

SCANNER INVARIANT MULTIPLE SCLEROSIS LESION SEGMENTATION FROM MRI

Shahab Aslani^{1,6}, Vittorio Murino^{1,4}, Michael Dayan², Roger Tam⁵, Diego Sona^{1,3}, Ghassan Hamarneh⁶

¹ Pattern Analysis and Computer Vision, Istituto Italiano di Tecnologia, Genoa, Italy

² Human Neuroscience Platform, Fondation Campus Biotech, Geneva, Switzerland

³ NeuroInformatics Laboratory, Fondazione Bruno Kessler, Trento, Italy

⁴ Computer Science Department, University of Verona, Italy

⁵ Department of Radiology and School of Biomedical Engineering,
University of British Columbia, Vancouver, Canada

⁶ School of Computing Science, Simon Fraser University, Vancouver, Canada

ABSTRACT

This paper presents a simple and effective generalization method for magnetic resonance imaging (MRI) segmentation when data is collected from multiple MRI scanning sites and as a consequence is affected by (site-)domain shifts. We propose to integrate a traditional encoder-decoder network with a regularization network. This added network includes an auxiliary loss term which is responsible for the reduction of the domain shift problem and for the resulting improved generalization. The proposed method was evaluated on multiple sclerosis lesion segmentation from MRI data. We tested the proposed model on an in-house clinical dataset including 117 patients from 56 different scanning sites. In the experiments, our method showed better generalization performance than other baseline networks.

Index Terms— Magnetic Resonance Imaging, Multiple Sclerosis, Lesion Segmentation, Domain Generalization,

1. INTRODUCTION

Deep learning models, in particular Convolutional Neural Networks (CNNs) [1] have shown excellent performance in a large variety of computer vision tasks, including image classification [2], object detection [3], semantic segmentation [4], etc. It is, therefore, common to expect that successful deep models can obtain good performances. However, it has been shown that, in practice, these approaches easily fail to generalize well [5]. According to the literature, the most important reasons for this failure are (i) the small size of the training data which causes overfitting and (ii) the large difference between training and test data which is typically addressed as domain shift. Therefore, one of the most important problems that arises is how to improve the quality of models so that they generalize well to unseen data from a different domain. During the last years, several algorithms have been proposed to tackle the mentioned problem, improving the models gen-

eralization through heuristic techniques such as dropout [6], early stopping [7], weight decay [8], data augmentation [9], and randomization methods [5].

Thanks to these regularization methods, deep models are reaching expert-level accuracy in medical image segmentation. However, they still have a limited clinical application due to the aforementioned challenge (i), which is also considered as one of the most relevant and common problem in medical image analysis tasks [10]. To tackle this challenge, several strategies have been proposed, such as using 2.5D (slices) rather than full-size 3D images [11], initializing parameters of the proposed model with pre-trained weights on natural images [12], and adopting special data augmentation techniques [13]. An effective solution, however, would be to merge datasets collected from different centers. This, however, introduces another important challenge. The medical data acquisition can vary significantly between different centers. For instance, in magnetic resonance imaging (MRI), this procedure is often subject to the variation of several specific properties such as scanner, magnet strength, and acquisition protocol. This causes high domain variability between datasets which eventually can result in poor generalization. In order to tackle this problem, several methods have been proposed such as scanner invariant representations for medical image harmonization [14], one-shot domain adaptation [15], and unsupervised methods [16] for medical image segmentation.

In this work, we propose a novel simple generalization method to enhance baseline models and diminish the effect of domain differences in the data. To this end, a regularization network, equipped with an auxiliary loss function, is proposed to incorporate regularization into a standard encoder-decoder segmentation network. We tested the model with a standard cross-validation procedure using an in-house dataset on Multiple Sclerosis (MS) lesion segmentation. Results show that the proposed regularization network has a significant impact on the generalization of the standard seg-

mentation network when data from multiple centers are used.

