On the local geometry of the pet reconstruction problem

Ramy Merabet, Florent Sureau, Alain Trouvé

ON THE LOCAL GEOMETRY OF THE PET RECONSTRUCTION PROBLEM

Ramy Merabet[⋆†] *Florent Sureau*[†] *Alain Trouvé*[⋆]

[⋆] Centre Borelli, ENS Paris-Saclay
[†] BioMaps, Université Paris-Saclay, CEA, CNRS, INSERM

ABSTRACT

Intrinsic uncertainties related to the ill-posedness of the PET reconstruction problem are investigated in this work. These uncertainties could lead to artifacts in the reconstructed images in particular when Deep Learning approaches are employed. We propose a framework enabling to define a distinguishability measure between PET images and propose a local analysis of such a measure in a neighborhood of a reference image. This approach is numerically tested on a synthetic experiment using a 3D brain PET phantom derived from a measured [18F]-FDG exam with 100 anatomo-functional regions extracted. Our analysis allows us to highlight the key factors impacting the detectability of variations and to exhibit concrete examples of clinically meaningful directions along which variations may not be detectable. In addition, we quantitatively analyze the role played by the injected radiotracer dose and show that low-dose scenarios are particularly prone to the presence of reconstruction artifacts.

***Index Terms*—** Aleatoric uncertainties, Bayesian inference, Fisher information, PET reconstruction

1. INTRODUCTION

Positron emission tomography (PET) is a functional and quantitative medical imaging modality that allows to follow, in three dimensions, the spatial distribution of a radiotracer previously injected into the patient. PET imaging can then provide information on a targeted physiological process (tumor uptake in oncology for example). During a PET scan, the radiotracer injected into the patient emits positrons which are annihilated by electrons resulting in the emission of two back-to-back photons that can be detected in coincidence by a PET camera located around the patient. The count data thus measured can be modeled by Poisson random variables whose parameters are given by the projections of the activity concentration [1]. The reconstruction of the image of this activity concentration from these count data constitutes an ill-posed tomographic inverse problem.

Over the last decades, many algorithms have been proposed to numerically solve this tomographic ill-posed inverse problem. Recent approaches using Deep Learning [1] open promising avenues and allow to obtain high-quality PET images, sometimes even in low-dose scenarios, thus reducing the patient's exposure to radiation and/or the cost of the PET scan [2]. However, the increase in performance over more traditional inverse problem approaches seems to come at the cost of poorly understood instability and bias phenomena that are potentially the source of false positives or false negatives, impeding clinical use, as illustrated in [3] [4].

These phenomena are related to the uncertainties on the inverse problem which can be categorized into epistemic and aleatoric uncertainties [5]. While the former has been widely studied with techniques such as the so-called Monte-Carlo dropout [6] and its application in medical imaging [7], we focus in this work on aleatoric uncertainties, intrinsic to the PET reconstruction problem and source of instabilities and biases as shown in [3] in the case of MRI.

Our contribution are as follows. First, from a theoretical point of view, we propose a framework based on Bayesian inference, allowing to define a distinguishability measure between PET images, namely the Bayesian risk. Secondly, from a numerical point of view, we analyze locally this Bayesian risk in a neighborhood of a reference image. This local analysis comes to the study of the Fisher information metric. Our analysis reveals a non-trivial geometry on the space of PET images and allows us to understand the key factors influencing the detectability of variations and the potential presence of artifacts in the reconstruction. In addition, we investigate different norms for measuring the perturbations produced, and propose a way of restricting the analysis to the clinically significant ones. Along the directions of these particular perturbations, we quantitatively analyze the role played by the injected radiotracer dose in the detectability of variations.

2. METHOD

In the sequel, we denote $\theta \in \mathbb{R}^J$ a PET image, $X \in \mathbb{N}^L$ the acquired data, $A \in \mathbb{R}_+^{L\times J}$ the forward operator, with lines $(A_l)_{l=1,\dots,L}$, and $b \in \mathbb{R}_{+*}^L$ collecting the expectation of scatter and random events. We denote $\lambda = A\theta + b \in \mathbb{R}_+^L$ the (deterministic) projection of an image θ. PET reconstruction consists in estimating θ from $X \sim \bigotimes_{l=1}^{L} \mathcal{P}(\lambda_l)$, $\mathcal{P}$ denoting the Poisson distribution.

