



First Order Locally Orderless Registration

Sune Darkner^(✉), José D. T. Vidarte, and François Lauze

Department of Computer Science, University of Copenhagen, Copenhagen, Denmark
{darkner, jota, francois}@di.ku.dk

Abstract. First Order Locally Orderless Registration (FLOR) is a scale-space framework for image density estimation used for defining image similarity, mainly for Image Registration. The Locally Orderless Registration framework was designed in principle to use zeroth-order information, providing image density estimates over three scales: image scale, intensity scale, and integration scale. We extend it to take first-order information into account and hint at higher-order information. We show how standard similarity measures extend into the framework. We study especially Sum of Squared Differences (SSD) and Normalized Cross-Correlation (NCC) but present the theory of how Normalised Mutual Information (NMI) can be included.

Keywords: Image registration · Locally Orderless Images · First order information

1 Introduction

Image similarity is generally based on zeroth-order information by a scalar to scalar comparison, e.g. Sum of Squared Differences (SSD), Normalised Cross-Correlation (NCC) or Mutual Information (MI) [9]. However, images have structure and they encode information that extends beyond zeroth-order, they do not look like random noise. MI and NCC do incorporate more than just pixel intensity but very weakly and indirectly. Higher-order information is seldom used with a few exceptions, notably the normalized gradient fields [5]. We aim to integrate high-order information for registration based on Locally Orderless Images (LOI) [10] and Locally Orderless Registration LOR [4].

LOI defines three fundamental scales for estimating a density from an image: the *spatial* scale, which is the “classical” scale-space one, the *intensity or information* scale, as “bin scale” and the *integration* scale, which define the localisation of the density estimates of intensity distributions. The key is to ‘marginalize over the geometry’ and leave only the correspondence of information. The locally orderless registration gives us a theoretical platform to perform this marginalization for scalar-valued images.

Locally Orderless Registration (LOR) [8] explored its application for Magnetic Resonance Diffusion-Weighted Imaging (DWI), which are images containing complicated geometries. Indeed, DWI images can be seen as functions

$\mathbf{I} : \Omega \times \mathbb{S}^2 \rightarrow \mathbb{R}$, with Ω an open subset of \mathbb{R}^3 , where \mathbb{S}^2 is seen as the space of directions (with orientation) in \mathbb{R}^3 . An extra directional scale is added before building and localizing densities.

In this work we extend the LOI and LOR [3,4,8] framework for images $\mathbf{I} : \Omega \rightarrow \mathbb{R}$, by lifting these images to images $\mathbf{I} : \Omega \times \mathbb{S}^2 \rightarrow \mathbb{R}$, where \mathbb{S}^2 this time parametrizes the local orientations of the image \mathbf{I} . This is performed through directional responses of derivatives of Gaussian. Other kernels could be used, for instance, non-symmetric ones. Lifts to second or higher-order structures can similarly be defined via higher kernel derivatives. Once the lifting has been performed, ideas similar to DWI image registration can be used. However, as opposed to the DWI case, this lifting comes already with its scale parameter. This lifting idea is not new, with especially works in the context of image smoothing and disentangling of directions ([7] and references therein). Tools and end goals in this work are different: classical Gaussian filters and image registration.

Given two images I and J , the registration problem is to find the transformation $\varphi : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that maps I onto J such that some similarity/dissimilarity $M(I \circ \varphi, J)$ is optimized. Registration is an ill-posed problem. Therefore, the deformation φ requires regularization. Typical regularizations use constraints on the family of admissible transformations e.g. diffeomorphisms. Other alternatives are to enforced local constraints by using additional smoothing (enforcing scale to the transformation). The LOI and LOR framework provides building blocks for similarity measures, and do not impose regularisers forms. We use a very simple ones here.

Organisation and Contributions. The paper is organized as follows. First we review previous work in Sect. 2 and recall the Locally Orderless Imaging and Locally Orderless Registration frameworks in Sect. 3. Our main contribution, the extension of the LOI and LOR frameworks to first order information, is presented in Sect. 4. Registration objective functions are also discussed in this section. We illustrate the effects first order extensions for SSD and NCC similarities on the quality, and convergence of the registration in Sect. 5. Finally, we summarise and discuss perspectives in Sect. 5.3.

2 Related Work

LOI was originally proposed by Koenderink and van Dorn [10] and describes the three inherent scales of images: spatial scale, intensity scale, and integration scale. This notion of images was used to describe image similarity in a variational framework [6] and formalized into a generalized framework for image registration and the image similarity measures as LOR in [4]. Some of the groundwork for LOR as well as the properties of the density estimators used for images in image registration where investigated in [3], revealing a ‘scale imbalance’ in the partial volume density estimator. The idea of marginalizing over more complex geometries than \mathbb{R}^n was proposed in [8].

