] demonstrated a monocular free-space estimation using appearance cues.

Stix-mantics [24], a recently introduced model, gives more flexibility compared to stixels. Instead of having only one stixel for every column, they allow multiple segments along every column in the image and also combine nearby segments to form superpixel-style entities with better geometric meaning. Using these stixel-inspired superpixels, semantic class labeling is addressed.

We focus on obtaining layer-aware semantic labels and depths jointly from street-view images. Our work is closely related to many existing algorithms in vision, and most notably with tiered scene labeling [9], joint semantic segmentation and depth estimation [20], stixels and more recently, stix-mantics [24]. Our approach achieves real-time processing speed and outperforms the competing algorithms [20] in accuracy. We also achieve this performance without using explicit depth estimation and temporal constraints, which can be obtained using visual odometry. Similar to layered street view constraint, Manhattan constraints have been useful in indoor scene understanding [21, 10].

The paper is organized as follow. In Section II, we present the problem formulation. Section III discusses the extraction of depth and appearance features. The inference algorithm and its implementation are described in Section IV and V. Experiments are presented in Section VI.




Fig. 2. **Illustration of the layered street view problem formulation:** Each column of the image is divided up to four horizontal layers. The four layers are ordered from the bottom of the image. The model is compact and effective in representing a wide variety of scenarios in a typical road scene.

II. LAYERED STREET VIEW

Our goal is to jointly estimate semantic labels and depth for each pixel in the street view image using both appearance and depth information. We adopt a layered image interpretation. An image is horizontally divided into four layers of different semantic and depth regions. The layers are ordered from the bottom to the top. In each image column, pixels belong to the same layer have the same semantic label and depth. The only exception is that depth of the pixels in the ground layer vary according to their vertical image coordinate, which is determined by the ground plane. The ground plane can either be obtained in an offline external calibration process or in an online estimation process such as using the v-disparity map [19]. We enforce a depth-order constraint, i.e., depth of a lower layer is always smaller than depth of a higher layer in each image column.

In our four layer model, the first layer can only have ground. The second layer can have pedestrians or vehicles. The third layer can have only buildings. The fourth layer can only have sky. Note that we do not enforce that each image column has exactly four layers. A column can have any number of layers between one to four. If a layer is absent at a particular image column, the bottom of its upper layer and the top of its bottom layer are next to each other. This implies that the curves defining these layers need not be a smooth continuous one. The four-layer model provides a flexible method for enforcing geometry and semantics to the scene. The only assumption required is the planar world assumption, which is not restrictive for many applications requiring street scene understanding. If necessary, the geometry of the layered street scene model can be further enhanced to dense depth map with additional computational cost. Similarly, the model can be improved with more layers and semantic classes.

A. Problem Formulation

Notations: We use $\mathcal{W} = \{1, 2, \dots, W\}$ and $\mathcal{H} = \{1, 2, \dots, H\}$ to refer the sets that hold the horizontal x and vertical y coordinates respectively. We consider five different semantic object classes; namely, ground, vehicle, pedestrian, building, and sky. They are denoted by the symbols \mathbb{G} , \mathbb{V} ,




Fig. 3. **Layered street view algorithm:** The proposed algorithm utilizes appearance features from a deep neural network and depth features from disparity costs. The features are used to jointly infer the dense semantic and depth labels.

\mathbb{P} , \mathbb{B} , and \mathbb{S} respectively. The set of the semantic class labels is denoted by $\mathcal{L} = \{\mathbb{G}, \mathbb{V}, \mathbb{P}, \mathbb{B}, \mathbb{S}\}$. We use \mathcal{D} for the set of disparity values. The words disparity and depth are used interchangeably in the paper for ease of presentation. It is understood that a one-to-one conversion can be easily obtained by using the parameters in the camera calibration matrix. The cardinality of the semantic label space and disparity values are denoted by $L = |\mathcal{L}|$ and $D = |\mathcal{D}|$ respectively.

We formulate the layered street view problem as a constrained energy minimization problem. The constraints encode the order of the semantic object class labels and depth values in each column. It limits the solution space of the variables associated with each image column. We solve the constrained energy minimization problem efficiently using an inference algorithm based on dynamic programming.

