Probabilistic ODF Estimation from Reduced HARDI Data with Sparse Regularization

Antonio Tristán-Vega and **Carl-Fredrik Westin**

Laboratory of Mathematics in Imaging, Brigham and Women's Hospital, Boston

Antonio Tristán-Vega: atriveg@bwh.harvard.edu; Carl-Fredrik Westin: westin@bwh.harvard.edu

Abstract

High Angular Resolution Diffusion Imaging (HARDI) demands a higher amount of data measurements compared to Diffusion Tensor Imaging (DTI), restricting its use in practice. We propose to represent the probabilistic Orientation Distribution Function (ODF) in the frame of Spherical Wavelets (SW), where it is highly sparse. From a reduced subset of measurements (nearly four times less than the standard for HARDI), we pose the estimation as an inverse problem with sparsity regularization. This allows the fast computation of a positive, unit-mass, probabilistic ODF from 14–16 samples, as we show with both synthetic diffusion signals and real HARDI data with typical parameters.

1 Introduction

Diffusion Magnetic Resonance Imaging (MRI) provides an unparalleled tool to probe the connectivity of the nerve fibers within the white matter of the brain *in vivo*. At the core of this technique is the relationship between the signal $E(\mathbf{q})$, acquired by the MRI scanner when a pulsed gradient encoded by the wave-vector $\mathbf{q}$ is applied, and the random process $P(\mathbf{R})$ driving the restricted diffusion of water molecules. Such relationship is modeled as the following Fourier transform [1]:

$$P(\mathbf{R}) = \int \int \int_{\mathbb{R}^3} |E(\mathbf{q})| \exp(-j2\pi \mathbf{q} \cdot \mathbf{R}) d\mathbf{q}. \quad (1)$$

The obvious way to recover $P(\mathbf{R})$ is to sample $E(\mathbf{q})$ and approximate eq. (1) as a discrete Fourier transform, what is known as Diffusion Spectrum Imaging [2]. The main drawback of this technique is the need to sample the entire space of wave-vectors $\mathbf{q}$, since each of them implies the acquisition of a whole MRI volume. Though it has been recently shown [3] how this amount of data can be drastically reduced using Compressed Sensing (CS, [4]), this approach still requires several hundreds of samples to attain a reliable reconstruction.

High Angular Resolution Diffusion Imaging (HARDI, [5]) permits reduction of the measurements to a sampling of the space of orientations with constant $\|\mathbf{q}\|$. This is achieved at the expense of losing the radial information in $P(\mathbf{R})$, computing some sort of projection or Orientation Distribution Function (ODF) [2]:

$$\Phi(\mathbf{r}) = \int_{-\infty}^{\infty} P(R\mathbf{r})R^2 dR = \frac{-1}{8\pi^2} \int\int_{\perp\mathbf{r}} \Delta E(\mathbf{q}) d\mathbf{q}. \quad (2)$$

There are some other alternative definitions for the orientation function (see [6] for an excellent review on this topic), but eq. (2) has the advantage of being a true probability law [7]. HARDI typically requires 60–100 wave-vectors **q**, but considerations similar to those in [3] allow relaxation of this requirement to a few tens (16–30) [8, 9], whenever a sparse representation of $E(\mathbf{q})$ may be achieved. This condition is met by representing $E(\mathbf{q})$ in a suitable vector frame, namely Spherical Ridgelets (SR), whose components closely resemble the shape of $E(\mathbf{q})$ at different scales. This frame is adequate in [8] since the object of interest is the diffusion signal. Unfortunately, these plate-like functions do not fit the cigar-shaped appearance of the probabilistic ODF, $\Phi(\mathbf{r})$, and cannot be expected to provide a sparse representation of it.

In **section 2**, we postulate Spherical Wavelets [10] as a suitable frame to sparsely represent the $\Phi(\mathbf{r})$ in eq. (2), as depicted in Fig. 1. In **section 3** we formulate the estimation of the ODF as an inverse problem, by relating the SW functions to their counterparts in the **q**-space. We propose two alternative solutions based on either ℓ_1 or ℓ_2 regularization, justifying that non-negative ℓ_2 provides sparse representations analogous to ℓ_1. Finally, the results in **section 4** illustrate how positive, unit-mass, probabilistic ODFs can be recovered from as few as 14 gradient directions. With ℓ_2 regularization, the minimization problem can be solved with Newton-like methods, becoming computationally efficient.