The idea of using higher order information for estimation of similarity between images is not new and normalized gradient fields NGF [5] were one of the first. In [14] an extension to the LDDMM using higher order information was presented.

There are few recent implementations of registration algorithms with NGF. The most noticeable uses NGF and a Gauss-Newton optimization scheme with locally rigid constraints [12]. This work was further evaluated on pelvis CT/CBCT images [11]. A recent first-order information approach adds another metric based on gradients to the registration cost function with NGF [15]. This metric is defined as the sum of three gradients norms, i.e. the transformed moving image, the fixed image, and the difference between moving and fixed while offering a small increase in registration accuracy.

3 Background on Locally Orderless Image Information

3.1 Notations

$\Omega \subset \mathbb{R}^3$ is the spatial domain of the images we use in the sequel. A scalar image is a function $f : \Omega \rightarrow \mathbb{R}$. We assume that images can be extended out of Ω to \mathbb{R}^3 – typically by 0 – as it is necessary for convolution. Convolution of two images $I, J : \mathbb{R}^3$ is defined by $I * J(\mathbf{x}) = \int_{\mathbb{R}^3} I(\mathbf{y})J(\mathbf{x} - \mathbf{y}) d\mathbf{y}$. This actually extends to the case where one of the images is vector-valued directly. G_σ denotes a 3D isotropic Gaussian of standard deviation σ .

3.2 Lebesgue Integration and Histograms

Consider a function integrable $I : \mathbb{R}^n \rightarrow \mathbb{R}$. Its integral $\int I d\mu$ with respect to the Lebesgue measure μ of \mathbb{R}^n , denoted in the sequel as $\int_{\mathbb{R}^n} I(\mathbf{x}) d\mathbf{x}$, can be computed as the limit over all subdivisions $0 \leq i_0 < \dots < i_N$, $\sum_{n=0}^{N-1} i_n \mu(I^{-1}([i_n, i_{n+1}]))$. At the limit, when $i_{n+1} - i_n \rightarrow 0$, this can be rewritten as $\int_{\mathbb{R}} i h_I(i) di$ where $h_I(i)$ is the length of isophote $I^{-1}(i)$. The function $i \mapsto h_I(i)$ is a generalized histogram of the values of I . Many standard integrals can be rewritten using this form. For instance $\int_{\mathbb{R}^n} I(x)^2 dx = \int_{\mathbb{R}} i^2 h_I(i) di$. This generalizes to joint histograms: given two images $I, J : \mathbb{R}^n \rightarrow \mathbb{R}$, $(I, J) : \mathbb{R}^n \rightarrow \mathbb{R}^2$, $\mathbf{x} \mapsto (I(\mathbf{x}), J(\mathbf{x}))$ and its integral can be written as $\int_{\mathbb{R}^2} (i, j) h_{I,J} di dj$ where $h_{I,J}$ is the joint histogram of I and J . Classical similarities can be rewritten using histograms, for instance, Sum of Square Differences (SSD): $\int_{\mathbb{R}^n} (I(x) - J(x))^2 dx = \int_{\mathbb{R}^2} (i - j)^2 h_{I,J}(i, j) didj$. Normalised Cross-Correlation, (Normalised) Mutual information, etc. can be written in terms of image histograms and their normalisations.

3.3 LOI and LOR Framework

LOI is a way to map images into local histograms, with three inherent scales: the spatial or image scale, the intensity scale, and integration scale. The image

or spatial scale σ is used to smooth input images I and obtain $I_\sigma = I * G_\sigma$. A localised histogram over the values of I_σ is computed as

$$h_{I,\sigma\beta\alpha}(i|\mathbf{x}) := \int_{\Omega} P_\beta(I_\sigma(\mathbf{y}) - i) W_\alpha(\mathbf{y} - \mathbf{x}) d\mathbf{y} \tag{3.1}$$

where P_β is a Parzen window of scale β , which provides the *intensity* scale and $W_\alpha(x)$ is an integration window which provides the *integration* scale α . The histogram $h_{I,\sigma\beta\alpha}(\cdot|\mathbf{x})$ is defined over \mathbb{R} or at least over an interval Λ containing the range of values of I_σ . Normalising it, we obtain the image density

$$p_{I,\sigma\beta\alpha}(i|\mathbf{x}) = \frac{h_{I,\sigma\beta\alpha}(i|\mathbf{x})}{\int_{\Lambda} h_{I,\sigma\beta\alpha}(j|\mathbf{x}) dj}. \tag{3.2}$$

