

On Removing Interpolation and Resampling Artifacts in Rigid Image Registration

Iman Aganj*, B.T. Thomas Yeo, Mert R. Sabuncu, and Bruce Fischl

Abstract— We show that image registration using conventional interpolation and summation approximations of continuous integrals can generally fail due to resampling artifacts. These artifacts negatively affect the accuracy of registration by producing local optima, altering the gradient, shifting the global optimum, and making rigid registration asymmetric. In this work, after an extensive literature review, we demonstrate the causes of the artifacts by comparing inclusion and avoidance of resampling analytically. We show the sum-of-squared-differences cost function formulated as an integral to be more accurate compared to its traditional sum form in a simple case of image registration. We then discuss aliasing that occurs in rotation, which is due to the fact that an image represented in the Cartesian grid is sampled with different rates in different directions, and propose the use of oscillatory isotropic interpolation kernels that allow better recovery of true global optima by overcoming this type of aliasing. Through our experiments on brain, fingerprint, and white noise images, we illustrate the superior performance of the integral registration cost function in both Cartesian and spherical coordinates, and

also validate the introduced radial interpolation kernel by demonstrating the improvement in registration.

Index Terms— Image registration, interpolation artifacts, aliasing, image resampling. (EDICS: TEC-ISR)

I. INTRODUCTION

Image resampling is indispensable in many image processing applications. Given that digital images are represented as discrete values on a regular grid, virtually any image transformation, such as translation, rotation, zooming, or nonlinear warping, requires mapping of the grid points (pixels) to a new set of coordinate points that do not generally correspond to the original grid. To present the transformed image on the same regular grid, it is thus necessary to estimate, i.e. interpolate, the image values between the grid points where no sampled data are available.

Accurately speaking, *interpolation* is fitting a continuous function passing through the image pixels or voxels (see [1-3] for reviews). Interpolation is required in advance of *resampling* this continuous function on the desired points [4]. In practice, these two processes are commonly combined, sometimes simply referred to as resampling. It has been known that both interpolation and resampling can produce artifacts in the resulting image (see Sec. II for a literature review). *Interpolation* with non-ideal kernels not only blurs the image, but also creates artificial high-frequency components. In addition, aliasing distortions can appear during *resampling* of such a non-ideally interpolated image (and even an ideally interpolated one, if rotation is involved).

In addition to helping to represent transformed images on a regular grid, interpolation and resampling are used in other procedures, such as image alignment or *registration*, as discussed in this paper. Registration provides a transformation maximizing a similarity measure between the transformed version of an image and a second reference image (see [5-8] for reviews). Interpolation and resampling are commonly used in the transformation step of the sub-pixel registration, and in turn produce the aforementioned artifacts – notably patterns of local extrema in the similarity measure – making the optimization challenging. The artifacts may in some cases cause the registration to fail completely, for example by lowering the signal-to-noise ratio of the image gradient leading the algorithm to wrong optima, by shifting the global optimum significantly, or by making a false stronger global optimum appear in another region of the image. Additionally, resampling generally breaks the inherent symmetry (inverse-consistency [9]) of rigid image registration [10], since the

Manuscript received January 6, 2012. Support for this research was provided in part by the National Center for Research Resources (P41-RR14075, and the NCRR BIRN Morphometric Project BIRN002, U24-RR021382), the National Institute for Biomedical Imaging and Bioengineering (R01EB006758), the National Institute on Aging (AG022381), the National Center for Alternative Medicine (RC1 AT005728-01), the National Institute for Neurological Disorders and Stroke (R01 NS052585-01, 1R21NS072652-01, 1R01NS070963), and was made possible by the resources provided by Shared Instrumentation Grants 1S10RR023401, 1S10RR019307, and 1S10RR023043. Additional support was provided by The Autism & Dyslexia Project funded by the Ellison Medical Foundation, and by the NIH Blueprint for Neuroscience Research (5U01-MH093765), part of the multi-institutional Human Connectome Project. Dr. Sabuncu received support from a KL2 Medical Research Investigator Training (MeRIT) grant awarded via Harvard Catalyst, The Harvard Clinical and Translational Science Center (NIH grant #1KL2RR025757-01 and financial contributions from Harvard University and its affiliated academic health care centers), and an NIH K25 grant (NIBIB 1K25EB013649-01).

Copyright (c) 2012 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

I. Aganj (iman@nmr.mgh.harvard.edu, +1 (612) 202-3019), B.T.T. Yeo (ythomas@csail.mit.edu), M.R. Sabuncu (msabuncu@csail.mit.edu), and B. Fischl (fischl@nmr.mgh.harvard.edu) are with the Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School, Charlestown, MA 02129 USA. I. Aganj is also with the Laboratory of Information and Decision Systems, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA. In addition, M.R. Sabuncu and B. Fischl are with the Computer Science and Artificial Intelligence Laboratory, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA. B.T.T. Yeo is also with the Department of Neuroscience and Behavioral Disorders, Duke-NUS Graduate Medical School, Singapore. B. Fischl is also with the Division of Health Sciences and Technology, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

artifacts alter only one of the images. The fact that the transformed image is merely used to compute the similarity measure raises the following question: Are interpolation and resampling both necessary in registration, or would interpolation alone be sufficient? As we will see in Sec. III, the latter is true, in which case aliasing artifacts due to resampling can be avoided by using an integral in the registration cost function instead of a sum, without resorting to costlier closer-to-ideal interpolation kernels. Another question addressed in this paper is whether using ideal interpolants can eliminate resampling artifacts and make the two cases of including and avoiding the resampling stage equivalent.

In this work, we first review the literature extensively and provide a comprehensive list of references where resampling artifacts have been observed in image registration, and also point out several studies in which an integral cost function has been used to avoid resampling. Although these artifacts have been explored in a number of papers and there are instances in the literature where they are mostly prevented by avoiding resampling, to the best of our knowledge, the two cases of keeping and avoiding the resampling step have not been analytically compared and contrasted. We (a) juxtapose inclusion and avoidance of resampling by showing the resulting cost functions side by side in a simple case of registration, and compare their corresponding errors analytically. In particular, we show that the resampling errors are asymmetric, resulting in the inverse-inconsistency of registration, and that by avoiding resampling the rigid image registration is kept symmetric. Next, we recall that contrary to the popular belief that employing the ideal (*sinc*) interpolator would eliminate the artifacts in rigid registration, aliasing still occurs when rotation is involved, and show that using multiple ideal low-pass filters to avert it biases the registration. We then (b) propose to use alternative interpolation kernels corresponding to disk/ball-shaped anti-aliasing filters that preserve as much signal energy as possible without biasing the registration, while preventing aliasing. With this approach such filters are implicitly incorporated in the already existing step of interpolation in image registration. We show experimental results on several images, including examples of triangular-mesh image representation on the sphere, while taking two different analytical and stochastic implementation approaches to compute the integral similarity measure. We demonstrate the improvement in registration accuracy achieved by both eschewing resampling and employing the new interpolation kernel, and in registration symmetry in the former case.

We continue by reviewing the relevant prior work in Sec. II. Section III compares the two cases of including and avoiding the resampling step, and Sec. IV demonstrates rotation-induced aliasing while introducing an unbiased kernel to prevent it. Practical considerations are mentioned in Sec. V, and experimental results are provided in Sec. VI. Section VII concludes the paper with a few final remarks.

II. LITERATURE REVIEW

We review four categories of the related prior work here, presenting them mostly in a chronological order in each subsection.

A. General Interpolation and Resampling Artifacts

Artifacts arising from signal interpolation and resampling have been initially studied outside the context of registration [3, 4, 11, 12], and further exploited in exposing digital forgeries [13]. Ref. [4] clearly distinguishes resampling from interpolation, and is the first paper that we are aware of to mention that resampling aliases the higher frequencies of the *non-ideal interpolant* into the lower frequencies (see Sec. III), a concept which was subsequently restated by the authors of [14, 15]. We will see in Sec. III why this is the major cause of the artifacts in image registration.

B. Instances of Artifacts in Sub-Pixel Registration

Sub-pixel image registration artifacts appear mostly as scalloping patterns with the period of one pixel in the rigid registration objective function. One of the first attempts to alleviate this issue was to remove the periodic elements with such a frequency, and was applied in realignment of functional magnetic resonance imaging (fMRI) time series [16]. These artifacts were primarily noticed in the entropy-based mutual information (MI, [17-19]) methods. In [20], the authors studied the MI metric and noticed sudden changes in the metric for grid-aligning transformations. It was shown that the two commonly used linear and partial-volume [21] interpolation methods produced dissimilar patterns of artifact, and that a slight resampling of one of the images might make these patterns smoother. Furthermore, they combined gradient information in the match metric to compensate for the artifacts [22]. Random resampling and inclusion of the prior joint probability were shown in [23] to reduce such artifacts. The authors of [24] presented a generalized partial volume estimation that by properly choosing the kernel function, reduces the artifacts in MI-based registration. Various interpolators for MI were compared in [25] and several strategies to reduce the artifacts were proposed. Ref. [26] presented an estimation error cancel method which reduces the peak estimation error when the sum of squared differences (SSD) or the cross-correlation (CC) is computed discretely and then interpolated. The extent and the nature of artifactual displacements produced by non-rigid SSD-based registration techniques for different interpolators were compared in [27]. In [28], it was shown that quasi-random sampling based on Halton sequences alleviates the grid effect in MI registration. The authors of [29] observed reduced interpolation artifacts when using a full-image mutual histogram as compared to only a sub-volume one. In [30], the noise reduction filtering that occurs when image samples are interpolated was hypothesized to be the reason behind the artifacts, and a constant-variance filter for linear interpolation was presented to remove them in CC- and MI-based registration. The authors of [31] proposed the use of Euclidean distance in partial volume interpolation to improve the artifacts in MI-based registration. Refs. [7, 19, 32] further report observing these artifacts in registration.

