

Scale Factor Point Spread Function Matching: Beyond Aliasing in Image Resampling

M. Jorge Cardoso, Marc Modat, Tom Vercauteren, and Sebastien Ourselin

Translational Imaging Group, CMIC, University College London, London, UK

Abstract. Imaging devices exploit the Nyquist-Shannon sampling theorem to avoid both aliasing and redundant oversampling by design. Conversely, in medical image resampling, images are considered as continuous functions, are warped by a spatial transformation, and are then sampled on a regular grid. In most cases, the spatial warping changes the frequency characteristics of the continuous function and no special care is taken to ensure that the resampling grid respects the conditions of the sampling theorem. This paper shows that this oversight introduces artefacts, including aliasing, that can lead to important bias in clinical applications. One notable exception to this common practice is when multi-resolution pyramids are constructed, with low-pass "anti-aliasing" filters being applied prior to downsampling. In this work, we illustrate why similar caution is needed when resampling images under general spatial transformations and propose a novel method that is more respectful of the sampling theorem, minimising aliasing and loss of information. We introduce the notion of scale factor point spread function (sfPSF) and employ Gaussian kernels to achieve a computationally tractable resampling scheme that can cope with arbitrary non-linear spatial transformations and grid sizes. Experiments demonstrate significant ($p < 10^{-4}$) technical and clinical implications of the proposed method.

1 Introduction

Image resampling is ubiquitous in medical imaging. Any processing pipeline that requires coordinate mapping, correction for imaging distortions, or simply altering the resolution of an image, needs a representation of the image outside of the digital sampling grid provided by the initial image. The theoretical foundations from the Nyquist-Shannon sampling theorem provide conditions under which perfect continuous reconstruction can be achieved from regularly spaced samples. This theoretical model thus provides a means of interpolating between the discrete samples and underpins the typical resampling procedure used in medical imaging. The Whittaker-Shannon interpolation formula, or sinc interpolation, is optimal for bandlimited signals but the unbounded support of the sinc function calls for approximations. Many interpolating methods, ranging from low order, e.g. piecewise constant and linear interpolation, to high order, e.g. polynomial, piecewise-polynomial (spline) and windowed sinc (Lanczos) methods have thus been developed. Most of these have direct extension to N dimensional data, with different methods being optimal for specific types of signal.

Resampling typically relies on one such interpolation method to represent, on a given (discrete) target sampling grid, a continuous image derived from the input source image. In most cases, the (discrete) resampled image is simply assumed to be a faithful representation of the continuous derived image, with no special care taken to ensure the sampling theorem conditions. Within the context of medical imaging, Meijering *et al.* [1] compared multiple resampling kernels, concluding that the quality of the resampling was proportional to the degree of the kernel, with linear, cubic and Lanczos resampling kernels being the best with a 1st, 2nd and 3rd order neighbourhood respectively. This comparison was done with regular grid aligned data under translation, thus respecting the Nyquist criterion. Thus, their conclusions cannot easily be extrapolated to the general resampling case.

Resampling is known to introduce artefacts in medical imaging. One common example arises within the context of image registration [2,3,4] where a source image is typically resampled to the sampling grid of a target image and a similarity function is used to compare the images based on the discrete samples at the target grid location only, thereby discarding the continuous image representation. Pluim *et al.* [2] ameliorated the problem by using partial volume (PV) sampling, Klein *et al.* [3] by random off grid sampling, and more recently Aganj *et al.* [4] by replacing the cost function summation term with an approximate integral. Another main source of artefacts is the fact that aliasing is introduced when resampling an image to lower resolution. While, in multi-resolution pyramids [5], aliasing is well addressed by applying a Gaussian filter prior to downsampling, the problem in the general resampling case has surprisingly received little attention in the medical imaging community. Unser *et al.* [6] propose a least-squares spline-based formulation that achieves aliasing-free resampling for affine spatial transformations but to the best of our knowledge no solution has been proposed for general non-linear spatial transformations. Inverse approximation-based resampling has also been proposed. It relies on the inverse of the standard spatial transformation to explicitly transport samples from the source image to the target space and then uses a least-squares grid-based spline reconstruction [7]. While this approach would avoid aliasing problems and can deal with non-linear spatial transformations, it can not easily be adapted to take into account the local anisotropy of the transformations and has the drawback of requiring the inverse transformation which is not always available or easy to compute.