By letting the integration scale $\alpha \rightarrow \infty$, we obtain *global* histograms and densities $h_{I,\sigma\beta}(i) := \int_{\Omega} P_\beta(I_\sigma(\mathbf{x}) - i) dx$ and $p_{I,\sigma\beta}(i)$. This will be the case in this paper. This construction extends to the definition of joint histograms and densities, at the heart of Locally Orderless Registration by

$$h_{I,J,\sigma\beta\alpha}(i,j|\mathbf{x}) := \int_{\Omega} P_\beta(I_\sigma(\mathbf{y}) - i) P_\beta(J_\sigma(\mathbf{y}) - j) W_\alpha(\mathbf{y} - \mathbf{x}) d\mathbf{y} \tag{3.3}$$

$$p_{I,J,\sigma\beta\alpha}(i,j|\mathbf{x}) = \frac{h_{I,J,\sigma\beta\alpha}(i,j|\mathbf{x})}{\int_{\Lambda} h_{I,J,\sigma\beta\alpha}(u,v|\mathbf{x}) du dv} \tag{3.4}$$

and similar formulas in the global case. Single histograms and densities can also be obtained from them by marginalisation. LOR Image similarities are defined through single and joint density estimates Eq. (3.2) and Eq. (3.4). Similarity measures are defined as

$$M_L(I, J) = \int_{\Omega} \int_{\Lambda^2} f(i, j, p_{I,J,\sigma\beta\alpha}(i, j|\mathbf{x})) di dj d\mathbf{x}, \tag{3.5}$$

$$M_G(I, J) = \int_{\Lambda^2} f(i, j, p_{I,J,\sigma\beta}(i, j)) di dj \tag{3.6}$$

with M_L built from localised densities and M_G from global ones. Among them, p -linear ones are characterized by $f(i, j, p) = g(i, j)p$, while nonlinear ones take more complex forms. We already mentioned in the previous section how SSD can be simply written using joint histograms. By normalising it, it can be written via densities (3.4). Another classical similarity, normalised cross-correlation (NCC), can also easily be written in term of histograms and densities.

$$NCC(I, J) = \frac{\langle I - \bar{I}, J - \bar{J} \rangle}{\|I - \bar{I}\| \|J - \bar{J}\|}$$

where \bar{I} and \bar{J} are the average values of I and J on Ω and the inner product and norms are L^2 ones. The inner product $\langle I, J \rangle$ is $\int_{\mathbb{R}^2} ij h_{I,J}(i, j) didj$. Replacing $h_{I,J}$ by $h_{I,J;\sigma\beta\alpha}$ provides its LOI counterpart expression. The average \bar{I} is $\int_{\mathbb{R}} ih_I(i) di / (\int_{\mathbb{R}} h(i) di)$. Again, we replace h_I by $h_{I,\sigma\beta\alpha}$ to obtain its LOI counterpart expression.

To use it in registration, the setting is typically the following. One chooses a hold on domain $D \subset \mathbb{R}^3$ large enough, with $\Omega \subset D$ and mappings $\varphi : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, with $\varphi \equiv \text{id}_3$ out of D , where id_3 is the identity transform. Here, we assume $D = \Omega$. These transformations are usually of class \mathcal{C}^k , $k \geq 1$, often more. They are often, but not always, constrained to be diffeomorphic. A goodness of fit functional is obtained by evaluation the (dis)similarity $\varphi \mapsto M(I \circ \varphi, J)$.

4 Extension of LOI and LOR to Higher Information

In this section, we introduce a straightforward way to extend the LOI to incorporate higher order image information in histogram and density formulations. We focus on first order, as higher order may be limited in practice because of the complexity and resulting memory footprint.

4.1 First Order Locally Orderless Registration (FLOR)

In this paper we probe and use first order differential information of an image $I : \mathbb{R}^3 \rightarrow \mathbb{R}$ (with effective spatial domain Ω). It is obtained by lifting it to image $\mathbf{I}_\sigma : \mathbb{R}^3 \times \mathbb{S}^2 \rightarrow \mathbb{R}$ which encodes gradient responses at different directions in a straightforward way. The differential $d_{\mathbf{x}}G_\sigma : \mathbb{R}^3 \rightarrow \mathbb{R}$ is, for each \mathbf{x} linear, it is enough to know it on $\mathbb{S}^2 \subset \mathbb{R}^3$.

$$\mathbf{I}_\sigma(\mathbf{x}, \mathbf{v}) = \left(\int_{\mathbb{R}^3} I(\mathbf{y}) d_{(\mathbf{y}-\mathbf{x})} G_\sigma d\mathbf{y} \right) \mathbf{v} = d_{\mathbf{x}} (I * G_\sigma) \mathbf{v} \quad (4.1)$$

