

Uncertainty driven probabilistic voxel selection for image registration

Boris N. Oreshkin, Tal Arbel *Member, IEEE*,

Abstract—This paper presents a novel probabilistic voxel selection strategy for medical image registration in time-sensitive contexts, where the goal is aggressive voxel sampling (e.g. using less than 1% of the total number) while maintaining registration accuracy and low failure rate. We develop a Bayesian framework whereby, first, a voxel sampling probability field (VSPF) is built based on the uncertainty on the transformation parameters. We then describe a practical, multi-scale registration algorithm, where, at each optimization iteration, different voxel subsets are sampled based on the VSPF. The approach maximizes accuracy without committing to a particular fixed subset of voxels. The probabilistic sampling scheme developed is shown to manage the tradeoff between the robustness of traditional random voxel selection (by permitting more exploration) and the accuracy of fixed voxel selection (by permitting a greater proportion of informative voxels).

I. INTRODUCTION

Image registration is one of the critical problems in the field of medical imaging, spanning a wide range of applications including, but not limited to, diagnostic imaging [1], image guided interventions [1], [2], [3], image guided therapy [4] and patient to atlas mapping [5]. Typically, the evaluation of the similarity measure and/or its derivatives are required to perform the optimization over the transformation parameters. Many registration strategies have been developed under the implicit assumption that all of the voxels (or, at least, as many as necessary to achieve desirable result) can be used for calculating the similarity measure — the ultimate goal being the highest accuracy at any computational cost. However, in time sensitive applications, such as in image guided interventions [1], [2], [3] and image guided therapy [4], image registration must be performed both accurately and quickly in order for the techniques to be adopted by clinicians. An example of this type of domain is image guided neurosurgery, where ultrasound images acquired during the operation must be quickly matched to pre-operative MRI or CT in order to correct for brain shift [1], [2], [6], [7], [8]. Performing computations based on all the available image voxels can be prohibitively costly mainly due to the large number of voxel intensity values involved in the calculations. Time-sensitive applications generally benefit from techniques that speed up direct image registration by utilizing only a subset of available voxels during registration. In these contexts, a decrease in

accuracy of several percentages can be tolerated in order to assure the preservation of robustness and a significant decrease in registration time. However, *significant* speedups attained via a very aggressive reduction in the number of selected voxels (e.g. less than 1% of the total number of voxels) often result in a significant deterioration in robustness (i.e. an increase in failure rate) and relatively rapid increase in the resulting registration error. In this paper, our main goal is to demonstrate the promise of a voxel sampling technique based on the Bayesian model for registration parameter uncertainty, by showing that a well designed sampling strategy is capable of maintaining the functionality of the registration procedure even if the average number of voxels used in similarity metric calculations is very low. In other words, we explore the limits to which image registration can be pushed in terms of voxel selection. Specifically, the goal is to develop a voxel selection (sampling) scheme that would work well even when a very small percentage of voxels (on the order of 0.1%) is used for registration, thus attaining fast registration in contexts where time is the most valuable resource. However, the registration framework based on this voxel selection scheme is also required to guarantee reliable and accurate results, a goal that is key to its adoption in the context of clinical practice (time sensitive domains, e.g. image guided therapy [4] or image guided interventions [1], [2], [3]).

Randomized *voxel selection* is one commonly used technique for reducing the computational cost of registration, where a randomly selected subset of all voxels are used to drive the optimization [9], [10]. This technique gained popularity due to its simplicity. Randomized voxel selection, which avoids any bias towards particular image patterns, does lead to a reduction in registration accuracy. Thus the utility of optimizing voxel selection has been well recognized. Early works that attempted to optimize voxel selection presented heuristics for the relatively simple context of template matching [11], [12] and demonstrated the advantage of optimized selection. More recently, a threshold-based optimized voxel selection method was proposed by Reeves and Hezar [13]. They proposed a threshold based k-space voxel selection scheme based on the mean squared reconstruction error criterion in the context of the magnetic resonance image reconstruction. However, voxels selected by the threshold based approach tend to form spatially localized clusters leading to reduced robustness. Dallaert and Collins [14] attempted to overcome this effect and proposed a heuristic voxel selection strategy for the context of template based tracking. It involves computing the Jacobian of the similarity metric for every voxel and randomly selecting M out of the top 20 percent voxels with the highest Jacobian [14]. Brooks and Arbel [15] extended

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Boris N. Oreshkin and Tal Arbel are with the Centre for Intelligent Machines, Department of Electrical and Computer Engineering, McGill University, QC, H3A 0E9, Canada (e-mail: boris.oreshkin@mail.mcgill.ca, arbel@cim.mcgill.ca).

Manuscript received October 19, 2012; revised May 17, 2013.

the approach of Dellaert and Collins [14] by proposing an information theoretic selection criterion and addressing the issue of the Jacobian scale. In practice, many methods simply favour voxels with high intensity gradient magnitudes since they better reveal internal structure of the image and thus are deemed more useful for obtaining a good alignment. Furthermore, Sabuncu and Ramadge proposed the gradient magnitude subsampling (GMS) approach, where the moving image is probabilistically subsampled using non-uniform grid generated based on the probabilities proportional to the gradient magnitude [16]. This approach permits diversification and spread of the subsampled voxels while still focusing on image details. The same paper discusses the fact that the randomized subsampling is superior to the subsampling on a regular grid (decimation). An extension of this work employing the mixture of the previous approach with the uniformly random subsampling was proposed by Oreshkin and Arbel in [17]. Here it is proposed to sample from the convex combination of the probability fields generated via the GMS and the uniformly random sampling approaches and the mixing parameter is identified from learning procedure using a small training dataset. Finally, curvlet based sampling, recently proposed by Freiman et al. [18] was tested on Vanderbilt RIRE (Retrospective Image Registration Evaluation) dataset [19] and resulted in approximately the same level of accuracy as the GMS approach [16]. All these approaches work well in some contexts, but most fail in the low sampling rate scenario that is the focus of the current paper. Moreover, existing approaches often lack a direct link to the task at hand, namely, the estimation of the registration parameters. Overall, developing a formal, general stochastic voxel selection framework for low sampling rate image registration is still an open research topic.

In this paper, we propose a new probabilistic voxel selection strategy based on a multi-resolution voxel sampling probability field (VSPF), formed by the set of probabilities associated with image voxels. At every scale, we optimize the VSPF to minimize the transformation parameter uncertainty calculated from the Bayesian generative model for the approximate moving image. In the proposed framework, the most informative voxels are selected more frequently, while the exploration of the image continues via sampling different voxel subsets at each optimization iteration. The fact that each scale has its own VSPF allows to further exploit this feature by emphasizing different voxel subsets at different resolution scales. We derive the analytical solution of the VSPF optimization problem and illustrate the adaptive nature and generality of the approach. We show that deterministic selection rules similar to e.g. one presented in Reeves and Hezar [13] and the random voxel selection based on uniform probabilities of Wells et al. [10] can be cast as particular cases of our framework. Specifically, the probabilistic sampling scheme developed manages the tradeoff between the robustness of random voxel selection (by permitting more exploration) and the accuracy of fixed voxel selection (by permitting a greater proportion of informative voxels). Finally, we propose a practical algorithm based on our framework and demonstrate significant registration performance gains through multi-modal (MRI-CT) image registration experiments with the popular, publicly available RIRE

Vanderbilt dataset [19]. One of the limitations of this dataset is that validation is restricted to rigid multi-modal registration only. The rest of the paper is organized as follows. In Section II we describe the relationship between direct image registration and voxel selection and introduce the notion of the VSPF. In Section III-A we define voxel utility, which is used to optimize the VSPF and to build a practical multi-scale algorithm in Sections III-B and IV. Results of numerical experiments are then presented in Section V. Section VI discusses our findings, limitations inherent to the proposed methodology and to the validation method. Section VII concludes the paper.

II. PROBLEM STATEMENT

Direct parametric image registration is formulated for a reference image, $I(\mathbf{x}) = [I_1(\mathbf{x}_1), \dots, I_N(\mathbf{x}_N)]$, and a transformed moving image, $V(T_\theta(\mathbf{x})) = [V_1(T_\theta(\mathbf{x}_1)), \dots, V_N(T_\theta(\mathbf{x}_N))]$ both defined as collections of N voxels. For clarity and simplicity we denote the images as $I = [I_1, \dots, I_N]$ and $V = [V_1, \dots, V_N]$. Each of the voxel intensities is defined as a mapping from the coordinate space $\mathcal{X} \subseteq \mathbb{R}^d$ to the intensity space $\mathcal{I} \subseteq \mathbb{R}$, $I_i, V_i : \mathcal{X} \rightarrow \mathcal{I}, i = 1 \dots N$, where d is the dimensionality of coordinate space. The transformation function, $T_\theta : \mathcal{X} \rightarrow \mathcal{X}$, is characterized by the parameters $\theta \in \Theta$. Typical examples of transformation functions include affine transformation containing rotation, translation, scale and skew as special cases; and non-rigid transformation based on e.g. B-spline modeling that allows capturing non-linear free-form image transforms. The goal is to find the transformation parameters that maximize the similarity metric $D_N : \mathcal{I}^{N \times 2} \rightarrow \mathbb{R}$ that establishes correspondence among voxels of the moving and the reference images and then maps the intensity values of the reference and the moving images into a single number representing the degree of image similarity. Widely used similarity metrics are mutual information [10] and normalized mutual information (NMI) [20]¹.

The registration parameter optimization problem associated with similarity metric D_N can be formally stated as:

$$\theta_{\text{opt}} = \arg \max_{\theta \in \Theta} D_N[I(\mathbf{x}), V(T_\theta(\mathbf{x}))]. \quad (1)$$

One of the major challenges encountered when solving this problem is the fact that the computational cost of computing D_N increases with N (sometimes linearly and sometimes quadratically). This could be a major obstacle when using registration algorithms in computer assisted interventions and other real time or soft real time applications. One of the most popular approaches to alleviating this problem is voxel selection, when only a subset of voxels are used to calculate the approximation to the similarity metric. Formally, voxel selection can be defined with a selection operator $\mathcal{S}^M : \mathcal{I}^N \rightarrow \mathcal{I}^M$ reducing the dimensionality of image intensity space from N to M ,

$$\theta_{\text{sel}} = \arg \max_{\theta \in \Theta} D_M[I(\mathbf{x}), \mathcal{S}^M \circ V(T_\theta(\mathbf{x}))]. \quad (2)$$

¹These metrics have been popular in general literature as well as in the context of the RIRE dataset

The approximate solution θ_{sel} based on the correspondences of $M < N$ voxels is thus less computationally expensive.

