# Tree species classification from hyperspectral data using graph-regularized neural networks

Debmita Bandyopadhyay, Subhadip Mukherjee, James Ball, Grégoire Vincent, David A. Coomes, and Carola-Bibiane Schönlieb

*Abstract*—We propose a novel graph-regularized neural network (GRNN) algorithm for tree species classification. The proposed algorithm encompasses superpixel-based segmentation for graph construction, a pixel-wise neural network classifier, and the label propagation technique to generate an accurate and realistic (emulating tree crowns) classification map on a sparsely annotated data set. GRNN outperforms several state-of-the-art techniques not only for the standard Indian Pines HSI but also achieves a high classification accuracy (approx. 92%) on a new HSI data set collected over the heterogeneous forests of French Guiana (FG) when less than 1% of the pixels are labeled. We further show that GRNN is competitive with the state-of-the-art semi-supervised methods and exhibits a small deviation in accuracy for different numbers of training samples and over repeated trials with randomly sampled labeled pixels for training.

*Index Terms*—Hyperspectral image, forest species classification, neural networks, superpixel graph.

## I. Introduction

TREE species classification of forests is important as several commercially important and vulnerable species are fast depleting. Using non-invasive technologies, like aerial hyperspectral sensors, tree species can be distinguished at scale even with very similar spectral curves. Most of the studies on tree-species classification using HSI data have either employed pixel-level supervised classification methods (e.g., support vector machines (SVMs), linear discriminant analysis, or random forests) or utilized crown cover boundaries to provide realistic classification maps [1, 2, 3]. However, these techniques rely heavily on apriori information with limited capability to extract spectral-spatial features from the raw data. Deep learning-based models like CNNs, RNNs and GANs can learn from the spatial texture information [4] but have a high probability of misalignment between the network's objectives of achieving high classification accuracy and providing a realistic tree species classified image. The resultant image is highly accurate but unrealistically classified, especially for the unlabeled pixels in the image [5].

For tropical forests, apart from having spectral similarity between species but heterogeneity in species distribution, there is an additional challenge of the limited availability of manual labels owing to the inaccessibility of these regions and the cost of establishing forest plots. Semi-supervised graph learning

can tackle the problem of sparsely annotated data labels and extract useful information from labeled and unlabeled pixels in an image [6]. Hence, they are a preferred choice over deep learning methods as the latter depends on high proportions of annotated data. Besides, hyperspectral data involves dealing with high-resolution and high-dimensional spectral-spatial data sets. Various feature extraction methods like the image fusion and recursive filter by Kang et al. [7], utilization of local binary patterns by Li et al. [8], local matrix representation by Fang et al. [9] and the hyper-manifold simple linear iterative clustering combined with the graph learning technique by Sellars et al. [10] have been shown to successfully address this high-dimensionality problem. Additionally, the superpixel-based classification emulates the tree crown, providing a realistic classification map. In summary, the superpixel-based semi-supervised method deals with the (i) limited ground information challenge by clustering pixels with similar features, (ii) maintains a competitive accuracy alongside the state-of-the-art techniques, and (iii) provides a realistic (tree-crown scale) classification map. However, to preserve the pixel-based information and augment label propagation to unlabeled pixels, a neural net (NN)-based classifier is also required owing to their expressive capacity and the ability to generate an accurate pixel-level classification. **Our contributions**: We propose a novel graph-regularized neural network (GRNN) classifier, which is built on the following three mechanisms: (1) First, to combine pixel-based classification and superpixel feature-based learning to extract maximum information from both labeled and unlabeled pixels in the HSI, (2) Second, to generate an accurate and realistic classification map by jointly optimizing the accuracy of a pixel-level NN and a superpixel-level graph energy potential that seeks to promote smoothness in the predicted classification map, and (3) third, the NN output predictions (exceeding certain confidence) are used to augment the ground-truth labels, on which we apply the classical semi-supervised label propagation method [11] using the superpixel graph to generate the final classification map. To the best of our knowledge, GRNN is the first framework that combines neural networks with superpixel graphs for the semi-supervised classification of tree species. For a data set where the tree species have similar spectral responses, the proposed GRNN algorithm (i) can accurately classify the top ten most commonly occurring species, (ii) provide a smooth superpixel-based classification map with a realistic view of the tree crowns, and (iii) works well even with $< 1\%$ of the pixels being annotated. Details of GRNN and our experimental findings are reported in Sections II and III, respectively.

## II. PROPOSED METHOD

The key idea behind the GRNN classification approach is to use a pixelwise fully-connected NN classifier together with a superpixel graph to learn effectively from a few labeled pixels and a large number of unlabeled pixels. The output probability distribution of the NN is regularized using a graph energy function that smooths variation of the NN output. The details of the GRNN algorithm are explained in the following, and a schematic of the entire pipeline is shown in Figure 1.

