

Match-And-Deform: Time Series Domain Adaptation through Optimal Transport and Temporal Alignment

François Painblanc¹, Laetitia Chapel², Nicolas Courty², Chloé Friguet²,
Charlotte Pelletier², and Romain Tavenard¹

¹ Université Rennes 2, LETG, IRISA, Rennes, France

² Université Bretagne Sud, IRISA, UMR CNRS 6074, Vannes, France

Abstract. While large volumes of unlabeled data are usually available, associated labels are often scarce. The unsupervised domain adaptation problem aims at exploiting labels from a source domain to classify data from a related, yet different, target domain. When time series are at stake, new difficulties arise as temporal shifts may appear in addition to the standard feature distribution shift. In this paper, we introduce the Match-And-Deform (MAD) approach that aims at finding correspondences between the source and target time series while allowing temporal distortions. The associated optimization problem simultaneously aligns the series thanks to an optimal transport loss and the time stamps through dynamic time warping. When embedded into a deep neural network, MAD helps learning new representations of time series that both align the domains and maximize the discriminative power of the network. Empirical studies on benchmark datasets and remote sensing data demonstrate that MAD makes meaningful sample-to-sample pairing and time shift estimation, reaching similar or better classification performance than state-of-the-art deep time series domain adaptation strategies.

Keywords: domain adaptation, time series, optimal transport, dynamic time warping.

Introduction

A standard assumption in machine learning is that the training and the test data are drawn from the same distribution. When this assumption is not met, trained models often have degraded performances because of their poor generalization ability. Domain adaptation (DA) is the branch of machine learning that tackles this generalisation problem when the difference in distribution can be expressed as a shift, allowing for improving task efficiency on a target domain by using all available information from a source domain. When dealing with time series data, that are ubiquitous in many real-world situations, being able to learn across time

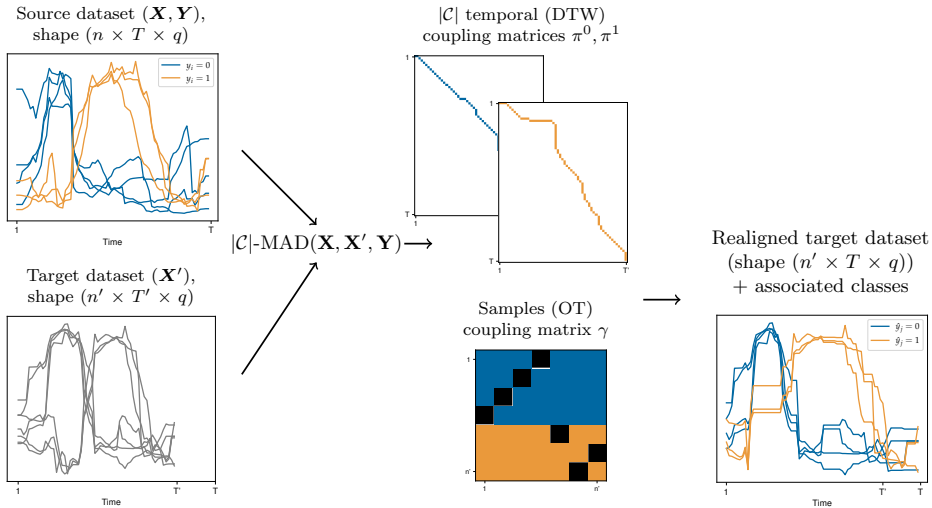


Fig. 1. Match-And-Deform ($|\mathcal{C}|$ -MAD) takes two time series datasets as inputs: a source (labelled) dataset and a target (unlabelled) dataset. It jointly computes an optimal transport (OT) coupling matrix γ and $|\mathcal{C}|$ class-wise dynamic time warping (DTW) paths $\{\pi^{(c)}\}_{c \in \mathcal{C}}$. The OT cost is derived from the pairwise distances yielded by the DTW paths while the DTW cost is weighted by the OT plan. These outputs are then used to improve classification in the target dataset.

series domains is a challenging task as temporal deformations between domains might occur in addition to the feature distribution shift.

Most of the literature in DA focuses on bridging the distribution shift, either by aligning distributions or by finding a common representation space [20,16]. In unsupervised DA frameworks, training is performed on the source domain data relying on this common representation so that effective performance on the target domain can be expected. A standard approach consists in using adversarial training to push a deep neural network model into learning domain-invariant intermediate representations [7]. This approach has been extended to the time series context by using dedicated network architectures. In VRADA [15], variational recurrent neural networks are considered while in the state-of-the-art method CoDATS [21], the feature extraction block is made of convolutional layers. However, these methods operate on pooled features in which the time dimension has vanished and, as such, ignore the impact of potential temporal dynamics across multiple yet related time series.

Optimal transport (OT) is an efficient tool for DA, in both unsupervised and semi-supervised learning. It can be used to evaluate the distribution shift as a non-linear function [3] or, as done in more recent works, it can help design deep neural network losses that take into account the dissimilarity between the domains [4]. OT-based DA methods have proven useful in many applications but

do not encode any temporal coherence that should be kept when dealing with time series data.

The approach presented in this paper, coined Match-And-Deform (MAD), combines OT with dynamic time warping (DTW) to achieve time series matching and timestamp alignment. This yields a dissimilarity measure that relies on a temporally realigned sample-to-sample matching. In other words, MAD evaluates the feature distribution shift between domains up to a global temporal alignment. MAD can then be used as a loss function in a neural network to learn a domain-invariant latent representation.

In the following, the background on DTW and OT are introduced. MAD and its use as a loss for DA are then presented and experimentally compared to standard competitors, considering classification tasks on benchmark time series datasets but also on remote sensing time series datasets as an application framework. Finally, we discuss related works and present perspectives.

Background

In the sequel, we aim at finding a coupling between samples of two datasets, up to a global temporal shift. For that purpose, we jointly solve two alignment problems: (i) a temporal realignment (of the timestamps) of the datasets and (ii) a matching between time series (samples). Both rely on related linear optimization problems that differ in the nature of the coupling involved.

