Improved Noise Reduction with NORDIC PCA for Submillimetre BOLD fMRI via Non-Local Patch Formation using Voxel Similarity Matching 'Voxel-Matching (VM) NORDIC'

Alessandro Nigi[1]*; Email:* a.nigi@students.uu.nl*; Student number:* 6358292
[1]*Department of Radiology, Center for Image Sciences, Vascular Brain Imaging Group (PI Jeroen C.W. Siero) University Medical Center Utrecht, Utrecht, the Netherlands*
Examiner: Dr. ir. Jeroen C.W. Siero *Second Reviewer: Dr. Natalia Petridou*

ABSTRACT

Submillimeter functional magnetic resonance imaging (fMRI) based on blood-oxygenation-level-dependent (BOLD) signal enables the study of brain function at the submillimeter level, uncovering insights into fine-scale organizations like cortical layers and columns. However, its inherently low contrast-to-noise ratio (CNR) and signal-to-noise ratio (SNR) often limit its reliability and applicability. Noise Reduction with Distribution Corrected Principal Components Analysis (NORDIC PCA) is a locally low-rank denoising algorithm that reduces thermal noise levels in BOLD fMRI in a local patch manner. However, local patches often contain a mixture of signals from multiple tissues that negatively affects the low-rank structure of the patches, which limits the denoising capabilities of the algorithm. We propose an alternative approach for patching formation by gathering similar non-local voxels, dubbed voxel-matching (VM) NORDIC. The results on submillimeter resolution BOLD fMRI data indicate that VM-NORDIC effectively promotes the low rankness of the patches by boosting signal redundancy, allowing for more efficient noise attenuation. Moreover, the method barely affects spatial smoothness due to the non-local voxel selection. In particular, VM-NORDIC outperforms NORDIC with default local patching (Standard-NORDIC) in terms of temporal SNR (tSNR) (~9-90% larger than Standard-NORDIC; ~23-250% than the original) and spatial smoothness estimates (~20% of the smoothness induced by Standard-NORDIC). These improvements are fundamental to improve the validity and precision of fMRI studies at submillimeter resolutions.

INTRODUCTION

Functional magnetic resonance imaging (fMRI)[1–3] based on blood oxygenation level-dependent (BOLD) contrast is indispensable for depicting brain activity and functional connectivity. The development of ultra-high field (UHF) MRI systems (>= 7T) allowed pushing the spatial resolution to the sub-millimetre scale by increasing the strength of the detectable signal[4,5]. High-resolution fMRI is especially essential in studying fine-scale structures at the mesoscopic level, like cortical columns and layers[4,6–8].

Nevertheless, fMRI short acquisition times, the small voxel sizes enabled by UHF, and the inherently small BOLD responses (~0.5-3%) lead to a low contrast-to-noise ratio (CNR) and signal-to-noise ratio (SNR), limiting the applicability and reliability of submillimetre fMRI[5,9–11]. These drawbacks become prominent at sub-millimetre resolution regimes, where the detectable MR signal is relatively weak, making noise effects more dominant[9,10,12–14].

The prevalent type of noise in high-resolution fMRI is thermal noise, an i.i.d. zero-mean Gaussian distributed noise (e.g. white noise) arising from the random fluctuations in the electrical resistance of the detectors or the magnetic field strength due to the thermal energy of the atomic nuclei in the body[15,16]. Thermal noise is practically and theoretically different from physiological noise. The influence of physiological noise decreases with increasing image resolution and derives from physiological phenomena like the heartbeat and respiration, which occur periodically over time[17–20]. As such, it falls within the category of structured, non-white noise[21–24], and is the target of dedicated techniques such as independent components analysis (ICA) and its applications (e.g. ICA-AROMA)[25]. On the other hand, the influence of thermal noise *increases* at higher resolutions, shorter TRs and lower magnetic fields[12,13]. In particular, it becomes dominant over physiological noise at ~0.8mm isotropic voxel size, but it can still affect the quality of lower-resolution scans[12,13]. Also, the application of parallel imaging to accelerate fMRI data acquisition introduces a non-uniform spatial amplification of thermal noise according to the geometry of the receive coils (g-factor)[26]. Overall, thermal noise lowers CNR, image quality and temporal SNR (tSNR) by increasing the signal variance[9]. Specifically, in BOLD fMRI, functional activity is identified via subtle changes in the voxels signal. Therefore, the random signal fluctuations associated with thermal noise interfere with the detection of the underlying signals of interest related to brain activation, leading to false positives and false negatives[9]. Thermal noise removal is therefore crucial in fMRI analysis to ensure clinically and research-wise high-quality data that lead to accurate, reliable results.

A way to overcome the penalties due to thermal noise at high resolutions is to increase the static magnetic field strength even further[4,14]. Nevertheless, hardware costs

and high field-related artefacts strictly limit this strategy[4]. An alternative is to decrease the noise level with appropriate data processing techniques. Spatially blurring the dataset with a smoothing filter is one popular method to increase the SNR[27]. Smoothing is fast and easily applicable, but it comes at the expense of lowering spatial specificity, which is undesirable if the goal is to depict the activity of fine-scale structures with high precision[9,27]. Temporal averaging is another approach often applied to cancel out random signal fluctuations. However, it also degrades spatial and temporal precision.

More recent denoising techniques use Principal Component Analysis (PCA) to identify and nullify the data components indistinguishable from thermal noise by exploiting signal redundancy across volumes. The repetitive acquisitions inherent to diffusion MRI (dMRI) or fMRI provide the data with explicit signal redundancy. That is, the multiple volumes report the same underlying biological environment, allowing the dataset to have a locally low-rank (LLR) structure[11,28,29]. A locally low-rank structure means that in small patches of an image, just a few principal components contain most of the signal-related variance. This low rankness allows approximating the intensities of the voxels by a smaller amount of signal principal components to preserve the details and structure of the original image while removing noise[11].

In 2016, Veraart et al. introduced Marchenko-Pastur PCA (MPPCA) to denoise dMRI. MPPCA automatically estimates a threshold to suppress thermal noise principal components by exploiting the asymptotic properties of the eigenspectum of local data matrices corrupted by thermal noise. Such spectra follow the well-known Marchenko-Pastur (MP) distribution[30], whose bounds depend on the variance of the data and is used in MPPCA as the threshold to nullify the noise components[28,31]. MPPCA processes the datasets patch-wise assuming the noise level to be constant within the patch. However, accelerated data do not always meet this assumption, potentially leading to an incorrect estimation of the noise variance[26]. Further, the bounds of the MP distribution are well defined only in the asymptotic limit (i.e. with a large amount of data), which is often an unrealistic condition with a finite amount of data[10,31]. The prevalent consequences of these drawbacks are an incorrect truncation of the singular values and the spatial smoothing of the data due to an excessive dimensionality reduction of the patches, which decreases the characteristic differences between the time series of adjacent voxels[29]. Smoothing is highly unwanted in sophisticated denoising methods such as PCA-based techniques as it can lead to erroneous results and biased statistical analyses. Furthermore, as unwanted smoothing can seemingly increase SNR at the expense of spatial integrity, one could simply spatially smooth the data rather than applying advanced PCA-based denoising techniques.

Moeller et al. partly overcame the MPPCA limitations by introducing NOise Reduction with Distribution Corrected (NORDIC) PCA for dMRI[29]. The main difference between the two methods is that NORDIC estimates the noise variance either from a noise scan or the complex data via MPPCA and applies it to compute the noise threshold via Monte Carlo Simulations of the Gaussian noise in the data[9,29]. Moreover, NORDIC uses MPPCA to compute local noise levels to spatially flatten the noise due to the g-factor and ensure constant noise levels within any patch. The noise flattening is eventually reversed after denoising to re-establish the original signal spatial variation. The authors reported that NORDIC surpassed MPPCA in terms of the degree of noise reduction and preservation of detail on dMRI data[29].

The NORDIC authors argued that fMRI can also benefit from LLR denoising thanks to the high redundancy of the data[9,29]. In this case, the method aims at denoising the time courses of the voxels. Several studies reported that NORDIC effectively removed the noise in the fMRI data while preserving the temporal structure of the underlying brain activity, leading to improved results in fMRI-based analyses, such as brain activation mapping and connectivity analysis[9,29,32]. Especially, Vizioli et al. showed that NORDIC outperforms MPPCA on fMRI data regarding tSNR levels, spatial smoothness and functional activation maps on data with a wide range of resolutions, scanning parameters and tasks[9].

However, local patches may represent a downside for NORDIC as they often contain heterogeneous signals from multiple tissues that can contaminate signal redundancy, degrading the low-rank structure of the data[33]. Another disadvantage is that, as for MPPCA, NORDIC slightly smoothens the data due to its LLR approach. Recently, Zhao et al. tackled these problems by developing a *non-local* low-rank (NLLR) PCA denoising method for dMRI. The proposed method uses matrices of 3D (2D spatial + 1D diffusion direction) similar non-local patches[33–35]. They argued that gathering similar patches into a single matrix better exploits signal redundancy and promotes the low-rankness of the dMRI data. The authors reported that their NNLR-based patching method led to a more effective dMRI denoising than MPPCA[33].

As mentioned, NORDIC for fMRI also relies on signal redundancy across volumes to define principal components and thus could benefit from non-local patch formation[9,29]. In the present work, we determine the denoising performance of an alternative patching method for NORDIC that uses non-local similar voxels

to construct low-rank temporal patches (voxel-matching, VM; dubbed VM-NORDIC). We hypothesize that the higher homogeneity of the resulting patches compared to the default local patches (Standard-NORDIC) leads to a more robust estimation of the low-rank structure of the dataset after PCA. Gathering voxels with similar time series boosts the degree of redundancy within the patch, amplifying the principal components related to the signal. Additionally, denoising together voxels from different locations in the dataset avoids spatial over-smoothing. As such, we expect VM-NORDIC to produce a higher noise attenuation and greater preservation of spatial detail than Standard-NORDIC while applying the same parameter-free threshold on high-resolution fMRI data.

We assess and compare the performance of VM-NORDIC with Standard-NORDIC on a group of sub-millimetre resting state BOLD fMRI datasets at 7 tesla in terms of fundamental metrics for fMRI, including tSNR and global spatial smoothness.

THEORY

PCA and SVD

Let $\mathbf{Y}$ be an $M\times Q$ complex-valued volumetric fMRI measurement Casorati matrix, with the columns representing M voxels and the rows representing Q MR signal samples for each voxel (e.g. time points of the BOLD data). We expect the fMRI data to have a low-rank representation due to the multiple volumes acquired through time. That is, Q is large enough to represent $\mathbf{Y}$ with a combination of a few linearly independent sources, or principal components $\mathbf{P}<<\mathbf{R}$, with $\mathbf{R} = \text{rank}(\mathbf{Y}) = \min\{M, Q\} >> 1$ and $\mathbf{P} = \text{rank}(\mathbf{Y}_L)$ where $\mathbf{Y}_L$ is the low-rank representation of $\mathbf{Y}$. NORDIC estimates the principal components via singular value decomposition (SVD) of the data $\mathbf{Y} = \mathbf{U}\cdot\mathbf{S}\cdot\mathbf{V}^H$, where $\mathbf{U}$ is a unitary matrix whose columns are the left singular vectors of $\mathbf{Y}$ and contains information about the spatial structure of the signal. $\mathbf{V}^H$ is the Hermitian unitary matrix whose columns are the right singular vectors of $\mathbf{Y}$, representing the temporal structure of the signal. The diagonal elements $S_{1,1}, \dots, S_{Q,Q}$ of the matrix $\mathbf{S}$ (size M×Q) are the singular values that represent the contribution of each source of signal or noise. Here, $S^2(j) = \lambda(j)$, $j \in \{1, \dots, Q\}$, is the j^{th} eigenvalue of the $Q\times Q$ sample covariance matrix $\mathbf{\Sigma} = \text{I}/Q\cdot\mathbf{Y}\mathbf{Y}^T = \mathbf{U}\cdot\mathbf{S}^2\cdot\mathbf{U}^T$, where $\mathbf{U}$ is the transposed orthogonal matrix whose columns are the principal components.

For *noise-free* measurement data, $\mathbf{Y}$ can be effectively low rank approximated by $\mathbf{P}<<\mathbf{R}$ singular values as the remaining $\mathbf{R}-\mathbf{P}$ singular values are zero[28,31]. However, in a realistic case, $\mathbf{Y}$ is a *noisy* data matrix described by the model $\mathbf{Y} = \mathbf{X} + \mathbf{N}$, where $\mathbf{X}$ is the signal of interest and $\mathbf{N}$ is additive i.i.d. zero-mean Gaussian noise and which propagates through *all* components making all $\mathbf{R}$ eigenvalues non-zero, so that $\text{rank}(\mathbf{Y}) = \mathbf{R}$.

