

Recurrently exploring class-wise attention in a hybrid convolutional and bidirectional LSTM network for multi-label aerial image classification



Yuansheng Hua¹, Lichao Mou¹, Xiao Xiang Zhu*

Remote Sensing Technology Institute (IMF), German Aerospace Center (DLR), Oberpfaffenhofen, 82234 Wessling, Germany
Signal Processing in Earth Observation (SiPEO), Technical University of Munich (TUM), Arcisstr. 21, 80333 Munich, Germany

ARTICLE INFO

Keywords:

Multi-label classification
High-resolution aerial image
Convolutional Neural Network (CNN) 1 Class
Attention Learning
Bidirectional Long Short-Term Memory (BiLSTM)
Class dependency

ABSTRACT

Aerial image classification is of great significance in the remote sensing community, and many researches have been conducted over the past few years. Among these studies, most of them focus on categorizing an image into one semantic label, while in the real world, an aerial image is often associated with multiple labels, e.g., multiple object-level labels in our case. Besides, a comprehensive picture of present objects in a given high-resolution aerial image can provide a more in-depth understanding of the studied region. For these reasons, aerial image multi-label classification has been attracting increasing attention. However, one common limitation shared by existing methods in the community is that the co-occurrence relationship of various classes, so-called class dependency, is underexplored and leads to an inconsiderate decision. In this paper, we propose a novel end-to-end network, namely class-wise attention-based convolutional and bidirectional LSTM network (CA-Conv-BiLSTM), for this task. The proposed network consists of three indispensable components: (1) a feature extraction module, (2) a class attention learning layer, and (3) a bidirectional LSTM-based sub-network. Particularly, the feature extraction module is designed for extracting fine-grained semantic feature maps, while the class attention learning layer aims at capturing discriminative class-specific features. As the most important part, the bidirectional LSTM-based sub-network models the underlying class dependency in both directions and produce structured multiple object labels. Experimental results on UCM multi-label dataset and DFC15 multi-label dataset validate the effectiveness of our model quantitatively and qualitatively.

1. Introduction

With the booming of remote sensing techniques in the recent years, a huge volume of high resolution aerial imagery is now accessible and benefits a wide range of real-world applications, such as urban mapping (Marmanis et al., 2018; Audebert et al., 2018; Marcos et al., 2018; Mou and Zhu, 2018a), ecological monitoring (Zarco-Tejada et al., 2014; Wen et al., 2017), geomorphological analysis (Mou and Zhu, 2018b; Lucchesi et al., 2013; Weng et al., 2018; Cheng et al., 2017), and traffic management (Mou and Zhu, 2018c; Mou and Zhu, 2016; Li et al., 2018). As a fundamental bridge between aerial images and these applications, image classification, which aims at categorizing images into semantic classes, has obtained wide attention, and many researches have been conducted recently (Nogueira et al., 2017; Yang and Newsam, 2010; Xia et al., 2017; Zhu et al., 2017; Demir and Bruzzone, 2016; Hu et al., 2015, 2018; Zhang et al., 2015; Huang et al., 2018; Mou et al., 2017). However, most existing studies assume that each image

belongs to only one label (e.g., scene-level labels in Fig. 1), while in reality, an image is usually associated with multiple labels (Tan et al., 2017). Furthermore, a comprehensive picture of objects present in an aerial image is capable of offering a holistic understanding of such image. With this intention, numerous researches, i.e., semantic segmentation (Ren et al., 2015; Long et al., 2015; Badrinarayanan et al., 2015) and object detection (Ren et al., 2015; Viola and Jones, 2001; Lin et al., 2017; Ren et al., 2017), have emerged recently. Unfortunately, it is extremely labor- and time-consuming to acquire ground truths for these studies (i.e., pixel-wise segmentation masks and bounding-box-level annotations). Compared to these expensive labels, image-level labels (cf. multiple object-level labels in Fig. 1) are at a fair low cost and readily accessible. To this end, multi-label classification, aiming at assigning an image with multiple object labels, is arising in both remote sensing (Karalas et al., 2016; Zeggada et al., 2017; Koda et al., 2018; Zeggada et al., 2018) and computer vision communities (Patterson and Hays, 2012; Chua et al., 2009; Everingham et al., 2010). In this paper,

* Corresponding author at: Remote Sensing Technology Institute (IMF), German Aerospace Center (DLR), Oberpfaffenhofen, 82234 Wessling, Germany.

E-mail addresses: yuansheng.hua@tum.de (Y. Hua), lichao.mou@dlr.de (L. Mou), xiaoxiang.zhu@dlr.de (X.X. Zhu).

¹ The first two authors contributed equally to this work.



Fig. 1. Example high resolution aerial images with their scene labels and multiple *object* labels. Common label pairs are **highlighted**. (a) Free way: *bare soil*, *car*, *grass*, *pavement* and *tree*. (b) Intersection: *building*, *car*, *grass*, *pavement* and *tree*. (c) Parking lot: *car* and *pavement*.

we deploy our efforts in exploring an efficient multi-label classification model.

1.1. The challenges of multi-label classification

Benefited from the fast-growing remote sensing technology, large quantities of high-resolution aerial images are available and widely used in many visual tasks. Along with such huge opportunities, challenges have come up inevitably.

On one hand, it is difficult to extract high-level features from high-resolution images. Considering its complex spatial structure, conventional hand-crafted features, and mid-level semantic models (Yang and Newsam, 2010; Shao et al., 2013; Risojevic and Babic, 2013; Lowe, 2004; Zhu et al., 2016) suffer from the poor performance of capturing holistic semantic features, which leads to an unsatisfactory classification ability.

On the other hand, underlying correlations between dependent labels are required to be unearthed for an efficient prediction of multiple object labels. E.g., the existence of ships infers to a high probable co-occurrence of the sea, while the presence of buildings is almost always accompanied by the coexistence of pavement. However, the recently proposed multi-label classification methods (Karalas et al., 2016; Zeggada et al., 2017; Koda et al., 2018; Zeggada et al., 2018) assumed that classes are independent and employed a set of binary classifiers (Karalas et al., 2016) or a regression model (Zeggada et al., 2017; Koda et al., 2018; Zeggada et al., 2018) to infer the existence of each class separately.

To summarize, a well-performed multi-label classification system requires powerful capabilities of learning holistic feature representations and should be capable of harnessing the implicit class dependency.

1.2. The motivation of our work

As our survey of related work shows above, recent approaches make few efforts to exploit the high-order class dependency, which constrains the performance in multi-label classification. Besides, direct utilization of CNNs pre-trained on natural image datasets (Zeggada et al., 2017; Koda et al., 2018; Zeggada et al., 2018) leads to a partial interpretation of aerial images due to their diverse visual patterns. Moreover, most state-of-the-art methods decompose multi-label classification into separate stages, which cuts off their inter-correlations and makes end-to-end training infeasible.

To tackle these problems, in this paper, we propose a novel end-to-end network architecture, class attention-based convolutional and bidirectional LSTM network (CA-Conv-BiLSTM), which integrates feature extraction and high-order class dependency exploitation together for multi-label classification. Contributions of our work to the literature are detailed as follows:

- We regard the multi-label classification of aerial images as a structured output problem instead of a simple regression problem. In this manner, labels are predicted in an ordered procedure, and the prediction of each label is dependent on others. As a consequence, the implicit class relevance is taken into consideration, and structured outputs are more reasonable and closer to the real-world case as compared to regression outputs.
- We propose an end-to-end trainable network architecture for multi-label classification, which consists of a feature extraction module (e.g., a modified network based on VGG-16), a class attention learning layer, and a bidirectional LSTM-based sub-network. These components are designed for extracting features from input images, learning discriminative class-specific features, and exploiting class dependencies, respectively. Besides, such a design makes it feasible to train the network in an end-to-end fashion, which enhances the compactness of our model.
- Considering that class dependencies are diverse in both directions, a bidirectional analysis is required for modeling such correlations. Therefore, we employ a bidirectional LSTM-based network, instead of a one-way recurrent neural network, to dig out class relationships.
- We build a new challenging dataset, DFC15 multi-label dataset, by reproducing from a semantic segmentation dataset, GRSS_DFC_2015 (DFC15) (IEEE GRSS data fusion contest, 2015). The proposed dataset consists of aerial images at a spatial resolution of 5 cm and can be used to evaluate the performance of networks for multi-label classification.

The following sections further introduce and discuss our network. Specifically, Section 2 provides an intuitive illustration of the class dependency and then details the structure of the proposed network in terms of its three fundamental components. Section 3 describes the setup of our experiments, and experimental results are discussed from quantitative and qualitative perspectives. Finally, the conclusion of this paper is drawn in Section 4.

