



Nonparametric Density Flows for MRI Intensity Normalisation

Daniel C. Castro^(✉) and Ben Glocker

Biomedical Image Analysis Group, Imperial College London,
London, UK
{dc315,b.glocker}@imperial.ac.uk

Abstract. With the adoption of powerful machine learning methods in medical image analysis, it is becoming increasingly desirable to aggregate data that is acquired across multiple sites. However, the underlying assumption of many analysis techniques that corresponding tissues have consistent intensities in all images is often violated in multi-centre databases. We introduce a novel intensity normalisation scheme based on density matching, wherein the histograms are modelled as Dirichlet process Gaussian mixtures. The source mixture model is transformed to minimise its L^2 divergence towards a target model, then the voxel intensities are transported through a mass-conserving flow to maintain agreement with the moving density. In a multi-centre study with brain MRI data, we show that the proposed technique produces excellent correspondence between the matched densities and histograms. We further demonstrate that our method makes tissue intensity statistics substantially more compatible between images than a baseline affine transformation and is comparable to state-of-the-art while providing considerably smoother transformations. Finally, we validate that nonlinear intensity normalisation is a step toward effective imaging data harmonisation.

1 Introduction

Many medical image analysis methods rely on the hypothesis that corresponding anatomical structures present similar intensity profiles. Unlike computed tomography, magnetic resonance imaging does not produce scans in an absolute standard scale, in general. Even when using the same imaging protocols, there can be significant variation between different scanners. Acquisition parameters have a complex effect on the luminance of the acquired images, therefore a simple linear rescaling of intensities is usually insufficient for effective data harmonisation [5]. Therefore, a crucial factor for enabling the construction of large-scale image databases from multiple sites is accurate nonlinear intensity normalisation.

A number of different approaches have been introduced for this task (cf. [1]), the most widely-adopted of which is that of Nyúl et al. [7]. The authors proposed to normalise intensities by matching a set of histogram quantiles, using these as landmarks for a piecewise linear transformation. Despite its apparent simplicity, it has proven very effective in clinical applications [9].




Fig. 1. Comparison of two MRI scans, before and after the proposed NDFLOW normalisation. *Right:* histograms (shaded) and fitted mixture models (*dotted:* likelihood, *solid:* mixture components).

Our proposed method, nonparametric density flows (NDFLOW), is perhaps conceptually closest to [5], which involves matching Gaussian mixture models (GMMs) fitted to a pair of image histograms. The author used a finite mixture to represent a pre-defined set of five tissues classes, whereas we propose to use nonparametric mixtures, focusing on accurately modelling the density rather than discriminating tissue types, and sidestepping the problem of pre-selecting the number of components. A further difference to our work is that, instead of polynomially interpolating between the means of corresponding components, we build a smooth transformation model based on density flows.

2 Method

We begin by justifying and describing the density model used to represent the intensity distributions to be matched. We then introduce the chosen objective function with its gradients for optimisation. Finally, we present our flow-based transformation model, which deforms the data so it conforms to the matched density model. Note that we focus here on single-modality intensity normalisation, although the entire formulation below extends naturally to the multivariate case.

2.1 Intensity Model

In order to be able to match the intensity distributions of a pair of images, a suitable probability density model is required. Typically, finite mixture models are considered for this task [5, 8]. However, a well-known limitation of these is the requirement to specify a priori a fixed number of components, which may in addition call for an iterative model selection loop (e.g. [8]).

On the opposite end of the spectrum, another approach is to use kernel density estimation, which is widespread for shape registration (e.g. [4, 6]). However, this formulation would result in an unwieldy optimisation problem, involving thousands or millions of parameters and all pairwise interactions. Furthermore,

the derived transformation would likely not be satisfactorily smooth without additional regularisation.

To overcome both issues we propose to use Dirichlet process Gaussian mixture models (DPGMMs) [3]. Instead of specifying a fixed number of components, they rely on a vague concentration parameter, which regulates the expected amount of clustering fragmentation and enables them to adapt their complexity to the data at hand. By allowing an unbounded number of components and setting a versatile prior on the mixture proportions, they appear as a parsimonious middle ground for flexibility and tractability.

We fit the DPGMMs to each image’s intensities using variational inference [2]. More specifically, we implemented an efficient weighted variant to fit a mixture directly to each 1D histogram.