This can of course be rewritten as $\mathbf{I}_\sigma(\mathbf{x}, \mathbf{v}) = \nabla I_\sigma(\mathbf{x})^T \mathbf{v}$. Note that $\mathbf{I}_\sigma(\mathbf{x}, -\mathbf{v}) = -\mathbf{I}_\sigma(\mathbf{x}, \mathbf{v})$ due to our lifting choice. Using a higher order operator, such as, for instance, the Hessian of Gaussian $\nabla^2 G_\sigma$ would allow us to probe second order structure as a $\tilde{\mathbf{I}}_\sigma(\mathbf{x}, \mathbf{v}) = \text{Hess } I_\sigma(\mathbf{x})(\mathbf{v}, \mathbf{v})$.

Once the lifting is performed, we can now define local histograms and densities. They are spatially localised, not directionally.

$$h_{\mathbf{I};\sigma\beta\alpha}(i|\mathbf{x}) = \int_{\mathbb{R}^3 \times \mathbb{S}^2} P_\beta(\mathbf{I}_\sigma(\mathbf{y}, \mathbf{v}) - i) W_\alpha(\mathbf{x} - \mathbf{y}) d\mathbf{v} d\mathbf{y} \quad (4.2)$$

$$p_{\mathbf{I};\sigma\beta\alpha}(i|\mathbf{x}) = \frac{h_{\mathbf{I};\sigma\beta\alpha}(i|\mathbf{x})}{\int_A h_{\mathbf{I};\sigma\beta\alpha}(j|\mathbf{x}) dj} \quad (4.3)$$

where this time A is an interval containing the range of \mathbf{I}_σ . As in the zeroth order case, global histograms and densities can be obtained by letting $\alpha \rightarrow \infty$. Given two images $I, J : \mathbb{R}^3 \rightarrow \mathbb{R}$, we can lift them to \mathbf{I}_σ and \mathbf{J}_σ and define joint histograms and densities

$$h_{\mathbf{I},\mathbf{J};\sigma\beta\alpha}(i, j|\mathbf{x}) = \int_{\mathbb{R}^3 \times \mathbb{S}^2} P_\beta(\mathbf{I}_\sigma(\mathbf{y}, \mathbf{v}) - i) P_\beta(\mathbf{J}_\sigma(\mathbf{y}, \mathbf{v}) - j) W_\alpha(\mathbf{x} - \mathbf{y}) d\mathbf{v} d\mathbf{y} \quad (4.4)$$

$$p_{\mathbf{I},\mathbf{J};\sigma\beta\alpha}(i, j|\mathbf{x}) = \frac{h_{\mathbf{I},\mathbf{J};\sigma\beta\alpha}(i, j|\mathbf{x})}{\int_{A^2} h_{\mathbf{I},\mathbf{J};\sigma\beta\alpha}(u, v|\mathbf{x}) du dv} \quad (4.5)$$

Here again, by letting $\alpha \rightarrow 0$, we obtain global histograms and densities.

4.2 First Order Deformation Model

Let $\varphi : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ a deformation. By the chain rule, $\frac{d}{dt} \Big|_{t=0} I_\sigma(\varphi(\mathbf{x} + t\mathbf{v})) = \nabla I_\sigma(\varphi(\mathbf{x}))^T J_x \varphi(\mathbf{v})$, with $J_x \varphi$ the Jacobian of φ . This implies of course that φ acts on the first order information via its differential. Here comes the problem that, as we have limited the directional probing space to \mathbb{S}^2 , there is no guarantee that $J_x \varphi(\mathbf{v}) \in \mathbb{S}^2$, let alone non zero. This is however the case if we restrict φ to be a diffeomorphism, and this is what we assume from now. From its very definition, the mapping of directions at $\mathbf{x} \in \Omega$ is given by

$$\psi_x : \mathbb{S}^2 \rightarrow \mathbb{S}^2, \quad v \mapsto \frac{J_x \varphi(\mathbf{v})}{|J_x \varphi(\mathbf{v})|}, \quad (4.6)$$

This lead to define the action of φ^1 on the lifted image $\mathbf{I}_\sigma(\mathbf{x}, \mathbf{v})$ as

$$(\varphi.\mathbf{I}_\sigma)(\mathbf{x}, \mathbf{v}) = |J_x \varphi(\mathbf{v})| \mathbf{I}_\sigma(\varphi(\mathbf{x}), \psi_x(\mathbf{v})). \quad (4.7)$$

It clearly satisfies $(\varphi.\mathbf{I}_\sigma)(\mathbf{x}, -\mathbf{v}) = -(\varphi.\mathbf{I})(\mathbf{x}, \mathbf{v})$, thus respecting the structure of lifted images. Alternatively, one could consider another first order deformation model, where the Jacobian scaling factor is ignored, i.e.

$$(\varphi.\mathbf{I}_\sigma)(\mathbf{x}, \mathbf{v}) = \mathbf{I}_\sigma(\varphi(\mathbf{x}), \psi_x(\mathbf{v})). \quad (4.8)$$

This may apply to images of more categorical nature. This can be the case for two images showing similar anatomical structures, with same tissue density, but which cannot be registered by a (local) rigid motion.