We use the variables, h_{i1} , h_{i2} , h_{i3} , and h_{i4} , to denote the y coordinates of the top pixels of the four layers in the image column i . Let l_{i1} , l_{i2} , l_{i3} , and l_{i4} be the semantic object class labels for the four layers and let d_{i1} , d_{i2} , d_{i3} , and d_{i4} be the depths of the four layers in the image column i . The ordering constraint and the knowledge of the ground plane allow us to fix some parameters. The actual number of unknowns is only 5 given by $\mathbf{x}_i = [h_{i1}, h_{i2}, h_{i3}, l_{i2}, d_{i3}]$. Hence the label assignment for the entire image is given by $\mathbf{x}_I = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_W]$. The number of possible assignments for an image column is in the order $O(H^3LD)$ since $h_{i1}, h_{i2}, h_{i3} \in \mathcal{H}$, $l_{i2} \in \mathcal{L}$, and $d_{i3} \in \mathcal{D}$.

To rank the likelihood of the label assignment, we use evidence from image appearance features and stereo disparity matching features. We aggregate evidence from all the pixels in a column to compute the evidence. Let $U_i^A(\mathbf{x}_i)$ and $U_i^D(\mathbf{x}_i)$ be the data terms representing the semantic and depth label cost, respectively, incurred as assigning \mathbf{x}_i to the image column i . The two terms are summed to yield the data term

$$U_i(\mathbf{x}_i) = U_i^A(\mathbf{x}_i) + U_i^D(\mathbf{x}_i), \quad (1)$$

denoting the cost for assigning \mathbf{x}_i to the image column i . Instead of working on the standard 2D Markov Random Field space where each pixel can have a depth value and a semantic label as independent variables, we reduce the problem to a

constrained energy minimization problem given by

$$\min_{\mathbf{x}} \sum_{i=1}^W U_i(\mathbf{x}_i) \quad (2)$$

$$\text{s.t. } h_{i1} \geq h_{i2} \geq h_{i3} \geq h_{i4} = 1, \quad (3)$$

$$d_{i1} < d_{i2} < d_{i3} < d_{i4}, \quad (4)$$

$$l_{i1} = \mathbb{G}, l_{i2} \in \{\mathbb{V}, \mathbb{P}\}, l_{i3} = \mathbb{B}, l_{i4} = \mathbb{S}, \forall i \quad (5)$$

where the constraint (3) gives the layer structure, the constraint (4) enforces the depth order, and the constraint (5) takes into account the possible semantic labels for each layer. The variable d_{i2} is a function of h_{i1} , the top pixel location of the ground layer, because we assume the dynamic object is standing upright on the ground surface at the h_{i1} th row of the image. The energy function in Equation (2) models the relation of pixels in the same column but not pixels in the same row. We use image patches centering around a pixel as the reception fields for the feature computation at the pixel location. Neighboring pixels have similar reception fields and thus have similar features.

The data term $U_i(\mathbf{x}_i)$ is the cost of assigning label \mathbf{x}_i to the column i . It is the sum of the pixel-wise data terms given by

$$\begin{aligned} U_i(\mathbf{x}_i) = & \sum_{y=h_{i4}}^{h_{i3}-1} (E^A(i, y, \mathbb{S}) + E^D(i, y, \infty)) \\ & + \sum_{y=h_{i3}}^{h_{i2}-1} (E^A(i, y, \mathbb{B}) + E^D(i, y, d_{i3})) \\ & + \sum_{y=h_{i2}}^{h_{i1}-1} (E^A(i, y, l_{i2}) + E^D(i, y, d_{i2}(h_{i1}))) \\ & + \sum_{y=h_{i1}}^H (E^A(i, y, \mathbb{G}) + E^D(i, y, d_{i1})) \end{aligned} \quad (6)$$

where the per pixel appearance data term $E^A(x, y, l)$ is the cost of assigning label l to the pixel (x, y) and the per pixel depth data term $E^D(x, y, d)$ is the cost of assigning depth d to the pixel (x, y) . We use a deep neural network for obtaining the per pixel appearance data term $E^A(x, y, l)$ and use a standard disparity cost for obtaining the per pixel depth data