Although the noise in each measurement is random, its effect on the spectrum of singular values becomes deterministic in the limit of $\mathbf{R}>>\mathbf{P}$. In other words, if a noisy fMRI data matrix contains enough time points Q, the histogram of the noise-only components will follow the Marchenko-Pastur (MP) distribution, a well-understood asymptotic universal law from random matrix theory for random covariance matrices[30]. The objective of NORDIC PCA is to estimate and eliminate the components within the MP distribution of a noisy dataset to attenuate thermal noise contributions. This result is achieved by numerically estimating the boundary between noise and signal singular values of a noisy data matrix and applying it as threshold λ_{thr} to nullify all the ordered singular values $\lambda(j)<\lambda_{\text{thr}}$ within the MP distribution (i.e. to cancel out all eigenvalues indistinguishable from zero-mean Gaussian distributed noise). The truncated singular value matrix $\mathbf{S}_{\lambda thr}$ is then recombined with the original eigenvectors matrices to reconstruct the low-rank denoised matrix $\mathbf{Y}_L = \mathbf{U}\cdot\mathbf{S}_{\lambda thr}\cdot\mathbf{V}^H$.

Locally Low-Rank (LLR) model

A locally low-rank model is a mathematical framework used to approximate a high-dimensional data set with a series of low-rank matrices[36]. This type of model allows for a more efficient and compact representation of the data and can be used for tasks such as dimensionality reduction, data compression, and denoising. The method of locally low-rank approximation typically involves breaking the matrix (here, the fMRI dataset) into smaller submatrices (local patches) and applying SVD to each submatrix[36]. The locally low-rank property allows approximating each local region in the data by a low-dimensional subspace of fewer principal components. Of these components, some are dominated by the signal and others by noise - this separability is the basic principle of PCA-based denoising. Standard-NORDIC exploits the LLR properties of fMRI data to reduce the noise level via hard thresholding the ordered singular values of local patches[36].

Degree of noise removal

In VM-NORDIC, the threshold estimation and SVT are the same as for Standard-NORDIC[7]. Nevertheless, the same threshold can be more effective at attenuating noise from a highly homogeneous patch rather than an inhomogeneous one. A higher homogeneity leads to a more defined low-rank structure of the data by making signal components more defined, which together allow

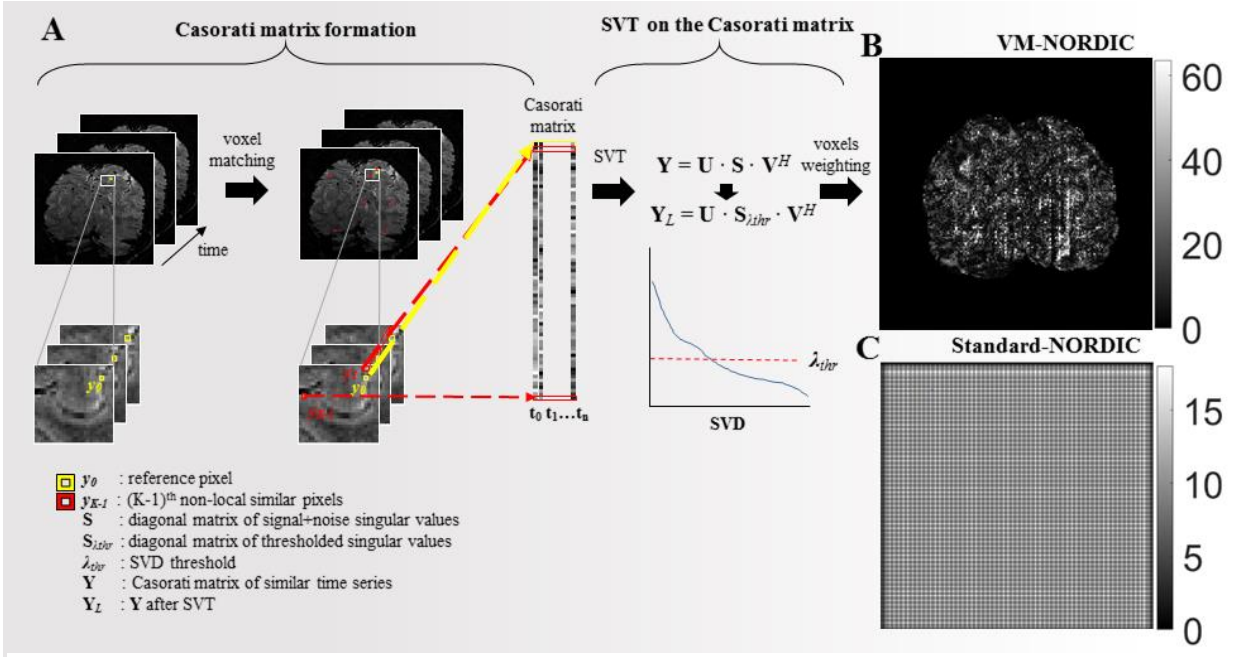

Fig. 1. A) Flowchart of VM-NORDIC. For each chunk, *N* reference voxels y_0 are compared to all the voxels y_m in the chunk. For each reference voxel, the *K*-1 voxels with the highest similarity scores are grouped together with the reference voxel into a *K*×*Q* Casorati matrix **Y**. The matrix undergoes SVT to produce a denoised matrix $\mathbf{Y}_L$. Panels **B** and **C** show the weighting matrices and indicate the amount of time each voxel has been denoised for VM-NORDIC and Standard-NORDIC. The two panels also highlight the main differences in the patching approach between the two methods.

approximating the data with a lower number of principal components. Better-defined principal components also diminish the risk of accidentally canceling signal components and leaving noise components intact. However, as for Standard-NORDIC, noise propagates through all components of the spectrum of singular values[9,31]. Therefore, VM-NORDIC may lead to a higher noise attenuation than Standard-NORDIC but does not completely removes it from the data.

METHODS AND MATERIALS

In the following sections, we first explain the implementation details of VM-NORDIC and highlight its differences from Standard-NORDIC. Then, we move to the methods we applied to acquire the data and analyse and compare the results from the two methods.

Standard-NORDIC LLR model

Standard-NORDIC relies on the LLR properties of fMRI data to reduce the noise level via hard thresholding the ordered singular values of local and spatially overlapping temporal patches. Neighbouring voxels within a fixed $k_1 \times k_2 \times k_3$, usually with $k_1 = k_2 = k_3$, sliding window are vectorized as $\mathbf{y}_t$ and gathered together to form a set of complex $M \times Q$ Casorati matrices of the form $\mathbf{Y} = [\mathbf{y}_1, \dots \mathbf{y}_t, \dots \mathbf{y}_Q]$, where Q is the length of the time series and $M = k_1 \times k_2 \times k_3 \sim Q \cdot 11$, as reported by Veraart et al (see Fig. 2A for an exemplar local patch) [9,31]. Each Casorati matrix undergoes SVT to truncate its singular values spectrum according to a parameter-free threshold λ_{thr}[9,29]. The result is a low-rank representation of the original Casorati matrix with a lower noise level. Finally, each matrix is reshaped into a 4D patch and relocated to its original spatial location.

Proposed Voxel-Matching NORDIC

In VM-NORDIC, patch formation occurs via grouping similar time series of non-local voxels. First, the dataset is divided into chunks of *X* slices. Chunking the dataset is necessary to speed up the denoising process since processing all voxels at once may be too computationally heavy for most machines. For each chunk, VM-NORDIC computes the Euclidian distance between the time series of *N* reference voxels and all the brain voxels in the chunk to assess their level of similarity. For each reference voxel, the *K*-1 most similar voxels are vectorised and grouped with their reference voxel to form *N* Casorati matrices of *K*×*Q* elements (Figure 1A), where *K* is the (optimal) patch size as explained in the next section (see Figure 2B for an exemplar non-local patch). Each of the *N* Casorati matrices undergoes SVT and is then transformed back into image space.

Patch size optimization

In VM-NORDIC, the patch size represents the number of vectorized time series per patch. Optimizing the patch

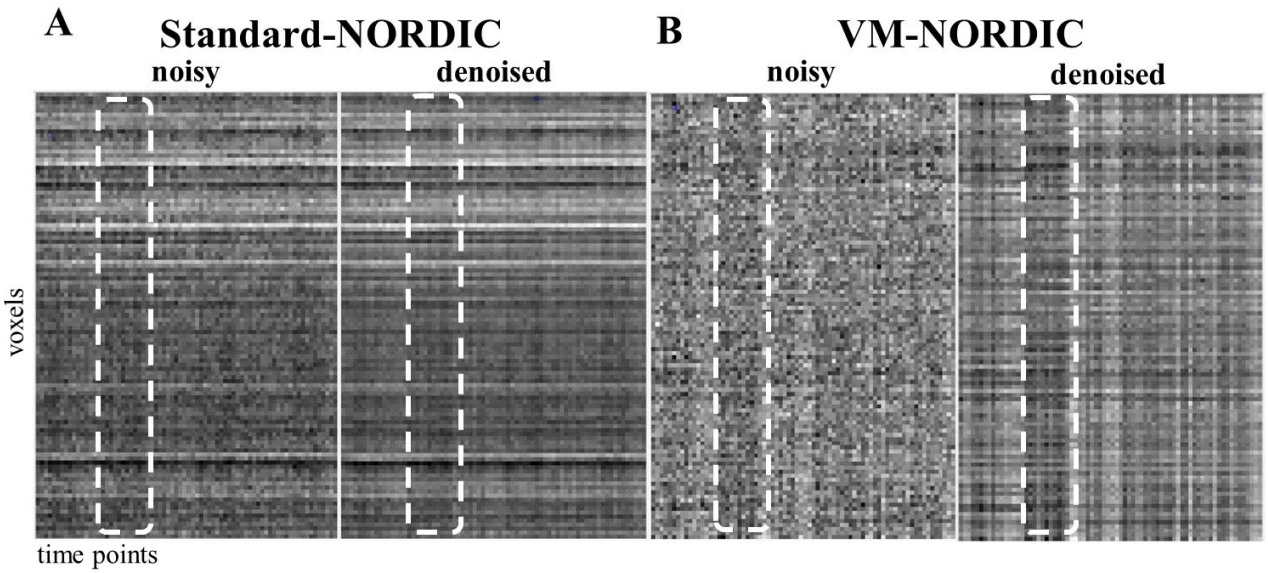

Fig. 2. Exemplar Standard-NORDIC **(A)** and VM-NORDIC **(B)** patches before and after denoising. The dashed boxes indicate one of the temporal signal patterns shared by the timeseries (visible as vertical fading stripe). The VM-NORDIC patch is more homogenous than the Standard-NORDIC patch, allowing the denoising to better reveal and highlight the underlying common signal patterns.

size is a fundamental step in VM-NORDIC. Large patches may boost redundancy by gathering a higher number of similar time series. However, it also implies selecting time series that are less similar to the reference one, degrading the low rankness of the patch. On the other hand, small patches may be more homogeneous but not exhibit enough redundancy for reliable denoising. Therefore, in VM-NORDIC, the optimal patch size is a data-driven trade-off between the level of redundancy and the degree of similarity between time series. VM-NORDIC first makes an overestimated initial guess by multiplying the length of the time series by a factor $F > 1$. This step follows from the observation that the optimal patch size increases with increasing time points in the dataset (see Suppl. Fig. 6 and 8). Then, the algorithm fine-tunes the initial guess via multiple dummy denoising runs to find the patch size K that produces the highest mean tSNR score.

Threshold estimation

The method for estimating the threshold λ_{thr} is analogous in Standard-NORDIC and VM-NORDIC and is based on Monte Carlo simulations of noise matrices[9,29]. In VM-NORDIC, once the optimal patch size K is selected, thermal noise is numerically simulated by generating multiple $K \times Q$ matrices filled with complex zero-mean Gaussian distributed entries and with the same variance as the noise of the data. Then, each random matrix undergoes SVD and the average highest singular value per decomposition is used as the threshold to separate noise from signal components in the data patches.