Virtually all of the above references propose alterations in the registration algorithm that lessen the artifacts, or at least make them less noticeable by reducing their regularity. Next, we mention a rather different category in the literature, analyzing the artifacts from the *resampling* view point.

C. Avoiding the Culprit: Resampling

Formulating the registration cost function using an integral over the continuous interpolated images – instead of a sum over discrete resampled values – eliminates the resampling step (see Sec. III). The integral SSD has been employed outside the registration framework to assess interpolation, for instance when there is a jump in the original signal at a non-grid location [11], or in parameter optimization for cubic convolution interpolation [33]. Conversely, the resampling error when using the discrete SSD has been shown to depend on the lattice location [12].

As for sub-pixel registration, omitting the resampling step was first suggested yet again in MI maximization [34], where computing the joint histogram from continuous interpolated images was argued to be, in principal, better than doing so from resampled images, because it is free of sampling effects. The authors also showed that the partial volume interpolator [21] is a special case of the continuous evaluation of the MI from interpolated images, when a boxcar kernel is used. The continuous approximation of MI was performed by computing the MI discretely from oversampled images. Alternative implementations of this method were later provided via nonparametric windows [35], auto-correlated kernels [36], and numerical approximation of the entropy [28, 37].

Similarly to [12], it was demonstrated in [38] that resampling the image using finite-length interpolants, such as linear and cubic convolution, acts as a low-pass filter which affects the registration differently depending on the grid location. In [39], the authors defined the CC similarity measure as the integral of the product of the interpolated images and computed it analytically in translation-only registration, thereby circumventing resampling. Later, they pointed out in [40] that the energy of the resampled image appearing in registration objective functions (SSD, CC, or MI) can have local optima, and quantified the oscillation artifacts similarly to [12]. They showed that these artifacts disappear if the resampling step is skipped, which they proposed to do in rigid registration by computing the objective function using stochastic integration (instead of summation) [40].

Thorough mathematical analysis of the resampling artifacts in MI-based registration would be a complex task, especially because the computation of the image histogram is involved. Concerning the SSD- and CC-based registration algorithms, noteworthy investigation has been conducted in the above references. Nonetheless, we have not encountered a clear analytical comparison between the errors introduced in the two cases of performing resampling (using the sum cost function) and leaving it out (using the integral cost function), which we aim to provide in Sec. III. We also discuss the implementation of the more accurate resampling-free registration in sections V and VI.A.

D. Rotation Artifacts

As we will explain in Sec. IV, aliasing occurs in resampling of a rotated image, even when the ideal interpolant is employed. To our knowledge, this was originally demonstrated in [41], and afterwards mentioned again in [42-44]. Low-pass filtering the image has been suggested to eliminate the resampling artifacts [41, 44] (outside the registration context), and similar ideas have been employed in

the steerable pyramids framework [45-47] as well. Multi-pass rotation algorithms, such as those decomposing the rotation into sequences of 1D translations through shearing [15, 48-54], have been shown to introduce even further aliasing in the image [42-44] (although negligible [43]).

Regarding the image registration application, aliasing in the image (as opposed to the common frequency) domain due to rotation was discussed in [55] for phase-correlation methods. In a similar context, the authors of [56] addressed the issue of ringing artifacts in rigid registration of fMRI datasets by post-registration filtering. In Sec. IV, we will propose the use of specific interpolation kernels in image registration to avoid aliasing in the frequency domain.

III. INTEGRAL VS. SUM COST FUNCTION

In this section, we consider the sum of squared differences (SSD) as the registration metric to analyze the effect of resampling in the simple case of translation-only registration. For the mutual information (MI) metric, we refer the interested reader to [34-37].

Let $f, g: \mathbb{R}^d \rightarrow \mathbb{R}$ be two continuous d -dimensional images, the axes of which can be aligned via translating one of them, f , to minimize the following SSD cost function:

$$\bar{\Delta}_{id} := \operatorname{argmin}_{\Delta} \int_{\mathbb{R}^d} (f(x - \Delta) - g(x))^2 dx. \quad (1)$$

The subscript id in the optimal shift $\bar{\Delta}_{id}$ stands for *ideal*, meaning that we consider the registration results of the original unsampled images as the ground truth. This is justified by the fact that acquiring higher-resolution images is expected to increase the precision of the registration, and at the limit, the original continuous version of the image (with infinite resolution) produces the ideal solution that can be considered as the ground truth while comparing the two approximate (integral and sum) cost functions.

By expanding the integral in Eq. (1) and taking into account that the integral of the shifted $f^2(\cdot)$ over the entire space is independent of the shift, the equation can be simplified to:

$$\begin{aligned} \bar{\Delta}_{id} &= \operatorname{argmin}_{\Delta} \int_{\mathbb{R}^d} f^2(x - \Delta) dx + \int_{\mathbb{R}^d} g^2(x) dx \\ &\quad - 2 \int_{\mathbb{R}^d} f(x - \Delta) g(x) dx \\ &= \operatorname{argmax}_{\Delta} \int_{\mathbb{R}^d} f(x - \Delta) g(x) dx, \end{aligned} \quad (2)$$

which is equivalent to maximizing the cross correlation (CC) of the two images. Using Parseval's theorem and the shift property of the Fourier transform, $f(x - \Delta) \leftrightarrow \hat{f}(v) e^{-i2\pi v \Delta}$, this can also be written as:

$$\begin{aligned} \bar{\Delta}_{id} &= \operatorname{argmax}_{\Delta} \int_{\mathbb{R}^d} \hat{f}^*(v) \hat{g}(v) e^{i2\pi v \Delta} dv \\ &= \operatorname{argmax}_{\Delta} \int_{[-\frac{1}{2}, \frac{1}{2}]^d} \hat{f}^*(v) \hat{g}(v) e^{i2\pi v \Delta} dv, \end{aligned} \quad (3)$$

with $*$ being the complex conjugate, and \hat{f} and \hat{g} the Fourier transforms of f and g , which for the reasons that we will see below are assumed to be nonzero only in $[-\frac{1}{2}, \frac{1}{2}]^d$.

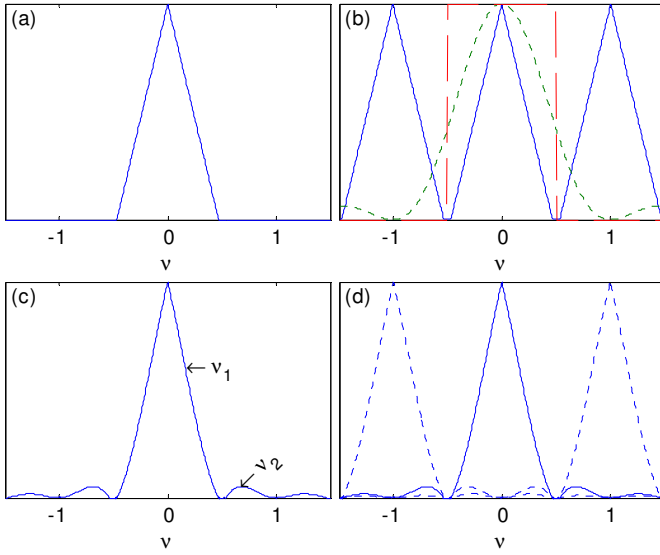


Fig. 1. a) A continuous bandlimited signal in the frequency domain. b) The signal after being sampled. The red dashed and the green dotted filters are the ideal and the non-ideal (linear interpolation) low-pass filters. c) The signal after linear interpolation. Some frequency components (\mathbf{v}_1) are attenuated and some components (\mathbf{v}_2) have been artificially created. d) Aliasing as a result of resampling the non-ideally interpolated signal.

A. Integral Cost Function

Now, suppose that the two images have been sampled on discrete points with the rate of one sample per unit length in each spatial direction. We try to approximate the same integral cost function as (1) by interpolating the two sampled images, i.e. convolving them with the kernel h . Note that the polynomial (e.g. cubic) spline interpolation is also applicable here, since it is equivalent to a convolution with the corresponding cardinal spline function [57].

Sampling a continuous function $s(x)$ (1D for simplicity), in mathematical terms, is multiplying s by the impulse train (Dirac comb), $\sum_{n \in \mathbb{Z}} \delta(x - n)$. This multiplication can be viewed in the frequency domain as the convolution of \hat{s} (Fig. 1a) with the Fourier transform of the impulse train, also an impulse train, resulting in replicas of \hat{s} placed at the frequencies $n \in \mathbb{Z}$, as in $\sum_{n \in \mathbb{Z}} \hat{s}(v - n)$ (Fig. 1b). According to the sampling theorem [58], if s is bandlimited, i.e., $\hat{s}(v) = 0$ for $|v| > 1/2$ (as in Fig. 1a), it can be reconstructed exactly from its sampled points using the ideal low-pass filter (Fig. 1b, red dashed line), i.e. by interpolation with the *sinc* function. Otherwise, the aliasing effects make it impossible to recover the original signal [58], negatively impacting the accuracy of the alignment [59, 60]. For simplicity, here we assume f and g to be bandlimited, meaning that they only have frequency components in $[-1/2, 1/2]^d$.

The infinite impulse response (IIR) of the sinc function makes the ideal low-pass filter difficult to implement (although some IIR filters can be efficiently evaluated using, e.g., recursive filtering), and the use of finite impulse response (FIR) interpolants (e.g. linear interpolation; the green dotted curve in Fig. 1b) more attractive. Nevertheless, in contrast to their low computational cost, such non-ideal kernels have two undesirable properties: attenuation of low-frequency components which are supposed to remain intact (\mathbf{v}_1 in Fig.

1c), and leaving behind some high-frequency components which are supposed to be removed (\mathbf{v}_2 in Fig. 1c). The latter is the major cause of the resampling artifacts in general, as explained here and in [4, 14, 15].