The given HSI $X \in \mathbb{R}^{H \times W \times B}$, having spatial dimensions $H \times W$ and $B$ spectral bands, is first reduced via principal component analysis (PCA) to $x \in \mathbb{R}^{H \times W \times b}$, where $b \ll B$. The number of bands $b$ is chosen such that 99.90% of variance in the HSI data is preserved. The labeled and unlabeled pixels in the given HSI are denoted by $\mathcal{L}$ and $\mathcal{U}$, respectively. Following PCA, we apply the *simple linear iterative clustering*
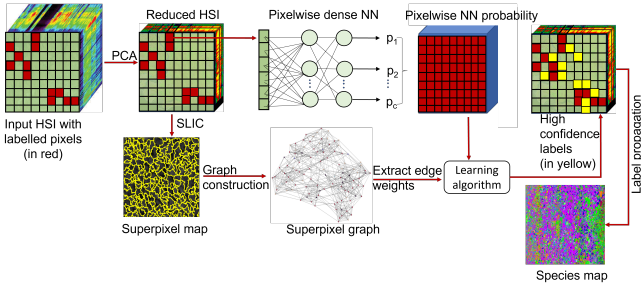


Fig. 1: Framework for the proposed method, GRNN.

(SLIC) algorithm [12] to the first principal component of $X$ to extract $N$ superpixels, denoted as $\{\mathcal{S}_k\}_{k=1}^{N}$, which are subsequently utilized to construct a graph. An essential step for superpixel graph construction is *feature extraction*, for which we follow the procedure outlined in [10, Sec. III.B]. The nodes of the graph represent the superpixels, and the adjacency weights $W_{k\ell}$ capture pairwise similarity between superpixels $\mathcal{S}_k$ and $\mathcal{S}_\ell$ in terms of the extracted features. We follow the same procedure as reported in [10] for constructing the graph adjacency matrix $W \in \mathbb{R}^{N \times N}$. Small adjacency weights are set to 0 by retaining only the top $N_s$ elements in each row of $W$. This thresholding step sparsifies $W$ and leads to efficient training. We choose $N_s = 20$ in our experiments, i.e., each node can have at most 20 neighboring nodes. A superpixel $\mathcal{S}_k$ is said to be labeled if it contains at least one labeled pixel in it, i.e., $\mathcal{S}_k \cap \mathcal{L} \neq \emptyset$, and the set of all labeled superpixels is denoted as $\mathcal{L}^{\mathrm{sp}}$. For a $c$-class classification problem, let $y_{\mathbf{i}} \in \{1, 2, \cdots, c\}$ be the true label of any pixel $\mathbf{i} = (i_1, i_2) \in \mathcal{L}$. The one-hot encoding of $y_{\mathbf{i}}$ is denoted by $Y_{\mathbf{i}} \in \mathbb{R}^c$, and is defined as follows:

$$Y_{\mathbf{i}}(q) = \begin{cases} 1 & \text{if } y_{\mathbf{i}} = q, \text{ where } q \in \{1, 2, \cdots, c\}, \text{ and} \\ 0 & \text{if } y_{\mathbf{i}} \neq q. \end{cases} \quad (1)$$

For any labeled superpixel $\mathcal{S}_k \in \mathcal{L}^{\mathrm{sp}}$, we compute two labels, referred to as the *soft* and the *hard* labels, respectively. The soft label $t_{\mathcal{S}_k}$ represents the relative frequency of different classes within it, and is computed as $t_{\mathcal{S}_k} = \frac{1}{|\mathcal{S}_k \cap \mathcal{L}|} \sum_{\mathbf{j} \in \mathcal{S}_k \cap \mathcal{L}} Y_{\mathbf{j}}$. Note that for any $\mathcal{S}_k \in \mathcal{L}^{\mathrm{sp}}$, the soft label $t_{\mathcal{S}_k} \in \Delta_c$, where $\Delta_c$ is the

probability simplex in $\mathbb{R}^c$. The hard label $T_{\mathcal{S}_k}$ corresponding to $\mathcal{S}_k$ represents the most frequent class in $\mathcal{S}_k$ using a one-hot encoding. Specifically, $T_{\mathcal{S}_k} \in \mathbb{R}^c$ is a $c$-length vector whose $q_*^{\mathrm{th}}$ element is 1, and every other element is 0, where $q_* = \arg\max_{q \in \{1, 2, \cdots, c\}} t_{\mathcal{S}_k}(q)$. For an unlabeled superpixel $\mathcal{S}_k \notin \mathcal{L}^{\mathrm{sp}}$, both $t_{\mathcal{S}_k}$ and $T_{\mathcal{S}_k}$ are defined as $c$-length vectors with all zeros. Let, $T^{(\mathcal{L})}$ be the $N \times c$ matrix whose $k^{\mathrm{th}}$ row is $T_{\mathcal{S}_k}$. The superscript $\mathcal{L}$ is used to indicate explicitly that $T$ depends on the underlying set of labeled pixels $\mathcal{L}$.