2 Spherical Wavelets to Sparsely Represent Φ(r)

The shape of the ODF in Fig. 1 implies that a combination of *a few* Spherical Wavelets (SW, [10]) with different widths and orientations may accurately represent $\Phi(\mathbf{r})$ for an arbitrary ensemble of diffusion compartments, i.e. the coefficients of $\Phi(\mathbf{r})$ in its SW expansion are sparse. The domain of definition of SW is the set $\Omega \equiv \{\mathbf{r} \in \mathbb{R}^3 : \|\mathbf{r}\| = 1\}$ of unit directions (i.e. diffusion gradients). Their definition is founded on the basis of Spherical Harmonics (SH), for which the following summation formula holds:

$$\sum_{m=-l}^{l} Y_l^m(\mathbf{u}) Y_l^m(\mathbf{v}) = \frac{2l+1}{4\pi} P_l(\mathbf{u} \cdot \mathbf{v}), \quad (3)$$

where Y_l^m is the real SH basis function of degree $l \geq 0$ and order $-l \leq m \leq l$, and P_l is the Legendre polynomial of degree l. Note $P_l(\mathbf{u} \cdot \mathbf{v})$ is a rotation of $Y_l^0(\mathbf{u})$ to an auxiliary system for which the 'z' axis is aligned with **v**. Any symmetric (i.e. $f(\mathbf{u}) = f(-\mathbf{u})$), square-integrable function $f \in \mathbb{L}_{\text{sym}}^2(\Omega)$, can be written as [10]:

$$f(\mathbf{u}) = \int\int_{\Omega} f(\mathbf{v}) \sum_{l=0,2,\ldots}^{\infty} \frac{2l+1}{4\pi} P_l(\mathbf{u} \cdot \mathbf{v}) d\mathbf{v} = \int\int_{\Omega} f(\mathbf{v}) \Theta_{\mathbf{v}}(\mathbf{u}) d\mathbf{v}, \quad (4)$$

where $\Theta_{\mathbf{v}}(\mathbf{u})$ is a spherical convolution kernel. To achieve a scale-space representation of $f(\mathbf{u})$, the following scaling kernel $K_{\mathbf{v},j}(\mathbf{u})$ and **semi**-discrete frame $\cup_{j=-1}^{\infty}\cup_{\mathbf{v}\in\Omega}\Psi_{\mathbf{v},j}(\mathbf{u})$, that spans $\mathbb{L}_{sym}^2(\Omega)$, are introduced in [10]:

$$K_{\mathbf{v},j}(\mathbf{u})=\sum_{l=0,2,...}^{\infty}\frac{2l+1}{4\pi}\kappa_\rho(-2^{-j}l)P_l(\mathbf{u}\cdot\mathbf{v}), \quad (5)$$

$$\Psi_{\mathbf{v},j}(\mathbf{u})=\begin{cases} K_{\mathbf{v},0}(\mathbf{u}), & j=-1; \\ K_{\mathbf{v},j+1}(\mathbf{u})-K_{\mathbf{v},j}(\mathbf{u}), & j\geq 0, \end{cases} \quad (6)$$

where we choose $\kappa_\rho(x)=e^{-\rho x(x+1)}$. When the resolution $j\to\infty$, $\kappa_\rho(2^{-j}l)\to 1$ for each l, and we have the exact recovery in eq. (4). When $j\to-\infty$, $\kappa_\rho(2^{-j}l)\to 0$ for $l \neq 0$, and we have the DC component of $f(\mathbf{u})$. This reasoning clarifies the multiscale nature of $K_{\mathbf{v},j}$, from coarse ($j\to-\infty$) to fine details ($j\to\infty$).

3 ODF Reconstruction from Sparse Representations

From the representation provided by SW as described in the previous section, we aim to estimate the ODF, $\Phi(\mathbf{r})$, as the sparsest combination of SW that can explain the data acquired in the **q**-space, $E(\mathbf{q})$. Therefore, we need to compute the **q**-space response of each SW defined in $\mathbf{r}\in\Omega$ (section 3.1), and then discretize the problem to account for the discrete nature of the measurements (section 3.2).

3.1 Response of SW $\Psi_{\mathbf{v},j(\mathbf{r})}$ in the Space of Wave-Vectors q

We are interested in probabilistic ODF estimators [7], so that the response pursued obeys the following form:

$$\Phi(\mathbf{r})=\frac{1}{4\pi}+\frac{1}{16\pi^2}\mathscr{G}\{\Delta_b\zeta(E(\mathbf{q}))\}(\mathbf{r}), \quad (7)$$

Among the chances provided in [7], we choose $\zeta(\cdot)=-E_i(-\log(\cdot))$ as it seems to be a trade-off between angular resolution and noise sensitivity. Also, the unit-mass constraint comes for free with this model, since the Beltrami operator Δ_b eliminates all DC components so that we keep only the $1/4\pi$ factor. We denote by $\mathscr{G}$ the Funk-Radon Transform (FRT, [5]).