2.1. The Bayesian risk

Let θ_0 and θ_1 be two PET images, and X acquired data. Consider the problem of determining which of these two PET images generated the acquired data. For such problems, the Bayesian paradigm enables to define optimal decisions, an a priori distribution and a loss function being specified. The risk of such decisions, called the Bayesian risk, is a measure of the intrinsic difficulty of the problem, i.e. a measure of the irreducible error incurred.

Let us consider the non-discriminative prior :

$$\pi = \frac{1}{2}(\delta_{\theta_0} + \delta_{\theta_1}).$$

The Poisson likelihood being :

$$f(x|\theta) = \prod_{l=1}^{L} \frac{(\langle A_l, \theta\rangle + b_l)^{x_l}}{x_l!} \exp(-\langle A_l, \theta\rangle + b_l),$$

the posterior is given by :

$$\pi(\theta|x) = \frac{f(x|\theta_0)}{f(x|\theta_0) + f(x|\theta_1)}\delta_{\theta_0} + \frac{f(x|\theta_1)}{f(x|\theta_0) + f(x|\theta_1)}\delta_{\theta_1}.$$

For the binary loss defined by $l(\theta_0, \theta_1) = \mathbf{1}_{\theta_0 \neq \theta_1}$, the Bayesian estimator is the MAP, giving the Bayesian risk $\mathcal{R}_0 = \frac{1}{2}\int f(x|\theta_0) \wedge f(x|\theta_1)dx$, which is not smooth with respect to θ_1. In the sequel, we work with the smooth approximation $\mathcal{R} = \mathbb{E}_{\theta_0}\left[\frac{f(x|\theta_1)}{f(x|\theta_0)+f(x|\theta_1)}\right]$, which verifies $\mathcal{R}_0 \leq \mathcal{R} \leq 2\mathcal{R}_0$. This quantity corresponds to the risk incurred by a Bayesian oracle trying to solve the proposed discrimination problem, and as such, encodes a notion of similarity between the PET images.

2.2. Local analysis

Consider now θ_0 a fixed reference image, and R as a function of θ_1. In a neighborhood of θ_0, we can compute the Taylor expansion of $\mathcal{R}$ and obtain (with $\lambda_0 = A\theta_0 + b$) :

$$\mathcal{R}(\theta_0 + \delta\theta) \approx \frac{1}{2}\left(1 - \frac{1}{4}\delta\theta^T I(\theta_0)\delta\theta\right),$$

with $I(\theta_0) = A^T D(\frac{1}{\lambda_0})A$ the Fisher information in θ_0 ($D(\frac{1}{\lambda_0})$ denoting the diagonal matrix with coefficients $\frac{1}{\lambda_{0,l}}$).

Thus, locally, the Bayesian risk geometry reduces to the Fisher-Rao geometry of probabilistic models.

Moving from the space of PET images to the space of projections, we obtain (with $\delta\lambda = A\delta\theta$) :

$$\mathcal{R}(\lambda_0 + \delta\lambda) \approx \frac{1}{2}\left(1 - \frac{1}{4}\sum_{l=1}^{L}\frac{\delta\lambda^2}{\lambda_{0,l}}\right).$$

This formula provides an understanding of the key factors influencing the detectability of perturbations in the reference image. To make $\mathcal{R}(\theta_0 + \delta\theta)$ as large as possible, we need to make $\sum_{l=1}^{L}\frac{\delta\lambda^2}{\lambda_{0,l}}$ as small as possible which reveals :

- A volume/geometry effect : the more the perturbation $\delta\theta$ is localized in a small/flat region, and the fewer lines of response it impacts (or the fewer it impacts lines of response in the case of a flat region), the smaller $\sum_{l=1}^{L}\frac{\delta\lambda^2}{\lambda_{0,l}}$ is.
- A relative activities effect : the more the perturbation $\delta\theta$ is localized in a region whose contribution to the lines of response passing through it is negligible compared to the contributions of the other regions, the lower the detectability of a variation of constrast caused by this perturbation.
- A tomographic effect : the more the perturbation $\delta\theta$ is localized in a central region, the more the lines of response passing through it will be impacted by the other regions, and the more this pertubation will be prone to the relative activities effect mentioned above.