2.2 Density Matching

The first step is to perform a coarse affine alignment by matching the moving density’s first and second moments to the target’s, accounting for arbitrary translation and rescaling of the values. This same affine transformation is then also applied to the data before the nonlinear warping takes place.

We quantify the disagreement between two probability density functions q and p on a probability space \mathcal{X} by means of the L^2 divergence:

$$D_{L^2}[q, p] = \frac{1}{2} \|q - p\|^2 = \frac{1}{2} \|q\|^2 + \frac{1}{2} \|p\|^2 - \langle q, p \rangle, \quad (1)$$

where $\langle q, p \rangle = \int q(x)p(x)dx$ is the L^2 inner product and $\|q\| = \sqrt{\langle q, q \rangle}$ is its induced norm. Aside from being symmetric, this quantity is positive and reaches zero iff $q \stackrel{\text{a.e.}}{=} p$. Crucially, unlike the usual Kullback–Leibler divergence, it is expressible in closed form for Gaussian mixture densities.

Let $q = \sum_k \pi_k q_k$ and $p = \sum_m \tau_m p_m$ denote two Gaussian mixtures, with components $q_k(x) = \mathcal{N}(x | \mu_k, \lambda_k^{-1})$ and $p_m(x) = \mathcal{N}(x | \nu_m, \omega_m^{-1})$. Equation (1) has tractable gradients w.r.t. the parameters of q , which we use to optimise its components’ means $\{\mu_k\}_k$ and precisions $\{\lambda_k\}_k$ (cf. extended version).¹

We have found, in practice, that it is largely unnecessary to adapt the mixing proportions, $\{\pi_k\}_k$, to get an excellent agreement between mixture densities. In fact, changing the mixture weights would require transferring samples *between* mixture components. Although surely possible, we point out that in the context of histogram matching this would imply altering their semantic value (e.g. consider a mixture of two well-separated components representing different tissue types).

2.3 Warping

After matching one GMM to another, we also need a way to transform the data modelled by that GMM so it matches the target data. To this end, we draw

¹ Available at: <http://arxiv.org/abs/1806.02613>.

inspiration from fluid mechanics and define the warping transformation, f , as the trajectories of particles under the effect of a velocity field u over time, taking the probability density q for the mechanical mass density. The key property that such flow must satisfy is *conservation of mass*: $\partial_t q + \partial_x(qu) = 0$, where $t \mapsto q^{(t)}$ is specified directly from the density matching.

Let us first consider the case of warping a single mixture component. A random variable $x \sim \mathcal{N}(\mu_k, \lambda_k^{-1})$ can be expressed via a diffeomorphic reparametrisation of a standard Gaussian, with $x = \psi_k(\epsilon) = \mu_k + \epsilon/\sqrt{\lambda_k}$ and $\epsilon \sim \mathcal{N}(0, 1)$. Assuming its mean and precision are changing with rates $\dot{\mu}_k$ and $\dot{\lambda}_k$, respectively, we can introduce a velocity field $u_k = \dot{\psi}_k \circ \psi_k^{-1}$ for its samples so that they agree with this evolving density. The instantaneous velocity at ‘time’ t is thus given by

$$u_k^{(t)}(x) = \dot{\mu}_k^{(t)} - \frac{\dot{\lambda}_k^{(t)}}{2\lambda_k^{(t)}}(x - \mu_k^{(t)}). \quad (2)$$

In the case of a mixture with constant weights $\{\pi_k\}_k$, we can construct a smooth, *mass-conserving* global velocity field u as

$$u^{(t)}(x) = \sum_k \frac{\pi_k q_k^{(t)}(x)}{q^{(t)}(x)} u_k^{(t)}(x), \quad (3)$$

which is simply a point-wise convex combination of each component’s velocity field, u_k , weighted by the corresponding posterior assignment probabilities.

Finally, the warping transformation $f^{(t)}$ is given by the solution to the following ordinary differential equation (ODE):

$$\partial_t f^{(t)}(x) = u^{(t)}(f^{(t)}(x)), \quad f^{(0)}(x) = x. \quad (4)$$

With f defined as above, we can prove that $q^{(t)}$ is indeed the density of samples from $q^{(0)}$ transformed through $f^{(t)}$, i.e. $q^{(0)} = |\partial_x f^{(t)}| q^{(t)} \circ f^{(t)}$ (cf. extended version). Crucially, the true solution to Eq. (4) is diffeomorphic by construction, and can be numerically approximated (and inverted) with arbitrary precision. In particular, we employ the classic fourth-order Runge–Kutta ODE solver (RK4).