Since SH are eigenfunctions for both $\mathscr{G}$ and Δ_b, it is straightforward to prove this property holds also for the rotated versions of $Y_l^0(\mathbf{u})$, i.e. $P_l(\mathbf{u}\cdot\mathbf{v})$. If we insert eqs. (5) and (6) into (7), we get the following duality between the SW, $\Psi_{\mathbf{v},j}(\mathbf{r})$, and their counterparts in the **q**-space, $\Xi_{\mathbf{v},j}(\mathbf{q})$ (see Fig. 1):

$$\Psi_{\mathbf{v},j}(\mathbf{r})=\sum_{l=2,4,...}^{\infty}\frac{2l+1}{4\pi}\nu_{\rho,j}(l)P_l(\mathbf{r}\cdot\mathbf{v}) \iff \Xi_{\mathbf{v},j}(\mathbf{q})=\sum_{l=2,4,...}^{\infty}\frac{4\pi(2l+1)}{-l(l+1)}\frac{\nu_{\rho,j}(l)}{\lambda_l}P_l(\mathbf{q}\cdot\mathbf{v}),$$

$$\nu_{\rho,j}(l) = \begin{cases} \kappa_\rho(l), & j = -1; \\ \kappa_\rho(2^{-j-1}l) - \kappa_\rho(2^{-j}l), & j \geq 0, \end{cases} \quad (8)$$

where λ_l and $-l(l+1)$ are respectively the eigenvalues for $\mathscr{G}$ and $_b$. Note we drop the DC term $l = 0$ from the sums, according to our previous statement.

3.2 Optimization Problem with ℓ_p Regularization

Like in [8], we need a fully discrete version of the frame we use for sparse representation: let $\{\mathbf{q}_n\}_{n=1}^N$ be the set of wave-vectors in the HARDI data-set; let $\{\mathbf{r}_m\}_{m=1}^M$ be the set of directions for which the ODF is computed; let $\{\mathbf{v}_{j(k),k}\}_{k=1}^K$ be a collection of discretizations of Ω for each resolution level $-1 \quad j \quad J$. The estimation of $\mathbf{x} \equiv \Phi(\mathbf{r})$ is arranged as the ℓ_p inverse problem:

$$\min_{\mathbf{a}} \|\mathbf{a}\|_p, \text{s.t.} \|A\mathbf{a} - \mathbf{y}\|_2^2 < \eta \quad \left| \quad A_{n,k} = \Xi_{\mathbf{v}_{j(k),k},j(k)}(\mathbf{q}_n); \mathbf{y}_n = E(\mathbf{q}_n); \right.$$
$$\text{with:} \mathbf{x} = B\mathbf{a} + \frac{1}{4\pi} \quad \left| \quad B_{m,k} = \Psi_{\mathbf{v}_{j(k),k},j(k)}(\mathbf{r}_m); \mathbf{x}_m = \Phi(\mathbf{r}_m). \right. \quad (9)$$

In brief, we discretize the responses $\Xi_{\mathbf{v},j}(\mathbf{q})$ into a $N \times K$ *sampling matrix* A to find that $\mathbf{a}$ best describing the $N \times 1$ vector $\mathbf{y} \equiv E(\mathbf{q})$. These coefficients are used to recover $\mathbf{x}$ from the discretized $M \times K$ *sparsifying matrix* B. We deliberately choose the same names used in [4], given the formal similarity of eq. (9) with CS for $p = 1$. Typically, $N \ll K$, so that $A\mathbf{a} = \mathbf{y}$ is highly underdetermined. This is the reason why $\mathbf{a}$ has to be assumed sparse *a priori*. Though sparsity is only enforced in case $p \quad 1$ [4], we also consider here the case $p = 2$ for the reasons detailed below. Note the ℓ_2 problem has a closed-form solution for eq. (9):

$$\mathbf{a} = A^T (AA^T + \tau(\eta) \cdot I_N)^{-1} \mathbf{y}, \quad (10)$$

3.3 Positivity Constraints on the Probabilistic ODF

Since $\Phi(\mathbf{r})$ is a probability law, it has to be positive for all $\mathbf{r}$. Given a set of directions $\{\mathbf{r}'_{m'}\}_{m'=1}^{M'}$, we add M' additional linear constraints to eq. (9):

$$B'\mathbf{a} + \frac{1}{4\pi} \geq 0, B_{m',k} = \Psi_{\mathbf{v}_{j(k),k},j(k)}(\mathbf{r}'_{m'}). \quad (11)$$

Two important remarks arise here regarding the ℓ_2 problem: 1) the solution in eq. (10) no longer holds, but eq. (9) becomes a quadratic program which is still very efficient to solve compared to ℓ_1. 2) Non-negative least squares in fact provide sparse solutions [11], so that we may reasonably hypothesize (see [9]) that the solution of the constrained ℓ_2 will approach that of the constrained ℓ_1.

4 Results

Synthetic Experiments

We generate a two-tensor model simulating two fibers crossing in angles 45°, 60°, or 90°. The results are the average for 300 trials with random rotations of the tensor ensemble, with the diffusion signal generated as:

$$S(\mathbf{u}) = \sqrt{\left(\frac{1}{2}\exp\left(-b\mathbf{u}^T\mathbf{D}_1\mathbf{u}\right) + \frac{1}{2}\exp\left(-b\mathbf{u}^T\mathbf{D}_2\mathbf{u}\right) + \sigma \cdot n_c\right)^2 + (\sigma \cdot n_s)^2}, \quad (12)$$

for eigenvalues of $\mathbf{D}_n$: $[1.7, 0.3, 0.3] \cdot 10^{-3}$ mm²/s, $b = 2,000$ s/mm², $n_{c,s} \sim \mathcal{N}(0, 1)$, and Signal to Noise Ratios (SNR): $1/\sigma = 100, 40$.

