



Harmonizing flows: Leveraging normalizing flows for unsupervised and source-free MRI harmonization

Farzad Beizaei^{a,b,c,*}, Gregory A. Lodygensky^{c,d}, Chris L. Adamson^e, Deanne K. Thompson^{e,f,g},
Jeanie L.Y. Cheong^{e,g,h,i}, Alicia J. Spittle^{e,h,j}, Peter J. Anderson^{e,f}, Christian Desrosiers^{a,b},
Jose Dolz^{a,b}

^a LIVIA, ÉTS, Montreal, Quebec, Canada

^b ILLS, McGill - ETS - Mila - CNRS - Université Paris-Saclay - CentraleSupélec, Canada

^c CHU Sainte-Justine, University of Montreal, Montreal, Canada

^d Canadian Neonatal Brain Platform, Montreal, Canada

^e Murdoch Children's Research Institute, Parkville, Victoria, Australia

^f School of Psychological Sciences, Monash University, Clayton, Victoria, Australia

^g Department of Paediatrics, The University of Melbourne, Victoria, Australia

^h The Royal Women's Hospital, Melbourne, Parkville, Victoria, Australia

ⁱ Department of Obstetrics and Gynaecology, The University of Melbourne, Victoria, Australia

^j Department of Physiotherapy, The University of Melbourne, Victoria, Australia

ARTICLE INFO

Keywords:

MRI harmonization
Normalizing flows
Test-time adaptation
Brain MRI

ABSTRACT

Lack of standardization and various intrinsic parameters for magnetic resonance (MR) image acquisition results in heterogeneous images across different sites and devices, which adversely affects the generalization of deep neural networks. To alleviate this issue, this work proposes a novel unsupervised harmonization framework that leverages normalizing flows to align MR images, thereby emulating the distribution of a source domain. The proposed strategy comprises three key steps. Initially, a normalizing flow network is trained to capture the distribution characteristics of the source domain. Then, we train a shallow harmonizer network to reconstruct images from the source domain via their augmented counterparts. Finally, during inference, the harmonizer network is updated to ensure that the output images conform to the learned source domain distribution, as modeled by the normalizing flow network. Our approach, which is unsupervised, source-free, and task-agnostic is assessed in the context of both adults and neonatal cross-domain brain MRI segmentation, as well as neonatal brain age estimation, demonstrating its generalizability across tasks and population demographics. The results underscore its superior performance compared to existing methodologies. The code is available at <https://github.com/farzad-bz/Harmonizing-Flows>.

1. Introduction

Magnetic Resonance Imaging (MRI) serves as an indispensable tool in modern medical diagnostics, enabling clinicians to obtain detailed insights into anatomical structures and pathological conditions. However, the inherent variability in MRI data acquisition protocols across different imaging sites poses significant challenges in achieving consistent and reliable image interpretation. This variability can stem from differences in scanner hardware, imaging parameters, and patient populations (Takao et al., 2011), leading to inconsistencies in image appearance and potentially confounding downstream analysis. For instance, MRIs acquired from two different scanners or with different sets

of parameters and configurations will often have noticeable appearance differences, which can be considered as domain shift. Therefore, pooling multi-centric clinic trials to address specific questions does not necessarily enhance statistical power, as the introduced variance may partially stem from non-clinical sources.

On the other hand, despite the considerable progress observed in deep learning, these models still face challenges in coping with distributional shifts. The performance of deep neural networks in fundamental visual problems, such as classification, segmentation, and regression, largely degrades when they are applied to data acquired under varied conditions, consequently limiting their broad applicability. Specifically, models trained on data from a specific site often struggle to achieve

* Corresponding author.

E-mail address: farzad.beizaei.1@ens.etsmtl.ca (F. Beizaei).

<https://doi.org/10.1016/j.media.2025.103483>

Received 11 July 2024; Received in revised form 21 January 2025; Accepted 25 January 2025

Available online 2 February 2025

1361-8415/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

similar performance when applied to images from other centers.

To mitigate this challenge, image harmonization tackles the problem of distributional drifts by mapping images from one domain to another, aiming at transferring contrast characteristics across diverse datasets. MRI harmonization ensures the comparability of MRI data collected from different scanners, facilitating accurate and consistent analysis for multi-center studies involving diverse imaging datasets. However, many existing harmonization approaches rely on assumptions that could impede their feasibility and scalability in real-world applications. For example, some methods involve acquiring imaging data of the same anatomical targets from multiple sites or locations. These methods, often referred to as using *traveling subjects*, aim to identify and quantify the transformations needed to harmonize the data across different acquisitions settings (Dewey et al., 2019; Durrer et al., 2023). Another family of approaches requires access to the source images during harmonization or knowing the target domain in advance (Pomponio et al., 2020; Modanwal et al., 2020; Liu et al., 2021; Cackowski et al., 2023), which might not be feasible in practical scenarios. It is worth mentioning that in this context, the source domain refers to the domain that serves as the reference for harmonization, representing the desired appearance or distribution that other domains (referred to as target domains) are aligned to after the harmonization process. Furthermore, several of these harmonization strategies require annotated data associated with the downstream task (Delisle et al., 2021; Dinsdale et al., 2021; Karani et al., 2021). This poses an additional challenge to the harmonization, as acquiring labeled data can be resource-intensive and time-consuming, especially when dealing with large datasets and dense tasks such as segmentation. Finally, many harmonization techniques require knowledge of the target domains during the training phase, despite the common occurrence of unknown target domains in real-world scenarios. Based on the limitations exposed above, we present the following contributions in this work:

- We alleviate the aforementioned constraints on MR harmonization and introduce a novel harmonization approach that is *unsupervised*, *source-free* (*SF*), *task-agnostic* (*TA*) and can cope with *unknown-domains* (*UD*) without necessitating retraining for every target distribution. In fact, our approach only requires MRIs from one modality of the source domain during training, as opposed to existing approaches.
- Specifically, we propose to leverage a modern family of generative models, known as normalizing flows, which have proven to be highly effective in modeling data distributions for generative purposes.
- Alongside the methodological novelty of the proposed method, our empirical findings illustrate that it yields significant improvements over existing harmonization techniques, while effectively mitigating their limitations. More importantly, the comprehensive experimental section on multiple tasks and datasets demonstrates that the proposed approach successfully generalizes across target tasks and population demographics.

A preliminary conference version of this work has been presented at IPMI 2023 (Beizae et al., 2023). This manuscript provides a substantial extension of the conference version, which includes (i) an extended literature review on methods addressing the problem of distribution shift, and a comprehensive empirical validation of the proposed approach, including (ii) additional recent approaches for harmonization, (iii) extensive ablation studies to validate our choices, (iv) assessing the performance of our approach in multiple tasks and population demographics, (v) including additional adult datasets in the experiments, (vi) employing additional evaluation metrics to assess the performance from a harmonization standpoint, (vii) and complimentary plots and results to better understand the overall performance of the different studied methods.

2. Related work

Image harmonization. In the medical domain, various methods have been proposed to harmonize images, with a particular focus on MRI data. Traditional post-processing procedures like intensity histogram matching (Nyúl et al., 2000; Shinohara et al., 2014) help mitigate biases across scanners, but may also eliminate informative local intensity variations. Statistical harmonization approaches, on the other hand, can model both image intensity and dataset bias at the voxel level (Fortin et al., 2016, 2017; Beer et al., 2020). However, when the variations in data distribution are more complex and localized, they typically lead to sub-optimal harmonization outcomes. Additionally, these methods must often be adjusted each time images from new sites are provided, further compromising their performance in real-world applications. Modern strategies for image harmonization, using deep learning techniques, hold significant promise as an alternative solution for this problem (Dewey et al., 2019; Zhu et al., 2017; Liu et al., 2021; Zuo et al., 2021; Delisle et al., 2021; Dinsdale et al., 2021). Yet, these approaches often rely on unrealistic assumptions, which pose significant barriers to scalability when applied to extensive multi-site harmonization endeavors. First, some methods require images of the same target anatomy across different sites, known as *traveling subjects*, to identify intensity transformations among different sites (Dewey et al., 2019; Durrer et al., 2023). This means that a given number of subjects being scanned at every site or scanner is employed for training, a condition rarely met in practice. The most widely used models for MR harmonization are GANs and autoencoders, which have shown promising results in reducing multi-site variation through image-to-image synthesis. GANs perform domain translation by learning domain invariant features. One issue of such models is that they are limited to mapping between two specific scanners for most studies. Also, target domains are required to be known at the training time (Zhu et al., 2017; Liu et al., 2021) which is a limiting factor for the scalability of the harmonization process. Additionally, each time a new domain is added, these approaches must be fine-tuned to accommodate the characteristics of this domain. Autoencoder-based methods (Torbaty et al., 2021; Dewey et al., 2020; Zuo et al., 2021, 2023; Wu et al., 2023; Cackowski et al., 2023), on the other hand, aim to harmonize data in terms of disentangled representations. This group of harmonization methods attempted to extract scanner-related features for harmonization. Similar to GANs, data from multiple sources and target domains are required for training. CALAMITI (Zuo et al., 2021), and HACA3 (Zuo et al., 2023), which are two key methods of this category, require paired multi-modal MRIs for training, thereby restricting their applicability, particularly in single-modality scenarios. DLEST (Wu et al., 2023) and Imunity (Cackowski et al., 2023) try to solve the problem of requiring multi-modal data for latent disentanglement. However, training these methods can be challenging due to the instability of their adversarial learning strategy. Lastly, task-dependent methods leverage labels associated with each image for a particular downstream task to optimize the harmonization for this specific problem (Delisle et al., 2021; Dinsdale et al., 2021). These methods often rely on task-specific features or assumptions, making them less effective when applied to new tasks or unseen data. Additionally, task-dependent harmonization approaches require large amounts of annotated data for training, which can be costly and time-consuming to acquire.

Test-time Domain Adaptation. Traditional solutions to the problem of distributional shift use labeled samples from a source domain and unlabeled ones from the target domain for adapting a source-trained model to perform well on the target. Several strategies for this task, known as unsupervised domain adaptation (UDA), work by explicitly aligning the feature distributions of the source and target domains (Wang et al., 2023; Wu and Zhuang, 2020). Another popular approach consists in learning a domain-agnostic representation, for example using adversarial networks (Dou et al., 2019; Kannitsas et al., 2017). Generative

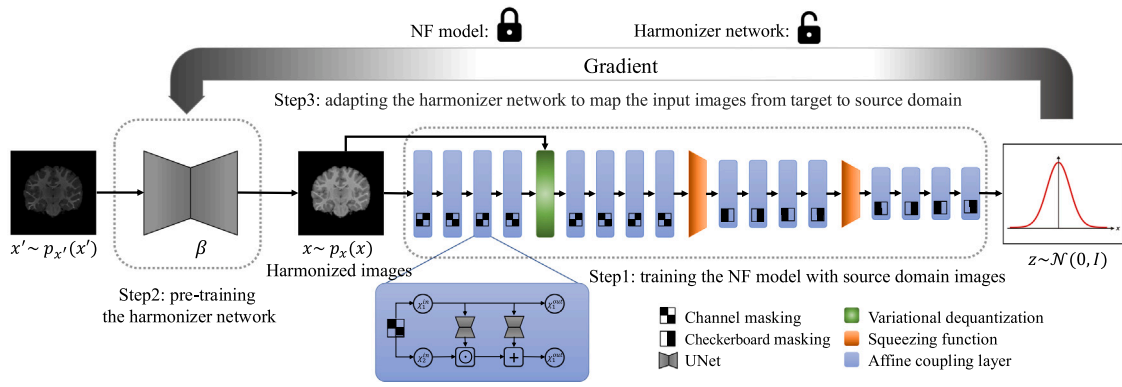


Fig. 1. Harmonizing flows pipeline. Our method comprises three primary steps. Initially, normalizing flow (NF) is utilized to capture the distribution of the source domain. Subsequently, a harmonizer network undergoes pre-training to reconstruct the original images from augmented counterparts, facilitating initial harmonization. In the third stage (test time), the trained NF is leveraged to update the parameters of the harmonizer network, ensuring maximal alignment between the harmonized outputs and the learned NF distribution. Notably, steps 1 and 2 are independent of each other and can be executed interchangeably.

adversarial networks (GANs) can also be employed in style transfer methods to change the appearance of images from the target to the source while also preserving their semantic structures (Chen et al., 2020; Zhao et al., 2021).