Voxel averaging

The proposed patching method implies that if a voxel is highly similar to L reference voxels, it will undergo the denoising cycle L times. The final voxel value is the weighted sum of its L values generated after each denoising cycle it underwent. Figure 1B shows a weighting matrix for a representative slice. The values in the matrix represent the number of times the voxels are visited for denoising. This averaging procedure also functions as an additional denoising step by reducing the residual contributions of noise[7]. A similar phenomenon occurs with Standard-NORDIC due to the overlapping patches (Fig. 1C). The resulting 2D weighting matrices for the two methods highlight the core difference in the way they handle voxels.

Brain masking

VM-NORDIC applies a binary bran mask to the dataset to exclude the background from the denoising process using the FSL "bet" function on the time-averaged datasets[37]. Masking reduces the computational time by enabling to process fewer voxels and ensures that the algorithm does not compute similarity scores between brain and background voxels, which could bias the SVD of the patch.

Data chunks

VM NORDIC divides the dataset into multiple chunks of 2D slices to further decrease the computation burden. The number of slices per chunk depends on a hardcoded value (~10^5) standing for the approximate number of voxels allowed per chunk. We chose this value based on the performance of the machines used for testing the algorithm and can be adjusted for different needs. With these procedures, VM-NORDIC takes ~5-10 minutes to denoise a submillimetre BOLD fMRI dataset on our machines.

Participants

We acquired six resting state data sets on three (two females and one male) healthy right-handed subjects (age range: 22-25), (see the "MRI imaging acquisition and processing" sections). All subjects had normal, or corrected vision and provided written informed consent. The study complied with all relevant ethical regulations for work with human participants.

MRI acquisition and processing

All fMRI acquisitions were performed at 7T (Philips) using a 2×16-channel surface coil and covered the occipital lobe. A segmented 3D GE BOLD-EPI sequence was used. Participants were instructed to stay still, minimize movements and close their eyes.

For the first dataset (A), the acquisition covered 30 coronal slices with TR/TE = 50/25ms, flip angle = 18°, segments = 1, SENSE factor (right-left, anterior-posterior) = 3/1, in-plane voxel size = 0.71mm^2, slice thickness = 0.8mm, volumes = 75, matrix size = 240×240, scan time = 4 min. For the second dataset (B), the acquisition covered 30 coronal slices with TR/TE = 65/31ms, flip angle = 18°, segments = 1, SENSE factor (right-left, anterior-posterior) = 4/1, in-plane voxel size = 0.71mm^2, slice thickness = 0.8mm, volumes = 75, matrix size = 240×240, scan time = 5.3 min. For the third dataset (C), the acquisition covered 44 coronal slices with TR/TE = 81/28ms, flip angle = 23°, segments = 1, SENSE factor (right-left, anterior-posterior) = 3/1, voxel size = 0.55mm^3 isotropic, volumes = 115, matrix size = 320×320, scan time = 15 min. For the fourth dataset (D), the acquisition covered 30 coronal slices with TR/TE = 96/31ms, flip angle = 20°, segments = 1, SENSE factor (right-left, anterior-posterior) = 3/1, in-plane voxel size = 0.5mm, slice thickness = 0.8mm, volumes = 75, matrix size = 352×352, scan time = 8 min. For the fifth dataset (E), the acquisition covered 33 coronal slices with TR/TE = 51/25ms, flip angle = 18°, segments = 1, SENSE factor (right-left, anterior-posterior) = 4/1, voxel size = 0.71mm^3 isotropic, volumes = 160, matrix size = 240×240, scan time = 10 min. For the sixth dataset (F), the acquisition covered 24 coronal slices with TR/TE = 98/28ms, flip angle = 23°, segments = 1, SENSE factor (right-left, anterior-posterior) = 3/1, voxel size = 0.45mm^3 isotropic, volumes = 100, matrix size = 320×320, scan time = 15 min.

Experimental design

This study consists of running VM-NORDIC and Standard-NORDIC on datasets with different resolutions and time points to assess which of the two methods leads to an overall more effective denoising in terms of key metrics like tSNR and spatial blurring. To ensure that the new patch size used in VM-NORDIC is not the major cause of differences in results between the two methods, we also performed multiple denoising runs with VM-NORDIC and Standard-NORDIC with different patch sizes. We then plotted Smoothness and tSNR scores Vs patch size to assess if the two methods could produce similar results by only modifying the patch sizes.

Metrics to assess denoising performance

We assessed the denoising performance by comparing tSNR and spatial smoothness estimates between Standard-NORDIC and VM-NORDIC. The tSNR comes from dividing the mean temporal signal by the temporal standard deviation voxel-wise. Smoothness was estimated through the degree of spatiotemporal autocorrelation (FWHM) using the 3dFWHMx function from AFNI with the '-ACF', 'detrend' and 'automask' commands[38]. This function gives back a value in mm, representing the voxel spread (blurring). The spatial autocorrelation was estimated using a Gaussian + monoexponential decay mixed model to account for possible long-tail autocorrelations[9,38].

RESULTS

Local vs. non-local patches

Figure 2 displays two representative patches for Standard-NORDIC and VM-NORDIC before and after denoising. The Standard-NORDIC patch was cropped vertically in the figure to have the same size as the VM-NORDIC patch for easier comparison. The noisy Standard-NORDIC patch in Fig. 2A includes similar voxels as well as voxels with different intensities and time courses. On the other hand, the non-local patch from VM-NORDIC in Fig. 2B appears highly homogenous and with clear common temporal patterns between time series. The denoised VM-NORDIC patch

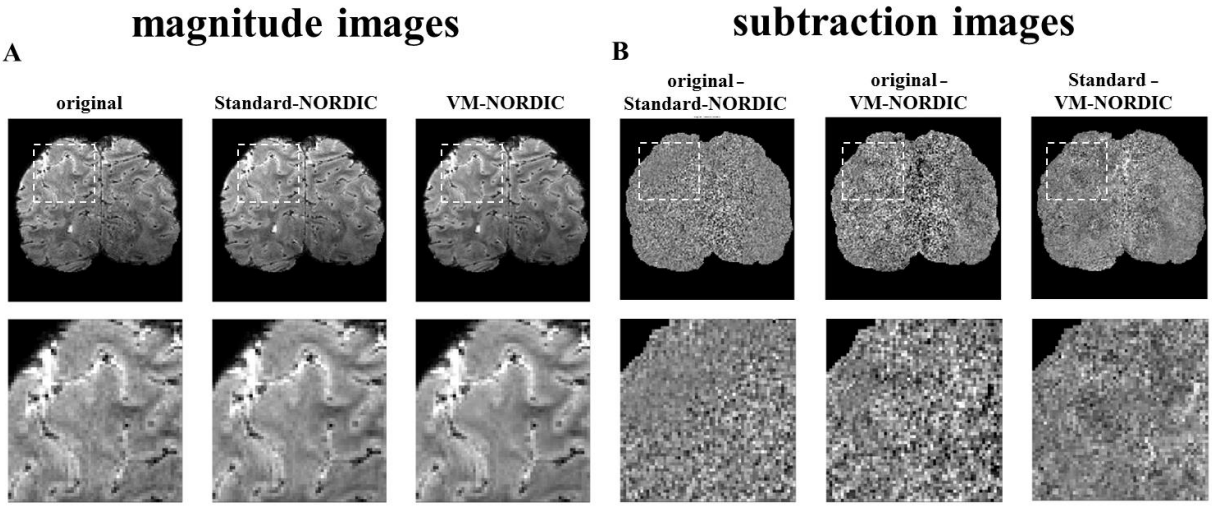

Fig. 3. **A)** Magnitude images and their zoomed views for a representative slice of dataset B before and after denoising with Standard-NORDIC and VM-NORDIC. VM-NORDIC does not interfere with the structural information and leads to more noticeable visual improvements. **B)** Subtraction images and their zoomed views of Standard-NORDIC versus the original data; VM-NORDIC versus the original data; and VM-NORDIC versus Standard-NORDIC. The random patterns indicate that VM-NORDIC leads to a more intense random noise removal especially in central areas with high g-factor values without affecting structural details.

shows enhanced common temporal patterns compared to the Standard-NORDIC denoised patch.

Magnitude images

Fig. 3A shows a reconstructed example slice before denoising and after denoising with Standard-NORDIC and VM-NORDIC, respectively. Both denoising methods lead to improvements upon visual inspection. Particularly, VM-NORDIC reduces noise more effectively in highly noisy central regions (see Fig. 4B). The zoomed insets show a more detailed view of the denoising performance of VM-NORDIC: a more drastic noise attenuation than Standard-NORDIC without degrading structural detail.

Subtraction images

Figure 3B reports the subtraction images for the same representative slice in Fig. 3A of the Standard-NORDIC versus the original image and VM-NORDIC image versus the original image. The two images only show random noise with higher intensities in central regions with high g-factors (Fig. 4). In particular, as also visible in the zoomed views, VM-NORDIC leads to a higher degree of noise removal globally and in central areas. The subtraction of VM-NORDIC from the Standard-NORDIC image shows that VM-NORDIC removes random noise with an emphasis on noisier central areas.

tSNR

Table 1 shows the mean tSNR scores and percentage increases for VM-NORDIC compared to the original data and Standard-NORDIC data for different resolutions and SENSE factors. On average, VM-NORDIC leads to tSNR levels ~9-90% larger than Standard-NORDIC and ~23-250% larger than the original. The highest noise removal occurs for the data with the largest voxel size ($0.7\times0.7\times0.8mm^3$, datasets A and B), especially for the dataset with the lower initial tSNR due to the higher SENSE factor (dataset B). Fig. 4A shows tSNR maps for two representative slices of dataset B before and after denoising with Standard-NORDIC and VM-NORDIC. VM-NORDIC leads to a globally higher tSNR compared to Standard-NORDIC, with a higher increase in central areas coinciding with high g-factor values (Fig. 4B). Fig. 4B shows the normalized tSNR difference map between Standard-NORDIC and VM-NORDIC, the g-factor map estimated in NORDIC via MPPCA, the g-factor map as measured by the scanner and the noise map obtained by acquiring the MR signal without RF excitation of the sample.

Smoothness estimates

On average, VM-NORDIC increased spatial smoothness by 4.8 ± 3.9% (std) compared to the original smoothness estimate. Standard-NORDIC increased spatial smoothness on average by 24 ± 11.6% (std). The smoothness patterns were similar across all datasets. Figure 5 reports the spatial smoothness estimates for dataset B before and after data denoising with Standard-NORDIC and VM-NORDIC. Here, Standard-NORDIC leads to a substantial smoothness increase, while VM-NORDIC manages to keep smoothness to a level similar to the original.

VOXEL SIZE	SENSE FACTOR	TSNR			TSNR INCREASE (%)	
		original	Standard NORDIC	VM-NORDIC	original vs. VM-NORDIC	Standard NORDIC vs. VM-NORDIC
0.5×0.5×0.5mm^3 (C)	3	7	8	8.6	23 %	9 %
0.45×0.45×0.45mm^3 (F)	3	7.3	8.6	9.7	33 %	13 %
0.5×0.5×0.8mm^3 (D)	3	8	10	12.6	57.5 %	27.4 %
0.7×0.7×0.7mm^3 (E)	4	5.4	7.4	9.3	72.2 %	25.5 %
0.7×0.7×0.8mm^3 (A)	3	14	21.3	32.2	130 %	51.3 %
0.7×0.7×0.8mm^3 (B)	4	10	18.4	35	250 %	90.2 %

Table 1. tSNR levels for the original data, after standad NORDIC denoising and after VM-NORDIC denoising and the percentate tSNR increase after VM-NORDIC with respect to the original data and Standard-NORDIC processed data at different spatial resolitions and SENSE factors.