Let h be the kernel that we use to interpolate the two sampled images. Convolution of the sampled versions of images f and g with h results in continuous images that we use in place of f and g in the cost function of Eq. (1). As in the previous case, the integrals of such continuous functions in the entire space are independent of the shift [40], and therefore SSD minimization reduces again to CC maximization (similarly to Eqs. (2,3)). Convolution of f with h can be seen in the frequency domain as multiplication of the replicas of \hat{f} by the frequency response of the kernel, \hat{h} , resulting in $\hat{h}(v) \sum_{n \in \mathbb{Z}^d} \hat{f}(v - n)$. The CC in this case becomes the following integral, which we simplify by dividing it into finite-length segments (see also [39]):

$$\begin{aligned}
 & \int_{\mathbb{R}^d} \left(\hat{h}^*(v) \sum_{n \in \mathbb{Z}^d} \hat{f}^*(v - n) \right) \left(\hat{h}(v) \sum_{m \in \mathbb{Z}^d} \hat{g}(v - m) \right) e^{i2\pi v \cdot \Delta} dv \\
 &= \int_{\mathbb{R}^d} \sum_{m, n \in \mathbb{Z}^d} \hat{f}^*(v - n) \hat{g}(v - m) e^{i2\pi v \cdot \Delta} |\hat{h}(v)|^2 dv \\
 &= \sum_{k \in \mathbb{Z}^d} \int_{[-1/2, 1/2]^d} \sum_{m, n \in \mathbb{Z}^d} \hat{f}^*(v + k - n) \hat{g}(v + k - m) e^{i2\pi(v+k) \cdot \Delta} |\hat{h}(v + k)|^2 dv \\
 &= \sum_{k \in \mathbb{Z}^d} \int_{[-1/2, 1/2]^d} \hat{f}^*(v) \hat{g}(v) e^{i2\pi(v+k) \cdot \Delta} |\hat{h}(v + k)|^2 dv \\
 &= \int_{[-1/2, 1/2]^d} \hat{f}^*(v) \hat{g}(v) e^{i2\pi v \cdot \Delta} \sum_{k \in \mathbb{Z}^d} |\hat{h}(v + k)|^2 e^{i2\pi k \cdot \Delta} dv,
 \end{aligned} \tag{4}$$

where we used the bandlimitedness of f and g to reach line 4 from line 3 of this equation. Thus, the optimum shift using the *integral* cost function, $\bar{\Delta}_I$, becomes:

$$\bar{\Delta}_I := \operatorname{argmax}_{\Delta} \int_{[-1/2, 1/2]^d} \hat{f}^*(v) \hat{g}(v) e^{i2\pi v \cdot \Delta} \sum_{k \in \mathbb{Z}^d} |\hat{h}(v + k)|^2 e^{i2\pi k \cdot \Delta} dv. \tag{5}$$

By comparing Eqs. (3) and (5), it becomes apparent that the error factor $\hat{\epsilon}_v(\Delta) := \sum_{k \in \mathbb{Z}^d} |\hat{h}(v + k)|^2 e^{i2\pi k \cdot \Delta}$ for $v \in [-1/2, 1/2]^d$, which is periodic with the period of 1 pixel in every dimension, is introduced as a result of non-ideal interpolation. $\hat{\epsilon}_v(\Delta)$ can be seen as the discrete-time Fourier transform (DTFT) of $h_v[k] := |\hat{h}(v - k)|^2$. The error-free case corresponds to $\hat{\epsilon}_v^{ef}(\Delta) = 1$, meaning that $h_v^{ef}[k] = \text{DTFT}^{-1}(1) = \delta_k$, or:

$$\begin{aligned}
 |\hat{h}^{ef}(v - k)| &= \begin{cases} 1 & k = 0 \\ 0 & k \neq 0 \end{cases}, \quad v \in [-1/2, 1/2]^d, \\
 |\hat{h}^{ef}(v)| &= \begin{cases} 1 & v \in [-1/2, 1/2]^d \\ 0 & \text{o. w.} \end{cases},
 \end{aligned} \tag{6}$$

which is the ideal low-pass filter. Note that there is no restriction on the phase of $\hat{h}^{ef}(v)$, since the filter is applied to both images and the phase is cancelled in the CC. The intrinsic symmetry of rigid registration also becomes apparent when images are treated as continuous, i.e. in Eqs. (3) and (5), meaning that swapping f and g and negating Δ does not change the cost function.

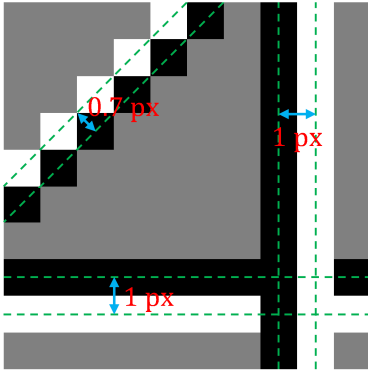


Fig. 2. For the lines to be distinguishable in the 2D Cartesian grid, they can be $\sqrt{2}$ times closer in the diagonal direction than in the horizontal and the vertical directions. Some fine details are therefore observed only in the diagonal directions (e.g. see Fig. 7), making the resolution higher in those directions than in the axial ones.

B. Sum Cost Function

It is common in image registration to compute the SSD merely on the grid points, while interpolating and resampling only *one* of the images, f , as needed:

$$\bar{\Delta}_S := \operatorname{argmin}_{\Delta} \sum_{p \in \mathbb{Z}^d} (f_{\Delta}[p] - g[p])^2, \quad (7)$$

where $f_{\Delta}[\cdot]$ is the image f after being sampled, interpolated, shifted with Δ , and resampled on the original grid, and $g[\cdot]$ is the image g sampled on the reference grid. The optimum shift of the above *sum* cost function, $\bar{\Delta}_S$, can be expanded as:

$$\bar{\Delta}_S = \operatorname{argmax}_{\Delta} \sum_{p \in \mathbb{Z}^d} f_{\Delta}[p]g[p] - \frac{1}{2} \sum_{p \in \mathbb{Z}^d} f_{\Delta}^2[p]. \quad (8)$$

Contrary to the previous case, we cannot trivially remove the energy of the resampled image, since as also demonstrated in [12, 40], $\sum_{p \in \mathbb{Z}^d} f_{\Delta}^2[p]$ depends on the position of the transformed image relative to the grid, and therefore on Δ . The discrete case of Parseval's theorem states that the inner product of two discrete signals can be computed in the frequency domain as the inner product of their DTFTs in $[-\frac{1}{2}, \frac{1}{2}]^d$. Image g is bandlimited and not interpolated, therefore according to the sampling theorem, $\text{DTFT}(g[p]) = \hat{g}(\nu)$ for $\nu \in [-\frac{1}{2}, \frac{1}{2}]^d$. Regarding $f_{\Delta}[\cdot]$, we saw in Sec. III.A that the Fourier transform of f after being sampled, interpolated, and shifted, is $\hat{h}(\nu) \sum_{n \in \mathbb{Z}^d} \hat{f}(\nu - n) e^{-i2\pi\nu\Delta}$ (Fig. 1c). Resampling, results in replicas of this function at $k \in \mathbb{Z}^d$ (Fig. 1d), i.e.,

$$\begin{aligned} \text{DTFT}(f_{\Delta}[p]) &= \sum_{k \in \mathbb{Z}^d} \hat{h}(\nu + k) \sum_{n \in \mathbb{Z}^d} \hat{f}(\nu + k - n) e^{-i2\pi(\nu+k)\Delta} \\ &= \hat{f}(\nu) e^{-i2\pi\nu\Delta} \sum_{k \in \mathbb{Z}^d} \hat{h}(\nu + k) e^{-i2\pi k\Delta}, \end{aligned} \quad (9)$$

for $\nu \in [-\frac{1}{2}, \frac{1}{2}]^d$, where the bandlimitedness of f was once again exploited. Finally, using the DTFTs of $f[\cdot]$ and $g[\cdot]$ and applying Parseval's theorem to Eq. (8) leads to:

$$\bar{\Delta}_S = \operatorname{argmax}_{\Delta} \int_{[-\frac{1}{2}, \frac{1}{2}]^d} \hat{f}^*(\nu) \hat{g}(\nu) e^{i2\pi\nu\Delta} \sum_{k \in \mathbb{Z}^d} \hat{h}^*(\nu + k) e^{-i2\pi k\Delta} d\nu \quad (10)$$

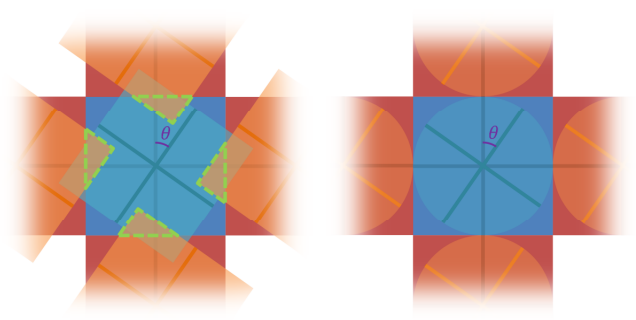


Fig. 3. (Left) Rotation in the Cartesian grid results in high-frequency components in the corners of the frequency domain (identified by green dashed lines) to be aliased into the low-frequency regions. (Right) A disk-shaped filter removes the frequency components that can potentially be aliased during rotation.