All the methods are coded in MatLab®, using the ℓ_1-magic library (http://www.acm.caltech.edu/l1magic/) to solve ℓ_1 and quadprog to solve ℓ_2.

The parameters are fixed following the guidelines in [8]: a width $\rho = 0.75$ of the kernel in eq. (5) is chosen by trial and error for best performance, and so it is the error threshold η in eq. (9) (its value depends on the scenario). For the directions $\mathbf{r}$, $\mathbf{r}'$, and $\mathbf{q}$, we choose regular samplings of Ω in all cases, with respective set sizes $M = 214$, $M' = 50$, and N (the number of HARDI measurements) ranging from 12 to 28. For the discretization of the SW frame, we uniformly sample Ω at each level $-1 \le j \le J$, taking $(2^{j+1}m_0 + 1)^2$ evenly spaced samples, with $m_0 = 4$ and $J = 1$.

For the sake of comparison, we have adapted the approach in [8] (namely, SR) to interpolate $\zeta(E)$, and used eq. (7) to evaluate the probabilistic ODF. We have also adapted the method described in section 3 by alternatively representing the ODF in the basis of SH with $l \le 4$, where $\Phi(\mathbf{r})$ **is not** sparse[1]. The comparison is based on the relative Mean Squared Error (rMSE) between the recovered ODF and the ground-truth, computed using eq. (7) over a SH fitting ($l \le 8$) of 214 noise-free $\mathbf{q}$-space samples, see [7, Table 2] for details.

Discussion

Numerical results are shown in Fig. 2, using the optimal parameters in each case. It is worth to stress the following remarks:

1)   With ℓ_1 regularization, our approach (SW) provides the lowest reconstruction errors (see Fig. 2.a, top). Though this numeric difference may seem subtle, the ODFs reconstructed in Fig. 2.c suggest it might be critical to correctly resolve the fiber crossings in poor SNR scenarios.

2)   As the number of gradients acquired increases, the rMSE with SW and SR rapidly decreases, but not with SH. Since the ODFs are not actually sparse in the SH basis (see Fig. 2.b), the *a priori* sparse behavior does not hold.

---

[1]We cannot reproduce the respective derivations due to space limits, but they are straightforward by updating eq. (8) using that SH are eigenfunctions of $\mathcal{G}_{\cdot,b}$.

3) From Fig. 2.a, bottom, our SW approach provides very similar rMSE (and visually identical ODFs, see Fig. 2.c) with either ℓ_1 or ℓ_2 regularization. As stated before, non-negative ℓ_2 yield sparse solutions [11], close to those with ℓ_1 in the space of coefficients. We show in Appendix A that SW, as used here, nearly meet the Parseval property: if the ℓ_1 and ℓ_2 solutions are alike, the Euclidean distance between the corresponding ODFs will also be small. The SR in [8] do not exhibit this feature. Moreover, according to eq (7), a differential operator $\mathcal{L}_b$ is applied to compute Φ from $\zeta(E)$, so the difference in the reconstruction of $\zeta(E)$ (which **is not** bounded by the distance between the ℓ_1 and ℓ_2 solutions) is amplified when computing the ODF. Accordingly, the corresponding errors in Fig. 2.a, bottom (ℓ_2) for [8] are far larger compared to the ℓ_1 case.

4) The minimum number of gradients to attain a reliable ODF reconstruction seems to be 14–16 (Fig. 2.a, top and bottom), and below this threshold the rMSE blows. With SW, this range coincides with the minimum number of non-zero coefficients of the solution (see remark 5), so we may hypothesize that a minimum of 3% of the SW coefficients are required to properly describe ODFs.

5) Our method, SW, provides also the sparsest solution (Fig. 2.b). As the number of measurements grows, we have more information on the ODF, so the solution may be refined with more SW coefficients and it becomes less sparse.

In vivo Experiments

A HARDI data set with 8 unweighted T2 baseline images and 16 evenly spaced gradients ($b = 3,000$ s/mm^2) is used in this section.

Fig. 3 shows an axial slice intersecting the corpus callosum (red lobes aligned with the 'x' axis) and the uppermost part of the cingulum (green lobes, 'y' axis).

We are able to correctly resolve the crossings between these structures in the top row, or even the complex three-compartment crossings in the bottom row (blue lobes along the 'z' axis correspond to the corona radiata).

While for ℓ_1 the improvement with our method over [8] is not dramatic, the reconstructed ODFs are still sharper (see blue square), according to remark 1.

With ℓ_2 regularization, the method in [8] completely fails, while our method outputs accurate ODFs, and even the three-compartment crossings are better resolved than with ℓ_1 (see green square).

This is in agreement with remark 3. Our final remark is that:

6) ℓ_1-based reconstruction takes ~200 s, while ℓ_2-based takes ~1.5 s, for Fig. 3 (the speedup is ~130), which is an additional argument supporting ℓ_2.