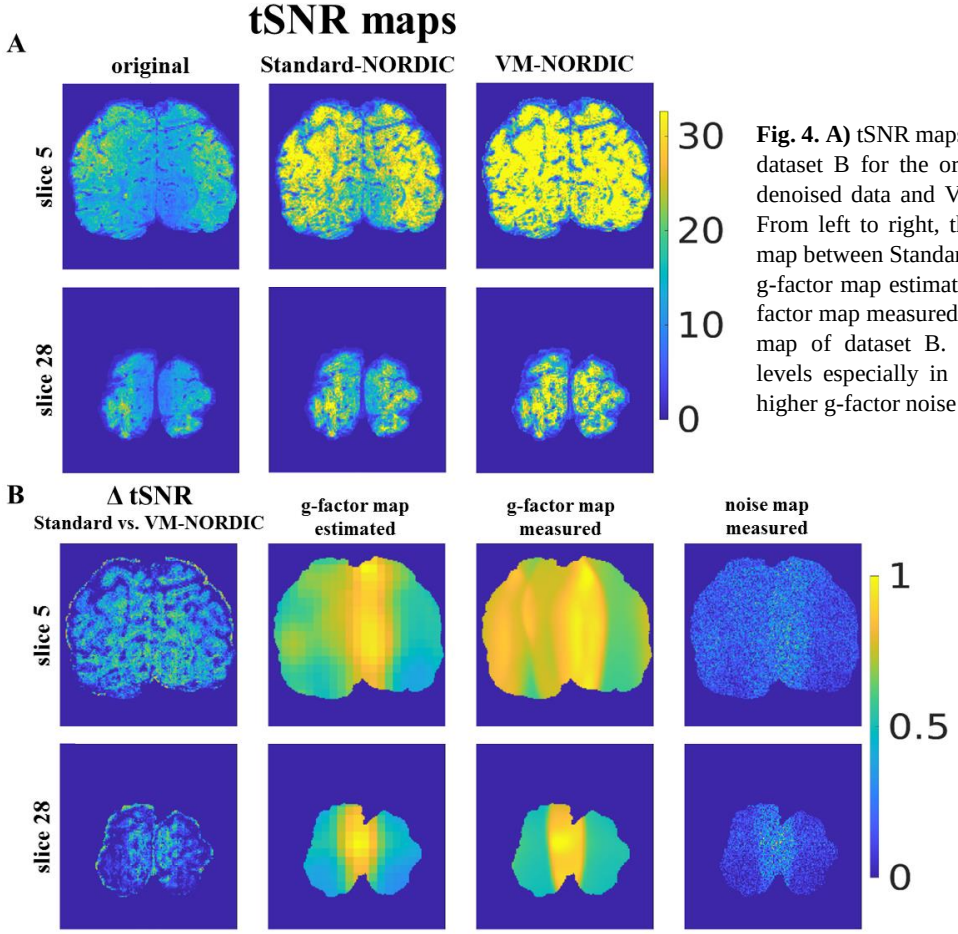

Fig. 4. A) tSNR maps for two representative slices of dataset B for the original data, Standard-NORDIC denoised data and VM-NORDIC denoised data. **B)** From left to right, the normalized tSNR difference map between Standard-NORDIC and VM-NORDIC; g-factor map estimated in NORDIC via MPPCA; g-factor map measured by the scanner; measured noise map of dataset B. VM-NORDIC increases tSNR levels especially in central areas of the brain with higher g-factor noise amplification.

Patch size finetuning

Supplementary figures 6 and 7 show how different patch sizes affect the denoising performance in terms of mean tSNR and smoothness scores for the two methods. The "tSNR vs patch size" plots (Suppl. Fig. 6) show steep curves with clear maxima for VM-NORDIC within the range of 20-240 time series per patch, with small fluctuations due to the length of the time series (Suppl. Fig. 8). Suppl. Fig. 7 shows that VM-NORDIC keeps spatial smoothness estimates similar to or slightly higher than the original one but always lower than after Standard-NORDIC. Suppl. Fig. 8 and 9 show the tSNR vs. patch size and the smoothness vs. patch size plots for datasets with different time series lengths. All plots in Suppl. Fig. 8 show a slight shift of the maxima towards

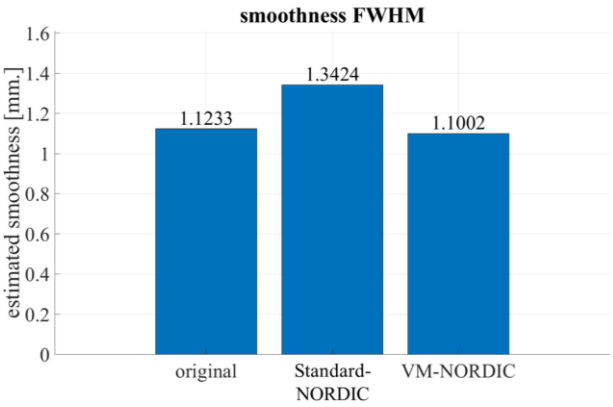

Fig. 5. Spatial smoothness estimates (using 3dFWMHx AFNI) for the representative dataset B. On average, the proposed VM-NORDIC barely affects the original spatial smoothness of the dataset (~5%), indicating improved preservation of structural detail compared to Standard-NORDIC. Specifically, for this dataset, VM-NORDIC results in the spatial smoothness estimate being lower than for the original data. This result follows from the particularly small patch size used on this dataset, which counteracts signal leakage-induced smoothness in the original data.

smaller patch sizes with decreasing time series length. Further, shorter datasets have an initially higher tSNR and show a larger tSNR gain following VM-NORDIC. Suppl. Fig. 9 shows consistent but non-significant smoothness fluctuations with increasing time series lengths.

tSNR/smoothness

Suppl. Fig. 10 displays the plots of the normalized tSNR/smoothness ratio scores vs. patch size. Here, the maxima in the curves indicate the patch sizes with which the methods better remove random noise without excessively smoothing the data. The maxima fall within the same range reported previously (20-240). Also, VM-NORDIC always surpasses Standard-NORDIC in terms of this metric.

DISCUSSION

We introduced VM-NORDIC, an extension to the NORDIC PCA algorithm that further decreases noise in BOLD fMRI data. VM-NORDIC achieves this result via non-local patch formation using voxel similarity matching to boost the low-rank properties of the patches, which facilitates denoising performance using SVT. The Standard-NORDIC and VM-NORDIC patches before and after denoising in Fig. 2 show how the higher homogeneity of the VM-NORDIC patch allows emphasizing the signal fluctuations that the time series have in common, visible as fading vertical stripes. Better-defined signal fluctuations enable the uncovering of the underlying temporal structure of the brain activity which otherwise would be too contaminated by noise (Fig. 2B). Overall, VM-NORDIC reaches a higher degree of thermal noise removal, as indicated by substantial increases in tSNR (up to 2-fold the original tSNR), while better preserving spatial specificity and structural detail, which are key advantages in sub-millimeter resolution BOLD fMRI studies. Altogether, these results indicate that non-local patching promotes a superior low-rank structure of the data, allowing to adequately represent the data with fewer principal components than with local patches in Standard-NORDIC. These improvements are in line with the results reported by Zhao et al., where patching using non-local similarities to denoise dMRI and DTI data significantly reduced noise while preserving structural details compared to MPPCA[33]. As mentioned by the original NORDIC authors[7], NORDIC is beneficial and complementary to other pre-processing steps that target different aspects of the dataset, such as motion correction and physiological noise removal.

The tSNR scores in Table 1 report that the effectiveness of VM-NORDIC increases with increasing voxel size. We argue that these results derive from the higher signal strength at greater voxel sizes, leading to larger and more pronounced principal components containing signal, which benefits SVT. Notably, among the two datasets with the largest and equal voxel size (datasets A and B), the one with the initially higher noise level (dataset B, due to the higher SENSE factor) shows the largest tSNR gain. This result implies that VM-NORDIC targets thermal noise only while leaving the target signal intact. However, the higher noise level of dataset B due to the greater SENSE factor makes its tSNR lower than in dataset A. Consequently, as visible in Table 1, the two datasets reach similar tSNR levels upon VM-NORDIC, implying that denoising removed a larger amount of thermal noise from the noisier dataset B.

That VM-NORDIC effectively removes more noise while preserving the structural integrity of the images is visually perceivable in the reconstructed images (Fig. 3A). Moreover, the difference between VM-NORDIC and the original image shows only noise without notable edge effects (Fig. 3B). The difference between VM-NORDIC and Standard-NORDIC also shows mostly noise, indicating that non-local patching in VM-NORDIC removes additional noise than Standard-NORDIC. Moreover, the same image also exhibits spatial patterns in correspondence of high g-factor values (Fig. 4B). These results suggest that the higher mean tSNR after VM-NORDIC mainly derives from stronger denoising of central areas. Nonetheless, also peripheral areas exhibit a less pronounced yet tangible tSNR increase after VM-NORDIC, indicating a globally stronger noise removal (see figure 4).

These converging findings point out that in local patches from highly noisy regions, noise irreversibly propagates through all signal components, leading to less effective SVT compared to areas where noise is less dominant. This scenario is never obtainable with VM-NORDIC since non-local patches inherently have well-pronounced signal components by construction. Always having well- pronounced signal components allows VM-NORDIC to denoise all voxels across the dataset with comparable efficiency, as long as enough similar voxels are present to generate sufficient redundancy. Therefore, the outcome is, on average, a high and consistent tSNR level across the whole dataset independently of the local noise levels.

VM-NORDIC also induces lower levels of spatial smoothing than Standard-NORDIC (Fig. 5, Suppl. Fig. 7). As mentioned, Standard-NORDIC applies an LLR method that inevitably increases spatial smoothing by increasing the similarity between adjacent voxels (see Methods and Materials, section "*Standard-NORDIC LLR model*"). Importantly, we observe that VM-NORDIC minimally smoothens the dataset by denoising groups of voxels from different locations across the dataset, increasing the original smoothness by less than 5% on average (about 20% of that induced by Standard-NORDIC). Minimal contamination of the spatial integrity is crucial for sub-millimeter BOLD fMRI applications.

The steep tSNR curves as a function of patch size in Suppl. Fig. 6 suggest that the performance of VM-NORDIC is more sensitive to the patch size compared to Standard-NORDIC. Notably, for VM-NORDIC, the sharp maxima of mean tSNR as a function of patch size reveal the existence of an optimal patch size per dataset. We argue that these results derive from the notion that in Standard-NORDIC, adding or subtracting a few voxels from an inhomogeneous local patch does not strongly modify its low-rank representation and thus the effectiveness of denoising (Suppl. Fig. 6). Contrarily, in VM-NORDIC, having larger patches can lower the overall similarity ranking, which potentially reduces the patch homogeneity and hampers efficient SVT. However, larger patches can also improve denoising by boosting signal redundancy if the additional voxels are sufficiently similar to the reference one. Oppositely, small patches may be highly homogeneous but not exhibit enough redundancy for a reliable SVT. These concepts suggest that the optimal patch size is a data-driven trade-off between the degree of similarity of the voxels and the level of signal redundancy. Hence, a data-specific fine-tuning process is necessary to find the optimal patch size (see Methods and Materials section "*Patch size optimization*"). The figure also shows that VM-NORDIC produces higher tSNR scores than Standard-NORDIC at suboptimal patch sizes, meaning that the patch size alone is not the sole promoter of the improvements. Suppl. Fig. 8 illustrates that the optimal patch size per dataset slightly fluctuates or decreases with decreasing time points, and that data with fewer time points reach a higher tSNR increase thanks to the higher initial tSNR.

Smoothness scores per different patch sizes indicate that, overall, the degree of smoothing upon VM-NORDIC remains lower than for Standard-NORDIC independently of the size of the time series (Suppl. Fig. 7 and 8). In particular, in VM-NORDIC, small patch sizes generally preserve spatial detail more effectively than large ones. We suggest that the reduced number of time series allowed in small patches limit the chance of denoising the same time series too many times, which could potentially remove the characteristic underlying signal. Further, signal leakage due to, for instance, susceptibility artefacts can slightly increase the smoothness of the original data by spreading the magnetic signals from one brain region into neighbouring voxels, leading to a decrease in spatial resolution and blurring[39]. Consequently, particularly small patches avoid further signal spreading or, in a few cases, even counteract signal leakage effects by decreasing the artificial similarities between adjacent voxels, leading to comparable or lower smoothness estimates than in the original data (e.g. Fig. 5, Suppl. Fig. 7, second panel).

Since the goal of a denoising method is maximizing the tSNR while minimizing spatial smoothing, combining tSNR and smoothness scores into a single metric allows for a direct and global assessment of the denoising performance. The normalized tSNR/smoothness ratio vs. patch size plots for different time series in Suppl. Fig. 10 shows that VM-NORDIC reaches higher tSNR/smoothness ratios than Standard-NORDIC, thus a more convenient trade-off between noise removal and induced smoothness. Further, also these curves emphasize a range of dataset-specific optimal patch sizes (20-240) that guarantee a higher degree of noise removal while preserving structural information.