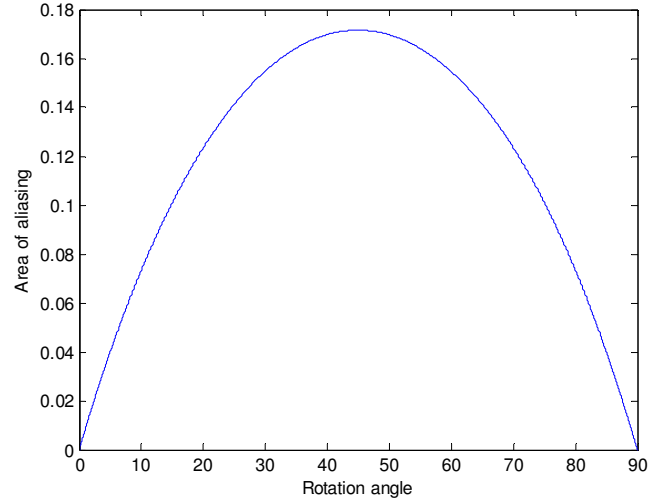


Fig. 4. Area of the aliasing regions in the frequency domain with respect to the rotation angle.

$$-\frac{1}{2} \int_{[-\frac{1}{2}, \frac{1}{2}]^d} |\hat{f}(\nu)|^2 \left| \sum_{k \in \mathbb{Z}^d} \hat{h}(\nu + k) e^{-i2\pi k\Delta} \right|^2 d\nu.$$

The CC (first) integral of the above equation is similar to that of the previous case, Eq. (5), with the exception that \hat{h}^* , as opposed to $|\hat{h}|^2$, appears in the periodic error term, meaning that for a given interpolation kernel, the CC after resampling has *more aliasing* artifacts, but *less smoothing* effects. This is because the larger magnitude of \hat{h}^* (compared to $|\hat{h}|^2$) results in larger unwanted aliased frequency components (ν_2 in Fig. 1c), but at the same time prevents too much attenuation of the desired frequencies (ν_1 in Fig. 1c). However, as shown in [12, 40], the second integral of Eq. (10) is an additional periodic error term that makes the sum cost function (Eq. (10)) deviate from the ground truth (Eq. (3)) even further.¹ This error in the energy of the resampled image, which for instance breaks the symmetry inherent to rigid image registration, does not appear when resampling is avoided (Eq. (5)). Note that the aliased frequencies (ν_2 in Fig. 1c) are results of using a non-ideal interpolant; aliasing happens even though the original images are assumed bandlimited.

¹ This error term is however constant for shiftable transforms [46].

With similar analysis as in Sec. III.A, both sources of error in Eq. (10) can be seen to disappear (become independent of Δ) if the ideal interpolant (sinc) is used. This raises the question of whether employing ideal (or more feasibly, close to ideal) kernels generally eliminates the resampling artifacts in image registration. Although we just proved this to be true in translation-only registration, we will see in the next section that this is in general not the case in image registration, when rotation is involved. We will adopt a rather intuitive analysis, as rigorous comparison of the integral and the sum cost functions with non-ideal interpolation in the presence of rotation would be tedious and outside the scope of this paper.

IV. ALIASING DUE TO ROTATION

Images are often represented on *radially non-symmetric* Cartesian grids. An instance of such asymmetry, as illustrated in Fig. 2 (see also [61]), would be the fact that a d -dimensional Cartesian grid provides a resolution \sqrt{d} times higher in the diagonal directions than in the axis directions, which can also be perceived as higher frequency components in the corners of the DTFT domain. Rotation of an interpolated image may align its high-resolution diagonal features with the axes of the grid, which do not have the capacity of representing such high-frequency components. Thus, if the rotated image is resampled on the original grid, aliasing will inherently occur (Fig. 3, left), even if the ideal interpolation has been used [41-44].

This problem can be particularly serious in image registration, where the dependence of the aliasing error in the (sum) cost function on the rotation angle may bias the registration. The area of the aliased regions in the frequency domain (dashed triangles on Fig. 3, left) for a rotation of angle θ can be calculated geometrically and shown to be $A(\theta) = (\sec\theta - 1)(\csc\theta - 1)$, plotted in Fig. 4 for $\theta \in [0^\circ, 90^\circ]$. Although $A(\theta)$ in this interval is maximal at $\theta = 45^\circ$, we cannot simply deduce that the bias is worst at this angle (and at $135^\circ, 225^\circ$, etc.), since the *minimum*, and not the zeros, of the cost function is searched for. This minimum happens at a zero of the *derivative* of the cost function, in which the error would depend on both $A(\theta)$ and $A'(\theta)$, with the magnitude of the latter being largest when the grids are aligned, i.e. at $\theta = 0^\circ, \pm 90^\circ, 180^\circ$.

Note again that computing the integral cost function instead of the sum eliminates the resampling step and consequently the resulting aliasing artifacts. Nonetheless, with rotation involved, such an analytical computation becomes much more complicated than the translation-only case of Sec. III.A, as various overlap configurations for the two meshes will have to be considered. Monte-Carlo methods may be employed to approximate the integral, yet a good such approximation requires oversampling the image, thereby increasing the computational cost. (An experimental comparison can be seen in [40, Fig. 12] for small rotation angles, although the authors do not mention the type of aliasing discussed here). Two alternatives to using the integral cost function to avoid aliasing would be: 1) applying the ideal low-pass filter a second time on the rotated interpolated image to clip off the corners [44], or 2) upsampling the d -dimensional images \sqrt{d} times in each direction, thus increasing the total number of sample points by

a factor of $d^{d/2}$. Besides the additional computational cost imposed by these approaches to avoid aliasing, one can see that the cost function would still be biased, since the intersection area of the two registered images (in the frequency domain), which is their non-aliased overlapping regions, $1 - A(\theta)$, still depends on the rotation angle. Normalizing the cost function by this area would not in general be helpful, either, since the frequency components are not necessarily distributed uniformly.

It has been suggested in [41] (outside the registration context) to filter an image with a circular disk (Fig. 3, right) before rotation is applied, to avoid resampling-induced aliasing for any arbitrary angle. Additionally, the corners outside of such a disk have been ignored in *phase-correlation* registration in [55], which would be theoretically equivalent to applying such a filter. The authors of [56] have also used this filter, however as a post-processing step to remove “ringing artifacts” (as opposed to aliasing).

In this work, we propose to make use of the disk-shaped filter in rigid image registration by incorporating it in the interpolation step. We start by identifying the interpolation kernel h with the Fourier transform $\hat{h}(v) = \begin{cases} 1 & \|v\| \leq \frac{1}{2} \\ 0 & o.w. \end{cases}$. Such functions have been computed in [62, 63] for different dimensions; in 2D and 3D spaces they are:

$$\begin{aligned} h_{2D}(r) &= \frac{J_1(\pi r)}{\pi r}, \\ h_{3D}(r) &= \frac{1}{2\pi r^2} \left(\frac{\sin(\pi r)}{\pi r} - \cos(\pi r) \right), \end{aligned} \quad (11)$$

with $r := \|x\|$, and J_1 the first order Bessel function. Such oscillatory interpolation kernels have been considered in [63] as *radial basis functions* for interpolation, however, neither for convolution-based interpolation, nor in the registration framework.

Interpolating the image with the IIR kernels in (11) – instead of with the ideal interpolant (sinc) – filters the image with the disk/ball-shaped kernel of a diameter equal to the sampling rate. Such a filter keeps the maximal frequency components in the image, while preventing the occurrence of aliasing at every arbitrary rotation in an *unbiased* fashion, meaning that the same amount of information is preserved for all the possible rotations. Therefore, with no aliasing coming from resampling, the two sum and integral cost functions theoretically produce the same results.

Nevertheless, for the same practical considerations as mentioned in Sec. III.A, the use of radial windows such as the radial Hanning window becomes indispensable to make FIR filters out of kernels (11). The radial Hanning window, for instance, is defined as $w(n_x, n_y, n_z) := 0.5 \left(1 + \cos \frac{2\pi n}{N-1} \right) \mathbf{1}_{(n < \frac{N-1}{2})}$, where N is the diameter of the window, $\mathbf{1}$ the characteristic function, and $n := \sqrt{n_x^2 + n_y^2 + n_z^2}$ the distance from the center of the window. The use of the window function, however, makes the kernels non-ideal which may consequently cause aliasing artifacts again as described in Sec. III, making the sum and integral cost functions in practice dissimilar for smaller windows. This is a general disadvantage of the FIR filters; the fact that they cannot exactly produce an arbitrary desired (especially bandlimited) frequency response.