2. Methodology

2.1. An observation

Current aerial image multi-label classification methods (Zeggada et al., 2017; Koda et al., 2018; Zeggada et al., 2018) consider such problem as a regression issue, where models are trained to fit a binary sequence, and each digit indicates the existence of its corresponding class. Unlike one-hot vectors, a binary sequence is allowed to contain more than one ‘hot’ value for indicating the joint existence of multiple candidate classes in one image. Besides, several researches (Karalas et al., 2016) formulate multi-label classification into several single-label classification tasks, and thus, train a set of binary classifiers for each class. Notably, one common assumption of these studies is that classes

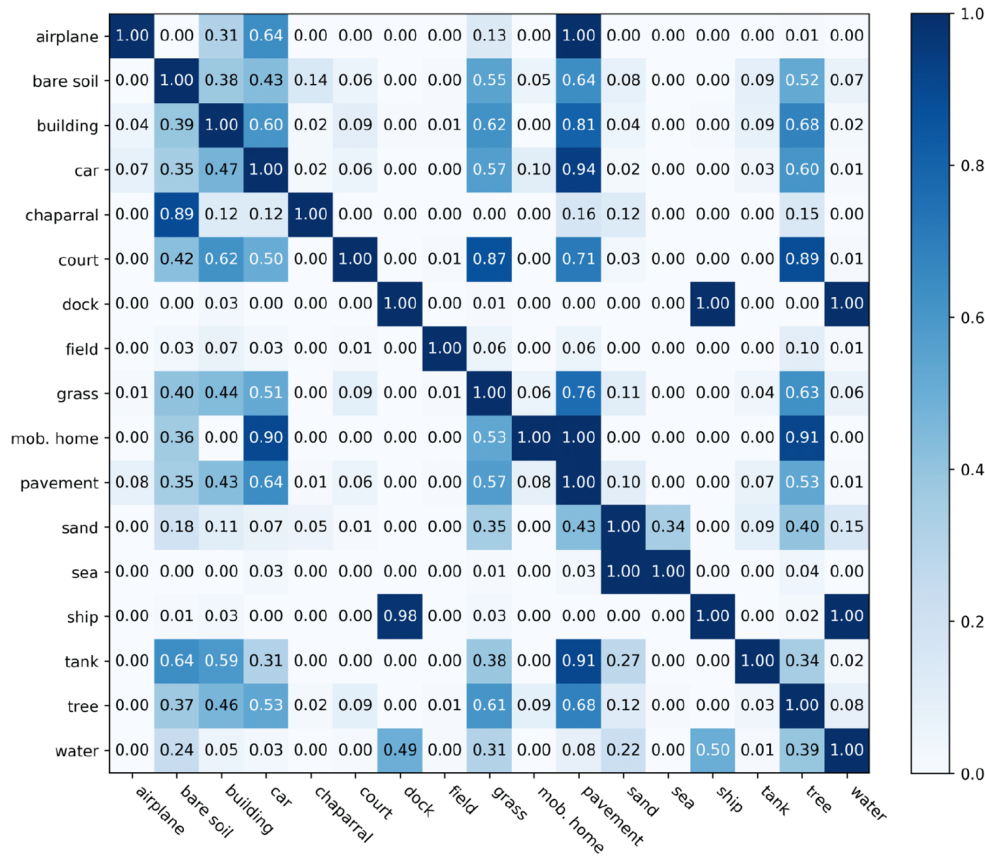


Fig. 2. The co-occurrence matrix of labels in UCM multi-label dataset. Notably, all images are taken into consideration when calculating this matrix. Labels at Y-axis represent referenced classes C_r , while labels at X-axis are potential co-occurrence classes C_p . The conditional probability $P(C_p|C_r)$ of each class pair is presented in the corresponding block.

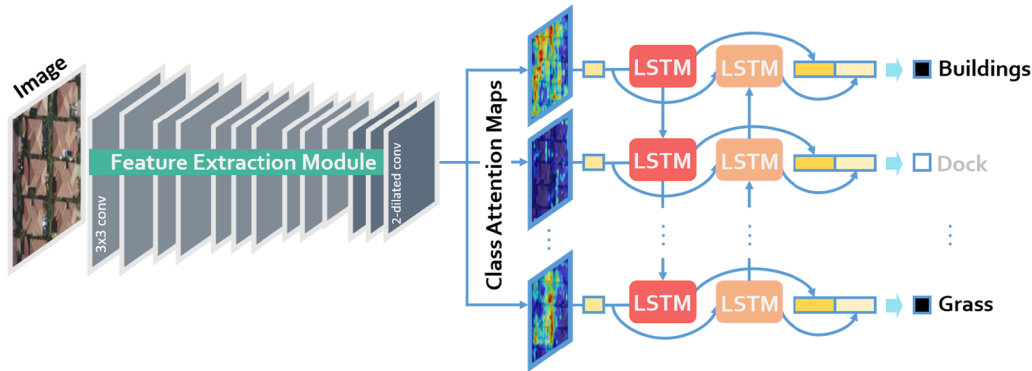


Fig. 3. The architecture of the proposed CA-Conv-BiLSTM for the multi-label classification of aerial images.

are independent of each other, and classifiers predict the existence of each category independently. However, this is violent and not accord with real life. As illustrated in Fig. 1, although images obtained in diverse scenes are assigned with multiple different labels, there are still common classes, e.g., car and pavement, coexisting in each image. This is because, in the real-life world, some classes have a strong correlation, for example, cars are often driven or parked on pavements. To further demonstrate the class dependency, we calculate conditional probabilities for each of the two categories. Let C_r denote referenced class, and C_p denote potential co-occurrence class. Conditional probability $P(C_p|C_r)$, which depicts the possibility that C_p exhibits in an image, where the existence of C_r is priorly known, can be solved with Eq. (1),

$$P(C_p|C_r) = \frac{P(C_p, C_r)}{P(C_r)}. \tag{1}$$

$P(C_p, C_r)$ indicates the joint occurrence probability of C_p and C_r , and $P(C_r)$ refers to the priori probability of C_r . Conditional probabilities of all class pairs in UCM multi-label datasets are listed in Fig. 2, and it is intuitive that some classes have strong dependencies. For instance, it is highly possible that there are pavements in images, which contain airplanes, buildings, cars, or tanks. Moreover, it is notable that class dependencies are not symmetric due to their particular properties. For example, $P(water|ship)$ is twice as $P(ship|water)$ due to the reason that the occurrence of ships always infer to the co-occurrence of water, while not vice versa. Therefore, to thoroughly dig out the correlation among various classes, it is crucial to model class probabilistic dependencies bidirectionally in a classification method.

To this end, we boil the multi-label classification down into a structured output problem, instead of a simple regression issue, and

employ a unified framework of a CNN and a bidirectional RNN to (1) extract semantic features from raw images and (2) model image-label relations as well as bidirectional class dependencies, respectively.

2.2. Network architecture

The proposed CA-Conv-BiLSTM, as illustrated in Fig. 3, is composed of three components: a feature extraction module, a class attention learning layer, and a Bidirectional LSTM-based recurrent sub-network. More specifically, the feature extraction module employs a stack of interleaved convolutional and pooling layers to extract high-level features, which are then fed into a class attention learning layer to produce discriminative class-specific features. Afterwards, a bidirectional LSTM-based recurrent sub-network is attached to model both probabilistic class dependencies and underlying relationships between image features and labels.

Section 2.2.1 details the architecture of the feature extraction module, and Section 2.2.2 describes the explicit design of the class attention learning layer. Finally, Section 2.2.3 introduces how to produce structured multi-label outputs from class-specific features via a bidirectional LSTM-based recurrent sub-network.

2.2.1. Dense high-level feature extraction

Learning efficient feature representations of input images is extremely crucial for the image classification task. To this end, a modern popular trend is to employ a CNN architecture to automatically extract discriminative features, and many recent studies (Hua et al., 2018; Mou and Zhu, 2018c; Xia et al., 2017; Kang et al., 2018; Zhu et al., 2017; Mou et al., 2017) have achieved great progress in a wide range of classification tasks. Inspired by this, our model adapts VGG-16 (Simonyan and Zisserman, 2014), one of the most welcoming CNN architectures for its effectiveness and elegance, to extract high-level features for our task.

Specifically, the feature extraction module consists of 5 convolutional blocks, and each of them contains 2 or 3 convolutional layers (as illustrated in the left of Fig. 3). Notably, the number of filters is equivalent in a common convolutional block and doubles after each pooling layer, which is utilized to reduce the spatial dimension of feature maps. The purpose of such design is to enable the feature extraction module to learn diverse features at a less computational expense. The receptive field of all convolutional filters is 3×3 , which increases nonlinearities inside the feature extraction module. Besides, the convolution stride is 1 pixel, and the spatial padding of each convolutional layer is set as 1 pixel as well. Among these convolutional blocks, max-pooling layers are interleaved for reducing the size of feature maps and meanwhile, maintaining only local representative, such as maximum in a 2×2 -pixel region. The size of pooling windows is 2×2 pixels, and the pooling stride is 2 pixels, which halves feature maps in width and length.