5 Conclusion

We have shown that as few as 14–16 diffusion directions suffice to obtain accurate probabilistic ODFs estimates. Positivity constraints, which are not considered in [7], can be easily incorporated into our model. Although sparse regularization has been previously used in [8] to fit the diffusion signal, our approach is, to our knowledge, the very first to use this technique to reconstruct the ODF from HARDI data with reduced measurements. Though we have discussed in the results section a work-around to compute the ODF from the diffusion signal provided by [8], the results with our own approach are clearly better with ℓ_1, and the only usable at all with ℓ_2. This is specially important, since ℓ_2 reconstruction is two orders of magnitude faster than ℓ_1: for this latter approach, the overwhelming duration of the estimation for an entire volume (in the order of 100 days with our Matlab implementation) makes the overall acquisition/reconstruction time unattractive. Our ℓ_2 approach provides virtually identical results as ℓ_1 with reasonable reconstruction times, while allowing an acquisition speedup in the order of 4.

Acknowledgments

Appendix

A Parseval Property in the Semi-Discrete SW Frame

A frame is said to be tight or Parseval if the projections of a vector f in each frame element conserve the energy of f. If C is the frame bound, c_l^m the coefficient of $f(\mathbf{u})$ for the SH basis function $Y_l^m(\mathbf{u})$, and $f_l(\mathbf{u})$ the orthogonal projection of $f(\mathbf{u})$ in the subspace of SH of degree l, this condition reads:

$$\sum_{j=-1}^{\infty} \int \int_{\Omega} |\langle f(\mathbf{u}), \Psi_{\mathbf{v},j}(\mathbf{u})\rangle|^2 d\mathbf{v} \overset{?}{=} C\|f(\mathbf{u})\|_2^2 \quad (13)$$

for the semi-discrete frame of SW. From its definition in eq. (6), eq. (13) reads:

$$\sum_{j=-1}^{\infty} \int \int_{\Omega}^{d\mathbf{v}} \left| \left\langle f(\mathbf{u}), \sum_{l=0,2,\ldots}^{\infty} \nu_{\rho,j}(l) \sum_{m=-l}^{l} Y_l^m(\mathbf{u})Y_l^m(\mathbf{v})\right\rangle\right|^2$$

$$= \sum_{j=-1}^{\infty} \int \int_{\Omega}^{d\mathbf{v}} \left| \sum_{l=0,2,\ldots}^{\infty} \nu_{\rho,j}(l) \sum_{m=-l}^{l} c_l^m Y_l^m(\mathbf{v})\right|^2 = \sum_{j=-1}^{\infty} \sum_{l,l'=0,2,\ldots}^{\infty} \nu_{\rho,j}(l)\nu_{\rho,j}(l') \quad (14)$$

$$\sum_{m,m'=-l}^{l} c_l^m c_{l'}^{m'} \left\langle Y_l^m, Y_{l'}^{m'}\right\rangle = \sum_{j=-1}^{\infty} \sum_{l=0,2,\ldots}^{\infty} \nu_{\rho,j}(l)^2 \sum_{m=-l}^{l} (c_l^m)^2 = \sum_{l=0,2,\ldots}^{\infty} s_l\|f_l\|_2^2.$$

The condition (13) is equivalent to $s_l = \sum_{j=-1}^{\infty} \nu_{\rho,j}(l)^2$ (see eq. (8)) being constant for all l. Fig. 4 shows the partial sums for the first few l. SW are *near* Parseval (i.e. the curves

approach the same value when $j \to \infty$) if $l = 0$ is ignored, but this is the DC component, unnecessary for our derivations.

References

1. Callaghan, P. Principles of Nuclear Magnetic Resonance Microscopy. Oxford: Clarendon Press; 1991.

2. Wedeen VJ, Hagmann P, Tseng W-YI, Reese TG, Weisskoff RM. Mapping complex tissue architecture with DSI. Mag. Res. Med. 2005; 54:1377–1386.

3. Merlet, S.; Deriche, R. MICCAI Workshop Comp. Beijing (China): Diffusion MRI; 2010 Sep. Compressed sensing for accelerated EAP recovery in diffusion MRI.

4. Donoho DL. Compressed sensing. IEEE Trans. Info. Th. 2006; 52(4):1289–1306.

5. Tuch DS. Q–Ball imaging. Mag. Res. Med. 2004; 52:1358–1372.

6. Assemlal H-E, Tschumperlé D, Brun L, Siddiqui K. Recent advances in diffusion MRI modeling: Angular and radial reconstruction. Med. Im. Anal. 2011

7. Tristán-Vega A, Westin C-F, Aja-Fernández S. A new methodology for the estimation of fiber populations in the white matter of the brain with the Funk-Radon transform. NeuroIm. 2010; 49:1301–1315.