A major limitation of UDA methods is their need to have source examples during the adaptation phase, which may be impracticable in medical applications due to data sharing restrictions. Source-free domain adaptation (SFDA) approaches (Bateson et al., 2022; Yang et al., 2022; Stan and Rostami, 2024) relax this constraint and instead adapt the source-trained model using only unlabeled data from the target domain. However, these approaches are usually task-dependent and require an explicit adaptation process for each new target domain, hence are not compatible with the harmonization setting investigated in this work.

Another strategy for addressing distribution shift, more closely related to our method, is test-time adaptation (TTA) (Boudiaf et al., 2022; Mummadi et al., 2022; Liang et al., 2020; Wang et al., 2021a; Niu et al., 2022, 2023). Unlike SFDA, which performs adaptation in an offline step, this strategy adapts a pre-trained deep neural network to domain shifts encountered *during inference* on test samples. One of the earliest TTA approaches, called TENT (Wang et al., 2021a), updates the normalization layers of the network by minimizing the Shannon entropy of predictions for test samples. In Mummadi et al. (2022), entropy minimization has been changed by optimizing a log-likelihood ratio and considering the normalization statistics of the test batch. EATA (Niu et al., 2022) instead introduces an active sample selection criterion to identify reliable and non-redundant samples. The model is then updated based on these samples to minimize entropy loss for test-time adaptation. In Niu et al. (2023), authors propose a sharpness-aware and reliable entropy minimization method called SAR, further stabilizing the TTA process. Moreover, SHOT (Liang et al., 2020) adapts the entire feature extractor with a mutual information loss, while using pseudo-labels to provide additional test-time guidance. Instead of updating the network parameters, LAME (Boudiaf et al., 2022) uses Laplacian regularization to do a post-hoc adaptation of the softmax predictions. Recent works have explored the potential of TTA for cross-site/modality segmentation of medical images. Authors of Hu et al. (2021) propose a TTA method for segmentation using a regional nuclear-norm loss to improve the discriminability and diversity of predictions and a contour regularization term to enforce segmentation consistency between nearby pixels. Contrary to this paradigm, where typically the target-task network (e.g., classification or segmentation) is adapted at inference based on surrogate losses on the network predictions, our work focuses on modifying the image appearance instead, which offers a more general solution, which is agnostic to the task at hand.

Furthermore, while Karani et al. (2021) uses the reconstruction error of an auto-encoder applied on segmentation outputs to normalize input images, it requires segmentation masks for the adaptation, which makes of this strategy task and annotation dependent.

Normalizing flows. Popular methods for generative tasks include generative adversarial networks (GANs) (Goodfellow et al., 2020) and Variational Auto-encoders (Kingma and Welling, 2014). Despite their popularity and wide acceptance, these methods present several important limitations, including mode (Salimans et al., 2016) and posterior collapse (Lucas et al., 2019), training instability (Salimans et al., 2016), and the incapability of providing an exact evaluation of the probability density of new data points. Recently, normalizing flows (NF) have emerged as a popular approach for constructing probabilistic and generative models due to their ability to model complex distributions (Dinh et al., 2017). Normalizing flows involves mapping a complex distribution, often unknown or poorly characterized, to a simpler distribution, typically the standard normal distribution. This is accomplished through a series of invertible and differentiable transformations. These transformations allow for efficient density estimation, sampling, and generative modeling. One of the key advantages of normalizing flows is their ability to capture intricate dependencies within data while providing tractable likelihood estimation, enabling a wide range of applications across domains. While the majority of current literature has utilized NFs for generative purposes (e.g., image generation (Ho et al., 2019; Kingma and Dhariwal, 2018), noise modeling (Abdelhamed et al., 2019), graph modeling (Zang and Wang, 2020)) and anomaly detection (Gudovskiy et al., 2022; Kirichenko et al., 2020), recent findings also indicate their effectiveness in aligning a given set of source domains (Grover et al., 2020; Usman et al., 2020; Osowiecki et al., 2023). Closely related to our problem, and up to the best of our knowledge, only a few attempts have investigated using normalizing flows to harmonize MR images. In particular, Wang et al. (2021b) presented a strategy that harmonizes pre-extracted features, i.e., brain ROI volume measures, and not image harmonization as in this work. In addition, extracting these ROIs requires pixel-wise labels, making of this approach task-dependent, contrary to our method which is task-agnostic. Furthermore, Jeong et al. (2023) is a concurrent approach that appeared after the conference version of this work was published. As discussed in their work, and shown empirically in our evaluation, although this method also leverages normalizing flows to harmonize images, it fails in the presence of large domain drifts between source and target sites.

3. Methodology

We first define the problem addressed in this study. Consider $\mathcal{X}_s = \{\mathbf{x}_n\}_{n=1}^N$ as a collection of unlabeled images from the source domain

S , where each image i is represented by $\mathbf{x}_i \in \mathbb{R}^{|\Omega|}$, with Ω indicating its spatial domain (i.e., $W \times H$). Likewise, let $\mathcal{X}_{\mathcal{T}} = \{\mathbf{x}_n\}_{n=1}^M$ be the set of unlabeled images within a target domain \mathcal{T} .¹ The objective of unsupervised data harmonization is to discover a mapping function $f_{\theta} : \mathcal{T} \rightarrow S$ without relying on labeled images or paired data from either domain.

We introduce a solution based on normalizing flows to address this problem. The proposed framework, which comprises three separate steps, is illustrated in Fig. 1. In Step 1, we first utilize Normalizing Flows, renowned for their ability to accurately learn data likelihoods, to capture the source domain's distribution. In Step 2, we then employ an auto-encoder as a harmonizer network, pre-training it by reconstructing original images from the augmented source domain counterparts. Finally, during testing, we update the parameters of the harmonizer network using images from the unseen target domain, ensuring that the harmonized outputs align with the learned distribution using NF model. The following sections present each of these steps in greater detail.

3.1. Learning the source domain distribution

We have leveraged normalizing flows (Dinh et al., 2017) to model the source domain distribution. NFs are a modern family of generative models capable of modeling complex probability density $p_x(\mathbf{x})$ (i.e., the source domain distribution) through applying a sequence of transformation functions, denoted as $g_{\phi} = g_1 \circ g_2 \circ \dots \circ g_T$, on known simple probability density $p_u(\mathbf{u})$ such as standard normal distribution. Source image can be represented as $\mathbf{x} = g_{\phi}(\mathbf{u})$, where $\mathbf{u} \sim p_u(\mathbf{u})$ and $p_u(\mathbf{u})$ denotes the base distribution of the normalizing flow model. An essential condition for the transformation function g_{ϕ} is its requirement to be *invertible*, with both g_{ϕ} and g_{ϕ}^{-1} being *differentiable*. With these prerequisites met, the density of the original variable \mathbf{x} is well-defined, allowing for the exact computation of its likelihood using the change of variables rule, expressed as:

$$\begin{aligned} \log p_x(\mathbf{x}) &= \log p_z(g_{\phi}^{-1}(\mathbf{x})) + \log \left| \det(\mathbf{J}_{g_{\phi}^{-1}}(\mathbf{x})) \right| \\ &= \log p_z(g_{\phi}^{-1}(\mathbf{x})) + \sum_{t=1}^T \log \left| \det(\mathbf{J}_{g_t^{-1}}(\mathbf{u}_{t-1})) \right| \end{aligned} \quad (1)$$

The first component of the right side corresponds to the log-likelihood within the simple distribution and $\mathbf{J}_{g_t^{-1}}(\mathbf{u}_{t-1})$ indicates the Jacobian matrix corresponding to the transformation g_t . In order to train the Normalizing flow model and learn the source domain distribution, model parameters ϕ are learned by maximizing the likelihood of the transformed data under the simpler distribution. This is achieved by minimizing the negative log-likelihood in Eq. (1) which leads to the following loss function:

$$\mathcal{L}_{NF} = -\log p_x(\mathbf{x}) \quad (2)$$

Building the Normalizing Flow. Constructing a bijective transformation function neural network for the Normalizing Flow (NF) model often involves the stacking of affine coupling layers, as highlighted by Dinh et al. (2017) and Kingma and Dhariwal (2018). This approach has been established as an efficient strategy. Coupling layers offer computational symmetry, meaning they are equally rapid in both evaluation and inversion processes. This characteristic addresses usability concerns inherent in asymmetric flows, such as masked autoregressive flows, which makes the coupling layers a preferred choice. Their balanced computational efficiency enables smoother integration into various applications, contributing to their widespread adoption in NF architectures. Suppose $\mathbf{z} \in \mathbb{R}^D$ serves as the input to the coupling layer, which is partitioned disjointly into $(\mathbf{z}^A, \mathbf{z}^B) \in \mathbb{R}^d \times \mathbb{R}^{D-d}$. The

¹ For simplicity, we assume a single domain exists here. However, our formulation can be readily extended to accommodate T distinct domains.

partitioning can be done along spatial dimensions (e.g. checkerboard masking strategy) or channels (channel masking strategy). Then, the transformation function $g(\cdot) : \mathbb{R}^D \rightarrow \mathbb{R}^D$ can be expressed as:

$$\mathbf{y}^A = \mathbf{z}^A, \quad \mathbf{y}^B = \mathbf{z}^B \odot \exp(s(\mathbf{z}^A)) + t(\mathbf{z}^A) \quad (3)$$

where \mathbf{y}^A and \mathbf{y}^B represent the transformed parts of the input data and \odot is element-wise multiplication. This formulation divides the input into two parts (\mathbf{z}^A and \mathbf{z}^B), and transforms only the latter part, leaving the former unchanged. This setting offers simplicity for calculating the Jacobian determinant, which makes it possible to use complex neural networks as shift $s(\cdot)$ and scale $t(\cdot)$ networks. Note that the transformation in Eq. (3) is invertible and therefore allows for efficient recovery of \mathbf{z}^A and \mathbf{z}^B from \mathbf{y}^A and \mathbf{y}^B . The work in Dinh et al. (2017) presented coupling flows on simpler tasks and datasets which demanded less complex representations. However, our current task necessitates pixel-to-pixel mappings on more challenging data. Therefore, we replace the simple convolutional blocks in Dinh et al. (2017) with shallow U-shaped convolutional neural networks to find the scale and shift parameters of the affine transformation, as they capture broader contextual information and provide higher representational capacity. Moreover, given that NFs rely on the change of variables rule, which operates within continuous space, it is important to ensure that the input is continuous. Traditionally, dequantization involves adding uniform noise $u \in U[0, 1]$ to discrete values to convert them into continuous representations. However, it could lead to a hypercube representation with sharp borders. Such sharp borders pose a challenge for modeling with a flow, as it relies on smooth transformations. Recently, a variational framework was introduced (Ho et al., 2019) to expand dequantization to more sophisticated distributions. This was achieved by substituting the uniform distribution with a learnable distribution. This learnable distribution can be optimized alongside other parameters of the normalizing flow, allowing for end-to-end training and seamless integration into the density estimation process.