Ideally, a reference voxel should be compared to all brain voxels across the entire dataset to maximize the chance of finding high similarities and grouping voxels from different slices. However, computing the similarity scores of the whole dataset at once is computationally expensive. A solution was to process fewer data per cycle by dividing the dataset into equally large chunks of temporal slices and denoise each chunk individually (for more details, see Methods and Materials section "*Chunks Size*"). As long as the chunks contain enough data and slices, there will still be a high chance of finding highly similar non-local voxels. After selecting the appropriate chunk size based on the performance of

our machines, we saw that chunking significantly speeds up the processing without interfering with the quality of denoising.

Finally, a recent study reported that Standard-NORDIC and MPPCA occasionally generate artificial functional activation depending on the selected patch size[11]. In particular, they found that larger patch sizes provide higher sensitivity to BOLD responses with both MPPCA and Standard-NORDIC, but with significant activation "spreading" and increasing false positives rate due to the local bleeding of active signal components[11]. They concluded that for both methods, the optimal patch size for each experiment depends on data tSNR and functional CNR. Similarly, our findings also show an increase in the induced spatial smoothness (hence signal spreading) with increasing patch sizes for VM-NORDIC (Supp. Fig. 7). However, as discussed, the patch size optimization process in VM-NORDIC already relies on the data tSNR and usually estimates patch sizes smaller than those used in Standard-NORDIC. Hence, we suggest that the lower spatial smoothness upon VM-NORDIC with small patch sizes can limit the activation spreading by keeping unaltered or even decreasing the artificial spatiotemporal similarities between adjacent voxels[39]. Yet, a more dedicated investigation is necessary to understand if a similar effect occurs during VM-NORDIC even after the patch size optimization.

Among the limitations of VM-NORDIC, there is the inevitable increase in the processing time (on average ~1.8 times longer than with NORDIC standard, but still within ~5-10 minutes on our machines depending on the size of the data) due to the additional steps of voxel matching and patch size fine-tuning. Further, the present study did not include analyses of functional activity such as t-statistic activation maps. Since the end goal of VM-NORDIC is to improve the reliability and further push sub-millimeter BOLD fMRI studies, these statistical analyses are of paramount importance. However, we suggest that because the voxels in a VM-NORDIC patch are already highly correlated owing to their similar time courses, increasing their similarity may help reveal unspotted activation.

CONCLUSION

A large part of the neuroscientific community is devoted to studying mesoscopic cortical organizations like cortical columns and layers in the visual cortex[8]. So far, most of the detailed studies were possible only by invasively mapping animal cortices. Expanding the investigation to in-vivo human brains requires non-invasive, precise and reliable imaging methods[3]. Specifically, the BRAIN Initiative Working group reported that the voxel size (e.g. resolution) necessary to resolve these small-scale structures spans from ~0.46mm and ~0.55mm isotropic[7,40]. Reducing noise while preserving spatial integrity is essential to enable high-resolution fMRI studies revealing new neuroscientific insights at the mesoscopic scale. Together with the latest advancements in hardware technologies, such as the development of dedicated receive head coils, powerful gradient inserts, and even stronger static magnetic fields[41–43], VM-NORDIC will allow the community to further push the boundaries of fMRI resolution to an unprecedented degree. This can potentially pave the way to highly relevant discoveries within UHF submillimeter fMRI[4,5,14]. Future steps should investigate the effects of VM-NORDIC on functional mapping and apply it to higher-resolution data acquired with the aforementioned specialized hardware.

ACKNOWLEDGEMENTS

I would like to thank Dr Natalia Petridou for providing the high-resolution fMRI data to test the algorithm in its initial stage and for accepting to be the second reader of my project thesis, Stanley D.T. Pham for helping me get familiar with the 7T MRI apparatus and application software, and Jason van Schoor for tips and trickss. Finally, thanks to the ISMRM committee and reviewers for accepting my abstract and allowing me to present my findings during the ISMRM 2023 Annual Meeting and Exhibition in Toronto, Canada and Benelux Chapter 2023 in Brussels, Belgium (abstract at the bottom of the document).

REFERENCES

1. Ogawa, S. *et al.* Intrinsic signal changes accompanying sensory stimulation: Functional brain mapping with magnetic resonance imaging. *Proc Natl Acad Sci U S A* **89**, (1992).
2. Kwong, K. K. *et al.* Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proc Natl Acad Sci U S A* **89**, (1992).
3. Bandettini, P. A., Wong, E. C., Hinks, R. S., Tikofsky, R. S. & Hyde, J. S. Time course EPI of human brain function during task activation. *Magn Reson Med* **25**, (1992).
4. Uğurbil, K. Imaging at ultrahigh magnetic fields: History, challenges, and solutions. *Neuroimage* **168**, (2018).
5. Viessmann, O. & Polimeni, J. R. High-resolution fMRI at 7 Tesla: challenges, promises and recent developments for individual-focused fMRI studies. *Current Opinion in Behavioral Sciences* vol. 40 Preprint at https://doi.org/10.1016/j.cobeha.2021.01.011 (2021).

6. Sudlow, C. *et al.* UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age. *PLoS Med* **12**, (2015).
7. Jorgenson, L. A. *et al.* The BRAIN initiative: Developing technology to catalyse neuroscience discovery. *Philosophical Transactions of the Royal Society B: Biological Sciences* **370**, (2015).
8. Dumoulin, S. O., Fracasso, A., van der Zwaag, W., Siero, J. C. W. & Petridou, N. Ultra-high field MRI: Advancing systems neuroscience towards mesoscopic human brain function. *Neuroimage* **168**, (2018).
9. Vizioli, L. *et al.* Lowering the thermal noise barrier in functional brain mapping with magnetic resonance imaging. *Nat Commun* **12**, (2021).
10. Dowdle, L. T. *et al.* Improving Sensitivity to Functional Responses without a Loss of Spatiotemporal Precision in Human Brain Imaging. doi:10.1101/2021.08.26.457833.
11. Fernandes, F. F., Olesen, J. L., Jespersen, S. N. & Shemesh, N. *MP-PCA denoising of fMRI time-series data can lead to artificial activation 'spreading'*.
12. Triantafyllou, C. *et al.* Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *Neuroimage* **26**, 243–250 (2005).
13. Triantafyllou, C., Polimeni, J. R. & Wald, L. L. Physiological noise and signal-to-noise ratio in fMRI with multi-channel array coils. *Neuroimage* **55**, (2011).
14. Wald, L. L. & Polimeni, J. R. Impacting the effect of fMRI noise through hardware and acquisition choices – Implications for controlling false positive rates. *Neuroimage* **154**, (2017).
15. Hoult, D. I. & Richards, R. E. The signal-to-noise ratio of the nuclear magnetic resonance experiment. *Journal of Magnetic Resonance* **213**, (2011).
16. Edelstein, W. A., Glover, G. H., Hardy, C. J. & Redington, R. W. The intrinsic signal-to-noise ratio in NMR imaging. *Magn Reson Med* **3**, (1986).
17. Bianciardi, M. *et al*. Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: a 7 T study. *Magn Reson Imaging* **27**, (2009).
18. Hu, X. & Kim, S. -G. Reduction of signal fluctuation in functional MRI using navigator echoes. *Magn Reson Med* **31**, (1994).
19. Bianciardi, M., van Gelderen, P., Duyn, J. H., Fukunaga, M. & de Zwart, J. A. Making the most of fMRI at 7 T by suppressing spontaneous signal fluctuations. *Neuroimage* **44**, (2009).
20. Fox, M. D., Snyder, A. Z., Zacks, J. M. & Raichle, M. E. Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. *Nat Neurosci* **9**, (2006).
21. Kay, K. N., Rokem, A., Winawer, J., Dougherty, R. F. & Wandell, B. A. GLMdenoise: A fast, automated technique for denoising task-based fMRI data. *Front Neurosci* (2013) doi:10.3389/fnins.2013.00247.
22. Caballero-Gaudes, C. & Reynolds, R. C. Methods for cleaning the BOLD fMRI signal. *Neuroimage* **154**, (2017).
23. Lund, T. E., Madsen, K. H., Sidaros, K., Luo, W. L. & Nichols, T. E. Non-white noise in fMRI: Does modelling have an impact? *Neuroimage* **29**, (2006).
24. Murphy, K., Birn, R. M. & Bandettini, P. A. Resting-state fMRI confounds and cleanup. *Neuroimage* **80**, (2013).
25. Pruim, R. H. R. *et al.* ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data. *Neuroimage* **112**, (2015).
26. Pruessmann, K. P., Weiger, M., Scheidegger, M. B. & Boesiger, P. SENSE: Sensitivity encoding for fast MRI. *Magn Reson Med* **42**, (1999).
27. Triantafyllou, C., Hoge, R. D. & Wald, L. L. Effect of spatial smoothing on physiological noise in high-resolution fMRI. *Neuroimage* **32**, (2006).
28. Veraart, J., Fieremans, E. & Novikov, D. S. Diffusion MRI noise mapping using random matrix theory. *Magn Reson Med* **76**, 1582–1593 (2016).
29. Moeller, S. *et al.* NOise reduction with DIstribution Corrected (NORDIC) PCA in dMRI with complex-valued parameter-free locally low-rank processing. *Neuroimage* **226**, (2021).
30. Marčenko, V. A. & Pastur, L. A. DISTRIBUTION OF EIGENVALUES FOR SOME SETS OF RANDOM MATRICES. *Mathematics of the USSR-Sbornik* **1**, (1967).
31. Veraart, J. *et al.* Denoising of diffusion MRI using random matrix theory. *Neuroimage* **142**, (2016).
32. Knudsen, L., Bailey, C. J., Blicher, J. U., Yang, Y. & Zhang, P. Feasibility of 3T layer-dependent fMRI with GE-BOLD using NORDIC and 2 phase regression 3. *Palle Juul-Jensens Boulevard* **99**,.
33. Zhao, Y. *et al.* Joint denoising of diffusion-weighted images via structured low-rank patch matrix approximation. *Magn Reson Med* (2022) doi:10.1002/mrm.29407.
34. Coupe, P. *et al.* An optimized blockwise nonlocal means denoising filter for 3-D magnetic resonance images. *IEEE Trans Med Imaging* **27**, (2008).
35. Dabov, K., Foi, A., Katkovnik, V. & Egiazarian, K. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on Image Processing* **16**, (2007).
36. Lee, J., Kim, S., Lebanon, G. & Singer, Y. Local low-rank matrix approximation. in *30th International Conference on Machine Learning, ICML 2013* (2013).
37. Popescu, V. *et al.* Optimizing parameter choice for FSL-Brain Extraction Tool (BET) on 3D T1 images in multiple sclerosis. *Neuroimage* **61**, (2012).
38. Cox, R. W., Chen, G., Glen, D. R., Reynolds, R. C. & Taylor, P. A. FMRI Clustering in AFNI: False-Positive Rates Redux. *Brain Connect* **7**, (2017).
39. Risk, B. B., Kociuba, M. C. & Rowe, D. B. Impacts of simultaneous multislice acquisition on

sensitivity and specificity in fMRI. *Neuroimage* **172**, (2018).
40. Bargmann, C., Newsome, W. & Family, H. *BRAIN 2025: A scientific vision Final report of the ACD BRAIN working group*.
41. Sengupta, S. *et al.* A specialized multi-transmit head coil for high resolution fMRI of the human visual cortex at 7T. *PLoS One* **11**, (2016).
42. Petridou, N. *et al.* Pushing the limits of high-resolution functional MRI using a simple high-density multi-element coil design. *NMR Biomed* **26**, (2013).
43. Versteeg, E. *et al.* A plug-and-play, lightweight, single-axis gradient insert design for increasing spatiotemporal resolution in echo planar imaging-based brain imaging. *NMR Biomed* **34**, (2021).