One popular approach to solving (2), provided we have a sufficiently good initial guess θ_0 , is the iterative Gauss-Newton algorithm:

$$\theta_{n+1} = \theta_n + \mathbf{H}_M^{-1} \nabla_{\theta_n} D_M[I(\mathbf{x}), \mathcal{S}^M \circ V(T_{\theta_n}(\mathbf{x}))], \quad (3)$$

where the Hessian \mathbf{H}_M of the similarity metric with respect to the transformation parameters is given by:

$$\mathbf{H}_M = \nabla_{\theta_n} \nabla_{\theta_n} D_M[I(\mathbf{x}), \mathcal{S}^M \circ V(T_{\theta_n}(\mathbf{x}))]. \quad (4)$$

We propose to extend the voxel selection framework by introducing the optimized voxel sampling probability field (VSPF) $\mathcal{P} \subseteq [0, 1]^N$, $\mathcal{P} = [p_1, p_2, \dots, p_N]$, where p_i in the column vector \mathcal{P} determines the probability that we select i -th voxel for the calculation of the similarity metric and $\sum_{i=1}^N p_i$ is the average number of selected voxels. Here each p_i is related to the randomized voxel selection indicator $d_i \in \{0, 1\}$ such that a given voxel is selected if $d_i = 1$, discarded if $d_i = 0$ and the expectation of the indicator is equal to p_i , $\mathbb{E}\{d_i\} = p_i$. We also propose to use a unique voxel selection \mathcal{S}_n^M at every optimization iteration n obtained by sampling from the VSPF. This results in the continued exploration of the images being registered. In the following sections we investigate how the notion of VSPF is related to the parameter uncertainty and develop a framework for optimizing the VSPF.

III. VSPF OPTIMIZATION BASED ON BAYES UNCERTAINTY

In this section, we present a generative model and a set of assumptions underlying our analysis of the utility of each voxel for the registration procedure and then optimize the VSPF based on the utilities of voxels. One of the core foundations underlying our model is the Bayesian approach. Within the Bayesian framework we assume that the transformation parameters are random variables. The natural justification for this assumption is the fact that the values of the registration parameters realized for a given pair of images depend on a large variety of factors that are impossible or hard to model explicitly. Uncertainty of patient positioning in both modalities and equipment calibration errors are just two of those factors. The most general and rigorous way to handle this uncertainty is to assume that the multiple unknown factors affecting the registration parameters and thus the registration parameters themselves are stochastic (random). The fact that there are many stochastic factors affecting the registration parameters suggests that one could choose a Gaussian distribution as an appropriate prior distribution for the registration parameters in the Bayesian model.

In Section III-A we create a generative model for the voxels of the moving image based on the Taylor expansion of the voxel values with respect to the registration transformation parameters. This model basically quantifies how a small perturbation in a registration parameter vector distorts the given voxel. Based on this model and on the Bayesian formulation of the registration parameter estimation problem, given voxels from the moving image, we quantify the utility of each voxel.

In other words, we quantify the potential that each voxel has towards improving the estimation of registration parameters.

In Section III-B we optimize the VSPF in such a manner that the total registration parameter estimation error calculated from our generative model is minimized via proper assignment of VSPF values. We initially pose the VSPF optimization problem as a constrained problem and then use Lagrange relaxation to find a simpler formulation. Our solution to the problem results in the VSPF that assigns higher sampling probabilities to the voxels that have higher utility. In other words, the proposed sampling strategy based on the optimized VSPF tends to favor those voxels that have better chance in helping to improve the estimation of registration parameters based on our generative model for the moving image.

Note that this section addresses exclusively the theoretical foundations of the proposed framework. Our implementation of these ideas as well as recommendations for setting the parameters are discussed in Section IV.

A. Analysis of Bayes Uncertainty

In this section we describe the generative model that we use to assess the utility of each voxel during the voxel selection procedure. Generative models used to construct practical similarity metrics $D_N(\cdot)$ may be found e.g. in [10]. In this section we briefly state main analysis results without concentrating on technical details. The technical details (including the definitions of expectation and variance operators) are presented in Appendix A.

We propose the following first-order Taylor approximation of the moving image $V_i(T_\theta(\mathbf{x}))$ defined in Section II:

$$V_i(T_\theta(\mathbf{x})) \approx V_i(T_{\mu_\theta}(\mathbf{x})) + \mathbf{g}_i^T (\theta - \mu_\theta) + \xi_i, \quad (5)$$

where $\mathbf{g}_i = \frac{\partial V_i(T_\theta(\mathbf{x}))}{\partial T_\theta(\mathbf{x})} \frac{\partial T_\theta(\mathbf{x})}{\partial \theta} \big|_{\theta=\mu_\theta}$ and ξ_i is modeling noise.

Our generative model and the ensuing analysis of Bayes uncertainty are based on the following set of assumptions:

- \mathcal{A}_1 Registration parameters θ are random variables with prior Gaussian distribution having mean μ_θ and covariance matrix $\mathbf{R}_{\theta\theta}$.
- \mathcal{A}_2 The modeling noise ξ is Gaussian zero-mean with variance σ_ξ^2 , independent of θ and such that for voxels $i \neq j$ ξ_i is independent of ξ_j .
- \mathcal{A}_3 The cross terms of the covariance matrix \mathbf{R}_{VV} are negligible: $|\mathbf{g}_i^T \mathbf{R}_{\theta\theta} \mathbf{g}_j| \ll \sqrt{(\mathbf{g}_i^T \mathbf{R}_{\theta\theta} \mathbf{g}_i + \sigma_\xi^2)(\mathbf{g}_j^T \mathbf{R}_{\theta\theta} \mathbf{g}_j + \sigma_\xi^2)}, \forall i \neq j$.

Now we introduce the set $\mathcal{D} = [d_1, d_2, \dots, d_N]$, composed of the randomized voxel selection decisions $d_i \in \{0, 1\}$. This is related to the vector of selected voxels $V_{\mathcal{D}} = [\dots V_i, \dots]^T$, where V_i is included in $V_{\mathcal{D}}$ if $d_i = 1$. Recall that the decision vector is related to the VSPF through expectation: $\mathbb{E}\{d_i\} = p_i$. In other words, the i -th element of VSPF defines the frequency of making the decision $d_i = 1$ corresponding to selecting the i -th voxel for similarity metric evaluation.

Let the optimal Bayesian estimator $\hat{\theta}_{\mathcal{D}}$ obtained using a realization of the random vector $V_{\mathcal{D}}$ of voxel selection decisions. We chose the variance of this optimal Bayesian estimator as a natural characterization of the Bayesian uncertainty. Under

our modeling assumptions listed above it can be shown (please refer to Appendix A for details) that this variance has the following form:

$$\text{var}(\theta - \hat{\theta}_D) = \mathbf{R}_{\theta\theta} - \sum_{i=1}^N p_i \mathbf{R}_{\theta V_i} \mathbf{R}_{V_i V_i}^{-1} \mathbf{R}_{\theta V_i}^T, \quad (6)$$

where $\mathbf{R}_{V_i V_i} = \text{var}(V_i)$ and $\mathbf{R}_{\theta V_i} = \text{cov}(\theta, V_i)$. Each term in the sum weighted by p_i represents the average contribution of every voxel towards decreasing the variance of estimating the transform parameters on the basis of known moving image values $\{V_i\}_{i=1}^N$.

B. Optimization of VSPF

We now optimize the VSPF by posing the problem of optimizing registration parameter uncertainty characterized by the trace of the error covariance matrix. We chose to characterize registration parameter vector uncertainty by the trace of the parameter covariance matrix, which is a common choice in statistical parameter estimation literature: for a given estimator, the trace of its error covariance matrix is equal to the mean squared estimation error (see e.g. [21], Example 3.6). Consequently, based on expression (6) we introduce $U_i = \text{tr}(\mathbf{R}_{\theta V_i} \mathbf{R}_{V_i V_i}^{-1} \mathbf{R}_{\theta V_i}^T)$ and call this voxel utility. This value characterizes the potential that each voxel has towards decreasing the variance of estimating the transform parameters on the basis of known moving image values $\{V_i\}_{i=1}^N$ under model (5).

Our optimization problem is constrained by the voxel sampling probability p_i being non-negative, upper bounded by $0 < P_h \leq 1$ and is subject to the constraint on the average cost C_{ave} of processing the selected voxels (C is the single voxel processing cost). One could consider assigning different processing costs to different voxels by introducing C_i instead of C . At the same time, in practice we have the case where $C_i = C$. Moreover, there is no loss of generality in supposing that $C_i = C$ as our current results can be easily extended in the case of different voxel processing costs. Because of this and because it makes our derivations more succinct and simpler to understand, we prefer to assume that $C_i = C$. Mathematically this can be formulated as follows:

$$\begin{aligned} \mathcal{P}^* = \min_{\mathcal{P}} \text{tr var}(\theta - \hat{\theta}_D) \\ \text{such that } p_i \geq 0, \quad p_i \leq P_h, \quad \sum_{i=1}^N p_i C = C_{\text{ave}}. \end{aligned} \quad (7)$$

Using Lagrange multipliers and skipping the constant term $\text{tr } \mathbf{R}_{\theta\theta}$ we obtain the following Lagrange function:

$$\begin{aligned} \mathcal{J}(\mathcal{P}) = & - \sum_{i=1}^N p_i U_i - \sum_{i=1}^N \gamma_i p_i \\ & + \sum_{i=1}^N \eta_i (p_i - P_h) + \lambda (C_{\text{ave}} - \sum_{i=1}^N p_i C). \end{aligned} \quad (8)$$

It can be shown (see Appendix B) that the problem (7) can be reformulated as follows using the Lagrange relaxation method [22] and introducing the auxiliary variable $A > 0$:

$$\begin{aligned} \mathcal{P}^* = \min_{A, \mathcal{P}} \mathcal{J}(\mathcal{P}) \\ \text{such that } p_i \geq 0, \quad p_i \leq P_h, \quad \sum_{i=1}^N p_i C = C_{\text{ave}}, \\ p_i = A(U_i + \gamma_i - \eta_i + \lambda C). \end{aligned} \quad (9)$$

The solution of the reformulated problem has the following form:

$$p_i^* = \begin{cases} A(U_i + \lambda^* C) & \text{if } 0 \leq A(U_i + \lambda^* C) \leq P_h \\ P_h & \text{if } A(U_i + \lambda^* C) > P_h \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

Here the auxiliary variable A should be determined from the reformulated minimization problem and we will later show how this can be accomplished by analyzing the properties of the cost function. The lagrange multiplier λ^* is selected to satisfy the average processing cost constraint, $\mathcal{P}^{*T} \mathbf{C} = C_{\text{ave}}$ (here $(\cdot)^T$ is the vector transpose operator and $\mathbf{C} = [C, \dots, C]$ is the vector of costs).

