

ROBUST HYPERSPECTRAL IMAGE CLASSIFICATION WITH REJECTION FIELDS

Filipe Condessa ^{a,b,c,e}, José Bioucas-Dias ^{a,b}, and Jelena Kovačević ^{c,d,e}

^aInstituto de Telecomunicações, Lisboa, Portugal

^bInstituto Superior Técnico, University of Lisbon, Lisboa, Portugal

^cDepartment of ECE, Carnegie Mellon University, Pittsburgh, PA, USA

^dDepartment of BME, Carnegie Mellon University, Pittsburgh, PA, USA

^eCenter for Bioimage Informatics, Carnegie Mellon University, Pittsburgh, PA, USA

ABSTRACT

In this paper we present a novel method for robust hyperspectral image classification using context and rejection. Hyperspectral image classification is generally an ill-posed image problem where pixels may belong to unknown classes, and obtaining representative and complete training sets is costly. Furthermore, the need for high classification accuracies is frequently greater than the need to classify the entire image.

We approach this problem with a robust classification method that combines classification with context with classification with rejection. A rejection field that will guide the rejection is derived from the classification with contextual information obtained by using the SegSALSA [1] algorithm. We validate our method in real hyperspectral data and show that the performance gains obtained from the rejection fields are equivalent to an increase the dimension of the training sets.

Index Terms— Hyperspectral image classification, hidden fields, robust classification, classification with rejection.

1. INTRODUCTION

Hyperspectral image classification is a challenging problem in remote sensing [2]. Due to generally ill-posed nature of hyperspectral image segmentation and classification, spatial regularization is often used (*e.g.* by promoting piecewise smooth classifications) which provides context to the classification. However, context alone cannot deal with difficulties arising from the existence of pixels belonging to unknown classes, unrepresentative and incomplete training sets, or overlapping classes. We propose a method that, combined with contextual classification, mitigates these difficulties through the inclusion of a reject option, thus achieving robust classification.

This paper was submitted to IEEE WHISPERS 2015: 7th Workshop on Hyperspectral Image and Signal Processing: Evolution on Remote Sensing. The authors gratefully acknowledge support from the Portuguese Science and Technology Foundation under projects UID/EEA/50008/2013, PTDC/EEI-PRO/1470/2012, the Portuguese Science and Technology Foundation and the CMU-Portugal (ICTI) program under grant SFRH/BD/51632/2011, NSF through award 1017278, and the CMU CIT Infrastructure Award.

In applications where classification performance is critical, performance gains can be obtained at the expense of not classifying all the samples. This can be achieved by selectively abstaining from classification in situations where misclassifications are expected. Classification with rejection was firstly analyzed in [3], where a rejection rule for optimum error-reject trade-off was designed for binary classification. Whereas the design of systems for classification with rejection is a rich area (see [4] and references therein for state of the art systems for classification with rejection), the application of these systems is rare in pixelwise image classification and in hyperspectral image classification.

In this paper we are interested in combining classification with context with classification with rejection to obtain a robust classification scheme. This means combining the option to reject when evidence for a classification is not enough (*i.e.* reject when the classifier is likely to misclassify) with the cues that arise from spatial context information (*i.e.* classification under assumption of piecewise smooth labeling). By associating spatial context with rejection, context cues influence the decision whether to reject or not (*e.g.* a sample is less likely to be rejected if all the neighboring samples have the same label), and rejection cues influence the context (*e.g.* a sample is more likely to be rejected if all the neighboring samples are also rejected). The robust classification idea was applied to tissue classification in stained microscopy images [5], where rejection is considered an extra class and Markov random fields are used as spatial contextual prior, with significant performance improvements. A major drawback of this approach is its rigidity with regard to the relative importance of rejection: if the amount of desired rejection is changed, the context has to be recomputed.

We propose a robust classification scheme that computes the rejection after the context, allowing us to change amount of samples rejected on the fly. By using the hidden fields resulting from segmentation via the constrained split augmented Lagrangian shrinkage algorithm (SegSALSA) [1, 6], we are able to infer a rejection field that reflects an ordering of the image pixels according to the degree of confidence associ-

ated with the contextual information, thus providing a simple and effective way to classify with rejection and context.

The paper is organized as follows: Section 2 provides the background on the contextual classification algorithm (SegSALSA) and performance measures for classification with rejection. Section 3 introduces the rejection field and describes their construction and properties. Section 4 presents experimental results and Section 5 concludes the paper.