2. METHOD

Generally, the performance of models suffers when they are applied to domains other than the ones they were trained upon. In this work, the goal is to improve the generalizability of a backbone segmentation model by training on a collection of datasets presenting high domain variability. The selected backbone encoder-decoder model has a traditional loss function used for image segmentation. However, to handle the domain shift problem, we propose a regularization network including an auxiliary loss function that is designed to encourage the model to ignore domain-specific information. This property emerges from optimizing cross entropy or correlation coefficient as detailed below. Training the backbone segmentation network incorporated with regularization network reduces the domain differences problem across the datasets.

2.1. Network Architecture

The overall architecture of the proposed model includes three main components (Figure 1). The first component is a feature extractor network consisting of an encoder ϕ_E which is fed by an input image $x_i \in \mathcal{X}$. The output of the encoder ϕ_E is a p -dimensional vector $r_i = \{r_{ij} \in \mathbb{R}\}_{j=1}^p \in \mathcal{R}$ representing the latent features. The second component is a segmentation network consisting of a decoder ϕ_D which reconstruct from the latent features $r_i \in \mathcal{R}$ a feature representation with the resolution of the input image $x_i \in \mathcal{X}$. The output layer then produces a dense pixel-wise prediction output $s_i \in \mathcal{S}$ using a softmax activation. This network includes a traditional loss term used to update the weights to improve the segmentation performance (see section 2.2). Note that the described architecture composed of the mentioned encoder ϕ_E and decoder ϕ_D is very similar to the model presented in [17], using 3D operations rather than 2D ones and removing the regularization layers (dropout). For simplicity, we removed skip-connections in Figure 1. The third component of the presented model is a regularization network ϕ_R including three perceptron layers and a softmax layer. The network receives the latent features $r_i \in \mathcal{R}$ to produce category-wise prediction $c_i \in \mathcal{C}$, which in our case corresponds to the prediction of the input’s domain. Our observation is that during training, the model without ϕ_R learns how to segment the input images also encoding their source domain. This results in overfitting of the model with a domain bias. As a result, the segmentation performance of the network on data coming from unseen source domains degrades. The goal of the regularization network is therefore to steer the whole model to reduce the domain bias, to obtain a better generalization and, hence, a fairer segmentation performance on seen and unseen domains. To this aim, we introduce an auxiliary loss term whose aim is to

confuse the model about the dataset domains, thus forcing the model to learn how to segment the image while maximally reducing the domain bias.

2.2. The Loss Functions

Our method has been tested using multiple loss terms to enable the network to precisely segment the input image while generalizing over the domains. Specifically, the proposed model was optimized according to the loss function formulated as:

$$\mathcal{L}(\mathcal{X}, \mathcal{S}, \mathcal{C}, \mathcal{H}, \mathcal{G}) = \mathcal{L}_{\text{seg}}(\mathcal{X}, \mathcal{S}, \mathcal{G}) + \lambda \mathcal{L}_{\text{reg}}(\mathcal{X}, \mathcal{C}, \mathcal{H}) \quad (1)$$

where $\lambda \in [0, 1]$ is a hyperparameter controlling the trade-off between segmentation and regularization losses. \mathcal{H} and \mathcal{G} are the domain-wise and pixel-wise ground truth, respectively. In this work, we propose three different regularization approaches using loss functions \mathcal{L}_{reg} based on two well-established measures, namely cross-entropy and the Pearson correlation coefficients. Hence, \mathcal{L}_{reg} can take on either of these three options:

$$\mathcal{L}_{\text{reg}}(\mathcal{X}, \mathcal{C}, \mathcal{H}) = \{\mathcal{L}_{\text{pc}}(\mathcal{X}, \mathcal{C}, \mathcal{H}), \mathcal{L}_{\text{rand}}(\mathcal{X}, \mathcal{C}, \mathcal{H}), \mathcal{L}_{\text{du}}(\mathcal{X}, \mathcal{C}, \mathcal{H})\} \quad (2)$$