Constraining the source-distribution learning. Optimizing the objective in Eq. (2) solely with source images could potentially bias the model towards emphasizing characteristics of subjects, such as age and gender, rather than focusing on source-specific attributes like contrast and brightness. To address this concern, we propose a strategy aimed at facilitating the learning of the source domain distribution. For this purpose, in each iteration, we randomly select N' images of the source dataset \mathcal{X}_S and apply a series of augmentations $f_{aug}(\cdot)$ in such a way that the resulting image exhibits a dissimilarity in appearance compared to the original image, as measured by the mean squared distance, surpassing a predetermined threshold. These images can be served as out-of-distribution samples, to guide the normalizing flow model to learn source-specific characteristics. In particular, we employ different types of contrast augmentation, brightness changes, multiplicative transformations, and random monotonically increasing mapping functions to augment these images. Then, the overall learning objective of our model can be defined as follows.

$$\mathcal{L}_{\mathcal{T}} = \underbrace{-\sum_{n=1}^{N-N'} \log p_x(\mathbf{x}_n)}_{\text{Source distribution modeling}} - \underbrace{\sum_{n=1}^{N'} \min(c, -\log p_x(f_{aug}(\mathbf{x}_n)))}_{\text{Guiding term}}. \quad (4)$$

The first term is the learning objective in Eq. (2) over the source images, while the second one encourages the NF model to reduce the likelihood of the augmented images, which facilitates the learning of domain-specific characteristics rather than subject-related features. Furthermore, to prevent divergence of the negative log-likelihood for an augmented sample to infinity, we employ a constant margin in the second term denoted as c .

3.2. Achieving image harmonization

Harmonizer network. The goal of our harmonizer network $h_\theta(\cdot)$ is to perform image-to-image translation of MRIs from the target to the source domain in such a way that: $p_x(\mathbf{x}) = p_{x'}(h_\theta(\mathbf{x}'))$. It is important to highlight that the input and output of the harmonizer network must share the same spatial dimensions. Here θ denotes the set of learnable parameters within the harmonizer network, and \mathbf{x} and \mathbf{x}' represent the samples from the source domain and target domains, respectively. This implies that the harmonizer network aims to shift the distribution of the target images so that they align with the distribution of source images. Nonetheless, we want the proposed method to operate effectively on unseen domains, necessitating that the target domains remain unknown during training. To this end, at first, the harmonizer network was trained to restore the original MRIs of the source domain from its augmented version. As in the previous step, we used different types of contrast augmentation, brightness changes, multiplicative transformations, and random monotonically increasing mapping functions. In contrast to the first step, there are no restrictions on how much the image can be altered as long as the augmented image is logical and details are not eliminated. To train this model, we employed the sum of two commonly used standard reconstruction loss functions: SSIM (Structural Similarity Index Measure) loss (Wang et al., 2004) and L1 loss (mean absolute error). SSIM loss is more appropriate when preserving structural similarity and perceptual quality is important, while L1 loss is suitable for tasks where exact pixel-wise accuracy is required, regardless of perceptual differences. The learning objective for the harmonizer network thus becomes:

$$\theta^{init} = \underset{\theta}{\operatorname{argmin}} \frac{1}{N} \sum_{n=1}^N \left(\left\| \mathbf{x}_n - h_\theta(f_{aug}(\mathbf{x}_n)) \right\| + SSIM(\mathbf{x}_n, h_\theta(f_{aug}(\mathbf{x}_n))) \right) \quad (5)$$

A simple UNet has been considered as the harmonizer network. Also, it is important to emphasize that the conducted augmentations may not perfectly represent potential unseen target domains. So, directly using the trained parameters θ^{init} for image-to-image translation yields sub-optimal harmonization. Nevertheless, it provides a good starting point for the next phase.

Adapting the harmonizer network leveraging the Normalizing Flow. So far, we have obtained an initial harmonizer network that gives us the possibility of transforming image appearance across domains, while bearing in mind that the output is sub-optimal. Also, we have learned the exact distribution of the source domain using a normalizing flow model, where we ensured that it focuses on domain-specific characteristics. The final step involves refining the harmonizer network to effectively map images from the target domain onto the distribution of the source domain. For this purpose, firstly, we stack the trained NF model at the top of the pre-trained harmonizer network as Fig. 1. Note that as we aim to leverage the learned distribution by the NF model which is already trained with the source data, its parameters should remain frozen during the adaptation of the harmonizer. Then we try to optimize the parameters of the harmonizer network, such that harmonizer network outputs exhibit a high likelihood of aligning with the source domain distribution under the supervision of the trained NF model. The learning objective of the adaptation stage is to increase the likelihood of the harmonizer outputs for images from the target domain, based on the density estimation provided by the NF model. This objective is encapsulated in a defined loss function, which can be expressed as follows:

$$\mathcal{L}_{Adap} = - \sum_{m=1}^M \log p_x(g_\phi(h_\theta(\mathbf{x}_m))) \quad (6)$$

As a stopping condition for updating the harmonizer, we assess two potential alternatives: One criterion involves assessing the Shannon entropy of the predictions generated by the target downstream task using the harmonized images (e.g., classification or segmentation), and stopping the adaptation when the entropy plateaus. Additionally, we take into account the bits per dimension (BPD), which is a scaled variation of the *negative log-likelihood* commonly utilized for assessing generative models:

$$BPD = - \log p_x(\mathbf{x}) \cdot \left(\log 2 \cdot \prod_i \Omega_i \right)^{-1} \quad (7)$$

where $\Omega_1, \dots, \Omega_T$, are the spatial dimensions of the input images. More concretely, we can stop updating the harmonizer network when the reached BPD value matches the BPD observed for the source images using the trained NF model. In practice, the source BPD value can be determined during the training time using a validation set.

4. Experiments

4.1. Experimental setting

First, we resort to the cross-site brain MRI segmentation task to evaluate the harmonization performance of different methods. We chose this task as it allows us to assess not only how effectively the proposed method aligns target domain images with the source domain, but also to evaluate whether the structural details are preserved well during the harmonization process. Furthermore, we investigate how well the proposed approach performs across different populations, encompassing both neonates and adults, and imaging modalities (i.e., T1-weighted and T2-weighted MRI). Last, to demonstrate the generalizability of our method, we explore its performance on the distinct task of neonatal brain gestational age estimation.

Datasets

MRI harmonization. It is important to recall that even though the empirical validation showcases the results across all the available sites, the proposed model only has access to a unique domain, i.e., the source, during the training steps and a unique target domain during the harmonization step. The details of the different datasets used for each task are provided below.

Adult brain MRI segmentation. In the context of adult brain MRI segmentation, we utilized data from a total of seven different sites.² Four of these sites are drawn from the Autism Brain Imaging Data Exchange (ABIDE) (Di Martino et al., 2014) dataset, which includes: California Institute of Technology (CALTECH), Kennedy Krieger Institute (KKI), University of Pittsburgh School of Medicine (PITT), and NYU Langone Medical Center (NYU) sets. Out of the remaining sites, Staten Island (SI) and Rutgers University (RU) sites are sourced from the Healthy Brain Network (HBN) (Alexander et al., 2017) dataset, which we refer to as HBNSI and HBNRI in this paper, along with data from the Open Access Series of Imaging Studies (OASIS) (Marcus et al., 2007). We selected T1-weighted MRIs of a healthy control population from each site, which were skull-stripped, motion-corrected, and quantized to 256 intensity levels. For each site, 60% of the images are used as the training set, 15% as the validation set, and the remaining 25% for testing, which are exploited in a 2D manner using the coronal plane slices. Also, the dimensionality of each 2D brain MRI slice is 256×256 with resolutions mentioned in Table 1. Moreover, following other large-scale studies (Dolz et al., 2018), we used Freesurfer (Fischl, 2012) to obtain the segmentations and grouped them into 15 labels: background, cerebral GM, cerebral WM, cerebellum GM, cerebellum WM,

² Please note that in the conference version of this work (Beizae et al., 2023), only four datasets were employed in the experiments, which explains the differences in the empirical results from both versions.

Table 1

Acquisition parameters across different sites. Scanner details and phenotypic information for each site used in this study. y: years; gw: gestation weeks.

Sites	# used MRIs	Scanner	TR (ms)	TE (ms)	Flip angle	Voxel size (mm ³)	Age
CALTECH	19	Siemens	1590	2.73	10	1.0 × 1.0 × 1.0	17–56 y
KKI	20	Phillips	8	3.70	8	1.0 × 1.0 × 1.0	8–13 y
NYU	20	Siemens	2530	3.25	7	1.3 × 1.0 × 1.3	6–39 y
PITT	20	Siemens	2100	3.93	7	1.1 × 1.1 × 1.1	9–35 y
HBNSI	77	Siemens	2730	1.64	7	1.0 × 1.0 × 1.0	5–21 y
HBNRU	57	Siemens	2500	3.15	8	0.8 × 0.8 × 0.8	5–21 y
OASIS	117	Siemens	9.7	4.0	10	1.0 × 1.0 × 1.25	18–96 y
DHCP T1-w	280	Philips	4795	8.7	N/A	0.8 × 0.8 × 0.8	37–45 gw
DHCP T2-w	333	Philips	12000	156	N/A	0.8 × 0.8 × 0.8	37–45 gw
V2LP T1-w	122	Siemens	2100	3.39	9	1.0 × 1.0 × 1.0	37–45 gw
V2LP T2-w	122	Siemens	8910	152	120	1.0 × 1.0 × 1.0	37–45 gw

CSF, ventricles, brainstem, thalamus, hippocampus, putamen, caudate, pallidum, amygdala and ventral DC.

Neonatal brain MRI segmentation and age estimation. We employed T1-weighted and T2-weighted MRIs from the developing Human Connectome Project (DHCP) (Makropoulos et al., 2018), VIBES2 (Spittle et al., 2014) and LaPrem (Cheong et al., 2021) datasets. As VIBES2 and LaPrem datasets are acquired with the same imaging device and parameters, we have combined them and considered them as one site, which we refer to as V2LP hereafter in this paper. All MRIs were sourced from a healthy control population aged between 37 and 45 weeks of gestational age. For neonatal brain MRI segmentation, we utilized the coronal view 2D slices, while the axial view 2D slices were chosen for brain age estimation due to their richer information content. Also, similar to adults, the dimensionality of 2D slices is 256×256 . The preprocessing and data splitting procedures for neonates were consistent with those used for adults, with the only difference being the inclusion of contours of 35 regional structures obtained from M-CRIBS (Adamson et al., 2020). Image acquisition parameters and device, number of used MRIs, and population age can be found in Table 1 for both adults and neonatal datasets. These values showcase how the selected datasets have distinct imaging devices and parameters.

Harmonization baselines. The proposed method is compared to a set of harmonization and image-to-image translation approaches. First, we apply either the segmentation or age regression network directly on non-harmonized images, which we refer to as a “Baseline”, so we can assess the improvement gained using the harmonized images. Furthermore, our comparison also includes the following harmonization strategies: Histogram Matching (Nyúl et al., 2000), Combat (Pomponio et al., 2020), SSIMH (Guan et al., 2022), two popular generative-based approaches, i.e., Cycle-GAN (Modanwal et al., 2020) and Style-Transfer (Liu et al., 2021), a source free latent-disentanglement harmonization (Imunity) (Cackowski et al., 2023), and a recent method for harmonization based on normalizing flows (BlindHarmony) (Jeong et al., 2023).

Test time domain adaptation and generalization baselines. In addition to the existing harmonization methods, the proposed approach is benchmarked against several test time domain adaptation and generalization methods. These methods include: aleatoric uncertainty estimation (AUE) (Wang et al., 2019), which uses test time augmentation to adapt to the target domain; BigAug (Zhang et al., 2020), which uses heavy augmentations on MRIs for generalization; and TENT (Wang et al., 2021a) and SAR (Niu et al., 2023), which are test-time adaptation methods based on the segmentation’s output confidence, i.e., entropy. The comparison with these methods aims to shed light on the importance of MRI harmonization and to reveal whether MRI harmonization can be replaced by domain generalization or test-time adaptation methods.