Features directly learned from a conventional CNN (e.g., VGG-16) are proved to be high-level and semantic, but their spatial resolution is significantly reduced, which is not favorable for generating high-dimensional class-specific features in the subsequent class attention learning layer. To address this, max-pooling layers following the last

two convolutional blocks are discarded in our model, and atrous convolutional filters with dilation rate 2 are employed in the last convolutional block for preserving original receptive fields. Consequently, our feature extraction module is capable of learning high-level features with finer spatial resolution, so-called “dense”, compared to VGG-16, and it is feasible to initialize our model with pre-trained VGG-16, considering that all filters have equivalent receptive fields.

Moreover, it is noteworthy that other popular CNN architectures can be taken as prototypes of the feature extraction module, and thus, we extend researches to GoogLeNet (Szegedy et al., 2015) and ResNet (He et al., 2016) for a comprehensive evaluation of CA-Conv-BiLSTM. Regarding GoogLeNet, i.e., Inception-v3 (Szegedy et al., 2016), the stride of convolutional and pooling layers after “mixed7” is reduced to 1 pixel, and the dilation rate of convolutional filters in “mixed9” is 2. For ResNet (we use ResNet-50), the convolution stride in last two residual blocks is set as 1 pixel, and the dilation rate of filters in the last residual block is 2. Besides, layers after global average pooling layers, as well as itself, are removed to ensure dense high-level feature maps.

2.2.2. Class attention learning layer

Although Features extracted from pre-trained CNNs are high-level and can be directly fed into a fully connected layer for generating multi-label predictions, it is infeasible to learn high-order probabilistic dependencies by recurrently feeding it with identical features. Therefore, extracting discriminative class-wise features plays a key role in discovering class dependencies and effectively bridging CNN and RNN for multi-label classification tasks.

Here, we propose a class attention learning layer to explore features with respect to each category, and the proposed layer, illustrated in the middle of Fig. 3, consists of the following two stages: (1) generating class attention maps via a 1×1 convolutional layer with stride 1, and (2) vectorizing each class attention map to obtain class-specific features. Formally, given feature maps X , extracted from the feature extraction module, with a size of $W \times W \times K$, and let w_l represent the l -th convolutional filter in the class attention learning layer. The attention map M_l for class l can be obtained with the following formula:

$$M_l = X * w_l, \quad (2)$$

where l ranges from 1 to the number of classes, and $*$ represents convolution operation. Considering that the size of convolutional filters is 1×1 , a class attention map M_l is intrinsically a linear combination of all channels in X . With this design, the proposed class attention learning layer is capable of learning discriminative class attention maps. Some examples are shown in Fig. 4. An aerial image (cf. Fig. 4a) in UCM multi-label dataset is first fed into the feature extraction module, adapted from VGG-16, and outputs of its last convolutional block are considered as the feature maps X in Eq. (2). Thus, X is abundant in high-level semantic information, and the size of X is $14 \times 14 \times 512$. Afterwards, a class attention learning layer, where the number of filters is equivalent to that of classes, is appended to generate class-specific feature representations with respect to all categories. With sufficient training, they are supposed to learn class-wise attention maps. It is observed that class attention maps highlight discriminative areas for different categories and exhibit almost no activations with respect to

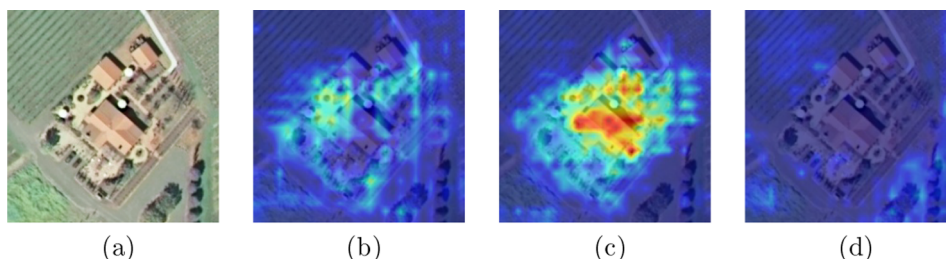


Fig. 4. Example class attention maps of an (a) aerial image, with respect to different classes: (b) bare soil, (c) building, and (d) water.

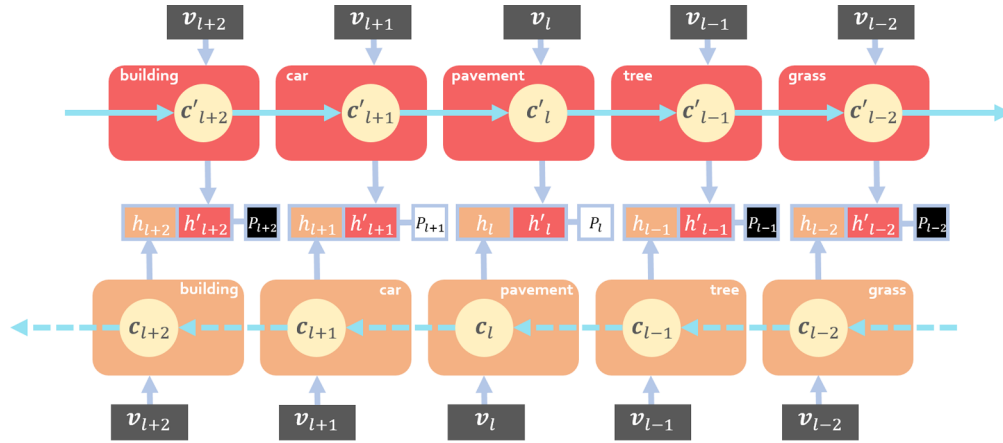


Fig. 5. Illustration of the bidirectional structure. The direction of the upper stream is opposite to that of the lower stream. Notably, h'_{l-1} , c'_{l-1} denotes the activation and memory cell in the upper stream at the time step, which corresponds to class $l - 1$ for convenience (considering that the subsequent time step is usually denoted as $l + 1$).

Table 1
The number of images in each object class.

Class No.	Class Name	Total	Training	Test
1	Airplane	100	80	20
2	Bare soil	718	577	141
3	Building	691	555	136
4	Car	886	722	164
5	Chaparral	115	82	33
6	Court	105	84	21
7	Dock	100	80	20
8	Field	104	79	25
9	Grass	975	804	171
10	Mobile home	102	82	20
11	Pavement	1300	1047	253
12	Sand	294	218	76
13	Sea	100	80	20
14	Ship	102	80	22
15	Tank	100	80	20
16	Tree	1009	801	208
17	Water	203	161	42
-	All	2100	1680	420

absent classes (as shown in Fig. 4c).

Subsequently, class attention maps M_l are transformed into class-wise feature vectors v_l of W^2 dimensions by vectorization. Instead of fully connecting class attention maps to each hidden unit in the following layer, we construct class-wise connections between class attention maps and their corresponding hidden units, i.e., corresponding time steps in an LSTM layer in our network. In this way, features fed into different units are retained to be class-specific discriminative and significantly contribute to the exploitation of the dynamic class dependency in the subsequent bidirectional LSTM layer.

2.2.3. Class dependency learning via a BiLSTM-based sub-network

As an important branch of neural networks, RNN is widely used in dealing with sequential data, e.g., textual data and temporal series, due to its strong capabilities of exploiting implicit dependencies among inputs. Unlike CNN, RNN is characterized by its recurrent neurons, of which activations are dependent on both current inputs and previous hidden states. However, conventional RNNs suffer from the gradient vanishing problem and are found difficult to learn long-term dependencies. Therefore, in this work, we seek to model class dependencies with an LSTM-based RNN, which is first proposed in Hochreiter and Schmidhuber (1997) and has shown great performance in processing long sequences (Graves, 2013; Gers et al., 1999; Xu et al., 2015; Mou et al., 2017, 2019).