SUPPLEMENTARY MATERIALS

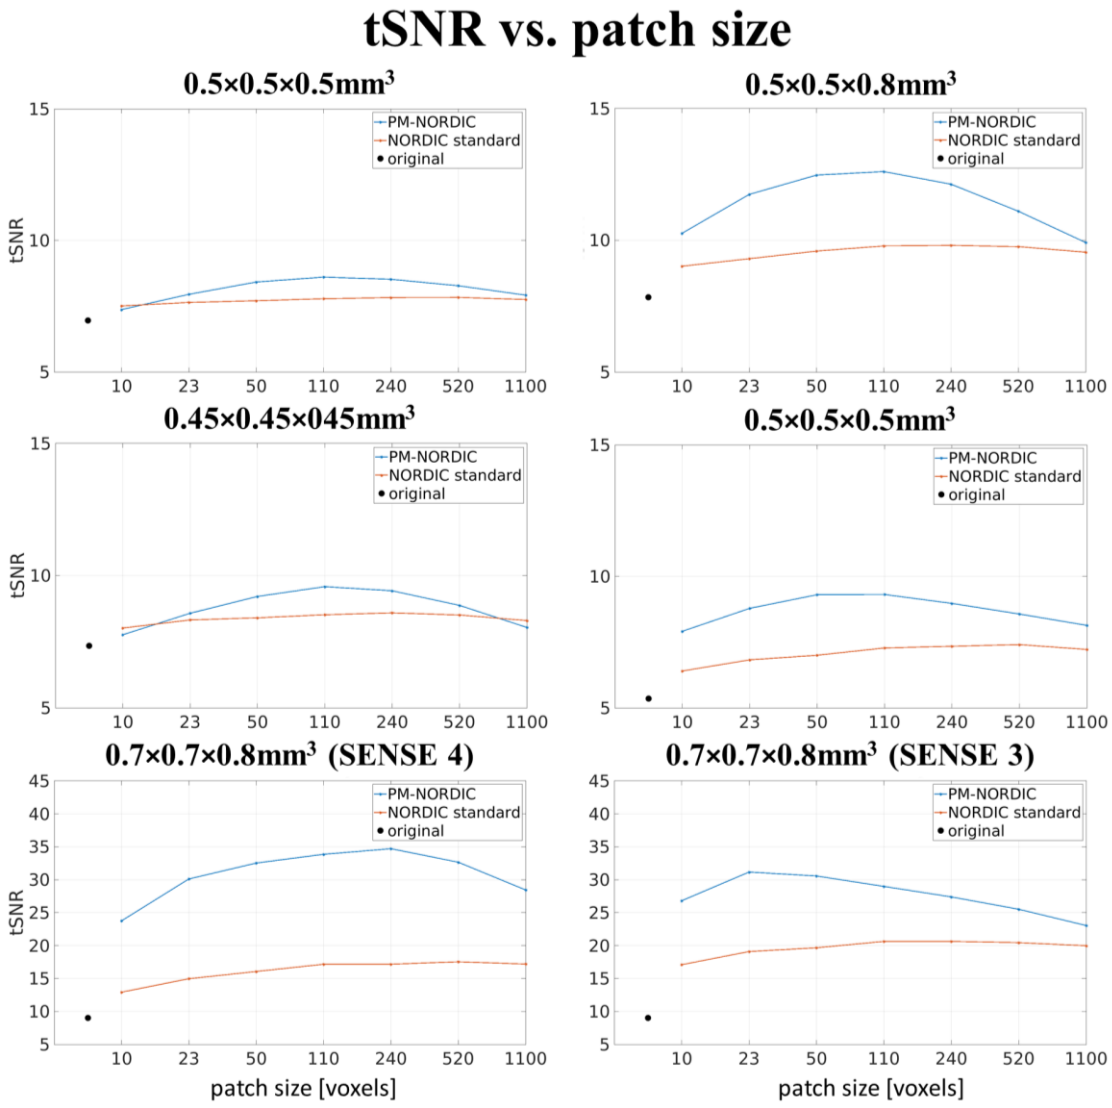

Suppl. Fig 6. Mean tSNR as a function of patch size for Standard-NORDIC and VM-NORDIC. The denoising performance of VM-NORDIC is more sensitive to the chosen patch size. Also, The curves for VM-NORDIC show clear maxima representing the optimal patch sizes.

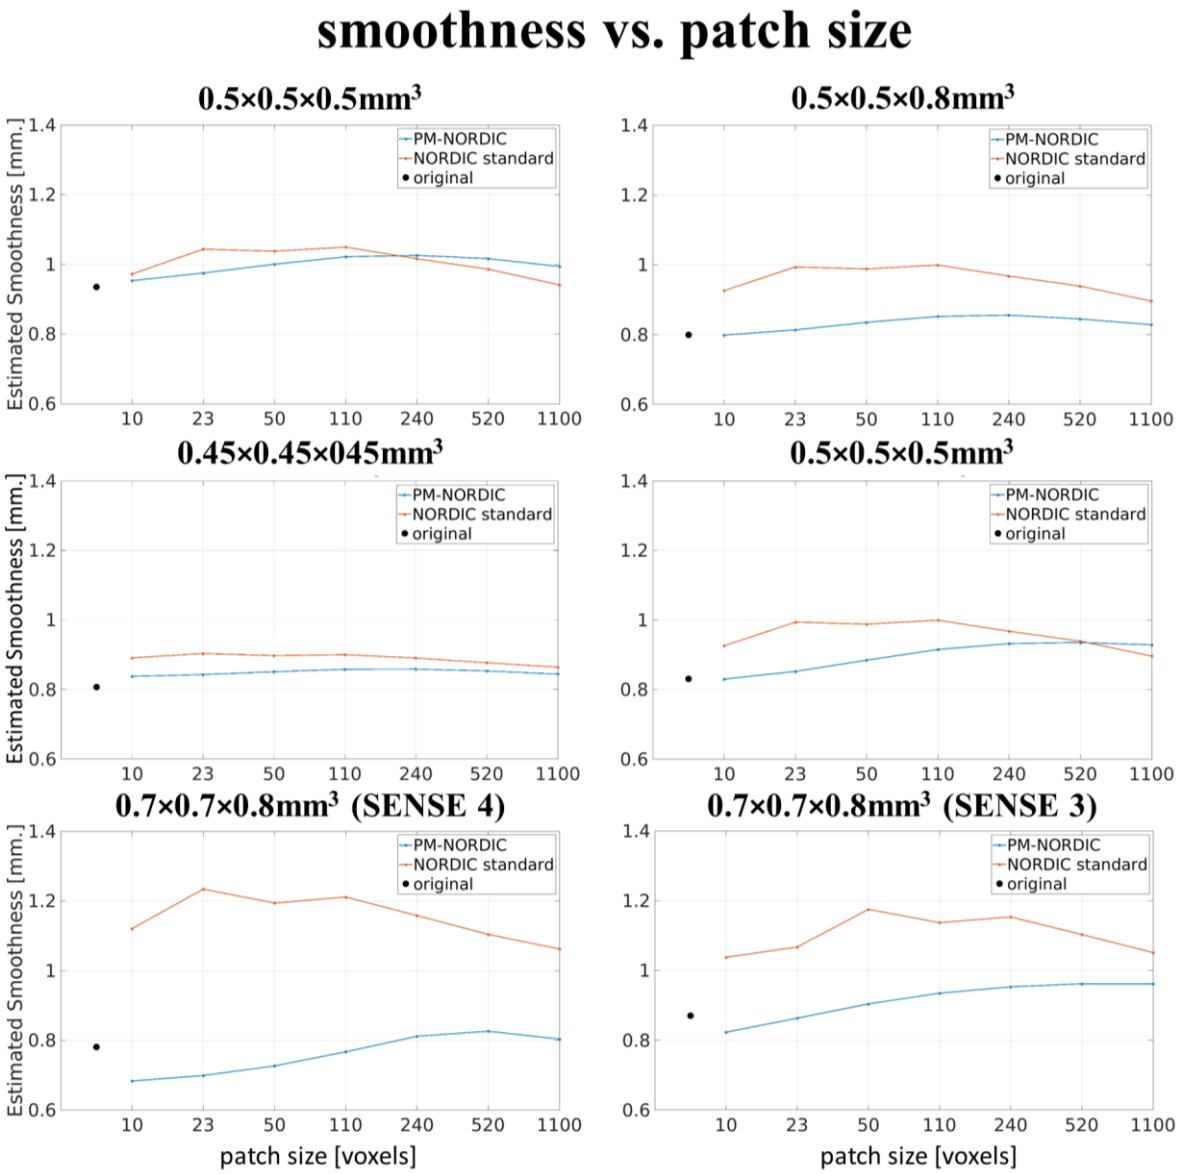

Suppl. Fig 7. Global smoothness estimates (FWHM) as a function of patch size for Standard-NORDIC and VM-NORDIC. VM-NORDIC does not significantly increase spatial smoothing and, for certain patch sizes, even leads to smoothness estimates lower than the originals by counteracting smoothing signal leakage.

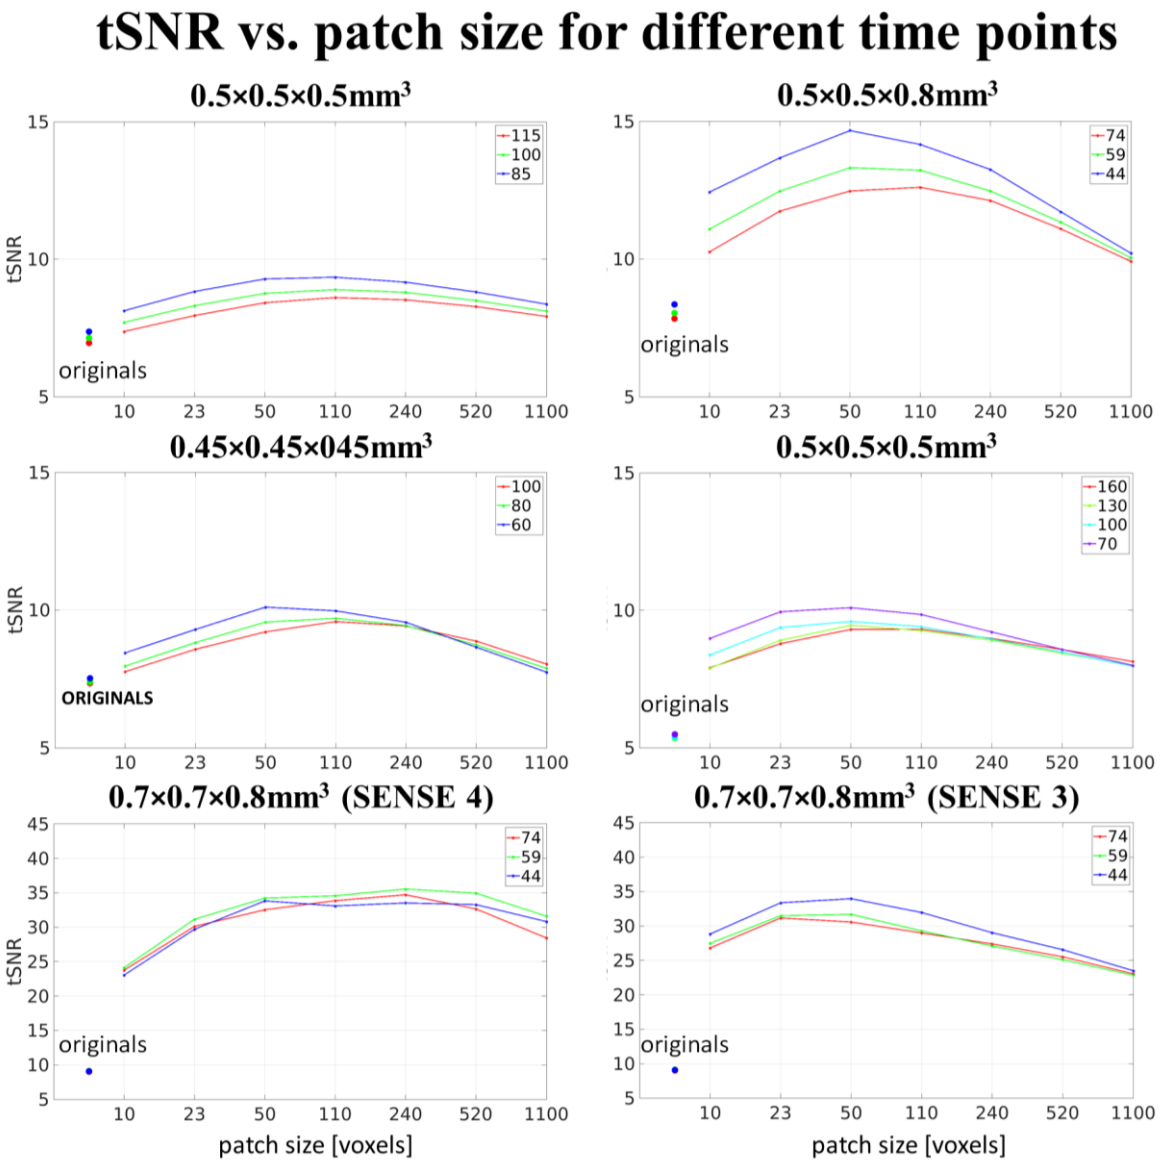

Suppl. Fig 8. Mean tSNR as a function of patch size for VM-NORDIC denoised datasets with different numbers of time points. The optimal patch size slightly increases with increasing time points. Nonetheless, the shift is not significant as long as the number of time points is within a regular range. Additionally, shorter datasets show the largest tSNR gains.