In a general way, the optimisation problem for comparing two objects (either time series or distributions) \mathbf{x} and \mathbf{x}' can be stated as:

$$J(\mathbf{C}(\mathbf{x}, \mathbf{x}'), \Pi) = \arg \min_{\pi \in \Pi} \langle \mathbf{C}(\mathbf{x}, \mathbf{x}'), \pi \rangle, \quad (1)$$

in which Π is a set of admissible couplings. A coupling will either be a temporal alignment if \mathbf{x} and \mathbf{x}' are time series or a matching between samples if they are distributions, with appropriate constraint sets, as detailed later in this section. The solution $J(\cdot, \cdot)$ of the optimization problem is called the optimal coupling matrix. The cost matrix $\mathbf{C}(\mathbf{x}, \mathbf{x}') = \{d(x^i, x'^j)\}_{ij}$ stores distances $d(x^i, x'^j)$ between atomic elements x^i and x'^j , respectively from \mathbf{x} and \mathbf{x}' .

Dynamic time warping (DTW) is an instance of problem (1) whose goal is to match two (multivariate) time series $\mathbf{x} \in \mathbb{R}^{T \times q}$, and $\mathbf{x}' \in \mathbb{R}^{T' \times q}$ as:

$$\text{DTW}(\mathbf{x}, \mathbf{x}') = J(\mathbf{C}(\mathbf{x}, \mathbf{x}'), \mathcal{A}(T, T')). \quad (2)$$

Here, $\mathbf{C}(\mathbf{x}, \mathbf{x}')$ stores Euclidean distances between q -dimensional atomic elements and $\mathcal{A}(T, T')$ is the set of admissible timestamps alignments between series of respective lengths T and T' . An admissible alignment $\pi \in \mathcal{A}(T, T')$ is a binary $T \times T'$ -matrix that matches timestamps between time series \mathbf{x} and \mathbf{x}' . Its non-zero entries should form a connected path between cells $(1, 1)$ and (T, T') . This

definition allows for efficient computation in quadratic time using dynamic programming [18]. Note that DTW can be seen as seeking a temporal transformation (in the form of repeated samples) such that the Euclidean distance between transformed series is minimized.

Optimal transport (OT) is another instance of the same general optimization problem (1) that defines a distance between two probability measures. The discrete probability measures \mathbf{X} and \mathbf{X}' are sets of weighted samples in \mathbb{R}^q : $\{(\mathbf{x}^i, w^i)\}_{i=1}^n$ with $\sum_i w^i = 1$ and $\{(\mathbf{x}'^j, w'^j)\}_{j=1}^{n'}$ with $\sum_j w'^j = 1$. When no prior information is available, weights are set uniformly. OT defines a distance between \mathbf{X} and \mathbf{X}' by seeking the transport plan $\gamma \in \Gamma(\mathbf{w}, \mathbf{w}')$ that minimises the transport cost:

$$\text{OT}(\mathbf{X}, \mathbf{X}') = J\left(\mathbf{C}(\mathbf{X}, \mathbf{X}'), \Gamma(\mathbf{w}, \mathbf{w}')\right). \quad (3)$$

$\Gamma(\mathbf{w}, \mathbf{w}')$ is the set of admissible transport plans, *i.e.* the set of linear transports such that all the mass from \mathbf{X} is transported toward all the mass of \mathbf{X}' :

$$\Gamma(\mathbf{w}, \mathbf{w}') = \left\{ \gamma \mid \gamma \geq \mathbf{0}, \gamma \mathbf{1}_{n'} = \mathbf{w}, \gamma^\top \mathbf{1}_n = \mathbf{w}' \right\}$$

with $\mathbf{1}_n$ a n -vector of ones. The transport plan γ is a $n \times n'$ -matrix whose elements γ_{ij} indicate the amount of mass *transported* from x^i to x'^j . The transport plan is sparse, with at most $n+n'-1$ elements that are non-zeros. The most common algorithmic tools to solve the discrete OT problem are borrowed from combinatorial optimisation and linear programming (see [14] for a thorough review).

Match-And-Deform (MAD)

In an unsupervised time series domain adaptation context, let us consider a source dataset $(\mathbf{X}, \mathbf{Y}) \in \mathbb{R}^{n \times T \times q} \times \mathcal{C}^n$ and a dataset from a target domain $\mathbf{X}' \in \mathbb{R}^{n' \times T' \times q}$. $\mathbf{X} = (x_1^i; \dots; x_T^i)_{i=1}^n$ and $\mathbf{X}' = (x_1'^j; \dots; x_{T'}'^j)_{j=1}^{n'}$ represent multidimensional time series data ($q \geq 1$) with lengths T and T' , respectively. Both domains share the same label space \mathcal{C} but only source labels \mathbf{Y} are observed. Our goal is to classify target data, using knowledge transferred from source data.

Despite its widespread use, DTW is limited to finding a temporal alignment between two single time series $\mathbf{x} \in \mathbb{R}^{T \times q}$, and $\mathbf{x}' \in \mathbb{R}^{T' \times q}$, without any consideration for the alignment of sets of series. On the other hand, OT is designed to match datasets regardless of any temporal dimension. A simple approach to combine both matchings is to use DTW as the inner cost $\mathbf{C}(\mathbf{X}, \mathbf{X}')$ of an OT problem. However, we argue that this would result in spurious matchings, due to the many degrees of freedom introduced by individual DTW computations (see Experiments section). Instead, we introduce a new metric, coined Match-And-Deform (MAD), that jointly optimizes a global DTW alignment and an OT coupling to match two sets of time series, as illustrated in Fig. 1. Let us therefore

define MAD as:

$$\begin{aligned} \text{MAD}(\mathbf{X}, \mathbf{X}') &= \arg \min_{\substack{\gamma \in \Gamma(\mathbf{w}, \mathbf{w}') \\ \pi \in \mathcal{A}(T, T')}} \langle \mathbf{L}(\mathbf{X}, \mathbf{X}') \otimes \pi, \gamma \rangle \\ &= \arg \min_{\substack{\gamma \in \Gamma(\mathbf{w}, \mathbf{w}') \\ \pi \in \mathcal{A}(T, T')}} \sum_{i,j} \sum_{\ell,m} d(x_\ell^i, x_m^j) \pi_{\ell m} \gamma_{ij}. \end{aligned} \quad (4)$$

Here, $\mathbf{L}(\mathbf{X}, \mathbf{X}')$ is a 4-dimensional tensor whose elements are $L_{\ell,m}^{i,j} = d(x_\ell^i, x_m^j)$, with $d: \mathbb{R}^q \times \mathbb{R}^q \rightarrow \mathbb{R}^+$ being a distance. \otimes is the tensor-matrix multiplication. π is a global DTW alignment between timestamps and γ is a transport plan between samples from \mathbf{X} and \mathbf{X}' .