Fig. 4. **Deep neural network for extracting the appearance features:** Our deep neural network contains two parts: multi-scale convolutional neural network and recursive context propagation network. The multi-scale convolutional neural network contains three convolutional layers and is applied to the Gaussian pyramid of the input image. It extracts multi-scale appearance features, which are fed into the recursive context propagation network. The recursive context propagation network consists of three sub-networks: the semantic mapper, the semantic combiner, and the semantic de-combiner networks. It embeds rich context information to the features and enhances their discriminative power. Note that the layers are color-coded. The same color is assigned to the layers sharing the filter weights. A dotted line indicates that the feature map is upsampled before feeding to the succeeding layer. For further details, please refer to Section III.

term $E^D(x, y, d)$ detailed in Section III. We summarize our layered street view algorithm in Figure 3.

III. FEATURES

We rely on two types of features: depth and appearance.

A. Depth

We use the smoothed absolute intensity difference for the per pixel depth data term, which is commonly used in stereo reconstruction algorithms. We first compute the pixel wise absolute intensity difference for each disparity value in D , which renders a cost volume representation. A box filter is then applied to smooth the cost volume. The per pixel depth data term is given by

$$E^D(x, y, d) = \frac{1}{N} \sum_{(\tilde{x}, \tilde{y}) \in P(x, y)} |I_L(\tilde{x}, \tilde{y} - d) - I_R(\tilde{x}, \tilde{y})| \quad (7)$$

where I_L and I_R refer to the intensity values of the left and right images, $P(x, y)$ is an image patch centered at (x, y) , and $N = |P(x, y)|$ denotes the cardinality of the patch, serving as a normalization constant. The patch size is fixed to 11-by-11 in our experiments.

B. Appearance

We compute per pixel appearance data term using a deep neural network. Our network, shown in Figure 4, consists of two parts: a multi-scale convolutional neural network (MSCNN) and a recursive context propagation network (RCPN). The MSCNN allows us to extract multi-scale appearance cues and RCPN allows us to extract rich contextual information.

MSCNN: Our MSCNN is a close variant to the neural network proposed in Farabet et al. [8]. It has 3 convolutional layers. The first convolutional layer has 16 filters of size 8×8 followed by rectified linear unit (ReLU) and 2×2 max-pooling

processing. The second layer has 64 filters of size 7×7 followed by ReLU and 2×2 max-pooling processing. The third layer has 256 filters of size 7×7 followed by ReLU. The stacking of the convolutional layers yields a reception field of 47-by-47 pixels. Due to max-pooling, the resolution of the output feature map is smaller than that of the input image. We upsample the feature map to have the same resolution.

The 3-layer convolutional neural network (CNN) is applied separately to three scales of the Gaussian pyramid of the input image. Specifically, we downsample the input image using three different scales (1, 2, and 4) and use the 3-layer CNN to extract features at each scale. We use upsampling to bring all the feature maps to the same resolution. The feature maps from the three scales are concatenated to obtain the MSCNN feature map. Note that the filter weights are constrained to be the same for all the scales. The MSCNN extracts multi-scale appearance information for each pixel, which is then passed to the RCPN. For further details about MSCNN, please refer to [8].

RCPN: Our RCPN is a variant of the network proposed in Sharma et al. [25]. It consists of three sub-networks: the semantic mapper network, the semantic combiner network, and the semantic de-combiner network. We use the RCPN to embed rich context information to the output appearance features. The semantic mapper network is a 1-layer CNN with 128 filters of size 1×1 followed by ReLU. It maps each 768-dimensional feature of the MSCNN feature map to a 128-dimensional semantic feature, which is then fed into the semantic combiner network.