Instead of directly summing up inputs as in a conventional recurrent layer, an LSTM layer relies on specifically designed hidden units, LSTM units, where information, such as the class dependency between category l and $l - 1$, is “memorized”, updated, and transmitted with a memory cell and several gates. Specifically, given a class-specific feature v_l obtained from the class attention learning layer as an input of the LSTM memory cell c_l at time step l , and let h_l represent the activation of c_l . New memory information \tilde{c}_l , learned from the previous activation h_{l-1} and the present input feature v_l , is obtained as follows:

$$\tilde{c}_l = \tanh(W_{cv}v_l + W_{ch}h_{l-1} + b_c), \quad (3)$$

where W_{cv} and W_{ch} denote weight matrix from input vectors to memory cell and hidden-memory coefficient matrix, respectively, and b_c is a bias term. Besides, $\tanh(\cdot)$ is the hyperbolic tangent function. In contrast to conventional recurrent units, where the \tilde{c}_l is directly used to update the current state h_l , an LSTM unit employs an input gate i_l to control the extent to which \tilde{c}_l is added, and meanwhile, partially omits uncorrelated prior information from c_{l-1} with a forget gate f_l . The two gates are performed by the following equations:

$$\begin{aligned} i_l &= \sigma(W_{iv}v_l + W_{ih}h_{l-1} + W_{ic}c_{l-1} + b_i), \\ f_l &= \sigma(W_{fv}v_l + W_{fh}h_{l-1} + W_{fc}c_{l-1} + b_f). \end{aligned} \quad (4)$$

Consequently, the memory cell c_l is updated by

$$c_l = i_l \odot \tilde{c}_l + f_l \odot c_{l-1}, \quad (5)$$

where \odot represents element-wise multiplication. Afterwards, an output gate o_l , formulated by

$$o_l = \sigma(W_{ov}v_l + W_{oh}h_{l-1} + W_{oc}c_l + b_o), \quad (6)$$

is designed to determine the proportion of memory content to be exposed, and eventually, the memory cell c_l at time step l is activated by

$$h_l = o_l \tanh(c_l). \quad (7)$$

Although it is not difficult to discover that the activation of the memory cell at each time step is dependent on both input class-specific feature vectors and previous cell states. However, taking into account that the class dependency is bidirectional, as demonstrated in Section 2.1, a single-directional LSTM-based RNN is insufficient to draw a comprehensive picture of inter-class relevance. Therefore, a bidirectional LSTM-based RNN, composed of two identical recurrent streams but with reversed directions, is introduced in our model, and the hidden units are updated based on signals from not only their preceding states but also subsequent ones.

In order to practically adapt a bidirectional LSTM-based RNN to modeling the class dependency, we set the number of time steps in our bidirectional LSTM-based sub-network equivalent to that of classes

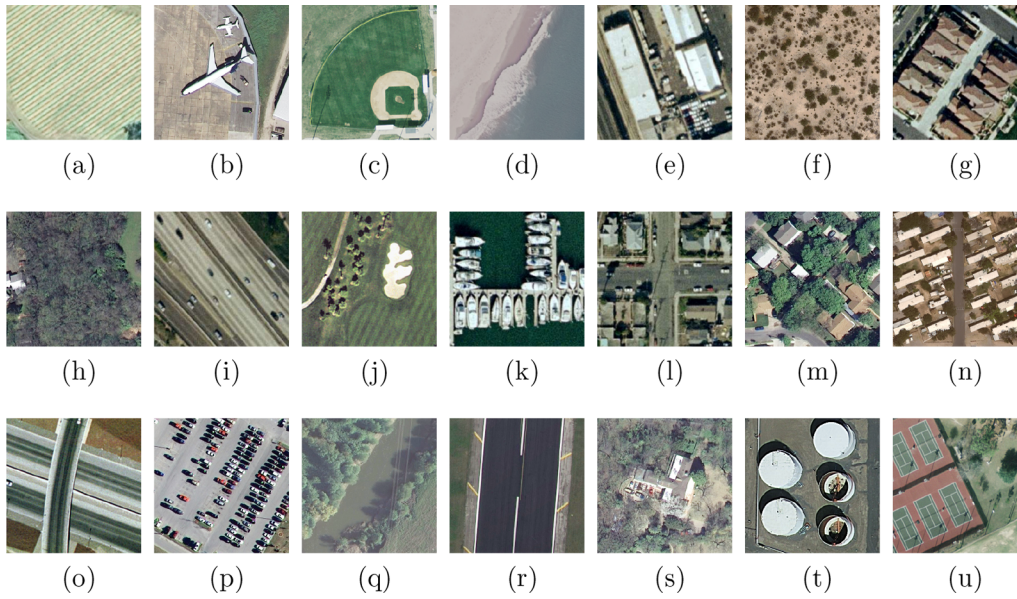


Fig. 6. Example images from each scene category and their corresponding multiple *object* labels in UCM multi-label dataset. Each image is 256×256 pixels with a spatial resolution of one foot, and their scene and *object* labels are introduced: (a) Agricultural: *field* and *tree*. (b) Airplane: *airplane*, *bare soil*, *car*, *grass* and *pavement*. (c) Baseball diamond: *bare soil*, *building*, *grass*, and *pavement*. (d) Beach: *sand* and *sea*. (e) building: *building*, *car*, and *pavement*. (f) Chaparral: *bare soil* and *chaparral*. (g) Dense residential: *building*, *car*, *grass*, *pavement*, and *tree*. (h) Forest: *building*, *grass*, and *tree*. (i) Free way: *bare soil*, *car*, *grass*, *pavement*, and *tree*. (j) Golf course: *grass*, *pavement*, *sand*, and *tree*. (k) Harbor: *dock*, *ship*, and *water*. (l) Intersection: *building*, *car*, *grass*, *pavement*, and *tree*. (m) Medium residential: *building*, *car*, *grass*, *pavement*, and *tree*. (n) Mobile home park: *bare soil*, *car*, *grass*, *mobile home*, *pavement*, and *tree*. (o) Overpass: *bare soil*, *car*, and *pavement*. (p) Parking lot: *car*, *grass*, and *pavement*. (q) River: *grass*, *tree*, and *water*. (r) Runway: *grass* and *pavement*. (s) Sparse residential: *bare soil*, *building*, *car*, *grass*, *pavement*, and *tree*. (t) Storage tank: *bare soil*, *pavement*, and *tank*. (u) Tennis court: *bare soil*, *court*, *grass*, and *tree*.

Table 2

The number of images in each object class.

Class No.	Class Name	Total	Training	Test
1	Impervious	3133	2532	602
2	Water	998	759	239
3	Clutter	1891	1801	90
4	Vegetation	1086	522	562
5	Building	1001	672	330
6	Tree	258	35	223
7	Boat	270	239	31
8	Car	705	478	227
–	All	3342	2674	668

under the assumption that distinct classes are predicted at respective time steps. Validated in Sections 3.3 and 3.4, such design enjoys two outstanding characteristics: on one hand, the LSTM memory cell at time step l , c_l , focuses on learning dependent relationship between class l and others in dual directions (cf. Fig. 5), and on the other hand, the occurrence probability of class l , P_l , can be predicted from outputs $[\mathbf{h}_l, \mathbf{h}_l']$ with a single-unit fully connected layer:

$$P_l = \sigma(\mathbf{w}_l[\mathbf{h}_l, \mathbf{h}_l'] + \mathbf{b}_l), \quad (8)$$

where \mathbf{h}_l' denotes the activation of c_l in the other direction, and σ is used as the activation function.

3. Experiments and discussion

In this section, two high-resolution aerial datasets of different resolution used for evaluating our network are first described in Section 3.1, and then, the training strategies are introduced in Section 3.2. Afterwards, the performance of the proposed network on the two datasets is quantitatively and qualitatively evaluated in the following sections.

3.1. Data description

3.1.1. UCM multi-label dataset

UCM multi-label dataset (Chaudhuri et al., 2018) is reproduced from UCM dataset (Yang and Newsam, 2010) by reassigning them with multiple object labels. Specifically, UCM dataset consists of 2100 aerial images of 256×256 pixels, and each of them is categorized into one of 21 scene labels: airplane, beach, agricultural, baseball diamond, building, tennis courts, dense residential, forest, freeway, golf course, mobile home park, harbor, intersection, storage tank, medium residential, overpass, sparse residential, parking lot, river, runway, and chaparral. For each of them, there are 100 images with a spatial resolution of one foot collected by cropping manually from aerial ortho imagery provided by the United States Geological Survey (USGS) National Map.

In contrast, images in UCM multi-label dataset are relabeled by assigning each image sample with one or more labels based on their primitive objects. The total number of newly defined object classes is 17: airplane, sand, pavement, building, car, chaparral, court, tree, dock, tank, water, grass, mobile home, ship, bare soil, sea, and field. It is notable that several labels, namely, airplane, building, and tank, are defined in both datasets but with variant level. In UCM dataset, they are scene-level labels, since they are predominant objects in an image and used to depict the whole image, while in UCM multi-label dataset, they are object-level labels, regarded as candidate objects in a scene. The numbers of images related to each object category are listed in Table 1, and examples from each scene category are shown in Fig. 6, as well as their corresponding object labels. To train and test our network on UCM multi-label dataset, we select 80% of sample images evenly from each scene category for training and the rest as the test set.

