

Geometric Scene Parsing with Hierarchical LSTM

Zhanglin Peng¹, Ruimao Zhang¹, Xiaodan Liang¹, Xiaobai Liu², Liang Lin^{1*}

¹Sun Yat-sen University, Guangzhou, China

²San Diego State University, U.S.

Abstract

This paper addresses the problem of geometric scene parsing, i.e. simultaneously labeling geometric surfaces (e.g. sky, ground and vertical plane) and determining the interaction relations (e.g. layering, supporting, siding and affinity) between main regions. This problem is more challenging than the traditional semantic scene labeling, as recovering geometric structures necessarily requires the rich and diverse contextual information. To achieve these goals, we propose a novel recurrent neural network model, named Hierarchical Long Short-Term Memory (H-LSTM). It contains two coupled sub-networks: the Pixel LSTM (P-LSTM) and the Multi-scale Super-pixel LSTM (MS-LSTM) for handling the surface labeling and relation prediction, respectively. The two sub-networks provide complementary information to each other to exploit hierarchical scene contexts, and they are jointly optimized for boosting the performance. Our extensive experiments show that our model is capable of parsing scene geometric structures and outperforming several state-of-the-art methods by large margins. In addition, we show promising 3D reconstruction results from the still images based on the geometric parsing.

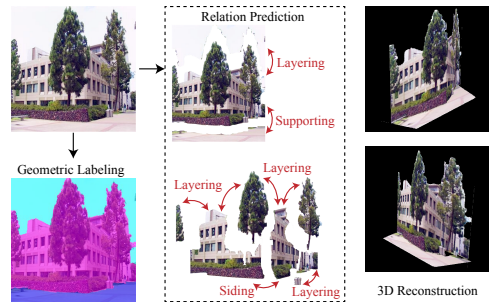


Figure 1: An illustration of our geometric scene parsing. Our task aims to predict the pixel-wise geometric surface labeling (first column) and the interaction relations between main regions (second column). Then the parsing result is applied to reconstruct a 3D model (third column).

1 Introduction

Humans can naturally sense the geometric structures of a scene by a single glance, while developing such a system remains to be quite challenging in several intelligent applications such as robotics [Kanji, 2015] and automatic navigation [Nieuwenhuisen *et al.*, 2010]. In this work, we investigate a novel learning-based approach for geometric scene parsing, which is capable of simultaneously labeling geometric surfaces (e.g. sky, ground and vertical) and determines

the interaction relations (e.g. layering, supporting, siding and affinity [Liu *et al.*, 2014]) between main regions, and further demonstrate its effectiveness in 3D reconstruction from a single scene image. An example generated by our approach is presented in Figure 1. In the literature of scene understanding, most of the efforts are dedicated for pixel-wise semantic labeling / segmentation [Long *et al.*, 2015][Pineiro and Collobert, 2015]. Although impressive progresses have been made, especially by the deep neural networks, these methods may have limitations on handling the geometric scene parsing due to the following challenges.

- The geometric regions in a scene often have diverse appearances and spatial configurations, e.g. the vertical plane may include trees and buildings of different looks. Labeling these regions generally requires fully exploiting image cues from different aspects ranging from local to global.
- In addition to region labeling, discovering the interaction relations between the main regions is crucial for recovering the scene structure in depth. The main difficulties for the relation prediction lie in the ambiguity of multi-scale region grouping and the fusion of hierarchical contextual information.

To address these above issues, we develop a novel **Hierarchical LSTM** (H-LSTM) recurrent network that simultaneously parses a still image into a series of geometric regions

*Corresponding author is Liang Lin (Email: linliang@ieee.org). This work was supported in part by Special Program for Applied Research on Super Computation of the NSFC-Guangdong Joint Fund (the second phase), in part by Guangdong Natural Science Foundation under Grant S2013050014548, in part by the Fundamental Research Funds for the Central Universities.

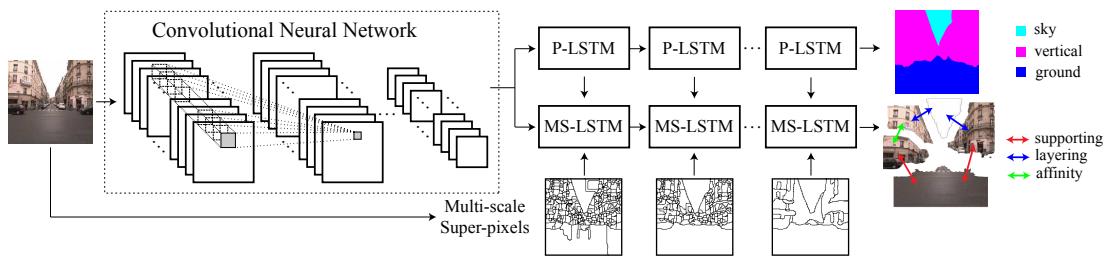


Figure 2: The proposed recurrent framework for geometric scene parsing. Each still image is first fed into several convolutional layers. Then these feature maps are passed into the the stacked Pixel LSTM (P-LSTM) layers and Multi-scale Super-pixel LSTM (MS-LSTM) to generate the geometric surface labeling of each pixel and interaction relations between regions, respectively.