Evaluation metrics and protocol. *Segmentation.* To evaluate the performance of the proposed MRI harmonization approach on adult cross-site brain MRI segmentation, we train a segmentation neural network

$S_{\phi}(\cdot)$ using the training set of the source domain. Then for each target domain, the segmentation performance is evaluated using the harmonized images, where the segmentation performance is measured with the Dice Similarity Coefficient (DSC) and 95th percentile Hausdorff Distance (HD95). We repeat these steps considering each site as a source domain and the remaining sites as the target domains in each iteration, and then we report the average of these metrics. The evaluation on neonatal brain MRI segmentation follows the same protocol as that of the adults, which is repeated for each of the two modalities considered, i.e., T1-weighted and T2-weighted.

Neonatal brain age estimation. The neonatal brain gestational age estimation task follows the same procedure, except that instead of the segmentation model, a deep regression network has been trained. Following the nature of the task, the metrics used to assess its performance are the Mean Absolute Error (MAE) and the Mean Square error (MSE).

Harmonization performance. Last, following Parida et al. (2024), we resort to the Wasserstein distance (WD) (Kantorovich, 1960) between the normalized intensity histograms of harmonized images and those of the source domain. The WD calculates the smallest “cost” to change one distribution into the other, taking into account both the amount of change needed and how far the changes need to be shifted. Thus, the Wasserstein distance between two normalized histograms tells us how different they are by considering not only the differences in terms of change magnitudes (such as the KL divergence does), but also how far apart the differing parts are. In our experiments, and as stated earlier, we consider one single site as a source domain (while all the remaining domains/sites remain unknown). Furthermore, we report the average of the WD for all source-domain pairs.

Implementation details.

Normalizing Flow (FL). The NF model has been trained for 20,000 iterations with Adam optimizer and a batch-size of 32, an initial learning-rate of 1×10^{-3} , and using a weight-decay of 0.5 every 2000 iterations. We employ a U-shaped architecture within the coupling layers, comprising four different scales with a scaling factor of 2. Each layer consists of an activation function followed by a convolutional layer and a normalizing layer. The activation function used in each scale is a modified version of ELU, i.e., $\text{concat}[\text{ELU}(x), \text{ELU}(-x)]$, which makes it easier for the NF model to map to normal distribution due to its symmetry properties. Additionally, there are 16, 32, 48, and 64 kernels in each scale, respectively. To construct the NF model, we initially employ four sequential coupling layers with checkerboard masking, aimed at capturing the noise distribution through variational dequantization. Then, this is followed by four identical coupling layers and a squeezing function as explained in Dinh et al. (2017) to decrease the spatial dimension. Then, we subsequently incorporate four more coupling layers employing a channel-masking strategy, and an additional layer of feature squeezing, followed by a final series of four coupling layers employing channel-masking. The overall architecture of the flow model is shown in Fig. 1. Also, the margin c used for guiding

Table 2

Performance overview on the cross-site adult MRI segmentation task. Segmentation performance, in terms of DSC and HD95 metrics, across different harmonization approaches. To facilitate the strengths and weaknesses of different methods, we also indicate whether they are *source-free* (SF), *task-agnostic* ($\mathcal{T}\mathcal{A}$), and can handle *unknown-domains* ($\mathcal{U}\mathcal{D}$), as well as the different strategy they fall in. The best results are highlighted in **bold**.

Method	Strategy	SF	$\mathcal{T}\mathcal{A}$	$\mathcal{U}\mathcal{D}$	DSC (%)	HD95 (mm)
Baseline	–	–	–	–	36.3 ±4.8	43.1 ±5.3
Hist matching (Nýúl et al., 2000)	Harmonization	✓	✓	✓	63.7 ±4.8	12.3 ±2.7
Combat (Pomponio et al., 2020)	Harmonization	✗	✓	✗	73.0 ±4.3	6.3 ±2.6
Cycle-GAN (Modanwal et al., 2020)	Harmonization	✗	✓	✗	75.0 ±2.8	5.5 ±1.7
Style-transfer (Liu et al., 2021)	Harmonization	✗	✓	✗	70.8 ±5.7	8.2 ±2.6
SSIMH _{MLMI'22} (Guan et al., 2022)	Harmonization	✗	✓	✓	59.4 ±5.1	12.3 ±2.9
ImUnity _{MedIA'23} (Cackowski et al., 2023)	Harmonization	✗	✓	✓	58.2 ±4.6	17.0 ±3.3
BlindHarmony _{ICCV'23} (Jeong et al., 2023)	Harmonization	✓	✓	✓	62.2 ±6.1	13.7 ±3.2
AUE (Wang et al., 2019)	Test time augmentation	✓	✗	✓	35.4 ±3.6	24.5 ±3.6
BigAug _{TMI'20} (Zhang et al., 2020)	Generalization	✓	✗	✓	82.0 ±2.1	3.2 ±0.8
TENT _{ICLR'21} (Wang et al., 2021a)	Test time adaptation	✓	✗	✓	72.8 ±3.4	7.4 ±2.2
SAR _{ICLR'23} (Niu et al., 2023)	Test time adaptation	✓	✗	✓	70.9 ±4.0	9.9 ±2.6
Harmonizing flows	Harmonization	✓	✓	✓	82.9 ±2.3	3.1 ±1.0

the NF model (Eq. (4)) is empirically set to 1.2. For adapting the harmonizer network using images of the target domain during test time, the Normalizing Flow model is frozen, and the harmonizer model is updated slightly using Adam optimizer with a learning rate of 5×10^{-7} , and a batch-size of 32 until the epoch where the stopping criterion has been reached. All the models were implemented on PyTorch using two NVIDIA RTX A6000 GPU cards.

Harmonizer. The utilized UNet network as harmonizer consists of five different scales with a scaling factor of 2, where each scale includes a layer of the modified ELU activation function, i.e., $\text{concat}[\text{ELU}(x), \text{ELU}(-x)]$, followed by two convolutional layers. The number of kernels of the convolutional layers for each scale is 16, 32, 48, 64, and 64 respectively.

Segmentation and regression models. The segmentation network employed in our empirical validation is nn-UNet (Isensee et al., 2021) with batch-normalization, stride 2, and kernel size of 3, whereas ResNet18 (He et al., 2016) serves as the age estimation network. Furthermore, all the segmentation, age estimation, and harmonizer networks are trained for 5000 iterations with Adam optimizer with an initial learning rate of 1×10^{-3} , a weight decay of 0.5 every 500 iterations, and a batch-size of 64. Last, the model at the best iteration, based on an independent validation set, is utilized.

It is worth mentioning that we did not use any type of contrast or intensity augmentation in the training of segmentation and regression networks, to better capture the effect of the harmonization methods and make them more sensitive to distribution shifts.

4.2. Results

In this section, we reported the empirical results of the experiments performed. In particular, we first resorted to the cross-site brain MRI segmentation results to evaluate the performance of different harmonization approaches, as the segmentation task is a good indicator of harmonization performance. Following the main results, we conducted a series of comprehensive ablation studies to empirically support our choices. Then, to show the task-agnostic nature of our method, we evaluated different harmonization strategies on the task of neonatal cross-site brain age estimation. Last, we took a closer look at the distance between the intensity histograms of the harmonized images with the source images. Surprisingly, closer intensity histograms between harmonized images with the source domain does not necessarily correlate with better harmonization performance or higher target task performance, such as segmentation or regression.

4.2.1. Performance on the segmentation task

Main results To evaluate the proposed harmonization method, cross-site brain MRI segmentation performance has been obtained before and

after the harmonization. As the segmentation networks are fixed, the segmentation results' improvement indicates the impact of the harmonization. Segmentation results obtained with the images harmonized by different methods are reported in Table 2.

Comparison to harmonization methods. Before looking at the segmentation results, we would like to highlight that most existing methods rely on assumptions that could limit their scalability and practicality in real-life scenarios. First, some methods must access at least one image of the source domain during the adaptation, thereby not being completely *source-free* (SF). Also, most harmonization techniques need to access the target domains during training, while ideally, the potential target domains should remain unknown, which we refer to as *unknown-domains* ($\mathcal{U}\mathcal{D}$). We relax all these assumptions by proposing a method that is *source-free*, *task-agnostic*, and unaware of potential target domains during training. First, from Table 2, we can observe that the proposed approach improves the segmentation results by more than 45% in terms of DSC over the baseline, i.e., without harmonization. Furthermore, it consistently outperforms other harmonization approaches by a significant margin, in terms of both segmentation metrics. To statistically validate the performance of our method, we conducted a paired t-test between our method and each of the comparing harmonization methods. The test was performed using the results from different source datasets as the data samples for each comparison. Across all comparisons, the maximum p -value observed was $p = 0.012$, which is below the commonly accepted significance threshold of 0.05, indicating that the observed improvements are consistently statistically significant. Particularly, the average gained improvement compared to CycleGAN, the next best-performing method, is larger than 8% in terms of DSC, and 2.4 mm smaller in terms of HD95. Considering that CycleGAN requires the source domain, as well as all the target domains to adapt, the differences in performance are even more important. Furthermore, if we compare it to approaches that offer the same benefits, i.e., (SF), ($\mathcal{T}\mathcal{A}$) and ($\mathcal{U}\mathcal{D}$), such as Hist matching, and BlindHarmony these differences increase up to 20%.

Comparison to test-time domain adaptation and generalization approaches. To better evaluate the method, we have compared it against common strategies to tackle distributional shift, including: domain generalization (Zhang et al., 2020), Test time Augmentation (Wang et al., 2019), and Test time adaptation (Wang et al., 2021a; Niu et al., 2023). This comparison highlights the importance of image harmonization and examines whether harmonization can be replaced by any of these strategies. These results, which are depicted at the bottom of Table 2, demonstrate that the proposed harmonization method also outperforms well-known test-time domain adaptation and generalization strategies in the task of cross-site brain segmentation. Furthermore, individual cross-site brain MRI segmentation results are depicted in Fig. 2, for a better interpretability of the per-site results. In every matrix, the



Fig. 2. Cross-site brain MRI segmentation matrix across the compared methods. Each cell indicates the segmentation result (DSC %) when the source dataset (in the rows) is used to harmonize each target dataset (in columns).

Table 3

Performance overview on the cross-site neonatal MRI segmentation task. Segmentation performance, in terms of DSC and HD95 metrics, across different harmonization approaches. To facilitate the strengths and weaknesses of different methods, we also indicate whether they are *source-free* (SF), *task-agnostic* (TA), and can handle *unknown-domains* (UD), as well as the different strategy they fall in. The best results are highlighted in **bold**.