Equation (10) provides a general intermediate solution. This intermediate result has interesting implications. From the set of inequalities $0 \leq A(U_i + \lambda^* C) \leq P_h$ in (10) we can see that for an arbitrary value of A the voxel selection distribution \mathcal{P}^* is comprised of three types of voxels defined by two thresholds: $-\lambda^* C_i$ and $P_h A^{-1} - \lambda^* C_i$. If the utility, U_i , of a given voxel is below the first threshold this voxel is never selected ($p_i^* = 0$). If the utility U_i of the voxel is between the two thresholds, the voxel may or may not be selected with probability proportional to its utility, $p_i^* = A(U_i + \lambda^* C)$. Finally, for voxels deemed very informative, for which utility U_i exceeds the second threshold we have $p_i^* = P_h$ and thus voxels from this class are most frequently selected.

In the following we provide two theoretical results that characterize (i) the behaviour of $\mathcal{J}(\mathcal{P}^*)$ as a function of the auxiliary variable A and (ii) the behaviour of the left hand side of the average processing cost constraint equation as a function of the Lagrange multiplier λ^* . These results will help us in specifying the optimal value of A and advising an efficient numerical procedure to find λ^* .

The optimal value of A can be found by minimizing $\mathcal{J}(\mathcal{P}^*)$. The following proposition, whose proof can be found in Appendix C, establishes that the cost function \mathcal{J} evaluated at the optimized solution \mathcal{P}^* is a non-increasing function of A .

Proposition 1: $\mathcal{J}(\mathcal{P}^*)$ is a monotonically non-increasing function of $A, \forall A \geq 0$.

This implies that the formal optimal solution of the voxel selection problem is thus achieved when $A \rightarrow \infty$. It is straightforward to deduce that it has the following form (recall that λ^* is a function of A):

$$p_i^* = \begin{cases} \lim_{A \rightarrow \infty} A(U_i + \lambda^* C) & \text{if } -\lambda^* C \rightarrow U_i \\ P_h & \text{if } U_i > -\lambda^* C \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

Before presenting another important result leading to a procedure to efficiently compute the value of λ^* we define $\varphi(\lambda^*) \triangleq \mathcal{P}^{*T} \mathbf{C}$.

Proposition 2: Function $\varphi(\lambda^*)$ is a monotonically non-decreasing function of λ^* , $\forall \lambda^*$.

The proof of this result appears in Appendix D. Since the function defining constraint equation $\mathcal{P}^{*T} \mathbf{C} = C_{\text{ave}}$ is monotonic, we can use efficient algorithms, for example, golden ratio search to find λ^* .

Solution (11), obtained in the case $A \rightarrow \infty$, assigns the highest probability, P_h , to all the voxels for which voxel utility is greater than the threshold $-\lambda^* C$. Additionally, there is exactly one voxel, whose probability $\lim_{A \rightarrow \infty} A(U_i + \lambda^* C)$ is selected to exactly satisfy the constraint $\mathcal{P}^{*T} \mathbf{C} = C_{\text{ave}}$. Note that without the contribution of this voxel $\mathcal{P}^{*T} \mathbf{C}$ can only be incremented by $P_h C$ with each new voxel whose utility exceeds the threshold $-\lambda^* C$.

According to our experiments, the value of the cost function saturates as $A \rightarrow \infty$. Hence, large increments of A tend to result in small increments of the cost function. Thus, the value of the cost function is almost the same at the optimal solution and at a suboptimal solution obtained with $A < \infty$. Typically, voxel selection strategies that define a fixed voxel subset tend to reduce the optimization capture range, robustness to local minima and exploratory power of the sampling scheme i.e. its ability to visit different image regions during optimization. This is true even if the voxel selections are optimized. One of the reasons for this is that it is difficult to correctly estimate the true value (utility) of a given voxel for a registration procedure. For any voxel selection criterion it is safe to say that noise and modeling errors result in highly unreliable estimates of the true value that a given voxel has for the registration problem. As a result, fixing voxel selection based even on the "most optimal" criterion may have catastrophic results for the registration problem if voxel value estimate is erroneous. On the other hand, the uniformly random sampling might decrease registration accuracy, but is more robust, since sampling induces image exploration and does not imply any bias towards particular voxel subsets. The general solution (10) with $A < \infty$ results in a more diverse VSPF since it assigns non-zero probabilities to a greater variety of voxels. At the same time, according to this solution, the additional voxels assigned with non-zero probabilities are defined based on their utilities and their probabilities are proportional to the respective utilities. This unique feature has potential to endow a sampling scheme with (i) vast exploratory power, (ii) ability to pick voxels most useful for achieving an accurate registration, (iii) ability to avoid catastrophic registration errors due to voxel utility modeling and estimation artifacts. These properties would help a sampling scheme achieve the robustness of randomized sampling approaches and improve accuracy by utilizing a greater proportion of meaningful voxels for similarity metric calculations.

We thus propose a sampling scheme based on (10), where A is taken as small as possible to diversify VSPF and as large as necessary in the view of being able to satisfy the average cost constraint. Here the latter requirement is due to the fact that for

too small values of A no λ^* exists that could satisfy constraint equation $\mathcal{P}^{*T} \mathbf{C} = C_{\text{ave}}$. Another parameter that regulates the exploratory power of the VSPF is the upper bound on the probability values, P_h . Lower values of this parameter tend to favour image exploration. This parameter could be chosen to be a universal constant whose reasonable value could be fixed based on experience. We also propose a more accurate strategy to learn the value of this parameter based on a small training dataset. This way it could be adapted to the problem at hand, optimization strategy and other specifics of the registration system. In the following section, where we describe a practical multi-scale algorithm, we discuss both variants.

Finally, let us examine how our approach relates to other approaches discussed in the literature. Consider the case where $A \rightarrow \infty$, $P_h = 1$, C_{ave} is an integer number and all the values of U_i are unique (recall that N is the total number of voxels in the image). According to (11), all the VSPF values will be either 0 or 1, i.e. voxels are either selected or ignored. In such context, the proposed voxel sampling methodology is similar in spirit to the selection strategy proposed by Reeves and Hezar [13], i.e. calculate some measure of voxel utility and compare it to a threshold. Consider an alternative case where $P_h = 1$ and all U_i values are equal (voxels are deemed equally important/informative). According to (11), the optimal solution then is to assign every voxel equal probability and $p_i^* = \lim_{A \rightarrow \infty} A(U_i + \lambda^* C) = C_{\text{ave}}/N$. In other words, such context yields an optimal voxel selection policy similar to the random uniform sampling proposed in [10].

IV. PRACTICAL DESIGN GUIDELINES

We now describe the full practical implementation of our voxel sampling scheme, whose high-level algorithmic summary is outlined in Fig. 1. The first stage involves creating an image pyramid of K -scales, where each level involves low-pass filtering and subsampling the original reference and moving images. Optimization is then performed in order of scale. Hence, the level with coarsest scale is optimized first, and the result of each optimization is fed as an initial point to the following level.

The VSPF is calculated from the moving image once at each scale using initial transformation parameter values and following the framework described in previous sections. In practice VSPF calculations can be performed either on the fixed or on the moving image, whichever is easiest. The choice mainly depends on the interpolation strategy (partial volume or linear/cubic) and the point at which voxels are sampled (in the moving image coordinate frame or in the fixed image coordinate frame). The computation of VSPF can be performed in different registration contexts. For example, in the context of the image guided neurosurgery one of the images (e.g. MRI) is often available before an operation and VSPFs can be computed ahead of time using parameter initializations corresponding to the identity transformation (no registration) to save time during registration. At the same time, if data are not available ahead of time, the proposed framework can be applied online, once per scale, (the way it is depicted in Fig. 1), since the computational expense of computing

- 1) Calculate K -scale image pyramid for the moving and the reference images;
- 2) Pick initialization θ_0 and set parameters C_{ave}, P_h as discussed in Section IV;
- 3) For every scale $k = K, K-1, \dots, 1$ do
 - a) Calculate utilities $\{U_i\}_{i=1}^N$ for all voxels of the moving image;
 - b) Estimate VSPF \mathcal{P}_k as discussed in Sections III-B and IV;
 - c) For iteration $n = 1, 2, \dots, L$ do
 - Generate new voxel selection by sampling \mathcal{P}_k ;
 - Calculate Jacobian and Hessian based on the voxel selection;
 - Update parameters using trust region Gauss-Newton iteration;
 - d) Pick θ_0 for the next scale based on the last iteration (account for scale change);

Fig. 1. High level summary of the proposed algorithm

the VSPF is small and consists of matrix manipulations with spatial image gradients and transform derivative matrices.

A. Parameter Selections

Voxel utility, U_i , is computed according to its definition in Section III-A as $U_i = (\mathbf{g}_i^T \mathbf{R}_{\theta\theta} \mathbf{g}_i + \sigma_\xi^2)^{-1} \text{tr}(\mathbf{R}_{\theta, V_i} \mathbf{R}_{\theta, V_i}^T)$, where $\mathbf{R}_{\theta, V_i} = \mathbf{R}_{\theta\theta} \mathbf{g}_i$. Recall that $\mathbf{g}_i = \frac{\partial V_i(\mathbf{T}_\theta(\mathbf{x}))}{\partial \mathbf{T}_\theta(\mathbf{x})} \frac{\partial \mathbf{T}_\theta(\mathbf{x})}{\partial \theta} \big|_{\theta=\mu_\theta}$ is the moving image intensity derivative with respect to the transformation parameters and that $\mathbf{R}_{\theta\theta}$ is the cov(θ, θ). The derivative can be obtained using analytical differentiation of $\mathbf{T}_\theta(\cdot)$ and the numerically calculated gradient of the moving image with respect to the coordinates. The gradient of the image with respect to coordinates is computed on the image lightly smoothed with gaussian kernel to alleviate gradient estimation noise.