Besides the superpixel graph, the other key component of GRNN is a pixel-wise, fully-connected NN-based classifier $\phi_\theta : \mathbb{R}^b \mapsto \Delta_c$ with parameters $\theta$. The NN has two intermediate layers with leaky-ReLU activations (with negative slope 0.1)[1] and a softmax layer in the end. The output of the NN gives a probability map over $c$ classes, i.e., for any pixel location $\mathbf{i} \in \mathbb{R}^2$, $\phi_\theta(x_{\mathbf{i}}) \in \Delta_c$ represents the predicted probabilities of the pixel belonging to any of the $c$ classes. The predicted probability map for a superpixel is computed by averaging the probability maps of all the pixels within the given superpixel: $\phi_{\mathcal{S}_k} = \frac{1}{|\mathcal{S}_k|} \sum_{\mathbf{j} \in \mathcal{S}_k} \phi_\theta(x_{\mathbf{j}})$. The overall training loss seeks to achieve the following objectives simultaneously: (i) minimize classification errors, both for the pixels and the superpixels for which the ground-truth labels are available, (ii) ensure that the predicted superpixel probability map varies smoothly over the graph, (iii) minimize the standard deviation in the predicted pixelwise probability maps within each superpixel, and (iv) discourage uniform labeling of the unlabeled superpixels, i.e., make sure that all unlabeled superpixels do not end up with the same class and that all classes are represented overall. To accomplish these objectives, the training problem is formulated as

$$\min_\theta \left[ \sum_{\mathbf{j} \in \mathcal{L}} \mathcal{D}_1 \left( Y_{\mathbf{j}}, \phi_\theta(x_{\mathbf{j}}) \right) + \lambda_{\mathrm{spc}} \sum_{k:\mathcal{S}_k \in \mathcal{L}^{\mathrm{sp}}} \mathcal{D}_2 \left( t_{\mathcal{S}_k}, \phi_{\mathcal{S}_k} \right) \right.$$
$$+ \lambda_{\mathrm{g}} \sum_{\mathcal{S}_k, \mathcal{S}_\ell} \left\| \frac{\phi_{\mathcal{S}_k}}{\sqrt{d_k}} - \frac{\phi_{\mathcal{S}_\ell}}{\sqrt{d_\ell}} \right\|_2^2 + \lambda_{\mathrm{v}} \sum_{\mathcal{S}_k} \mathrm{var} \left\{ \phi_\theta(x_{\mathbf{j}}) \right\}_{\mathbf{j} \in \mathcal{S}_k}$$
$$\left. - \lambda_{\mathrm{en}} \mathcal{H} \left( \frac{1}{N} \sum_{k=1}^{N} \phi_{\mathcal{S}_k} \right) \right]. \quad (2)$$

Here, $d_k = \sum_{\ell:\ell \neq k} W_{k\ell}$ denotes the degree of node-$k$ of the superpixel graph, $\mathcal{D}_1$ and $\mathcal{D}_2$ measure pixel- and superpixel-wise classification losses, respectively, and $\mathcal{H} : \Delta_c \mapsto \mathbb{R}$ denotes the entropy functional. The first two terms in (2) promote high classification accuracy on the labeled pixels and superpixels, respectively. We choose $\mathcal{D}_1$ to be the cross-entropy loss, while $\mathcal{D}_2$ is taken as the standard squared Euclidean distance. The third term is a graph energy function that enforces a smoothness prior on the predicted superpixel probability map, ensuring that it does not vary abruptly as a function defined on the superpixel graph. The fourth term in the training loss seeks to minimize the variance of prediction probabilities for all pixels inside each superpixel, which promotes high intra-superpixel similarity of the prediction. The final negative entropy term penalizes uniform

---

[1]leaky-relu$(u) = \max(0, u) + \gamma \min(0, u)$, where $\gamma$ = negative slope.

labeling of the unlabeled superpixels, thereby encouraging high inter-superpixel variability of the predicted classification map (which is well-suited for the highly heterogeneous French Guiana data considered in Sec. III). The $\lambda$ parameters in the overall training loss trade off different components of the loss function and need to be selected optimally. The training problem (2) is solved iteratively using the Adam optimizer [13], the de facto standard for training NNs, with a learning rate of $10^{-3}$ and $(\beta_1, \beta_2) = (0.9, 0.999)$.

The trained NN, although optimized simultaneously for accuracy and low intra-superpixel variability of classification, can still produce a classification map with differently labeled pixels inside a given superpixel. Such a prediction might be physically unrealistic since the superpixels are generated based on spectral similarity and any given superpixel should ideally belong to a single tree species. To avoid this, the final step of GRNN generates a smooth label map by applying the classical label propagation algorithm for semi-supervised learning on graphs [11], after combining the NN prediction with the ground-truth labels. The label propagation step is carried out by augmenting the given ground-truth labels with the labels predicted by the NN with high confidence. The set of pixels for which the trained NN classifies with confidence exceeding a given threshold $\tau \in (0, 1)$ is defined as

$$\mathcal{I} = \{\mathbf{i} : \max\left(\phi_{\theta^*}(x_\mathbf{i})\right) \geq \tau\}, \tag{3}$$

where $\theta^*$ is the optimal NN parameters obtained by solving (2). In other words, if the probability of the most likely class is greater than $\tau$ (we set $\tau = 0.4$) then the corresponding predicted label is added to the existing ground-truth labels. The smoothed superpixel probability map is generated by

$$T^* = \left(\text{Id} - \alpha\, D^{-\frac{1}{2}} W D^{\frac{1}{2}}\right)^{-1} T^{(\mathcal{M})}, \tag{4}$$

where Id denotes the $N \times N$ identity matrix, $T^{(\mathcal{M})}$ is the $N \times c$ one-hot superpixel probability matrix computed using the augmented set of labeled pixels $\mathcal{M} = \mathcal{L} \cup \mathcal{I}$, and $W \in \mathbb{R}^{N \times N}$ and $D \in \mathbb{R}^{N \times N}$ are the graph adjacency and the (diagonal) degree matrices, respectively. For any superpixel $\mathcal{S}_k$, the final classification label is produced as $q_k = \arg\max\limits_{q \in \{1, 2, \cdots, c\}} T^*_{\mathcal{S}_k}(q)$, where $T^*_{\mathcal{S}_k}$ is the $k^{\text{th}}$ row of $T^*$. All pixels within a superpixel are assigned the same class label as their parent superpixel.