and predicts the interaction relations among these regions. The parsing results can be directly used to reconstruct the 3D structure from a single image. As shown in Figure 2, the proposed model collaboratively integrates the Pixel LSTM (P-LSTM) [Liang *et al.*, 2015] and Multi-scale Super-pixel LSTM (MS-LSTM) sub-networks into a unified framework. First, the P-LSTM sub-network produces the geometric surface regions, where local contextual information from neighboring positions is imposed on each pixel to better exploit the spatial dependencies. Second, the MS-LSTM sub-network generates the interaction relations for all adjacent surface regions based on the multi-scale super-pixel representations. Benefiting from the diverse levels of information captured by hierarchical representations (i.e. pixels and multi-scale super-pixels), the proposed H-LSTM can jointly optimize the two tasks based on the hierarchical information, where different levels of contexts are captured for better reasoning in local area. Based on the shared basic convolutional layers, the parameters in P-LSTM and MS-LSTM sub-networks are jointly updated during the back-propagation. Therefore, the pixel-wise geometric surface prediction and the super-pixel-wise relation categorization can mutually benefit from each other.

The proposed H-LSTM is primarily inspired by the success of Long Short-Term Memory Networks (LSTM) [Graves *et al.*, 2007][Kalchbrenner *et al.*, 2015] on the effective incorporation of long and short range dependencies from the whole image. Different from previous LSTM structure [Byeon *et al.*, 2014][Byeon *et al.*, 2015] that simply operates on each pixel, our H-LSTM exploits hierarchical information dependencies from different levels of units such as pixels and multi-scale super-pixels. The hidden cells are treated as the enhanced features and the memory cells can recurrently remember all previous contextual interactions for different levels of representations from different layers.

Since the geometric surface labeling needs the fine prediction results while the relation prediction cares more about the coarse semantic layouts, we thus resort to the specialized P-LSTM and MS-LSTM to separately address these two tasks. In terms of **geometric surface labeling**, the P-LSTM is used to incorporate the information from neighboring pixels to guide the local prediction of each pixel, where the local contextual information can be selectively remembered and then guide the feature extraction in the later layer. In terms of

interaction relation prediction, the MS-LSTM effectively reduces the information redundancy by the natural smoothed regions and different levels of information can be hierarchically used to extract interaction relations in different layers. Particularly, in each MS-LSTM layer, the super-pixel map with a specific scale is used to extract the smoothed feature representation. Then, the features of adjacent super-pixels are fed into the LSTM units to exploit the spatial dependencies. The super-pixel map with larger scale is used in the deep layer to extract the higher-level contextual dependencies. After passing through all of the hierarchical MS-LSTM layers, the final interaction relation prediction can be obtained by the final relation classifier based on the enhanced features benefiting from the hierarchical LSTM units.

This paper makes the following contributions. (1) A novel recurrent neural network model is proposed for geometric scene parsing, which jointly optimizes the geometric surface labeling and relation prediction. (2) Hierarchically modeling image contexts with LSTM units over super-pixels is original to the literature, which can be extended to similar tasks such as human parsing. (3) Extensive experiments on three public benchmarks demonstrate the superiority of H-LSTM over other state-of-the-art geometric surface labeling approaches. Moreover, we show promising 3D reconstruction results from the still images based on the geometric parsing.

2 Related Work

Semantic Scene Labeling. Most of the existing works focused on the semantic region labeling problem [Krähenbühl and Koltun, 2011][Socher *et al.*, 2011][Long *et al.*, 2015], while the critical interaction relation prediction is often overlooked. Based on the hand-crafted features and models, the CRF inference [Ladicky *et al.*, 2009][Krähenbühl and Koltun, 2011] refines the labeling results by considering the label agreement between similar pixels. The fully convolutional network (FCN) [Long *et al.*, 2015] and its expansion [Chen *et al.*, 2015] have achieved great success on the semantic labeling. [Liu *et al.*, 2015] incorporates the markov random field (MRF) into deep networks for pixel-level labeling. Most recently, the multi-dimensional LSTM [Byeon *et al.*, 2015] has also been employed to capture the local spatial dependencies. However, our H-LSTM differs from these works in that we train a unified network to collaboratively address the ge-

ometric region labeling and relation prediction. The novel P-LSTM and MS-LSTM can effectively capture the long-range spatial dependencies benefiting from the hierarchical feature representation on the pixels and multi-scale super-pixels.