Pearson Correlation Loss \mathcal{L}_{pc} : Given an input image $x_i \in \mathcal{X}$, the regularization network ϕ_R generates the corresponding output vector $c_i = \{c_{ij} \in [0, 1]\}_{j=1}^n \in \mathcal{C}$ which shows the probabilities for x_i to be in one of the n domains. For each input image x_i the corresponding one-hot encoded vector as ground truth domain labeling is also given by $h_i = (0, 0, \dots, 1, 0, \dots, 0) \in \mathcal{H}$. The Pearson correlation coefficient measures the strength of linear correlation or similarity between two variables, where higher values correspond to higher similarity. Hence, to remove the domain bias, the model can be trained to minimize the Pearson correlation between \mathcal{C} and \mathcal{H}

$$\mathcal{L}_{\text{pc}}(x_i, c_i, h_i) = \frac{\sum_j (c_{ij} - \bar{c}_i)(h_{ij} - \bar{h}_i)}{\sqrt{\sum_j (c_{ij} - \bar{c}_i)^2 \sum_j (h_{ij} - \bar{h}_i)^2}} \quad (3)$$

where \bar{h}_i and \bar{c}_i denote the mean values of elements in the vectors h_i and c_i , respectively.

Randomized Cross-Entropy Loss $\mathcal{L}_{\text{rand}}$: The most commonly used loss function for image classification is the cross-entropy:

$$\mathcal{L}_{\text{rand}}(x_i, c_i, h_i) = - \sum_j h_{ij} \log c_{ij} \quad (4)$$

which allows comparing the class predictions vector $c_i \in \mathcal{C}$ and the ground truth one-hot encoded vector $h_i \in \mathcal{H}$, penalizing the correct classes having a probability diverging from

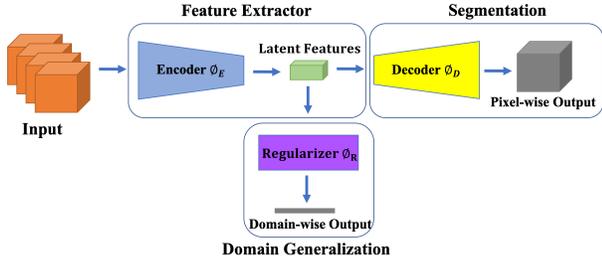


Fig. 1. Overall architecture of the proposed method including three main components: An encoder network ϕ_E for extracting latent features, a decoder network ϕ_D for segmentation and a regularization network ϕ_R for domain generalization.

the expected value. Our ultimate goal is to push the network to filter unnecessary domain information during training. This can be easily obtained by shuffling the ground truth $h_i \in \mathcal{H}$ at each training iteration and for every single input. This encourages the model to learn the wrong classes selected randomly.

Discrete Uniform Loss \mathcal{L}_{du} : Analyzing the problem from a different prospective, to remove the domain bias, the encoder network ϕ_E should generate a representation $r_i \in \mathcal{R}$ from which the domain classifier in ϕ_R cannot extract information. This should correspond to a classifier in ϕ_R that classify any class with equal probability, independently from the input. We can obtain this result training the model with the cross-entropy loss (Equation 4), forcing the domain ground truth to be a uniform distribution $h_i : \{h_{ij} = \frac{1}{n}\}_{j=1}^n \in \mathcal{H}$

Segmentation Loss \mathcal{L}_{seg} : To fit the model according to the segmentation task, we used a well-known loss function for image segmentation, namely the soft-Dice loss function:

$$\mathcal{L}_{seg}(x_i, s_i, g_i) = 1 - \frac{2 \sum_i s_i g_i}{\sum_i s_i^2 + \sum_i g_i^2} \quad (5)$$

where $s_i \in \mathcal{S}$ is the dense pixel-wise prediction and $g_i \in \mathcal{G}$ is the corresponding ground truth segmentation. This function penalizes ϕ_E and ϕ_D based on the overlap between the prediction and the ground truth segmentation.