8. Michailovich, O.; Rathi, Y. Fast and accurate reconstruction of HARDI data using compressed sensing. In: Jiang, T.; Navab, N.; Pluim, JPW.; Viergever, MA., editors. MICCAI. Vol. 6361. Heidelberg: Springer; 2010. p. 607-614.LNCS

9. Jian, B.; Vemuri, BC. Multi-fiber reconstruction from diffusion MRI using mixture of wisharts and sparse deconvolution. In: Karssemeijer, N.; Lelieveldt, B., editors. IPMI. Vol. 4584. Heidelberg: Springer; 2007. p. 384-395.LNCS

10. Freeden W, Schreiner M. Orthogonal and non-orthogonal multirresolution analysis, scale discrete and exact fully discrete wavelet transform on the sphere. Const. Approx. 1998; 14:493–515.

11. Lawson, CL.; Hanson, RJ. Solving Least Squares Problems. Englewood Cliffs: Prentice-Hall; 1974.
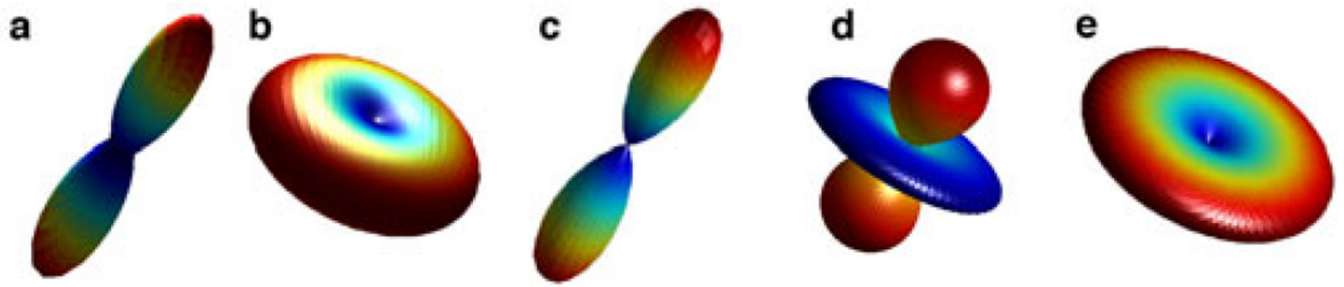
Fig. 1.
Examples of: **a)** the true $\Phi(\mathbf{r})$ when $P(\mathbf{R})$ is Gaussian; **b)** its corresponding diffusion signal $E(\mathbf{q})$; **c)** a SW, $\Psi_{\mathbf{v},j}(\mathbf{r})$. These wavelets aim to sparely represent signals of the type a), which seems quite reasonable; **d)** the response $\Xi(\mathbf{q})$ to $\Psi_{\mathbf{v},j}(\mathbf{r})$ in the space of measurements; **e)** a SR, which in [8] aims to sparsely represent signals of the type b). Note this shape prevents the use of SR to represent ODFs like a).
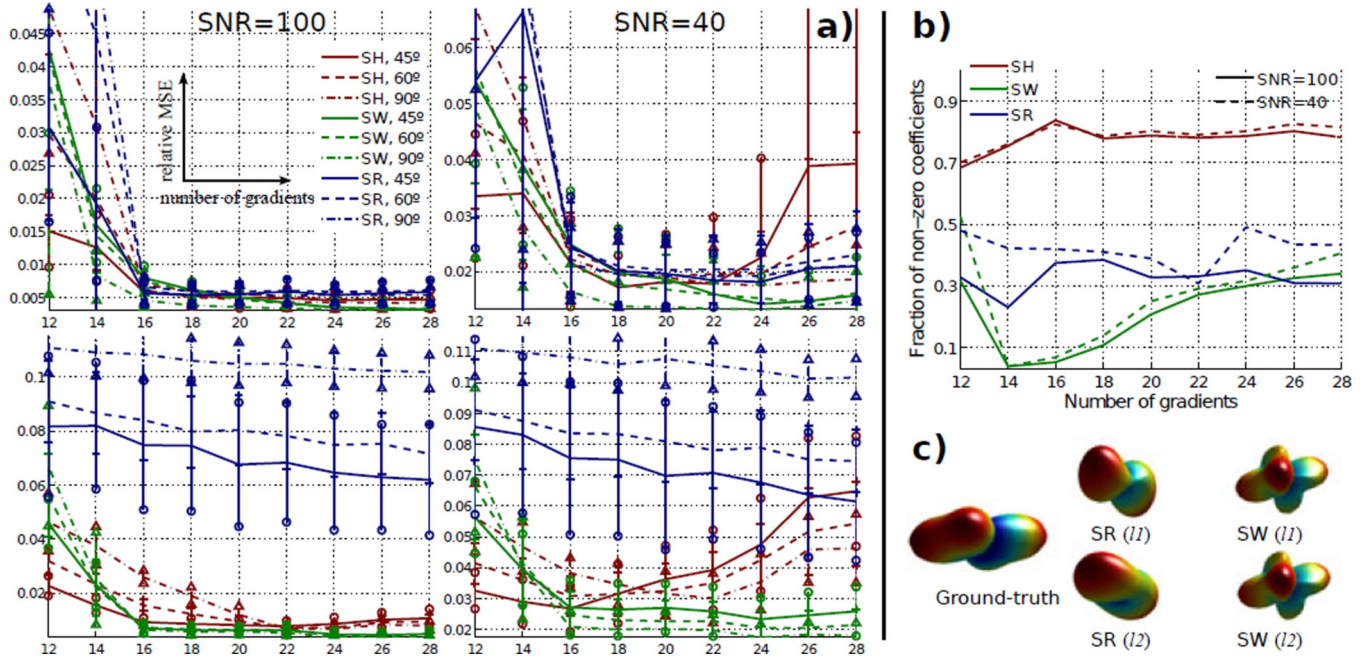
Fig. 2.
Results obtained with our approach (SW), our adaptation of [8] to estimate ODFs (SR), and
an implementation of our method based on SH. **a)** Mean and variance of the rMSE for ℓ_1