Single View 3D Reconstruction. The 3D reconstruction from the single view image is an under explored task and only a few researches have made some efforts on this task. Mobahi et al. [Mobahi *et al.*, 2011] reconstructed the urban structures from the single view by transforming invariant low-rank textures. Without the explicit assumptions about the structure of the scene, Saxena et al. [Saxena *et al.*, 2009] trained the MRF model to discover the depth cues as well as the relationships between different parts of the image in a fully supervised manner. An attribute grammar model [Liu *et al.*, 2014] regarded super-pixels as its terminal nodes and applied five production rules to generate the scene into a hierarchical parse graph. Differed from the previous methods, the proposed H-LSTM predicts the layout segmentation and the spatial arrangement with a unified network architecture, and thus can reconstruct the 3D scene from a still image directly.

3 Hierarchical LSTM

Overview. The geometric scene parsing aims to generate the pixel-wise geometric surface labeling and relation prediction for each image. As illustrated in Figure 2, the input image is first passed through CNN to generate a set of feature maps. Then the P-LSTM and MS-LSTM take these feature maps as inputs in a shared mode, and their outputs are the geometric surface labeling and interaction relations between adjacent regions respectively.

Notations. Each LSTM [Hochreiter and Schmidhuber, 1997] unit in i -th layer receives the input \mathbf{x}_i from the previous state, and determines the current state which is comprised of the hidden cells $\mathbf{h}_{i+1} \in R^d$ and the memory cells $\mathbf{m}_{i+1} \in R^d$, where d is the dimension of the network output. Similar to the work in [Graves *et al.*, 2013], we apply g^u, g^f, g^v, g^o to indicate the input, forget, memory and output gate respectively. Define W^u, W^f, W^v, W^o as the corresponding recurrent gate weights. Thus the hidden and memory cells for the current state can be calculated by,

$$\begin{aligned} g^u &= \phi(W^u * \mathbf{H}_i) \\ g^f &= \phi(W^f * \mathbf{H}_i) \\ g^o &= \phi(W^o * \mathbf{H}_i) \\ g^v &= \tanh(W^v * \mathbf{H}_i) \\ \mathbf{m}_{i+1} &= g^f \odot \mathbf{m}_i + g^u \odot g^v \\ \mathbf{h}_{i+1} &= \tanh(g^o \odot \mathbf{m}_i) \end{aligned} \quad (1)$$

where \mathbf{H}_i denotes the concatenation of input \mathbf{x}_i and previous state \mathbf{h}_i . ϕ is a sigmoid function with the form $\phi(t) = 1/(1 + e^{-t})$, and \odot indicates the element-wise product. Following [Kalchbrenner *et al.*, 2015], we can simplify the expression Eqn.(1) as,

$$(\mathbf{m}_{i+1}, \mathbf{h}_{i+1}) = \text{LSTM}(\mathbf{H}_i, \mathbf{m}_i, W) \quad (2)$$

where W is the concatenation of four different kinds of recurrent gate weights.

3.1 P-LSTM for Geometric Surface Labeling

Following [Liang *et al.*, 2015], we use the P-LSTM to propagate the local information to each position and further discover the short-distance contextual interactions in pixel level. For the feature representation of each position j , we extract $N = 8$ spatial hidden cells from N local neighbor pixels and one depth hidden cells from previous layer. Note that the ‘‘depth’’ in a special position indicates the features produced by the hidden cells at that position in the previous layer. Let $\{\mathbf{h}_{j,i,n}^s\}_{n=1}^N$ indicate the set of hidden cells from neighboring positions to pixel j , which are calculated by the N spatial LSTMs updated in i -th layer. And $\mathbf{h}_{j,i}^t$ denotes the hidden cells computed by the i -th layer depth LSTM on the pixel j . Then the input states of pixel j for the $(i + 1)$ -th layer LSTM can be expressed by,

$$\mathbf{H}_{j,i} = [\mathbf{h}_{j,i,1}^s \ \mathbf{h}_{j,i,2}^s \ \dots \ \mathbf{h}_{j,i,n}^s \ \mathbf{h}_{j,i}^t]^T \quad (3)$$

where $\mathbf{H}_{j,i} \in R^{(N+1) \times d}$. By the same token, let $\{\mathbf{m}_{j,i,n}^s\}_{n=1}^N$ be the memory cells for all N spatial dimensions of pixel j in the i -th layer and $\mathbf{m}_{j,i}^t$ be memory cell for the depth dimension. Then the hidden cells and memory cells of each position j in the $(i + 1)$ -th layer for all $N + 1$ dimensions are calculated as,

$$(\mathbf{m}_{j,i+1,n}^s, \tilde{\mathbf{h}}_{j,i+1,n}^s) = \text{LSTM}(\mathbf{H}_{j,i}, \mathbf{m}_{j,i,n}^s, W_i^s) \quad (4)$$