The semantic combiner network has three recursive layers, and each contains 128 filter of size 4×4 followed by ReLU. The semantic combiner fuses input features in a 4-by-4 region of the input semantic feature map to an output semantic feature. This process is non-overlapping; hence, each semantic combiner layer generates an output feature map that is 16 times

smaller than the input one. Applying the 3-layer recursive combiner network renders an output feature map that is 4096 times smaller than the input feature map of the semantic mapper. The semantic combiner network recursively embeds context information from image regions with larger and larger spatial support. The output feature maps from the three layers form a context feature pyramid, which is fed into the semantic de-combiner network.

Similar to the semantic combiner network, the semantic de-combiner network has three recursive layers and each contains 128 filter of size 1×1 followed by ReLU. It is used to recursively distribute context information residing in the context pyramid back to the individual pixels, from higher to lower levels of the context pyramid. Each de-combiner layer fuses the feature map from the previous de-combiner layer with that from the corresponding level of the context pyramid. Note that our RCPN implementation differs from [25] in two major places: 1) we use square patches for context propagation while [25] uses superpixels[?], and 2) we use pyramids to represent hierarchy of context information while [25] uses superpixel trees. Our design choices allow a more efficient implementation because we do not use superpixel segmentation.

Training: We use grayscale images. The pixel intensity values are scaled between 0 to 1 and centered by subtracting 0.5 before being fed into the deep neural network. To train the network, we connect the output layer to a fully connected layer having 5 neurons, corresponding to ground, pedestrian, vehicle, building, and sky classes. The fully connected layer is followed by the softmax layer. The network is trained by minimizing the cross-entropy error via stochastic gradient descent with momentum. We use the Caffe library Jia et al. [17] for training. The number of pixels in the semantic classes can be quite different. To avoid the bias from dominant classes (ground and building), we weight the cross-entropy loss based on the semantic class distribution, which yields better performance in practice.

We use the negative logarithm of the softmax scores of the semantic classes as the per-pixel appearance data terms. Let $f(x, y, l)$ be the softmax score of the deep neural network at pixel location (x, y) , which represents the probability of assigning the label l to the pixel. The per pixel appearance data term is given by

$$E^A(x, y, l) = -\beta \log f(x, y, l) \quad (8)$$

where β is a parameter controlling the relative weight of the appearance and depth data terms.

IV. EFFICIENT INFERENCE ALGORITHM

We decompose the energy minimization problem in Equation (2) into W sub-problems where the i th sub-problem is

given by

$$\min_{\mathbf{x}_i} U_i(\mathbf{x}_i) \quad (9)$$

$$\text{s.t. } h_{i1} \geq h_{i2} \geq h_{i3} \geq h_{i4} = 1 \quad (10)$$

$$d_{i1} < d_{i2} < d_{i3} < d_{i4} \quad (11)$$

$$l_{i1} = \mathbb{G}, l_{i2} \in \{\mathbb{V}, \mathbb{P}\}, l_{i3} = \mathbb{B}, l_{i4} = \mathbb{S}. \quad (12)$$

We solve each of the sub-problems optimally and combine their solutions to construct the semantic labeling and depth map of the image. For simplicity, we will drop the subscript i in the discussion below.

Each of the sub-problems can be mapped to a 1D chain labeling problem. The chain has 4 nodes where the first node contains the variables (h_1, d_1, l_1) , the second node contains the variables (h_2, d_2, l_2) , the third node contains the variables (h_3, d_3, l_3) , and the fourth node contain the variables (h_4, d_4, l_4) . Utilizing the recursion in the label cost evaluation, a standard dynamic programming algorithm can solve the inference on the 1D chain with a complexity of $O((HDL \cdot H)^2)$ where the product HDL represents the size of label space at each node and the second H comes from the label cost evaluation at each node. Unfortunately, the complexity is too high for real-time applications.

We propose a variant of the dynamic programming algorithm to reduce the complexity of solving the sub-problem in (9) to $O(H^2D)$ and achieve real-time performance. We first note that some of the variables are known from our street view setup as discussed in the problem formulation section. We only need to search the values for h_1, h_2, h_3, l_2, d_3 . For any combination of h_1, h_2 and l_2 , we need to find the best combination of d_3 and h_3 . In the following, we show that pre-computing the best combination of d_3 and h_3 for any h_1, h_2 and l_2 can be achieved in $O(H^2D)$ time using recursion.