By using either the local or global histograms and densities, higher order similarities $\mathbf{M}(\mathbf{I}, \mathbf{J})$ are obtained exactly the same way as discussed in the previous section. Finally one may combine zeroth and first order to get new similarity measures, and use them in a registration framework via

$$\varphi \mapsto M(I \circ \varphi, J) + \lambda \mathbf{M}(\varphi.\mathbf{I}_\sigma, \mathbf{J}_\sigma). \quad (4.9)$$

4.3 Registration Objectives and Deformations

The similarities used in this paper are 1) SSD for zeroth and first order information, and 2) NCC for zeroth and first order information. Free-form B-spline deformation models [13] are used, with simple control point grid motion limitation as regularisation. We also use simple translation deformations on some experiments.

4.4 Implementation

The implementation has been made in PyTorch 1.7.1 and the basis consists of a Cubic B-spline from which both the image interpolation and deformation field can be estimated. Analytical Jacobians of both image and deformation has been

¹ Note that this is not stricto sensu a group action here.

implemented which allow us to use the backpropagation of PyTorch and optimizer for finding the solution. The action of the Jacobian on the directional derivative have 2 implementations, given by Eq. (4.7) and Eq. (4.8). The implementations ensures that all scales are consistent and to change image scale we simply blur the images prior to registration with the desired kernel. Objectives are optimised using PyTorch Adam implementation. The code runs both on CPU and GPU; a full 3D registration takes 2–3 minutes on a laptop and around 1 min on an RTX3090.

5 Experiments and Results

We conduct 2 main experiments. First we investigate the properties of the 1^{st} -order information compared to the 0^{th} -order using translation only. Secondly we show that we can perform 3D non-rigid registration with convincing results. We have used two 3D T1 weighted magnetic resonance images (MRI) from two separate individuals for our proof of principle. The images are shown in Fig. 1.

5.1 The Similarity Properties

To illustrate the effects of including higher-order (1^{st} -order) information in the similarity-measure, we map the 0^{th} -order and 1^{st} -order information as a function of translation in 2D (x, y) -plane.

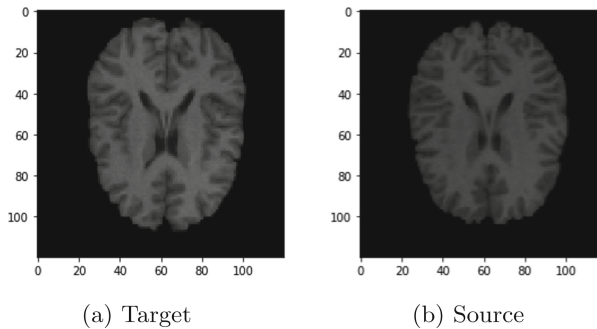


Fig. 1. A Slice of the target and source image used for our experiments

Our first experiment shows how the information from the images using SSD and NCC respectively appears in the simple case where the deformation φ is a pure translation in 2D, for an MRI compared with itself. As Fig. 2 illustrates, the similarity in the 1^{st} -order information has a significantly steeper slope close to the optimum, compared to that of the 0^{th} -order information for both SSD and NCC. This indicates that including 1^{st} -order information may improve registration close to the optimum. However, when comparing 2 different images in Fig. 3 we observe

that multiple minima exist with 1st-order only, and that 0th-order has a better and a wider basin of attraction. Furthermore NCC seems to be more suitable compared to SSD. Therefore a combination of 0th-order, 1st-order information and NCC seems more appropriate for image registration applications.

