Metric Selection and Diffusion Tensor Swelling

Ofer Pasternak, Nir Sochen and Peter J. Basser

Abstract The measurement of the distance between diffusion tensors is the foundation on which any subsequent analysis or processing of these quantities, such as registration, regularization, interpolation, or statistical inference is based. Euclidean metrics were first used in the context of diffusion tensors; then geometric metrics, having the practical advantage of reducing the "swelling effect," were proposed instead. In this chapter we explore the physical roots of the swelling effect and relate it to acquisition noise. We find that Johnson noise causes shrinking of tensors, and suggest that in order to account for this shrinking, a metric should support swelling of tensors while averaging or interpolating. This interpretation of the swelling effect leads us to favor the Euclidean metric for diffusion tensor analysis. This is a surprising result considering the recent increase of interest in the geometric metrics.

1 Introduction

Metric selection defines how we compare entities and is the basic step for most, if not all, processing methods. In diffusion tensor imaging (DTI), the entity of interest is the diffusion tensor [5], and a metric for the diffusion tensors has to be explicitly provided or implicitly assumed for all related processing, including averaging, interpolations, registration methods, clustering and any statistical inference being done with the tensor quantities. DTI was introduced over 15 years ago, and in that

O. Pasternak
Brigham and Women's Hospital. Harvard Medical School, Boston MA., USA, e-mail: ofer@bwh.harvard.edu

N. Sochen
Department of Applied Mathematics, Tel Aviv University, Tel Aviv, Israel

P.J. Basser
The Eunice Kennedy Shriver National Institute of Child Health & Human Development (NICHD), National Institutes of Health, Bethesda, MD, USA

time the method has successfully delineated white matter structures in the brain, recognized many types of brain disorders, and established connectivity measures that might help us understand how the brain functions [3]. As the use of DTI increases, the accuracy of the measurements becomes more important; researchers would now like to distinguish subtle differences reflected in DTI derived quantities, or that require the grouping of a large cohort of subjects to increase the significance of the observed results [17]. Making an appropriate metric selection therefore becomes more critical in order to support a more robust basis for any subsequent analysis.

A rotation-invariant Euclidean metric for diffusion tensors was the initial metric used for proposing a tensor-variate statistical framework [6]. It was followed by the introduction of affine-invariant geometric metrics for diffusion tensors that were designed to account for tensors with only positive eigenvalues, and that had the practical advantage of reducing the "swelling effect" [8, 24, 21, 18, 13]. Swelling is described as an increase in the determinant of a tensor, when this tensor is obtained by interpolating or averaging tensors with a lower determinant.

In our previous work we were able to link a Euclidean metric with diffusion tensors through the expected distribution caused by thermal acquisition noise [23]. At the same time our theoretical framework predicted that geometric metrics, such as the affine-invariant metric and the log-Euclidean metric, would cause bias in diffusion tensor estimations and processing. Nevertheless, there is one aspect – swelling – in which the geometric metric has consistently been shown to provide superior results compared with the Euclidean metric [2]. In this chapter we extend our previous work to find why swelling occurs in the Euclidean analysis; how is it circumvented using the geometric metrics; and which of these two approaches is more appropriate for the analysis of diffusion tensors. The chapter briefly introduces diffusion tensor metrics, the problem of metric selection, and tensor swelling. In Sect. 3 we discuss the determinant and trace as invariants of the different metrics, and provide physical considerations for their preservation. In Sect. 4 we relate the invariants to the swelling problem through the effect of acquisition noise, and demonstrate that trace preservation through the Euclidean metric reduces statistical biases which are encountered when preserving the determinant using the geometric metrics. We conclude in Sect. 5, and provide insights into possible future improvements to diffusion metrics.

2 Riemannian Metrics for Diffusion Tensors

A Riemannian metric, $\mathbf{G}(\mathbf{x}) = \{g_{ij}(\mathbf{x})\}$, defines the infinitesimal distance over a Riemannian manifold [11] as:

$$ds^2 = d\mathbf{x}^T \mathbf{G}(\mathbf{x}) d\mathbf{x},$$

where x is the coordinate of a point on the manifold for a chosen coordinate system. Any positive-definite and symmetric metric is admissible. The distance function is defined as the geodesic, i.e., the shortest path on the manifold. To define the geometric distance between tensors, a metric and a local coordinate system for tensor representation are chosen.