By combining metrics that operate between time series and samples in a single optimization problem, MAD gets the best of both worlds: it seeks a transport plan γ that matches samples up to a global temporal transformation π (in the form of repeated samples, see the DTW presentation in Background section), hence defining a metric between datasets that is invariant to dataset-wide time shifts. Note that MAD seeks for a global temporal alignment between datasets, which implies that all series inside a dataset have the same length, otherwise matching timestamps through π would be meaningless.

The optimization problem in Eq. (4) can be further extended to the case of distinct DTW mappings for each class c in the source data. This results in the following optimization problem, coined $|\mathcal{C}|$ -MAD:

$$|\mathcal{C}|\text{-MAD}(\mathbf{X}, \mathbf{X}', \mathbf{Y}) = \arg \min_{\substack{\gamma \in \Gamma(\mathbf{w}, \mathbf{w}') \\ \forall c, \pi^{(c)} \in \mathcal{A}(T, T')}} \sum_{i,j} \sum_{\ell,m} L_{\ell,m}^{i,j} \pi_{\ell m}^{(y^i)} \gamma_{ij}. \quad (5)$$

In that case, $|\mathcal{C}|$ DTW alignments are involved, one for each class c . $\pi^{(y^i)}$ denotes the DTW matrix associated to the class y^i of x^i . This more flexible formulation allows adapting to different temporal distortions that might occur across classes.

Properties. Let us now study some of the properties of this new similarity measure. Our first property links Optimal Transport, MAD and $|\mathcal{C}|$ -MAD problems.

Property 1. Let $\mathbf{X} \in \mathbb{R}^{n \times T \times q}$ and $\mathbf{X}' \in \mathbb{R}^{n' \times T' \times q}$ be time series datasets. Let $\text{OT}_{\text{DTW}}(\mathbf{X}, \mathbf{X}')$ be the solution of the OT problem with cost $\mathbf{C}(\mathbf{X}, \mathbf{X}') = \{\text{DTW}(\mathbf{x}, \mathbf{x}')\}_{\mathbf{x}, \mathbf{x}' \in \mathbf{X}, \mathbf{X}'}$, and let us denote $\text{cost}(\cdot)$ the cost associated to a solution of any optimization problem. We have:

$$\text{cost}\left(\text{OT}_{\text{DTW}}(\mathbf{X}, \mathbf{X}')\right) \leq \text{cost}\left(|\mathcal{C}|\text{-MAD}(\mathbf{X}, \mathbf{X}', \mathbf{Y})\right) \leq \text{cost}\left(\text{MAD}(\mathbf{X}, \mathbf{X}')\right)$$

Moreover, MAD inherits properties from the OT field. Typically, the transport plans γ resulting from MAD and $|\mathcal{C}|$ -MAD are sparse. If $n = n'$ and uniform weights are considered, there always exists a transport plan solution of the MAD (resp. $|\mathcal{C}|$ -MAD) problem that is a permutation matrix, as stated in the following property.

Algorithm 1 $|\mathcal{C}|$ -MAD optimization via BCD

```

1: Input: weighted time series datasets  $\mathbf{X}$  and  $\mathbf{X}'$ , initial DTW paths  $\{\boldsymbol{\pi}^{(c)}\}$ 
2: repeat
3:   Compute  $\mathbf{C}_{\text{OT}}(\mathbf{X}, \mathbf{X}', Y, \{\boldsymbol{\pi}^{(c)}\}_c)$  using Eq. (6)
4:    $\boldsymbol{\gamma} \leftarrow J\left(\mathbf{C}_{\text{OT}}(\mathbf{X}, \mathbf{X}', Y, \{\boldsymbol{\pi}^{(c)}\}_c), \Gamma(\mathbf{w}, \mathbf{w}')\right)$ 
5:   for  $c = 1 \dots |\mathcal{C}|$  do
6:     Compute  $\mathbf{C}_{\text{DTW}}^c(\mathbf{X}, \mathbf{X}', \boldsymbol{\gamma})$  using Eq. (7)
7:      $\boldsymbol{\pi}^{(c)} \leftarrow J\left(\mathbf{C}_{\text{DTW}}^c(\mathbf{X}, \mathbf{X}', \boldsymbol{\gamma}), \mathcal{A}(T, T')\right)$ 
8:   end for
9: until convergence

```

Property 2. Let \mathbf{X} and \mathbf{X}' be datasets each composed of n time series, and let us assume uniform weights, *i.e.* $\mathbf{w} = \mathbf{w}' = (1/n, \dots, 1/n)$. There exists a transport plan solution to the MAD (resp. $|\mathcal{C}|$ -MAD) problem that is a one-to-one matching, *i.e.* each sample from \mathbf{X} is matched to exactly one sample in \mathbf{X}' (and conversely).

Proofs for these properties are provided as Supplementary Material.

Optimization. Let us consider the joint optimization problem introduced in Eq. (5), which involves $|\mathcal{C}|$ finite sets of admissible DTW paths and a continuous space with linear constraints for the OT plan. Extension to solving the problem of Eq. (4) is straightforward. We perform a Block Coordinate Descent (BCD) to optimize the corresponding loss. BCD is a simple iterative algorithm for non-convex optimization problems in which one set of parameters is optimized while the others are held fixed at each step; this algorithm has already been used in similar contexts, *e.g.* in [17]. In our case, one of the involved maps is optimized with all the other ones fixed, giving $|\mathcal{C}| + 1$ intertwined optimization problems in the form of Eq. (1). Indeed for a given set of DTW paths $\{\boldsymbol{\pi}^{(c)}\}_{c=1}^{|\mathcal{C}|}$, the problem in Eq. (5) boils down to a linear OT problem with a cost matrix defined as:

$$\mathbf{C}_{\text{OT}}(\mathbf{X}, \mathbf{X}', \mathbf{Y}, \{\boldsymbol{\pi}^{(c)}\}_c) = \left\{ \sum_{\ell, m} L_{\ell, m}^{i, j} \pi_{\ell m}^{(y^i)} \right\}_{i, j}. \quad (6)$$

Similarly, for a fixed transport plan $\boldsymbol{\gamma}$, solving the problem in Eq. (5) with respect to a given class-specific DTW path $\boldsymbol{\pi}^{(c)}$ is a DTW problem in which the cost matrix is:

$$\mathbf{C}_{\text{DTW}}^c(\mathbf{X}, \mathbf{X}', \boldsymbol{\gamma}) = \left\{ \sum_{i \text{ s.t. } y^i = c, j} L_{\ell, m}^{i, j} \gamma_{ij} \right\}_{\ell, m}. \quad (7)$$

Hence, it leads to optimizing independently $|\mathcal{C}|$ DTW alignment problems.