We approximate $\mathbf{R}_{\theta\theta}$ with the aid of Hessian \mathbf{H} of the similarity metric as $\mathbf{R}_{\theta\theta} \propto \mathbf{H}_M^{-1}$. A detailed discussion on the validity and applicability of such an approximation can be found in [23]. The Hessian \mathbf{H} is often available from the optimization routine. If not, its numerical approximation can be easily constructed. If Hessian \mathbf{H} is not available from the optimization routine or is too costly to compute, it is possible to replace it with the second order derivatives of the log-likelihood function resulting from model (5). We set $\mathbf{R}_{\theta\theta} = \mathbf{H}_M^{-1}$ and $\sigma_\xi^2 = 1$ in our implementation. We have experimented with different settings of σ_ξ^2 ranging from 0.1 to 100 and found that the performance of the algorithm is not at all sensitive to the value of σ_ξ^2 . To initialize the parameter μ_θ at a given resolution scale we use its estimate obtained from the registration at the previous resolution scale.

Once utilities are calculated for each voxel we can compute the VSPF in accordance with (10). As was mentioned in the previous section we keep the value of the free parameter λ^* minimal to diversify voxel selections and sufficient so that λ^* exists such that the constraint equation $\mathcal{P}^{*T} \mathbf{C} = C_{\text{ave}}$ can be satisfied. Furthermore, we compute the value of Lagrange multiplier λ^* by finding the root of equation $\mathcal{P}^{*T} \mathbf{C} = C_{\text{ave}}$ via bisection search (using the result stated in Proposition 2). In the constraint equation (see (7)), C_{ave} is the total average cost of processing the selected voxels. It is specified by the user e.g. in the form $C_{\text{ave}} = cN$ where c is the user specified average proportion of voxels being selected by the algorithm

(the sampling rate). The high-level outline of the registration algorithm employing proposed voxel sampling algorithm is shown in Fig. 1.

Finally, the last parameter of the algorithm is P_h , the highest probability that could be assigned to a voxel. In the first approach that we explore, for every scale we simply adjust P_h heuristically such that more image exploration is induced in lower resolution scales while more emphasis is put on selecting voxels with higher utility in higher resolution scales. We found empirically, via experiments with the Brainweb [24] dataset, that there exist simple heuristic schedules according to which P_h is set proportional to C_{ave}/N that produce universally good results. We present these schedules in more detail in Section V.

The second approach that we propose is to learn P_h using a small dataset representative of the particular problem context. In the application domain, the specifics of the particular registration problem often affect the choice of similarity metric, optimization strategy and interpolation scheme. This implies that at least some training information in the form of the small set of exemplar image pairs from the problem-specific modalities using certain acquisition and post-processing protocols should be available to the registration algorithm designer to guide the algorithm development.

Based on the assumption that we have a training data set and the gold standard registration parameters for the image pairs in this dataset we formulate the empirical learning criterion $Q^r(P_h^r)$ for each resolution level r . We define the empirical target registration error (ETRE) as the average over V image pairs in the training dataset and U Monte-Carlo trials:

$$Q^r(P_h^r) = \frac{1}{V} \frac{1}{U} \sum_{v=1}^V \sum_{u=1}^U \|\mathbf{X}_v - \hat{\mathbf{X}}_{u,v}^r(P_h^r)\|_2^2. \quad (12)$$

Here \mathbf{X}_v is the set of transformed coordinates obtained using gold standard registration parameters for image pair indexed by v and $\hat{\mathbf{X}}_{u,v}^r(P_h^r)$ is the set of transformed coordinates for image pair v and Monte-Carlo trial u found using the empirical estimate of the registration parameters obtained via the optimization of the similarity metric at resolution scale r using the proposed voxel sampling algorithm with a given value of P_h^r . As the voxel sampling algorithm is randomized, some degree of Monte-Carlo averaging could be beneficial

if V is relatively small (3...5 images). Thus we repeat the registration procedure for the same candidate value P_h^r , level r and image pair v U times and calculate $\hat{X}_{u,v}^r(P_h^r)$ based on the new registration parameter estimate each time.

We propose to learn the value of P_h^r by minimizing the ETRE $Q^r(P_h^r)$:

$$\hat{P}_h^r = \arg \min_{P_h^r \in [0;1]} Q^r(P_h^r). \quad (13)$$

The function $Q^r(\cdot)$ is generally extremely irregular and non-smooth, because of the possible registration failures and because of complex dependence of the ETRE on the value of P_h^r . At the same time, the domain of this function is well defined and restricted. Thus any optimizer capable of performing global or quasi-global search on a restricted interval using only the objective function values will suffice to solve this problem. In order to find \hat{P}_h^r we propose using either the particle swarm optimization (PSO) [25] or the exhaustive grid search with grid step 0.01. Our algorithm proceeds by finding \hat{P}_h^K , the value for the scale with the lowest resolution. The algorithm then uses the identified value of \hat{P}_h^K to find the estimate of the registration parameters at resolution level $K - 1$. This procedure iterates until the values of highest probabilities for all resolution levels $K, K - 1, \dots, 1$ are found.

V. EXPERIMENTS WITH THE RIRE VANDERBILT DATASET

A. RIRE Dataset

To test the proposed algorithm we made use of the real clinical data available in the popular, publicly available RIRE Vanderbilt dataset [19]. The performance of algorithms was evaluated by registering 3D volumes corresponding to CT images to geometrically corrected MR images. MR image set included images acquired using T1, T2 and PD acquisition protocols. The total number of different image pairs used was 19. Those pairs were taken from patients 001, 002, 003, 004, 005, 006, 007 for which geometrically corrected images are available. Patients 003 and 006 did not have geometrically corrected PD and MR-T1 images respectively. According to the data exchange protocol established by the RIRE Vanderbilt project, registration results obtained via algorithms under the test were uploaded to the RIRE Vanderbilt web-site. Algorithm evaluation results were calculated by the RIRE Vanderbilt remote computer using the gold standard transformation not available to us and published on their web-site in the form of tables containing target registration errors (TREs) calculated over 6 to 10 volumes of interest (VOIs) for each image pair. For patient 000 geometrically corrected MR-T1, MR-T2, MR-PD images and corresponding CT image are available along with the set of transformed coordinates obtained using gold standard registration parameters. Three image pairs from patient 000 were used to learn the values of P_h in one of the versions of the proposed algorithm according to the procedure described in Section IV-A.

B. Experimental Setup

All images were first resampled to a common 1mm grid using bicubic interpolation. We used 2-scale registration based

on the low-pass filtered and downsampled image pyramid containing two resolution scales. Resolution level number two had grid spacing 2 mm along each axis and resolution level number one had grid spacing 1mm along each axis. The estimate of the registration parameters obtained at a lower resolution level was used as a starting point for the registration at the next higher resolution level; level 2 had all its parameters initialized to zero values. We concentrated on recovering 6 rigid registration parameters (3 translations and 3 rotations) using the NMI similarity metric [20]. Histogram for the evaluation of the similarity metric was calculated using the partial volume approach with Hanning windowed sinc kernel function [26]. The similarity metric was optimized using the trust region Gauss-Newton approach [27].

We evaluated the performance of the following voxel sampling approaches. The uniformly random sampling (**URS**) technique consists of randomly selecting voxels with equal probabilities at every iteration [10]. At a given resolution level r all voxels have equal probability of being selected, M/N_r if $M < N_r$ and 1 if $M \geq N_r$; the average number of selected voxels is thus equal to $\min(M, N_r)$ at each resolution level. Note that we used equal number of selected voxels for all resolution scales. As a variation of this scheme, we also considered the uniformly random sampling that samples voxels once per scale and keeps them fixed during iterations (**fURS**) [9]. Gradient magnitude sampling (**GMS**), a slight modification of gradient based subsampling originally proposed by Subuncu and Ramadge [16], consists of calculating the spatial gradient magnitude $\|\nabla V_i\|_2 = \sqrt{(\partial V_i / \partial x_i)^2 + (\partial V_i / \partial y_i)^2 + (\partial V_i / \partial z_i)^2}$ and sampling voxels at every optimization iteration according to the probabilities proportional to it, where the proportionality coefficient is chosen so that the average number of voxels selected at every resolution scale is equal to M . The extension of the previous approach, **GMS+URS**, [17], samples voxels from the VSPF defined as a convex combination of VSPFs generated via **GMS** and **URS** approaches. The mixing parameter is optimized using a small training dataset (in our case this is the dataset obtained from patient 000). For completeness we also present results for the gradient magnitude (**GM**) approach that deterministically selects a portion of voxels with highest gradient magnitudes once per scale. The proposed method consists of evaluating voxel utilities and assigning voxel probabilities based on these utilities. We evaluated two versions of the proposed algorithm. For the first option (**ProposedH**) we set P_h to $\min(1, 3M/N_2)$ for the resolution level 2 (coarse resolution) and $\min(1, 10M/N_1)$ for the resolution level 1 (fine resolution). These are the simple heuristic schedules (see Section IV-A). For the second option (**ProposedL**) we learn the value of P_h for every scale and voxel sampling rate using methodology described in Section IV-A with the aid of the training dataset consisting of the image pairs taken from patient 000. For completeness, we also present the **ProposedD** method, which is based on the thresholding of the proposed VSPF (ProposedD deterministically selects a portion of voxels with highest voxel utilities). We evaluate these algorithms for the following values of voxel sampling rates (given in %): $M/N \in \{0.06, 0.1, 0.3, 1, 3, 10\}$ (sampling rate is calculated with respect to the image size at the highest