2.3. Implementation Details

We designed a backbone model based on the U-Net [17], built by concatenating a down-sampling encoder ϕ_E made by 4 stages with an up-sampling decoder ϕ_D made by 4 stages. However, differently from the U-Net, the regularization layers (dropout) were removed and all 2D operations were replaced by their 3D counterparts. During the training procedure, the backbone model was combined with the proposed regularization network. On the contrary, during the testing phase, the regularization network was removed addressing only the segmentation task. The model was executed on 3D patches with size $64 \times 64 \times 64$ cropped from each volume (with 50% overlap). The model was trained only using patches containing

at least one voxel labeled as a lesion. During the test, on the contrary, all patches were used. The evaluation of the model was performed on the reconstructed full-size volumes, fusing the predictions for all patches. The proposed model was implemented in Python using Keras with Tensorflow backend. We trained our model using Adam optimizer with an initial learning rate of 0.0001. The size of batch and the maximum number of training epochs were fixed respectively at 15 and 500 (with 300 steps per epoch). Regarding the model initialization, all blocks were randomly initialized from a Gaussian distribution. The hyperparameter λ in Equation 1 was selected through grid search with values equal to 0.2, 0.3 and 0.1 for \mathcal{L}_{pc} , \mathcal{L}_{du} and \mathcal{L}_{rand} , respectively.

3. EXPERIMENTS

We evaluate the performance of the proposed method on an in-house clinical dataset, collected from 56 different centers with a variable number of MS patients per center. Each patient had several scanning sessions, with each session including 4 MRI modalities (T1w, T2w, PDw, and FLAIR). Each volume is composed by 60 slices with FOV= 256×256 and $1mm \times 1mm \times 3mm$ voxel resolution. All volumes were already segmented manually representing the ground truth lesion masks. All images were skull-stripped using Brain Extraction Tool (BET) [18], and rigidly registered to the $1mm^3$ MNI-ICBM152 template [19] using FMRIB’s Linear Image Registration tool (FLIRT) [20]. Ethics approval for data collection and secondary analysis was granted by a Research Ethics Board at the University of British Columbia.

Experimental Protocol: We considered each site as a separate domain and to keep data balanced over all available sites, a single subject including one time point with four MRI modalities (T1w, T2w, PDw, and FLAIR) was selected from each site. We implemented 5-fold cross-validation over the whole data (60%:20%:20% for training, validation, and test, respectively). For comparison purpose, we repeated the above-mentioned experiment (using exactly the same folds) for the backbone model without any regularization (denoted as the BM) and the same backbone model with additional dropout layers (denoted as the BDM). Note that BDM model is equivalent the U-net model [17] with replacing the 2D operations by their 3D counterparts.

4. RESULTS

We evaluated our model using four measures: Dice Similarity Coefficient (DSC), Lesion-wise True Positive Rate (LTPR), Lesion-wise False Positive Rate (LFPR), and Positive Prediction Value (PPV) (refer to [10] for more details). Table 1 summarizes the results of our experiment on the test set comparing our model with other baseline methods. The Table shows the mean value of DSC, LTPR, LFPR, and PPV. As can

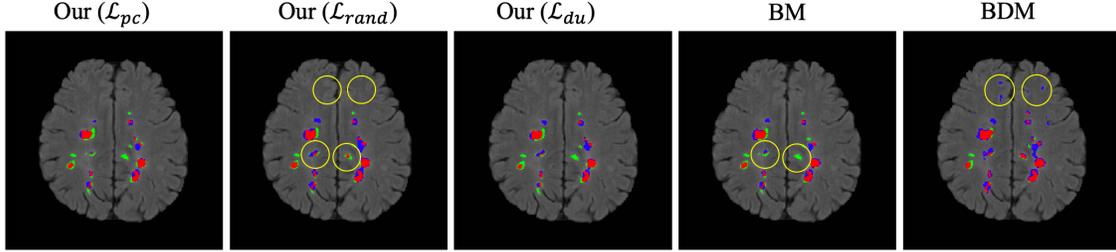


Fig. 2. Segmentation of a random subject obtained by different methods against ground truth annotation. On all images, true positives, false negatives, and false positives are marked in red, green and blue, respectively (refer to Sec. 5 for yellow circles).