3.1.2. DFC15 multi-label dataset

Considering that images collected from the same scene may share similar patterns, alleviating task challenges, we build a new multi-label dataset, DFC15 multi-label dataset, based on a semantic segmentation

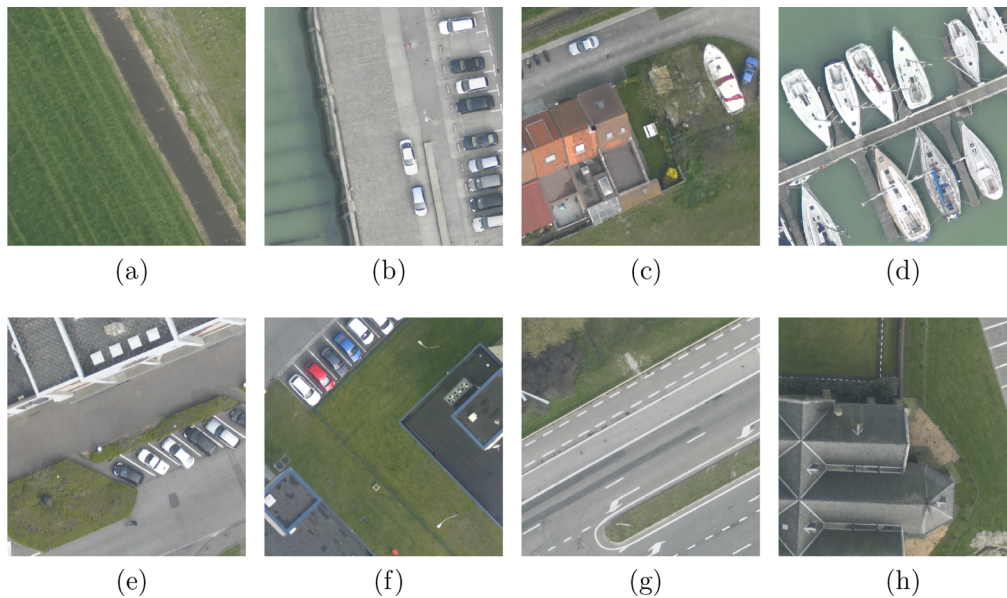


Fig. 7. Example images in DFC15 multi-label dataset and their multiple object labels. Each image is 600 × 600 pixels with a spatial resolution of 5 cm. (a) Water and vegetation. (b) Impervious, water, and car. (c) Impervious, water, vegetation, building, and car. (d) Water, clutter, and boat. (e) Impervious, vegetation, building, and car. (f) Impervious, vegetation, building, and car. (g) Impervious, vegetation, and tree. (h) Impervious, vegetation, and building.

Table 3
Configurations of CA-Conv-LSTM architectures.

Model	CNN model	Class Attention Map	Bi.
CA-VGG-LSTM	VGG-16	28 × 28 × N	✗
CA-VGG-BiLSTM	VGG-16	28 × 28 × N	✓
CA-GoogLeNet-LSTM	Inception-v3	17 × 17 × N	✗
CA-GoogLeNet-BiLSTM	Inception-v3	17 × 17 × N	✓
CA-ResNet-LSTM	ResNet-50	28 × 28 × N	✗
CA-ResNet-BiLSTM	ResNet-50	28 × 28 × N	✓

N indicates the number of classes in the dataset.

Bi. indicates whether the model is bidirectional or not.

dataset, DFC15 (IEEE GRSS data fusion contest, 2015), which was published and first used in 2015 IEEE GRSS Data Fusion Contest. DFC15 dataset is acquired over Zeebrugge with an airborne sensor, which is 300 m off the ground. In total, 7 tiles are collected in DFC dataset, and each of them is 10000 × 10000 pixels with a spatial resolution of 5 cm. Unlike UCM dataset, where images are assigned with image-level

labels, all tiles in DFC15 dataset are labeled in pixel-level, and each pixel is categorized into 8 distinct object classes: impervious, water, clutter, vegetation, building, tree, boat, and car. Notably, vegetation refers to low vegetation, such as bushes and grasses, and has no overlap with trees. Impervious indicates impervious surfaces (e.g., roads) excluding building rooftops.

Considering our task, the following processes are conducted: First, we crop large tiles into images of 600 × 600 pixels with a 200-pixel-stride sliding window. Afterwards, images containing unclassified pixels are ignored, and labels of all pixels in each image are aggregated into image-level multi-labels. An important characteristic of images in DFC15 multi-label dataset is lower inter-image similarity due to that they are cropped from vast regions consecutively without specific preferences, e.g., seeking images belonging to a specific scene. Moreover, extremely high resolution makes it more challenging as compared to UCM multi-label dataset. The numbers of images containing each object label are listed in Table 2, and example images with their image-level object labels are shown in Fig. 7. To conduct the evaluation, 80% of images are randomly selected as the training set, while the others are utilized to test our network.

Table 4
Quantitative results on UCM multi-label dataset (%).

Model	m.F ₁	m.F ₂	m.P _e	m.R _e	m.P _l	m.R _l
VGGNet (Simonyan and Zisserman, 2014)	78.54	80.17	79.06	82.30	86.02	80.21
VGG-RBFNN (Zeggada et al., 2017)	78.80	81.14	78.18	83.91	81.90	82.63
CA-VGG-LSTM	79.57	80.75	80.64	82.47	87.74	75.95
CA-VGG-BiLSTM	79.78	81.69	79.33	83.99	85.28	76.52
GoogLeNet (Szegedy et al., 2015)	80.68	82.32	80.51	84.27	87.51	80.85
GoogLeNet-RBFNN (Zeggada et al., 2017)	81.54	84.05	79.95	86.75	86.19	84.92
CA-GoogLeNet-LSTM	81.78	85.16	78.52	88.60	86.66	85.99
CA-GoogLeNet-BiLSTM	81.82	84.41	79.91	87.06	86.29	84.38
ResNet-50 (He et al., 2016)	79.68	80.58	80.86	81.95	88.78	78.98
ResNet-RBFNN (Zeggada et al., 2017)	80.58	82.47	79.92	84.59	86.21	83.72
CA-ResNet-LSTM	81.36	83.66	79.90	86.14	86.99	82.24
CA-ResNet-BiLSTM	81.47	85.27	77.94	89.02	86.12	84.26

m.F₁ and m.F₂ indicate the mean F₁ and F₂ score.

m.P_e and m.R_e indicate mean example-based precision and recall.

m.P_l and m.R_l indicate mean label-based precision and recall.

Table 5
Example Predictions on UCM and DFC15 Multi-label Dataset.

<p>Images in UCM Multi-label Dataset</p>					
<p>Ground Truths</p>	dock, ship, and water	building, car, pavement, and tree	building, court, pavement, grass, and tree	car, pavement, mobile-home, and tree	bare soil, car, grass, and pavement
<p>Predictions</p>	dock, ship, and water	building, car, pavement, and tree	building, court, pavement, grass, and tree	car, pavement, mobile-home, tree and grass	bare soil, car, grass, tree, and pavement
<p>Images in DFC15 Multi-label Dataset</p>					
<p>Ground Truths</p>	impervious, water, and building	impervious, vegetation, car, and building	impervious, vegetation, building, clutter, and car	water, vegetation, tree	impervious, vegetation, building, car, and tree
<p>Predictions</p>	impervious, water, and building	impervious, vegetation, car, and building	impervious, vegetation, building, clutter, and car	impervious, water, tree, vegetation	impervious, vegetation, building, car, and tree
<p>Red predictions indicate false positives, while blue predictions are false negatives.</p>					

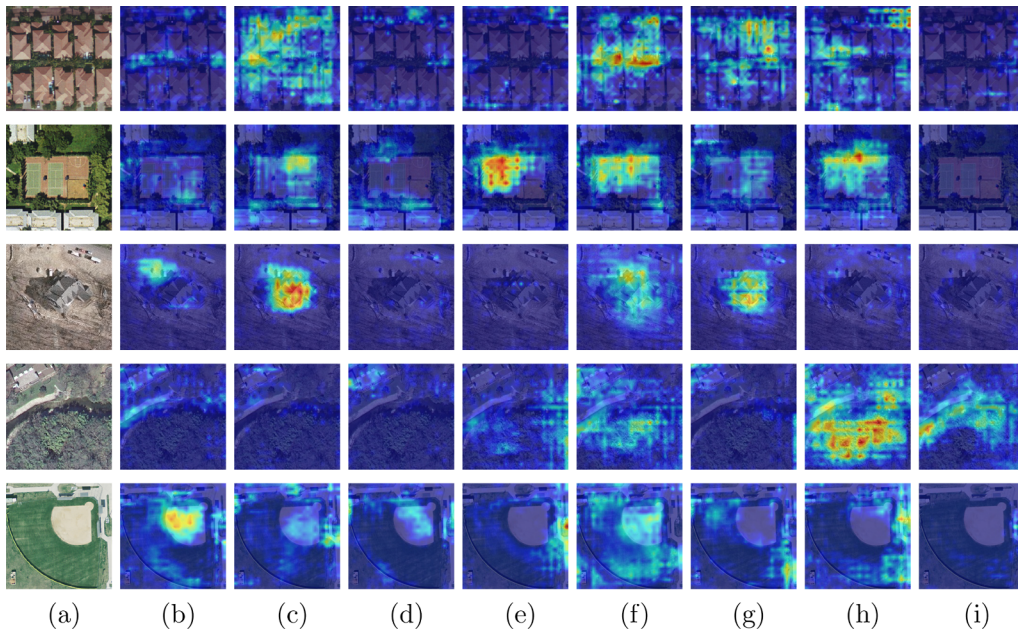


Fig. 8. Example class attention maps of (a) images in UCM multi-label dataset with respect to (b) bare soil, (c) building, (d) car, (e) court, (f) grass, (g) pavement, (h) tree, and (i) water. Red indicates strong activations, while blue represents non-activations. Besides, normalization is performed based on each row for a fair comparison among class attention maps of the same images. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 6
Quantitative results on DFC15 multi-label dataset (%).