In this paper, we propose to address aliasing artefacts in the general context of resampling with non-linear spatial transformations by associating a scale factor point spread function (sfPSF) to the images and formulating the resampling process as an sfPSF matching problem. In the case of an original image from a known imaging device, the sfPSF is simply the PSF of the device. Our approach models local sfPSF transformations as a result of coordinate mapping, providing a unified and scale-aware way to resample images between grids of arbitrary sizes, under arbitrary spatial transformations.

2 Methods and Methodological Considerations

Nominal Scale Factor Point Spread Function (sfPSF). Observed samples (pixels/voxels) in medical images are not direct Dirac impulse measurements of a biological scene, but are typically samples from the convolution of the unobservable biological scene with an observation kernel, i.e. they have an associated point spread function (PSF). This PSF is commonly a property of the acquisition or reconstruction system. For example, 3D MRI images have approximately Gaussian PSFs due to the low-pass filters in the receiver coils, while CT and PET images have spatially variant PSFs dependent on many parameters including the image reconstruction algorithm.

In this work, we propose to associate a PSF to each image, whether it directly arises from an imaging device or is the result of further transformation or processing, and refer to this as the scale factor PSF (sfPSF) by analogy with the scale factor used in scale space theory [5]. In the case where the image directly arises from an imaging device with known PSF, the sfPSF is simply taken as the closest Gaussian kernel to the PSF of the device. When the actual PSF associated with a given image is unknown, the nominal sfPSF is taken as a Gaussian \mathcal{G}_Σ , where the covariance matrix Σ is diagonal and the full-width-half-maximum (FWHM) of the Gaussian PSF is classically and without loss of generality matched to the voxel size $D = \{D_x, D_y, D_z\}$: $\Sigma = (2\sqrt{2\ln 2})^{-2} \text{diag}(D_x^2, D_y^2, D_z^2)$. Similarly, when resampling to a target grid not already associated with an image, we associate the same voxel-size matched nominal sfPSF with the grid. The relevance of the sfPSF will become clear in the following section. Intuitively, when an image with an associated sfPSF is smoothed, the resulting smoothed image will be associated with a larger sfPSF. Conversely, when an image is upsampled, the resulting image will be associated with the same sfPSF as the source image if expressed in world (rather than pixel) coordinates.

Compensating for Aliasing When Downsampling. In medical images, frequencies commonly approach the Nyquist limit of their representation (e.g. full k-space sampling in MRI). If one affinely transforms an image onto its own grid with an affine transformation which has a determinant greater than 1, i.e. spatially compressing the samples, or transforms an image into a lower resolution grid, then the Nyquist criterion is not satisfied with typical interpolation. This sub-Nyquist sampling introduces frequency aliasing and loss of information.

In scale-space and pyramidal approaches, when downsampling an image by a factor of $k = 2$, the original image is typically pre-convolved with a Gaussian filter with variance $\sigma_P = 0.7355$. Applying the notion of sfPSF and the notations introduced above, we specify \mathcal{G}_{Σ_S} and \mathcal{G}_{Σ_T} as the source image and target grid nominal sfPSFs respectively. Following the scale space methodology [5], sfPSF matching consists of finding the best scale-matching PSF \mathcal{G}_{Σ_P} that, when combined with the source sfPSF, provides the best approximation of the target sfPSF. If the sfPSFs are Gaussian, then the covariance of \mathcal{G}_{Σ_P} can be obtained through simple covariance addition $\Sigma_T = \Sigma_S + \Sigma_P$. In the downsampling example above, without additional knowledge, we have $\text{FWHM}_T = k = 2$