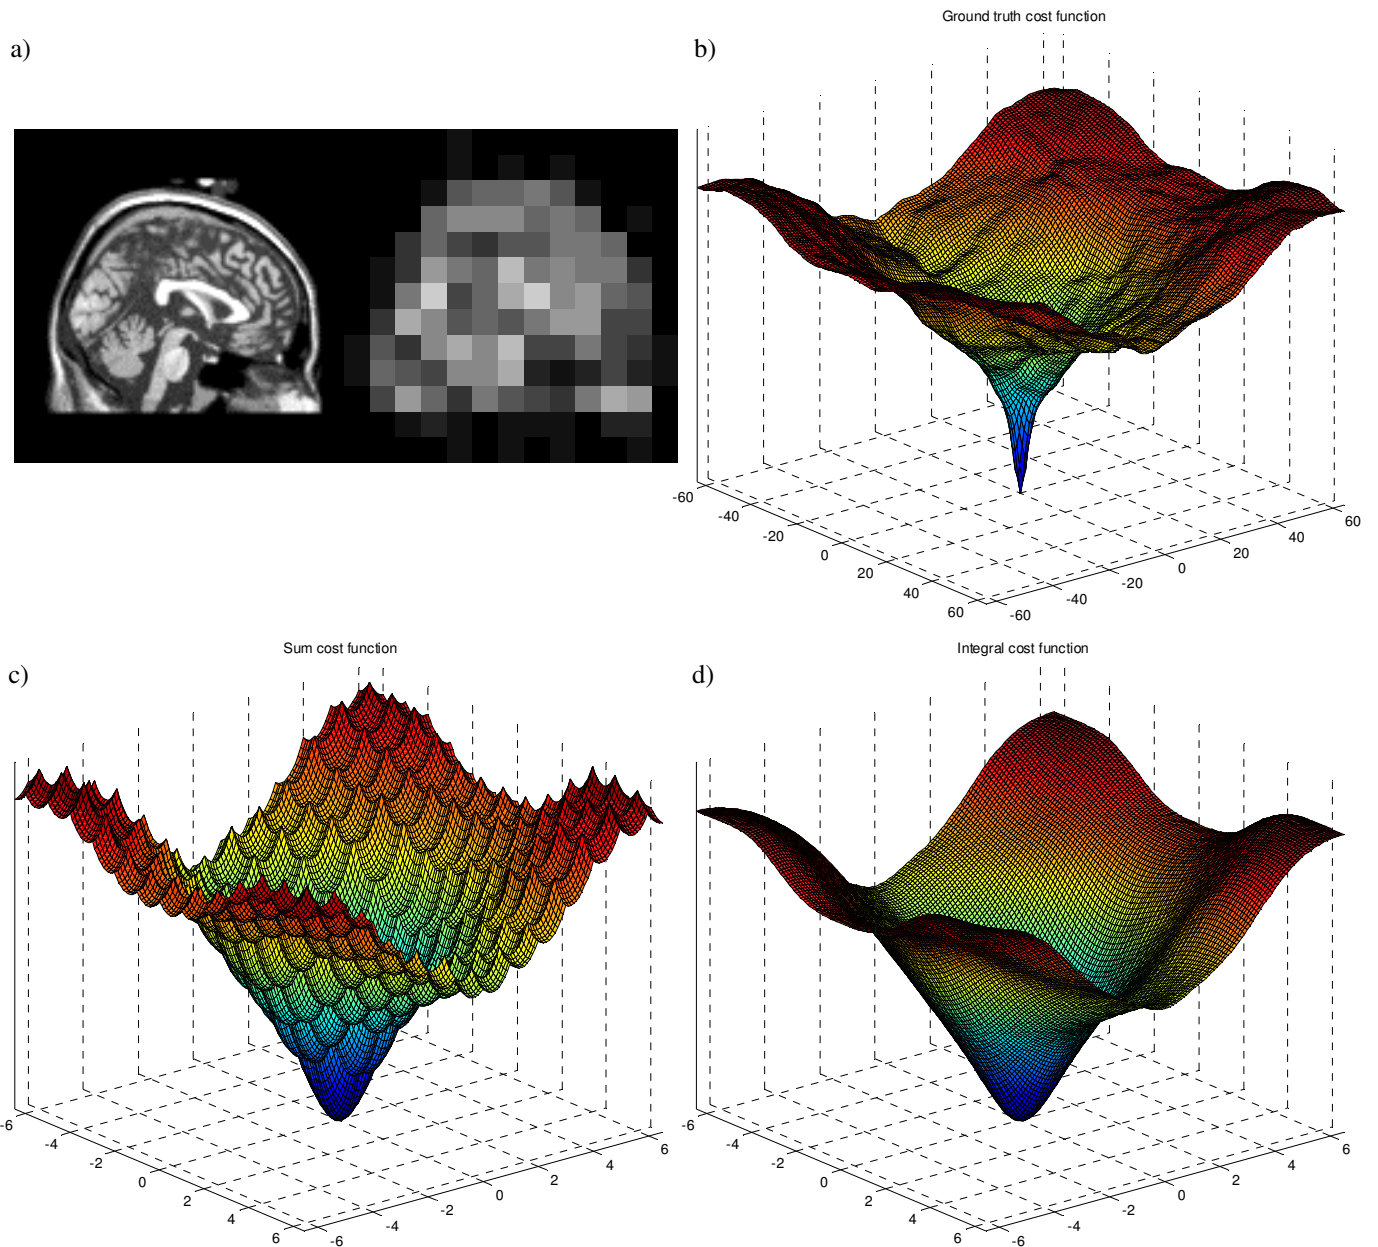


Fig. 5. (a) High-res and low-res versions of the sagittal slice. (b) 2D registration cost function of the high-res image. Registration sum (c) and integral (d) cost functions of the low-res image.

V. PRACTICAL CONSIDERATIONS

Analytical computation of the integral cost function requires solving definite integrals in irregular d -dimensional regions defined by the intersections of the pixels (or voxels) of the two meshes. This can be complicated particularly when rotation is involved, however, a simple Monte-Carlo approximation of the integral cost function is achieved by oversampling the image on random [40] or quasi-random [28] sample points.

Nevertheless, the particular case of translation-only registration with the integral cost function discussed in Sec. III.A (Eq. (5)) can be shown to be equivalent to maximizing the CC after interpolating its discrete version with the self-convolution of the kernel, $h * h$ (see also [39]). This can be done quickly by computing the CC using the fast Fourier transform approaches (e.g. as in [64, 65]),

interpolating and upsampling it to the desired resolution, and finding the maximum of the upsampled CC. Otherwise, computing the integral cost function would require dividing each voxel into 2^d subvoxels depending on the subvoxel shift between the two images, and subsequently calculating the integral analytically in each subregion. Once the analytical formula for a subregion is derived, the formula for the rest of the subregions can be computed by swapping the two images and shifting them. For instance, if the integral cost function is desired to be computed in the 1D case with $0 \leq \Delta \leq 1$, the pixel $[0,1]$ needs to be divided into the two subpixel intervals of $[0,\Delta]$ and $[\Delta,1]$. Once the analytical formula $I(f[n], g[n], \Delta)$ is derived for the integral of the first interval, the integral of the second interval can be seen to be $I(g[n+1], f[n], 1-\Delta)$, and therefore does not need to be computed from scratch.

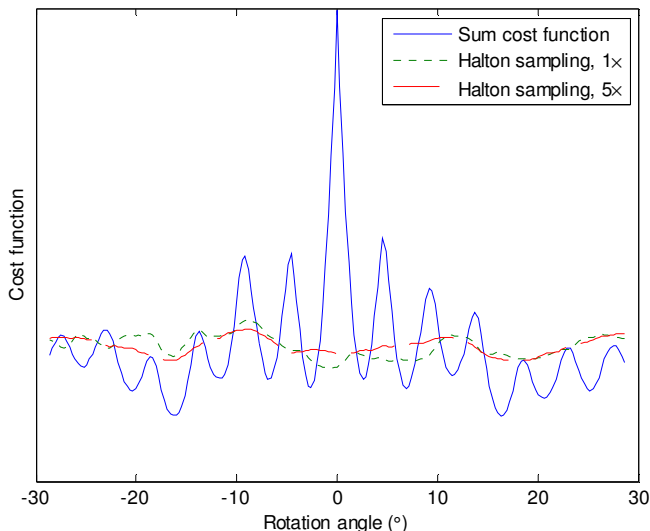


Fig. 6. Comparison of sum (solid blue) and integral cost functions with the same number of Halton sample points (green dotted) and five times more points (red dashed) in spherical registration of white noise.

For our experiments on the sphere, we estimated the integral cost function via the Monte-Carlo integral approximation method. We generated spatially-uniform quasi-random Halton sample points on the 2D Cartesian interval of $(z, \phi) \in [-1, 1] \times [0, 2\pi]$, using the `haltonset` command of Matlab. We then transferred the points onto the sphere by computing $\theta = \arccos z$. One can see that the surface element $|\sin \theta d\theta d\phi| = |d(\cos \theta) d\phi| = |dz d\phi|$ remains unchanged, preserving the uniformity of the sample points.

Regarding the disk-shaped kernels (11) in Sec. IV, a drawback of their employment is the fact that they cannot be implemented in a separable fashion, i.e. as a product of 1D functions, thereby making them computationally more expensive. In fact, besides the Gaussian kernel, no other radially symmetric kernel is separable [4]. To make up for the computational complexity, one might want to pre-filter the image with the disk-shaped kernel offline to remove the corner frequency components and subsequently interpolate it at each iteration using a separable kernel such as bicubic or windowed sinc, the downside of which, yet, would be the introduction of undesirable artifacts (aliasing, blurring, ringing, etc.) twice over, instead of only once.

In Sec. IV, we made the general assumption that images have nonzero components in the corners of the frequency domain, since they are represented on the Cartesian lattice, and that they have not been upsampled to avoid diagonal aliasing. In case the interpolated image is known not to contain any components in the corners of the frequency domain, such as in MRI with spiral trajectories in the k-space, aliasing due to rotation will naturally not occur and the use of disk-shaped kernels is not expected to bring about any improvement.

VI. RESULTS AND DISCUSSIONS

A. Integral vs. Sum Cost Function

Cartesian Grid

In this section, we first compare the integral and sum cost functions by visualizing them in 2D self-registration of a

synthesized brain magnetic resonance image (MRI) slice taken from the BrainWeb simulated brain database [66, 67].² A T1 image of a normal brain with isotropic 1-mm³ voxels was generated, and a sagittal slice of it with the dimensions 129×129 was taken for the task of 2D translation-only registration. The ideal low-pass filter was then applied to the image which was subsequently downsampled 10 times (to 13×13). Both high-res and low-res versions of the images are shown on Fig. 5(a). Next, the high-res image was registered to itself with the resolution of 1 pixel, the cost function of which is illustrated on Fig. 5(b) with respect to the 2D shift. This cost function is considered as the ground truth for the low-res downsampled data, which were next registered with the resolution of 0.1 pixels using linear interpolation. The sub-pixel registration of the low-res data was performed once using the sum cost function (Fig. 5(c)), and again using the integral cost function by analytically computing the integral of the square of the difference of the continuous interpolated images on the sub-pixel regions created by the overlapping pixels of the fixed and the shifted meshes (Fig. 5(d)). Resampling artifacts can be clearly seen in the sum cost function as periodic piecewise-convex regions. Their convex nature is due to the quadratic convexity of the square of the linearly interpolated image, appearing in the SSD. Such artifacts can, for instance, alter the direction of the gradient in gradient-based optimization approaches, potentially leading to local optima and slower convergence. The artifacts, however, are much smaller and almost invisible in the smooth integral cost function. Furthermore, when we computed the correlation coefficient between each low-res cost function and the ground truth, it was slightly higher (0.3%) for the integral case ($r = 0.996$) than for the sum case ($r = 0.993$).