Method	SF	TA	UD	T1-w		T2-w	
				DSC (%)	HD95 (mm)	DSC (%)	HD95 (mm)
Baseline	–	–	–	34.7±6.5	38.4±15.8	44.2±17.1	30.8±17.0
Hist matching (Nyúl et al., 2000)	✓	✓	✓	42.9±8.0	32.9±17.6	84.3±4.1	4.7±2.3
Combat (Pomponio et al., 2020)	✗	✓	✗	83.0±4.5	2.7±1.4	88.0±1.5	2.0±1.1
Cycle-GAN (Modanwal et al., 2020)	✗	✓	✗	77.2±6.5	4.8±3.1	88.3±1.5	1.8±0.9
Style-transfer (Liu et al., 2021)	✗	✓	✗	56.7±9.0	19.8±9.9	88.6±1.4	1.5±0.2
SSIMH _{MLMF22} (Guan et al., 2022)	✗	✓	✓	38.7±9.7	30.2±13.3	61.9±8.5	17.9±6.6
Imunity _{MedIA'23} (Cackowski et al., 2023)	✓	✓	✓	41.8±7.9	23.8±9.5	87.1±3.0	2.4±1.8
BlindHarmony _{ICCV'23} (Jeong et al., 2023)	✓	✓	✓	56.1±9.2	28.4±12.6	88.0±1.8	2.9±1.6
AUE (Wang et al., 2019)	✓	✗	✓	38.0±6.1	25.5±7.0	68.7±5.2	15.5±1.9
BigAug _{TMI'20} (Zhang et al., 2020)	✓	✗	✓	84.1±3.8	2.2±53.9	90.0±1.3	1.4±0.3
TENT _{ICLR'21} (Wang et al., 2021a)	✓	✗	✓	76.3±3.8	4.7±2.2	88.1±1.6	2.4±1.6
SAR _{ICLR'23} (Niu et al., 2023)	✓	✗	✓	72.3±5.1	5.1±2.2	88.9±1.7	1.7±0.6
Harmonizing flows	✓	✓	✓	84.4±2.5	2.1±1.0	89.6±1.4	1.4±0.4

diagonal elements represent the segmentation of intra-site brain MRI and establish the upper bound of segmentation results when test images originate from the source domain. Then, the elements outside the diagonal indicate the segmentation result when a given dataset (indicated in the rows) is used to harmonize the target datasets (in columns). As the non-diagonal elements approach the diagonal, it can be interpreted as an enhanced capacity of the method to handle distributional shifts. Based on these results, we can state that our approach proved the most effective in this aspect, and consistently across all source-domain datasets.

Validation on a different population and image modality

To investigate the scalability of our harmonization method, we evaluate it using different population demographics and image modalities. Particularly, we resort to the neonatal brain MRI segmentation task, which differs from the more traditional adult brain MRI segmentation. Furthermore, this evaluation has been repeated for both T1-weighted and T2-weighted MRIs to explore whether the proposed approach can

yield satisfactory performance for both modalities. According to the results from this experiment, which are reported in Table 3, our method demonstrates comparable performance for both modalities in neonatal brain MRI segmentation. More concretely, the proposed harmonization strategy obtains the best results in T1-weighted, whereas it ranks first among all compared methods and second by a small margin among all compared methods for both metrics in T2-weighted images. These values underscore the effectiveness of our proposed approach across diverse populations and modalities.

Ablation studies

I-Impact of normalizing flows. This section assesses the impact of each component of the proposed method, which is achieved by comparing the cross-site brain MRI segmentation results obtained: (i) when images are not harmonized, (ii) when harmonized just with the proposed pre-trained harmonizer network θ^{init} , or (iii) harmonized with the proposed method. The results of this ablation study, which are depicted in Fig. 3, empirically support that the proposed NF-based models




Fig. 3. Effect of each component of the harmonizing flow. Particularly, we depict the improvement gained using the proposed pre-trained harmonizer network (~44 DSC%), and the adaptation using normalizing flows (~2.0 DSC%).

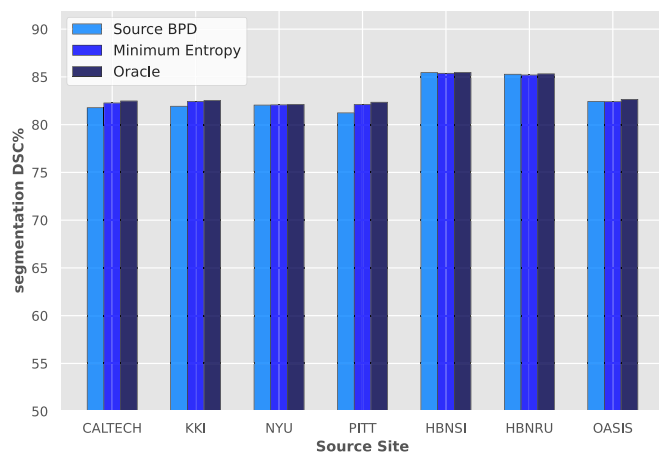


Fig. 4. The effect of different stopping criteria to stop the harmonizer network adaptation. Oracle represents selecting the best iteration based on the performance of the target task (i.e., segmentation in this example), which serves as the upper bound. Both criteria, *source BPD* and minimum entropy, provide good stopping points, with a slight advantage of the minimum entropy criterion.

can serve as an effective mechanism to adapt the harmonizer network. First, the proposed approach to pre-train the harmonizer network yields significant enhancements compared to non-harmonized images (nearly 44% of DSC on average) while, despite its simplicity, there is no need to know the target domain in advance. Furthermore, adapting the harmonizer network using the proposed NF model further improves the cross-site MRI segmentation results, with nearly 2.0% of DSC on average, illustrating the effectiveness of the proposed harmonization strategy.

II-Adaptation stopping criterion. Now we explore the crucial issue of determining the appropriate iteration to stop the adaptation, specially as updating the harmonizer network is conducted in an unsupervised manner. Three different criteria to stop adapting the harmonizer network are explored. The first criterion is to stop in an iteration where the Shannon entropy of the target task predictions (segmentation in this case) reaches its minimum. According to Wang et al. (2021a), the Shannon entropy of the predictions of the target task is highly correlated with its performance. As it does not require any labeled data, it provides an *a priori* reliable stopping criterion for fine-tuning the harmonizer network. As a second alternative, we stopped the harmonizer network adaptation when the target *BPD* reaches the observed *BPD* on the source

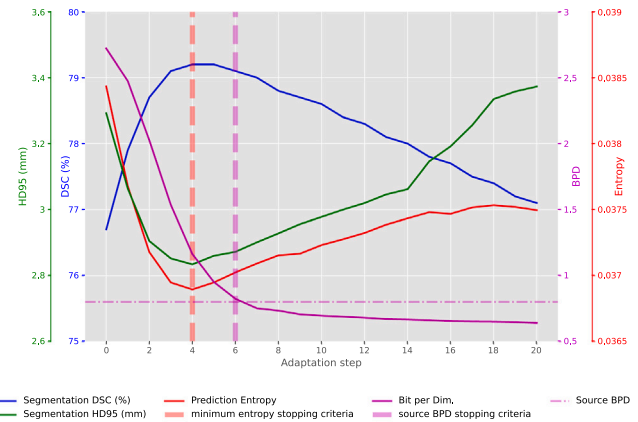


Fig. 5. Which is the best metric as stopping criteria? This plot depicts different metrics during the adaptation of the harmonizer network (from HBNSI to NYU). Step zero corresponds to using the initial harmonizer network without adaptation. The vertical lines show the stopping time-points based on two proposed stopping criteria: minimum entropy of the predictions (*red*) and reaching source BPD (*purple*).

domain (which can be computed during training time using a validation set). As opposed to the first criterion, this criterion is task-agnostic and well-suited for unsupervised tasks or scenarios where entropy calculations are not applicable (i.e., regression problems). Finally, we directly used the target task (segmentation) performance and stopped the adaptation once it achieved the highest DSC score which we refer to as the *Oracle*. Please bear in mind that this criterion is impractical in real-world scenarios, as one may not have access to segmentation labels,³ and it only intends to serve as an upper bound to show how much we can gain using a well-defined stopping criterion. As shown in Fig. 4, despite minimum entropy being a slightly better criterion compared to *source BPD*, both yield similar performances, with the *source BPD* providing a more general strategy, as it is not tailored to a task that requires probabilistic outputs. In summary, both stopping criteria prove to be viable options, as their outcomes closely resemble those of the *Oracle*.

Furthermore, in Fig. 5 we show, for a given scenario – harmonizing HBNSI dataset to NYU – the strong correlation between the segmentation metrics with both the segmentation prediction entropy and proximity to the source BPD. In particular, the evolution of the segmentation metrics is depicted in green (HD95) and blue (DSC) curves, with the values of the entropy of the segmentation predictions in red, and the target BPD in purple. We can observe that the point for which these metrics are optimal (i.e., minimum entropy and target BPD matching source BPD), is actually close. Nevertheless, we advocate that, for more general use, looking at the source BPD should be a preferred option, as it does not depend on the target task.

III-Ablation studies on components and hyper-parameters. This section aims to empirically support the choices made in the proposed harmonization strategy.

III.a-Flow Depth: We explored the effect of the number of coupling layers in the normalizing flow network in learning the distribution of the source domain. Particularly, we investigated using 6, 12, and 18 coupling layers in the normalizing flow network and its effect on the distribution learning capacity of the model and harmonization process. As can be seen in Fig. 6(a), 12 coupling layers were the optimal choice and resulted in better harmonization. We believe that using 6 coupling layers does not fully capture the source domain distribution. On the other hand, using 18 coupling layers makes the training process harder,

³ Please note that using segmentation, or other kind of labels, for the stopping criteria would make the model task-dependent.

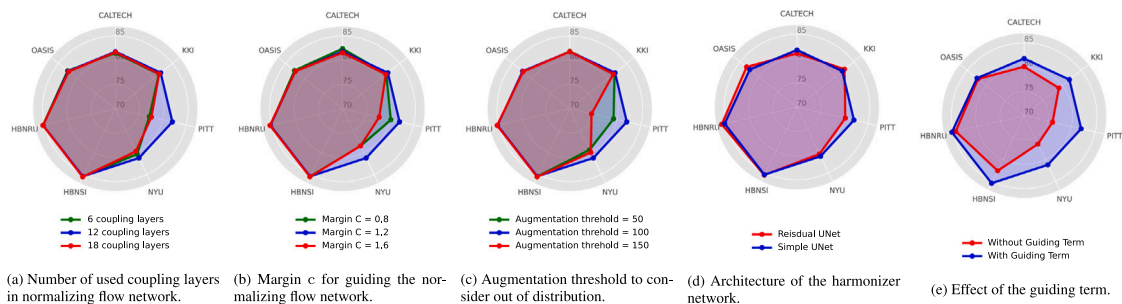


Fig. 6. Ablation study on architecture components and hyper-parameters of the harmonizing flow. The best performance is obtained when the number of coupling layers in the normalizing flow is set to 12, the guiding term is used and the guiding margin c to 1.2, the augmentation threshold is 100, and a simple UNet is chosen as the harmonizer network.