$n \in \{1, 2, \dots, N\};$

$$(\mathbf{m}_{j,i+1}^t, \mathbf{h}_{j,i+1}^t) = \text{LSTM}(\mathbf{H}_{j,i}, \mathbf{m}_{j,i}^t, W_i^t)$$

where W_i^s and W_i^t indicate the weights for spatial and depth dimension in the i -th layer, respectively. Note that $\tilde{\mathbf{h}}_{j,i+1,n}^s$ should be distinguished from $\mathbf{h}_{j,i+1,n}^s$ by the directions of information propagation. $\tilde{\mathbf{h}}_{j,i+1,n}^s$ represents the hidden cells position j to its n -th neighbor, which is used to generate the input hidden cells of n -th neighbor position for the next layer. In contrast, $\mathbf{h}_{j,i+1,n}^s$ is the neighbor hidden cells fed into Eqn.(3) to calculate the input state of pixel j .

In particular, the P-LSTM sub-network is built upon the modified VGG-16 model [Simonyan and Zisserman, 2015]. We remove the last two fully-connected layers in VGG-16, and replace with two fully-convolutional layers to obtain the convolutional feature maps for the input image. Then the convolutional feature maps are fed into the transition layer [Liang *et al.*, 2015] to produce hidden cells and memory cells of each position in advance, and make sure the number of the input states for the first P-LSTM layer is equal to that of following P-LSTM layer. Then the hidden cells and memory cells are passed through five stacked P-LSTM layers. By this way, the receptive field of each position can be considerably increased to sense a much larger contextual region. Note that the intermediate hidden cells generated by P-LSTM layer are also taken as the input to the corresponding Super-pixel LSTM layer for relation prediction. Please check more details of this part in Sec. 3.2. At last, several 1×1 feed-forward convolutional filters are applied to generate confidence maps for each geometric surface. The final label of each pixel is returned by a softmax classifier with the form,

$$y_j = \text{softmax}(\mathcal{F}(\mathbf{h}_j; W_{label})) \quad (5)$$




Figure 3: An illustration of super-pixel maps with different scales. In each scale, the orange super-pixel is the one under the current operation, and the blue ones are adjacent super-pixels, which propagate the neighboring information to the orange one. More contextual information can be captured by the larger-scale super-pixels.

where y_j is the predicted geometric surface probability of the j -th pixel, and W_{label} denotes the network parameter. $\mathcal{F}(\cdot)$ is a transformation function.

3.2 MS-LSTM for Interaction Relation Prediction

The Multi-scale Super-pixel LSTM (MS-LSTM) is used to explore high-level interaction relation between pair-wise super-pixels, and predict the functional boundaries between geometric surfaces. The hidden cells of j -th position in i -th MS-LSTM layer are the concatenation of hidden cells $\mathbf{h}_{j,i}^t \in R^d$ from previous layer (same as the depth dimension in P-LSTM) and $\mathbf{h}_{j,i}^r \in R^d$ from the corresponding P-LSTM layer. For simplicity, we rewrite the enhanced hidden cells as $\tilde{\mathbf{h}}_{j,i} = [\mathbf{h}_{j,i}^t, \mathbf{h}_{j,i}^r]$. In each MS-LSTM layer, an over-segmentation algorithm [Liu *et al.*, 2011b] is employed to produce the super-pixel map S^i with a specific scale c_i . To obtain the compact feature representation for each super-pixel, we use *Log-Sum-Exp*(LSE) [Boyd and Vandenberghe, 2004], a convex approximation of the *max* function to fuse the hidden cells of pixels in the same super-pixel,

$$\mathbf{h}_{\Lambda,i} = \frac{1}{\pi} \log \left[\frac{1}{Q_{\Lambda}} \sum_{j \in \Lambda} \exp(\pi \tilde{\mathbf{h}}_{j,i}) \right] \quad (6)$$

where $\mathbf{h}_{\Lambda,i} \in R^{2d}$ denotes the hidden cells of the super-pixel Λ in the i -th super-pixel layer, $\tilde{\mathbf{h}}_{j,i}$ denotes the enhance hidden cells of the j -th position, Q_{Λ} is the total number of pixels in Λ , and π is a hyper-parameter to control smoothness. With higher value of π , the function tends to preserve the max value for each dimension in the hidden cells, while with lower value the function behaves like an averaging function.

