



INTERPRETABLE DYSARTHIC SPEAKER ADAPTATION BASED ON OPTIMAL-TRANSPORT

Rosanna Turrisi^{1,2,3}, Leonardo Badino⁴

¹University of Genoa, ²Italian Institute of Technology, ³University of Ferrara, ⁴PerVoice

rosanna.turrisi@edu.unige.it, leonardo.badino@pervoice.it

Abstract

This work addresses the mismatch problem between the distribution of training data (*source*) and testing data (*target*), in the challenging context of dysarthric speech recognition. We focus on Speaker Adaptation (SA) in command speech recognition, where data from multiple sources (i.e., multiple speakers) are available. Specifically, we propose an unsupervised Multi-Source Domain Adaptation (MSDA) algorithm based on optimal-transport, called MSDA via Weighted Joint Optimal Transport (MSDA-WJDOT). We achieve a Command Error Rate relative reduction of 16% and 7% over the speaker-independent model and the best competitor method, respectively. The strength of the proposed approach is that, differently from any other existing SA method, it offers an interpretable model that can also be exploited, in this context, to diagnose dysarthria without any specific training. Indeed, it provides a closeness measure between the target and the source speakers, reflecting their similarity in terms of speech characteristics. Based on the similarity between the target speaker and the healthy/dysarthric source speakers, we then define the healthy/dysarthric score of the target speaker that we leverage to perform dysarthria detection. This approach does not require any additional training and achieves a 95% accuracy in the dysarthria diagnosis.

Index Terms: unsupervised speaker adaptation, dysarthric speech, optimal transport, dysarthria detection

1. Introduction

Many Automatic Speech Recognition (ASR) models assume that test and training datasets are sampled from the same distribution. However, new data can exhibit a distribution change and drastically decrease the recognition performance. This is due, for example, to different characteristics between speakers (e.g., different accents, different speaking style). Such a mismatch is even more evident in speech-impaired people as the speech characteristics also depend on the type and the severity of the disorder. This problem can be addressed through supervised speaker adaptation (SA), where transcribed utterances of the target speaker are used to adapt a speaker-independent (SI) model to better recognize the speech of the target speaker.

Here we consider the more frequent unsupervised SA scenario, where the transcriptions of the adaptation utterances are not available. Specifically, we focus on unsupervised SA for spoken command recognition of dysarthric speech. Dysarthria is a motor speech disorder consisting in the disruption of the normal control of the vocal tract musculature, common in conditions such as Parkinson Disease (PD), Amyotrophic Lateral Sclerosis (ALS) and post-stroke motor impairments [1, 2]. Dysarthric patients tend to have mumbled speech, acceleration or deceleration in the speaking rate, abnormal pitch and rhythm. However, dysarthric speech is characterized by high variability

as the impairment varies based on the location and severity of the neurological damage. This, beside the limited amount of available data, makes dysarthric speech recognition very challenging. Indeed, although many studies improved the ASR performance on dysarthric speech, they still poorly perform on low intelligible speech [3, 4, 5, 6]. That points out the necessity of *ad hoc* SA techniques for dysarthric speakers.

In [7], authors suggest that different adaptation techniques should be adopted based on the impairment severity of the speaker. They employ two large datasets to train two different SI models. The first dataset only includes healthy speakers, while the second dataset contains additional recordings from dysarthric speakers. Adaptation is then performed to better recognize the utterances from a dysarthric corpus. Results show that adaptation methods for dysarthric patients with mild or moderate dysarthria are more effective when the initial model is trained on healthy speech, whereas the adaptation of severely dysarthric patients improves when the SI model is trained on impaired-speech. Following this direction, [8] introduces a two-stage adaptation approach in which a SI model is first adapted to the general speaking style of multiple dysarthric speakers, and then the adapted model is further adapted to the speaker of interest. However, the described techniques require a large amount of data for fine-tuning, that is usually not available. To overcome this problem, [9] proposes to only fine-tune a subset of the network layers to better adapt an ASR model to the dysarthric speech. Differently, [10] adopts confusion matrix transducers within the weighted finite state transducer decoding paradigm, to handle phonetic-articulatory errors made by the target speakers. However, all the aforementioned studies are limited to the supervised framework, whereas the unsupervised speaker adaptation is the most common scenario.