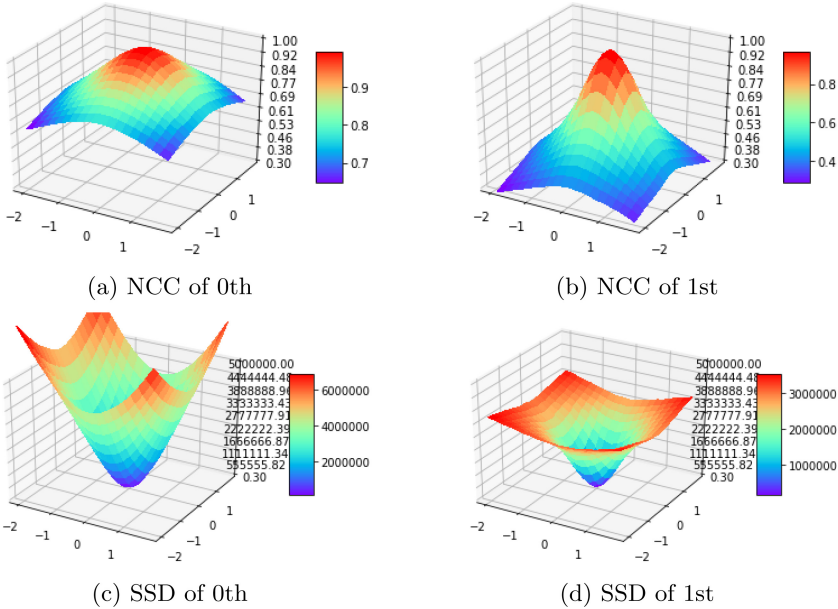


Fig. 2. The self similarity of the 0th and 1st order image information for NCC and SSD respectively under translation in 2D around the identity.

5.2 Non-rigid Registration

We perform 3 non-rigid registrations of the source and the target using only 1st-order information and using only 0th-order information and a combination of both respectively. We used a free-form deformation cubic B-spline [13] with 5 voxel spacing between the knots and evaluation points for every second voxel. We discretized the 1st-order information with 26 normalized directions, pointing to each neighbouring voxel in a 3 × 3 × 3 local grid. We weighted 0th-order and 1st-order terms by the ratio between the number evaluation of the 0th-order and the 1st-order ($\frac{1}{26}$). As can be seen from the convergence plots (Fig. 4), this ratio will align both gradient information and intensity information, in contrast to optimizing only the 0th-order or the 1st-order.

The final registration results are shown in Fig. 5. We have used both formulations from Eq. (4.7), with Jacobian normalisation, and Eq. (4.8), without Jacobian normalisation. As Fig. 5 shows, the results are quite convincing, and the difference between the two is very small. However, in this MRI registration case, the version not including scaling seems to be more suitable compared to SSD.

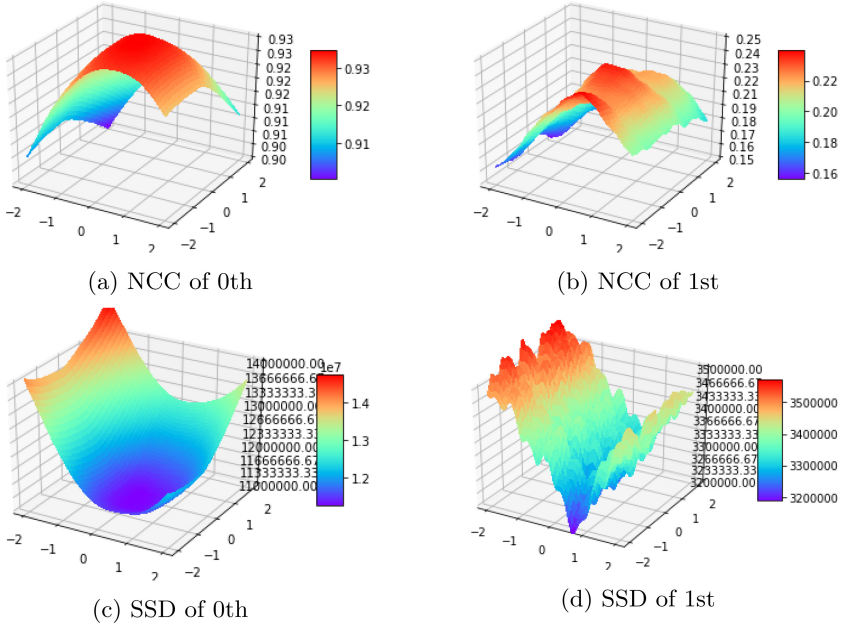


Fig. 3. The similarity of the 0^{th} - and 1^{st} -order image information for NCC and SSD respectively under translation in 2D around the identity for the source and target image. Clearly, multiple local minima exist in the 1^{st} -order information, in contrast to 0^{th} -order that only has one.