(top row) and ℓ_2 (bottom row). **b)** Sparsity of the solutions for a 60° crossing. **c)** Examples
of the ODFs reconstructed for SNR=40 and a 60° crossing.
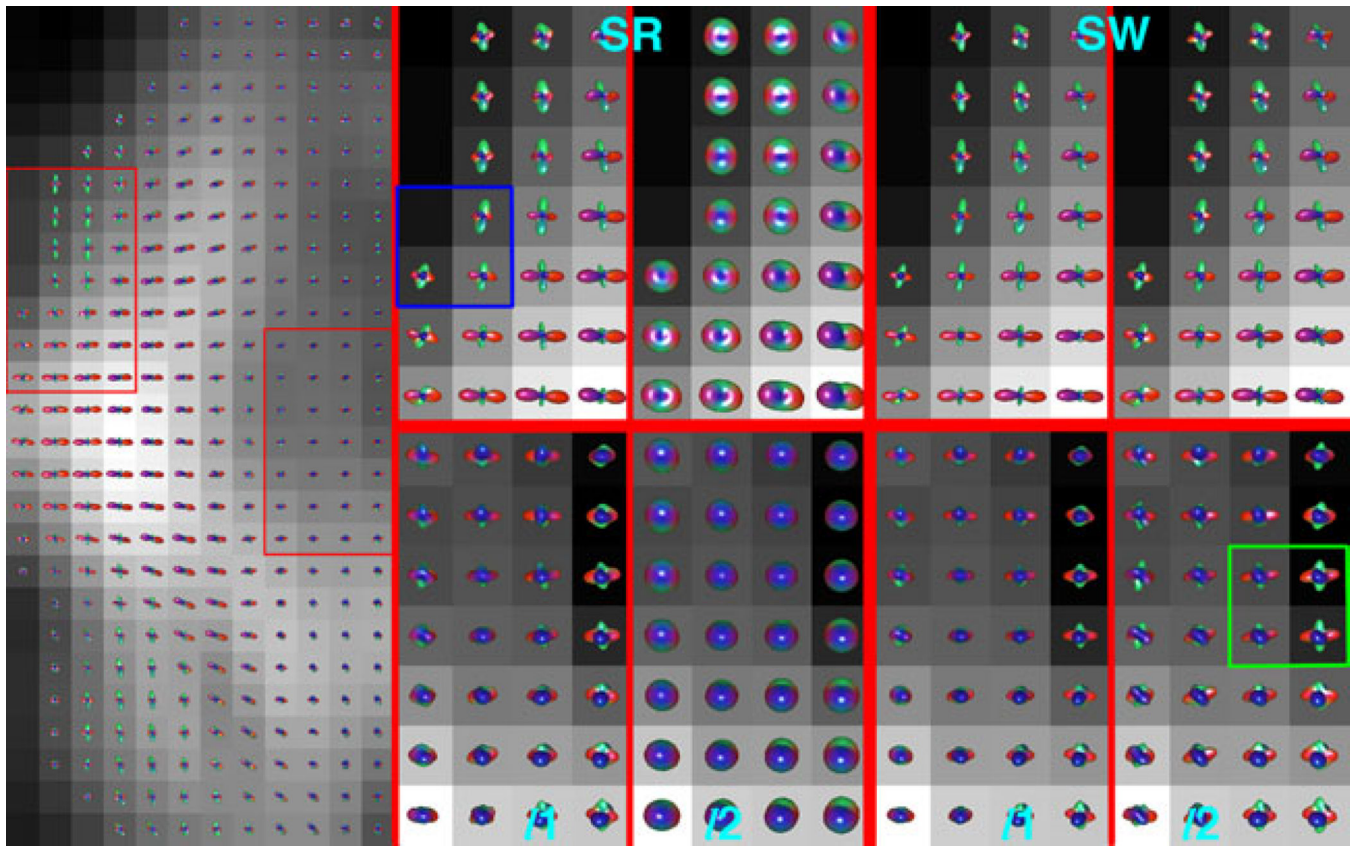
Fig. 3.
Axial slice of a real data set with 16 gradients. The ODFs field has been recovered with our method. The highlighted regions have been further analyzed with both our method (SW) and the method in [8] (SR), with either ℓ_1 or ℓ_2 regularization.
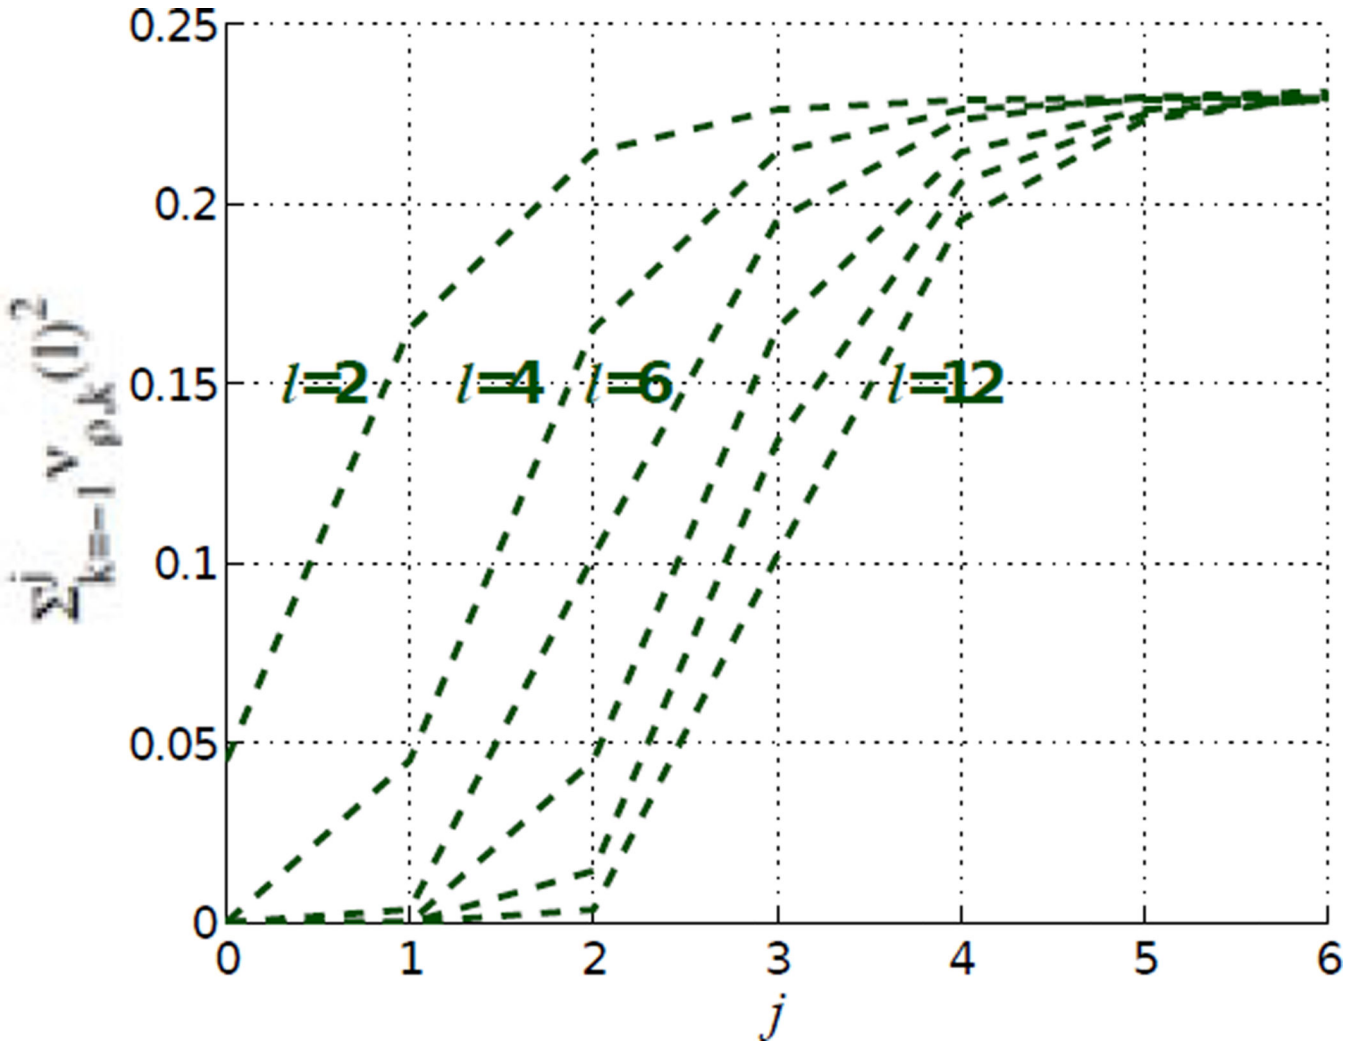
Fig. 4.

Partial sums $\sum_{k=-1}^{J} \nu_{\rho,k}(l)^2$ for different degrees of l (for $l = 0$, the curve is a line with constant value 1). We use $\rho = 0.75$ as in all the experiments throughout the results section (this behavior does not hold for all ρ).