and $\text{FWHM}_S = 1$. We therefore get $\Sigma_P = \Sigma_T - \Sigma_S = (k^2 - 1^2)(2\sqrt{2\log 2})^{-2} \text{Id}$ or, in this case, $\sigma_P = 0.7355$. This is equivalent to finding the best Fourier domain filter that limits the bandwidth of the source image frequencies to the representable Nyquist frequencies on the target grid. Similar filtering would be necessary to avoid aliasing when rigidly resampling high resolution images to a lower resolution, e.g. high resolution anatomical MRI and CT images or binary segmentations are transformed to lower resolution MRI (2D FLAIR), metabolic (PET), microstructural (DWI) and functional (fMRI) images. In this work we propose to compensate for the differences in sfPSF when resampling to arbitrary grid sizes under global (affine) or local (non-linear) spatial transformations.

Before presenting our methodological contribution in detail, we illustrate the sfPSF performance in comparison with standard interpolation on a synthetic phantom (Fig. 1) with a low frequency object (square) overlaid on a high-frequency pattern (to highlight aliasing). This phantom is resampled to a grid with 3 times less resolution *per axis*. Note that the proposed sfPSF method appropriately integrates out high frequencies without introducing aliasing.

Pre-Convolution sfPSF Matching. In the restricted affine transformation \mathcal{A} scenario going from the space of T to S , the source image can be pre-convolved with an *anti-aliasing* filter with covariance such that $\Sigma_P^S + \Sigma_S = \mathcal{A} \cdot \Sigma_T \cdot \mathcal{A}^\top$, where we now emphasize that Σ_P^S is defined in the space of S , and then resampled using classical interpolation.

In the general non-linear spatial transformation case, a similar but spatially-variant pre-filtering procedure could be applied by relying on a local affine approximation of the spatial transformation, i.e. using the jacobian matrix. This solution might provide useful results but because the sfPSF matching, or smoothing, is performed before the spatial transformation, and the sfPSF potentially has a large spatial extent, it can be seen as disregarding the non-homogeneous

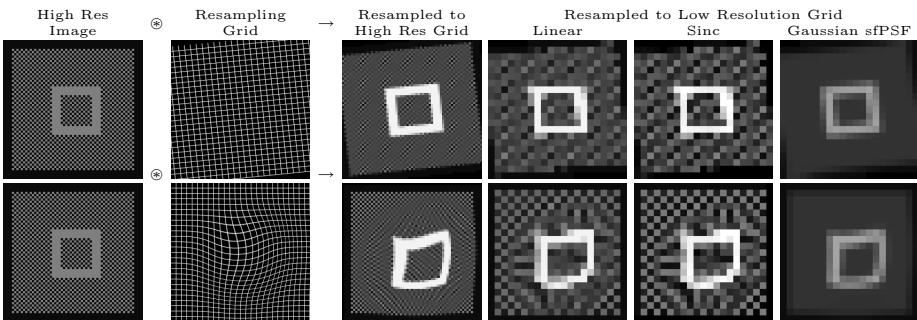


Fig. 1. Left to right: a high resolution (HR) image (63×63 voxels at $1 \times 1\text{mm}$ voxel size) with a square on a high frequency pattern; the resampling grid (transformation); the HR image resampled to the original HR grid and to a lower resolution (21×21 voxels at $3 \times 3\text{mm}$ voxel size) grid using linear, sinc and a Gaussian sfPSF (proposed) interpolation strategy. Note the aliasing in the linear and sinc examples.

distribution of the samples in target space T . Instead, we propose to model and perform the sfPSF matching in the space of T and thus rely on Σ_P^T .