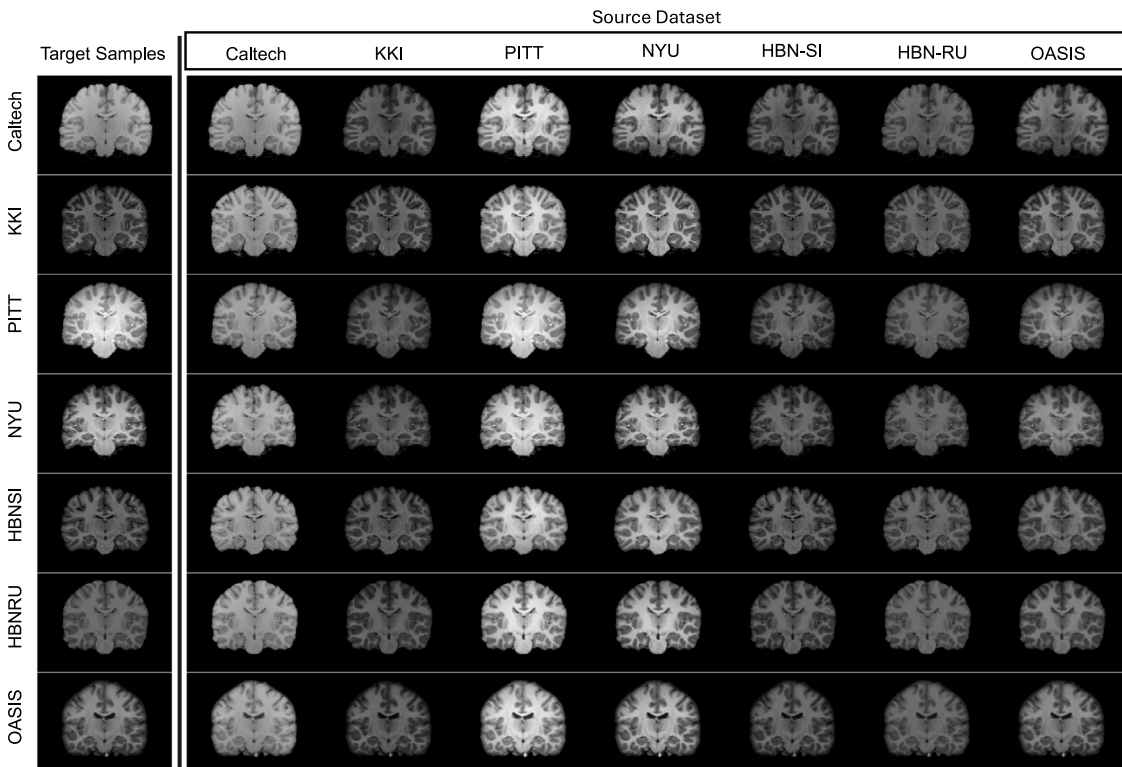


Fig. 7. This figure showcases examples of harmonized images generated by the proposed method. The first column presents the sample images from various target domains. The subsequent columns display the images harmonized to specific source domains, as indicated at the top of each column. Therefore, each row maintains consistent anatomical structures, while each column shares the same visual characteristics.

and it probably needs more data to be fully trained.

III.b-Margin c : In this section, we investigate the influence of the margin c on the guidance of the normalizing flow process. Our initial choice of 1.2 for c at the first level stemmed from observing an average source BPD of approximately 0.8 when the normalizing flow was trained without guidance. Consequently, we opted for a value of 1.2. To further investigate the effect of this value, we conducted additional experiments with margin values c set to 0.8 and 1.6. As illustrated in Fig. 6(b), the initial choice yielded better performance in comparison to 0.8 and 1.6.

III.c-Augmentation threshold: Defining how much augmentation is enough for an augmented image to be considered as out of distribution to train the normalizing flow is an important step. This threshold, which is defined in terms of mean squared distance with original images, was first selected as 100, by visual inspection of the images. After the initial choice, we explored how adjusting this hyper-parameter affects the constraint on the normalizing flow network and the harmonization process. We conducted experiments using augmentation

thresholds of 50 and 150, whose results (Fig. 6(c)) demonstrate 100 to be an optimal choice for this hyperparameter.

III.d-Harmonizer Network: We investigated two U-shaped architectures for the harmonizer network. First, a conventional UNet architecture was employed. Subsequently, like the conference version of this paper (Beizae et al., 2023), we utilized a modified UNet to extract two separate sets of values. The final layer of the network (β) serves as a bias value, matching the input image’s dimensions. Additionally, a scalar value α , derived from the network’s middle layer, acts as a scale parameter. In this way, the harmonizer’s output can be expressed as $h_{\theta}(x) = \alpha * x + \beta$. As depicted in Fig. 6(d), both options are effective, with a slight superiority of the simple UNet. We believe that this marginal superiority is due to the greater degrees of freedom in simple UNet, which might be better for transferring complex distributions.

III.e-Guiding Term: In this section, we examine the impact of the guiding term by analyzing the effects of its removal from the training objective of the NF model (Eq. (6)). As illustrated in Fig. 6(e), incorporating the guiding term significantly enhances the results, underscoring




Fig. 8. Visual examples of neonatal MRI brain images harmonized using the proposed method. The left section displays the coronal and axial planes for both modalities of a sample image from the DHCP dataset (target) alongside their harmonized counterparts to the V2LP dataset (source). The right section shows the reverse, with V2LP images harmonized to the DHCP dataset.

its critical role in enabling the NF model to capture domain-specific characteristics of the source dataset. Conversely, omitting the guiding term often leads to suboptimal performance, adversely affecting the harmonization process. In many cases, the NF model trained without the guiding term produces results that are inferior even to those obtained with the initial harmonizer alone.

Qualitative results.

Fig. 7 showcases the instances of harmonized images using the proposed method across different source and target sites, for the adult brain MRIs. In particular, we randomly picked a sample from each site and then mapped it to different target sites. As can be seen, the harmonized samples in each column share the same visual characteristics, while on each row, the details of the harmonized images are preserved. These qualitative results illustrate that, regardless of the source or the target domains, the proposed method consistently produces reliable harmonized images, which is supported quantitatively by the comprehensive empirical validation conducted. Furthermore, visual examples depicted in Fig. 8 (neonatal brains) demonstrate the effectiveness of our approach in harmonizing inter-site images, regardless of the modality and plane used.

In Fig. 10, we have visualized a harmonized sample from the target domain (CALTECH here) to the source domain (KKI here) using different harmonization methods. The harmonized image using our proposed method appears to have the closest visual characteristics with the source domain compared to other harmonization methods. Combat Pomponio et al. (2020) is excluded from this figure as it attempts to remove variations between domains and cannot transform the target domain to match the source domain.

4.2.2. Performance on neonatal brain age estimation

In the previous section, we demonstrated the generalization capabilities of the proposed approach, by showing its superiority when evaluating the harmonized images in adults and neonatal brain MRI segmentation, and across multiple modalities. In this section, we will take one step further and evaluate the quality of the harmonization based on a different task, i.e., neonatal brain age estimation, which

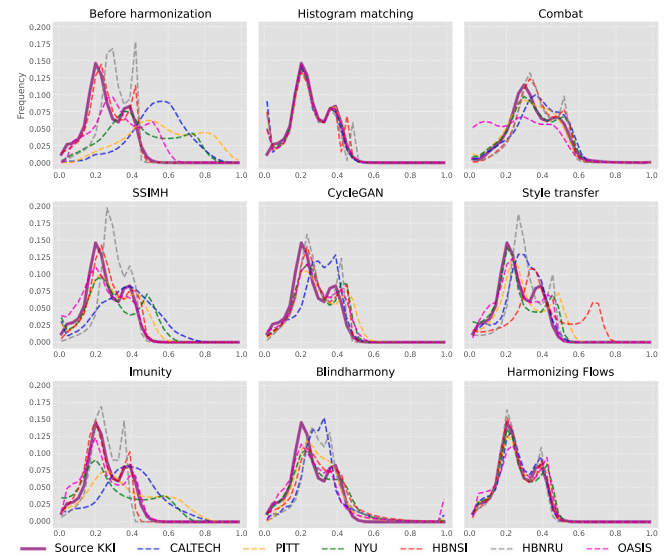


Fig. 9. Histograms of the harmonized MRIs from multiple target datasets compared to the histogram of the source MRIs (KKI dataset, in purple) for all harmonization methods.

necessitates regression neural networks and the utilization of different 2D slices (specifically the axial view). As depicted in Table 4, our method also outperforms all compared methods for both metrics and both modalities, highlighting its effectiveness for various tasks and applications, as well as on different 2D planes. Thus, based on the results of the segmentation and regression tasks, we can state that the proposed harmonization strategy leads to a significantly more flexible solution with substantial improvement gains.

4.2.3. A closer look at the harmonization performance

In previous sections, we assessed the quality of the harmonization based on the performance of different target tasks, i.e., adult and neonatal brain segmentation and neonatal brain age estimation. A natural question that arises is whether the harmonization performance can be quantitatively evaluated without requiring further labeled target tasks or *traveling subjects*. Indeed, recent literature (Parida et al., 2024) in evaluating harmonization techniques has proposed the use of the Wasserstein distance, to measure the similarity between harmonized and source image intensity histograms. The reasoning behind using this kind of divergence is two-fold. First, standard harmonization metrics, such as mean absolute error, mean squared error, or peak signal-to-noise ratio, can provide high-quality measurements, but at the price of requiring paired harmonized data, i.e., *traveling subjects*. Second, while other metrics, such as structural similarity (SSIM), can help mitigate the need for paired data, they primarily compare structures at a higher level, potentially overlooking smaller artifacts or hallucinations introduced by generative models. In this section, we follow up on the recent work in Parida et al. (2024), and assess the harmonization performance of different strategies based on the Wasserstein distance between intensity histograms.

To this end, we first depict in Fig. 9 the histogram distributions of the harmonized MRIs from multiple target domains compared to the MRIs from the source site (the KKI dataset in this example), across all harmonization methods. Looking at these plots, we can observe that Histogram Matching, as well as our approach, yields to the closest intensity histogram distributions to the source domain (in purple), after harmonization. Furthermore, this behavior is consistent for all the target sites for both Histogram Matching and the proposed approach. Thus, from a pure harmonization standpoint, these plots tell us that Histogram Matching and our harmonization strategy are capable of mapping images from a target to a source site such that the harmonized

Table 4

Performance overview on the cross-site neonatal age estimation task. Brain age estimation performance, in terms of Mean Absolute Error (MAE) and Mean Squared Error (MSE) metrics, across different harmonization approaches and modalities. To facilitate the strengths and weaknesses of different methods, we also indicate whether they are *source-free* (*SF*), *task-agnostic* (*TA*), and can handle *unknown-domains* (*UD*), as well as the different strategy they fall in. The best results are highlighted in **bold**.

Method	<i>SF</i>	<i>TA</i>	<i>UD</i>	T1-w		T2-w	
				MAE	MSE	MAE	MSE
Baseline	–	–	–	1.15 ± 0.91	2.23	1.67 ± 1.15	4.67
Hist. matching (Nyúl et al., 2000)	✓	✓	✓	1.04 ± 0.78	1.77	1.09 ± 0.84	2.05
Combat (Pomponio et al., 2020)	✗	✓	✗	1.33 ± 0.84	2.49	1.28 ± 0.82	2.33
Cycle-GAN (Modanwal et al., 2020)	✗	✓	✗	1.05 ± 0.71	1.62	0.89 ± 0.63	1.24
Style-transfer (Liu et al., 2021)	✗	✓	✗	1.75 ± 1.17	5.15	0.96 ± 0.67	1.41
SSIMH (Guan et al., 2022)	✗	✓	✓	1.18 ± 0.96	2.33	1.41 ± 0.85	2.35
Imunity (Cackowski et al., 2023)	✗	✓	✓	1.10 ± 0.81	1.89	1.55 ± 1.00	3.82
BlindHarmony (Jeong et al., 2023)	✓	✓	✓	1.19 ± 0.98	2.45	1.29 ± 0.90	2.62
AUE (Wang et al., 2019)	✓	✗	✓	1.32 ± 1.01	2.91	1.4 ± 1.01	3.51
BigAug (Zhang et al., 2020)	✓	✗	✓	1.05 ± 0.62	1.55	1.09 ± 0.80	1.88
Harmonizing flows	✓	✓	✓	1.01 ± 0.69	1.51	0.82 ± 0.68	1.21




Fig. 10. Example of harmonized images using different methods. The first image is the source domain sample (KKI here), the second one shows a sample from the target domain (CALTECH here), and the rest, shows the harmonized target sample using different methods.




Fig. 11. Segmentation (DSC%) and Age estimation (MAE) results vs. WD of intensity histograms for compared Harmonization methods.

intensity histograms almost match perfectly those of the source domain.