We adopt an unsupervised SA approach based on Optimal Transport (OT) [11, 12], named Multi-Source Domain Adaptation Weighted Joint Distribution Optimal Transport (MSDA-WJDOT), firstly introduced in [13] for other applications. In this work, we assume to have access to spoken commands of several speakers (called *source*) and aim at finding a classification function for spoken commands of a new speaker (called *target*) through a two-step procedure. The first step learns a global feature extractor from a SI model, e.g., the last hidden layer of a neural network. The second step learns a target-specific classifier, from the feature space to the output space (i.e., the commands to classify). Specifically, we define i) a weighted combination (with weights α) of the joint probability distributions of the source speakers on the feature-output product space; ii) a proxy joint probability distribution of the target speaker on the feature-output product space, in which the output space is represented by the predictions of the target classifier. The algorithm is trained to jointly look for the optimal weights α and the optimal target classifier that minimize the Wasserstein distance between the distributions defined in i) and ii). Intuitively, at each

iteration, this method simultaneously selects the sources most similar to the target (by attributing them higher weights) and updates the target classifier to increase this similarity. The advantage of MSDA-WJDOT over the other adaptation approaches is that only the relevant source speakers are used, avoiding redundant or misleading information. Mostly important, the proposed approach offers an interpretable model as we found the source weights α reflecting a similarity between the source speakers and the target speaker in terms of speech characteristics. Contrary to [7] in which the sources are manually chosen based on the impairment severity of the target speaker, MSDA-WJDOT is able to perform this selection automatically considering also other important similarity aspects (e.g., the gender).

2. MSDA-WJDOT

In this section, we recall MSDA-WJDOT, firstly proposed in [13]. In the classical machine learning problem, the aim is to learn a feature extractor and a classifier, such that the combination of these two functions describes the input-output relationship. MSDA-WJDOT assumes to have access to the feature extraction function g from the input space to an embedding space \mathcal{G} and looks for a target classifier defined on \mathcal{G} . This is the case in which, for example, a pre-trained network is used to extract meaningful features. Once g is known, MSDA-WJDOT performs adaptation on the feature-output product space. However, in the considered framework the feature extractor is not available and, hence, the proposed method becomes a two-step procedure, as described in the following.

2.1. Global feature extractor

The first step consists in estimating a global feature extractor, e.g., a function that extracts features from the audio of a *general* speaker. To do that, we learn a SI model, composed by a feature extraction function g and a classifier f_{SI} . Let us suppose to have N_j utterances (X_j, Y_j) of J source speakers, representing the utterance acoustic features X_j and the corresponding command label Y_j . We can train the model on the source speakers $\{(X_j, Y_j)\}_{j=1}^J$ by

$$\min_{g, f_{SI}} \sum_{j=1}^J \frac{1}{N_j} \sum_{i=1}^{N_j} \mathcal{L}(f_{SI} \circ g(x_j^i), y_j^i), \quad (1)$$

where \mathcal{L} is the cross-entropy loss function. Note that in the second step, f_{SI} will not be used.

2.2. Adaptation

Once g is given, we can define the joint distributions p_{S_j} , for each source speaker j , with support on the product space $\mathcal{G} \times \mathcal{Y}$, where \mathcal{Y} is the label space. We then define a convex combination of the source distributions

$$p_S^\alpha = \sum_{j=1}^J \alpha_j p_{S_j}$$

with α element of the simplex Δ^J (i.e., $0 \leq \alpha_j \leq 1$ for each j and $\sum_{j=1}^J \alpha_j = 1$). Finally, we define the proxy target joint distribution p_T^f , with support on $\mathcal{G} \times f(\mathcal{G})$, where f is the target classifier to learn. Here, the upper index f indicates that the joint distribution is over predicted labels and not reference labels. MSDA-WJDOT aims at finding the α weights and the

classifier f that align the source combination of distributions p_S^α and the proxy target distribution p_T^f . This can be expressed as

$$\min_{\alpha, f} W_D(p_T^f, p_S^\alpha) = \min_{\alpha, f} W_D\left(p_T^f, \sum_{j=1}^J \alpha_j p_{S_j}\right), \quad (2)$$

where W_D is the Wasserstein distance [11, 14]. Roughly speaking, the source joint distributions can be seen as proxy joint distributions with support on $\mathcal{G} \times f_j(\mathcal{G})$, where f_j is the true classifier (i.e., $f_j(X_j) \equiv Y_j$). This approach dynamically selects the most similar sources and re-adapts f by “imitating” the best f_j .