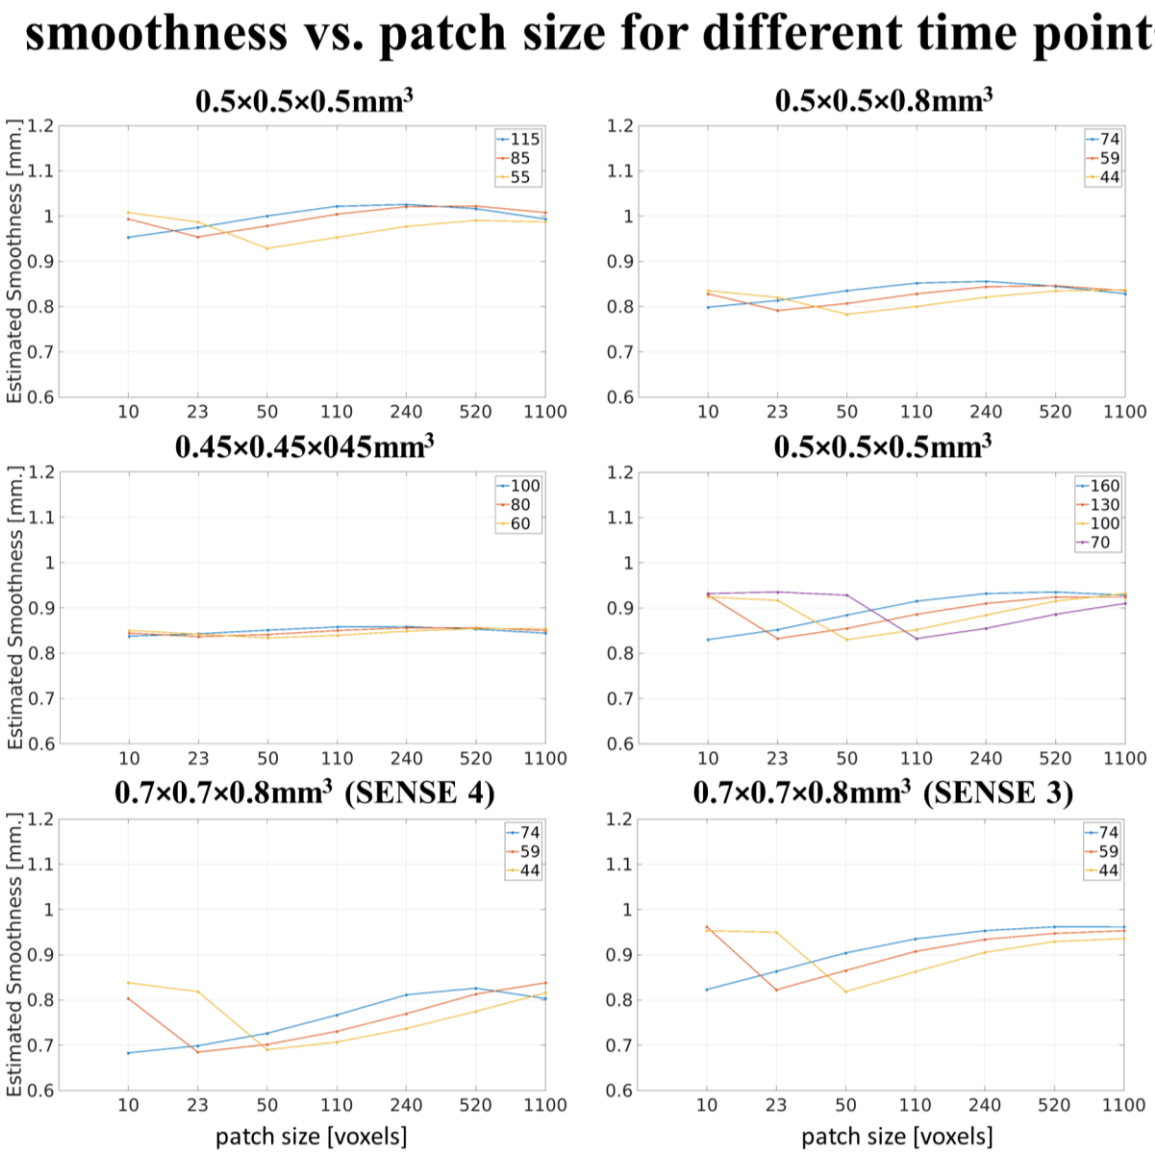

Suppl. Fig 9. Global smoothness estimates (FWHM) as a function of patch size for VM-NORDIC denoised datasets with different numbers of time points. The degree of spatial blurring exhibits small but irrelevant fluctuations as the number of time points of the dataset changes.

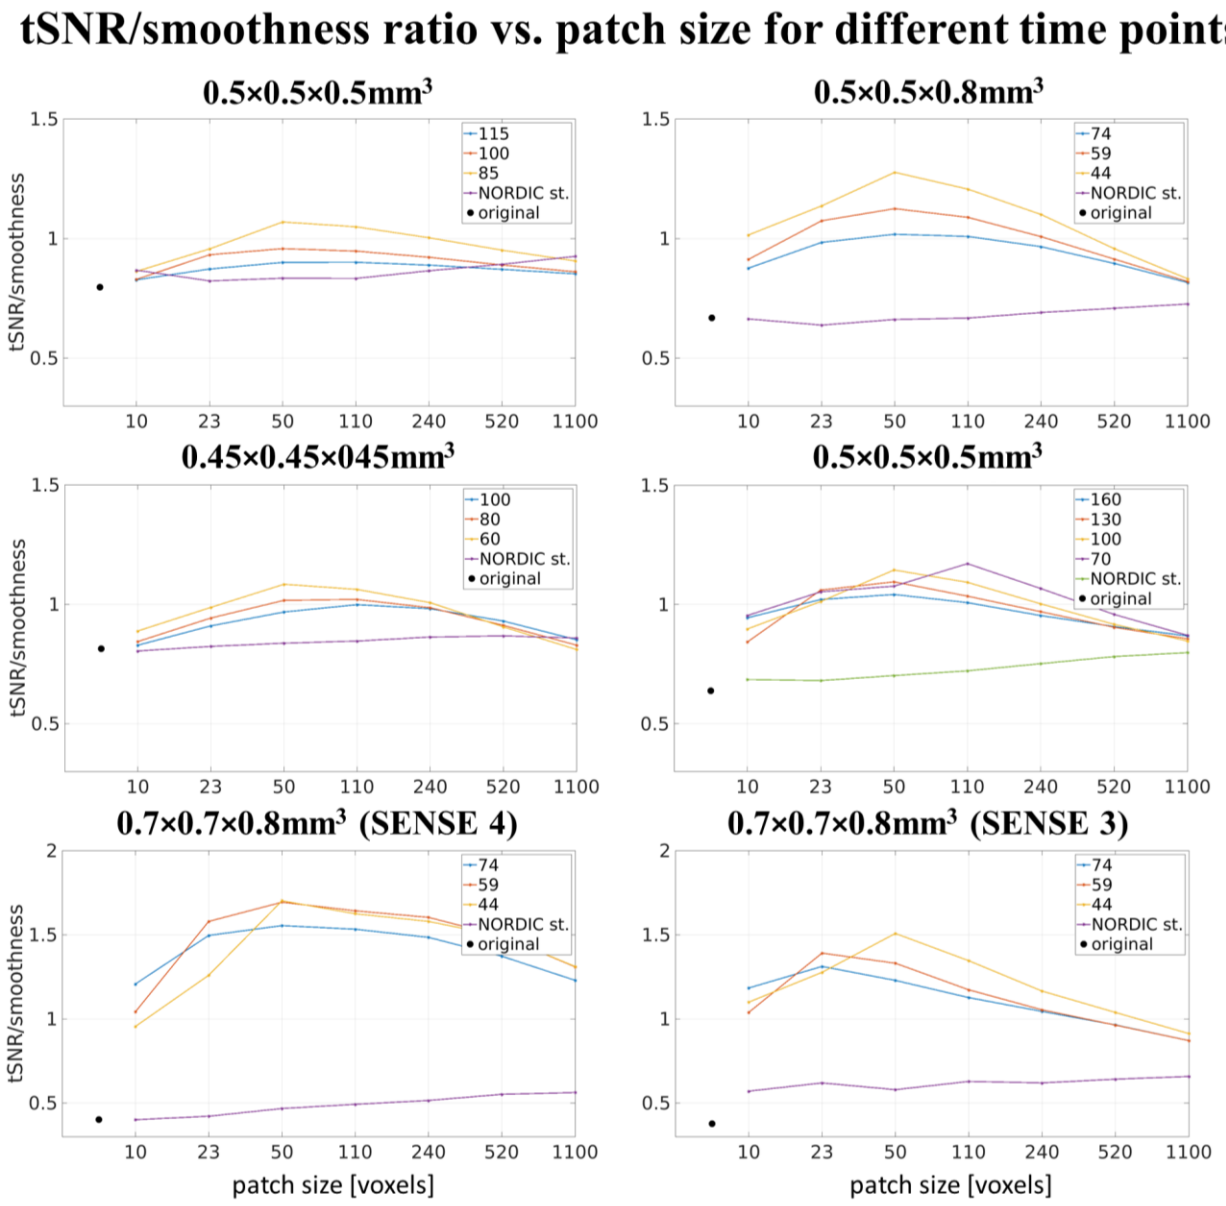

Suppl. Fig 10. Normalized tSNR over normalized global smoothness estimate (FWHM) ratio as a function of patch size for VM-NORDIC denoised datasets with different number of time points. This metric shows which patch size guarantees the best denoising performance in terms of the trade-off between nosie removal and induced spatial smoothness. The curves indicate that VM-NORDIC has a clear advantage over Standard-NORDIC, and that the optimal patch sizes range from 20 up to 240 timeseries per patch.

ISMRM ABSTRACT

Below is the abstract that allowed me to get accepted in the ISMRM 2023 Annual Meeting and Exhibition as well as the ISMRM Benelux Chapter 2023.

Improved NORDIC denoising for submillimetre BOLD fMRI using patch formation via non-local pixels similarity - pixel-matching (PM) NORDIC

Alessandro Nigi[1] and Jeroen C.W. Siero[1]

[1]*Radiology, UMCU, Utrecht, Netherlands*

Synopsis

Submillimetre BOLD fMRI enables studying brain function at the mesoscopic level but is limited by low SNR. The NORDIC PCA algorithm reduces thermal noise levels in fMRI in a patch-wise manner via singular value thresholding (SVT). However, the NORDIC patch formation uses adjacent pixels that often contain signals from multiple tissues, which can degrade the denoising performance. We propose an alternative patch formation using the similarity between non-local pixels, dubbed pixel-matching (PM) NORDIC. PM-NORDIC outperforms standard NORDIC in terms of temporal SNR and spatial smoothness estimates.

Introduction

BOLD fMRI is an indispensable tool for depicting brain function. However, its low contrast-to-noise ratio and low SNR for high-resolution data limit its applicability and reliability[1]. Recently, Vizioli et *al*. (2021) introduced NOise Reduction with Distribution Corrected (NORDIC) PCA to attenuate thermal noise levels in fMRI via a principal component analysis (PCA) approach[2]. The algorithm divides the dataset into consecutive 4D (adjacent 3D spatial + 1D temporal) patches and denoises each patch by truncating the distribution of its singular values according to a parameter-free threshold. A potential downside of this patch formation through adjacent pixels is that the patches often contain signals from multiple tissues. Mixing of different tissues within a patch can negatively affect the identification of true signal components and may not exploit all information redundancy[3,4,5]. Zhao et *al*. (2022) introduced PCA-based denoising for diffusion MRI data by constructing patch matrices from similar non-local patches[3,6,7]. They argued that grouping similar patches exploits more information redundancy and promotes the low-rankness of the patch matrices, resulting in improved denoising for diffusion MRI. NORDIC denoising for BOLD fMRI utilizes signal redundancy over time and could thus also benefit from patch formation based on the non-local similarity of pixels[2]. Here, we apply an alternative patching method based on non-local pixel similarity (pixel-matching) for NORDIC (dubbed PM-NORDIC) on high-resolution fMRI data and show improved denoising compared to standard NORDIC in terms of temporal SNR (tSNR) and spatial smoothness estimates.