Nevertheless, perfect alignment between intensity histograms is not the ultimate goal. It should be noted that different individuals’ brain structures (for example between a child and an old person) necessitate different intensity histograms. Therefore, aligning intensity histograms is not an ideal solution. Also, we believe that the ability to enhance the performance of subsequent tasks, such as segmentation, classification, or brain age estimation in our context, which relies on harmonized images when dealing with multi-centric data, serve as a better indicator of the harmonization potential. However, if we analyze the performance on target tasks when employing harmonized data, we can easily observe that histogram matching indeed generates harmonized images that lead to poor performance on multiple target tasks. For example, in adult brain segmentation, resorting to histogram matching as a harmonization method results in an average DSC of 63.6, substantially lower than other methods which, a priori, had the highest misalignment between intensity histograms of the harmonized images with the source one in terms of WD (e.g., SSIMH or Imunity).

To better understand whether a correlation between the intensity histogram distance and the target task performance exists, we plot *task-performance vs. intensity-histogram-distance (WD)* in Fig. 11. Indeed, we can observe from this figure that, in most cases, a closer intensity histogram (i.e., lower WD) corresponds to better performance in the target task (i.e., higher DSC in segmentation and lower MAE in brain age estimation). A method warranting further discussion is histogram matching, which shows a weak correlation, if any, between intensity histogram distance in terms of WD and target task performance. Its

good performance in the WD metric is somehow expected, as it directly optimizes the histogram of intensities for MRI harmonization. Nevertheless, we believe that solely forcing the histograms to be close is not necessarily a good condition to yield usable harmonized images, as the spatial variation across intensities is not considered when matching distributions. Harmonized images depicted in Fig. 10 could further illustrate this point visually. The remaining harmonization strategies, however, seem to be generally correlated, with approaches obtaining lower WD values also achieving better performances on the respective target tasks. In particular, our proposed approach always ranks among the top-three approaches in terms of WD, yielding the best performance in segmentation and age estimation tasks among compared harmonization methods. These findings are based on the metrics proposed in the recent work (Parida et al., 2024), and may not hold for a more in-depth list of harmonization metrics. Ideally, while some other metrics could have been explored to evaluate the harmonization performance, some of the metrics used depend on *traveling-subjects*, which limits their applicability to many scenarios. We also stress that having an in-depth evaluation of the harmonization performance and assessing the best harmonization metric, is not within the scope of this work. In contrast, we believe that based on the results reported in this work, solely relying on the closeness of intensity histograms may not be sufficient, and evaluating the quality of the harmonization on target task indicators (e.g., DSC in segmentation or MAE in regression tasks) could serve as a more general manner to assess the harmonization quality.




Fig. 12. Visualization of harmonization applied to a sample of neonatal injured (atrophy and large ventricles) brain MRI from the V2LP dataset, harmonized to the DHCP dataset. The results demonstrate that the model preserves the structural integrity of the input while performing harmonization, even in the presence of injury-related anomalies.

4.2.4. Structural integrity

One of the main concerns regarding learning-based MRI harmonization methods is anatomy hallucination (Cohen et al., 2018). Specifically, one issue that arises is whether training the normalizing flow model on a population different from the target domain with distinct anatomical variations could result in unintended changes to anatomical structures. As such, it is essential to implement strategies that ensure that structural integrity is preserved throughout the harmonization process. We have employed several techniques to mitigate this risk effectively.

First, to ensure robustness to structural variations, we applied a range of geometric augmentations while training the normalizing flow and harmonization model. These augmentations included resizing, scaling along the X, Y, and Z axes, shearing, perspective transformations, and deformations like piece-wise affine. These augmentations help the model focus on domain-characteristic features, such as brightness and contrast, while reducing sensitivity to structural differences. Second, during the pre-training of the harmonization network, we used the Structural Similarity Index (SSIM) loss, which, combined with geometric augmentation could help promote structural consistency and alignment between the input and output images. Finally, To further validate structural consistency, we specifically selected a segmentation task, which contains detailed regional contours, as a downstream evaluation. despite differences in age ranges between adult datasets, our method demonstrated superior performance compared to other harmonization techniques, providing empirical evidence of its structural fidelity.

It is important to note that evaluating the performance of harmonization in the presence of brain injuries, where the target domain contains injured brains and the source domain consists of normal brain MRIs, is not straightforward in our setting. Downstream tasks such as segmentation and age estimation, which are trained on normal brains, cannot be applied to abnormal brains, as these models are likely to fail. To demonstrate the structural integrity of the proposed harmonization method in the presence of brain injuries, we visually examined a neonatal brain MRI with atrophy and large ventricles from the V2LP domain before and after harmonization to the DHCP dataset. The results, depicted in Fig. 12 for both T1 and T2 modalities, reveal that the structures of the injured areas are preserved, highlighting the structural integrity of the proposed harmonization approach.

4.2.5. Friedman ranking

To fairly compare the performance of the different harmonization methods across various metrics and tasks, we resort to the Friedman Rank (Friedman, 1937, 1940), which has been employed for this purpose in the literature (Wang et al., 2022; Murugesan et al., 2024). The **Friedman Rank** is defined as:




Fig. 13. Friedman Rank for the compared harmonization methods: (a) Harmonizing Flows (Ours), (b) BlindHarmony, (c) Imunity, (d) SSIMH, (e) Style-transfer, (f) CycleGAN, (g) Combat, (h) Hist-matching.

$$\text{rank} = \frac{1}{S_m} \sum \text{rank}_i$$

where S_m is the number of evaluation settings and rank_i is the rank of a method in the i th setting. Thus, the lower the rank obtained by an approach, the better the method is. In our scenario, we have 13 different settings: DSC and HD95 in 3 segmentation tasks, MAE and MSE in 2 regression problems, and 3 WD values in harmonization (one for adults and two for neonatal brains). The results from the Friedman Ranking across all analyzed methods are depicted in Fig. 13. As it can clearly observed, our proposed method (*last column*) achieves the best Friedman rank among all compared harmonization methods, demonstrating their overall superiority across different scenarios.

5. Conclusion

In this work, we proposed a novel harmonization method that leverages Normalizing Flows to guide the adaptation of a harmonizer network. Our approach is source-free, task-agnostic, and works with unseen domains. These characteristics make our model applicable in real-life problems where the source domain is not accessible during adaptation, target domains are unknown at training time and harmonization is task-independent. Furthermore, another advantage of our method over the existing approaches is that it only requires images from one source domain, and one modality, at the training time.

Through extensive comparisons with other harmonization methods, as well as test-time domain adaptation and generalization approaches, our method consistently proved its superiority in multiple medical image problems, yet relaxing the strong assumptions made by existing harmonization strategies. Furthermore, we validated the scalability of our proposed method by evaluating it on a different population (neonates), different modalities (T1-weighted and T2-weighted MRIs), and different tasks (neonatal brain age estimation and segmentation). The results demonstrated comparable performance across diverse populations and modalities, highlighting the robustness and versatility of our approach. It is worth noting that while this paper applies harmonization to 2D brain MRI slices, the method can be readily extended to 3D by replacing 2D convolutional layers with 3D ones. However, this extension would require a larger dataset and more computational resources due to the increased number of parameters. Additionally, although no visible inconsistencies between harmonized slices were observed in our experiments, averaging harmonization results across all 2D planes (i.e., axial, coronal, and sagittal) could further mitigate potential inconsistencies. Qualitative results further supported the reliability and effectiveness of our method, illustrating consistent and reliable image-to-image mappings across different target domains. Last, even compared to recent test-time adaptation strategies, empirical

results suggest that the proposed method is a powerful alternative to deal with the presence of domain drifts, more particularly for MRI multi-site harmonization.

In conclusion, our proposed harmonization method offers a promising solution for addressing distributional shifts in medical image analysis, paving the way for improved performance and generalizability across diverse datasets, and enabling the use of large-scale multi-centric studies.

Limitations: Although the proposed harmonization method demonstrates promising potential, there are a few considerations. First of all, it requires additional inference time for test-time adaptation, which may impact efficiency in certain time-sensitive applications. Moreover, while efforts have been made to prevent the normalizing flow model from capturing biases or anatomical structures in the source domain, it may still inadvertently learn some biases, such as image quality variations. Furthermore, our approach assumes that the source and target domains have close resolutions, which may limit its flexibility when applied to datasets with varying resolutions.

CRedit authorship contribution statement

Farzad Beizae: Writing – original draft, Visualization, Validation, Methodology, Investigation. **Gregory A. Lodygensky:** Writing – review & editing, Supervision, Resources, Funding acquisition, Data curation. **Chris L. Adamson:** Software, Data curation. **Deanne K. Thompson:** Writing – review & editing, Validation, Data curation. **Jeanie L.Y. Cheong:** Writing – review & editing, Data curation. **Alicia J. Spittle:** Writing – review & editing, Data curation. **Peter J. Anderson:** Writing – review & editing, Data curation. **Christian Desrosiers:** Writing – review & editing, Writing – original draft, Supervision, Resources, Methodology, Investigation. **Jose Dolz:** Writing – review & editing, Writing – original draft, Supervision, Resources, Methodology, Investigation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work is supported by the National Science and Engineering Research Council of Canada (NSERC), via its Discovery Grant program and IVADO - Grants for Fundamental Research Projects. We also thank Calcul Quebec and Compute Canada.

Data availability

I have used total of 10 different datasets for this manuscripts. Among these datasets, 8 of them are publicly available for research purpose upon accepting the agreement, and 2 datasets are not public.