## III. Experiments and Results

*1) Data description/set up:* The widely used Indian Pines (IP) HSI benchmark data set is used to compare the proposed method with the existing state-of-the-art (SoTA) classification techniques comprising 16 classes. The data was collected over an agricultural site by the airborne visible/infrared imaging spectrometer (AVIRIS) sensor. We chose the IP data for performance comparison and benchmarking since all the competing SoTA methods were previously applied to this data and optimized for their respective best performances.

To check the robustness of GRNN, we apply it to a new data set collected with Hyspex VNIR 1600 and Hyspex SWIR 384 sensor-mounted alongside Riegl scanner. The study site is Paracou Research Station (5°18'N, 52°55'W) in French Guiana (FG), which is in the tropical rainforest. The data set comprises 76 tree species [14]. The ground truth labels were manually confirmed from the field survey, and the crowns were manually segmented. Only the top 10 most dominant species (comprising 6% of the total number of pixels) are considered for classification, and the other species were excluded (not assigned any label). Our experiment is performed in two stages. First, we develop and test the algorithm on the benchmark IP HSI data and compare the overall accuracy (OA) and the Cohen kappa coefficient ($\kappa$) with competing classification methods. The mean and standard deviation of classification accuracy were evaluated over ten random trials for this benchmark HSI data. Subsequently, we test GRNN on the large, highly heterogeneous FG HSI data with very limited ground-truth labels. Among the labeled pixels, the test and training samples were randomly chosen to estimate the OA and the $\kappa$ score.

*2) Hyperparameter selection:* GRNN involves two sets of hyperparameters, one for graph construction and the other for training. The optimal selection of these hyperparameters is crucial for the success of GRNN. The specific values chosen for the hyperparameters corresponding to two different data sets are reported in Table I. We chose these values based on cross-validation on a small validation set kept aside from the labeled data used for training. A more principled approach to learning these hyperparameters from the data would be highly desirable, and we leave it as future work.

TABLE I: The hyperparameters for graph construction and GRNN (chosen through cross-validation). $h$, $\beta$, $\sigma_s$ and $\sigma_l$ are as defined in (5), (8), and (9) of [10]. $\xi =$ connectivity defining superpixel adjacency. $n_\theta =$ total number of NN parameters. $N =$ number of superpixels extracted using SLIC, the $\lambda$ parameters control different components of the training objective (2), $\alpha =$ label-smoothing parameter in (4), $\tau =$ threshold in (3), and $n_{\text{iter}} =$ # training iterations.

| superpixel graph construction and NN parameters | | | | | | |
|---|---|---|---|---|---|---|
| | $h$ | $\beta$ | $\sigma_s$ | $\sigma_l$ | $\xi$ | $n_\theta$ |
| Indian Pines | 15.0 | 0.9 | 2.0 | 1.0 | 8 | 73888 |
| French Guiana | 15.0 | 0.5 | 5.0 | 40.0 | 8 | 369778 |

| GRNN parameters | | | | | | |
|---|---|---|---|---|---|---|
| | $N$ | $\lambda_{\text{spc}}$ | $\lambda_{\text{g}}$ | $\lambda_{\text{v}}$ | $\lambda_{\text{en}}$ | $\alpha$ | $n_{\text{iter}}$ |
| Indian Pines | 1200 | 0.15 | $10^5$ | 2.0 | 20.0 | 0.5 | 500 |
| French Guiana | 5000 | 0.1 | 0.2 | 0.1 | 20.0 | 0.5 | 400 |

*3) Comparison with different methods:* We compare and benchmark the performance of the proposed GRNN framework with respect to several SoTA methods, namely superpixel graph learning (SGL) [10], the local covariance matrix representation (LCMR) [9], superpixel-based classification via multiple kernels (SC-MK) [15], edge-preserving filter (EPF) [16], local binary pattern (LBP) method [8], the image fusion and recursive filtering (IFRF) [7], and the traditional SVM [1].