Similar to the Eqn.(3), let $\{\mathbf{h}_{\Lambda,i,k}\}_{k=1}^{K_{\Lambda}}$ indicate the set of hidden cells from K_{Λ} adjacent super-pixels of Λ . Then the input states of super-pixel Λ for the $(i+1)$ -th MS-LSTM layer can be computed by,

$$\mathbf{H}_{\Lambda,i} = \left[\frac{1}{K_{\Lambda}} \sum_k \mathbf{h}_{\Lambda,i,k} \quad \mathbf{h}_{\Lambda,i} \right]^T \quad (7)$$

where $\mathbf{H}_{\Lambda,i} \in R^{4d}$. The hidden cells and memory cells of super-pixel Λ in the $(i+1)$ -th layer can be calculated by,

$$(\mathbf{m}_{\Lambda,i+1}, \mathbf{h}_{\Lambda,i+1}) = \text{LSTM}(\mathbf{H}_{\Lambda,i}, \mathbf{m}_{\Lambda,i}, W'_i) \quad (8)$$

where W'_i denotes the concatenation gate weights of i -th MS-LSTM layer. $\mathbf{m}_{\Lambda,i}$ is the average value of the memory cells

of each position in super-pixel Λ . Note that the dimension of $\mathbf{h}_{\Lambda,i+1}$ in Eqn.(8) is d , which is equal to the output hidden cells from the P-LSTM. In the $(i+1)$ -th layer, the values of $\mathbf{h}_{\Lambda,i+1}$ and $\mathbf{m}_{\Lambda,i+1}$ can be directly assigned to the hidden cells and memory cells of each position in super-pixel Λ . Then the new hidden states can be accordingly learned by applying MS-LSTM layer on the super-pixel map with larger scale.

In particular, the MS-LSTM layers share the convolutional feature maps with the P-LSTM. In total, five stacked MS-LSTM layers are applied to extract hierarchical feature representations with different scales of contextual dependencies. Therefore, five super-pixel maps with different scales (i.e. 16, 32, 48, 64 and 128) are extract by the over-segmentation algorithm. Note that the scale in here refers to the average number of pixels in each super-pixel. Thus these multi-scale super-pixel maps are employed by different MS-LSTM layers, and the hidden cells for each layer are enhanced by the output of the corresponding P-LSTM layer. After passing through these hierarchical MS-LSTM layers, the local inference of each super-pixel can be influenced by different degrees of context, which enables the model simultaneously taking the local semantic information into account. Finally, the relation prediction of adjacent super-pixels is optimized as,

$$z_{\{\Lambda,\Lambda'\}} = \text{softmax}(\mathcal{F}([\mathbf{h}_{\Lambda} \quad \mathbf{h}_{\Lambda'}]; W'_{relation})) \quad (9)$$

where $z_{\{\Lambda,\Lambda'\}}$ is the predicted relation probability vector between super-pixel Λ and Λ' , and $W'_{relation}$ denotes the network parameters. $\mathcal{F}(\cdot)$ is a transformation function.

3.3 Model Optimization

The total loss of H-LSTM is the sum of losses of two tasks: geometric surface labeling loss \mathcal{J}_C by P-LSTM and relation prediction loss \mathcal{J}_R by MS-LSTM. Given U training images with $\{(I_1, \hat{Y}_1, \hat{Z}_1), \dots, (I_U, \hat{Y}_U, \hat{Z}_U)\}$, where \hat{Y} indicates the groundtruth geometric surfaces for all pixels for image I , and \hat{Z} denotes the groundtruth relation labels for all of adjacent super-pixel pairs in different scales. The overall loss function is as follows,

$$\mathcal{J}(W) = \frac{1}{U} \sum_{i=1}^U (\mathcal{J}_C(W_P; I_i, \hat{Y}_i) + \mathcal{J}_R(W_S; I_i, \hat{Z}_i)) \quad (10)$$

where W_P and W_S indicate the parameters of P-LSTM and MS-LSTM, respectively, and W denotes all of the parameters with the form $W = \{W_P, W_S, W_{CNN}\}$. W_{CNN} is the parameters of Convolution Neural Network. We apply the back propagation algorithm to update all the parameters. $\mathcal{J}_C(\cdot)$ is the standard pixel-wise cross-entropy loss. $\mathcal{J}_R(\cdot)$ is the cross-entropy loss for all super-pixels under all scales. Each MS-LSTM layer with a specific scale of the super-pixel map can output the final interaction relation prediction. Note that $\mathcal{J}_R(\cdot)$ is the sum of losses after all MS-LSTM layers.