Methods

Imaging data: fMRI acquisitions were performed at 7T (Philips) using a 2x16-channel surface coil. A segmented 3D GE BOLD-EPI sequence was used, covering 40 coronal slices on the visual cortex with TR/TE=54/27ms, flipangle=20°, segments=3, SENSE factor (right-left, anterior-postior)=3.5/1.5, 0.55mm isotropic voxelsize, volumes=49, matrixsize=240x240.

Standard NORDIC denoising: NORDIC uses locally low-rank (LLR) properties of image patches across image series. The algorithm divides the dataset into a series of 4D k x k x k x Q adjacent patches, where k stands for the dimensions of the spatial patch and Q is the number of time points. Here k=8 was used. Each patch is denoised via singular value thresholding (SVT) by cancelling the components indistinguishable from zero-mean Gaussian distributed noise (via Monte Carlo simulations), producing a noise-free low-rank approximation of the original noisy patch.

Proposed pixel-matching NORDIC: Here, patch formation occurs via grouping similar non-local pixels (Figure1). For each 2D slice, a number (N) of reference pixels are selected. Each reference pixel is compared to M neighbouring pixels to compute a similarity score based on the Frobenius norm between the pixel's time series. The (k^2)-1 pixels exhibiting the highest similarity scores are vectorised and grouped with their reference pixel to form N Casorati matrices of (k^2) x Q elements (Figure1A). Finally, each Casorati matrix is denoised according to the standard NORDIC SVT and back distributed into image space (Figure1C).

Experiment design: In PM-NORDIC, the outcome of denoising highly depends on the patch size. A large patch size contains pixels with lower scores in the similarity ranking, thus results less homogenous compared to a smaller patch. Hence, we compared PM-NORDIC to standard NORDIC using two patch sizes: 1) the default computed by NORDIC (k x k x k x Q=8x8x8x49) when the user does not explicitly provide a specific patch size, and 2) a smaller patch of ([(k x k x k)^(⅓)]^2 x Q, i.e. 8x8x49)

Data analysis: The denoising performance was assessed by comparing tSNR and spatial smoothness estimates between standard NORDIC and PM-NORDIC. TSNR was computed by dividing the mean temporal signal by the temporal standard deviation pixel-wise. Smoothness was estimated through the degree of spatiotemporal autocorrelation (FWHM) using 3dFWHMx from AFNI using the '-ACF', 'detrend' and 'automask' commands.

Results

Compared to the non-denoised original data, the tSNR increased by a factor of 1.89, 2.2 and 2.5 for standard NORDIC with default patch size, for standard NORDIC with small patch size and PM-NORDIC, respectively (Figure2). Spatial smoothness increased by 11%, 30% and 7% for standard NORDIC with default patch size, standard NORDIC with small patch size, and PM-NORDIC, respectively (Figure3).

Discussion

We introduced PM-NORDIC, an alternative patch formation approach for NORDIC PCA based on the similarity between non-local pixels[3,6,7]. PM-NORDIC showed improved tSNR and spatial smoothness scores for ~0.5mm isotropic BOLD fMRI data compared to standard NORDIC. Hence, homogenous patches are more easily low-rank approximated via SVT, leading to more effective denoising. Moreover, for PM-NORDIC, a small patch size signifies selecting fewer pixels along the similarity ranking, increasing the homogeneity of the resulting patch and further improving the denoising performance (Figure4,5). Finally, smoothness scores show that the proposed method better preserves structural detail than the default approach (Figure4,5). However, the pixel-matching approach acts on groups of non-local pixels. Therefore, adjacent pixels will unlikely be present in the same patch, leading to low levels of autocorrelation across the denoised image.

Conclusion

The proposed PM-NORDIC denoising using patches of highly similar non-local pixels can boost denoising performance. PM-NORDIC led to substantial increases in tSNR and preservation of spatial smoothness. Reducing thermal noise contributions while preserving spatial integrity is essential for conducting BOLD fMRI at submillimeter resolutions, potentially revealing new neuroscientific insights at the mesoscopic scale[2].

Acknowledgements

Dr. Natalia Petridou for supplying high-resolution BOLD fMRI data.

References

1. Logothetis, N. K. (2008). What we can do and what we cannot do with fmri. Nature, 453(7197), 869–878. https://doi.org/10.1038/nature06976

2. Vizioli, L., Moeller, S., Dowdle, L., Akçakaya, M., De Martino, F., Yacoub, E., & Ugurbil, K. (2020). Lowering the thermal noise barrier in functional brain mapping with Magnetic Resonance Imaging. https://doi.org/10.1101/2020.11.04.368357
3. Zhao, Y., Yi, Z., Xiao, L., Lau, V., Liu, Y., Zhang, Z., Guo, H., Leong, A. T., & Wu, E. X. (2022). Joint denoising of diffusion-weighted images via structured low-rank patch matrix approximation. Magnetic Resonance in Medicine, 88(6), 2461–2474.https://doi.org/10.1002/mrm.29407
4. Veraart, J., Fieremans, E., & Novikov, D. S. (2015). Diffusion MRI noise mapping using random matrix theory. Magnetic Resonance in Medicine, 76(5), 1582–1593. https://doi.org/10.1002/mrm.26059
5. Veraart, J., Novikov, D. S., Christiaens, D., Ades-aron, B., Sijbers, J., & Fieremans, E. (2016). Denoising of diffusion MRI using random matrix theory. NeuroImage, 142, 394–406. https://doi.org/10.1016/j.neuroimage.2016.08.016
6. Coupe, P., Yger, P., Prima, S., Hellier, P., Kervrann, C., & Barillot, C. (2008). An optimized blockwise nonlocal means denoising filter for 3-D Magnetic Resonance images. IEEE Transactions on Medical Imaging, 27(4), 425–441. https://doi.org/10.1109/tmi.2007.906087
7. Dabov, K., Foi, A., Katkovnik, V., & Egiazarian, K. (2007). Image denoising by sparse 3-D transform-domain collaborative filtering. IEEE Transactions on Image Processing, 16(8), 2080–2095. https://doi.org/10.1109/tip.2007.901238

Figures

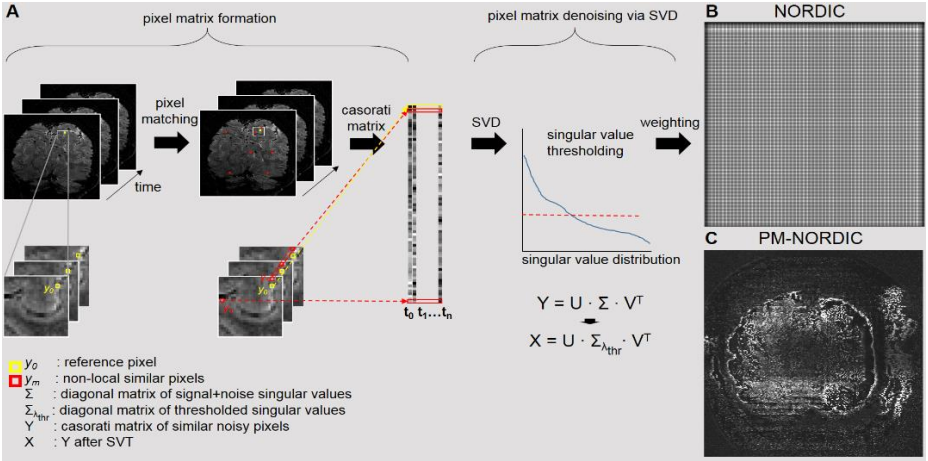

Fig 1. A) Flowchart of PM-NORDIC. For each 2D slice, reference pixels y_0 are compared to M pixels y_m. For each reference pixel, the (k^2)-1 pixels with the highest similarity scores are grouped together with the reference pixel into a (k^2)xQ Casorati matrix Y. The matrix undergoes NORDIC SVT to produce a denoised matrix X. **B) patch-formation pixel weighting matrix** for a 2D slice for standard NORDIC. **C) patch-formation pixel weighting matrix** for a 2D slice for PM-NORDIC.

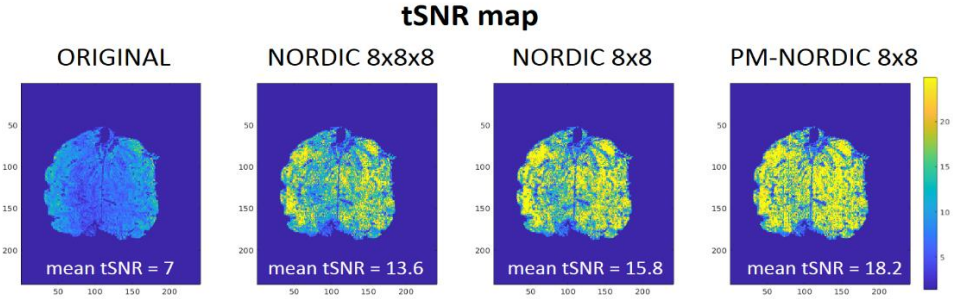

Fig 2. Mean tSNR scores for a representative slice. The proposed PM-NORDIC further increases tSNR compared to standard NORDIC, especially in central areas of the brain where the g-factor amplification is usually higher (and thus higher thermal noise).

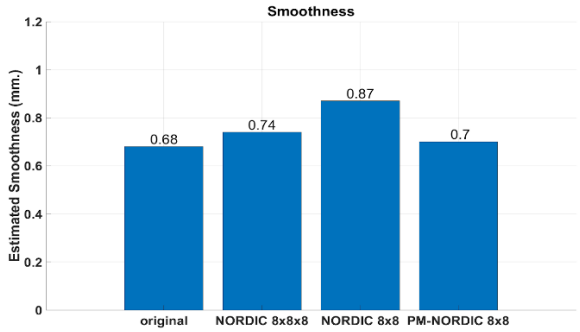

Fig 3. Spatial smoothness estimates (using 3dFWMHx AFNI) for a representative dataset. The proposed PM-NORDIC barely affects the original spatial smoothness of the dataset, indicating improved preservation of structural detail and potentially spatial specificity of functional BOLD activity compared to standard NORDIC.

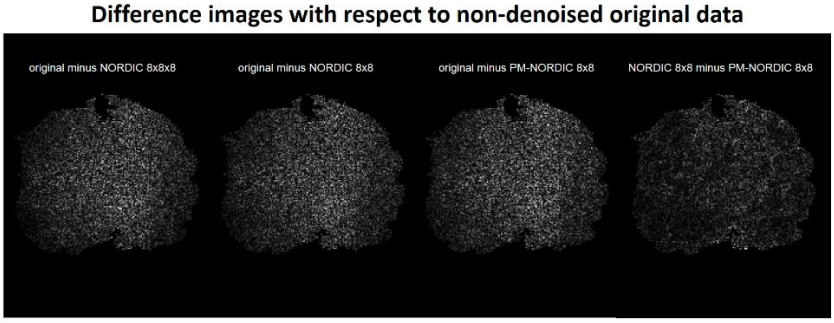

Fig 4. Difference images (with respect to the non-denoised original data) for the different patching

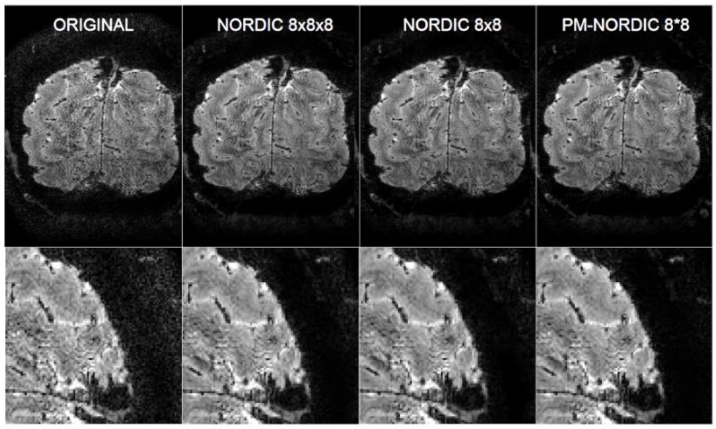

methods for a representative slice. PM-NORDIC shows increased noise reduction compared to standard NORDIC with default and small patching, indicating that no structural information is being removed. The difference between the proposed PM-NORDIC and standard NORDIC reveals that PM-NORDIC removes more thermal noise and subtle vessel outlines.

Fig 5. Magnitude images and zoomed view of a representative slice before and the after denoising. PM-NORDIC effectively removes thermal noise without interfering with the structural information.