We first observe that the problem in (9) can then be written as

$$\begin{aligned} \min_{h_1, h_2, l_2} & \sum_{y=h_2}^{h_1-1} (E^A(i, y, l_2) + E^D(i, y, d_2(h_1))) + \\ & \sum_{y=h_1}^H (E^A(i, y, \mathbb{G}) + E^D(i, y, d_1) + Q_i(h_1, h_2)) \end{aligned} \quad (13)$$

where Q_i is an intermediate cost table given by

$$\begin{aligned} Q_i(h_1, h_2) \equiv & \min_{\substack{h_3, d_3 \\ d_3 > d_2(h_1)}} \sum_{y=h_4}^{h_3-1} (E^A(i, y, \mathbb{S}) \\ & + E^D(i, y, \infty)) + \sum_{y=h_3}^{h_2-1} (E^A(i, y, \mathbb{B}) + E^D(i, y, d_3)) \end{aligned} \quad (14)$$

Note that depth of the second layer object d_2 is a function of h_1 because d_2 can be uniquely determined from h_1 and the ground plane equation. As a result, Q_i depends on both h_1 and h_2 .

By integrating $E^A(i, y, \mathbb{B}) + E^D(i, y, d_3)$ and $E^A(i, y, \mathbb{S}) + E^D(i, y, \infty)$ along the y direction, the sum given by

$$R_i(h_2, h_3, d_3) \equiv \sum_{y=h_4}^{h_3-1} (E^A(i, y, \mathbb{S}) + E^D(i, y, \infty)) + \sum_{y=h_3}^{h_2-1} (E^A(i, y, \mathbb{B}) + E^D(i, y, d_3)) \quad (15)$$

for all combination of h_2 and h_3 can be computed in $O(H^2)$ time for each d_3 . We further note that Q_i can be computed via a recursive update rule given by

$$Q_i(h_1 + 1, h_2) = \min(Q_i(h_1, h_2), \min_{h_3} R_i(h_2, h_3, \tilde{d})) \quad (16)$$

where \tilde{d} is an integer satisfying $d_2(h_1) \leq \tilde{d} < d_2(h_1 + 1)$ and is used to ensure that the depth ordering constraint between the second and third layers is met. Intuitively, we are computing a running min structure along the decreasing depth of the building layer. The recursive update rule allows us to compute Q_i for any h_1 and h_2 in $O(H^2D)$ time. As a result, the complexity of finding the best configuration for a partition can be reduced to $O(H^2L + H^2D) = O(H^2D)$ where H^2L is the time required for searching combinations of h_1 , h_2 and l_2 . We perform the 1D labeling algorithm to each image column. The overall complexity of the labeling algorithm is $O(WH^2D)$.

V. IMPLEMENTATION

Our algorithm is massively parrallelizable and can be implemented using CUDA, a general purpose parallel computing language for NVIDIA GPUs [5]. A GPU comprises a large number of Single-Instruction-Multiple-Data (SIMD) processor cores to allow many threads to execute common operations concurrently on large data arrays. In our implementation of the labeling algorithm, we exploit data level parallelism in all stages of computation.

- 1 Depth data term:** We use $W \times H$ threads to compute the disparity values for each pixel at (x, y) . We implemented the box filter using the sliding window approach. First, we use H threads to perform 1D sliding window on each row. As the window moves from left to right in the horizontal direction, the new pixel value on the right is added and the existing one on the left is subtracted. We then use W threads to carry out the same computation for each column.
- 2 Appearance data term:** We use the Caffe library [17] to compute the softmax scores output of the deep neural network. We use NVIDIA’s cuDNN library to achieve additional speed up.
- 3 Intermediate cost table:** We first compute the integral of $E^A(i, y, \mathbb{B}) + E^D(i, y, d_3)$ and $E^A(i, y, \mathbb{S}) + E^D(i, y, \infty)$ over y such that the two sum terms can be retrieved in constant time for any range. This can be done in parallel for each column i . We observe that for a fixed h_2 , computing $R_i(h_2, h_3, d_3)$ over each i and h_3 can be jointly parallelized. Therefore we use $h_2 \times W$ threads to

compute an intermediate table for d_3 and h_3 and use W threads to find the combinations that yield the minimum cost for each h_1 and h_2 .