Finding the least distinguishable local direction (with highest Bayesian risk) amounts to solve the following problem :

$$\min_{\delta\theta} \quad \delta\theta^T I(\theta_0)\delta\theta \qquad \text{s.t.} \qquad \|\delta\theta\|^2 = 1 \tag{1}$$

The geometry of the optimization landscape of problem (1) crucially depends on the choice of the metric on the PET image space.

In the sequel, we numerically investigate differents norms for measuring pertubrations $\delta\theta$ with the aim of observing in practice the three above-mentioned effects in clinically significant situations.

3. EXPERIMENTS

3.1. ROIs based pertubations

In order to construct interpretable metrics and to reduce to dimensionality of the problem, we chose to work with PET images segmented by anatomo-functional regions of interest (ROIs). We then consider a parametric model $u = (u_r) \in \mathcal{U} \subset \mathbb{R}^R \mapsto \theta = Mu$, where $\mathcal{U}$ is the parameter space, and R is the number of ROIs.

While information on the relative activities and locations of perturbations is preserved by this parameterization of the PET image space, the volume differences between the regions impacted by the perturbations considered are very significant, and the perturbations produced are therefore a priori incomparable.

In the sequel, we investigate different ways of taking into account ROIs volumes to construct comparable perturbations.

3.2. Results

We consider a synthetic experiment derived from a measured [^{18}F]-FDG brain exam of a healthy subject. The associated T1 MRI image was segmented into 100 regions using FreeSurfer[1]. The PET signal was measured in a frame between 30 minutes and 60 minutes after injection in each region using PETSurfer [8] to generate the anatomo-functional

[1] https://surfer.nmr.mgh.harvard.edu

phantom. 3-dimensional PET simulations for a Biograph 6 TruePoint TrueV PET system were then generated using an analytical simulator, including all corrections.

We denote $(e_1, \cdots, e_R)$ the canonical basis of the parameter space, and $D_r = \delta\theta_r^T I(\theta_0)\delta\theta_r$ for $\delta\theta_r = Mw_re_r$ *the distinguishability score* of ROI r, where $\|w_re_r\| = 1$ for different choice of the norm $\|.\|$.

Note that in the following analysis, we will consider absolute variations in activity and not relative variations in contrast (i.e. the perturbations produced will be independent of the activity in the reference image), in order to analyze the volumes effect independently of the relative activities effects.

Euclidean norm on the parameter space. For $\delta u = \sum_{r=1}^{R} \delta u_r e_r$, $\|\delta u\|_u^2 = \sum_r \delta u_r^2$. This amounts to comparing perturbations for which the same variation in mean activity has been achieved.

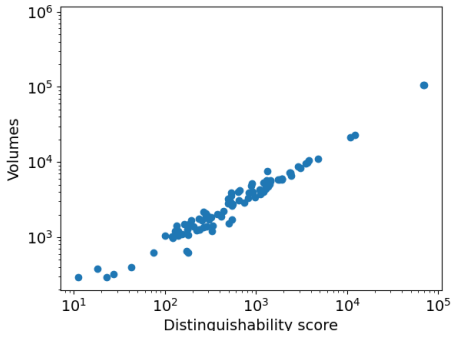

Fig. 1: Correlation between distinguishability score and volumes for the norm $\|.\|_u$.

Figure 1 shows the correlation between the distinguishability score and ROI volumes in the case of the norm $\|.\|_u$. Naturally, distinguishability is dominated by volume effects (which is well explained by the volume differences between ROIs), thus preventing the observation of the relative activities and location effects mentioned above.

Euclidean norm on the PET image space. For $\delta u = \sum_{r=1}^{R} \delta u_r e_r$, $\|\delta u\|_2^2 = \sum_r V_r \delta u_r^2$, where V_r denotes the volume of region r. This is the classical metric used in image processing.

Figure 2 shows that while the volume effect has been partially mitigated, detectability is still largely dominated by volume differences between ROIs.