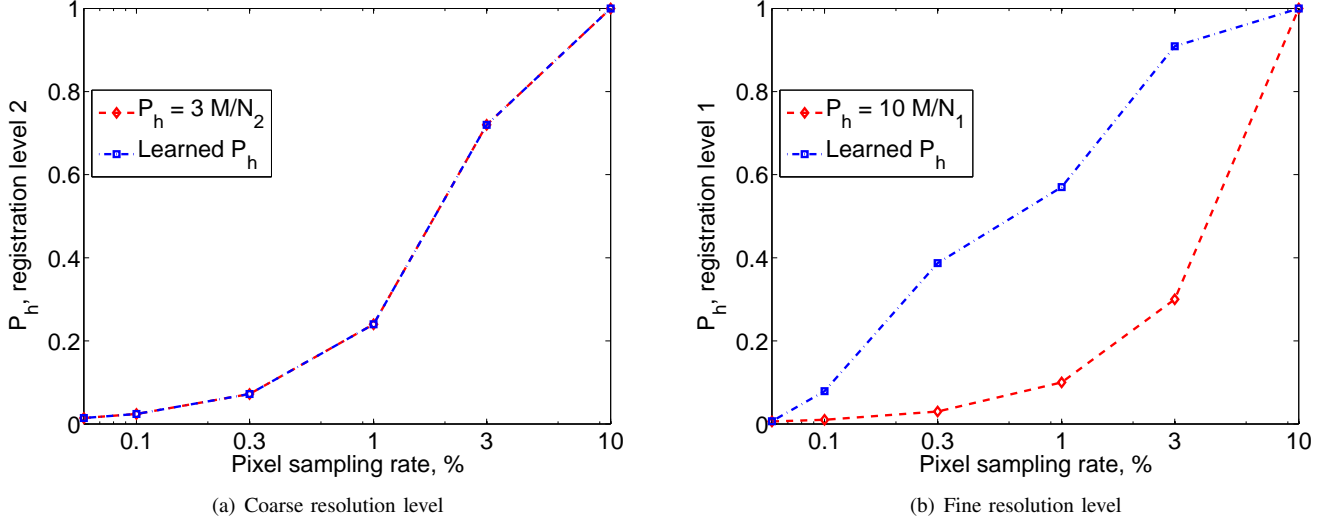


Fig. 2. Heuristic and learned schedules of P_h for the proposed algorithm. Note that abscissa is in the log scale.

resolution level, $N = N_1$). For the **GMS+URS** approach we learned the values of the mixing parameter via the exhaustive grid search procedure with grid step 0.01 and obtained the following results: $\beta^2 = [0.98; 0.93; 0.94; 0.97; 0.97; 0.95]$, $\beta^1 = [0.33; 0.29; 0.25; 0.22; 0.09; 0.00]$. The mixing parameter (denoted as β^r in the original paper [17]) is responsible for finding a good balance between the **GMS** and the **URS** approaches within the **GMS+URS** sampling scheme. Each value of β^r presented above corresponds to one of the sampling rates used to evaluate the algorithms.

C. Results

Figure 2 compares two schedules for setting the P_h parameter of the proposed algorithm. The heuristic schedule $\min(1, 3M/N_2)$ for coarse resolution and $\min(1, 10M/N_1)$ for high resolution level is contrasted with the schedule learned using small training dataset methodology presented in Section IV-A. Note that the curve corresponding to the heuristic schedule does not appear to be linear in the plot, because abscissa is in logarithmic scale. It can be seen from Fig. 2(a) that for the coarse resolution level there is not much difference in the heuristically proposed and the learned schedules. On the other hand, Fig. 2(b) showing results for the higher resolution level reveals that, according to the learned schedule, P_h grows at a much steeper rate than $10M/N_2$. This suggests that for the higher resolution level there is more variation in the utility of different voxels and there is a need for choosing higher values of P_h to enhance the stratification among voxel probabilities to boost selection frequencies for more informative voxels. Another interpretation of this result is that when the total average number of sampled voxels becomes large enough, we can allow the sampling scheme to focus almost exclusively on voxels with high utility values, because the exploratory power of this voxel subset becomes strong enough.

Qualitative results are presented in Figure 3 showing an example of VSPFs obtained via the ProposedL, ProposedD, **GMS+URS**, **GMS** and **GM** approaches (**URS** and **fURS** are not shown since their VSPFs are uniform and ProposedH

was visually similar to ProposedL). The figure is obtained for $C_{ave} = 0.01N$, corresponding to 1% of voxels being sampled on average. We can see that the proposed approach allows for the selection from a comparatively large variety of voxels (compared e.g. to the **GM** or **ProposedD** approaches). We also see that there are voxels in the proposed VSPF that are assigned significantly higher probabilities than some other voxels. These voxels are deemed more informative by our sampling scheme. The **GMS** sampling scheme assigns probabilities proportional to voxel gradient magnitudes. The **GMS+URS** scheme tries to find the compromise between **GMS** and **URS** schemes by adding uniform probabilities to the **GMS** VSPF. The **GM** scheme selects a subset of voxels with highest gradient magnitudes and assigns them probabilities equal to 1 (voxels are always selected). The selection of voxels produced by the **GM** scheme is not rich as they are mostly concentrated along the most prominent spatial gradient structures. Similar comments apply to the **ProposedD** mask that selects a subset of voxels having most prominent utility values. Moreover, note that the proposed approach favours more peripheral voxels useful for angular alignment, as opposed to the **GM**, **GMS** and **GMS+URS** approaches that do not account for parameter uncertainty and concentrate only on image gradient details. We argue that this characteristic, along with the ability to diversify voxel selections brings forward an improved registration robustness and accuracy. On the other hand, large areas populated with uninformative voxels have zero VSPF values. Thus, by concentrating image exploration only on potentially informative voxels, the proposed approach improves the accuracy over the **GMS+URS**, **URS** and **fURS** approaches.

Figure 4(a) shows registration failure rate for different voxel selection mechanisms as a function of the sampling rate shown in the log scale. We define a registration failure as any case with error exceeding 10mm in any of the VOIs. The x axis of this plot shows voxel sampling rates $\{0.06, 0.1, 0.3, 1, 3, 10\}$. We can see that both variants of the proposed approach employing sampling from the VSPF consistently outperform

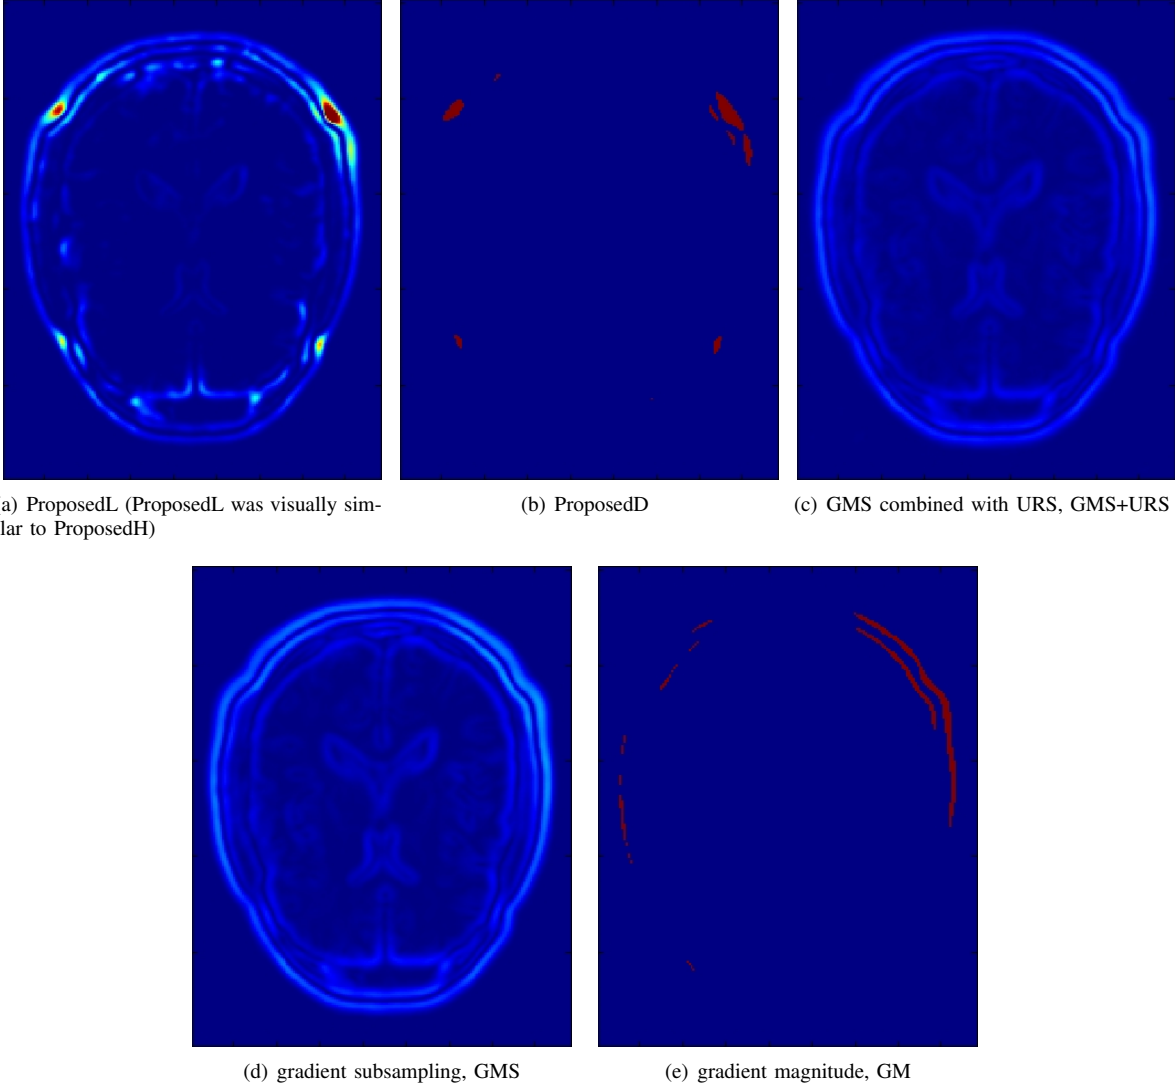


Fig. 3. VSPFs for the highest resolution level shown in jet colour scheme. The jet color scheme uses colour assignment as follows: black corresponds to zero ($p_i = 0$); cyan, yellow and orange correspond to intermediate values between 0 and 1; red corresponds to ($p_i = 1$).

other approaches in terms of robustness, because both methods employ heavy exploration of the voxel values by continuous sampling combined with emphasizing more important voxels.