4 Application to 3D Reconstruction

In this work, we apply our geometric scene parsing results for single-view 3D reconstruction. The predicted geometric surfaces and their relations are used to "cut and fold" the image

into a pop-up model [Hoiem *et al.*, 2005]. This process contains two main steps: (1) restoring the 3D spatial structure based on the interaction relations between adjacent super-pixels, (2) constructing the positions of the specific planes using projective geometry and texture mapping from the labelled image onto the planes. In practice, we first find the ground-vertical boundary according to the predicted supporting relations and estimate the horizon position as the benchmark of 3D structure. Then the algorithm uses the different kinds of predicted relations to generate the polylines and folds the space along these polylines. The algorithm also cuts the ground-sky and vertical-sky boundaries according to the layering relations. At last, the geometric surface is projected onto the above 3D structures to reconstruct the 3D model.

5 Experiment

5.1 Experiment Settings

Datasets. We validate the effectiveness of the proposed H-LSTM on three public datasets, including **SIFT-Flow** dataset [Liu *et al.*, 2011a], **LM+SUN** dataset [Tighe and Lazebnik, 2013] and **Geometric Context** dataset [Hoiem *et al.*, 2007]. The **SIFT-Flow** consists of 2,488 training images and 200 testing images. The **LM+SUN** contains 45,676 images (21,182 indoor images and 24,494 outdoor images). Following [Tighe and Lazebnik, 2013], we apply 45,176 images as training data and 500 images as test ones. For these two datasets, three geometric surface classes (i.e. *sky*, *ground* and *vertical*) are considered for the evaluation. The **Geometric Context** dataset includes 300 outdoor images, where 50 images are used for training and the rest for testing as [Liu *et al.*, 2014]. Except for the three main geometric surface classes as used in the previous two datasets, **Geometric Context** dataset also labels the five subclasses: *left*, *center*, *right*, *porous*, and *solid* for *vertical* class. For all of three datasets, four interaction relation labels (i.e. *layering*, *supporting*, *siding* and *affinity*) are defined and evaluated in our experiments.

Evaluation Metrics. Following [Long *et al.*, 2015], we use the pixel accuracy and mean accuracy metrics as the standard evaluation criteria for the geometric surface labeling. The pixel accuracy assesses the classification accuracy of pixels over the entire dataset while the mean accuracy calculates the mean accuracy for all categories. To evaluate the performance of relation prediction, the average precision metric is adopted.

Implementation Details. In our experiment, we keep the original size 256×256 of the input image for the SIFT-Flow dataset. The scale of input image is fixed as 321×321 for LM+SUN and Geometric Context datasets. During the training phase, the learning rates of transition layer, P-LSTM layers and MS-LSTM layers are initialized as 0.001 and that of pre-training CNN model is initialized as 0.0001. The dimension of hidden cells and memory cells, which is corresponding to the symbol d in Sec. 3, is set as 64 in both P-LSTM and MS-LSTM.

5.2 Performance Comparisons

Geometric Surface Labeling. We compare the proposed H-LSTM with three recent state-of-the-art approaches, including Superparsing [Tighe and Lazebnik, 2013], FCN [Long

Method	Sky	Ground	Vertical	Mean Acc.
Superparsing	-	-	-	89.2
FCN	96.4	93.1	91.8	93.8
DeepLab	96.1	93.8	93.4	94.4
Ours	96.4	95.1	93.1	94.9

Table 1: Comparison of geometric surface labeling performance with three state-of-the-art methods on SIFT-Flow dataset.

Method	Sky	Ground	Vertical	Mean Acc.
Superparsing	-	-	-	86.8
FCN	81.8	83.5	94.1	86.4
DeepLab	76.2	72.8	94.6	81.2
Ours	83.9	83.6	94.1	87.2

Table 2: Comparison of geometric surface labeling performance with three state-of-the-art methods over LM+SUN dataset.

et al., 2015] and DeepLab [Chen *et al.*, 2015] on the SIFT-Flow and LM+SUN datasets. Figure 4 gives the the comparison results on the pixel accuracy. Table 1 and Table 2 show the performance of our H-LSTM and comparisons with three state-of-the-art methods on the per-class accuracy. It can be observed that the proposed H-LSTM can significantly outperform three baselines in terms of both metrics. For the Geometric Context dataset, the model is fine-tuned based on the trained model on LM+SUN due to the small size of training data. We compare our results with those reported in [Hoiem *et al.*, 2008], [Tighe and Lazebnik, 2013] and [Liu *et al.*, 2014]. Table 3 reports the pixel accuracy on three main classes and five subclasses. Our H-LSTM can outperform the three baselines over 3.8% and 2.8% when evaluating on three main classes and five subclasses, respectively. This superior performance achieved by H-LSTM on three public datasets demonstrates that incorporating the coupled P-LSTM and MS-LSTM in a unified network is very effective in capturing the complex contextual patterns within images that are critical to exploit the diverse surface structures.