References

- Abdelhamed, A., Brubaker, M.A., Brown, M.S., 2019. Noise flow: Noise modeling with conditional normalizing flows. In: ICCV. pp. 3165–3173.
- Adamson, C.L., Alexander, B., Ball, G., Beare, R., Cheong, J.L., Spittle, A.J., Doyle, L.W., Anderson, P.J., Seal, M.L., Thompson, D.K., 2020. Parcellation of the neonatal cortex using surface-based melbourne children's regional infant brain atlases (M-CRIB-S). *Sci. Rep.* 10 (1), 4359.
- Alexander, L.M., Escalera, J., Ai, L., Andreotti, C., Febre, K., Mangone, A., Vega-Potler, N., Langer, N., Alexander, A., Kovacs, M., et al., 2017. An open resource for transdiagnostic research in pediatric mental health and learning disorders. *Sci. Data* 4 (1), 1–26.
- Bateson, M., Kervadec, H., Dolz, J., Lombaert, H., Ayed, I.B., 2022. Source-free domain adaptation for image segmentation. *Med. Image Anal.* 82, 102617.
- Beer, J.C., et al., 2020. Longitudinal ComBat: A method for harmonizing longitudinal multi-scanner imaging data. *Neuroimage* 220, 117129.
- Beizae, F., Desrosiers, C., Lodygensky, G.A., Dolz, J., 2023. Harmonizing flows: Unsupervised MR harmonization based on normalizing flows. In: *International Conference on Information Processing in Medical Imaging*. Springer, pp. 347–359.
- Boudiaf, M., et al., 2022. Parameter-free online test-time adaptation. In: *CVPR*. pp. 8344–8353.
- Cackowski, S., Barbier, E.L., Dojat, M., Christen, T., 2023. Imunity: a generalizable VAE-GAN solution for multicenter MR image harmonization. *Med. Image Anal.* 88, 102799.
- Chen, C., Dou, Q., Chen, H., Qin, J., Heng, P.A., 2020. Unsupervised bidirectional cross-modality adaptation via deeply synergistic image and feature alignment for medical image segmentation. *IEEE Trans. Med. Imaging* 39 (7), 2494–2505.
- Cheong, J., Cameron, K.L.L., Thompson, D., Anderson, P.J., Ranganathan, S., Clark, R., Mentiplay, B., Burnett, A., Lee, K., Doyle, L.W., et al., 2021. Impact of moderate and late preterm birth on neurodevelopment, brain development and respiratory health at school age: protocol for a longitudinal cohort study (LaPrem study). *BMJ Open* 11 (1), e044491.
- Cohen, J.P., Luck, M., Honari, S., 2018. Distribution matching losses can hallucinate features in medical image translation. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part I*. Springer, pp. 529–536.
- Delisle, P.-L., et al., 2021. Realistic image normalization for multi-domain segmentation. *Med. Image Anal.* 74, 102191.
- Dewey, B.E., et al., 2019. DeepHarmony: A deep learning approach to contrast harmonization across scanner changes. *Magn. Reson. Imaging* 64, 160–170.
- Dewey, B.E., et al., 2020. A disentangled latent space for cross-site MRI harmonization. In: *MICCAI*. pp. 720–729.
- Di Martino, A., et al., 2014. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Mol. Psychiatry* 19 (6), 659–667.
- Dinh, L., et al., 2017. Density estimation using real NVP. In: *ICLR*.
- Dinsdale, N.K., et al., 2021. Deep learning-based unlearning of dataset bias for MRI harmonisation and confound removal. *NeuroImage* 228, 117689.
- Dolz, J., Desrosiers, C., Ayed, I.B., 2018. 3D fully convolutional networks for subcortical segmentation in MRI: A large-scale study. *NeuroImage* 170, 456–470.
- Dou, Q., Ouyang, C., Chen, C., Chen, H., Glocker, B., Zhuang, X., Heng, P.-A., 2019. Pnp-adanet: Plug-and-play adversarial domain adaptation network at unpaired cross-modality cardiac segmentation. *IEEE Access* 7, 99065–99076.
- Durrer, A., Wolleb, J., Bieder, F., Sinnecker, T., Weigel, M., Sandkühler, R., Granziera, C., Yaldizli, Ö., Cattin, P.C., 2023. Diffusion models for contrast harmonization of magnetic resonance images. *arXiv preprint arXiv:2303.08189*.
- Fischl, B., 2012. FreeSurfer. *Neuroimage* 62 (2), 774–781.
- Fortin, J.-P., et al., 2016. Removing inter-subject technical variability in magnetic resonance imaging studies. *NeuroImage* 132, 198–212.
- Fortin, J.-P., et al., 2017. Harmonization of multi-site diffusion tensor imaging data. *Neuroimage* 161, 149–170.
- Friedman, M., 1937. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *J. Amer. Statist. Assoc.* 32 (200), 675–701.
- Friedman, M., 1940. A comparison of alternative tests of significance for the problem of m rankings. *Ann. Math. Stat.* 11 (1), 86–92.
- Goodfellow, I., et al., 2020. Generative adversarial networks. *Commun. ACM* 63 (11), 139–144.
- Grover, A., et al., 2020. Alignflow: Cycle consistent learning from multiple domains via normalizing flows. In: *AAAL*. pp. 4028–4035.
- Guan, H., Liu, S., Lin, W., Yap, P.-T., Liu, M., 2022. Fast image-level MRI harmonization via spectrum analysis. In: *International Workshop on Machine Learning in Medical Imaging*. Springer, pp. 201–209.
- Godovskiy, D., et al., 2022. Cflow-ad: Real-time unsupervised anomaly detection with localization via conditional normalizing flows. In: *WACV*. pp. 98–107.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 770–778.
- Ho, J., et al., 2019. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In: *ICML*. pp. 2722–2730.
- Hu, M., Song, T., Gu, Y., Luo, X., Chen, J., Chen, Y., Zhang, Y., Zhang, S., 2021. Fully test-time adaptation for image segmentation. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part III* 24. Springer, pp. 251–260.
- Isensee, F., et al., 2021. Nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* 18 (2), 203–211.
- Jeong, H., Byun, H., Kang, D.U., Lee, J., 2023. BlindHarmony: "blind" harmonization for MR images via flow model. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 21129–21139.
- Kamnitsas, K., Baumgartner, C., Ledig, C., Newcombe, V., Simpson, J., Kane, A., Menon, D., Nori, A., Criminisi, A., Rueckert, D., et al., 2017. Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. In: *Information Processing in Medical Imaging: 25th International Conference, IPMI 2017, Boone, NC, USA, June 25–30, 2017, Proceedings* 25. Springer, pp. 597–609.

- Kantorovich, L.V., 1960. Mathematical methods of organizing and planning production. *Manage. Sci.* 6 (4), 366–422.
- Karani, N., et al., 2021. Test-time adaptable neural networks for robust medical image segmentation. *Med. Image Anal.* 68, 101907.
- Kingma, D.P., Dhariwal, P., 2018. Glow: Generative flow with invertible 1x1 convolutions. *NeurIPS* 31.
- Kingma, D.P., Welling, M., 2014. Auto-encoding variational Bayes. *Stat.* 1050, 1.
- Kirichenko, P., Izmailov, P., Wilson, A.G., 2020. Why normalizing flows fail to detect out-of-distribution data. *NeurIPS* 33, 20578–20589.
- Liang, J., et al., 2020. Do we really need to access the source data? Source hypothesis transfer for unsupervised domain adaptation. In: *ICML*. pp. 6028–6039.
- Liu, M., et al., 2021. Style transfer using generative adversarial networks for multi-site MRI harmonization. In: *MICCAI*. pp. 313–322.
- Lucas, J., Tucker, G., Grosse, R., Norouzi, M., 2019. Understanding posterior collapse in generative latent variable models. In: *DGS@ICLR*.
- Makropoulos, A., Robinson, E.C., Schuh, A., Wright, R., Fitzgibbon, S., Bozek, J., Counsell, S.J., Steinweg, J., Vecchiato, K., Passerat-Palmbach, J., et al., 2018. The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. *Neuroimage* 173, 88–112.
- Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L., 2007. Open access series of imaging studies (OASIS): cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *J. Cogn. Neurosci.* 19 (9), 1498–1507.
- Modanwal, G., et al., 2020. MRI image harmonization using cycle-consistent generative adversarial network. In: *SPIE Medical Imaging 2020*. Vol. 11314, pp. 259–264.
- Mumtazi, C.K., et al., 2022. Test-time adaptation to distribution shift by confidence maximization and input transformation. In: *ICLR*.
- Murugesan, B., Vasudeva, S.A., Liu, B., Lombaert, H., Ayed, I.B., Dolz, J., 2024. Neighbor-aware calibration of segmentation networks with penalty-based constraints. *arXiv preprint arXiv:2401.14487*.
- Niu, S., Wu, J., Zhang, Y., Chen, Y., Zheng, S., Zhao, P., Tan, M., 2022. Efficient test-time model adaptation without forgetting. In: *International Conference on Machine Learning*. PMLR, pp. 16888–16905.
- Niu, S., Wu, J., Zhang, Y., Wen, Z., Chen, Y., Zhao, P., Tan, M., 2023. Towards stable test-time adaptation in dynamic wild world. *arXiv preprint arXiv:2302.12400*.
- Nyúl, L.G., Udupa, J.K., Zhang, X., 2000. New variants of a method of MRI scale standardization. *IEEE Trans. Med. Imaging* 19 (2), 143–150.
- Oswiecki, D., Hakim, G.A.V., Noori, M., Cheraghalikhani, M., Ben Ayed, I., Desrosiers, C., 2023. Tttflow: Unsupervised test-time training with normalizing flow. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 2126–2134.
- Parida, A., Jiang, Z., Packer, R.J., Avery, R.A., Anwar, S.M., Linguraru, M.G., 2024. Quantitative metrics for benchmarking medical image harmonization. In: *IEEE International Symposium on Biomedical Imaging*. ISBI.
- Pomponio, R., et al., 2020. Harmonization of large MRI datasets for the analysis of brain imaging patterns throughout the lifespan. *NeuroImage* 208, 116450.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., 2016. Improved techniques for training gans. *NeurIPS* 29.
- Shinohara, R., et al., 2014. Statistical normalization techniques for magnetic resonance imaging. *NeuroImage: Clin.* 6, 9–19.
- Spittle, A.J., Thompson, D.K., Brown, N.C., Treyvaud, K., Cheong, J.L., Lee, K.J., Pace, C.C., Olsen, J., Allinson, L.G., Morgan, A.T., et al., 2014. Neurobehaviour between birth and 40 weeks' gestation in infants born < 30 weeks' gestation and parental psychological wellbeing: predictors of brain development and child outcomes. *BMC Pediatr.* 14, 1–13.
- Stan, S., Rostami, M., 2024. Unsupervised model adaptation for source-free segmentation of medical images. *Med. Image Anal.* 95, 103179.
- Takao, H., et al., 2011. Effect of scanner in longitudinal studies of brain volume changes. *J. Magn. Reson. Imaging* 34 (2), 438–444.
- Torbati, M.E., Tudorascu, D.L., Minhas, D.S., Maillard, P., DeCarli, C.S., Hwang, S.J., 2021. Multi-scanner harmonization of paired neuroimaging data via structure preserving embedding learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3284–3293.
- Usman, B., et al., 2020. Log-likelihood ratio minimizing flows: Towards robust and quantifiable neural distribution alignment. *NeurIPS* 33, 21118–21129.
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* 13 (4), 600–612.
- Wang, Y., Chen, H., Fan, Y., Sun, W., Tao, R., Hou, W., Wang, R., Yang, L., Zhou, Z., Guo, L.-Z., et al., 2022. Usb: A unified semi-supervised learning benchmark for classification. *Adv. Neural Inf. Process. Syst.* 35, 3938–3961.
- Wang, P., Peng, J., Pedersoli, M., Zhou, Y., Zhang, C., Desrosiers, C., 2023. Shape-aware joint distribution alignment for cross-domain image segmentation. *IEEE Trans. Med. Imaging* 42 (8), 2338–2347.
- Wang, G., et al., 2019. Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. *Neurocomputing* 338, 34–45.
- Wang, D., et al., 2021a. TENT: Fully test-time adaptation by entropy minimization. In: *ICLR*.
- Wang, R., et al., 2021b. Harmonization with flow-based causal inference. In: *MICCAI*. pp. 181–190.
- Wu, M., Zhang, L., Yap, P.-T., Lin, W., Zhu, H., Liu, M., 2023. Structural MRI harmonization via disentangled latent energy-based style translation. In: *International Workshop on Machine Learning in Medical Imaging*. Springer, pp. 1–11.
- Wu, F., Zhuang, X., 2020. CF distance: A new domain discrepancy metric and application to explicit domain adaptation for cross-modality cardiac image segmentation. *IEEE Trans. Med. Imaging* 39 (12), 4274–4285.
- Yang, C., Guo, X., Chen, Z., Yuan, Y., 2022. Source free domain adaptation for medical image segmentation with fourier style mining. *Med. Image Anal.* 79, 102457.
- Zang, C., Wang, F., 2020. Moflow: an invertible flow model for generating molecular graphs. In: *Proceedings of the 26th ACM SIGKDD*. pp. 617–626.
- Zhang, L., et al., 2020. Generalizing deep learning for medical image segmentation to unseen domains via deep stacked transformation. In: *IEEE TMI*. pp. 2531–2540.
- Zhao, Z., Xu, K., Li, S., Zeng, Z., Guan, C., 2021. MT-UDA: Towards unsupervised cross-modality medical image segmentation with limited source labels. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 293–303.
- Zhu, J.-Y., et al., 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *ICCV*. pp. 2223–2232.
- Zuo, L., Liu, Y., Xue, Y., Dewey, B.E., Remedios, S.W., Hays, S.P., Bilgel, M., Mowry, E.M., Newsome, S.D., Calabresi, P.A., et al., 2023. HACA3: A unified approach for multi-site MR image harmonization. *Comput. Med. Imaging Graph.* 109, 102285.
- Zuo, L., et al., 2021. Information-based disentangled representation learning for unsupervised MR harmonization. In: *IPMI*. pp. 346–359.