2.1 The Euclidean and Geometric Metrics

Two main families of Riemannian metrics are commonly used with diffusion tensors: the Euclidean metric family and the geometric metric family. The Euclidean family places the diffusion tensors on a Euclidean manifold and its most common member defines the metric to be a constant $\mathbf{G}(\mathbf{x}) = \mathbf{I}$, resulting in

$$ds^2 = tr((d\mathbf{D})^T d\mathbf{D}),$$

where $\mathbf{D}$ is the tensor coordinates in the canonical tensor coordinate system, and tr denotes the matrix trace, which is equivalent to summing the eigenvalues:

$$tr(\mathbf{D}) = \sum_i \lambda_i.$$

The geodesic between any two tensors, $\mathbf{D}_1$ and $\mathbf{D}_2$, with this metric, is simply a straight line, or the Euclidean distance

$$Dist_{Euc}(\mathbf{D}_1, \mathbf{D}_2) = ||\mathbf{D}_1 - \mathbf{D}_2|| \,, \tag{1}$$

where $||\cdot||$ denotes the Frobenius norm. The Euclidean metric is defined over the entire space of symmetric matrices and is rotation-invariant, which makes it invariant for the selection of orthogonal coordinates, but not for the selection of non-orthogonal tensor coordinate systems.

The geometric metric restricts the distance function to be affine-invariant (which includes rotation, scale, shear, and inversion invariance), and operates only on tensors belonging to the space of positive definite symmetric matrices, S^+ [8, 24, 21, 18, 13, 14]. The Affine-invariant metric [24], a Riemannian metric that satisfies these requirements, has an infinitesimal distance [20]

$$ds^2 = tr((\mathbf{D}^{-1} d\mathbf{D})^2.$$

Since this distance is affine-invariant it does not depend on the choice of tensor coordinate system. The corresponding geodesic is found by integration [20]:

$$Dist_{Aff}(\mathbf{D}_1, \mathbf{D}_2) = \sqrt{tr\left(log^2(\mathbf{D}_1^{-1}\mathbf{D}_2)\right)} \,. \tag{2}$$

The notation $log(\mathbf{D})$ stands for the matrix logarithm. The Log-Euclidean metric with its corresponding geodesic [2],

$$Dist_{LogEuc} = ||log(\mathbf{D}_1) - log(\mathbf{D}_2)|| \,, \tag{3}$$

was proposed as an efficient approximation for the computationally demanding Affine-invariant metric.

2.2 Metric Selection

In our previous work we have shown that the Euclidean metrics are related to a normal distribution and that the geometric metrics are related to a log-normal distribution [23]. We further showed that many of the relevant variability sources in the acquisition lead to normal distribution of the diffusion coefficients or any linear combination of them. We concluded that an unbiased estimator for diffusion quantities should be based on the Euclidean metric, and demonstrated this point with synthetic and real datasets where diffusivity and variability estimations were biased when using a geometric metric, and less so when using the Euclidean metric. This lead us to the conclusion that the Euclidean metric is a more appropriate choice for diffusion tensor analysis.

But in practice, previous studies pointed out that the main effect of selecting a geometric metric rather than a Euclidean metric is encountered when interpolating or averaging between two anisotropic tensors [8]. The Euclidean metric does not preserve the determinant (which is proportional to the volume of the ellipsoid described by the tensor) and, as a result, the interpolated tensor may have a determinant larger than the initial tensors, i.e., it may be "swollen." With the introduction of the Affine-invariant and Log-Euclidean metrics it was shown that the swelling effect is reduced [2]. In practice, the swelling effect is usually obviated by applying piecewise smoothed operators, or pre-segmentation that will avoid interpolating initially distant tensors [14]. In theory, it is still interesting to understand why the swelling effect occurs.

3 Determinant vs. Trace

The determinant of a tensor is proportional to the volume of an ellipsoid. It is a rotation-invariant measure, and is equivalent to the product of the eigenvalues:

$$det(\mathbf{D}) = \prod_i \lambda_i \,.$$

It was noticed early on that the geometric metrics preserve the determinant of interpolated tensors, while a Euclidean metric preserves the trace of the tensors [8]. Figure 1 demonstrates the difference between preserving the trace and preserving the determinant. All blue ellipsoids have the same trace ($\mathrm{tr} = 0.8 \times 10^{-3}\ mm^2/s$), while all red ellipsoids have the same determinant ($\det = 0.24 \times 10^{-3}\ (mm^2/s)^3$). Clearly both trace and determinant do not preserve the fractional anisotropy (FA), as the tensors range from elongated to isotropic. When preserving the trace, the

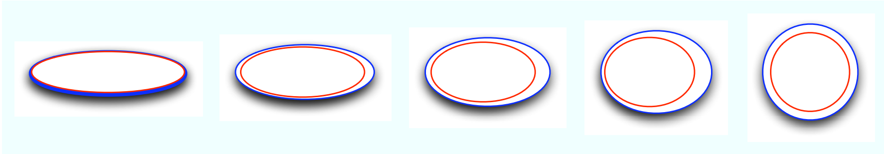

Fig. 1 Determinant Vs. Trace. All blue tensors have the same trace while all red tensors have the same determinant as the tensor on the left. Tensors with a preserved trace have a higher volume (swollen) than those with a preserved determinant.

tensors are indeed more "swollen", i.e., their volume, or determinant grows. Deciding between the geometric metric and the Euclidean metric therefore first requires preference for trace over determinant preservation.