Model	m.F ₁	m.F ₂	m.P _e	m.R _e	m.P _l	m.R _l
VGGNet (Simonyan and Zisserman, 2014)	73.86	74.09	76.16	74.95	62.57	59.95
VGGNet-RBFNN (Zeggada et al., 2017)	72.21	73.02	74.08	74.42	60.82	66.58
CA-VGG-LSTM	75.46	75.85	77.95	76.95	73.56	59.19
CA-VGG-BiLSTM	76.25	76.93	78.27	78.30	74.99	64.31
GoogLeNet (Szegedy et al., 2015)	74.99	73.41	81.01	73.01	71.80	53.95
GoogLeNet-RBFNN (Zeggada et al., 2017)	73.38	72.62	78.46	72.94	64.62	63.22
CA-GoogLeNet-LSTM	75.67	75.46	79.08	76.12	70.22	60.65
CA-GoogLeNet-BiLSTM	78.25	76.80	83.97	76.52	82.98	61.04
ResNet-50 (He et al., 2016)	78.10	76.21	84.89	75.64	81.50	59.99
ResNet-RBFNN (Zeggada et al., 2017)	78.36	78.08	82.64	78.76	72.01	69.85
CA-ResNet-LSTM	78.78	76.65	85.66	75.84	83.83	60.05
CA-ResNet-BiLSTM	83.65	80.61	91.93	79.12	94.35	62.35

3.2. Training details

The proposed CA-Conv-BiLSTM is initialized with separate strategies with respect to three dominant components: (1) the feature extraction module is initialized with CNNs pre-trained on ImageNet dataset (Deng et al., 2009), (2) convolutional filters in the class attention learning layer is initialized with a Glorot uniform initializer, and (3) all weights in the bidirectional 2048-d LSTM layer are randomly initialized in the range of $[-0.1, 0.1]$ with a uniform distribution. Notably, weights in the feature extraction module are trainable and fine-tuned during the training phase of our network.

Regarding the optimizer, we chose Adam with Nesterov momentum (Dozat, 2016), claimed to converge faster than stochastic gradient descent (SGD), and set parameters of the optimizer as recommended: $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 1e - 08$. The learning rate is set as $1e - 04$ and decayed by 0.1 when the validation accuracy is saturated. The loss of the network is defined as the binary cross entropy. We implement the network on TensorFlow and train it on one NVIDIA Tesla P100 16 GB GPU for 100 epochs. The size of the training batch is 32 as a trade-off between GPU memory capacity and training speed. To avoid overfitting, we stop training procedure when the loss fails to decrease in five epochs. Concerning ground truths, multiple labels of an image are

encoded into a multi-hot binary sequence, of which the length is equivalent to the number of all candidate labels. For each digit, 1 indicates the existence of its corresponding label, while 0 denotes the absent label.

3.3. Results on UCM multi-label dataset

3.3.1. Quantitative results

To evaluate the performance of CA-Conv-BiLSTM for multi-label classification of high resolution aerial imagery, we calculate both F_1 (Wu and Zhou, 2017) and F_2 (Planet: Understanding the Amazon from space) score as follows:

$$F_{\beta} = (1 + \beta^2) \frac{p_e r_e}{\beta^2 p_e + r_e}, \beta = 1, 2, \quad (9)$$

where p_e is the example-based precision (Tsoumakas and Vlahavas, 2007) of predicted multiple labels, and r_e indicates the example-based recall. They are computed by:

$$p_e = \frac{TP_e}{TP_e + FP_e}, r_e = \frac{TP_e}{TP_e + FN_e}, \quad (10)$$

where TP_e , FP_e , and FN_e indicate the numbers of positive labels, which are predicted correctly (true positives) and incorrectly (false positives), and negative labels, which are incorrectly predicted (false negatives) in an example (i.e., an image with multiple object labels in our case), respectively. Then, the average of F_2 scores of each example is formed to assess the overall accuracy of multi-label classification tasks. Besides, example-based mean precision as well as mean recall are calculated to assess the performance from the perspective of examples, while label-based mean precision and mean recall can help us understand the performance of the network from the perspective of object labels:

$$p_l = \frac{TP_l}{TP_l + FP_l}, r_l = \frac{TP_l}{TP_l + FN_l}, \quad (11)$$

where TP_l , FP_l , and FN_l represent the numbers of correctly predicted positive images, incorrectly predicted positive images, and incorrectly predicted negative images with respect to each label.

For a fair validation of CA-Conv-BiLSTM, we decompose the evaluation into two components: we compare (1) CA-Conv-LSTM with standard CNNs to validate the effectiveness of employing LSTM-based recurrent sub-network, and (2) CA-Conv-BiLSTM with CA-Conv-LSTM for further assess the significance of the bidirectional structure. The

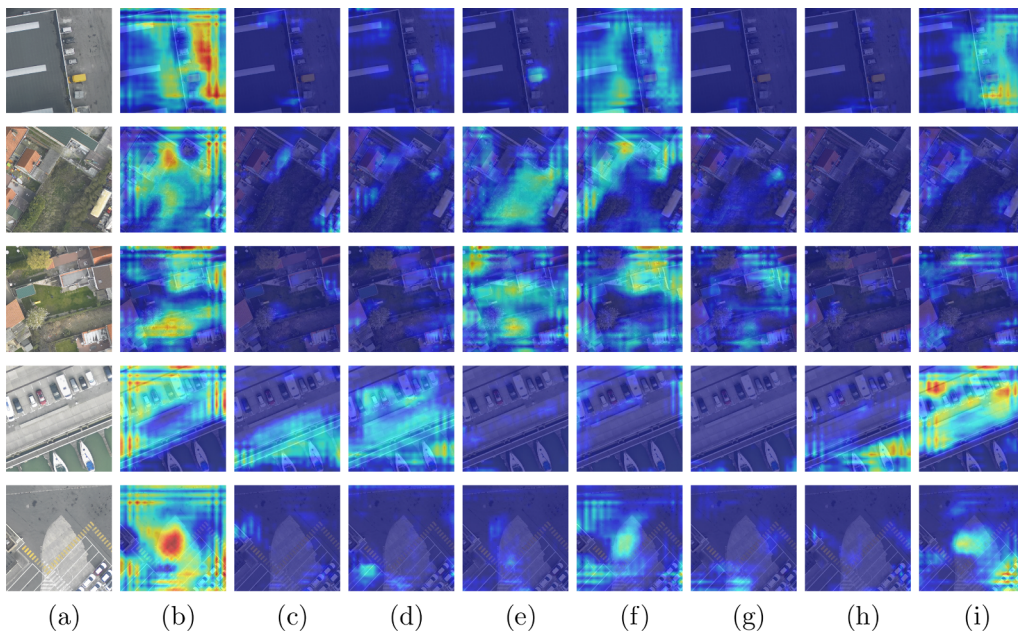


Fig. 9. Example class attention maps of (a) images in DFC15 dataset with respect to (b) impervious, (c) water, (d) clutter, (e) vegetation, (f) building, (g) tree, (h) boat, and (i) car. Red indicates strong activations, while blue represents non-activations. Besides, normalization is performed based on each row for a fair comparison among class attention maps of the same images. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

detailed configurations of these competitors are listed in Table 3. For standard CNNs, we substitute last softmax layers, which are designed for single-label classification, with sigmoid layers to predict multi-hot binary sequences, where each digit indicates the probability of the presence of its corresponding category. To calculate evaluation metrics, we binarize outputs of all models with a threshold of 0.5 for producing binary sequences. Besides, our model is compared with a relevant existing method (Zeggada et al., 2017) for a comprehensive evaluation of its performance.

Table 4 exhibits results on UCM multi-label dataset, and it can be seen that compared to directly applying standard CNNs to multi-label classification, CA-Conv-LSTM framework performs superiorly as expected due to taking class dependencies into consideration. CA-VGG-LSTM increases the mean F_1 score by 1.03% with respect to VGGNet, while for CA-ResNet-LSTM, an increment of 1.68%, is obtained compared to ResNet. Mostly enjoying this framework, CA-GoogLeNet-LSTM achieves the best mean F_1 score of 81.78% and an increment of 1.10% in comparison with other CA-Conv-LSTM models and GoogLeNet, respectively. Moreover, CA-ResNet-LSTM shows an improvement of 3.08% of the mean F_2 score in comparison with ResNet, while CA-GoogLeNet-LSTM obtains the best F_2 score of 85.16%. To summarize, all comparisons demonstrate that instead of directly using a standard CNN as a regression task, exploiting class dependencies plays a key role in multi-label classification.