The resulting BCD optimization is presented in Algorithm 1. Note that in each update step, we get an optimal value for the considered map given the other

ones fixed, hence the sequence of $|\mathcal{C}|$ -MAD losses is both decreasing and lower bounded by zero, thus converging to a local optimum.

Neural domain adaptation with a MAD loss

OT has been successfully used as a loss to measure the discrepancy between source and target domain samples embedded into a latent space. Similarly to DeepJDOT [4], our proposal considers a deep unsupervised temporal DA model that relies on MAD or $|\mathcal{C}|$ -MAD as a regularization loss function.

The deep neural network architecture of DeepJDOT is composed of two parts: (i) an embedding function g_Ω that maps the inputs into a given latent space, and (ii) a classifier f_θ that maps the latent representation of the samples into a label space shared by both source and target domains. OT is applied on the output of the embeddings such that g_Ω yields a discriminant yet domain invariant representation of the data. When dealing with temporal data, we minimize the following overall loss function over $\{\boldsymbol{\pi}^{(c)}\}_c$, $\boldsymbol{\gamma}$, Ω and θ :

$$\mathcal{L}(\mathbf{X}, \mathbf{Y}, \mathbf{X}') = \overbrace{\frac{1}{n} \sum_i \mathcal{L}_s(y^i, f_\theta(g_\Omega(\mathbf{x}^i)))}^{(A)} + \sum_{i,j} \gamma_{ij} \left(\underbrace{\alpha \sum_{\ell,m} \pi_{\ell m}^{(y^i)} L(g_\Omega(\mathbf{X}), g_\Omega(\mathbf{X}'))_{\ell,m}^{i,j}}_{(B)} + \underbrace{\beta \mathcal{L}_t(y^i, f_\theta(g_\Omega(\mathbf{x}'^j)))}_{(C)} \right) \quad (8)$$

where $\mathcal{L}_s(\cdot, \cdot)$ and $\mathcal{L}_t(\cdot, \cdot)$ are cross entropy losses, $\boldsymbol{\gamma}$ (resp. $\{\boldsymbol{\pi}^{(c)}\}_c$) is the transport plan (resp. the set of DTW paths) yielded by $|\mathcal{C}|$ -MAD. The first part (A) of the loss is a classification loss on the source domain; (B) and (C) rely on the OT plan γ_{ij} : (C) seeks to align the labels of the source time series with the predicted labels of their matched ($\gamma_{ij} > 0$) target time series. The difference with DeepJDOT lies in the term (B) of Eq. (8). This part aims at aligning the latent representations of the time series from the source and target domains that have been matched, the main difference here is that the embeddings of \mathbf{x}^i and \mathbf{x}'^j are additionally temporally realigned thanks to the DTW mapping $\boldsymbol{\pi}^{(y^i)}$. Finally, α and β are hyper-parameters to balance terms (B) and (C) in the global loss.

This loss can be minimized to zero if (i) the classifier achieves perfect accuracy on source data, (ii) the distribution of features at the output of the embedding for the source and target domains match up to a global temporal alignment per class, and (iii) source labels and target predictions are matched by the MAD transport plan.

Optimization. The loss function described in Eq. (8) is optimized over two groups of parameters: (i) the neural network parameters θ and Ω and (ii) MAD transport plan $\boldsymbol{\gamma}$ and DTW paths $\{\boldsymbol{\pi}^{(c)}\}_c$. Similar to what is done in [4], we use

an approximate optimization procedure that relies on stochastic gradients. For each mini-batch, a forward pass is performed during which mini-batch-specific γ and $\{\pi^{(c)}\}_c$ maps are estimated using Algorithm 1, for fixed network blocks g_Ω and f_θ . While the DTW cost matrix presented in Eq. (7) remains unchanged, the OT cost gets an extra term from the cross-domain label mapping:

$$\left\{ \alpha \sum_{\ell, m} \pi_{\ell m}^{(y^i)} L(g_\Omega(\mathbf{X}), g_\Omega(\mathbf{X}'))_{\ell, m}^{i, j} + \beta \mathcal{L}_t(y^i, f_\theta(g_\Omega(\mathbf{x}'^j))) \right\}_{i, j} .$$

To initialize the learning process, we use random DTW paths and, when repeating this process for consecutive mini-batches of data, we use the DTW paths from the previous mini-batch as initializers. Note that even though stochastic optimisation can be expected to converge towards optimal parameters Ω and θ , it is not the case for the MAD parameters: it is known from the OT literature [8] that the expected value of the OT plan over the mini-batches does not converge to the full OT coupling. However, authors of [4] claim that the resulting non-sparse estimated plan can act as a regularizer to enforce mass sharing between similar samples.

Complexity. Assuming time series lengths in both datasets are comparable, and relying on the sparsity of γ and $\pi^{(c)}$, computing the cost matrices \mathbf{C}_{OT} and $\{\mathbf{C}_{\text{DTW}}^c\}_c$ can be done in $O(b^2Te + bT^2e)$ where b is the number of series in a mini-batch and e is the dimension of the embedding $g_\Omega(\cdot)$. Given a pre-computed cost matrix, each DTW computation is of complexity $O(T^2)$ and similarly, the complexity for the OT solving step is $O(b^3 \log b)$. The overall complexity for each iteration of Algorithm 1 is then $O(b^2(b \log b + Te) + (be + |\mathcal{C}|)T^2)$. We observe in practice that, in all our experiments, a few iterations were sufficient to reach convergence (characterised by DTW paths not evolving anymore).