2. BACKGROUND

SegSALSA The SegSALSA algorithm performs a marginal maximum *a posteriori* (MMAP) segmentation through the marginalization, on the discrete labels, of a hidden field driving the probabilities [7] and applies a vectorial total variation (VTV) prior [8, 9] on the hidden field. This results on a convex segmentation formulation that is solved using the constraint split augmented Lagrangian shrinkage algorithm (SALSA) [10].

To describe the SegSALSA algorithm, we start by introducing notation. Let $\mathbf{x} \in \mathbb{R}^{d \times n}$ represent a n -pixel hyperspectral image with d bands and $\mathbf{x}_i \in \mathbb{R}^d$ represent the feature vector of the i th image pixel, with $\mathcal{S} = \{1, \dots, n\}$ a set indexing the image pixels. Let $\mathcal{L} = \{1, \dots, K\}$ denote the set of possible K labels, and $\mathbf{y} \in \mathcal{L}^n$ a labeling of the image with $y_i \in \mathcal{L}$ the label of the i th pixel.

Under a Bayesian perspective, the maximum *a posteriori* (MAP) labeling $\hat{\mathbf{y}}$ is given by

$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y} \in \mathcal{L}^n} p(\mathbf{y}|\mathbf{x}) = \arg \max_{\mathbf{y} \in \mathcal{L}^n} p(\mathbf{x}|\mathbf{y})p(\mathbf{y}), \quad (1)$$

where $p(\mathbf{y}|\mathbf{x})$ represents the posterior probability of the labeling \mathbf{y} given the feature vectors \mathbf{x} , $p(\mathbf{x}|\mathbf{y})$ the observation model, and $p(\mathbf{y})$ the prior probability of the labeling \mathbf{y} .

SegSALSA approaches the segmentation, or labeling, problem by introducing a hidden field [7] \mathbf{z} represented by a $K \times n$ matrix that, for each pixel $i \in \mathcal{S}$, contains the hidden random vectors $\mathbf{z}_i \in \mathbb{R}^K$. The joint probability of labels \mathbf{y} and field \mathbf{z} is defined as $p(\mathbf{y}, \mathbf{z}) = p(\mathbf{y}|\mathbf{z})p(\mathbf{z})$, with $p(\mathbf{y}|\mathbf{z}) = \prod_{i \in \mathcal{S}} p(y_i|\mathbf{z}_i)$, allowing the expression of the joint probability of the features, labels and fields $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ as $p(\mathbf{x}, \mathbf{y}, \mathbf{z}) = p(\mathbf{x}|\mathbf{y})p(\mathbf{y}|\mathbf{z})p(\mathbf{z})$. With the hidden field and the joint probability defined, the marginalization on the discrete labels is now possible:

$$p(\mathbf{x}, \mathbf{z}) = \prod_{i \in \mathcal{S}} \left\{ \sum_{y_i \in \mathcal{L}} p(\mathbf{x}_i|y_i)p(y_i|\mathbf{z}_i) \right\} p(\mathbf{z}),$$

with the MMAP estimate being $\hat{\mathbf{z}}_{\text{MMAP}} = \arg \min_{\mathbf{z} \in \mathbb{R}^{K \times n}} p(\mathbf{x}, \mathbf{z})$.

By modeling the conditional probability $p(y_i = k|\mathbf{z}_i)$ as the k th component of the i th random vector $[\mathbf{z}_i]_k$, two constraints are introduced in the hidden field \mathbf{z} : nonnegativity constraint (*i.e.*, $[\mathbf{z}_i]_k \geq 0$) and sum-to-one constraint (*i.e.*, $\mathbf{1}_K^T \mathbf{z}_i = 1$). As only the discriminative power of the conditional probabilities $\mathbf{p}_i := [p(\mathbf{x}_i|y_i = 1), \dots, p(\mathbf{x}_i|y_i = K)]^T$

is relevant to the segmentation problem, we model them with the multinomial logistic regression (MLR) and use the logistic regression via splitting and augmented Lagrangian (LOR-SAL) [11] algorithm to learn the regression weights.

By dealing with the MMAP problem instead of the MAP, the prior is no longer applied on the discrete labels \mathbf{y} but on the continuous hidden field \mathbf{z} . A convex VTV prior [8, 9] is applied on the hidden field leading to promote a smoothness along the spatial dimensions of the field, and preservation and alignment of discontinuities across the classes.