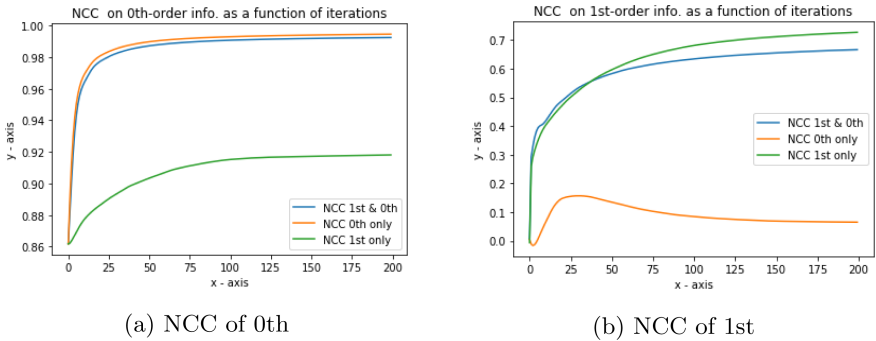


Fig. 4. Convergence plot of NCC for 0^{th} - and 1^{st} -order similarities, separately, as function of iterations. The experiment was performed optimizing only for 0^{th} -order information, only 1^{st} -order information and both with weight between 0^{th} - and 1^{st} -order terms as the ratio of the number of evaluation-orientations ($\frac{1}{26}$). Note how the 1st-order appears to also maximize the 0^{th} -order and how 1^{st} -order fails to maximize the gradient information.

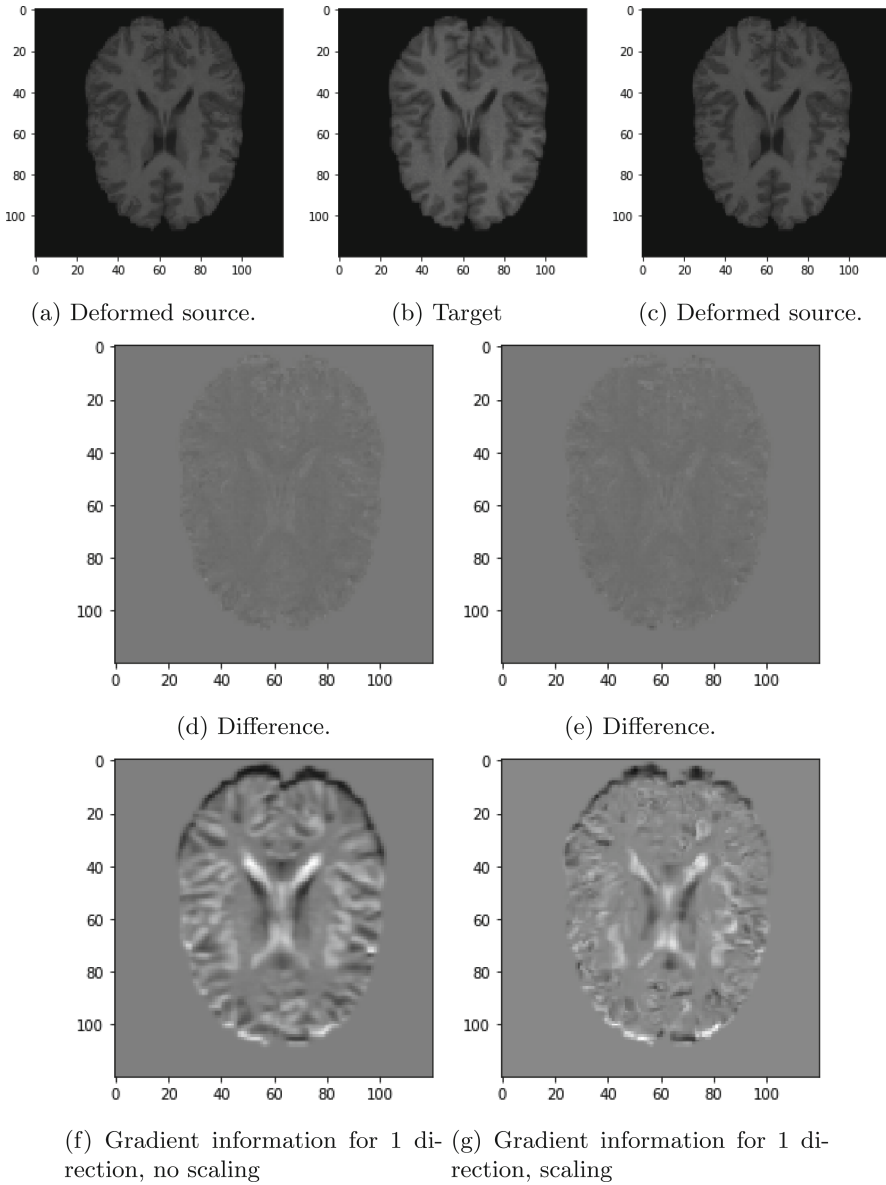


Fig. 5. (a) and (c) are the result of the registration to (b), where (a) has been registered with Eq. (4.7) and (c) with Eq. (4.8). (d) and (e) are the difference images of (a) and (b), and (c) and (b) respectively. (f) and (g) are examples of the 1st order information in (c) matched to (b) in the frame of (c)

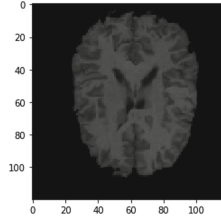


Fig. 6. An example of the deformation from pure 1st-order information.