Experiments

The use of MAD and $|\mathcal{C}|$ -MAD as losses for neural domain adaptation is now assessed considering a real remote sensing dataset, for which there exists a known global temporal shift between the (classes of the) domains due to different weather conditions. We further study its use in a motion capture context.³

Backbone architecture. In order to evaluate the impact of the MAD regularization loss for domain adaptation, we use the exact same deep neural network architecture as in CoDATS [21]. The feature extractor g_Ω is composed of a stack of 3 convolutional layers followed by batch normalization. The classification head f_θ then consists of a global average pooling followed by a single fully-connected layer, as shown in Fig. 2. The first layer of the feature extractor has 128 filters of size 8, the second layer has 256 filters of size 5 and the third layer has 128 filters of size 3. Note that, for CoDATS, the domain adversarial classification head is plugged after the pooling operator.

³ One can find code and datasets at <https://github.com/rtavenar/MatchAndDeform>.

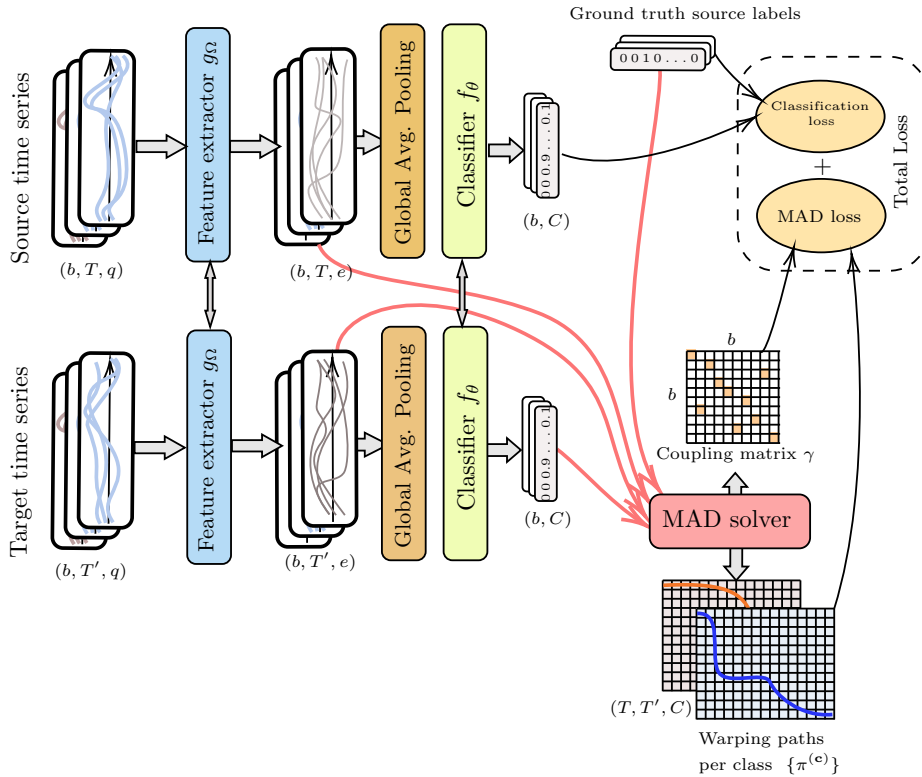


Fig. 2. Backbone architecture and schematic view of the loss computation for $|\mathcal{C}|$ -MAD.

Hyper-parameters. For a fair comparison, and for all the experiments, we choose a single learning rate (0.0001) and batch size ($b = 256$) for both the CoDATS baseline and our approach. The MAD loss in Eq. (8) introduces two extra hyper-parameters; their values are set to $\alpha = 0.01$ and $\beta = 0.01$ for all the experiments. In practice, we observe that a large range of values give similar performances (a deeper investigation is provided as Supplementary Material).

Experimental setup. Our proposed method MAD and its variant $|\mathcal{C}|$ -MAD are compared to CoDATS, which extends the domain adversarial neural network (DANN, [7]) framework to time series data and has been shown to reach state-of-the-art performance. As in [21], the proportion of each class in the target domain is assumed to be known during the learning phase. This corresponds to the CoDATS-WS variant. As an OT for DA baseline, we furthermore compare to DeepJDOT-DTW, considering DeepJDOT [4] adapted for time series: individual DTW between pairs of series are considered as transport cost matrix $C(\mathbf{X}, \mathbf{X}') = \{\text{DTW}(\mathbf{x}, \mathbf{x}')\}_{\mathbf{x}, \mathbf{x}' \in \mathbf{X}, \mathbf{X}'}$.

In MAD and DeepJDOT-DTW, the weights \mathbf{w} from the source mini-batches are set such that they reflect the proportion of the classes in the target domain;




Fig. 3. Per-class median-filtered average Normal Difference Vegetation Index (NDVI) profiles, showing the growth of 4 crop types (miniTimeMatch dataset, 1 domain pair (DK1 \rightarrow FR1)). Crops develop similarly in the 2 considered regions, but the patterns are temporally shifted, with a class-specific shift.

uniform weights \mathbf{w}' are used for target mini-batches. Note that integrating this extra information could be avoided by using unbalanced optimal transport [6], which is left for future work. For the sake of fairness, all competing methods are trained and tested on the same data splits with the same backbone architecture.

Classification accuracy from the same model trained on source data to predict the target labels (hereafter denoted *No adaptation*) is reported for each domain pair. This gives a proxy of the overall difficulty of the adaptation problem for this pair. The *Target only* baseline is trained and evaluated on the target domain, and can be seen as an upper bound estimation of the classification accuracy. For each domain, the data is split into train/validation/test sets with 64% / 16% / 20% of the data in each split, respectively. Note that the validation set is used only for CoDATS-WS for early stopping purpose. For each experiment, the average over all pairs of domains is provided together with the averaged standard deviations.

Remote Sensing data. We first focus on an application field for which the temporal shift between domains has already been documented. In Earth observation, and especially for land cover mapping applications, the differences in weather, soil conditions or farmer practices between study sites are known to induce temporal shifts in the data that should be dealt with, as discussed in [13] for example. Moreover, these temporal shifts can be global to the study sites or class specific (see Fig. 3), which is the setup for which MAD and $|\mathcal{C}|$ -MAD are designed.