GRNN performs both qualitatively (Figure 2) and quantitatively (Table II) better than the existing classifiers for the Indian Pines data set. Specifically, the difference in $\kappa$ from the nearest competitive algorithm SGL (which is also based on superpixel segmentation and graph-based semi-supervised label propagation, but does not use any NN) is about 0.08. The key difference between SGL and GRNN is the NN-assisted graph learning which utilizes both pixel- and superpixel-level

TABLE II: Comparison of different methods on Indian Pines in terms of overall accuracy and the $\kappa$ score, corresponding to ten samples per class. GRNN performed the best (highlighted in boldface).

|  | GRNN (Ours) | SGL | LCMR | SC-MK | EPF | LBP | IFRF | SVM |
|---|---|---|---|---|---|---|---|---|
| OA (%) | **96.3 ± 0.8** | 90.7 ± 2.2 | 82.7 ± 3.1 | 80.7 ± 2.5 | 67.3 ± 3.2 | 78.9 ± 2.7 | 80.3 ± 1.8 | 53.1 ± 1.9 |
| $\kappa$ | **0.96 ± 0.01** | 0.88 ± 0.03 | 0.81 ± 0.03 | 0.77 ± 0.03 | 0.65 ± 0.02 | 0.78 ± 0.02 | 0.78 ± 0.04 | 0.48 ± 0.02 |



| (a) RGB | (b) ground-truth | (c) SVM | (d) SGL | (e) GRNN |

Legend:
- Alfalfa
- Corn-notill
- Corn-min
- Corn
- Grass/pasture
- Grass/Trees
- Grass/pasture-mowed
- Hay-windrowed
- Oats
- Soybeans-notill
- Soybeans-min
- Soybeans-clean
- Wheat
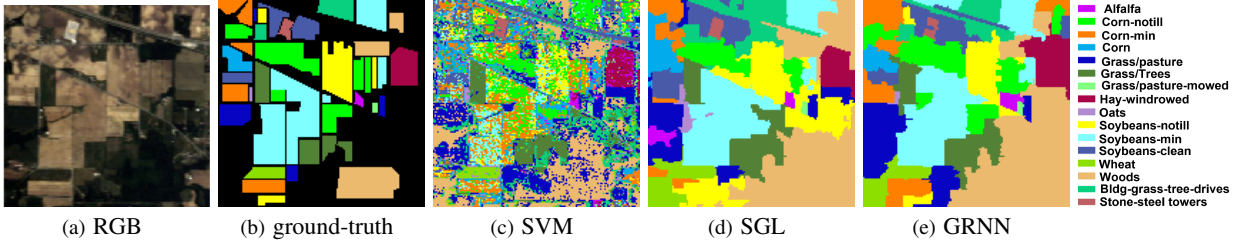- Woods
- Bldg-grass-tree-drives
- Stone-steel towers

Fig. 2: Comparison of SVM (supervised), SGL (semi-supervised), and the proposed GRNN classification maps on Indian Pines with ten samples per class. Qualitatively, GRNN produces a smoothed classification map with a higher accuracy as reported in Table II.

information for efficient semi-supervised classification while not compromising on pixelwise classification accuracy. The high-confidence NN predictions are used to augment the scarce ground-truth labels, and a visibly smooth classification is generated using the label propagation approach [10, 11].

To study the effect of the number of labeled samples on the overall accuracy, we compared GRNN with one of the best-performing semi-supervised algorithms (namely, SGL) and the classical supervised method of SVM (see Figure 1 in the supplementary document). We found that the OA in all three methods consistently increases as the number of samples increases from three to fifteen. However, GRNN does not exhibit a noticeable change in accuracy (varies between 95% and 96%) with a reduction in the number of training samples. The deviation in OA for GRNN is also significantly smaller in comparison with the other methods. The standard deviation (SD) for GRNN varies from 0.58 to 0.80, whereas the SD is as high as 3.85 for SGL and 5.30 for SVM.

Subsequently, we extend the study to a new data set from French Guiana for which the data has very limited labeled information (with a smaller fraction of labeled pixels than the standard IP data). The FG data set has highly heterogeneous classes which change as one moves from one crown to another. Further, there is an imbalance in class representation, and we, therefore, select the ten most commonly occurring classes to train the model and generate the classification map. The GRNN algorithm, after hyperparameter tuning (Table I), is trained using 10%, 30%, 50% and 70% of the labeled data. The 10% of the labeled data constitutes $\leq$ 1% of the total number of pixels in the HSI. Even with such few labeled pixels, the algorithm achieves a $\kappa$ of 0.79. As the training sample percentage increases, $\kappa$ improves up to 0.89. However, the performance plateaus beyond 50% training data usage. The experiments indicate that GRNN is particularly suitable when very few pixels are labeled. We compare the performance of GRNN with SVM (as a baseline pixelwise classifier) and the SGL method on the FG data (Table III). Notably, SGL (the closest competitor of GRNN on IP data) and SVM work solely at the superpixel- and pixel levels, respectively. The com-

parison demonstrates the shortcomings of methods that work purely on either the pixels or the superpixels and underscores the efficacy of a method such as GRNN that combines pixel-level classification with a superpixel-level regularization. Such a hybrid approach has the potential to achieve a good balance between high overall pixel-wise classification accuracy and the ability to accurately capture the tree crowns, as indicated by the final classification maps shown in Fig. 5. The superpixel graph regularization also prevents the NN from overfitting when trained on a limited number of labeled pixels and leads to better generalization.