Spherical Registration

We also looked into the patterns of artifacts in spherical image registration, which has applications in, e.g., registration of the cerebral cortex (see [68] and the references therein). We generated two separate images of independent and identically distributed white noise on the sphere with a triangular mesh based on the icosahedron subdivision [68], with 2562 vertices and 5120 triangles. We then computed the *sum* SSD registration cost function of the two images with barycentric (linear) interpolation while rotating one of them around a fixed axis. In the absence of artifacts, the cost function is expected to be roughly flat superimposed with an irregular noise pattern. However, structured artifacts with a period of 4.6° – about the angular size of the edges – were observed in the cost function (Fig. 6, blue solid curve), most strongly affecting the origin where the two triangular meshes are entirely aligned. Rotation of a sphere acts as *rotation* close to the poles and as *translation* close to the equator, which explains why the cost function appears to be a combination of patterns resulted by both rotation with a bias at the grid aligning origin (see for instance [40, Fig. 12]), and translation with periodic artifacts (similar to Fig. 5(c)).

² To avoid algorithmic biases in the results, we only change the interpolation strategy in each experiment while using the same registration algorithm.

TABLE I
REGISTRATION ERROR FOR DIFFERENT SAMPLING METHODS

| Ground truth | Regular sampling | Halton, 1× | Halton, 5× |
|-----------------------|------------------|------------|------------|
| Regular sampling, Ic6 | 1.64° | 1.45° | 1.29° |
| Halton sampling, Ic6 | 1.22° | 1.02° | 0.78° |

The spherical registration error for different sampling schemes on Ic4, where Ic6 is used as the ground truth.

To compute the *integral* cost function, we resorted to the Monte Carlo approach of stochastic estimation of an integral [40], and approximated the integral as a sum over randomly distributed points on the sphere. Moreover, since the uniform distribution produces sample points that do not fill the surface as uniformly as expected (due to the clustering effects), we followed the suggestion in [28] and used Halton sampling [69] to produce a set of quasi-random but more uniform set of points. Figure 6 shows the results with as many Halton sample points as the number of vertices (1×, green dotted curve) and with five times more points (5×, red dashed curve). As can be seen, there are significantly less artifacts compared to the case with the sum cost function. Additionally, a better approximation of the integral (more sample points) results in a smoother cost function. The standard deviations of the integral – 1× and 5× – and the sum cost functions over the rotation angle were 3.3, 2.5, and 16.7, respectively.

We then used the cortical sulcal maps of the five subjects included in the Spherical Demons package [68], originally from the OASIS public database [70], to test the spherical registration on real brain MRI data. The maps were projected on icosahedron meshes of orders four (low-res Ic4, 2562 vertices) and six (high-res Ic6, 40962 vertices), for both of which we computed the sum and integral cost functions similarly to the previous case, except that this time images were rotated about two axes, resulting in 2D cost functions in $[-10^\circ, 10^\circ]^2$. Every pair in the five subjects was forward and backward registered, totaling 20 experiments. We found the location of the minimum of each cost function, and used the high-res Ic6 minima as the ground truths for comparing the low-res Ic4 results of the different techniques. The error for each technique was computed as the L_2 distance between the locations of cost function minima in Ic4 and Ic6, the average of which across experiments is listed in Table I. As can be seen, low-res registration with Halton instead of regular sampling produces minima that are more compatible with those of the high-res image. In addition, a fivefold increase in the number of Halton samples further improves the accuracy of registration via better approximation of the integral. One can also see that when Halton sampling is used in the ground truth Ic6 registration, the low- and high-res experiments result in closer minima. This is expected, since the artifacts exist even in high-res registration, albeit with lower magnitude, and reducing them via Halton sampling eliminates one source of perturbation in the distance between the minima.

We also measured the inverse-consistency of registration [10] by computing the L_1 distance between the forward and backward registration cost functions corresponding to each pair of subjects, and observed that on average, compared to

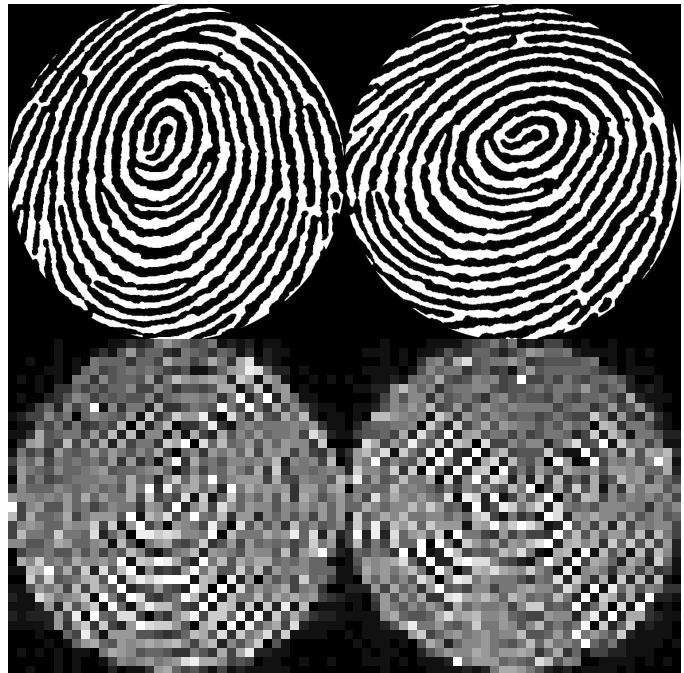


Fig. 7. Hi-res (top) and low-res (bottom) versions of the original (left) and the rotated (right) fingerprint image. Note how in the both low-res images the diagonal frequencies are better retained than the horizontal and vertical frequencies are.

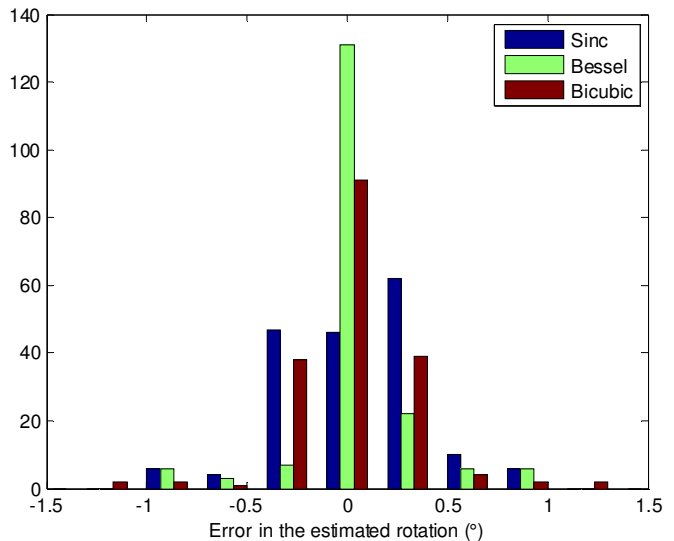


Fig. 8. Histogram of the registration error with sinc, Bessel, and bicubic interpolations.

regular sampling, the asymmetry error is reduced by 36% and 56% using Halton sampling with the same number of and five times more sample points, respectively.

B. Rotation-Induced Aliasing

Next, we performed rotation-only registration experiments to validate the interpolation kernel introduced in Sec. IV. We used a high-res 760×760 planar fingerprint image [71] (Fig. 7, top left) as the ground truth, and each time rotated it with a specific angle (Fig. 7, top right). Subsequently, we applied the ideal low-pass filter to both the original and rotated images and downsampled them 20 times (Fig. 7, bottom). Then we registered the low-res images by an exhaustive search in the

vicinity of the known rotation angle with a resolution of 0.1° , using the sum cost function. We performed the experiment three times, making use of the Bessel (h_{2D} , Eq. (11)) and the sinc kernels (both with a 15-tap Hanning window) and also the bicubic interpolation implemented in Matlab (`imrotate`). We repeated the experiment for 181 rotation angles from 0° to 180° , and noted that the Bessel kernel with a mean error of 0.18° outperformed the bicubic and the sinc interpolations with mean errors of 0.20° and 0.29° , respectively. The corresponding p-values derived using Wilcoxon's paired signed rank test for the equality of the median of errors were 3×10^{-19} , 3×10^{-2} , and 3×10^{-7} , for Bessel vs. sinc, Bessel vs. bicubic, and bicubic vs. sinc, respectively. The histogram of the signed errors is plotted in Fig. 8, where the Bessel kernel can clearly be seen to have a higher concentration of error around zero than the other two methods do.

Bicubic interpolation seems to be more robust than the windowed sinc, which may be explained by the fact that it is equivalent to convolution with the cardinal cubic kernel which approximates the IIR sinc [57], as opposed to the FIR 15-tap windowed sinc kernel used here. Also note that no matter what the rotation angle is, the low-res images (Fig. 7, bottom) resolve the high frequencies better in the diagonal than in the horizontal and vertical directions, in accordance with our discussion in Sec. IV and Fig. 2.