Figure 4(b) shows the mTRE (mean TRE) as a function of the average time spent per registration (note that the x-axis is in the log scale). Recall that RIRE provides us with 6 to 10 TRE values for each image pair. We compute the mTRE as the mean TRE of all the successful (non-failed) registrations corresponding to a particular experimental case defined by the selection of voxel sampling scheme and voxel sampling rate. For objective comparison we show only those points that correspond to zero failure rate, because we deem this scenario as the one that has most practical importance. Besides, this helps avoiding the situation where a method demonstrates anomalously low mTRE, because it fails many difficult cases. The markers in Figures 4(b) and 4(a) correspond to the same instances of the sampling rate. It can be seen that the average time per registration is roughly proportional to the sampling rate (for the proposed method the calculation of the VSPF induces only a small overhead). The mTRE is mini-

mal for the proposed approach with learned P_h (ProposedL) compared to other methods. The ProposedD method has the best performance for sampling rates 3 and 10%, however, it simply does not work for lower sampling rates because of failures. Same applies to the GMS and the GM approaches, both having only two points shown in the plot because of failures. The proposed approach retains high level of accuracy and robustness even with low voxel sampling rates. Moreover, the proposed approach with the heuristic setting of P_h (ProposedH) is not much worse than the proposed approach with learned P_h (ProposedL). This shows relative insensitivity of the proposed approach to the changes in parameter P_h . The properties of the proposed sampling approach allow to significantly reduce computational time and simultaneously guarantee high accuracy and robustness in a practical system. Such results support our conjecture that image exploration induced by VSPF based voxel sampling maintains the robustness due to continued image exploration and improves registration accuracy by selecting more informative voxels. Overall, the proposed technique at 0.06% sampling rate is better than other

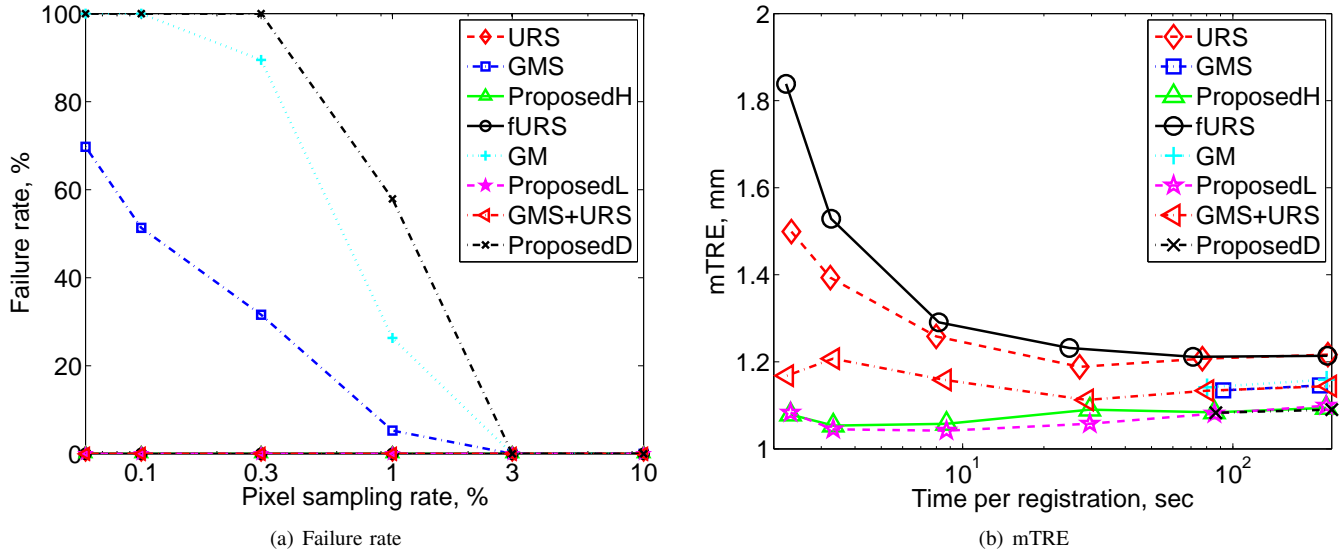


Fig. 4. Failure rate (a) and mTRE (b) for different voxel selection mechanisms: gradient magnitude (GM), gradient magnitude sampling (GMS), proposed with learned P_h (ProposedL) and heuristic P_h (ProposedH), ProposedD corresponding to the thresholded proposed VSPF, uniformly random sampling (URS), uniformly random sampling with fixed subsets (fURS), convex combination of GMS and URS (GMS+URS). Note that the proposed approach consistently outperforms in terms of robustness and accuracy.

techniques at 10% sampling rate, maintaining 0% failure rate and around 1 mm mTRE accuracy. This amounts to more than hundredfold acceleration for the registration procedure. None of the alternative techniques achieve this level of performance.

VI. DISCUSSION OF FINDINGS

We have demonstrated the promise of the proposed theoretical methodology based on the experiments with the 3D RIRE Vanderbilt dataset containing real clinical data. The main finding of the paper is that with the aid of a carefully designed VSPF (in our case relying on the principle of the Bayesian registration parameter uncertainty) we can maintain the appropriate levels of accuracy and robustness of a registration procedure at very low voxel sampling rates.

At the same time, our validation methodology has a number of limitations and calls for further investigation and analysis of the properties of the proposed theoretical framework. We conducted our experiments in the context of rigid registration. In some contexts rigid registration is appropriate even during image guided interventions [7]. However, some other contexts may call for the use of non-rigid registration approaches (e.g. in the context of image guided neurosurgery, open craniotomies cause non-rigid intra-operative brain movements also known as brain shift [28]). The proposed theoretical framework does not pose any principal limitations on the size or properties of the registration parameter set. At the same time, with the current validation setup we cannot firmly state if the proposed methodology is directly applicable to the registration problems with non-rigid parameterization. Intuitively, however, we may expect even better performance gains for the non-rigid registration problems. The reason for this is that the data (voxel) support for each of the registration parameters is highly spatially localized. The effects of differing voxel utilities (abilities to reduce registration error) are likely to be even more pronounced in this setting. This should result in

increased gain from a smart sampling strategy and may call for further extensions. For example, including parameters into the sampling scheme and deciding on a subset of parameters to update at a given iteration. Often, many parameters in complex non-rigid registration problems converge early and do not require frequent updates during the rest of the procedure.

Another limitation of the validation methodology is that the evaluation of the proposed method was performed on the RIRE Vanderbilt dataset, which does not necessarily target real-time applications. Real-time applications often use noisy images, resulting e.g. from ultrasound or fluoro modalities [3], [7]. Based on the current validation scheme it is not clear if the proposed method will provide the same performance gain in the case of the low quality images. The study of the proposed method in this context is thus a necessary important future research task. At the same time, we would like to note that our VSPF calculation methodology requires only one of the images to compute the sampling probabilities. In the common real-time applications outlined above one of the images: MRI, CT or another high quality image, is typically available before the application of the registration procedure. In this context the high quality image can be defined as a moving image and used to compute the VSPF off-line. This scenario is very close to our current validation setup in terms of the images used to compute the VSPF. Our validation setup has its limitations, but we believe that it is rigorous enough to demonstrate the promise of the proposed theoretical framework and to approximate some of the realistic real-time applications.

Another important finding is that from our experiments (see Fig 4(b)) it follows that, at least for the rigid registration problem, the calculation of the VSPF does not induce any significant overhead. Recall that we compute the VSPF only once per registration scale and do not recompute it at every iteration. With this scheme we achieve significant speed-up in the registration procedure and keep robustness and accuracy

of the method at very high levels. Although the complexity of computing the VSPF scales linearly with the number of registration parameters, so does the complexity of computing the similarity metric and its derivatives. Thus we have a reason to believe that the implementation of the proposed methodology in the context of non-rigid registration problems will not result in any disproportional overhead. Furthermore, for many problems the high-quality image (e.g. diagnostic MRI) is available well before the registration procedure must be executed. In this context we may further reduce the overhead by computing the VSPF off-line and then just reusing it during the operation. Moreover, we can imagine much more complicated sampling schemes being implementable in such scenario. For example, for a non-rigid problem we could define its own VSPF for every registration parameter and compute it off-line. During the application of the registration procedure we would sample optimal voxels for each parameter individually and then combine all of them into the common set to evaluate the similarity metric.

Our sampling approach uses different pixel subsets at every optimization iteration. This results in the continued exploration of the images being registered, but also induces noise in the approximations of the cost function and its derivatives. In the context of the Gauss-Newton optimization used in our paper, changing voxel selections over iterations does not pose any problem. It should be noted that although we presented the concepts of VSPF and continued voxel sampling within the Gauss-Newton optimization framework, they have a much wider applicability. For example, the global CMAES optimizer [29] widely used for complex optimization problems (including the non-rigid registration) is capable of handling the noise in cost function values. Different variations of gradient descent algorithms are not sensitive to noise in derivatives as well, especially when combined with stochastic optimization approaches [30]. Thus the proposed framework can be used with these optimization approaches directly. At the same time, the proposed framework is not directly applicable to the quasi-Newton (non-zero memory) optimizers where an estimate of the Hessian of the cost function is built up iteratively. It may well be that the proposed sampling strategy could be adapted for the use with such optimizers (e.g. via proper use of common random numbers techniques [31]). However, investigating this matter is beyond the scope of this paper.

Finally, the current approach is formulated for the direct registration methods using only voxel values. It would be interesting to look for the extensions of this method in the context of feature based registration. To directly extend the proposed method in this context one could formulate a Bayesian model for the analysis of registration parameter uncertainty as a function of selected feature set. The next step in this direction would be to define and optimize the feature sampling probability field based on the uncertainty analysis.

VII. CONCLUSIONS

We have presented a new uncertainty based multi-scale voxel selection framework for the context of multi-modal medical image registration. At each resolution scale, a voxel sam-

pling probability field (VSPF) is optimized based on the transformation parameter uncertainties derived from the Bayesian model and the optimal voxel subsets are then sampled from the VSPF. The proposed sampling scheme exhibits the advantages of both random sampling techniques, which are robust but not necessarily accurate, and fixed sampling techniques, which are accurate but not necessarily robust. Specifically, the proposed sampling strategy focuses on both informative image details that have potential to contribute towards reducing registration uncertainty whilst simultaneously maintaining the variability and exploratory power of the sampled voxel set. The experimental results on the 3D, multimodal, publicly available RIRE dataset demonstrate that our approach yields significant improvements in registration performance, compared to other voxel selection techniques such as random voxel sampling and gradient sampling, particularly at aggressive sampling levels (e.g. less than 1% of voxels). Overall, the proposed technique at 0.06% sampling rate was better than other tested techniques at 10% sampling rate, maintaining 0% failure rate and around 1 mm mTRE accuracy. This amounts to more than hundredfold reduction in the number of voxels used for similarity metric calculations, which translates to similar accelerations for the registration procedure. This leads to minimization of the computational time of the registration task while also minimizing the potential penalty in terms of accuracy and robustness (i.e. failure rate). The approach is general and can provide an efficient solution applicable to a wide array of important, time-sensitive medical imaging registration problems, such as is required in image guided interventions. Future work will focus on experimental validation of the registration scheme in various interventional contexts.