3.1 Physical Considerations

The Einstein equation for free diffusion caused by Brownian motion establishes the fundamental defining relationship between the diffusion coefficient and the mean-squared displacement along an axis [10]:

$$\sigma^2(t) = \mathrm{E}[(\mathbf{x}_t - \mathbf{x}_0)^2] = 2dt \ . \tag{4}$$

The position along the axis at time t is $\mathbf{x}_t$; $\mathbf{x}_0$ is the position at the origin. This relationship defines the diffusion coefficient, d, as proportional to the variance of particle displacements, $\sigma^2(t)$, at time t, and arises from the normal distribution of particles expected for Brownian motion, $\mathbf{x}_t \sim N(\mathbf{x}_0, \sigma^2(t))$. The diffusion tensor is a 3D generalization of the diffusion coefficient [9]: it compactly quantifies the variance of particle displacements along any axis, with eigenvalues that are the diffusion coefficients along a set of three orthogonal axes, described by the eigenvectors and aligned with the maximal variance (or maximal diffusivity) orientations [4]. Therefore the diffusion coefficient, any element of the diffusion tensor, the eigenvalue of the diffusion tensor, and the trace of the tensor are all measures of diffusion with units of $(distance^2/time)$. The determinant of a matrix is a coefficient that describes a scale factor. In the case where the matrix describes distances, the determinant describes a volume, but in the diffusion tensor case, the determinant – with units of $(distance^2/time)^3$ – describes the volume of the ellipsoid that represents the diffusion tensor, not the volume of diffusion itself. Thus there is a clear physical distinction between the trace and the determinant of a diffusion tensor, which should be accounted for when deciding which of these two to preserve.

When introducing the Log-Euclidean metric for diffusion tensors, Arsigny et al. [2] identified the determinant of the diffusion tensor as a *"direct measure of the dispersion of the local diffusion process,"* the preservation of which is critical since *"introducing more dispersion in computations amounts to introducing more diffu-*

sion, which is physically unrealistic." We agree that a computation should preserve the amount of dispersion of the local diffusion process, however, we believe that in the context of diffusion tensors, the determinant is not directly measuring dispersion. Moreover, the Einstein equation above (Eq. 4) shows that the direct measure of dispersion (variance) is the diffusion quantity itself captured by the diffusion tenor, its eigenvalues or their summation as the trace. Indeed the trace – not the determinant – quantifies dispersion, and introducing more diffusion by computations is physically unrealistic, encouraging the preservation of trace over the preservation of determinant.

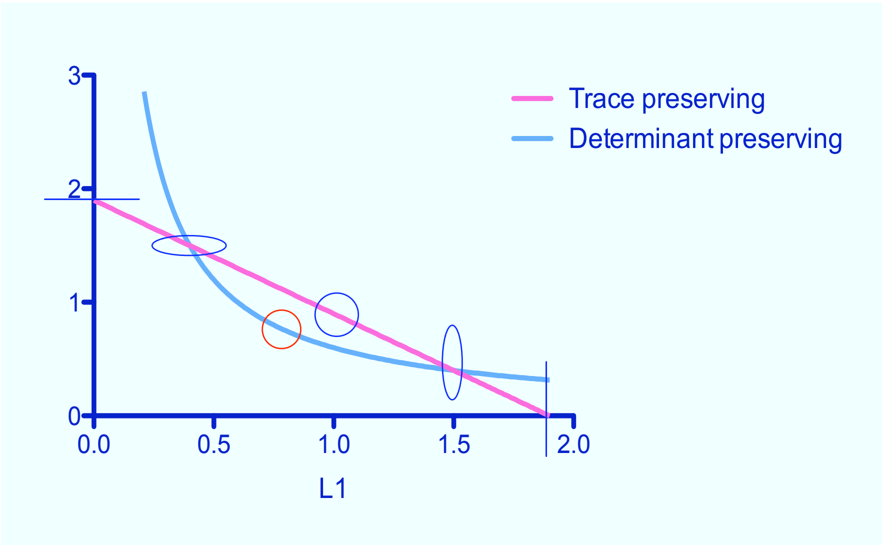

Fig. 2 Interpolations. The graph shows the first (λ_1) over the second (λ_2) eigenvalues of a 2D tensor interpolated while preserving trace, and while preserving determinant. The two original tensors are diag[1.5 0.4] and diag[0.4 1.5] and are shown along with the isotropic tensors that the interpolation goes through. When further extrapolating while preserving the trace, a degenerate tensor (line) with a single finite eigenvalue is found. When extrapolating while preserving the determinant, a tensor with an infinitely large principal-eigenvalue and infinitely small minor-eigenvalue is predicted; this kind of tensor is experimentally not feasible.