Target Space sfPSF Matching. Let $\mathcal{F}_{S\leftarrow T}(\mathbf{v})$ be a spatial transformation mapping the location $\mathbf{v} = (v_x, v_y, v_z)$ from the space of T to that of S , given either an affine or non-linear transformation. Let $\mathcal{A}_{S\leftarrow T}(\mathbf{v})$ be the jacobian matrix of \mathcal{F} at \mathbf{v} which provides the best linear approximation of \mathcal{F} at \mathbf{v} . Following our previous derivations we have $\Sigma_P^S(\mathbf{v}) + \Sigma_S \approx \mathcal{A}_{S\leftarrow T}(\mathbf{v}) \cdot \Sigma_T \cdot \mathcal{A}_{S\leftarrow T}^T(\mathbf{v})$. Thus, given $\Sigma_P^S(\mathbf{v}) = \mathcal{A}_{S\leftarrow T}(\mathbf{v}) \cdot \Sigma_P^T(\mathbf{v}) \cdot \mathcal{A}_{S\leftarrow T}^T(\mathbf{v})$ we now want $\Sigma_P^T(\mathbf{v}) + \mathcal{A}_{S\leftarrow T}^{-1}(\mathbf{v}) \cdot \Sigma_S \cdot \mathcal{A}_{S\leftarrow T}^T(\mathbf{v}) \approx \Sigma_T$. In other words, we want to find a *symmetric positive semi-definite* covariance $\Sigma_P^T(\mathbf{v})$ that, when combined with the affinely transformed covariance $\Sigma_S^T(\mathbf{v}) = \mathcal{A}_{S\leftarrow T}^{-1}(\mathbf{v}) \cdot \Sigma_S \cdot \mathcal{A}_{S\leftarrow T}^T(\mathbf{v})$, best approximates Σ_T . Under the typical assumption that the nominal sfPSF Σ_T is diagonal, this approximation is given by $\Sigma_P^T(\mathbf{v}) = \max(\Sigma_T - \lambda(\mathbf{v}), 0)$ with $\lambda(\mathbf{v})$ being a diagonal scaling matrix containing the components of $\Sigma_S^T(\mathbf{v})$ obtained through polar decomposition and $\max(\cdot, \cdot)$ is the element-wise maximum operator between two matrices. It is important to note that under regimes where the Nyquist limit is not violated, e.g. upsampling, rigid transformation between isotropic grids, or non-linear transformations with the Jacobian determinant below 1, then the proposed method leads to a Dirac in the sfPSF matching meaning that our procedure reverts to standard resampling.

Interpolation by Convolution with the sfPSF. Let $I_T(\mathbf{v})$ be the sought resampled intensity in the space of T at location \mathbf{v} , and $I_S(\mathbf{v})$ be the intensity of the source image. The interpolated and PSF matched value of I_T at location (\mathbf{v}) can be obtained by an *oversampled* discretised convolution

$$I_T(\mathbf{v}) = \frac{1}{Z} \sum_{v_x}^{N_x(\mathbf{v})} \sum_{v_y}^{N_y(\mathbf{v})} \sum_{v_z}^{N_z(\mathbf{v})} I_S(\mathcal{F}_{S\leftarrow T}(\mathbf{v} - v)) \mathcal{G}_{\Sigma_P^T(\mathbf{v})}(v)$$

where $\mathcal{G}_{\Sigma_P^T(\mathbf{v})}(v) = (2\pi)^{-\frac{3}{2}} |\Sigma_P^T(\mathbf{v})|^{-\frac{1}{2}} e^{-\frac{1}{2} v^T (\Sigma_P^T(\mathbf{v}))^{-1} v}$, Z is a normaliser that ensures the (discrete) sum over $\mathcal{G}_{\Sigma_P^T(\mathbf{v})}(v)$ equals 1. $N_x(\mathbf{v})$, $N_y(\mathbf{v})$ and $N_z(\mathbf{v})$ are sampling regions centered at \mathbf{v} in the x , y and z directions respectively. More specifically, $N_x(\mathbf{v})$ are homogeneously spaced samples between -3 and 3 standard deviations, i.e. $N_x(\mathbf{v}) = \{k\sigma_{P_x}(\mathbf{v}) \ \forall k \in [-3, 3]\}$, with $\sigma_{P_x}^2(\mathbf{v}) = \Sigma_P^T(\mathbf{v})^{xx}$ and equivalently for $N_y(\mathbf{v})$, $N_z(\mathbf{v})$, $\sigma_{P_y}(\mathbf{v})$ and $\sigma_{P_z}(\mathbf{v})$. As S is a discrete image, an appropriate interpolation method, such as linear, cubic or sinc interpolation is used to get the sample values at $\mathcal{F}_{S\leftarrow T}(\mathbf{v} - v)$. Similarly, interpolation is required to compute $\mathcal{F}_{S\leftarrow T}(\mathbf{v} - v)$ if the spatial transformation is represented as a displacement field. Note also that when σ_{P_x} (resp. σ_{P_y} or σ_{P_z}) becomes 0 or very small, care has to be taken to appropriately compute the limit of the Gaussian weight by resorting to an axis aligned Dirac function.