Now assume we obtain optimal parameter values $\{\mu_k^*\}_k$ and $\{\lambda_k^*\}_k$ after matching q to p . We can then warp the data using the above approach, for example linearly interpolating the intermediate parameter values, $\mu_k^{(t)} = t\mu_k^* + (1-t)\mu_k^{(0)}$ and $\lambda_k^{(t)} = t\lambda_k^* + (1-t)\lambda_k^{(0)}$, hence setting the rates in Eq. (2) to constant values, $\dot{\mu}_k = \mu_k^* - \mu_k^{(0)}$ and $\dot{\lambda}_k = \lambda_k^* - \lambda_k^{(0)}$, and integrating Eq. (4) for $t \in [0, 1]$.

2.4 Practical Considerations

Since each medical image in a dataset can have millions of voxels, computing the posteriors and flows for every voxel individually can be too expensive for batch processing. To mitigate this issue, we can compute the end-to-end transformation on a mesh in the range of interest, which is then interpolated for the intensities

in the entire volume. In the reported experiments, we have used a uniformly-spaced mesh of 200 points, which has proven accurate enough for normalisation purposes.

Note that the transformation could also be computed on the histogram of discrete intensity values and built into a look-up table. However, this would not scale well to two or more dimensions for multi-modal intensity normalisation, whereas a mesh would not need to be very fine nor require a regular grid layout.

3 Experiments

3.1 Dataset

Our experiments were run on 581 T1-weighted MRI scans from the IXI database, collected from three imaging centres with different scanners.² Each scan was bias field-corrected using SPM12³ with default settings and rigidly registered to MNI space. SPM12 was further used to produce grey matter (GM), white matter (WM) and cerebrospinal fluid (CSF) tissue probability maps. We obtained brain masks by adding the three probability maps and thresholding at 0.5. The statistics reported below were weighted by the voxel-wise tissue probabilities to account for partial-volume effects and segmentation ambiguities.

3.2 Setup

We firstly fitted the nonparametric mixture models to the full integer-value histograms of the raw images (inside the brain masks), as described in Sect. 2.1. We set the DP’s concentration parameter to 2 and used data-driven Normal–Gamma priors for the components. As an ad-hoc post-processing step, we pruned the leftover mixture components with weights smaller than 10^{-3} . In the absence of one global reference distribution, we affinely aligned these DPGMMs and the corresponding data to zero mean and unit variance (cf. Fig. 2, middle).




Fig. 2. Population densities, colour-coded by imaging centre

After this rough alignment, global and centre-wise average densities were computed. These were then considered as histograms to which we fitted global and centre-wise reference DPGMMs.

² <http://brain-development.org/ixi-dataset/>.

³ <http://www.fil.ion.ucl.ac.uk/spm/software/spm12/>.

For normalisation, we consider two scenarios. The first is to normalise each centre’s reference distribution to the global target, then to apply this same transformation to all subjects in that centre. In the other approach, each subject’s image is individually normalised to the global target density. These scenarios reflect different practical applications where the centre-wise normalisation aims to preserve intra-centre variation, which might be desired. On the other hand, the individual normalisation aims to make all scans as similar as possible.

We compare our technique to Nyúl et al.’s prevalent quantile-based, piecewise linear histogram matching method [7], considered state-of-the-art for intensity normalisation and referred here as NYUL. We acquired the default 11 landmarks (histogram deciles and upper/lower percentiles) from the affine-aligned data for all subjects, then normalised each subject to this set of average landmarks.

3.3 Results

Histogram Fitness. Fig. 3 illustrates the results of normalisation between the pair of images in Fig. 1, which have a notable dissimilarity in the CSF region of the histograms. We observe that both our NDFLOW- and NYUL-transformed histograms present substantially lower mean absolute and root mean squared errors (MAE and RMSE) than the affine-aligned one, and our method performed best by a small margin. This is confirmed in a number of trials with other images.