Interaction Relation Prediction. The MS-LSTM sub-network can predict the interaction relation results for two adjacent super-pixels. Note that we use five MS-LSTM layers and five scales of super-pixel maps are sequentially employed, including 128, 64, 48, 32, 16 super-pixels in five layers. The H-LSTM outputs the interaction relation prediction results after each MS-LSTM layer to enable the deep super-

Method	Subclasses	Main classes
Hoiem <i>et al.</i>	68.8	89.0
Superparsing	73.7	88.2
Liu <i>et al.</i>	76.3	-
Ours	80.1	91.8

Table 3: Comparison of geometric surface labeling performance with three state-of-the-arts methods in terms of mean accuracy on Geometric Context dataset.




Figure 4: Geometric surface labeling results (Pixel-wise Accuracy) on SIFT-Flow and LM+SUN datasets.

The number of MS-LSTM layers	SIFT-Flow	LM+SUN	G-Context
H-LSTM_1	85.8	89.1	87.8
H-LSTM_2	89.8	94.7	90.6
H-LSTM_3	90.3	95.6	89.8
H-LSTM_4	90.4	96.7	90.7
H-LSTM	91.2	95.8	90.8

Table 4: Comparisons of interaction relation prediction results (Average Precision) by using different number of MS-LSTM layers on three datasets. ‘‘H-LSTM_1’’, ‘‘H-LSTM_2’’, ‘‘H-LSTM_3’’, ‘‘H-LSTM_4’’ represent the results using 1,2,3,4 MS-LSTM layers, respectively.

vision for better feature learning. Table 4 shows the average precision after passing different number of MS-LSTM layers. The improvements can be observed on most of datasets by gradually using more MS-LSTM layers. It verifies well the effectiveness of exploiting more discriminative feature representation based on the hierarchical multi-scale super-pixel LSTM. The hierarchical MS-LSTM enables the model to simultaneously capture the global geometric structure information by increasingly sensing the larger contextual region and also keep track of local fine details by remembering the local interaction of small super-pixels.

5.3 Ablative Study

We further evaluate different architecture variants to verify the effectiveness of the important components in our H-LSTM, presented in Table 5.

Comparison with convolutional layers. To strictly evaluate the effectiveness of using the proposed P-LSTM layer, we report the performance of purely using convolutional layers, i.e. ‘‘convolution’’. To make fair comparison with P-LSTM layer, we utilize five convolutional layers, each of which contains $576 = 64 \times 9$ convolutional filters with size 3×3 , because nine LSTMs are used in a P-LSTM layer and each of them has 64 hidden cell outputs. Compared with ‘‘H-LSTM (ours)’’, ‘‘convolution’’ decreases the pixel accuracy. It demonstrates the superiority of using P-LSTM layers to harness complex long-distances dependencies over convolutional layers.

Multi-task learning. Note that we jointly optimize the geometric surface labeling and relation prediction task within a unified network. We demonstrate the effectiveness of multi-task learning by comparing our H-LSTM with the version that only predicts the geometric surface labeling, i.e. ‘‘P-LSTM’’. The supervision information for interaction relation and MS-

Model settings	SIFT-Flow	LM+SUN
Convolution	94.66	89.92
P-LSTM	94.68	90.13
P-LSTM + S-LSTM	95.24	91.06
H-LSTM (ours)	95.41	91.34

Table 5: Performance comparisons with different variants of our method in terms of pixel accuracy.




Figure 5: Some results of single-view 3D reconstruction. The first column is the original image. The second column is the geometric surface labeling result and the last two columns are the reconstruction results from two different views.

LSTM networks are discarded in ‘‘P-LSTM’’. The large performance decrease speaks well that these two tasks can mutually benefit from each other and help learn more meaningful and discriminative features.

Comparison with single scale of super-pixel map. We also validate the advantage of using multi-scale super-pixel representation in the MS-LSTM sub-network on interaction relation prediction. ‘‘S-LSTM’’ shows the results of using the same scale of super-pixels (i.e. 48 super-pixels) in each S-LSTM layer. The improvement of ‘‘H-LSTM’’ over ‘‘P-LSTM+S-LSTM’’ demonstrates that the richer contextual dependencies can be captured by using hierarchical multi-scale feature learning.

5.4 Application to 3D Reconstruction

Our main geometric class labels and interaction relation prediction over regions are sufficient to reconstruct scaled 3D models of many scenes. Figure 5 shows some scene images and the reconstructed 3D scenes generated based on our geometric parsing results. Besides the obvious graphic applications, e.g. creating virtual walkthroughs, we believe that extra valuable information could be provided by such models to other artificial intelligence applications.

6 Conclusion

In this paper, we have proposed a multi-scale and context-aware scene parsing model via recurrent Long Short-Term Memory neural network. Our approach have demonstrated a new state-of-the-art on the problem of geometric scene parsing, and also impressive results on 3D reconstruction from still images.