TimeMatch dataset [13] is a crop-type mapping dataset covering four different geographical areas in Austria, Denmark, and France. It contains time series of multi-spectral measurements from satellite imagery at each geo-location. Labels are available at a parcel level: hence we consider a modified version of TimeMatch, referred to as miniTimeMatch in the following, in which measurements are averaged over agricultural parcels. The aim is to recognize parcels' crop type, such as corn or wheat. Domains correspond to different geographical areas, with their own characteristics. Moreover, we remove data from classes that are not observed in all domains, or with less than 200 agricultural fields in at least one domain. We also remove the "unknown" class that gathers crops from many types. Finally, we sample n_{\min} series per class in each domain, where n_{\min} is the minimum frequency of the class across all domains. This results in a

dataset of 28,858 time series and 8 classes per domain, each being described by 10 features per timestamp. The same domain pairs as in [13] are considered for the evaluation.

Observations from the miniTimeMatch dataset may be noisy, with some measurements corresponding to cloudy observations. Moreover, the time series of the different domains have different temporal sampling due to the filtering of highly cloudy images (see Figure 3). Then, we introduce a new domain adaptation problem, considering another land crop mapping dataset called TarnBrittany, in which observations from cloudy images are removed. It is built from Sentinel-2⁴ time series for the year 2018, and is composed of two domains: (Brittany) in the mid-west of France on tile T30UVU and (Tarn) in the southwest of France on tile T31TDJ. We collect 72 images for Tarn and 58 images for the cloudier Brittany. Images have been orthorectified and corrected from atmospheric effects by using the MAJA processing chain[11] to yield Level-2A surface reflectance images. As for miniTimeMatch, we average the reflectance values over the agricultural fields to obtain multivariate time series of 10 spectral bands. Labels are obtained from the French land parcel identification system (IGN RPG⁵) for the year 2018. We keep only the five main crops present in both areas. 5,743 series are sampled in each area with similar class proportions in source and target domains. The two domain pairs are considered for the evaluation.

Motion Capture (MoCap) data. We further evaluate MAD and $|\mathcal{C}|$ -MAD on a MoCap setup: Human Activity Recognition (HAR) dataset [5], in which no temporal shift is documented. As such, we do not expect MAD and $|\mathcal{C}|$ -MAD to behave better than state-of-the-art algorithms. The dataset contains time series (128 timestamps, 9 features, 6 classes) of movements recorded thanks to mobile sensors such as accelerometer or gyroscope; the aim is to recognize human activities such as walking or sitting. Domains correspond to different persons with their own characteristics. The same 10 domain pairs as in [21] are considered for the evaluation. Details on the dataset are provided as Supplementary Material.

Results Tab. 1 shows that, on HAR dataset, MAD and $|\mathcal{C}|$ -MAD reach similar performance to CoDATS-WS. In more details, MAD and $|\mathcal{C}|$ -MAD outperform CoDATS-WS in five out of ten situations, achieving a slightly better average performance. Tab. 2 provides classification accuracy for remote sensing datasets, on which MAD and $|\mathcal{C}|$ -MAD are expected to work better thanks to their ability to handle (per class) time shifts. The discrepancies between performances of models trained only on source data (*No adaptation*) and those only trained on target data (*Target only*) also point out that all domain pairs constitute a more challenging adaptation problem compared to HAR dataset. One can notice that MAD and $|\mathcal{C}|$ -MAD outperform CoDATS-WS in 5 out of the 7 problems, sometimes with an important improvement (see DK1 \rightarrow FR1 for example), even in the presence of noise (that is present in the miniTimeMatch dataset). This illustrates the fact that MAD and $|\mathcal{C}|$ -MAD are of prime interest when global or class-specific temporal deformations occur between domains (see Fig. 3). For each

⁴ <https://theia-ide.org/>

⁵ <https://geoservices.ign.fr/rpg>

Problem	No adapt.	CoDATS-WS	DeepJDOT-DTW	MAD	$ \mathcal{C} $ -MAD	Target only
2 \rightarrow 11	83.3 \pm 0.7	81.3 \pm 4.4	99.5 \pm 0.7	98.4 \pm 1.3	99.0 \pm 1.5	100 \pm 0.0
7 \rightarrow 13	89.9 \pm 3.6	94.4 \pm 2.6	88.4 \pm 0.7	100.0 \pm 0.0	100 \pm 0.0	100 \pm 0.0
12 \rightarrow 16	41.9 \pm 0.0	64.0 \pm 0.6	56.8 \pm 3.8	64.0 \pm 0.6	64.0 \pm 0.6	100 \pm 0.0
12 \rightarrow 18	90.0 \pm 1.7	100 \pm 0.0	94.5 \pm 0.0	99.5 \pm 0.6	99.5 \pm 0.7	100 \pm 0.0
9 \rightarrow 18	31.1 \pm 1.7	78.1 \pm 4.9	51.6 \pm 10.0	71.7 \pm 0.6	71.2 \pm 1.1	100 \pm 0.0
14 \rightarrow 19	62.0 \pm 4.3	99.5 \pm 0.7	64.4 \pm 1.3	83.3 \pm 2.3	84.3 \pm 2.4	100 \pm 0.0
18 \rightarrow 23	89.3 \pm 5.0	89.8 \pm 0.6	97.3 \pm 0.0	98.2 \pm 0.6	97.8 \pm 0.6	100 \pm 0.0
6 \rightarrow 23	52.9 \pm 2.3	95.1 \pm 3.5	60.9 \pm 2.7	97.8 \pm 0.6	97.8 \pm 0.6	100 \pm 0.0
7 \rightarrow 24	94.4 \pm 2.7	99.6 \pm 0.6	98.3 \pm 1.2	100.0 \pm 0.0	100 \pm 0.0	100 \pm 0.0
17 \rightarrow 25	57.3 \pm 5.5	95.5 \pm 4.6	73.6 \pm 3.2	96.3 \pm 2.0	95.5 \pm 1.2	100 \pm 0.0
Average	69.2 \pm 2.8	89.7 \pm 2.3	78.5 \pm 2.4	90.9 \pm 0.9	90.9 \pm 0.9	100 \pm 0.0

Table 1. HAR dataset: classification performance (% accuracy, avg over 3 runs \pm std).