Figure 2 further illustrates that preserving the determinant may lead to physically unrealistic results, unlike preserving the trace. The figure shows trace-preserving and determinant-preserving interpolations between two 2D tensors, diag[1.5 0.4] and diag[0.4 1.5]. Extrapolating beyond these tensors, the trace preserving extrapolation predicts a degenerate tensor (line) restricted to a single orientation having a single and finite eigenvalue; physically it predicts that when restricting diffusion to a single orientation, the average displacement will be larger than any other non-restricted orientation, yet finite. On the other hand, the determinant-preserving extrapolation predicts that one of the eigenvalues diminishes to an infinitesimally small value while the other eigenvalue grows to an infinitesimally large value – physically this means that it is predicted that when restricting diffusion to a single orientation the diffusion coefficient grows to infinity, meaning that particles will on average diffuse infinitely fast. This is clearly physically infeasible.

4 Tensor Swelling

In the previous section we concluded that preserving the trace of a diffusion tensor has a better physical justification than preserving the determinant of a diffusion tensor. In this section we further explore the practical implications of using the trace preserving Euclidean metric versus the determinant-preserving Log-Euclidean metric, concentrating on the main effect of the Log-Euclidean metric [2], namely reducing tensor swelling.

4.1 Variability Caused by Johnson Noise

In order to test how acquisition noise affects the determinant and trace we ran a Monte Carlo simulation that creates multiple noisy realizations of a given tensor. The methods follow the synthetic experiments reported in [23, 22]: noisy replications of tensor images were generated by selecting a reference tensor and introducing Johnson noise (Rician distributed) using Monte Carlo simulations to reproduce noisy diffusion weighted images. Tensors were fitted to each replicate using a conventional tensor estimation [7] and using the fitting procedure in [12] that ensures positive definite tensors. We note that when the conventional fitting procedure yielded a positive tensor, this tensor was identical to the one obtained by the positive restricting fitting, hence the difference between the two fitting procedures is restricted only to tensors with one or more negative eigenvalues. Trace and determinant were then calculated for each noisy replicate and the collection of all values was plotted as a probability distribution function.

Figure 3 shows the distribution of determinants and traces for 100K noisy replicates of an anisotropic tensor representing white matter that has eigenvalues 1.5, 0.4 and 0.4×10^{-3} mm^2/s, i.e., with a determinant of 0.24×10^{-3} $(mm^2/s)^3$ and a trace of 2.3×10^{-3} mm^2/s and of an isotropic tensor, representing gray matter, that has eigenvalues of 0.8×10^{-3} mm^2/s, i.e., a determinant of 0.512×10^{-3} $(mm^2/s)^3$ and a trace of 2.4×10^{-3} mm^2/s. As expected noise can increase or decrease both trace and determinant. However, while the trace shows equal probability to be either higher or lower than the initial trace value, the determinant shows a tendency to shrink following the introduction of Johnson noise. For the anisotropic tensor, 51.22% of the trace values are below the initial trace value, while 70.62% of the tensors shrink with determinant values that are below the initial determinant value; for the isotropic tensor, 49.999% of the trace values are below the initial trace value

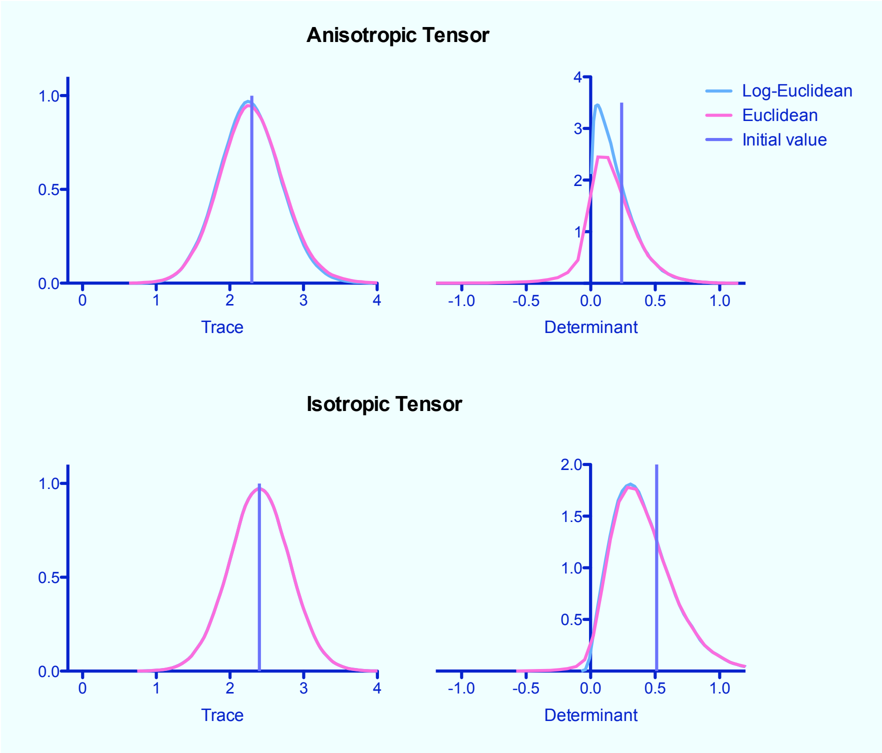

Fig. 3 Distributions. Trace and determinant probability density functions in the presence of noise for a given anisotropic tensor, and a given isotropic tensor. Trace tends to have equal proability to have higher or lower values than the initial value. The determinant is more likely to have a lower value compared to the initial value, i.e., the tensors are expected to shrink.