Normalized for volumes norm on the parameter space. For $\delta u = \sum_{r=1}^{R} \delta u_r e_r$, $\|\delta u\|_V^2 = \sum_r V_r^2 \delta u_r^2$. This amounts to comparing perturbations for which the same variation in total dose was achieved.

Figure 3 shows that the normalization induced by the use of the norm $\|.\|_V$ has reversed the influence of volumes on detectability: the variations in mean activity permitted in the largest regions by the norm $\|.\|_V$ are so insignificant that these become the least distinguishable regions, and vice versa for

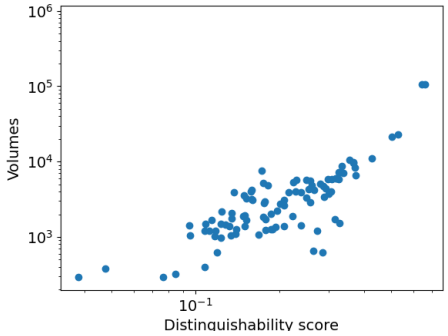

Fig. 2: Correlation between distinguishability score and volumes for the norm $\|.\|_2$.

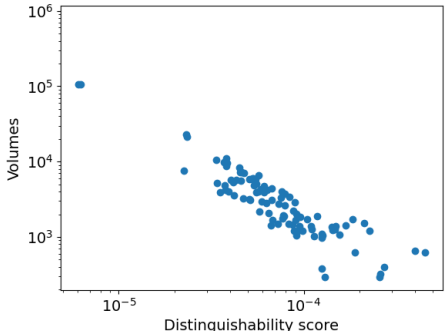

Fig. 3: Correlation between distinguishability score and volumes for the norm $\|.\|_V$.

the smallest regions.

The conclusion of this empirical study is that it is highly non-trivial to take appropriate account of ROIs volumes : this would involve tuning normalization coefficients (V_r^β here) that are likely to depend on ROIs geometries and whose tuned value can no longer be interpreted in terms of injected radiotracer dose.

In the sequel, we propose to restrict ourselves to a particular perturbations model, allowing us to intrinsically normalize volumes and geometries effects, and thus consider comparable perturbations.

3.3. Functional subregions based perturbations

The proposed perturbation model consists in restricting the perturbations considered to impact a functional sub-region of each ROI. These functional subregions are centered in each ROI and have a volume of the order of 2 cm³, which corresponds to real cases of cerebral lesions. Regions with smaller volumes are not considered. This restriction is used to normalize the perturbations in terms of volumes and geometries leading to a new perturbation model: $\delta\theta = M'\delta u$.

Note that, since the effects of volumes and geometries

have been mitigated, we can now work with a relative metric to measure variations in contrast (which is clinically more significant).

3.4. Results

As above, $D_r = \delta\theta_r^T I(\theta_0)\delta\theta_r$ for $\delta\theta_r = M' w_r e_r$ the distinguishability score of ROI r, where w_r is the variation of mean activity in the functional subregion of ROI r and $\|w_r e_r\| = 1$ for the relative norm on the parameter space.

Relative norm on the parameter space. For $\delta u = \sum_{r=1}^{R} \delta u_r e_r$, $\|\delta u\|_u^2 = \sum_r (\frac{\delta u_r}{u_{r,0}})^2$. This amounts to comparing perturbations for which the same contrast w.r.t the reference has been achieved.

In figure 4 below, we show concrete examples of such perturbations (the functional subregions shown correspond to three of the five worst cases, the other two being cortical regions and therefore difficult to represent in two dimensions).

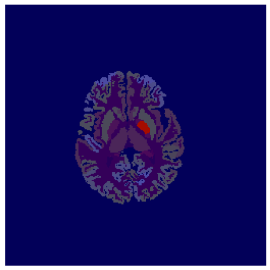

Axial plane

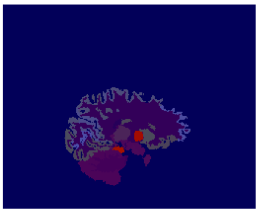

Sagittal plane

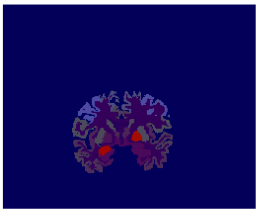

Coronal plane

Fig. 4: Concrete examples of three of the five less distinguishable perturbations (in red).