Problem	No adapt.	CoDATS-WS	DeepJDOT-DTW	MAD	$ \mathcal{C} $ -MAD	Target only
Tarn \rightarrow Brittany	88.5 \pm 3.7	96.0 \pm 1.6	89.4 \pm 2.2	98.8 \pm 0.3	98.9 \pm 0.4	99.7 \pm 0.1
Brittany \rightarrow Tarn	48.9 \pm 0.6	93.6 \pm 0.1	47.6 \pm 1.4	92.0 \pm 0.2	90.6 \pm 1.1	98.6 \pm 0.2
Average	68.7 \pm 1.7	94.8 \pm 0.9	68.5 \pm 1.8	95.4 \pm 0.2	94.7 \pm 0.8	99.2 \pm 0.1
DK1 \rightarrow FR1	69.2 \pm 1.3	74.8 \pm 1.5	79.0 \pm 1.6	88.4 \pm 0.4	88.3 \pm 0.9	95.8 \pm 0.9
DK1 \rightarrow FR2	62.2 \pm 3.5	87.0 \pm 3.4	76.6 \pm 2.5	82.5 \pm 1.1	81.0 \pm 1.1	94.2 \pm 1.7
DK1 \rightarrow AT1	73.9 \pm 0.2	71.6 \pm 15.4	78.6 \pm 0.6	93.1 \pm 1.2	92.3 \pm 2.2	96.7 \pm 0.7
FR1 \rightarrow DK1	61.9 \pm 5.2	78.0 \pm 10.7	71.3 \pm 2.9	88.2 \pm 0.3	88.2 \pm 0.5	96.2 \pm 0.3
FR1 \rightarrow FR2	78.8 \pm 0.9	82.1 \pm 8.2	77.1 \pm 1.3	90.5 \pm 0.2	89.6 \pm 0.4	94.2 \pm 1.7
Average	69.2 \pm 2.2	78.7 \pm 7.8	76.5 \pm 1.8	88.5 \pm 0.6	87.9 \pm 1.0	95.4 \pm 1.1

Table 2. Remote sensing datasets: classification performance (% accuracy, avg over 3 runs \pm std).

set of datasets, a one-sided Wilcoxon signed-rank test is performed to assert if the results of MAD and $|\mathcal{C}|$ -MAD are statistically greater than those of CoDATS-WS and deepJDOT-DTW. Significant results ($\alpha = 0.05$) are obtained for MAD and $|\mathcal{C}|$ -MAD over CoDATS-WS ($p = 0.005$ and $p = 0.01$, respectively), and over DeepJDOT-DTW ($p = 3.10^{-6}$ and $p = 6.10^{-6}$, respectively), confirming the superiority of our approach. There is no significant difference ($p = 0.18$) between CoDATS-WS and deepJDOT-DTW performances.

$|\mathcal{C}|$ -MAD as a trade-off between MAD and DeepJDOT-DTW. $|\mathcal{C}|$ -MAD can be seen as an intermediate configuration between MAD (in which a single global DTW alignment is performed) and DeepJDOT-DTW (which allows individual alignments between each pair of series). The good performance of $|\mathcal{C}|$ -MAD compared to DeepJDOT-DTW show that it successfully manages the specific intra-class global alignments. By computing pair-to-pair DTW alignments, DeepJDOT-DTW allows for too many degrees of freedom that lead to spurious matchings. On the contrary, ($|\mathcal{C}|$ -)MAD acts as a regularizer (see Property 1), allowing to meaningfully constrain the alignments for better performance.

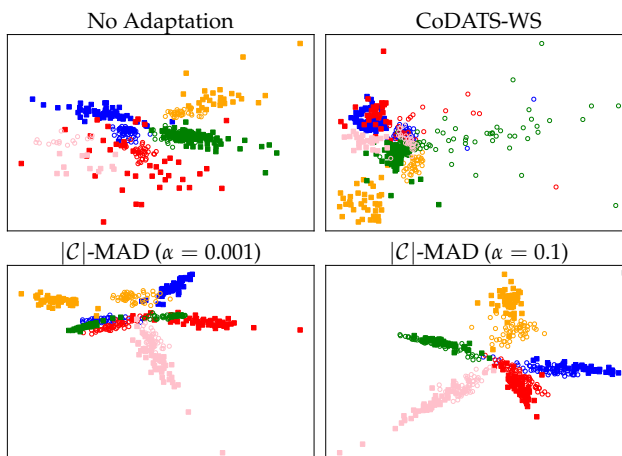


Fig. 4. Multi-Dimensional Scaling (MDS) representation of the latent space for several models on the Tarn (empty circles) \rightarrow Brittany (squares) adaptation problem: each colour represents a different class.

Latent space visualization. To further investigate the internal properties of the models, we use Multi-Dimensional Scaling (MDS) to visualize the latent space of the compared models. To do so, we focus on the models learned for the “Tarn \rightarrow Brittany” adaptation problem and visualize $g_{\Omega}(\mathbf{X})$ and $g_{\Omega}(\mathbf{X}')$ jointly for a model trained on source data only, as well as for CoDATS-WS and $|\mathcal{C}|$ -MAD that aim at learning a domain-invariant representation. While Euclidean distance is used to feed MDS for both baselines, we use the distances resulting from the DTW paths for $|\mathcal{C}|$ -MAD (as stored in \mathbf{C}_{OT}), in order to account for the temporal realignment on which our method relies. The resulting visualizations presented in Fig. 4 show a significant shift in distributions from target to source domain, when no DA strategy is employed. CoDATS-WS also fails to align domains in the presence of time shifts, whereas $|\mathcal{C}|$ -MAD successfully clusters series by class.

Related work

Several approaches proposed in the literature share some common ground with MAD, either because they aim at aligning complex objects to perform domain adaptation tasks or because they deal with the problem of aligning (sets of) time series.

Using OT for domain adaptation has been initially proposed in OTDA [2], which aims at aligning the source and target domains using OT. It relies on the hypothesis that the distribution shift between the two domains is an affine transformation and uses the transport plan to evaluate its coefficients. Classification is then performed on the realigned samples. DeepJDOT [4], on which the neural

domain adaptation with a MAD loss introduced in this paper relies, can be seen as an improvement over OTDA. CO-OT [17] jointly optimizes two OT plans γ_s and γ_f between samples (rows) and features (columns) of datasets \mathbf{X} and \mathbf{X}' , resulting in a loss of the form $\langle \mathbf{L}(\mathbf{X}, \mathbf{X}') \otimes \gamma_s, \gamma_f \rangle$. It has been benchmarked in a heterogeneous domain adaptation context but does not allow enforcing temporal coherence that should be met when time series are at stake. To the best of our knowledge, no OT-based approach has yet tackled the challenge of unsupervised DA for time series classification.