and 70.5% of the tensors shrink. When restricting the fitting of the noisy tensors to be positive, the determinant distribution for the anisotropic tensor changes considerably, compared with the non-restricting fitting, yet the tendency to have a lower determinant than the initial value remains, as 69.23% of the tensors shrink. For the isotropic tensor, representing gray matter, the number of negative eigenvalue occurrences decreases, and the two fits are very much aligned with each other both for the trace distribution and for the determinant distribution. Similar to the anisotropic case, the determinant in the isotropic case is more likely to have a lower value than the initial value following the introduction of Johnson noise.

The conclusion of this finding is that for these two types of tensors – representing white and gray matter – the introduction of noise causes shrinking, i.e., a decrease in the determinant. This finding is not dependent on the type of tensor fitting used.

4.2 The Extent of Tensor Shrinking in the Brain

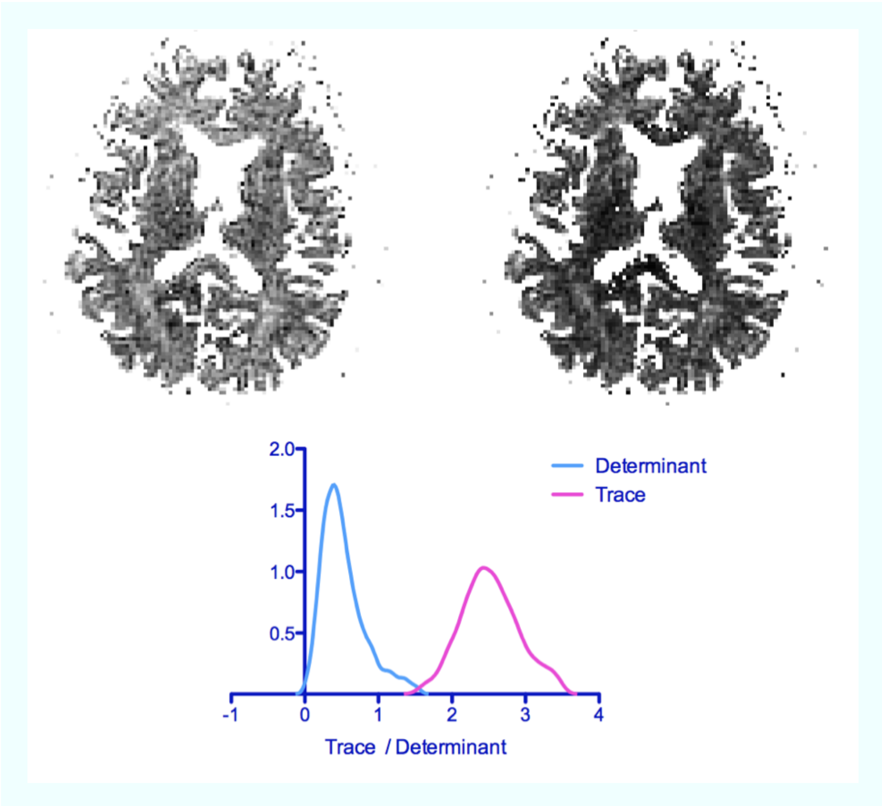

Fig. 4 Ground truth. A slice of diffusion data was taken and fitted to positive tensors, then masked to obtain only gray and white matter. The determinant (left) and trace (right) maps have similar contrast that does not distinguish gray matter from white matter. The slice contains tensors ranging around a trace value of 2.4 and a determinant value of 0.4 .

The finding in the previous section suggests that tensors have the tendency to shrink as a result of thermal acquisition noise, yet in order to get a better notion of the extent of the shrinking in a more realistic setting we will next look at the effect of noise over an entire brain slice. We took a diffusion imaging dataset, and fitted it with the positive restricting method to yield a positive tensor field. We treat this field as a given "ground truth" for a Monte Carlo simulation, where 100 noisy replicates for each voxel were generated.