5.3 Discussion and Limitations

The experiments presented in the previous section illustrate the gain obtained by including 1st-order information, but also some limitations. First, it is clear from our experiments that 1st-order information itself is not sufficient to make a proper registration in our setup. Figure 6 shows it clearly, where the registration of the image is inferior and suffers from undesirable deformations. Another point is that experiment here only serve as a proof of principle, and future work will include a thorough comparison over multiple large data sets similar to [9]. Furthermore, we will extend the work to use diffeomorphisms like previous methods, such as: Symmetric Normalization [1] or the Collocation for Diffeomorphic Deformations [2]. Finally, we will include an evaluation of information theory metrics and the exploration of all the scales presented in the formulation.

6 Conclusion

We introduced a framework for including higher-order information into image similarity and illustrated the application in image registration. We have shown that the method is able to match both 0th- and 1st-order information using SSD and NCC. The framework allows us to use all admissible measures from the LOR framework. We have shown that the framework is able to deliver high-quality non-rigid registration and that it has the potential to improve the accuracy of image registration in general.

Acknowledgements. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 764644. This paper only contains the author’s views and the Research Executive Agency and the Commission are not responsible for any use that may be made of the information it.



References

1. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* **12**(1), 26–41 (2008)
2. Darkner, S., Pai, A., Liptrot, M., Sporring, J.: Collocation for diffeomorphic deformations in medical image registration. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**(7), 1570–1583 (2018). <https://doi.org/10.1109/TPAMI.2017.2730205>
3. Darkner, S., Sporring, J.: Generalized partial volume: an inferior density estimator to Parzen windows for normalized mutual information. In: Székely, G., Hahn, H.K. (eds.) *IPMI 2011*. LNCS, vol. 6801, pp. 436–447. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22092-0_36
4. Darkner, S., Sporring, J.: Locally orderless registration. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(6), 1437–1450 (2013)
5. Haber, E., Modersitzki, J.: Intensity gradient based registration and fusion of multi-modal images. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) *MICCAI 2006*. LNCS, vol. 4191, pp. 726–733. Springer, Heidelberg (2006). https://doi.org/10.1007/11866763_89
6. Hermosillo, G., Chef d’Hotel, C., Faugeras, O.: Variational methods for multimodal image matching. *Int. J. Comput. Vision* **50**(3), 329–343 (2002)
7. Janssen, M.H., Janssen, A.J., Bekkers, E.J., Bescos, J.O., Duits, R.: Design and processing of 3D invertible orientation scores of 3D images. *J. Math. Imaging Vis.* **60**, 1427–1458 (2018)
8. Jensen, H.G., Lauze, F., Nielsen, M., Darkner, S.: Locally orderless registration for diffusion weighted images. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015*. LNCS, vol. 9350, pp. 305–312. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24571-3_37
9. Klein, A., et al.: Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *Neuroimage* **46**(3), 786–802 (2009)
10. Koenderink, J.J., Van Doorn, A.J.: The structure of locally orderless images. *Int. J. Comput. Vision* **31**(2), 159–168 (1999)
11. König, L., Derksen, A., Papenberg, N., Haas, B.: Deformable image registration for adaptive radiotherapy with guaranteed local rigidity constraints. *Radiat. Oncol.* **11**(1), 1–9 (2016)
12. König, L., Rühaak, J.: A fast and accurate parallel algorithm for non-linear image registration using normalized gradient fields. In: 2014 IEEE 11th International Symposium on Biomedical Imaging (ISBI), pp. 580–583. IEEE (2014)
13. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J.: Non-rigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Med. Imaging* **18**(8), 712–721 (1999)
14. Sommer, S., Nielsen, M., Darkner, S., Pennec, X.: Higher-order momentum distributions and locally affine LDDMM registration. *SIAM J. Imag. Sci.* **6**(1), 341–367 (2013)
15. Theljani, A., Chen, K.: An augmented lagrangian method for solving a new variational model based on gradients similarity measures and high order regularization for multimodality registration. *Inverse Probl. Imaging* **13**(2), 309–335 (2019)