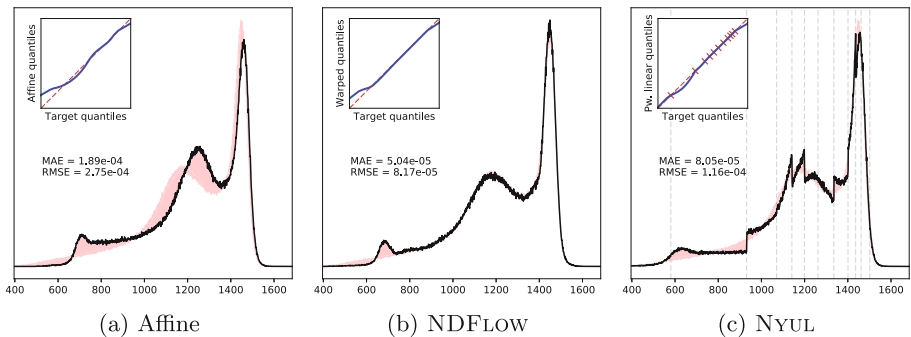


Fig. 3. Histograms and Q–Q plots of each of the methods against the target histogram. The shading shows the discrepancy between the transformed (black) and target histogram (light red). In the rightmost plot, the landmarks are indicated by vertical lines in the histogram and ticks in the Q–Q plot.

A noteworthy artefact of NYUL are abrupt jumps produced at the landmark values (e.g. Fig. 3c), which appear because intervals are uniformly compressed or dilated by different factors, and may be detrimental to downstream histogram-based tasks (e.g. mutual information registration). NDFLOW causes no such discontinuities due to the smoothness of the mass-conserving flows.

Tissue Statistics. In Table 1 we report the WM, GM and CSF intensity statistics for different normalisations. Firstly, we see that the centre-wise normalisation had a small but significant effect on the overall distribution statistics. More importantly, the variances of the statistics after individual NDFLOW and NYUL transformations were typically similar, and both were almost always substantially smaller than the variance after only affine alignment, with the exception of CSF.

Table 1. Tissue statistics after normalisation (mean \pm std. dev., $N = 581$)

	Method	1 st Quartile	Median	3 rd Quartile
WM	Affine	0.900 \pm 0.040	1.024 \pm 0.045	1.126 \pm 0.055
	NDFLOW: Centre	0.898 \pm 0.040	1.020 \pm 0.040	1.121 \pm 0.043
	NDFLOW: Indiv.	0.890 \pm 0.029	1.014 \pm 0.018	1.120 \pm 0.016
	NYUL	0.897 \pm 0.029	1.023 \pm 0.015	1.126 \pm 0.008
GM	Affine	-0.296 \pm 0.142	0.025 \pm 0.117	0.344 \pm 0.080
	NDFLOW: Centre	-0.297 \pm 0.139	0.025 \pm 0.114	0.344 \pm 0.076
	NDFLOW: Indiv.	-0.312 \pm 0.094	0.027 \pm 0.065	0.351 \pm 0.058
	NYUL	-0.309 \pm 0.106	0.027 \pm 0.070	0.350 \pm 0.064
CSF	Affine	-2.036 \pm 0.145	-1.486 \pm 0.140	-1.024 \pm 0.156
	NDFLOW: Centre	-2.035 \pm 0.143	-1.480 \pm 0.142	-1.018 \pm 0.160
	NDFLOW: Indiv.	-2.031 \pm 0.136	-1.484 \pm 0.170	-1.028 \pm 0.191
	NYUL	-2.025 \pm 0.111	-1.474 \pm 0.178	-1.029 \pm 0.207

Bold: $p < .01$, one-tailed Brown-Forsythe test for lower variance than ‘Affine’

It is known that the amount of intra-cranial fluid can vary substantially due to factors such as age and some neurodegenerative conditions, and this reflects on the distributions of intensities in brain MRI scans, which is evident in Fig. 2. As a result, normalising all subjects to a ‘mean’ distribution fails to identify a consistent reference range for CSF intensities.