Figure 4 shows the determinant and trace maps for the ground truth tensors (maps are masked to include only gray and white matter). The contrast of both maps is very similar, not distinguishing between gray (isotropic) and white (anisotropic) matter. The trace is distributed around the value of 2.4×10^{-3} mm^2/s, and the determinant

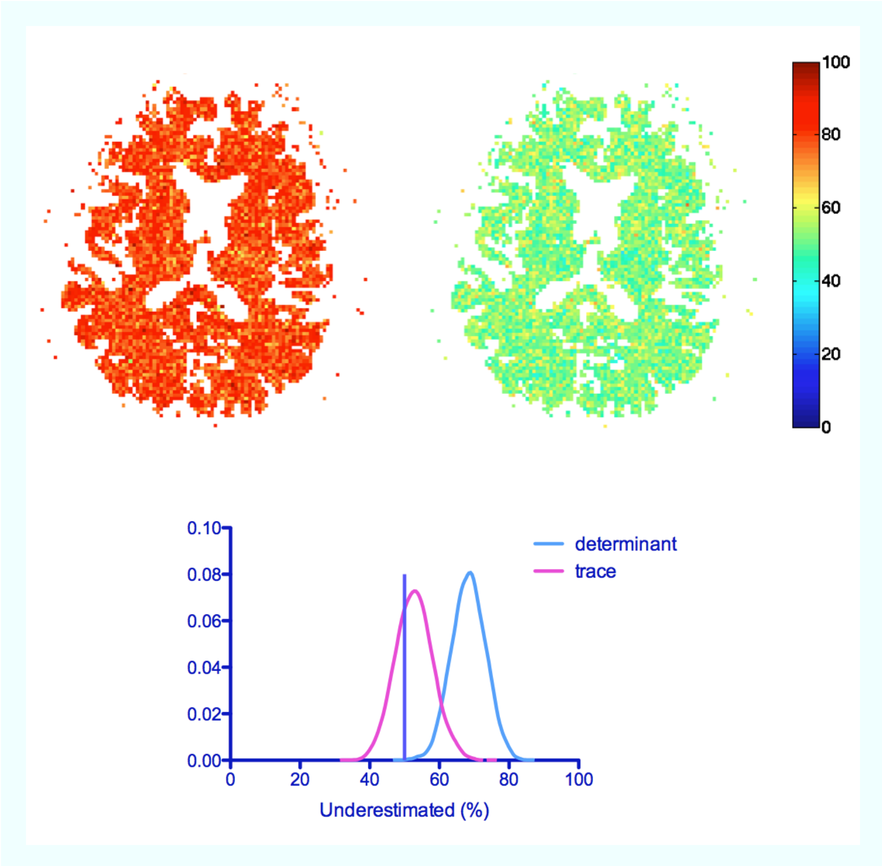

Fig. 5 Extent of Shrinking. The maps show the percent of noisy replicates that had a lower determinant (left) and trace value (right) than the ground truth values. The value for the determinant map for most voxels is larger than 50% indicating that the noisy replicates there tend to shrink. The values for the trace map are closer to 50%, indicating almost equal probability to find an increased or decreased trace in the noisy replicates.

around the value 0.4×10^{-3} $(mm^2/s)^3$. The original values were compared with the trace and determinant values of all noisy replicates to yield the maps in Figure 5. In these maps the value of each voxel is weighted by the percentage of noisy replicates that had a value lower than the ground truth value for the same voxel. For example, for the determinant map (left), a voxel with a value higher than 50% (yellow to red) means that most of the noisy replicates in this voxel were shrinking since they had a lower determinant value than the ground truth. For the determinant map it is clear that almost all tensors in this slice tend to shrink. The trace map shows a slight tendency to a decrease in value, yet most tensors are around the 50% mark (green), meaning that the noisy replicates had equal probability to either increase or decrease trace.

The conclusion of these findings is that following the type of noise expected in diffusion MRI acquisition, tensors are expected on average to decrease their determinants while preserving their trace. This conclusion provides an additional support for favoring a trace preserving metric over a determinant-preserving metric.

4.3 Why do Tensors Swell?

Having seen that tensors are expected to shrink when noise is introduced we can now better understand why tensors swell during averaging and interpolations. Averaging is usually performed when a single tensor is not considered informative enough to provide an accurate reading. When noise is considerable, averaging is used to increase SNR, which in turn allows an unbiased estimate of the true value. In the estimation of diffusion tensors from noisy replicates, our simulations predict that the noisy replicates will have a lower determinant than the original determinant. Therefore, in order to obtain an unbiased estimate of the tensor and its trace, the estimation method should cause the determinant to increase, i.e., an unbiased estimate for a diffusion tensor is one that causes noisy replicates of the tensor to swell. Indeed we have previously shown that the main ill-effect accompanied with using the determinant-preserving geometric metric is a consistent bias to the estimation [23]. We can now understand that the bias is there since the method preserves the determinant, while the noise properties dictate preservation of the trace. One way of interpreting our results would be saying that both the Euclidean and the Geometric metrics introduce biases, but as it happens the Euclidean metric bias happens to compensate for noise relating bias. This is a valid claim that could be further explored.