From the initial integer optimization problem in (1), the contextual classification problem is now formulated as a convex optimization problem

$$\begin{aligned} \hat{\mathbf{z}}_{\text{MMAP}} = \arg \min_{\mathbf{z} \in \mathbb{R}^{K \times n}} & - \sum_{i \in \mathcal{S}} \left(\ln(\mathbf{p}_i^T \mathbf{z}_i) \right) - \ln p(\mathbf{z}) \quad (2) \\ \text{subject to: } & \mathbf{z} \geq 0, \quad \mathbf{1}_K^T \mathbf{z} = \mathbf{1}_n^T. \end{aligned}$$

Based on $\hat{\mathbf{z}}_{\text{MMAP}}$, $p(\mathbf{y}|\hat{\mathbf{z}}_{\text{MMAP}})$ provides a soft classification, and its maximization with respect to \mathbf{y} a hard classification. The optimization (2) is solved with SALSA [10], an instance of the alternating direction method of multipliers, in $O(Kn \log n)$ time. 1

Performance measures for classification with rejection

To assess the performance of classification systems with rejection we use the *nonrejected accuracy* A , the *fraction of rejected samples* r , and the *classification quality* Q [12]. The *nonrejected accuracy* measures the accuracy on the subset of samples that are not rejected, the *rejected fraction* measures how much rejection is performed, and the *classification quality* jointly measures how accurate the classification on the nonrejected samples is and how inaccurate are classification on the rejected samples is.

Considering \mathcal{S} the set of pixel indexes, let \mathcal{R} denote the set of rejected pixels ($\bar{\mathcal{R}}$ the set of nonrejected samples) and \mathcal{C} the set of correctly classified samples ($\bar{\mathcal{C}}$ the set of incorrectly classified samples). We define the nonrejected accuracy A as

$$A = \frac{|\mathcal{C} \cap \bar{\mathcal{R}}|}{|\bar{\mathcal{R}}|}.$$

This measure, combined with the respective fraction of rejected samples, cannot compare directly the behavior of two classifiers with rejection with different rejected fractions.

The *classification quality* Q is defined as

$$Q = \frac{|\mathcal{C} \cap \bar{\mathcal{R}}| + |\bar{\mathcal{C}} \cap \mathcal{R}|}{|\mathcal{S}|}.$$

The classification quality measures the proportion of samples that are *either* correctly classified and not rejected *or* incorrectly classified and rejected, relative to the total number of samples.

A classifier with rejection with a classification quality of Q when rejecting a fraction of samples r will be equivalent, in terms of correct decisions performed, to a classifier with no rejection and accuracy numerically equal to Q . The classification quality allows us to directly compare the performance of classification systems with rejection working at different rejected fractions.

3. REJECTION FIELD

From the SegSALSA formulation and resulting hidden field we can derive a contextual rejection scheme — the rejection field. The hidden field \mathbf{z} that results from the optimization problem (2) provides an indication of the degree of confidence associated with each label in each pixel. This is, if $[\mathbf{z}_i]_k > [\mathbf{z}_j]_l$, we are led to believe that the label l in the j th pixel has a smaller degree of confidence associated with the classification than the label k in the i th pixel.

Considering the following labeling

$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y} \in \mathcal{L}^n} p(\mathbf{y} | \hat{\mathbf{z}}_{\text{MMAP}}),$$

and obtaining the associated maximum probabilities

$$[\mathbf{z}_{\hat{\mathbf{y}}}]_i = p(\hat{\mathbf{y}}_i | \hat{\mathbf{z}}_{\text{MMAP}}), \quad (3)$$

the probabilities associated with the MMAP labeling, we note that the same line of thought of the components of the hidden fields as an indication of confidence can be applied to the entire labeling. If $[\mathbf{z}_{\hat{\mathbf{y}}}]_i > [\mathbf{z}_{\hat{\mathbf{y}}}]_j$, there is strong evidence that a higher degree of confidence exists in the labeling of the i th pixel as $\hat{\mathbf{y}}_i$ than in the labeling of the j th pixel as $\hat{\mathbf{y}}_j$.

We denote the $\mathbf{z}_{\hat{\mathbf{y}}}$ field (3) associated with the labeling \mathbf{y} as *rejection field*. By sorting $\mathbf{z}_{\hat{\mathbf{y}}}$ we obtain an ordering of the samples according to their relative confidence. The selection of a fraction of the lowest confidence samples to be rejected yields a simple, yet very effective, scheme for rejection. This method allows not only to define, *a priori*, specific values of the rejected fraction, but also to change it instantly. Furthermore the optimal value of rejection (the rejected fraction that maximizes the classification quality) can be estimated from a subset of samples, a validation set.

The characteristics of the VTV prior used in SegSALSA indirectly impose context on the rejection field. As it promotes smooth hidden fields, preservation of discontinuities and their alignment among classes, it preserves the discontinuities on the maximum values of the hidden field, and consequently promotes smoothness and preservation of discontinuities on the rejection field.