Concerning the signification of employing a bidirectional structure, CA-Conv-BiLSTM performs better than CA-Conv-LSTM in the mean F_1 score, and compared to Conv-RBFNN, our models achieve higher mean F_1 and F_2 scores, increased by at most 0.98% and 2.80%, respectively. Another important observation is that our proposed model is equipped with higher example-based recall but lower example-based precision, which leads to a relatively higher mean F_2 score. Notably, the F_2 score is an evaluation index used in Kaggle Amazon contest (Planet: Understanding the Amazon from space) to assess the performance of recognizing challenging rare objects in aerial images, and a higher score indicates a stronger capability. Table 5 exhibits several example predictions in UCM multi-label dataset. Although our model successfully predicts most multiple object labels, it is observed that the grass and tree are prone to be misclassified due to their analogous appearances. In the 4th image, the grass is a false positive when there exist trees, while in the 5th image, the tree is a false positive when the grass presents. Likewise, the bare soil in the 5th image is neglected unfortunately for

its similar visual patterns with the grass.

3.3.2. Qualitative results

In addition to validate classification capabilities of the network by computing the mean F_2 score, we further explore the effectiveness of class-specific features learned from the proposed class attention learning layer and try to “open” the black box of our network by feature visualization. Example class attention maps produced by the proposed network on UCM multi-label dataset are shown in Fig. 8, where column (a) is original images, and columns (b)–(i) are class attention maps for different objects: (b) bare soil, (c) building, (d) car, (e) court, (f) grass, (g) pavement, (h) tree, and (i) water. As we can see, these maps highlight discriminative regions for positive classes, while present almost no activations when corresponding objects are absent in original images. For example, object labels of the image at the first row in Fig. 8 are building, grass, pavement, and tree, and its class attention maps for these categories are strongly activated. From images at the fourth row of Fig. 8, it can be seen that regions of the grassland, forest, and river are highlighted in their corresponding class attention maps, leading to positive predictions, while no discriminative areas are intensively activated in the other maps.

3.4. Results on DFC15 multi-label dataset

3.4.1. Quantitative results

Following the evaluation on UCM multi-label dataset, we assess our network on DFC15 multi-label dataset by calculating the mean F_1 and F_2 score as well as mean example- and label-based precision and recall. Table 6 shows experimental results on this dataset, and the conclusion can be drawn that modeling class dependencies with a bidirectional structure contributes significantly to multi-label classification. Specifically, the mean F_1 score achieved by CA-ResNet-BiLSTM is 4.87% and 5.55% higher than CA-ResNet-LSTM and ResNet, respectively. CA-VGG-BiLSTM obtains the best mean F_1 score of 76.25% in comparison with VGGNet and CA-VGG-LSTM, and the mean F_1 score of CA-GoogLeNet-BiLSTM is 78.25%, higher than its competitors. In comparison with Conv-RBFNN, CA-Conv-BiLSTM exhibits an improvement of at most 5.29% and 4.18% in terms of the mean F_1 and F_2 score, respectively. To conclude, all these increments demonstrate the effectiveness and robustness of our bidirectional structure for high-resolution aerial image multi-label classification. Several example predictions in DFC15 multi-

label dataset are shown in Table 5. The last two examples of DFC15 multi-label dataset show that trees are false negatives with the occurrence of vegetations due to their similar appearances. Moreover, we note the best result (Campos-Taberner et al., 2016) in 2015 IEEE GRSS Data Fusion Contest achieves 71.18% in the mean F1 score, which is reduced by 12.47% with respect to our best result. This is because predicting dense pixel-level labels is challenging in comparison with classifying multiple image-level labels.

3.4.2. Qualitative results

To study the effectiveness of class-specific features, we visualize class attention maps learned from the proposed class attention learning layer, as shown in Fig. 9. Columns (b)–(i) are example class attention maps with respect to (b) impervious, (c) water, (d) clutter, (e) vegetation, (f) building, (g) tree, (h) boat, and (i) car. As we can see, figures at column (b) of Fig. 9 show that the network pays high attention to impervious regions, such as parking lots, while figures at column (i) highlight regions of cars. However, some of class attention maps for negative object labels exhibit unexpected strong activations. For instance, the class attention map for the car at the third row of Fig. 9 is not supposed to highlight any region due to its absence of cars. This can be explained as the highlighted regions share similar patterns as cars, which also illustrates why the network made wrong predictions (cf. wrongly predicted car label in Fig. 9). Overall, the visualization of class attention maps demonstrates that the features captured from the proposed class attention learning layer are discriminative and class-specific. Besides, we note that there exist strong border artifacts in figures, especially those at column (b) of Fig. 9, which questions whether improving the quality of class attention maps benefits the effectiveness of the BiLSTM-based sub-network. Then we experimented with using the skip connection scheme in order to refine class attention maps. Experimental results demonstrated that this provides negligible improvements.

4. Conclusion

In this paper, we propose a novel network, CA-Conv-BiLSTM, for the multi-label classification of high-resolution aerial imagery. The proposed network is composed of three indispensable elements: (1) a feature extraction module, (2) a class attention learning layer, and (3) a bidirectional LSTM-based sub-network. Specifically, the feature extraction module is responsible for capturing fine-grained high-level feature maps from raw images, while the class attention learning layer is designed for extracting discriminative class-specific features. Afterwards, the bidirectional LSTM-based sub-network is used to model the underlying class dependency in both directions and predict multiple object labels in a structured manner. With such design, the prediction of multiple object-level labels is performed in an ordered procedure, and outputs are structured sequences instead of discrete values. We evaluate our network on two datasets, UCM multi-label dataset and DFC15 multi-label dataset, and experimental results validate the effectiveness of our model from both quantitative and qualitative respects. On one hand, the mean F_2 score is increased by at most 0.0446 compared to other competitors. On the other hand, visualized class attention maps, where discriminative regions for existing objects are strongly activated, demonstrate that features learned from this layer are class-specific and discriminative. Looking into the future, the application of our network can be extended to fields, such as weakly supervised semantic segmentation and object localization.

Acknowledgment

This work is jointly supported by the China Scholarship Council, the Helmholtz Association under the framework of the Young Investigators Group SiPEO (VH-NG-1018, www.sipeo.bgu.tum.de), and the European Research Council (ERC) under the European Union's Horizon 2020

research and innovation programme (grant agreement No. ERC-2016-StG-714087, Acronym: So2Sat). In addition, the authors would like to thank the National Center for Airborne Laser Mapping and the Hyperspectral Image Analysis Laboratory at the University of Houston for acquiring and providing the data used in this study, and the IEEE GRSS Image Analysis and Data Fusion Technical Committee.