- 4 Energy minimization and labeling:** Using previously computed Q , we use W threads to search the x_i with the minimum cost in each column in parallel.

Memory layout is an important factor for processing speed in GPU. By default, our image data is stored in row-major form. The GPU implementation naturally takes advantage of this memory layout by assigning threads to work on pixels on the same row, which resides in memory as a continuous array. This allows GPU to coalesce the memory accesses of the threads such that the GPU memory bandwidth is efficiently utilized. In addition, our algorithm avoids reshaping or transposing the data in the memory, which would take extra memory and time.

We execute our algorithm on a Windows 7 desktop computer equipped with NVIDIA Tesla K40 GPU along with Intel i7 processor. We set the size of one-dimension thread blocks to be 64 and the size of two-dimension thread blocks to be 32×32 to facilitate efficient scheduling of the threads on target GPU. To avoid register spills to local memory, we minimize local variable declarations. In our algorithm, no data needs to be shared among threads within a block, therefore shared memory is not used in the implementation.

VI. EXPERIMENTS

Benchmark: We evaluated our approach using the public Daimler Urban Segmentation dataset [23]. The dataset contains 500 stereo grayscale image pairs with pixel-wise semantic class annotations for the left images. While the image size in the dataset is 1024×440 , only the middle region from $(24, 40)$ to $(1000, 400)$ is fully labeled. Hence, the effective image size is 976×360 . The dataset is composed for evaluating only the semantic labeling using stereo image pairs. There is no ground truth for depth.

The semantic labels in the annotations include ground, sky, building, pedestrian, vehicle, curbs, bicyclist, motorcyclist, and background clutters. However, only the ground, sky, building, pedestrian, and vehicle are considered in the evaluation protocol. The performance metric for the semantic labeling is based on the PASCAL VOC intersection over union (IoU) measure [7], which is the ratio of cardinality of the intersection of the ground truth and estimated semantic segments over that of their union. Let \tilde{S} and S be the set of pixels labeled as sky in the computed semantic class label map and the ground truth label map, respectively. The IoU measure of the sky class is given by

$$IoU(\text{Sky}) = \frac{|\tilde{S} \cap S|}{|\tilde{S} \cup S|}. \quad (17)$$

The larger the IoU measure, the better the matching between the ground truth and estimated segments; and, hence, the better the semantic labeling accuracy.

Evaluation: We followed the evaluation protocol described in [24], which used the first 300 stereo image pairs in the

TABLE I

THE TABLE COMPARES THE PROPOSED METHOD WITH SEVERAL COMPETING ALGORITHMS ON SEMANTIC OBJECT CLASS LABELING ACCURACY USING THE DAIMLER URBAN SEGMENTATION DATASET. THE RESULTS ARE SHOWN IN PERCENTAGE USING THE IOU MEASURE.

Class\Method	Joint-Opt. ALE [20]	Stix-mantics [24]	Darwin-pairwise [14]	PN-RCPN [26]	Appearance features	Proposed
Ground	94.9	93.8	95.7	96.7	96.7	96.4
Vehicle	76.0	78.8	68.7	79.4	80.7	83.3
Pedestrian	73.1	66.0	21.2	68.4	61.3	71.1
Sky	95.5	75.4	94.2	91.4	87.6	89.5
Building	90.6	89.2	87.6	86.3	87.4	91.2
Avg (all)	86.0	80.6	73.5	84.5	82.8	86.3
Avg (dynamic)	74.5	72.4	44.9	73.8	71.0	77.2
Runtime per frame in second	111	0.05	N/A	2.8	0.07	0.11




Fig. 5. **Visualization:** The figure visualizes the output computed from the proposed method. From top to bottom, we show the left images, the right images, the ground truth semantic labeling, the semantic labeling, and depth. The black regions are the regions where the ground truths are not available.

dataset for training and the remaining 200 stereo image pairs for testing. During testing, we downsampled the input images by half in each dimension. This was necessary for our GPU implementation. During evaluation, we upsampled the image to the original size.