Combining OT and (soft-)DTW has also been proposed in [9] where the goal is to align brain imaging signals using soft-DTW with an OT-based ground cost. In this setting however, optimization is not performed jointly on γ and π but, rather, individual transport problems are solved independently for each pair of timestamps, hence not enforcing global spatial consistency. [10] propose an extension with the aim to average spatio-temporal signals. However, none of these works consider the DA scenario.

In the same line of work, [1] use a Gromov-Wasserstein-like similarity measure to compare time series that may not lie in the same ambient space (*e.g.* that do not share the same features). They solve a problem of the form $\langle \mathbf{L}(\mathbf{X}, \mathbf{X}') \otimes \pi, \pi \rangle$ (in which π is a coupling matrix that meets the DTW alignment constraints), coined Gromov-dynamic time warping. Optimization also relies on a BCD, however its convergence is not guaranteed. Moreover, it does not allow taking into account the temporal distortion that occurs between *sets* of time series.

Conclusion and perspectives

In this paper, we introduce Match-And-Deform (MAD) that combines optimal transport and dynamic time warping to match time series across domains given global temporal transformations. We embed MAD as a regularization loss in a neural domain adaptation setting and evaluate its performance in different settings: on MoCap datasets, MAD gives similar performance to state-of-the-art DA time series classification methods; when a temporal deformation is at stake (such as in a remote sensing scenario), we show that MAD reaches better performance, thanks to its ability to capture temporal shifts.

The $|\mathcal{C}|$ -MAD variant presented in this paper relies on source domain classes to form groups of temporally coherent series, yet an unsupervised strategy to form groups based on time series content alone is a worthy track for future works. Moreover, inspired by [19], time series weights could be learned instead of set *a priori*, in order to decrease the sensitivity to outliers. Finally, the use of MAD in other tasks, such as cross-domain missing data imputation [12], can also be considered.

Acknowledgements. François Painblanc and Romain Tavenard are partially funded through project MATS ANR-18-CE23-0006. Nicolas Courty is partially funded through project OTTOPIA ANR-20-CHIA-0030. Laetitia Chapel is partially funded through project MULTISCALE ANR-18-CE23-0022.

References

1. Cohen, S., Luise, G., Terenin, A., Amos, B., Deisenroth, M.: Aligning time series on incomparable spaces. In: International Conference on Artificial Intelligence and Statistics. pp. 1036–1044 (2021)
2. Courty, N., Flamary, R., Tuia, D., Rakotomamonjy, A.: Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**, 1853–1865 (2016)
3. Courty, N., Flamary, R., Habrard, A., Rakotomamonjy, A.: Joint distribution optimal transportation for domain adaptation. *Advances in Neural Information Processing Systems* **30** (2017)
4. Damodaran, B.B., Kellenberger, B., Flamary, R., Tuia, D., Courty, N.: DeepJDOT: Deep joint distribution optimal transport for unsupervised domain adaptation. In: European Conference on Computer Vision. pp. 447–463 (2018)
5. Dua, D., Graff, C.: UCI machine learning repository (2017), <http://archive.ics.uci.edu/ml>
6. Fatras, K., Séjourné, T., Flamary, R., Courty, N.: Unbalanced minibatch optimal transport; applications to domain adaptation. In: International Conference on Machine Learning. pp. 3186–3197 (2021)
7. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* **17**(1), 2096–2030 (2016)
8. Genevay, A., Peyré, G., Cuturi, M.: Learning generative models with Sinkhorn divergences. In: International Conference on Artificial Intelligence and Statistics. pp. 1608–1617 (2018)
9. Janati, H., Cuturi, M., Gramfort, A.: Spatio-temporal alignments: Optimal transport through space and time. In: International Conference on Artificial Intelligence and Statistics. pp. 1695–1704 (2020)
10. Janati, H., Cuturi, M., Gramfort, A.: Averaging spatio-temporal signals using optimal transport and soft alignments. *arXiv:2203.05813* (2022)
11. Lonjou, V., Desjardins, C., Hagolle, O., Petrucci, B., Tremas, T., Dejus, M., Makarau, A., Auer, S.: Maccs-atcor joint algorithm (maja). In: Remote Sensing of Clouds and the Atmosphere XXI. vol. 10001, p. 1000107 (2016)
12. Muzellec, B., Josse, J., Boyer, C., Cuturi, M.: Missing data imputation using optimal transport. In: International Conference on Machine Learning. pp. 7130–7140. PMLR (2020)
13. Nyborg, J., Pelletier, C., Lefèvre, S., Assent, I.: Timematch: Unsupervised cross-region adaptation by temporal shift estimation. *ISPRS Journal of Photogrammetry and Remote Sensing* **188**, 301–313 (2022)
14. Peyré, G., Cuturi, M.: Computational optimal transport. *Foundations and Trends® in Machine Learning* **11**(5-6), 355–607 (2019)
15. Purushotham, S., Carvalho, W., Nilanon, T., Liu, Y.: Variational recurrent adversarial deep domain adaptation. In: International Conference on Learning Representations (2017)
16. Redko, I., Habrard, A., Morvant, E., Sebban, M., Bannani, Y.: *Advances in Domain Adaptation Theory*. ISTE Press (2019)
17. Redko, I., Vayer, T., Flamary, R., Courty, N.: Co-optimal transport. In: *Advances in Neural Information Processing Systems*. NeurIPS (2020)
18. Sakoe, H., Chiba, S.: Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing* **26**(1), 43–49 (1978)

19. Vincent-Cuaz, C., Flamary, R., Corneli, M., Vayer, T., Courty, N.: Semi-relaxed gromov-wasserstein divergence and applications on graphs. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=RShaMexjc-x>
20. Wilson, G., Cook, D.: A survey of unsupervised deep domain adaptation. *ACM Journal* (2020)
21. Wilson, G., Doppa, J.R., Cook, D.J.: Multi-source deep domain adaptation with weak supervision for time-series sensor data. In: ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 1768–1778 (2020)