Along these particular directions, we can perform a quantitative analysis of the role played by the injected radiotracer dose in the detectability of the perturbations produced. Indeed, by fixing the Bayesian risk to a constant value, for example 30 %, we can study the maximum contrast achievable along these particular directions w.r.t the injected radiotracer dose.

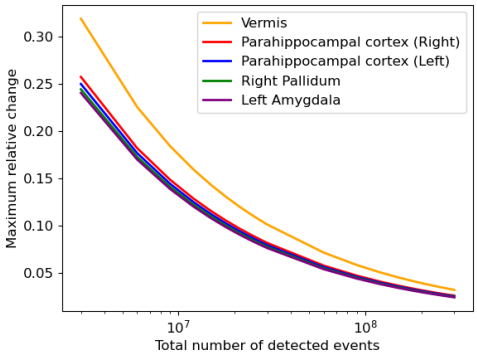

Fig. 5: Maximum relative change in particular functional subregions w.r.t the injected radiotracer dose at $\mathcal{R} = 30\%$.

Thus, according to our study, it is possible, in a low-dose scenario (with a total of three million events detected), to perturb the reference image by creating a contrast of more than thirty percent in the vermis, for example, while keeping a probability of not being able to distinguish this perturbation in the acquired data of more than thirty percent.

Note that the worst case in problem (1) corresponds to the eigenvectors of $I(\theta_0)$ associated with the lowest eigenvalues. The orthogonal nature of these eigenvectors implies that they are made up of mixtures of (functional subregions of) ROIs, which may not give rise to clinically relevant interpretation and is therefore not considered here. In particular, we restricted the analysis to directions given by the canonical basis of the parameter space (which correspond to localized perturbation of the reference image). It should be noted, however, that the point of view adopted here is optimistic in the sense that perturbations even less distinguishable than those proposed can be constructed by following the directions given by these eigenvectors.

4. CONCLUSION

This paper proposes a framework, based on a Bayesian inference problem, for investigating aleatoric uncertainties in the PET reconstruction problem. This framework allows us to define a distinguishability measure between PET images, which we have analyzed locally in the neighborhood of a reference image, which naturally amounts to the analysis of the Fisher information metric. This framework is quite general and applies beyond the particular PET reconstruction problem. This analysis has enabled us to identify the key factors controlling the detectability of a perturbation of the reference image, and to propose a perturbation model to illustrate them while avoiding undesirable effects. Along these particular perturbations, we analyzed the role of the injected radiotracer dose, and highlighted that low-dose scenarios are particularly prone to the presence of reconstruction artifacts.

5. COMPLIANCE WITH ETHICAL STANDARDS

This is a numerical simulation study for which no ethical approval was required.

6. REFERENCES

[1] A. J. Reader and al, "Deep learning for pet image reconstruction," in *IEEE Transactions on Radiation and Plasma Medical Sciences, 5(1), 1-25*, 2020.

[2] Y. Zhu S. Kaplan and J. Dig, "Full-dose pet image estimation from low-dose pet image using deep learning : a pilot study. journal of digital imaging," in *Journal of Digital Imaging 32(5), 32(5), 773-778,*, 2019.

[3] V. Antun and al, "On instabilities of deep learning in image reconstruction and the potential costs of ai," in *PNAS*, 2020.

[4] A. C. Hansen N. M. Gottschling, V. Antun and B. Adcock, "The troublesome kernel - on hallucinations, no free lunches and the accuracy-stability trade-off in inverse problems," in *Preprint*, 2023.

[5] A. D. Kiureghian and O. Ditlevsen, "Aleatory or epistemic? does it matter?," in *Structural Safety, 31(2), 105-112*, 2009.

[6] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *International Conference on Machine Learning, 1050-1059*, 2016.

[7] P. Seebock and al, "Exploiting epistemic uncertainty of anatomy segmentation for anomaly detection in retinal oct," in *IEEE Transactions on Medical Imaging, 39(1), 87-98,*, 2020.

[8] D.N. Greve and al, "Different partial volume correction methods lead to different conclusions: An (18)f-fdg-pet study of aging," in *Neuroimage, 132, 334-343*, 2016.