We use left images of the stereo pairs for training our deep neural network. During testing, the network outputs per-semantic-class softmax scores for each pixel. We compared our approach with several approaches. The competing approaches include the joint-optimal ALE (ALE) algorithm [20], the stix-mantics [24], the Darwin pairwise [14], and the PN-RCPN [26]. The algorithms in [20, 26] utilizes superpixel segmentations as input, which demands additional computation resources. The stix-mantics algorithm [24] uses depth, obtained using an FPGA chip, and temporal constraints from adjacent stereo images. We achieve better accuracy and comparable computational performance without using an FPGA chip and temporal constraints.

In Table I, we compare the semantic labeling accuracy of the

competing algorithms. The results of the competing algorithms are duplicated from the Daimler dataset website [1]. Note that the results in the website are different from those reported in the original paper [24] because the unlabeled pixels were initially not excluded from the IoU computation in the original paper [24].

Performance: In the table, we show that our method achieves the state-of-the-art performance of 86.3%, while ALE achieves an accuracy of 86.0%. The stix-mantics algorithm achieves an accuracy of 80.6%, which is significantly lower than our method and ALE. In terms of the performance for the dynamic objects (vehicles and pedestrians), the proposed algorithm achieves an accuracy of 77.2%, outperforming ALE, which gets 74.5%. In terms of speed, we are several magnitude faster than ALE. We generate both depth and semantic labels in 114.1 ms, while ALE takes 111,000 ms. The stix-mantics algorithm is slightly faster than our method requiring only 50 ms. However, they use an FPGA chip to precompute the depth before estimating the semantic labels, whereas we jointly




Fig. 6. **The advantages of the layered constraint:** Top: Semantic labeling using only the appearance features. Bottom: Semantic labeling using the layered constraint. This constraint avoids geometrically impossible labeling, which may occur while using only appearance features extracted using the deep neural network.

TABLE II
ALGORITHM COMPONENT EXECUTION TIME

Component name	Execution time in ms
Depth data term (cost volume)	5.2
Appearance data term (DNN)	70.0
Intermediate table Q	25.6
Inference	13.3
Overall	114.1

compute both semantic labels and depth.

In Figure 5, we visualize the semantic labels and depth from the proposed algorithm. Qualitatively, we find that we obtain visually accurate semantic labels and our depth map resembles a piece-wise planar approximation of the 3D scene.

We observe that our appearance features alone performs quite well. It achieves an average accuracy of 82.8%. The layered constraint allows us to avoid impossible labelings such as having ground regions in the middle of sky, or having vehicles in the middle of a building, etc. By incorporating this constraint, our performance improves to 86.3%, which corresponds to a 20.3% error reduction. This improvement can be qualitatively seen in a few examples in Figure 6.

We report the execution time required by the individual steps in our algorithm in Table II. All the computation is performed in the GPU. Overall, our algorithm takes 114.1 ms to infer the semantic labels and depth. The run time can be further reduced by half, by utilizing a second GPU card for processing the neural network computation.

VII. CONCLUSION

We propose a novel layered street view model and develop an efficient algorithm to jointly estimate semantic labels and depth for street view images. We obtain this result using appearance features, which can be computed from a deep neural network, and depth features, which can be derived from stereo disparity costs. Our algorithm outperforms the competing methods on the Daimler Urban Segmentation data set and runs at 8.8 frame-per-second using a GPU implementation.

ACKNOWLEDGEMENT

We thank Jay Thornton, Shumpei Kameyama, reviewers, and area chairs for their feedback and support of this work.