The computation of a rejection field and its use as a rejection rule is an approximation to the problem of contextual rejection approached in [5], where a joint optimization on the labels and on the reject option is performed. We perform a sequential optimization: first an optimization on the labels and then a binary optimization on the reject option through the use

Table 1. Classwise performance measures for classification with rejection of the Indian Pines scene (Fig. 1, top row). OA corresponds to the accuracy of the SegSALSA classification method with no rejection (Fig. 1 b), and A corresponds to nonrejected accuracy, Q to classification quality, and r to rejected fraction from classification with rejection (Fig. 1 c). n Is the number of samples per class.

| | OA (%) | A (%) | Q (%) | r (%) | n |
|------------------|----------|---------|--------------|--------------|------|
| alfafa | 71.74 | 0.00 | 26.09 | 97.83 | 46 |
| corn no-till | 66.67 | 76.57 | 79.06 | 13.94 | 1428 |
| corn min-till | 53.13 | 47.33 | 43.49 | 36.87 | 830 |
| corn clean | 100.00 | 100.00 | 96.62 | 3.38 | 237 |
| grass past. | 77.85 | 81.06 | 75.78 | 13.66 | 483 |
| grass trees | 90.55 | 90.83 | 90.00 | 1.37 | 730 |
| grass mowed | 0.00 | 0.00 | 0.00 | 0.00 | 28 |
| hay | 99.16 | 100.00 | 97.70 | 3.14 | 478 |
| oats | 0.00 | 0.00 | 100.00 | 100.00 | 20 |
| soybean no-till | 72.94 | 74.09 | 71.81 | 7.10 | 972 |
| soybean min-till | 72.38 | 88.54 | 89.53 | 19.67 | 2455 |
| soybean clean | 79.26 | 78.50 | 69.48 | 14.50 | 593 |
| wheat | 86.34 | 86.21 | 85.37 | 0.98 | 205 |
| woods | 74.55 | 81.56 | 82.29 | 9.96 | 1265 |
| bldg. | 66.06 | 80.70 | 84.20 | 18.13 | 386 |
| stone | 32.26 | 32.26 | 32.26 | 0.00 | 93 |

of a rejection field. Whereas the solution we obtain is an approximation to the contextual rejection problem (joint minimization), the sequential optimization we perform has a clear advantage over the joint optimization approach: the amount of rejection can be changed on the fly, whereas on the joint optimization approach the context has to be recomputed.

4. EXPERIMENTAL RESULTS

We illustrate the performance of our algorithm through the robust classification of the AVIRIS Indian Pine scene, and the ROSIS Pavia university scene. The Indian Pine scene was acquired with the AVIRIS sensor in NorthWest Indiana (USA), being a 145×145 pixel hyperspectral image with 200 spectral bands (excluding water absorption bands) containing 16 not mutually exclusive classes. The Pavia University scene was acquired with the ROSIS sensor in Pavia (Italy), being a 610×340 pixel hyperspectral image with 103 spectral bands containing 9 not mutually exclusive classes. We model the MLR weights with LORSAL and use the SegSALSA algorithm to include context in the classification.

Figure 1 illustrates the performance gains obtained by combining classification with context with classification with rejection. Using the rejection field, we are able to change the amount of rejected samples on the fly, without need to recompute the context. Table 1 shows that the performance gains are not equally distributed among all classes. The bulk of the performance gains is achieved by increasing the per-




Fig. 1. Top row: Robust classification of the Indian Pine scene. (a) Ground truth and (b) classification with 15 training samples per class using LORSAL and SegSALSA (73.5% accuracy), (c) classification with optimal rejected fraction (80.9% nonrejected accuracy at a rejected fraction of 14.7% with classification quality of 79.2%), and (d) associated rejection fields. (e) Nonrejected accuracy and classification quality variation with rejected fraction (maximum classification quality in red). Bottom row: Robust classification of the Pavia University scene. (f) Ground truth and (g) classification with 15 training samples per class using LORSAL and SegSALSA (69.8% accuracy), (h) classification with optimal rejected fraction (74.6% nonrejected accuracy at a rejected fraction of 12.9% with classification quality of 73.0%), and (i) associated rejection fields. (j) Nonrejected accuracy and classification quality variation with rejected fraction (maximum classification quality in red).

formance in highly populated classes. This is achieved either by a minor drop in nonrejected accuracy in small number of lesser populated classes, or by the entire rejection of lesser populated class.