2.3. The role of α and the similarity

We can consider α as a measure of similarity between the sources and the target, in terms of joint probability distribution. Indeed, in order to minimize Eq. 2, the algorithm will attribute higher weights to the source distributions closest to the proxy target distribution, and smaller (or zero) weights to far source distributions. The zero-weighted distributions will not contribute to the adaptation procedure. Consequently, misleading and unrelated source domains are in practice not used and this allows to learn a more accurate classifier. For example, if α were a one-hot vector with $\alpha_j = 1$, the method would consist in estimating f such that $p_T^f \simeq p_{S_j}$. Authors in [13] only offer a mathematical explanation of the weight alpha. In this study, we explore the interpretation of alpha in terms of speaker similarity and speech characteristics.

3. Experimental setup

3.1. Dataset

We employ the AllSpeak dataset [15], that consists of speech recordings from 29 Italian native speakers. Seventeen of these (thirteen males, four females) are affected by Amyotrophic Lateral Sclerosis, while the remaining twelve (six males, six females) are healthy control speakers. The dataset contains 25 commands in Italian, relative to basic needs such as “I am thirsty”, “I am hungry”. This dataset is very challenging due to the small amount of recordings. Indeed, only 2387 and 1857 examples have been recorded from control and dysarthric speakers, respectively. Specifically, each command has been recorded from 8 to 10 times by healthy speakers, while from 4 to 10 times by patients – depending on the patient medical condition. In the following we will adopt a speaker coding according to which F/M stands for female/male, and the letter C is added to refer to control subjects.

We perform speaker adaptation of each dysarthric speaker (target speaker) by using all the remaining speakers as source speakers. We trained the SI model on the source speakers, by splitting the dataset into training and validation sets. Specifically, the validation set consists of one example of each command of each speaker and it has been used to perform early stopping. The unlabelled target speaker data is split into adaptation set (80%) and testing (20%) set, which contains one example of each command. To train MSDA-WJDOT, we simultaneously employ the source training dataset and the adaptation dataset of the target speaker.

3.2. Model details

The feature extraction function g is represented by a Bidirectional Long Short-Term Memory (BLSTM) with five hidden

Table 1: *Command Error Rate (CER) for each dysarthric target speaker provided by the SI, SSA, MSDA-WJDOT and the competitor models.*

Speaker	SI	CJDOT [16]	MJDOT [16]	MSDA-WJDOT	SSA [17]
M01	35.79	32.77	31.93	31.93	25.00
M02	34.26	36.75	36.75	37.71	54.17
F01	63.16	52.63	49.12	49.12	60.00
F02	48.50	40.00	40.00	40.00	36.00
M03	64.44	68.89	71.11	57.78	55.56
M04	30.00	31.20	32.00	30.40	24.00
M05	18.62	17.46	15.08	14.29	17.39
F03	68.33	61.74	70.43	62.61	56.52
M06	48.67	34.78	33.91	35.65	26.09
M07	11.00	7.20	8.00	8.80	20.00
M08	39.50	36.00	41.60	33.60	32.00
M09	24.79	16.81	18.49	19.33	13.04
F04	48.07	38.60	38.60	38.60	40.91
M10	18.00	12.80	12.00	12.80	16.00
M11	56.50	47.20	48.80	45.60	40.00
M12	7.50	5.60	7.20	4.80	4.00
M13	30.91	45.45	43.64	21.82	9.09
A. CER	38.11	34.46	34.37	32.05	31.16

layers, each containing 250 memory blocks. Both the SI classifier f_{SI} and the target classifier f are represented by a softmax layer. All weights are initialized with Xavier initialization. Both SI model training and adaptation are performed with Adam optimizer with 0.9 momentum and $\epsilon = e^{-8}$. Learning rate is fixed to 0.001.