APPENDIX A

ANALYSIS OF BAYES UNCERTAINTY: TECHNICAL DETAILS

First, we introduce some probabilistic notation. A real valued vector random variable $X : \Omega \rightarrow \mathbb{R}^d$ is defined on the probability space (Ω, \mathcal{F}, P) equipped with sample space Ω , σ -algebra of events \mathcal{F} and probability measure P . The mathematical expectation is then defined as the following Lebesgue integral:

$$\mathbb{E}\{X\} = \int_{\Omega} X(\omega)P(d\omega). \quad (14)$$

For two random vectors Y and X we also define the covariance matrix $\mathbf{R}_{YX} = \text{cov}(X, Y)$ as

$$\text{cov}(X, Y) = \mathbb{E}\{(X - \mathbb{E}\{X\})(Y - \mathbb{E}\{Y\})^T\}, \quad (15)$$

and variance $\text{var}(X) = \text{cov}(X, X)$.

Under our assumptions \mathcal{A}_1 , \mathcal{A}_2 and by the definition of VSPF we have three primary random vectors in our paper: the vector of registration parameters θ , the sequence of modeling noises $\xi_i, i = 1 \dots N$ and the vector of decision variables \mathcal{D} . In what follows all the expectations are taken with respect to the joint probability measure induced by the collection of these three random vectors unless specifically stated otherwise

or indicated by the conditional expectation $\mathbb{E}\{X|Y\}$ defined as:

$$\mathbb{E}\{X|Y\} = \int_{B \subseteq \Omega: Y=Y(\omega)} X(\omega) P(d\omega). \quad (16)$$

According to the model (5) and under assumptions \mathcal{A}_1 and \mathcal{A}_2 , $V_i(T_\theta(\mathbf{x}))$ has a Gaussian distribution with mean $V_i(T_{\mu_\theta}(\mathbf{x}))$ and variance $\mathbf{g}_i^T \mathbf{R}_{\theta\theta} \mathbf{g}_i + \sigma_{\xi}^2$. If we pose the problem of optimally estimating the parameters θ from the generative model (5) then the Bayesian estimate of parameters θ is simply

$$\hat{\theta} = \mu_\theta + \mathbf{R}_{\theta V} \mathbf{R}_{V V}^{-1} (V - \mu_V). \quad (17)$$

where $\mu_V = \mathbb{E}\{V\}$, $\mathbf{R}_{V V} = \text{var}(V)$, $\mathbf{R}_{\theta V} = \text{cov}(\theta, V)$. On the other hand, the optimal Bayesian estimate $\hat{\theta}_D$ obtained using a given realization of random voxel selection decisions vector D and associated vector of selected voxels V_D is

$$\hat{\theta}_D = \mu_\theta + \mathbf{R}_{\theta V_D} \mathbf{R}_{V_D V_D}^{-1} (V_D - \mu_{V_D}). \quad (18)$$

where we denote $\mu_{V_D} = \mathbb{E}\{V_D|D\}$, $\mathbf{R}_{V_D V_D} = \text{cov}(V_D, V_D|D)$, $\mathbf{R}_{\theta V_D} = \text{cov}(\theta, V_D|D)$. The estimation error for $\hat{\theta}_D$ can be calculated using the law of total variance:

$$\text{var}(\theta - \hat{\theta}_D) = \mathbb{E}\{\text{var}(\theta - \hat{\theta}_D|D)\} + \text{var}(\mathbb{E}\{\theta - \hat{\theta}_D|D\}). \quad (19)$$

Since $\mathbb{E}\{V_D - \mu_{V_D}|D\} = 0$ implies $\mathbb{E}\{\theta - \hat{\theta}_D|D\} = 0$ we have $\text{var}(\theta - \hat{\theta}_D) = \mathbb{E}\{\text{var}(\theta - \hat{\theta}_D|D)\}$. For a fixed D we can calculate the conditional variance as

$$\text{var}(\theta - \hat{\theta}_D|D) = \mathbf{R}_{\theta\theta} - \mathbf{R}_{\theta V_D} \mathbf{R}_{V_D V_D}^{-1} \mathbf{R}_{\theta V_D}^T. \quad (20)$$

Under the assumption \mathcal{A}_3 implying that the covariance matrix $\mathbf{R}_{V_D V_D}$ is diagonal for any configuration of D the above expression can be simplified as follows:

$$\text{var}(\theta - \hat{\theta}_D|D) = \mathbf{R}_{\theta\theta} - \sum_{i=1}^N d_i \mathbf{R}_{\theta V_i} \mathbf{R}_{V_i V_i}^{-1} \mathbf{R}_{\theta V_i}^T. \quad (21)$$

Finally, upon evaluating the expectation with respect to D we obtain expression (6).

APPENDIX B

SOLUTION OF OPTIMIZATION PROBLEM (7)

After rearranging terms in (8) we can write:

$$\begin{aligned} \mathcal{J}(\mathcal{P}) &= - \sum_{i=1}^N p_i (U_i + \gamma_i - \eta_i + \lambda C) - \sum_{i=1}^N P_h \eta_i + \lambda C_{\text{ave}} \\ &= -\mathcal{P}^T (\mathbf{U} + \boldsymbol{\gamma} - \boldsymbol{\eta} + \lambda \mathbf{C}) - \sum_{i=1}^N P_h \eta_i + \lambda C_{\text{ave}} \\ &= -\|\mathcal{P}\|_2 \|\mathbf{U} + \boldsymbol{\gamma} - \boldsymbol{\eta} + \lambda \mathbf{C}\|_2 \cos \beta - \sum_{i=1}^N P_h \eta_i + \lambda C_{\text{ave}}. \end{aligned} \quad (22)$$

Here in the last two expressions we use vector notation $\mathbf{U} = [U_1, U_2, \dots, U_N]^T$, $\boldsymbol{\gamma} = [\gamma_1, \gamma_2, \dots, \gamma_N]^T$, $\boldsymbol{\eta} = [\eta_1, \eta_2, \dots, \eta_N]^T$, $\mathbf{C} = [C, C, \dots, C]^T$ and β is the angle between vectors \mathcal{P} and $\mathbf{U} + \boldsymbol{\gamma} - \boldsymbol{\eta} + \lambda \mathbf{C}$.

The Lagrange relaxation method [22] allows us to simplify the original minimization problem by first solving the unconstrained problem $\mathcal{P}^* = \min_{\mathcal{P}} \mathcal{J}(\mathcal{P})$ while treating the Lagrange multipliers as constants and then identifying the values of the Lagrange multipliers via original constraints.

In the first step of Lagrange relaxation, the last expression is minimized when $\cos \beta = 1$, i.e. vectors \mathcal{P} and $\mathbf{U} + \boldsymbol{\gamma} - \boldsymbol{\eta} + \lambda \mathbf{C}$ are collinear and proportional: $\mathcal{P}^* = A(\mathbf{U} + \boldsymbol{\gamma} - \boldsymbol{\eta} + \lambda \mathbf{C})$ with proportionality constant $A > 0$.

This can be proven easily by considering some arbitrary \mathcal{P}^* with $\|\mathcal{P}^*\|_2 = c$, $0 < c < \infty$ for which, necessarily, $\cos \beta < 1$. A better solution, \mathcal{P}^\sharp , can always be found by picking $\mathcal{P}^\sharp = \frac{c}{\|\mathbf{U} + \boldsymbol{\gamma} - \boldsymbol{\eta} + \lambda \mathbf{C}\|_2} (\mathbf{U} + \boldsymbol{\gamma} - \boldsymbol{\eta} + \lambda \mathbf{C})$. Inserting \mathcal{P}^\sharp into the expression for \mathcal{J} demonstrates that $\mathcal{J}(\mathcal{P}^\sharp) < \mathcal{J}(\mathcal{P}^*)$.

We can thus replace the original problem by introducing the positive auxiliary variable A and reformulating the problem as follows:

$$\begin{aligned} \mathcal{P}^* &= \min_{\mathcal{P}, A} \mathcal{J}(\mathcal{P}) \\ \text{such that} \quad p_i &\geq 0, \quad p_i \leq P_h, \quad \sum_{i=1}^N p_i C = C_{\text{ave}}, \\ p_i &= A(U_i + \gamma_i - \eta_i + \lambda C). \end{aligned} \quad (23)$$

In the second step of Lagrange relaxation we identify the Lagrange multipliers utilizing the constraint conditions derived from the original constraints:

$$\gamma_i \leq 0, p_i \geq 0, \gamma_i p_i = 0 \quad (24)$$

$$\eta_i \leq 0, p_i \leq P_h, \eta_i (P_h - p_i) = 0 \quad (25)$$

$$\mathcal{P}^T \mathbf{C} = C_{\text{ave}} \quad (26)$$

First we deduce the values of γ and η . If $0 \leq A(U_i + \lambda C) \leq P_h$, constraints are satisfied and thus $\gamma_i = 0$, $\eta_i = 0$. If $A(U_i + \lambda C) < 0$, constraint $p_i \leq P_h$ is satisfied and $\eta_i = 0$, but $\gamma_i \neq 0$. To satisfy condition $\gamma_i p_i = 0$, we set $\gamma_i = -(U_i + \lambda C)$ resulting in $p_i = 0$. If $A(U_i + \lambda C) > P_h$, constraint $p_i \geq 0$ is satisfied and $\gamma_i = 0$. To satisfy condition $\eta_i (P_h - p_i) = 0$, we set $\eta_i = P_h A^{-1} - (U_i + \lambda C)$ resulting in $p_i = P_h$.

The resulting solution can thus be represented as

$$p_i^* = \begin{cases} A(U_i + \lambda^* C) & \text{if } 0 \leq A(U_i + \lambda^* C) \leq P_h \\ P_h & \text{if } A(U_i + \lambda^* C) > P_h \\ 0 & \text{otherwise} \end{cases} \quad (27)$$

Here the value λ^* of the remaining Lagrange multiplier can be found by solving the constraint equation $\mathcal{P}^{*T} \mathbf{C} = C_{\text{ave}}$ for λ^* . The auxiliary variable A should be determined by the minimization problem (23) and we will later show how this can be accomplished by analyzing the properties of the cost function.