The performance gains obtained from the allocation of labeled samples to estimate the optimal rejected fraction (the rejected fraction that maximizes the classification quality) can be larger than the gains obtained from using those samples to extend the training set, retraining with LORSAL and classifying the image with SegSALSA. This effect is clearly illustrated on table 2, where, in the Indian Pine scene, for an initial training set of 30 samples the class, the effect of either estimating the optimal rejected fraction from 50 randomly selected samples or retraining the classifier with the extra 50 samples is shown. Whereas it is clear that the increased performance obtained by estimating the rejected fraction when compared to retraining the classifier will not hold for smaller training sets, for larger training sets it is a computationally cheaper and performance-wise better alternative to retraining the classifier.

Table 2. Effect of increasing the dimension of the training set with new samples vs. using the new samples as validation set to estimate the rejected fraction r in the Indian Pine scene. Comparison of average performance (classification quality Q , nonrejected accuracy A , and rejected fraction r) over 30 Monte Carlo runs.

| | r (%) | Q (%) | A (%) |
|--|---------|--------------|---------|
| initial – training set of 480 samples with no rejection | 0.00 | 84.21 | 84.21 |
| extended – training set of 480 + 50 samples with no rejection | 0.00 | 86.46 | 86.46 |
| estimated – training set of 480 samples, with optimal rejected fraction estimated from 50 samples | 12.77 | 87.02 | 91.16 |
| optimal – training set of 480 samples, with true optimal rejected fraction | 12.49 | 88.37 | 91.53 |

5. CONCLUDING REMARKS

We presented a simple and effective scheme for robust hyper-spectral image classification by combining classification with context and classification with rejection by deriving a rejection field from the hidden fields that drive the contextual clas-

sification. We moved from the joint optimization problem of context and rejection, to a faster separate optimization without losing the contextual effect on the rejection. The performance gains obtained by using robust classification are shown to be equivalent to training the classifier with larger training sets.

Acknowledgements The authors would like to thank D. Landgrebe at Purdue University for providing the AVIRIS Indian Pines scene, P. Gamba at Pavia University for providing the ROSIS Pavia University scene.

6. REFERENCES

- [1] J. Bioucas-Dias, F. Condessa, and J. Kovačević, “Alternating direction optimization for image segmentation using hidden Markov measure field models,” in *Proc. SPIE Conf. Image Process.*, San Francisco, Feb. 2014.
- [2] J. Bioucas-Dias, A. Plaza, G. Camps-Valls, P. Scheunders, N. Nasrabadi, and J. Chanussot, “Hyperspectral remote sensing data analysis and future challenges,” *Geoscience and Remote Sensing Magazine, IEEE*, vol. 1, no. 2, pp. 6–36, 2013.
- [3] C. K. Chow, “On optimum recognition error and reject tradeoff,” *IEEE Trans. Inf. Theory*, vol. 16, no. 1, pp. 41–46, Jan. 1970.
- [4] I. Pillai, G. Fumera, and F. Roli, “Multi-label classification with a reject option,” *Pattern Recogn.*, vol. 46, no. 8, pp. 2256 – 2266, 2013.
- [5] F. Condessa, J. Bioucas-Dias, C. Castro, J. Ozolek, and J. Kovačević, “Classification with rejection option using contextual information,” in *Proc. IEEE Int. Symp. Biomed. Imag.*, San Francisco, Apr. 2013.
- [6] F. Condessa, J. Bioucas-Dias, and J. Kovačević, “Supervised hyperspectral image segmentation: a convex formulation using hidden fields,” in *IEEE GRSS Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS’14)*, Lausanne Switzerland, June 2014.
- [7] J. Marroquin, E. Santana, and S. Botello, “Hidden Markov measure field models for image segmentation,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 25, no. 11, pp. 1380–1387, 2003.
- [8] B. Goldluecke, E. Strelakovsky, and D. Cremers, “The natural vectorial total variation which arises from geometric measure theory,” *SIAM Journal on Imaging Sciences*, vol. 5, no. 2, pp. 537–563, 2012.
- [9] L. Sun, Z. Wu, J. Liu, and Z. Wei, “Supervised hyperspectral image classification using sparse logistic regression and spatial-tv regularization,” in *IEEE Geoscience and Remote Sensing Symposium (IGARSS’13)*, Melbourne, Australia, 2013.
- [10] M. Afonso, J. Bioucas-Dias, and M. Figueiredo, “An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems,” *Image Processing, IEEE Transactions on*, vol. 20, no. 3, pp. 681–695, 2011.
- [11] J. Li, J. Bioucas-Dias, and A. Plaza, “Hyperspectral image segmentation using a new bayesian approach with active learning,” *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 10, pp. 3947–3960, Oct. 2011.
- [12] F. Condessa, J. Bioucas-Dias, and J. Kovačević, “Performance measures for classification systems with rejection,” *Preprint*, 2015, arxiv.org/abs/1504.02763 [cs.CV].