Another way of interpreting the results would consider the acquisition step it self as an averaging mechanism. When measuring diffusion images, the signal is affected by many different tissue compartments and by various noise sources, each of these components may be modeled as a tensor, and the final signal averages all of the components into a single tensor. We hypothesize that the Euclidean metric might be the native metric of this kind of averaging mechanism, yet testing this hypothesis requires further experimental results. At this time we can only point at the way diffusion itself is defined via Fick's Law

$$J = -D\nabla\phi(x),$$

where J is the diffusion flux and ϕ is the concentration, or the displacement probability function. In Fick's law the gradient is measuring variations of the probability function using a Euclidean norm. Analogous equations are believed to govern many other physical phenomena, such as electrical flux (Ohm's law) and heat flux (the heat equation). Taking the geometric claim further step ahead would claim that these laws themselves should be defined using geometric metrics to become

$$J = -D\nabla\phi(log(x)).$$

This claim was made, for example, in [28]. It is our belief that this is a highly controversial claim that is required to justify the relevancy of the geometric metric, yet at the same time it falsifies almost any theory that is based on physical observations.

5 Beyond Riemannian Metrics

We have shown that Johnson noise that is expected in the acquisition of diffusion images is more likely to cause a decrease in the determinant of the noisy tensors. Johnson noise is expected to be encountered in all voxels and under all scanning conditions [16, 1], but there are many other types of variability sources besides thermal noise. To name a few, there are eddy currents, which depend on the gradient magnitude and direction and specific acquisition sequence used [26]; reconstruction artifacts originating from the use of multiple surface coils [16]; and motion artifacts, due to rigid head motion or non-linear cardiac pulsation effects [27, 25]. In addition, there is biological variability, for example, tensors that are part of the same white matter tract are expected to have similar diffusivity values and similar shapes, but to express higher variability in the orientation of the tensors as they follow the trajectory of the tract. In this example, preserving the determinant of the tensors when interpolating or averaging among them is required. But not only the determinant must be preserved: Since the variability source dictates rotation of the reference frame, all rotation-invariant tensor properties must be preserved [8]. For this all eigenvalues, and their related quantities (such as trace, determinant, and FA) must be maintained. As we show above in Figure 1, preserving trace alone or determinant alone does not guarantee preservation of FA.

The Riemannian metrics such as those we have analyzed here have the property of being global, meaning the distance between two given tensors remains the same no matter where these tensors are located within the image or within the tissue. With the absence of any prior geometric information (such as the expected anisotropy or orientation), the swelling effect is predicted by the MR measurement, and there is no physical reason to preserve the determinant. But when additional geometric information is available, local metrics, especially designed for the type of variability expected within that tissue are needed. These tissue-specific metrics have yet to be developed, but we note that the family of metrics proposed in [15, 19] might be of interest, as they are able to separate the information encoded in the tensor to orthogonal features embedded in a Euclidean space, allowing weighting of the influence of each feature on the final distance measure. Some more relevant information can be found in Chapter [Cross Reference to Garcia's chapter].

The diffusion tensor is proportional to the variance of normally distributed particle displacements (4). But as experiments dictate, when a voxel contains multiple components, each normally distributed on its own, the composition may deviate from the gaussian distribution and the tensor model. This brings up an interesting question of whether the result of averaging diffusion tensors, which actually means mixing two separate diffusivity components, should yield a tensor, or a different statistical representation. Some interesting ideas as to what this type of representation can be are found in Chapter [Cross Reference to Westin's chapter].

6 Summary

We have tested how the determinant and trace of a diffusion tensor change when Johnson noise is introduced. We have found that the determinant is likely to be reduced in all types of tensors, and that the trace is equally likely to either reduce or increase following the introduction of Johnson noise. This lead us to the conclusion that in order to provide an unbiased estimation, an average or interpolation tensor operator should result in an increased determinant compared with the determinant of the initial noisy replicates. This provides a physical explanation for the swelling effect, and implies that in the most general case, where no additional geometric information is provided, there is no physical justification for a determinant-preserving metric, but there is a justification for a trace preserving metric. This explanation provides a practical reason to prefer the Euclidean metric over the geometric metrics, in addition to previous theoretic considerations [23] that are also in favor of the Euclidean choice.

Acknowledgements We thank Liz Salak for reviewing the manuscript. OP wishes to acknowledge with thanks that part of the research on which this publication is based was supported by a Fulbright Post-doctoral Scholar Fellowship, awarded by the Fulbright commission for Israel, the United States-Israel Educational Foundation.