The success of GRNN can be attributed to its ability to apply an effective regularization on the classification loss. Generating a classification map starting from only a few labeled pixels is a highly ill-posed problem, i.e., as long as the classifier correctly identifies the species on the pixels for which the ground-truth information is available, one would obtain a high accuracy regardless of what the classifier does on the unlabeled pixels, potentially producing an unrealistic classification overall (see the SGL output in Fig. 5(c)). However, our prior knowledge about the species distribution dictates that the unlabeled pixels cannot be labeled arbitrarily and must satisfy certain smoothness conditions. This prior information is encoded by using a superpixel graph and adding a penalty that encourages the classification map to vary smoothly over this graph. The high expressive capacity of a fully-connected neural network, combined with this strong regularization, allows GRNN to achieve the right balance between accuracy over pixels with known labels and smoothness of classification over a large number of unlabeled pixels.

## IV. CONCLUSIONS

We propose a novel algorithm based on superpixel graph energy minimization and a NN classifier for HSI pixel classification. The algorithm is validated on the standard Indian Pines region and shows an improved performance (varying between 6% and 16% increase in OA) compared to other state-of-the-art semi-supervised methods. The proposed GRNN method is also applied to a new high-resolution highly heterogeneous
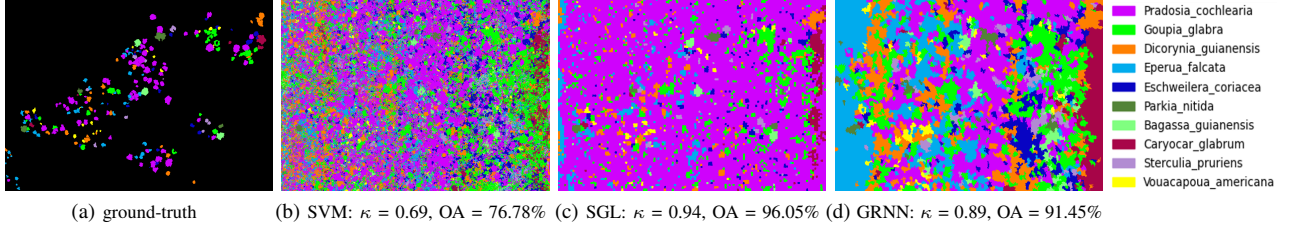
Fig. 3: SVM, SGL and GRNN applied to the FG data: (a) shows the labeled tree species. All three classification maps are shown: (b) SVM (c) SGL and (d) GRNN, with 50% of the available labeled pixels used for training. Although SGL attains a higher classification accuracy, it overwhelmingly assigns the most frequent species in the training data to almost all unlabeled pixels.

TABLE III: SVM and SGL vs. GRNN (OA, $\kappa$) on the FG data with 10%, 30%, 50% and 70% of the available labeled pixels. Classification maps for all results in Table III are presented in Figure 2 of the supplementary document. Despite higher pixelwise accuracy, SGL suffers from the propensity of generating an oversmoothed and unrealistic classification, while SVM produces a classification map that does not faithfully capture the tree crowns. (c.f. Fig. 5).

|  | 10% | 30% | 50% | 70% |
|---|---|---|---|---|
| SVM | 73.99%, 0.65 | 76.22%, 0.68 | 76.78%, 0.69 | 77.44%, 0.70 |
| SGL | 92.50 %, 0.90 | 95.61%, 0.94 | 96.05%, 0.94 | 96.26%, 0.95 |
| GRNN | 84.62%, 0.79 | 89.73%, 0.86 | 91.45%, 0.89 | 90.56%, 0.88 |

HSI data set and is consistent in performance (around 85% OA) using just 10% of the training samples (which is $< 1\%$ of the total pixels). GRNN can learn from a limited labeled data set by optimizing the neural network accuracy with a superpixel graph-based energy that imposes a strong prior on the predicted classification map. In addition to the enhanced accuracy, the variance in the estimates is also significantly less compared to other methods. The reduced variance highlights the robustness of the GRNN algorithm with respect to the randomly selected samples for training. On the new, sparsely annotated, and highly heterogeneous FG data, GRNN combines information at both pixel- and superpixel-levels, thereby achieving high pixelwise classification accuracy and accurate tree-crown identification simultaneously. In future work, we intend to explore automatic feature learning from the extracted superpixels instead of using handcrafted features to enhance the performance of GRNN. A more extensive comparison of GRNN with deep-learning-based SoTA classification methods, will also be undertaken in a future study.

## Supplementary material

This supplementary document contains additional experimental results that could not be accommodated in the main body of the paper.

In the first experiment, we compare the overall accuracy (OA) of GRNN with that of SVM (pixel-wise classifier) and

SGL (superpixel-level classifier) on the Indian Pines (IP) data for different numbers of labeled pixels per class in training (see Fig. 4). GRNN attains a higher classification accuracy over both SVM and SGL and shows little variance in the overall classification accuracy over 10 random trials.
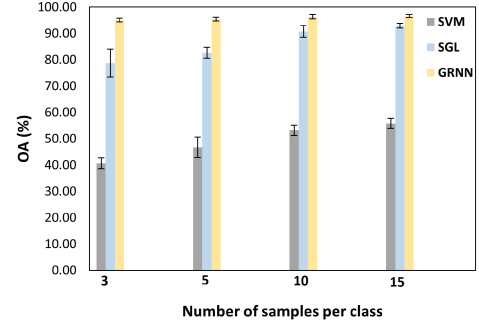


Fig. 4: Comparison of the OA of SVM, SGL and the proposed GRNN-based classification on IP data. The error bar represents SD in accuracy after ten iterations on the test samples. GRNN is more robust with higher accuracy as compared to SGL.