4. Results

4.1. Competitor models

We compare the proposed approach with an unadapted speaker-independent (SI) model, in which the target speaker is directly tested on a model trained on all source speakers. Notice that the literature lacks unsupervised MSDA specific for dysarthric speaker adaptation. Hence, we implemented two MSDA-WJDOT approaches - CJDOT and MJDOT (details can be found in [13]) - which are different extensions of a state-of-the-art unsupervised domain adaptation method - JDOT [16] - adapted to the multi-speaker adaptation case. In addition, we trained a supervised SA (SSA) model, in which we added a feed-forward linear layer atop the input to the SI model and trained it on the target speaker [17]. As this method uses the target labels, it must be considered as the lower bound of the unsupervised SA performance error.

4.2. Command speech recognition

Table 1 reports the results in terms of Command Error Rate (CER). A first remark is that, although the SI model always achieved a CER between 15% and 20% on the source validation set, it often had low accuracy on the target speaker. Once again, this emphasizes the difficulty of an ASR system to generalize to a new dysarthric speaker and the importance of the speaker adaptation in this context.

The unsupervised SA carried out by MSDA-WJDOT outperforms all the methods by providing the best Average CER. Indeed, it reduces the CER of 16% and 7% over the SI system and the best competitor model, respectively. Surprisingly, MSDA-WJDOT achieves an Average CER similar to the SSA approach, in which the labels are used.

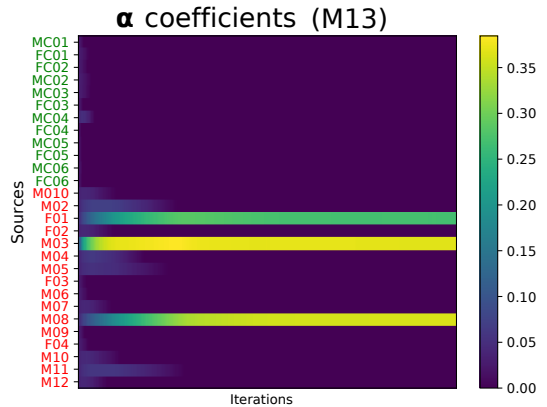


Figure 1: α coefficients recovered during MSDA-WJDOT training for dysarthric target speaker M13. The α_j coefficients are close to zero for healthy speakers (in green), while the highest weights are attributed to dysarthric speakers (in red).

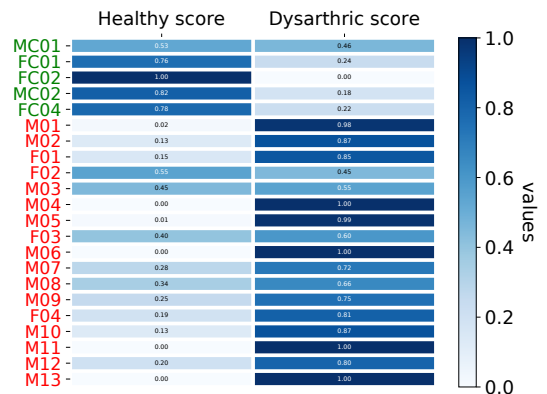


Figure 2: HS and DS computed for healthy (in green) and dysarthric (in red) speakers.

It is essential to recall that the proposed approach provides a measure of similarity between the source speakers and the target speaker, in terms of probability distribution. We found that the recovered α always attributes highest similarity scores to dysarthric speakers rather than healthy ones, when the target speaker is dysarthric, and vice versa for healthy target speakers. This means that α weights can catch the similarity in terms of speech characteristics and that this approach can realistically estimate speaker similarity. Fig. 1 shows the recovered weights for a dysarthric target speaker (M13) during the training. Specifically, we report on the x -axis the epochs of training, while the y -axis lists the healthy (in green) and dysarthric (in red) source speakers. As we can see, after few epochs α_j becomes sparse, using in practice only the data of the most similar speakers. Interestingly, all the healthy speakers have zero weights whereas the α mass is spread along dysarthric speakers. Also, the highest weights are attributed to M03 and M08, which are both males, while a lower importance is attributed to F01. This suggests that MSDA-WJDOT also incorporates similarity information about the gender.