3 Validation and Discussion

Data. 30 healthy control subjects were obtained from the ADNI2 database. All subjects had an associated T1-weighted MRI image, acquired at $1.1 \times 1 \times 1 \text{mm}$ voxel size, and a ^{18}F -FDG PET image, reconstructed at $3 \times 3 \times 3 \text{mm}$. For this selected subset of ADNI2, all data (MRI and PET) were acquired on the same scanner with the same scanning parameters, thus removing acquisition confounds. While the effective PSF of PET images can be between 3 and 6mm FWHM, in these experiments we will assume a PSF with 3mm FWHM.

Frequency Domain Analysis Under Non-linear Transformation. Frequency domain analysis was used to assess whether resampling images under non-linear transformation respects the Nyquist limit and to test whether the proposed method can mitigate any error. The 30 T1 MRIs were zero padded by a factor of 2 in the frequency domain ($0.55 \times 0.5 \times 0.5 \text{mm}$ voxel size), doubling the representable frequencies. Fig. 2 shows two upsampled images and their frequency magnitude. Note the empty power spectrum outside the Nyquist band (white box). 29 upsampled T1 images were affinely [8] and then non-rigidly [9] registered to the remaining image using standard algorithms. Each upsampled image was resampled to the space of the remaining image using both a three-lobed truncated sinc resampling, and the proposed Gaussian sPSF method. The Gaussian sPSF was set to $1.1 \times 1 \times 1 \text{mm}$ FWHM for both Σ_S and Σ_T . An example of an image resampled using both methods is shown in Fig. 2. Note that, as expected, supra-Nyquist frequencies are created when using sinc interpolation. These frequencies are greatly suppressed when using the proposed method. Analysis of the power spectra of the 29 resampled images showed an average power suppression of 94.4% for frequencies above the Nyquist band when using the Gaussian PSF instead of sinc interpolation. This power suppression results in an equivalent reduction of aliasing at the original resolution.




Fig. 2. Left to right: (Top) upsampled source S and target T images followed by S non-linearly resampled to T using sinc and Gaussian sPSF; (Bottom) their respective frequency domain log magnitude. The white box represents the Nyquist limit before zero-padding. Note the high supra-Nyquist magnitude after sinc resampling.




Fig. 3. Left to right: A segmented region overlaid on a high resolution T1 ($1.1 \times 1 \times 1 \text{mm}$) image (and zoomed), followed by the same segmentation resampled to the resolution of a PET image ($3 \times 3 \times 3 \text{mm}$) using linear and the proposed Gaussian sfPSF. The zoomed segmentations are in gray scale between 0 and 1. Given that the cortex is commonly thinner than 3mm , note the unrealistic large amount of voxels with segmentation probability equal to 1 at PET resolution when using linear resampling.

Volume Preservation and Partial Volume under Rigid Resampling. To demonstrate clinically relevant consequences of aliasing, we tested the effect of resampling within the context of segmentation propagation and partial volume estimation. Specifically, T1 images are segmented into 98 different regions using multi-atlas label propagation and fusion [10] based on the Neuromorphometrics, Inc. labels. The T1 images were then rigidly registered [8] to the PET data, and each one of the segmented regions was resampled to the PET data using both linear interpolation and the Gaussian sfPDF. Linear was chosen here due to the signal-limited $[0, 1]$ nature of probabilities. Example results are shown in Fig. 3.