References

1. Andersson, J.L.: Maximum a posteriori estimation of diffusion tensor parameters using a Rician noise model: Why, how and but. NeuroImage **42**(4), 1340–1356 (2008)
2. Arsigny, V., Fillard, P., Pennec, X., Ayache, N.: Log-Euclidean metrics for fast and simple calculus on diffusion tensors. Magn Reson Med **56**(2), 411–421 (2006). DOI 10.1002/mrm.20965
3. Assaf, Y., Pasternak, O.: Diffusion tensor imaging (DTI)-based white matter mapping in brain research: a review. J Mol Neurosci **34**(1), 51–61 (2008)
4. Basser, P.J., Jones, D.K.: Diffusion-tensor MRI: theory, experimental design and data analysis – a technical review. NMR in Biomedicine **15**, 456–467 (2002)
5. Basser, P.J., Mattiello, J., LeBihan, D.: MR diffusion tensor spectroscopy and imaging. Biophys J **66**, 259–267 (1994)
6. Basser, P.J., Pajevic, S.: A normal distribution for tensor-valued random variables to analyze diffusion tensor MRI data. IEEE TMI **22**, 785–794 (2003)
7. Basser, P.J., Pierpaoli, C.: A simplified method to measure the diffusion tensor from seven MR images. Magn. Reson. Med. **39**, 928–934 (1998)
8. Batchelor, P.G., Moakher, M., Atkinson, D., Calamante, F., Connelly, A.: A rigorous framework for diffusion tensor calculus. Magn Reson Med. **53**, 221–225 (2005)

9. Crank, J.: The mathematics of diffusion. Oxford Univ. Press (1975)
10. Einstein, A.: Investigations on the Theory of the Brownian Movement. Dover, New York (1926)
11. Eisenhart, L.: Differential Geometry. Princeton Univ. Press (1940)
12. Fillard, P., Arsigny, V., Pennec, X., Ayache, N.: Clinical DT-MRI estimation, smoothing and fiber tracking with log-Euclidean metrics. IEEE Transactions on Medical Imaging **26**(11), 1472–1482 (2007). DOI 10.1109/TMI.2007.899173. PMID: 18041263
13. Fletcher, P.T., Joshi, S.: Riemannian geometry for the statistical analysis of diffusion tensor data. Signal Process. **87**(2), 250–262 (2007). DOI http://dx.doi.org/10.1016/j.sigpro.2005.12.018
14. Gur, Y., Pasternak, O., Sochen, N.: Fast gl(n)-invariant framework for tensors regularization. Int J Comput Vis **85**(3) (2009)
15. Kindlmann, G., Ennis, D.B., Whitaker, R.T., Westin, C.F.: Diffusion tensor analysis with invariant gradients and rotation tangents. IEEE Transactions on Medical Imaging **26**(11), 1483–1499 (2007). DOI 10.1109/TMI.2007.907277
16. Koay, C.G., Basser, P.J.: Analytically exact correction scheme for signal extraction from noisy magnitude MR signals. Journal of Magnetic Resonance **179**(2), 317–322 (2006)
17. Kubicki, M., McCarley, R., Westin, C.F., Park, H.J., Maier, S., Kikinis, R., Jolesz, F.A., Shenton, M.E.: A review of diffusion tensor imaging studies in schizophrenia. J Psychiatr Res **41**, 15–30 (2007)
18. Lenglet, C., Rousson, M., Deriche, R., Faugeras, O.: Statistics on the manifold of multivariate normal distributions: Theory and application to diffusion tensor MRI processing. J. Math. Imaging Vision **25**(3), 423–444 (2006)
19. de Luis-Garcia, R., Alberola-Lopez, C., Kindlmann, G., Westin, C.F.: Automatic segmentation of white matter structures from dti using tensor invariants and tensor orientation. In: Proc 17th Annual Meeting ISMRM (2009)
20. Maaß, H.: Siegel's Modular Forms and Dirichlet Series. Springer, Berlin (1971)
21. Moakher, M.: On the averaging of symmetric positive-definite tensors. J. Elasticity **82**, 273–296 (2006)
22. Pajevic, S., Basser, P.J.: Parametric and non-parametric statistical analysis of DT-MRI data. J Magn Reson **161**(1), 1–14 (2003)
23. Pasternak, O., Sochen, N., Basser, P.: The effect of metric selection on the analysis of diffusion tensor MRI data. Neuroimage **49**, 2190–2204 (2010)
24. Pennec, X.: Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. J. Math. Imaging Vision **25**(1), 127–154 (2006). DOI 10.1007/s10851-006-6228-4
25. Pierpaoli, C., Marenco, S., Rohde, G., Jones, D., Barnett, A.: Analyzing the contribution of cardiac pulsation to the variability of quantities derived from the diffusion tensor. In: Proc 11th Annual Meeting ISMRM, p. 70 (2003)
26. Rohde, G., Barnett, A., Basser, P., Marenco, S., Pierpaoli, C.: Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. Magnetic Resonance in Medicine **51**(1), 103–114 (2004)
27. Skare, S., Anderson, J.: On the effects of gating in diffusion imaging of the brain using single shot EPI. Magnetic Resonance Imaging **19**, 1125–1128 (2001)
28. Tarantola, A.: Elements for Physics: Quantities, Qualities, and Intrinsic Theories. Springer, Berlin (2006)