4.3. The α weight and the dysarthria detection

We here propose a completely novel approach to perform dysarthria detection based on the α weights. As shown in Fig. 1, MSDA-WJDOT associates speakers with similar voice characteristics. We investigated the possibility of leveraging the α weights of the command classifier to infer if the target speaker is affected by dysarthria. Specifically, we propose to classify a speaker as healthy or dysarthric based on his/her similarity with the other subjects.

Let define I_c as the set indexing the control speakers and I_d as the set of indices related to dysarthric speakers. We then define the Healthy Score (HS) and the Dysarthric Score (DS) as follow

$$HS = \sum_{j \in I_c} \alpha_j, \quad DS = \sum_{j \in I_d} \alpha_j. \quad (3)$$

As $\alpha \in \Delta^J$, we have that $HS, DS \in [0, 1]$ such that $HS + DS = 1$. Fig. 2 reports the computed scores for all dysarthric speakers and for 5 control speakers. The control subjects always show a higher HS , while for the patients (except for F02), we have a higher DS . Moving a step forward, we can use these scores to perform dysarthria detection by stating that

A speaker is affected by dysarthria if $DS > HS$.

This results in a final accuracy of 95%. It is crucial to emphasize that, differently from standard dysarthria detection approaches [18, 19, 20], any specific training has been performed: we directly infer the speaker health status from the similarity coefficients α that provide us an interpretable model. Note that none of existing SA methods are able to also perform dysarthria detection. Hence, we cannot fairly compare our approach to any competitor model.

5. Conclusions

In this work, we addressed unsupervised speaker adaptation in the challenging context of pathological speech by employing MSDA-WJDOT algorithm. Our approach significantly improves the speaker-independent system, reducing the CER of 16%, and it also outperforms all the competitor models. The advantage of the proposed method is that it does not require the labels for the target speaker and, interestingly, it also provides source-target probability distribution similarity coefficients α , reflecting the source-target speakers characteristics similarity. From these coefficients, we derived the Healthy (HS) and Dysarthric (DS) Scores of a target speaker that allowed us to detect dysarthria with an accuracy of 95%.

The novelty of our work relies on several aspects. Firstly, this work brings a contribution in the field of dysarthric speaker adaptation, in which unsupervised methods dramatically lack while they are essential for a good recognition accuracy. The second type of novelty is the interpretation of the method. To the best of our knowledge, this is the first SA study in which the model provides a measure of similarity between speakers. Note that this approach may be transferred to other SA contexts in which, for instance, the algorithm may exploit speakers accents and provide a similarity in terms of region of origin of the speakers. Last but not least, we moved a step forward and we introduced new ranking indices that we used to perform dysarthria detection. This approach is completely new and can achieve high diagnosis performance, without requiring a specific model training.

Future directions could delve into the dysarthria assessment by individuating intervals of values in which the DS cor-

responds to dysarthria severity levels (e.g., mild, moderate, severe). This may lead to very efficient ASR systems that simultaneously improve their performance via SA, and compute the DS warning the subject when the score is close to the right endpoint of its interval. Such a device could perform an early detection of the disorder or predict the disease degeneration and allow the patient to act in time in order to prevent it. Furthermore, the dysarthric score may be used to track the speaker changes during a treatment and monitoring the benefits of the rehabilitation.

Finally, our MSDA-WJDOT approach was applied on a limited resource scenario where the initial unadapted speaker independent model was trained on the AllSpeak training set. In future work we will experiment with speaker independent models whose hidden layers are pre-trained on much larger datasets (e.g., datasets used for ASR).