As all the segmentations are probabilistic, volume was estimated as the sum of all probabilities for all voxels \mathbf{v} times the voxel size \mathcal{V} . The volume was estimated in the original T1 space (V_{HR}), and in the PET space after linear (V_{TRI}) and the Gaussian sfPSF (V_{sfPSF}) resampling. Similar volumes were obtained when the probabilistic segmentations were thresholded at 0.5. The relative volume difference $RVD = (V_{sfPSF} - V_{HR})/V_{HR}$ and its absolute value, $ARVD = |RVD|$, were estimated for each of the 30 subjects, 98 cortical regions and 2 resampling methods. The RVD and ARVD were averaged over all subjects for each region, resulting in 98 mean RVD and ARVD values per region, per method. These values are plotted in Fig. 4. A Parzen window mean and STD are also plotted in the




Fig. 4. The mean RVD (Left) and mean ARVD (Right) between a low and high resolution representation of a segmented region averaged over the population. 100 different regions are plotted against the mean volume of the region V_{HR} over the population.

same figure. Both resampling methods are unbiased according to the mean RVD. However, the Gaussian sfPSF method provides significantly lower ($p < 10^{-4}$ - Wilcoxon signed-rank test) errors in terms of mean ARVD for all regions, especially in smaller regions where linear resampling can introduce up to 8% mean absolute volume difference. A 8% mean difference implies even larger local errors, resulting in detrimental effects in partial volume/ compartment modelling, and when estimating biomarkers that rely on structural segmentations.

4 Conclusion

The presented work explores aliasing in medical image resampling and proposes a new interpolation technique that matches the scale factor PSF given arbitrary grid sizes and non-linear transformations. We demonstrate the advantages of using Gaussian sfPSF resampling, both in terms of Nyquist limit and volume preservation, when compared with common resampling techniques. Future work will involve deploying the proposed methodology within image registration algorithms and verifying the impact of the sfPSF in partial volume and compartment modelling. An implementation will be made available at the time of publication.

References

1. Meijering, E.H.W., Niessen, W.J., Pluim, J.P.W., Viergever, M.A.: Quantitative comparison of sinc-approximating kernels for medical image interpolation. In: Taylor, C., Colchester, A. (eds.) MICCAI 1999. LNCS, vol. 1679, pp. 210–217. Springer, Heidelberg (1999)
2. Pluim, J.P., Maintz, J.B.A., Viergever, M.A.: Mutual information matching and interpolation artifacts. In: SPIE Medical Imaging, vol. 3661, pp. 56–65. SPIE Press, May 1999
3. Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P.W.: ElastiX: A toolbox for intensity-based medical image registration. *IEEE Trans. Med. Imag.* 29(1), 196–205 (2010)
4. Aganj, I., Yeo, B.T.T., Sabuncu, M.R., Fischl, B.: On removing interpolation and resampling artifacts in rigid image registration. *IEEE Trans. Image Process.* 22(2), 816–827 (2013)
5. Lindeberg, T.: Scale-space for discrete signals. *IEEE Trans. Pattern Anal. Mach. Intell.* 12(3), 234–254 (1990)
6. Unser, M., Neimark, M., Lee, C.: Affine transformations of images: a least squares formulation. In: Proc. ICIP 1994, vol. 3, pp. 558–561 (1994)
7. Arigovindan, M., Sühling, M., Hunziker, P., Unser, M.: Variational image reconstruction from arbitrarily spaced samples: A fast multiresolution spline solution. *IEEE Trans. Image Process.* 14(4), 450–460 (2005)
8. Modat, M., Cash, D.M., Daga, P., Winston, G.P., Duncan, J.S., Ourselin, S.: Global image registration using a symmetric block-matching approach. *Journal of Medical Imaging* 1(2), 024003–024003 (2014)

9. Modat, M., Ridgway, G.R., Taylor, Z.A., Lehmann, M., Barnes, J., Hawkes, D.J., Fox, N.C., Ourselin, S.: Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine* 98(3), 278–284 (2010)
10. Cardoso, M.J., Wolz, R., Modat, M., Fox, N.C., Rueckert, D., Ourselin, S.: Geodesic information flows. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *MICCAI 2012, Part II*. LNCS, vol. 7511, pp. 262–270. Springer, Heidelberg (2012)