References

- Audebert, N., Saux, B.L., Lefèvre, S., 2018. Beyond RGB: very high resolution urban remote sensing with multimodal deep networks. *ISPRS J. Photogram. Remote Sens.* 140 (June), 20–32.
- Badrinarayanan, V., Kendall, A., Cipolla, R., 2015. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. arXiv:1511.00561.
- Campos-Taberner, M., Romero-Soriano, A., Gatta, C., Camps-Valls, G., Lagrange, A., Saux, B.L., Beaupère, A., Boulch, A., Chan-Hon-Tong, A., Herbin, S., Randrianarivo, H., Ferecatu, M., Shimoni, M., Moser, G., Tuia, D., 2016. Processing of extremely high-resolution LiDAR and RGB data: outcome of the 2015 IEEE GRSS data fusion contest—part A: 2-D contest. *IEEE J. Select. Top. Appl. Earth Observ. Remote Sens.* 9 (12), 5547–5559.
- Chaudhuri, B., Demir, B., Chaudhuri, S., Bruzzone, L., 2018. Multilabel remote sensing image retrieval using a semisupervised graph-theoretic method. *IEEE Trans. Geosci. Remote Sens.* 56 (2), 1144–1158.
- Cheng, G., Han, J., Lu, X., 2017. Remote sensing image scene classification: benchmark and state of the art. *Proc. IEEE* 105 (10), 1865–1883.
- Chua, T.-S., Tang, J., Hong, R., Li, H., Luo, Z., Zheng, Y.-T., 2009. NUS-WIDE: A real-world web image database from National University of Singapore. In: *ACM Conference on Image and Video Retrieval (CIVR)*.
- Demir, B., Bruzzone, L., 2016. Histogram-based attribute profiles for classification of very high resolution remote sensing images. *IEEE Trans. Geosci. Remote Sens.* 54 (4), 2096–2107.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Dozat, T., 2016. Incorporating Nesterov momentum into Adam, <http://cs229.stanford.edu/proj2015/054_report.pdf>, online.
- Everingham, M., Van Gool, L., Williams, C., Winn, J., Zisserman, A., 2010. The Pascal visual object classes (VOC) challenge. *Int. J. Comput. Vision* 88 (2), 303–338.
- Gers, F., Schmidhuber, J., Cummins, F., 1999. Learning to forget: Continual prediction with LSTM, in: *International Conference on Artificial Neural Networks*.
- Graves, A., 2013. Generating sequences with recurrent neural networks, arXiv:1308.0850.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput.* 9 (8), 1735–1780.
- Hua, Y., Mou, L., Zhu, X.X., 2018. LAHNet: A convolutional neural network fusing low- and high-level features for aerial scene classification. In: *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*.
- Huang, X., Chen, H., Gong, J., 2018. Angular difference feature extraction for urban scene classification using ZY-3 multi-angle high-resolution satellite imagery. *ISPRS J. Photogram. Remote Sens.* 135, 127–141.
- Hu, F., Xia, G., Hu, J., Zhang, L., 2015. Transferring deep convolutional neural networks for the scene classification of high-resolution remote sensing imagery. *Remote Sens.* 7 (11), 14680–14707.
- Hu, F., Xia, G., Y.W., Z.L., 2018. Recent advances and opportunities in scene classification of aerial images with deep models. In: *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, 2018.
- 2015 IEEE GRSS data fusion contest. <<http://www.grss-ieee.org/community/technical-committees/data-fusion>>, online.
- Kang, J., Wang, Y., Körner, M., Taubenböck, H., Zhu, X.X., 2018. Building instance classification using street view images.
- Karalas, K., Tsagkatakis, G., Zervakis, M., Tsakalides, P., 2016. Land classification using remotely sensed data: going multilabel. *IEEE Trans. Geosci. Remote Sens.* 54 (6), 3548–3563.
- Koda, S., Zeggada, A., Melgani, F., Nishii, R., 2018. Spatial and structured SVM for multilabel image classification. *IEEE Trans. Geosci. Remote Sens.* 1–13.
- Li, Q., Mou, L., Liu, Q., Wang, Y., Zhu, X.X., 2018. Hsf-net: Multi-scale deep feature embedding for ship detection in optical remote sensing imagery. *IEEE Trans. Geosci. Remote Sens.*
- Lin, T., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S., 2017. Feature pyramid networks for object detection. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Lowe, D., 2004. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision* 60 (2), 91–110.
- Lucchesi, S., Giardino, M., Perotti, L., 2013. Applications of high-resolution images and DTMs for detailed geomorphological analysis of mountain and plain areas of NW Italy. *Eur. J. Remote Sens.* 46 (1), 216–233.
- Marcos, D., Volpi, M., Kellenberger, B., Tuia, D., 2018. Land cover mapping at very high resolution with rotation equivariant CNNs: towards small yet accurate models. *ISPRS*

- J. Photogram. Remote Sens.
- Marmanis, D., Schindler, K., Wegner, J.D., Galliani, S., Datcu, M., Stilla, U., 2018. Classification with an edge: improving semantic image segmentation with boundary detection. *ISPRS J. Photogram. Remote Sens.* 135 (January), 158–172.
- Mou, L., Zhu, X.X., 2016. Spatiotemporal scene interpretation of space videos via deep neural network and tracklet analysis. In: *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*.
- Mou, L., Zhu, X.X., 2018a. RiFCN: recurrent network in fully convolutional network for semantic segmentation of high resolution remote sensing images, arXiv:1805.02091.
- Mou, L., Zhu, X.X., 2018b. IM2HEIGHT: height estimation from single monocular imagery via fully residual convolutional-deconvolutional network, arXiv:1802.10249.
- Mou, L., Zhu, X.X., 2018c. Vehicle instance segmentation from aerial image and video using a multi-task learning residual fully convolutional network. *IEEE Trans. Geosci. Remote Sens.*
- Mou, L., Zhu, X.X., Vakalopoulou, M., Karantzalos, K., Paragios, N., Le Saux, B., Moser, G., Tuia, D., 2017. Multitemporal very high resolution from space: outcome of the 2016 IEEE GRSS data fusion contest. *IEEE J. Select. Top. Appl. Earth Observ. Remote Sens.* 10 (8), 3435–3447.
- Mou, L., Ghamisi, P., Zhu, X., 2017. Deep recurrent neural networks for hyperspectral image classification. *IEEE Trans. Geosci. Remote Sens.* 55 (7), 3639–3655.
- Mou, L., Bruzzone, L., Zhu, X.X., 2019. Learning spectral-spatial-temporal features via a recurrent convolutional neural network for change detection in multispectral imagery, arXiv:1803.02642.
- Nogueira, K., Penatti, O., dos Santos, J., 2017. Towards better exploiting convolutional neural networks for remote sensing scene classification. *Pattern Recogn.* 61, 539–556.
- Patterson, G., Hays, J., 2012. Sun attribute database: discovering, annotating, and recognizing scene attributes. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Planet: Understanding the Amazon from space, <<https://www.kaggle.com/c/planet-understanding-the-amazon-from-space#evaluation>>, online.
- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster R-CNN: towards real-time object detection with region proposal networks. In: *Advances in Neural Information Processing Systems*.
- Ren, S., He, K., Girshick, R., Zhang, X., Sun, J., 2017. Object detection networks on convolutional feature maps. *IEEE Trans. Pattern Anal. Mach. Intell.* 39 (7), 1476–1481.
- Risojevic, V., Babic, Z., 2013. Fusion of global and local descriptors for remote sensing image classification. *IEEE Geosci. Remote Sens. Lett.* 10 (4), 836–840.
- Shao, W., Yang, W., Xia, G., Liu, G., 2013. A hierarchical scheme of multiple feature fusion for high-resolution satellite scene categorization. In: *International Conference on Computer Vision Systems*.
- Simonyan, K., Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition, arXiv:1409.1556.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2016. Rethinking the inception architecture for computer vision. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tan, Q., Liu, Y., Chen, X., Yu, G., 2017. Multi-label classification based on low rank representation for image annotation. *Remote Sens.* 9 (2), 109.
- Tsoumakas, G., Vlahavas, I., 2007. Random k-labelsets: An ensemble method for multi-label classification. In: *European Conference on Machine Learning*.
- Viola, P., Jones, M., 2001. Rapid object detection using a boosted cascade of simple features. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wen, D., Huang, X., Liu, H., Liao, W., Zhang, L., 2017. Semantic classification of urban trees using very high resolution satellite imagery. *IEEE J. Select. Top. Appl. Earth Observ. Remote Sens.* 10 (4), 1413–1424.
- Weng, Q., Mao, Z., Lin, J., Liao, X., 2018. Land-use scene classification based on a CNN using a constrained extreme learning machine. *Int. J. Remote Sens.* 0 (0), 1–19.
- Wu, X., Zhou, Z., 2017. A unified view of multi-label performance measures, arXiv:1609.00288.
- Xia, G., Hu, J., Hu, F., Shi, B., Bai, X., Zhong, Y., Zhang, L., Lu, X., 2017. AID: a benchmark data set for performance evaluation of aerial scene classification. *IEEE Trans. Geosci. Remote Sens.* 55 (7), 3965–3981.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., Bengio, Y., 2015. Show, attend and tell: neural image caption generation with visual attention. In: *International Conference on Machine Learning*.
- Yang, Y., Newsam, S., 2010. Bag-of-visual-words and spatial extensions for land-use classification. In: *International Conference on Advances in Geographic Information Systems (SIGSPATIAL)*.
- Zarco-Tejada, P., Diaz-Varela, R., Angileri, V., Loudjani, P., 2014. Tree height quantification using very high resolution imagery acquired from an unmanned aerial vehicle (UAV) and automatic 3D photo-reconstruction methods. *Eur. J. Agron.* 55, 89–99.
- Zeggada, A., Melgani, F., Bazi, Y., 2017. A deep learning approach to UAV image multilabeling. *IEEE Geosci. Remote Sens. Lett.* 14 (5), 694–698.
- Zeggada, A., Benbraika, S., Melgani, F., Mokhtari, Z., 2018. Multilabel conditional random field classification for UAV images. *IEEE Geosci. Remote Sens. Lett.* 15 (3), 399–403.
- Zhang, F., Du, B., Zhang, L., 2015. Saliency-guided unsupervised feature learning for scene classification. *IEEE Trans. Geosci. Remote Sens.* 53 (4), 2175–2184.
- Zhu, Q., Zhong, Y., Zhao, B., Xia, G.S., Zhang, L., 2016. Bag-of-visual-words scene classifier with local and global features for high spatial resolution remote sensing imagery. *IEEE Geosci. Remote Sens. Lett.* 13 (6), 747–751.
- Zhu, X.X., Tuia, D., Mou, L., Xia, S., Zhang, L., Xu, F., Fraundorfer, F., 2017. Deep learning in remote sensing: a comprehensive review and list of resources. *IEEE Geosci. Remote Sens. Magaz.* 5 (4), 8–36.