In Figure 5, we show a comparison of GRNN with SVM and SGL, two representative pixel- and superpixel-based classification algorithms, for different fractions of the available labeled pixels used for learning. Notably, the choice for SGL as a competing method is motivated by the fact that SGL was found to be the closest competitor of GRNN on the Indian Pines data.

### References

[1] F. Melgani and L. Bruzzone, "Classification of hyperspectral remote sensing images with support vector machines," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 42, no. 8, pp. 1778–1790, 2004.

[2] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[3] A. Laybros, M. Aubry-Kientz, J.-B. Féret, C. Bedeau, O. Brunaux, G. Derroire, and G. Vincent, "Quantitative airborne inventories in dense tropical forest using imaging spectroscopy," *Remote Sensing*, vol. 12, no. 10, 2020. [Online]. Available: https://www.mdpi.com/2072-4292/12/10/1577

[4] K. Safari, S. Prasad, and D. Labate, "A multiscale deep learning approach for high-resolution hyperspectral image classification," *IEEE Geoscience and Remote Sensing Letters*, vol. 18, no. 1, pp. 167–171, 2021.
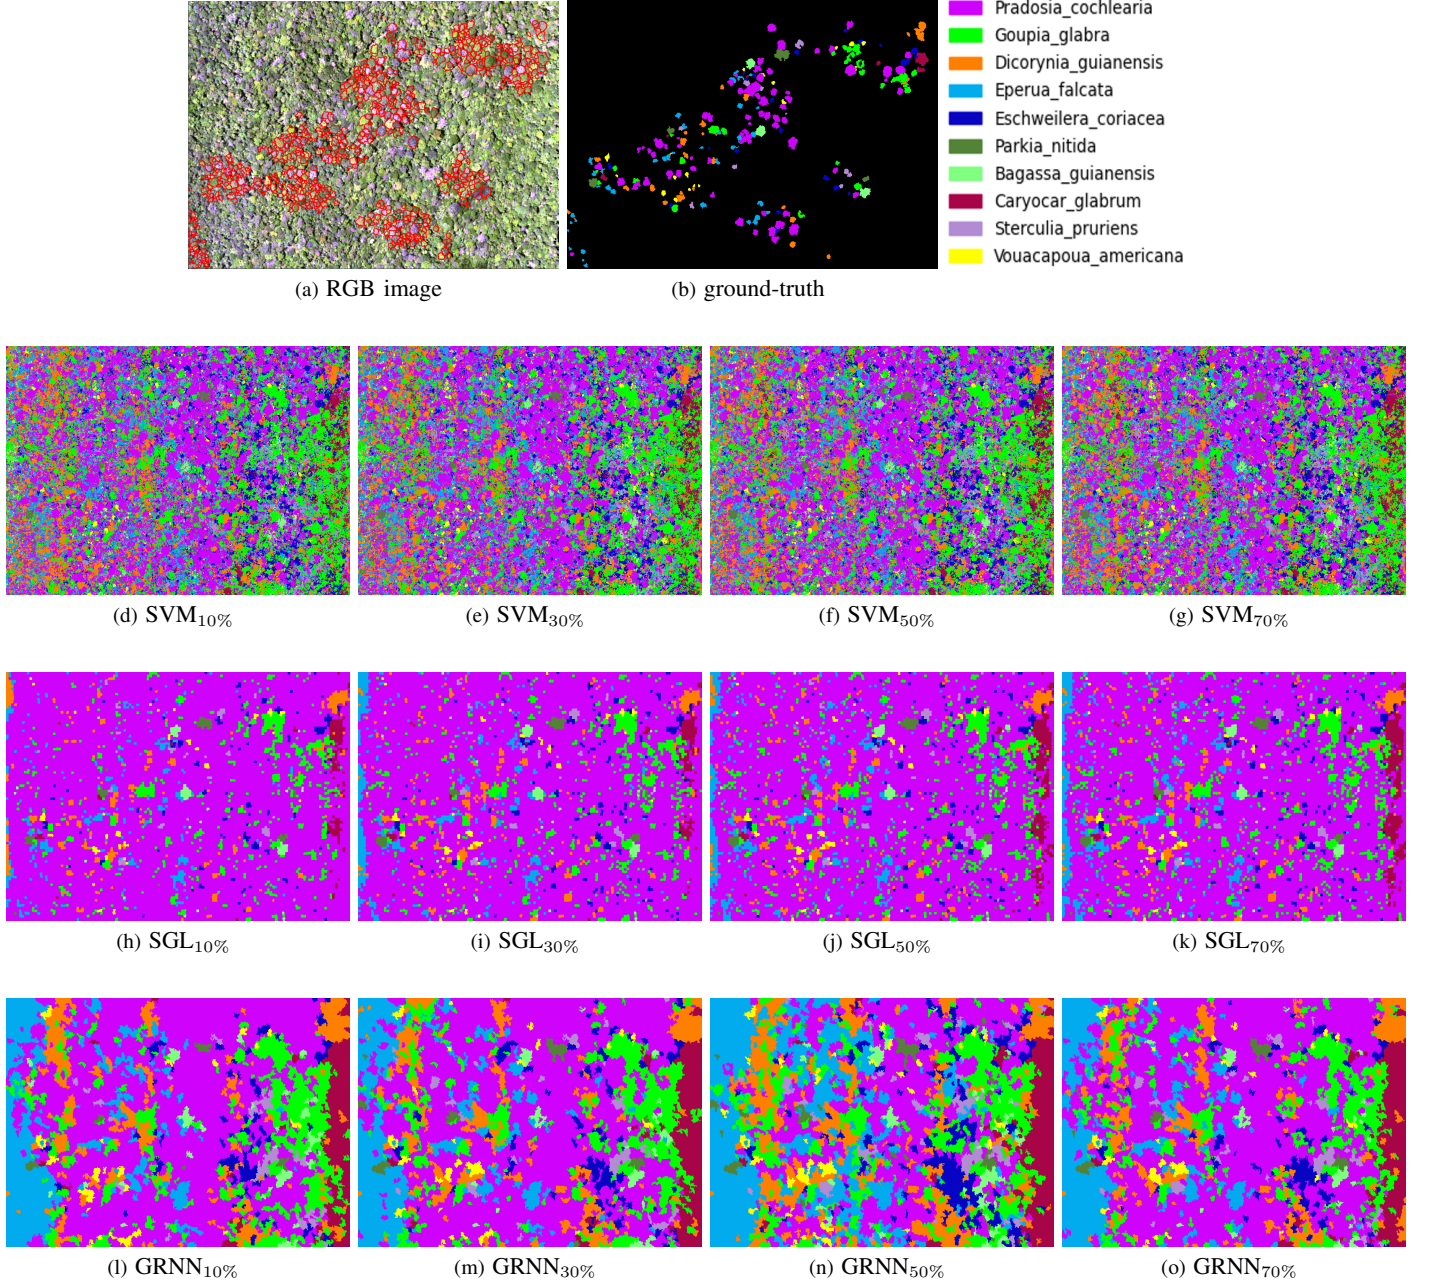
Fig. 5: (a) shows the composite RGB image of the forest region and the red polygons are the annotated tree crowns for the ten most frequent species, (b) shows the labeled tree species for which the key is provided in the right corner of the image. The subscripts 10, 30, 50 and 70 for all the three methods: SVM, SGL and GRNN, correspond to 10%, 30%, 50% and 70% of the labeled pixels used for training. The corresponding overall accuracy and $\kappa$ values for each of the clssification maps are mentioned in Table III.