References

- [Boyd and Vandenberghe, 2004] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [Byeon *et al.*, 2014] Wonmin Byeon, Marcus Liwicki, and Thomas M Breuel. Texture classification using 2d lstm networks. In *ICPR*, 2014.
- [Byeon *et al.*, 2015] Wonmin Byeon, Thomas M. Breuel, Federico Raue, and Marcus Liwicki. Scene labeling with LSTM recurrent neural networks. In *CVPR*, 2015.
- [Chen *et al.*, 2015] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *ICLR*, 2015.
- [Graves *et al.*, 2007] Alex Graves, Santiago Fernández, and Jürgen Schmidhuber. Multi-dimensional recurrent neural networks. In *ICANN*, 2007.
- [Graves *et al.*, 2013] A. Graves, A. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *ICASSP*, 2013.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [Hoiem *et al.*, 2005] Derek Hoiem, Alexei A. Efros, and Martial Hebert. Automatic photo pop-up. *ACM Trans. Graph.*, 24(3):577–584, 2005.
- [Hoiem *et al.*, 2007] Derek Hoiem, Alexei A Efros, and Martial Hebert. Recovering surface layout from an image. *International Journal of Computer Vision*, 75(1):151–172, 2007.
- [Hoiem *et al.*, 2008] Derek Hoiem, Alexei Efros, Martial Hebert, et al. Closing the loop in scene interpretation. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- [Kalchbrenner *et al.*, 2015] Nal Kalchbrenner, Ivo Danihelka, and Alex Graves. Grid long short-term memory. *arXiv preprint arXiv:1507.01526*, 2015.
- [Kanji, 2015] Tanaka Kanji. Unsupervised part-based scene modeling for visual robot localization. In *ICRA*, 2015.
- [Krähenbühl and Koltun, 2011] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2011.
- [Ladicky *et al.*, 2009] Lubor Ladicky, Christopher Russell, Pushmeet Kohli, and Philip H. S. Torr. Associative hierarchical crfs for object class image segmentation. In *ICCV*, 2009.
- [Liang *et al.*, 2015] Xiaodan Liang, Xiaohui Shen, Donglai Xiang, Jiashi Feng, Liang Lin, and Shuicheng Yan. Semantic object parsing with local-global long short-term memory. *arXiv preprint arXiv:1511.04510*, 2015.
- [Liu *et al.*, 2011a] Ce Liu, Jenny Yuen, and Antonio Torralba. Nonparametric scene parsing via label transfer. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(12):2368–2382, 2011.
- [Liu *et al.*, 2011b] Ming-Yu Liu, Oncel Tuzel, Srikumar Ramalingam, and Rama Chellappa. Entropy rate superpixel segmentation. In *CVPR*, 2011.
- [Liu *et al.*, 2014] Xiaobai Liu, Yibiao Zhao, and Song-Chun Zhu. Single-view 3d scene parsing by attributed grammar. In *CVPR*, 2014.
- [Liu *et al.*, 2015] Ziwei Liu, Xiao Xiao Li, Ping Luo, Chen Change Loy, and Xiaoou Tang. Semantic image segmentation via deep parsing network. In *ICCV*, 2015.
- [Long *et al.*, 2015] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [Mobahi *et al.*, 2011] Hossein Mobahi, Zihan Zhou, Allen Y. Yang, and Yi Ma. Holistic 3d reconstruction of urban structures from low-rank textures. In *ICCV Workshops*, 2011.
- [Nieuwenhuisen *et al.*, 2010] Matthias Nieuwenhuisen, Jörg Stückler, and Sven Behnke. Improving indoor navigation of autonomous robots by an explicit representation of doors. In *ICRA*, 2010.
- [Pinheiro and Collobert, 2015] Pedro H. O. Pinheiro and Roman Collobert. From image-level to pixel-level labeling with convolutional networks. In *CVPR*, 2015.
- [Saxena *et al.*, 2009] Ashutosh Saxena, Min Sun, and Andrew Y. Ng. Make3d: Learning 3d scene structure from a single still image. *IEEE Trans. Pattern Anal. Mach. Intell.*, 31(5):824–840, 2009.
- [Simonyan and Zisserman, 2015] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [Socher *et al.*, 2011] Richard Socher, Cliff Chiung-Yu Lin, Andrew Y. Ng, and Christopher D. Manning. Parsing natural scenes and natural language with recursive neural networks. In *ICML*, 2011.
- [Tighe and Lazebnik, 2013] Joseph Tighe and Svetlana Lazebnik. Superparsing - scalable nonparametric image parsing with superpixels. *International Journal of Computer Vision*, 101(2):329–349, 2013.