VII. CONCLUSION

In this paper, we discussed the artifacts in subpixel registration from a sampling point of view, and showed that resampling the interpolated image is the major cause of these artifacts. We compared the two cases of including and excluding the resampling step in image registration analytically, and also through experiments by calculating the *integral* cost function via both analytical and stochastic approaches. Computing the objective function in a continuous manner dramatically reduced the amplitude of the induced artifacts, thereby increasing the accuracy of the gradient and the global optimum of the cost function, and also the inverse-consistency of registration. We also demonstrated aliasing errors due to rotation, and proposed to use radially symmetric interpolation kernels to avoid them. We observed significant improvement in registration accuracy by choosing the proposed interpolation kernels over the traditional ones.

Applying the derived methods to functional connectivity MRI registration is part of an ongoing project, where we have noted that not addressing the resampling artifacts guaranteed that almost every registration would converge to the wrong optimum. Future work consists of devising practical implementations with lower computational complexity for the integral cost function.

REFERENCES

- [1] T. M. Lehmann, C. Gonner, and K. Spitzer, "Survey: interpolation methods in medical image processing," *IEEE Transactions on Medical Imaging*, vol. 18, no. 11, pp. 1049-75, Nov, 1999.
- [2] E. Meijering, "A chronology of interpolation: from ancient astronomy to modern signal and image processing," *Proceedings of the IEEE*, vol. 90, no. 3, pp. 319-342, 2002.
- [3] P. Thévenaz, T. Blu, and M. Unser, "Chapter 28 - Image Interpolation and Resampling," *Handbook of Medical Image Processing and Analysis (Second Edition)*, N. B. Isaac and I. N. Bankman, eds., pp. 465-493, Burlington: Academic Press, 2009.
- [4] J. Parker, R. V. Kenyon, and D. E. Troxel, "Comparison of interpolating methods for image resampling," *IEEE Transactions on Medical Imaging*, vol. 2, no. 1, pp. 31-9, 1983.
- [5] D. L. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes, "Medical image registration," *Physics in medicine and biology*, vol. 46, no. 3, pp. R1-45, Mar, 2001.
- [6] B. Zitová, and J. Flusser, "Image registration methods: a survey," *Image and Vision Computing*, vol. 21, pp. 977-1000, 2003.
- [7] L. Ng, and L. Ibanez, "Medical image registration: Concepts and implementation," *Insight into images: principles and practice for segmentation, registration, and image analysis*, T. S. Yoo, ed.: A K Peters, 2004.
- [8] O. Scherzer, *Mathematical models for registration and applications to medical imaging*: Springer, 2006.
- [9] G. E. Christensen, and H. J. Johnson, "Consistent image registration," *Medical Imaging, IEEE Transactions on*, vol. 20, no. 7, pp. 568-582, 2001.
- [10] M. Reuter, H. D. Rosas, and B. Fischl, "Highly accurate inverse consistent registration: A robust approach," *NeuroImage*, vol. 53, no. 4, pp. 1181-1196, 2010.
- [11] S. K. Park, and R. A. Schowengerdt, "Image sampling, reconstruction, and the effect of sample-scene phasing," *Appl. Opt.*, vol. 21, no. 17, pp. 3142-3151, 1982.
- [12] A. Schaum, "Theory and Design of Local Interpolators," *CVGIP: Graphical Models and Image Processing*, vol. 55, no. 6, pp. 464-481, 1993.
- [13] A. C. Popescu, and H. Farid, "Exposing digital forgeries by detecting traces of resampling," *Signal Processing, IEEE Transactions on*, vol. 53, no. 2, pp. 758-767, 2005.
- [14] Q. Tian, and M. N. Huhns, "Algorithms for subpixel registration," *Computer Vision, Graphics, and Image Processing*, vol. 35, no. 2, pp. 220-233, 1986.
- [15] R. Szeliski, S. Winder, and M. Uyttendaele, *High quality multi-pass image resampling*, MSR-TR-2010-10, 2010.
- [16] S. Grootoink, C. Hutton, J. Ashburner, A. M. Howseman, O. Josephs, G. Rees *et al.*, "Characterization and Correction of Interpolation Effects in the Realignment of fMRI Time Series," *NeuroImage*, vol. 11, no. 1, pp. 49-57, 2000.
- [17] W. M. Wells Iii, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis, "Multi-modal volume registration by maximization of mutual information," *Medical Image Analysis*, vol. 1, no. 1, pp. 35-51, 1996.
- [18] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *Medical Imaging, IEEE Transactions on*, vol. 16, no. 2, pp. 187-198, 1997.
- [19] A. A. Cole-Rhodes, and P. K. Varshney, "Image registration using mutual information," *Image Registration for Remote Sensing*, J. L. Moigne, N. S. Netanyahu and R. D. Eastman, eds.: Cambridge University Press, 2011.
- [20] J. P. W. Pluim, J. B. Antoine Maintz, and M. A. Viergever, "Interpolation Artefacts in Mutual Information-Based Image Registration," *Computer Vision and Image Understanding*, vol. 77, no. 2, pp. 211-232, 2000.
- [21] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal, "Automated multi-modality image registration based on information theory," *Information Processing in Medical Imaging*, pp. 263-274, 1995.
- [22] J. P. W. Pluim, J. B. A. Maintz, and M. A. Viergever, "Image registration by maximization of combined mutual information and gradient information," *Medical Imaging, IEEE Transactions on*, vol. 19, no. 8, pp. 809-814, 2000.
- [23] B. Likar, and F. Pernuš, "A hierarchical approach to elastic registration based on mutual information," *Image and Vision Computing*, vol. 19, no. 1-2, pp. 33-44, 2001.
- [24] H. M. Chen, and P. K. Varshney, "Mutual information-based CT-MR brain image registration using generalized partial volume joint histogram estimation," *IEEE Transactions on Medical Imaging*, vol. 22, no. 9, pp. 1111-9, Sep, 2003.
- [25] J. Tsao, "Interpolation artifacts in multimodality image registration based on maximization of mutual information," *IEEE Transactions on Medical Imaging*, vol. 22, no. 7, pp. 854-64, Jul, 2003.