[5] M. Zhang, W. Li, X. Zhao, H. Liu, R. Tao, and Q. Du, "Morphological transformation and spatial-logical aggregation for tree species classification using hyperspectral imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, no. 5501212, pp. 1–12, 2023.

[6] X. Zhu, "Semi-supervised learning literature survey, dept. computer. sci., univ. wisconsin-madison," Madison, Tech. Rep. 1530, Tech. Rep., 2005.

[7] X. Kang, S. Li, and J. A. Benediktsson, "Feature extraction of hyperspectral images with image fusion and recursive filtering," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 52, no. 6, pp. 3742–3752, 2013.

[8] W. Li, C. Chen, H. Su, and Q. Du, "Local binary patterns and extreme learning machine for hyperspectral imagery classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 7, pp. 3681–3693, 2015.

[9] L. Fang, N. He, S. Li, A. J. Plaza, and J. Plaza, "A New Spatial-Spectral Feature Extraction Method for Hyperspectral Images Using Local Covariance Matrix Representation," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 6, pp. 3534–3546, 2018.

[10] P. Sellars, A. I. Aviles-Rivero, and C. B. Schonlieb, "Superpixel Contracted Graph-Based Learning for Hyperspectral Image Classification," *IEEE Transactions on Geoscience and Remote*

*Sensing*, vol. 58, no. 6, pp. 4180–4193, 2020.

[11] D. Zhou, O. Bousquet, T. Lal, J. Weston, and B. Schölkopf, "Learning with local and global consistency," *Advances in neural information processing systems*, vol. 16, 2003.

[12] R. Achanta, A. Shaji, K. Smith, and A. Lucchi, "SLIC Superpixels Compared to State-of-the-Art Superpixel Methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 11, pp. 2274–2281, 2012.

[13] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv:1412.6980v9*, 2014.

[14] M. Aubry-Kientz, A. Laybros, B. Weinstein, J. G. C. Ball, T. Jackson, D. Coomes, and G. Vincent, "Multisensor data fusion for improved segmentation of individual tree crowns in dense tropical forests," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 3927–3936, 2021.

[15] L. Fang, S. Li, W. Duan, J. Ren, and J. A. Benediktsson, "Classification of hyperspectral images by exploiting spectral–spatial information of superpixel via multiple kernels," *IEEE transactions on geoscience and remote sensing*, vol. 53, no. 12, pp. 6663–6674, 2015.

[16] X. Kang, S. Li, and J. A. Benediktsson, "Spectral–spatial hyperspectral image classification with edge-preserving filtering," *IEEE transactions on geoscience and remote sensing*, vol. 52, no. 5, pp. 2666–2677, 2013.