REFERENCES

- [1] Scene labeling, 2015. URL <http://www.6d-vision.com/scene-labeling>.
- [2] H. Badino, U. Franke, and D. Pfeiffer. The stixel world - a compact medium level representation of the 3d-world. In *DAGM*, 2009.
- [3] R. Benenson, R. Timofte, and L.V. Gool. Stixels estimation without depth map computation. In *ICCV*, 2011.
- [4] João Carreira, Rui Caseiro, Jorge Batista, and Cristian Sminchisescu. Semantic segmentation with second-order pooling. In *European Conference on Computer Vision*, 2012.
- [5] NVIDIA Corporation. NVIDIA CUDA C programming guide. 2014.
- [6] A. Ess, T. Mueller, H. Grabner, and L.V. Gool. Segmentation-based urban traffic scene understanding. In *British Machine Vision Conference*, 2009.
- [7] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [8] Clément Farabet, Camille Couprie, Laurent Najman, and Yann LeCun. Learning hierarchical features for scene labeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1915–1929, 2013.
- [9] Pedro F Felzenszwalb and Olga Veksler. Tiered scene labeling with dynamic programming. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3097–3104, 2010.
- [10] Alex Flint, David Murray, and Ian Reid. Manhattan scene understanding using monocular, stereo, and 3d features. In *IEEE International Conference on Computer Vision*.
- [11] B. Fulkerson, A. Vedaldi, and S. Soatto. Class segmentation and object localization with superpixel neighborhoods. In *IEEE International Conference on Computer Vision*, 2009.
- [12] A. Geiger, M. Lauer, C. Wojek, C. Stiller, and R. Urtasun. 3d traffic scene understanding from movable platforms. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2013.
- [13] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. *arXiv preprint arXiv:1311.2524*, 2013.
- [14] Stephen Gould. Darwin: A framework for machine learning and computer vision research and development. *The Journal of Machine Learning Research*, 13(1):3533–3537, 2012.
- [15] H. Hirschmuller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2008.
- [16] D. Hoiem, A. A. Efros, and M. Hebert. Automatic photo pop-up. *ACM Transactions on Graphics*, 2005.
- [17] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadar-

- rama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [18] Philipp Krahenbuhl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Neural Information Processing Systems*, 2011.
 - [19] Raphael Labayrade, Didier Aubert, and J-P Tarel. Real time obstacle detection in stereovision on non flat road geometry through” v-disparity” representation. In *Intelligent Vehicle Symposium*, volume 2, pages 646–651. IEEE, 2002.
 - [20] Lubor Ladicky, Paul Sturgess, Chris Russell, Sunando Sengupta, Yalin Bastanlar, William Clocksin, Philip HS Torr, et al. Joint optimisation for object class segmentation and dense stereo reconstruction. In *British Machine Vision Conference*, 2010.
 - [21] David C Lee, Martial Hebert, and Takeo Kanade. Geometric reasoning for single image structure recovery. In *IEEE Conference on Computer Vision and Pattern Recognition*.
 - [22] Pedro Pinheiro and Ronan Collobert. Rich feature hierarchies for accurate object detection and semantic segmentation. *International Conference on Machine Learning*, 2014.
 - [23] Timo Scharwächter, MarkusENZweiler, Uwe Franke, and Stefan Roth. Efficient multi-cue scene segmentation. In *Pattern Recognition*, pages 435–445. Springer, 2013.
 - [24] Timo Scharwächter, MarkusENZweiler, Uwe Franke, and Stefan Roth. Stixmantics: A medium-level model for real-time semantic scene understanding. In *European Conference on Computer Vision*, pages 533–548. Springer, 2014.
 - [25] Abhishek Sharma, Oncel Tuzel, and Ming-Yu Liu. Recursive context propagation network for semantic scene labeling. In *Neural Information Processing Systems*, pages 2447–2455, 2014.
 - [26] Abhishek Sharma, Oncel Tuzel, and David W Jacobs. Deep hierarchical parsing for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2015.
 - [27] J. Shotton, J. Winn, C. Rother, and A. Criminisi. Textonboost: Joint appearance, shape and context modeling for multi-class object recognition and segmentation. In *European Conference on Computer Vision*, 2006.
 - [28] J. Shotton, J. Winn, C. Rother, and A. Criminisi. Textonboost for image understanding: multi-class object recognition and segmentation by jointly modeling texture, layout, and context. In *International Journal of Computer Vision*, 2009.
 - [29] C. Wojek, S. Walk, S. Roth, K. Schindler, and B. Schiele. Monocular visual scene understanding: Understanding multi-object traffic scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2013.