6. References

- [1] M. R. McNeil, *Clinical management of sensorimotor speech disorders*. Thieme, 2009.
- [2] F. L. Darley, A. E. Aronson, and J. R. Brown, "Differential diagnostic patterns of dysarthria," *Journal of speech and hearing research*, vol. 12, no. 2, pp. 246–269, 1969.
- [3] C. Espana-Bonet and J. A. R. Fonollosa, "Automatic speech recognition with deep neural networks for impaired speech," in *International Conference on Advances in Speech and Language Technologies for Iberian Languages*. Springer, 2016, pp. 97–107.
- [4] B. Vachhani, C. Bhat, B. Das, and S. K. Kopparapu, "Deep auto-encoder based speech features for improved dysarthric speech recognition," in *Interspeech*, 2017, pp. 1854–1858.
- [5] N. M. Joy and S. Umesh, "Improving acoustic models in torgo dysarthric speech database," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 26, no. 3, pp. 637–645, 2018.
- [6] X. Xie, R. Ruzi, X. Liu, and L. Wang, "Variational auto-encoder based variability encoding for dysarthric speech recognition," *Proc. Interspeech 2021*, pp. 4808–4812, 2021.
- [7] M. B. Mustafa, S. S. Salim, N. Mohamed, B. Al-Qatab, and C. E. Siang, "Severity-based adaptation with limited data for asr to aid dysarthric speakers," *PloS one*, vol. 9, no. 1, p. e86285, 2014.
- [8] R. Takashima, T. Takiguchi, and Y. Ariki, "Two-step acoustic model adaptation for dysarthric speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6104–6108.
- [9] J. Shor, D. Emanuel, O. Lang, O. Tuval, M. Brenner, J. Cattiau, F. Vieira, M. McNally, T. Charbonneau, M. Nollstadt *et al.*, "Personalizing asr for dysarthric and accented speech with limited data," *arXiv preprint arXiv:1907.13511*, 2019.
- [10] S. O. C. Morales and S. J. Cox, "Modelling errors in automatic speech recognition for dysarthric speakers," *EURASIP Journal on Advances in Signal Processing*, vol. 2009, no. 1, p. 308340, 2009.
- [11] G. Monge, "Mémoire sur la théorie des déblais et de remblais," in *Histoire de l'Académie Royale des Sciences de Paris, avec les Mémoires de Mathématique et de Physique pour la même année*, 1781.
- [12] L. V. Kantorovich, "On the translocation of masses," in *Journal of Mathematical Sciences*, 2006.
- [13] R. Turrisi, R. Flamary, A. Rakotomamonjy, and M. Pontil, "Multi-source domain adaptation via weighted joint distributions optimal transport," *arXiv preprint arXiv:2006.12938*, 2020.
- [14] A. Genevay, G. Peyré, and M. Cuturi, "Learning generative models with sinkhorn divergences," *arXiv preprint arXiv:1706.00292*, 2017.
- [15] C. Di Nardi, R. Turrisi, A. Inuggi, N. Riva, I. Mauri, and L. Badino, "An automatic speech recognition android app for als patients," *Collana AISV Speech in the natural context*, 2018.

- [16] N. Courty, R. Flamary, A. Habrard, and A. Rakotomamonjy, "Joint distribution optimal transportation for domain adaptation," in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. Curran Associates, Inc., 2017, pp. 3730–3739. [Online]. Available: <http://papers.nips.cc/paper/6963-joint-distribution-optimal-transportation-for-domain-adaptation.pdf>
- [17] J. Neto, L. Almeida, M. Hochberg, C. Martins, L. Nunes, S. Renals, and T. Robinson, "Speaker-adaptation for hybrid hmm-ann continuous speech recognition system," *EUROSPEECH*, 1995.
- [18] J. R. Williamson, T. F. Quatieri, B. S. Helfer, J. Perricone, S. S. Ghosh, G. Ciccarelli, and D. D. Mehta, "Segment-dependent dynamics in predicting parkinson's disease," in *Sixteenth annual conference of the international speech communication association*, 2015.
- [19] M. Tu, V. Berisha, and J. Liss, "Interpretable objective assessment of dysarthric speech based on deep neural networks." in *INTERSPEECH*, 2017, pp. 1849–1853.
- [20] A. Hernandez, E. J. Yeo, S. Kim, and M. Chung, "Dysarthria detection and severity assessment using rhythm-based metrics," in *Proceedings of the Annual Conference of the International Speech Communication Association (INTERSPEECH)*, Shanghai, China, 2020, pp. 25–29.