A fundamental limitation of any histogram matching scheme is that it is unclear how to proceed when the distributions are *genuinely* different. Intensity distributions can be strongly affected by anatomical differences; for example, we can observe large variations in the amounts of fluid and fat in brain or whole-body scans, which may heavily skew the overall distributions (moderate example in Fig. 3). The underlying assumption of these methods (including ours) is that the distributions are similar enough up to an affine rescaling and a mild nonlinear deformation of the values, thus handling histograms of truly different shapes remains an open challenge. For images with different fields of view, it may be beneficial to perform image registration before applying intensity normalisation.

Centre Classification. To evaluate the effectiveness of intensity normalisation for data harmonisation, we conducted a centre discrimination experiment with

random forest classifiers trained on the full images. We report the pooled test results from two-fold cross validation (detailed results in extended version).

Relative to affine normalisation, centre-wise and individual NDFLOW and NYUL showed a slight drop in overall classification accuracy (94.1% vs. 92.7%, 93.6%, 92.9%, resp.). On the other hand, the uncertainty, as measured by the entropy of the predictions, was significantly higher (paired t -test, all $p < .01$). Nonlinear intensity normalisation therefore seems to successfully remove some of the biasing factors which are discriminative of the origin of the images.

4 Conclusion

In this paper, we have introduced a novel method for MRI intensity normalisation, called *nonparametric density flows* (NDFLOW). It is based on fitting and matching Dirichlet process Gaussian mixture densities, by minimising their L^2 divergence, and on mass-conserving flows, which ensure that the empirical intensity distribution agrees with the matched density model.

We demonstrated that our normalisation approach makes tissue intensity statistics significantly more consistent across subjects than a simple affine alignment, and compares favourably to the state-of-the-art method of Nyúl et al. [7]. We have additionally verified that NDFLOW is able to accurately match histograms without introducing spurious artefacts produced by the competing method. Finally, we argued that both normalisation techniques can reduce some discriminative scanner biases, in a step toward effective data harmonisation.

By employing nonparametric mixture models, we are able to represent arbitrary histogram shapes with any number of modes. In addition, our formulation has the flexibility to match only part of the distributions, by freezing the parameters of some mixture components. This may be useful for ignoring lesion-related modes (e.g. multiple sclerosis hyperintensities), if the corresponding components can be identified (e.g., via anomaly detection). Evaluating this approach and its robustness against lesion load is a compelling direction for further research.

Acknowledgements. This project was supported by CAPES, Brazil (BEX 1500/2015-05), and by the European Research Council under the EU's Horizon 2020 programme (grant agreement No 757173, project MIRA, ERC-2017-STG).

References

1. Bergeest, J.P., Jäger, F.: A comparison of five methods for signal intensity standardization in MRI. In: Tolxdorff, T., Braun, J., Deserno, T.M., Horsch, A., Handels, H., Meinzer, H.P. (eds.) *Bildverarbeitung für die Medizin*, pp. 36–40. Springer, Heidelberg (2008)
2. Blei, D.M., Jordan, M.I.: Variational inference for Dirichlet process mixtures. *Bayesian Anal.* **1**(1), 121–144 (2006)
3. Ferguson, T.S.: Bayesian density estimation by mixtures of normal distributions. *Recent Adv. Stat.* **24**(1983), 287–302 (1983)

4. Hasanbelliu, E., Giraldo, L.S., Principe, J.C.: A robust point matching algorithm for non-rigid registration using the Cauchy-Schwarz divergence. In: IEEE International Workshop on Machine Learning for Signal Processing, pp. 1–6. IEEE (2011)
5. Hellier, P.: Consistent intensity correction of MR images. In: Proceedings of the 2003 International Conference on Image Processing (ICIP 2003). IEEE (2003)
6. Jian, B., Vemuri, B.C.: Robust point set registration using Gaussian mixture models. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(8), 1633–1645 (2011)
7. Nyúl, L.G., Udupa, J.K., Zhang, X.: New variants of a method of MRI scale standardization. *IEEE Trans. Med. Imaging* **19**(2), 143–150 (2000)
8. Roy, A.S., Gopinath, A., Rangarajan, A.: Deformable density matching for 3D non-rigid registration of shapes. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007. LNCS, vol. 4791, pp. 942–949. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75757-3_114
9. Shah, M., Xiao, Y., Subbanna, N., Francis, S., Arnold, D.L., Collins, D.L., Arbel, T.: Evaluating intensity normalization on MRIs of human brain with multiple sclerosis. *Med. Image Anal.* **15**(2), 267–282 (2011)