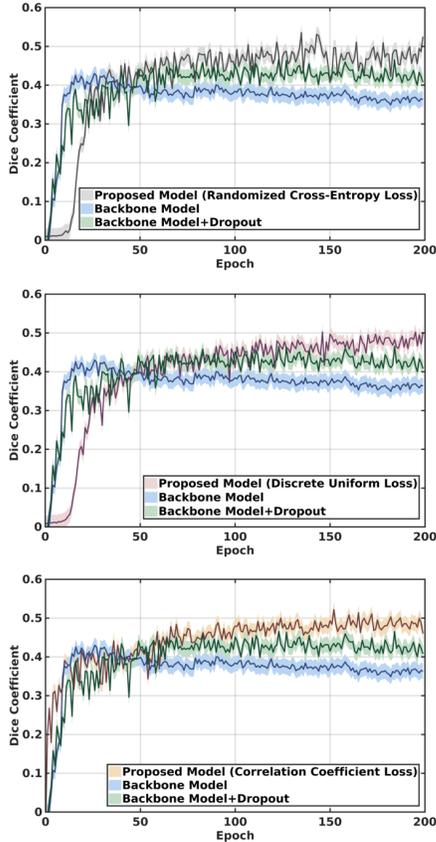


Fig. 3. Comparison of the DSC measure performance of the proposed methods with other baseline models on validation set during training.

be seen, our proposed methods outperformed baseline methods on the DSC measure. Moreover, in terms of LTPR and LFPR measures, our model with the randomized and discrete uniform auxiliary loss functions showed more balanced performance compared with the other models. Figure 2 shows an example of segmentation for all methods. Figure 3 compares the DSC performance of the proposed methods with other models on the validation set. Confirming the results

Table 1. Results related to our experiment. Mean values of DSC, LTPR, LFPR, and PPV were measured for different methods. Values in bold and italic indicate the first-best and second-best results.

| Method | DSC | LTPR | LFPR | PPV |
|------------------------------|---------------|---------------|---------------|---------------|
| Our (\mathcal{L}_{pc}) | 0.4638 | 0.4267 | 0.3954 | 0.4865 |
| Our (\mathcal{L}_{rand}) | 0.5001 | 0.4618 | 0.3348 | 0.5193 |
| Our (\mathcal{L}_{du}) | <i>0.4893</i> | <i>0.4670</i> | 0.3525 | <i>0.5182</i> |
| BM | 0.4540 | 0.4318 | <i>0.3383</i> | 0.5088 |
| BDM | 0.4598 | 0.5821 | 0.5151 | 0.4577 |

reported in the test set, as shown in this Figure, our model with all three possible auxiliary loss terms depicts better DSC performance than the baseline methods.

5. DISCUSSION AND CONCLUSION

In this paper, we have introduced a generalization methods implemented via an auxiliary loss with three variants. We tested this method on medical image segmentation, particularly MS lesion segmentation from MRI modalities in the presence of domain shift originating from multi-center dataset. The proposed model is the combination of a traditional encoder-decoder network for segmentation and an additional regularization network including an auxiliary loss term for domain generalization. Investigating the impact of the proposed method summarized in Table 1, our model always outperformed the baseline models when considering the DSC measures (regardless of which of the adopted auxiliary loss variant was used). However, the best performance in terms of DSC, LFPR, and PPV measures among all tested models is provided by our model with the randomized auxiliary loss function. The BDM model showed the best LTPR measure together with the worst LFPR measure showing that this model has very poor trade-off between LTPR and LFPR.

Confirming the above-mentioned point, Figure 2 shows that BDM model has over-segmented lesion regions (referring to the top two yellow circles). Moreover, it can be seen that BM model did not identify some lesions (referring to the bottom two yellow circles). However, the proposed method with randomized loss term shows a considerable good per-

formance by not only identifying the mentioned small lesions but also ignoring the false positives.

Acknowledgments. We gratefully acknowledge Medical Image Analysis Lab members for their helpful comments and Compute Canada for computational resources.