- [26] M. Shimizu, and M. Okutomi, "Sub-Pixel Estimation Error Cancellation on Area-Based Matching," *International Journal of Computer Vision*, vol. 63, no. 3, pp. 207-224, 2005.
- [27] P. Aljabar, J. V. Hajnal, R. G. Boyes, and D. Rueckert, "Interpolation artefacts in non-rigid registration," *MICCAI - International Conference on Medical Image Computing and Computer-Assisted Intervention*, vol. 8, no. Pt 2, pp. 247-54, 2005.
- [28] P. Thévenaz, M. Bierlaire, and M. Unser, "Halton sampling for image registration based on mutual information," *Sampling Theory In Signal And Image Processing*, 2006.
- [29] V. Walimbe, and R. Shekhar, "Automatic elastic image registration by interpolation of 3D rotations and translations from discrete rigid-body transformations," *Medical Image Analysis*, vol. 10, no. 6, pp. 899-914, Dec, 2006.
- [30] O. Salvado, and D. L. Wilson, "Removal of local and biased global maxima in intensity-based registration," *Medical Image Analysis*, vol. 11, no. 2, pp. 183-96, Apr, 2007.
- [31] E. Ardizzone, R. Gallea, O. Gambino, and R. Pirrone, "Effective and efficient interpolation for mutual information based multimodality elastic image registration," in *Computer Vision Workshops (ICCV Workshops)*, 2009 IEEE 12th International Conference on, 2009, pp. 376-381.
- [32] J. Ashburner, and K. J. Friston, "Rigid body registration," *Human Brain Function*, R. S. J. Frackowiak, K. J. Friston, C. Frith, R. Dolan, K. J. Friston, C. J. Price *et al.*, eds.: Academic Press, 2003.
- [33] S. K. Park, and R. A. Schowengerdt, "Image reconstruction by parametric cubic convolution," *Computer Vision, Graphics, and Image Processing*, vol. 23, no. 3, pp. 258-272, 1983.
- [34] J. X. Ji, H. Pan, and Z.-P. Liang, "Further analysis of interpolation effects in mutual information-based image registration," *Medical Imaging, IEEE Transactions on*, vol. 22, no. 9, pp. 1131-1140, 2003.
- [35] N. Dowson, T. Kadir, and R. Bowden, "Estimating the Joint Statistics of Images Using Nonparametric Windows with Application to Registration Using Mutual Information," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 30, no. 10, pp. 1841-1857, 2008.
- [36] M. Seppa, "Continuous Sampling in Mutual-Information Registration," *Image Processing, IEEE Transactions on*, vol. 17, no. 5, pp. 823-826, 2008.
- [37] L. Xu, "A study of efficiency, accuracy, and robustness in intensity-based rigid image registration," *Computer Science*, University of Waterloo, Waterloo, Canada, 2008.
- [38] J. Inglada, V. Muron, D. Pichard, and T. Feuvrier, "Analysis of Artifacts in Subpixel Remote Sensing Image Registration," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 45, no. 1, pp. 254-264, 2007.
- [39] G. K. Rohde, D. M. Healy, and A. Aldroubi, "Sampling and Reconstruction for Biomedical Image Registration," in *Signals, Systems and Computers, 2007. ACSSC 2007. Conference Record of the Forty-First Asilomar Conference on, 2007*, pp. 220-223.
- [40] G. K. Rohde, A. Aldroubi, and D. M. Healy, "Interpolation Artifacts in Sub-Pixel Image Registration," *Image Processing, IEEE Transactions on*, vol. 18, no. 2, pp. 333-345, 2009.
- [41] G. A. Reitineier, and C. H. Strolle, "Digital image processing techniques for video special effect," in *IEEE International Symposium on Circuits and Systems Conference Proceedings, 1983*, pp. 168—171.
- [42] T. Thong, "Frequency domain analysis of two-pass rotation algorithm," in *Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP '85., 1985*, pp. 1333-1336.
- [43] D. Fraser, and R. A. Schowengerdt, "Avoidance of additional aliasing in multipass image rotations," *Image Processing, IEEE Transactions on*, vol. 3, no. 6, pp. 721-735, 1994.
- [44] C. B. Owen, and F. Makedon, "High quality alias free image rotation," in *Signals, Systems and Computers, 1996. 1996 Conference Record of the Thirtieth Asilomar Conference on, 1996*, pp. 115-119 vol.1.
- [45] W. T. Freeman, and E. H. Adelson, "The design and use of steerable filters," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 13, no. 9, pp. 891-906, 1991.
- [46] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multiscale transforms," *Information Theory, IEEE Transactions on*, vol. 38, no. 2, pp. 587-607, 1992.
- [47] M. Unser, N. Chenouard, and D. Van De Ville, "Steerable Pyramids and Tight Wavelet Frames in L2(Rd)," *IEEE Transactions on Image Processing*, vol. 20, no. 10, pp. 1057-1149, 2011.
- [48] E. Catmull, and A. R. Smith, "3-D transformations of images in scanline order," *SIGGRAPH Comput. Graph.*, vol. 14, no. 3, pp. 279-285, 1980.
- [49] H. Kieseetter, and A. Graf, *Rotation in digital grids and corresponding models*, Zentral Institut fur Kybemetik und Informations Prozesse, 1985.
- [50] A. W. Paeth, "A fast algorithm for general raster rotation," *Graphics gems*, A. S. Glassner, ed.: Academic Press, 1993.
- [51] A. Tanaka, M. Kameyama, S. Kazama, and O. Watanabe, "A rotation method for raster image using skew transformation," in *IEEE Conference on Computer Vision and Pattern Recognition, 1986*, pp. 272-277.
- [52] P. Hanrahan, "Three-pass affine transforms for volume rendering," *SIGGRAPH Comput. Graph.*, vol. 24, no. 5, pp. 71-78, 1990.
- [53] P.-E. Danielsson, and M. Hammerin, "High-accuracy rotation of images," *CVGIP: Graphical Models and Image Processing*, vol. 54, no. 4, pp. 340-344, 1992.
- [54] M. Unser, P. Thevenaz, and L. Yaroslavsky, "Convolution-based interpolation for fast, high-quality rotation of images," *Image Processing, IEEE Transactions on*, vol. 4, no. 10, pp. 1371-1381, 1995.
- [55] H. S. Stone, B. Tao, and M. McGuire, "Analysis of image registration noise due to rotationally dependent aliasing," *Journal of Visual Communication and Image Representation*, vol. 14, no. 2, pp. 114-135, 2003.
- [56] L. C. Maas, and P. F. Renshaw, "Post-registration spatial filtering to reduce noise in functional MRI data sets," *Magnetic Resonance Imaging*, vol. 17, no. 9, pp. 1371-1382, 1999.
- [57] M. Unser, A. Aldroubi, and M. Eden, "B-spline signal processing. I. Theory," *Signal Processing, IEEE Transactions on*, vol. 41, no. 2, pp. 821-833, 1993.
- [58] A. V. Oppenheim, and R. W. Schaffer, *Discrete-time signal processing*: Prentice Hall, 2010.
- [59] S. P. Kim, and W. Y. Su, "Subpixel accuracy image registration by spectrum cancellation," in *Acoustics, Speech, and Signal Processing, 1993. ICASSP-93., 1993 IEEE International Conference on, 1993*, pp. 153-156 vol.5.
- [60] H. S. Stone, M. T. Orchard, C. Ee-Chien, and S. A. Martucci, "A fast direct Fourier-based algorithm for subpixel registration of images," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 39, no. 10, pp. 2235-2243, 2001.
- [61] D. Martindale. "Resampling and Digital Cameras - Using CCDs with 45 degree rotation," http://www.satsignal.eu/imaging/frequency_chart.htm.
- [62] S. Sýkora. "K-space images of n-dimensional spheres and generalized sinc functions," <http://dx.doi.org/10.3247/SL2Math07.002>.
- [63] B. Fornberg, E. Larsson, and G. Wright, "A new class of oscillatory radial basis functions," *Computers and Mathematics with Applications*, vol. 51, no. 8, pp. 1209-1222, 2006.
- [64] P. E. Anuta, "Spatial Registration of Multispectral and Multitemporal Digital Imagery Using Fast Fourier Transform Techniques," *Geoscience Electronics, IEEE Transactions on*, vol. 8, no. 4, pp. 353-368, 1970.
- [65] A. Wade, and F. Fitzke, "A fast, robust pattern recognition asystem for low light level image registration and its application to retinal imaging," *Opt. Express*, vol. 3, pp. 190-197, 1998.
- [66] "BrainWeb: Simulated Brain Database," <http://www.bic.mni.mcgill.ca/brainweb/>.
- [67] D. L. Collins, A. P. Zijdenbos, V. Kollokian, J. G. Sled, N. J. Kabani, C. J. Holmes *et al.*, "Design and construction of a realistic digital brain phantom," *Medical Imaging, IEEE Transactions on*, vol. 17, no. 3, pp. 463-468, 1998.
- [68] B. T. T. Yeo, M. R. Sabuncu, T. Vercauteren, N. Ayache, B. Fischl, and P. Golland, "Spherical Demons: Fast Diffeomorphic Landmark-Free Surface Registration," *Medical Imaging, IEEE Transactions on*, vol. 29, no. 3, pp. 650-668, 2010.
- [69] J. H. Halton, "On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals," *Numerische Mathematik*, vol. 2, no. 1, pp. 84-90, 1960.
- [70] D. S. Marcus, T. H. Wang, J. Parker, J. G. Csernansky, J. C. Morris, and R. L. Buckner, "Open Access Series of Imaging Studies (OASIS): Cross-sectional MRI Data in Young, Middle Aged, Nondemented, and Demented Older Adults," *Journal of Cognitive Neuroscience*, vol. 19, no. 9, pp. 1498-1507, 2007/09/01, 2007.
- [71] http://ridgesandfurrows.homestead.com/fingerprint_patterns.html.



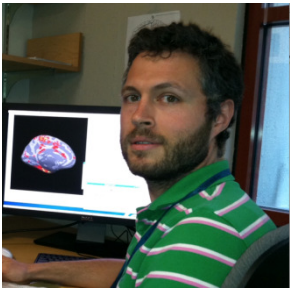
Iman Aganj received his B.S. from École Polytechnique of Paris in 2005, majoring in computer science, and then joined Prof. Guillermo Sapiro's image processing laboratory at the University of Minnesota. He worked on various problems in medical image processing and analysis, focusing on diffusion-weighted MRI, anatomical MRI, and electron microscopy, and obtained his Ph.D. in electrical engineering in 2010. Since 2011, Dr. Aganj has been a Research Fellow at the Martinos Center for Biomedical Imaging, Radiology Department, Massachusetts General

Hospital, Harvard Medical School, and a Research Affiliate of the Laboratory for Information and Decision Systems, Electrical Engineering and Computer Science Department, Massachusetts Institute of Technology. His current research at Prof. Bruce Fischl's Laboratory for Computational Neuroimaging focuses on anatomical and functional MRI image registration.



BT Thomas Yeo received the B.S. and M.S. from Stanford University and Ph.D. from the Massachusetts Institute of Technology in 2010. He is currently a Research Fellow at Duke-NUS Graduate Medical School. Dr. Yeo has worked on various problems in medical image analysis, including image registration and image segmentation. His current research capitalizes on emerging approaches to measure the structural and functional architecture of the human brain. By characterizing the brain systems of healthy human subjects, he seeks to understand how these systems support cognition and how they are disrupted as a result of psychiatric disorders. Dr. Yeo

has received several awards for his work, including the MICCAI Young Scientist award in 2007, MICCAI Young Scientist Runner-up Award in 2008, and the MICCAI Young Investigator Publication Impact Award in 2011.



Mert Sabuncu completed a Ph.D. in electrical engineering at Princeton University in 2006. For his dissertation, Dr. Sabuncu studied image processing and its application to various medical imaging problems. After earning his Ph.D., Dr. Sabuncu moved to Massachusetts Institute of Technology (MIT), where he spent three years as a post-doctoral researcher in Prof. Polina Golland's lab. His work at MIT continued to deal with medical image analysis

problems, and he started working with researchers in the Harvard Medical School (HMS). These collaborations prompted his move to HMS as a faculty member in 2009. Today, Dr. Sabuncu is an Assistant Professor at the Martinos Center for Biomedical Imaging, where he leads a research group that focuses on developing computational algorithms for examining the relationships between medical image data, genotype data, and clinical outcome.



Bruce Fischl's research involves the development of novel techniques for the generation of models of neuroanatomical structures using MRI. This includes the automatic construction of geometrically accurate and topologically correct models of the human cerebral cortex, as well as segmentation of over 35 subcortical and ventricular structures. The cortical models have been useful in a number of domains, including the development of a technique that exploits the correlation between cortical folding patterns and function to generate a more accurate mapping across different brains. This high-dimensional nonlinear registration procedure results in a substantial increase

in statistical power over more standard methods of inter-subject averaging, and allows the automatic labeling of many anatomical features of the cortex. The surface models also yield measures of the thickness of the cortical ribbon, which is of great clinical and research significance as many neurodegenerative diseases result in progressive, regionally specific atrophy of the cortical gray matter.

The confluence of these two techniques – the construction of highly accurate and topologically correct models of the cortex, as well as the capability to generate high-resolution mappings across subjects – allows the comparison of cortical atrophy in patient and control population. Dr. Fischl has been involved in a number of studies utilizing these tools to understand the pattern of cortical thinning in healthy aging, development and disorders such as Huntington's disease, Alzheimer's disease, multiple sclerosis, dyslexia, autism and schizophrenia. His research group freely distributes the tools via the FreeSurfer analysis suite to a community of over 12,000 researchers worldwide, which has resulted in hundreds of publications of clinically and neuroscientifically significant findings.