6. REFERENCES

- [1] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al., “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [2] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger, “Densely connected convolutional networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4700–4708.
- [3] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik, “Rich feature hierarchies for accurate object detection and semantic segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 580–587.
- [4] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla, “Segnet: A deep convolutional encoder-decoder architecture for image segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 12, pp. 2481–2495, 2017.
- [5] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals, “Understanding deep learning requires rethinking generalization,” *arXiv preprint arXiv:1611.03530*, 2016.
- [6] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [7] Grégoire Montavon, Geneviève Orr, and Klaus-Robert Müller, *Neural networks: tricks of the trade*, vol. 7700, springer, 2012.
- [8] Anders Krogh and John A Hertz, “A simple weight decay can improve generalization,” in *Advances in Neural Information Processing Systems*, 1992, pp. 950–957.
- [9] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu, “Spatial transformer networks,” in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, Cambridge, MA, USA, 2015, NIPS’15, pp. 2017–2025, MIT Press.
- [10] Aaron Carass, Snehashis Roy, Amod Jog, Jennifer L Cuzzocreo, Elizabeth Magrath, Adrian Gherman, Julia Button, James Nguyen, Ferran Prados, Carole H Sudre, et al., “Longitudinal multiple sclerosis lesion segmentation: Resource and challenge,” *NeuroImage*, vol. 148, pp. 77–102, 2017.
- [11] Shahab Aslani, Michael Dayan, Loredana Storelli, Massimo Filippi, Vittorio Murino, Maria A Rocca, and Diego Sona, “Multi-branch convolutional neural network for multiple sclerosis lesion segmentation,” *NeuroImage*, vol. 196, pp. 1–15, 2019.
- [12] Shahab Aslani, Michael Dayan, Vittorio Murino, and Diego Sona, “Deep 2D encoder-decoder convolutional neural network for multiple sclerosis lesion segmentation in brain mri,” in *International MICCAI Brainlesion Workshop*. Springer, 2018, pp. 132–141.
- [13] Ling Zhang, Xiaosong Wang, Dong Yang, Thomas Sanford, Stephanie Harmon, Baris Turkbey, Holger Roth, Andriy Myronenko, Daguang Xu, and Ziyue Xu, “When unseen domain generalization is unnecessary? rethinking data augmentation,” *arXiv preprint arXiv:1906.03347*, 2019.
- [14] Daniel Moyer, Greg Ver Steeg, Chantal MW Tax, and Paul M Thompson, “Scanner invariant representations for diffusion mri harmonization,” *arXiv preprint arXiv:1904.05375*, 2019.
- [15] Sergi Valverde, Mostafa Salem, Mariano Cabezas, Deborah Pareto, Joan C Vilanova, Lluís Ramió-Torrentà, Àlex Rovira, Joaquim Salvi, Arnau Oliver, and Xavier Lladó, “One-shot domain adaptation in multiple sclerosis lesion segmentation using convolutional neural networks,” *NeuroImage: Clinical*, vol. 21, pp. 101638, 2019.
- [16] Christian S Perone, Pedro Ballester, Rodrigo C Barros, and Julien Cohen-Adad, “Unsupervised domain adaptation for medical imaging segmentation with self-ensembling,” *NeuroImage*, vol. 194, pp. 1–11, 2019.
- [17] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical Image Computing and Computer-assisted Intervention*. Springer, 2015, pp. 234–241.
- [18] Stephen M Smith, “Fast robust automated brain extraction,” *Human Brain Mapping*, vol. 17, no. 3, pp. 143–155, 2002.
- [19] Kenichi Oishi, Karl Zilles, Katrin Amunts, Andreia Faria, Hangyi Jiang, Xin Li, Kazi Akhter, Kegang Hua, Roger Woods, Arthur W Toga, et al., “Human brain white matter atlas: identification and assignment of

common anatomical structures in superficial white matter,” *Neuroimage*, vol. 43, no. 3, pp. 447–457, 2008.

- [20] Mark Jenkinson, Peter Bannister, Michael Brady, and Stephen Smith, “Improved optimization for the robust and accurate linear registration and motion correction of brain images,” *Neuroimage*, vol. 17, no. 2, pp. 825–841, 2002.