APPENDIX C

PROOF OF PROPOSITION 1

Proof: We will now show that $\mathcal{J}(\mathcal{P}^*)$, $\forall A \geq 0$ is a monotonically non-increasing function of A by showing that $\frac{\partial}{\partial A} \mathcal{J}(\mathcal{P}^*) \leq 0$, $\forall A \geq 0$ since the latter implies the former by the mean value theorem. We will assume, without loss of

generality, that $C = 1/N$, which simply means that the processing cost is the same for any voxel and the total cost when all voxels are selected is normalized to unity. To facilitate the derivations we will use the right continuous Heaviside function $H(\cdot)$ defined as $H(x) = 1, x \geq 0$ and $H(x) = 0, x < 0$ and the following notation $\mathcal{A}_i = A(U_i + \lambda^*/N)$. With this notation, solution (27) can be written concisely as follows:

$$p_i^* = H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)\mathcal{A}_i + P_h(1 - H(P_h - \mathcal{A}_i)) \quad (28)$$

Since \mathcal{P}^* is the solution of (7) and it thus satisfies all the constraints we note that $\mathcal{J}(\mathcal{P}^*) = -\sum_{i=1}^N p_i^* U_i$. Taking these remarks into account we can rewrite the objective function as:

$$\mathcal{J}(\mathcal{P}^*) = -\sum_{i=1}^N U_i [H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)\mathcal{A}_i + P_h(1 - H(P_h - \mathcal{A}_i))] \quad (29)$$

and the constraint $\mathcal{P}^{*T} \mathbf{C} = C_{\text{ave}}$:

$$\frac{1}{N} \sum_{i=1}^N [H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)\mathcal{A}_i + P_h(1 - H(P_h - \mathcal{A}_i))] = C_{\text{ave}}. \quad (30)$$

The differentiation of $\mathcal{J}(\mathcal{P}^*)$ results in

$$\begin{aligned} \frac{\partial}{\partial A} \mathcal{J}(\mathcal{P}^*) &= -\sum_{i=1}^N U_i \frac{\partial \mathcal{A}_i}{\partial A} [P_h \delta(\mathcal{A}_i - P_h) \\ &\quad + \delta(\mathcal{A}_i)H(P_h - \mathcal{A}_i)\mathcal{A}_i - \delta(\mathcal{A}_i - P_h)H(\mathcal{A}_i)\mathcal{A}_i \\ &\quad + H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)]. \end{aligned} \quad (31)$$

Consider $P_h \delta(\mathcal{A}_i - P_h) + \delta(\mathcal{A}_i)H(P_h - \mathcal{A}_i)\mathcal{A}_i - \delta(\mathcal{A}_i - P_h)H(\mathcal{A}_i)\mathcal{A}_i$. It is straightforward to see that this expression is 0 whenever $\mathcal{A}_i < 0$, $0 < \mathcal{A}_i < P_h$ and $\mathcal{A}_i > P_h$. Consider the case $\mathcal{A}_i \rightarrow P_h$:

$$\begin{aligned} \lim_{\mathcal{A}_i \rightarrow P_h} (P_h \delta(\mathcal{A}_i - P_h) + \delta(\mathcal{A}_i)H(P_h - \mathcal{A}_i)\mathcal{A}_i \\ - \delta(\mathcal{A}_i - P_h)H(\mathcal{A}_i)\mathcal{A}_i) \\ = \lim_{\mathcal{A}_i \rightarrow P_h} P_h (\delta(\mathcal{A}_i - P_h) - \delta(\mathcal{A}_i - P_h)) = 0. \end{aligned} \quad (32)$$

Consider the case $\mathcal{A}_i \rightarrow 0$:

$$\begin{aligned} \lim_{\mathcal{A}_i \rightarrow 0} (P_h \delta(\mathcal{A}_i - P_h) + \delta(\mathcal{A}_i)H(P_h - \mathcal{A}_i)\mathcal{A}_i \\ - \delta(\mathcal{A}_i - P_h)H(\mathcal{A}_i)\mathcal{A}_i) = \lim_{\mathcal{A}_i \rightarrow 0} \delta(\mathcal{A}_i)\mathcal{A}_i = 0. \end{aligned} \quad (33)$$

The last equality is one of the fundamental properties of the Delta function. Thus we see that for any value of \mathcal{A}_i the expression just analyzed is zero and the derivative $\frac{\partial}{\partial A} \mathcal{J}(\mathcal{P}^*)$ is simply:

$$\frac{\partial}{\partial A} \mathcal{J}(\mathcal{P}^*) = -\sum_{i=1}^N U_i \frac{\partial \mathcal{A}_i}{\partial A} H(\mathcal{A}_i)H(P_h - \mathcal{A}_i). \quad (34)$$

Taking into account the fact that λ^* is a function of A and calculating the derivative

$$\frac{\partial \mathcal{A}_i}{\partial A} = \frac{\partial}{\partial A} A(U_i + \lambda^*/N) = [U_i + \lambda^*/N] + \frac{\partial \lambda^*}{\partial A} \frac{A}{N}, \quad (35)$$

we can rewrite the derivative $\frac{\partial}{\partial A} \mathcal{J}(\mathcal{P}^*)$ as follows:

$$\begin{aligned} \frac{\partial}{\partial A} \mathcal{J}(\mathcal{P}^*) &= -\sum_{i=1}^N H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)U_i [(U_i + \lambda^*/N) \\ &\quad + \frac{\partial \lambda^*}{\partial A} A/N]. \end{aligned} \quad (36)$$

To derive $\frac{\partial \lambda^*}{\partial A}$ we use the implicit differentiation of the constraint equation (30) which, upon differentiating its both sides and simplifying it using the same logic that was used to simplify (31), becomes:

$$\sum_{i=1}^N \frac{H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)}{N} \left[\left(U_i + \frac{\lambda^*}{N} \right) + \frac{\partial \lambda^*}{\partial A} \frac{A}{N} \right] = 0. \quad (37)$$

After grouping and moving the terms we obtain:

$$\begin{aligned} \frac{\partial \lambda^*}{\partial A} &= -\frac{N \sum_{i=1}^N H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)(U_i + \lambda^*/N)}{A \sum_{i=1}^N H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)} \\ &= -\frac{N \sum_{i=1}^N H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)U_i}{A \sum_{i=1}^N H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)} - \lambda^*/A. \end{aligned} \quad (38)$$

Inserting this into (36) produces:

$$\begin{aligned} \frac{\partial \mathcal{J}}{\partial A}(\mathcal{P}^*) &= -\sum_{i=1}^N H(\mathcal{A}_i)H(P_h - \mathcal{A}_i) \\ &\quad \times \left[U_i^2 - \frac{\sum_{j=1}^N H(\mathcal{A}_j)H(P_h - \mathcal{A}_j)U_i U_j}{\sum_{j=1}^N H(\mathcal{A}_j)H(P_h - \mathcal{A}_j)} \right] \\ &= -\sum_{i=1}^N \frac{H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)}{\sum_{k=1}^N H(\mathcal{A}_k)H(P_h - \mathcal{A}_k)} \\ &\quad \times \left[\sum_{j=1}^N U_i^2 H(\mathcal{A}_j)H(P_h - \mathcal{A}_j) \right. \\ &\quad \left. - \sum_{j=1}^N H(\mathcal{A}_j)H(P_h - \mathcal{A}_j)U_i U_j \right]. \end{aligned} \quad (39)$$

Taking into account that

$$\begin{aligned} \sum_{i=1}^N \sum_{j=1}^N H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)H(\mathcal{A}_j)H(P_h - \mathcal{A}_j)U_i^2 \\ = \sum_{i=1}^N \sum_{j=1}^N H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)H(\mathcal{A}_j)H(P_h - \mathcal{A}_j)U_j^2, \end{aligned} \quad (40)$$

we can finally write:

$$\begin{aligned} \frac{\partial}{\partial A} \mathcal{J}(\mathcal{P}^*) &= -\sum_{i=1}^N \sum_{j=1}^N \frac{H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)H(\mathcal{A}_j)H(P_h - \mathcal{A}_j)}{2 \sum_{k=1}^N H(\mathcal{A}_k)H(P_h - \mathcal{A}_k)} \\ &\quad \times [U_i^2 + U_j^2 - 2U_i U_j] \\ &= -\sum_{i=1}^N \sum_{j=1}^N \frac{H(\mathcal{A}_i)H(P_h - \mathcal{A}_i)H(\mathcal{A}_j)H(P_h - \mathcal{A}_j)}{2 \sum_{k=1}^N H(\mathcal{A}_k)H(P_h - \mathcal{A}_k)} \\ &\quad \times (U_i - U_j)^2. \end{aligned} \quad (41)$$

The last expression is clearly non-positive $\forall A$. \blacksquare

APPENDIX D

PROOF OF PROPOSITION 2

Proof: Using the definition of $\varphi(\lambda^*)$ in the left hand side of (30), differentiating it with respect to λ^* and using simplification methodology applied in Proposition 1 we arrive at:

$$\frac{\partial}{\partial \lambda^*} \varphi(\lambda^*) = C \sum_{i=1}^N \frac{\partial \mathcal{A}_i}{\partial \lambda^*} H(\mathcal{A}_i) H(P_h - \mathcal{A}_i). \quad (42)$$

Taking into account that

$$\frac{\partial \mathcal{A}_i}{\partial \lambda^*} = \frac{\partial}{\partial \lambda^*} A(U_i + \lambda^* C) = AC, \quad (43)$$

It is clear that $\frac{\partial}{\partial \lambda^*} \varphi(\lambda^*) \geq 0, \forall \lambda^*$ and the claim of the proposition follows. ■

ACKNOWLEDGMENT

The authors would like to thank Colm Elliott, Dante De Nigris and Dr. D.L. Collins for their insightful comments and for helpful discussions.

REFERENCES

- [1] W. Wein, S. Brunke, A. Khamene, M. R. Callstrom, and N. Navab, "Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention," *Medical Image Analysis*, vol. 12, no. 5, pp. 577–585, Oct. 2008. 1
- [2] W. Zhang, J. A. Noble, and J. M. Brady, "Adaptive non-rigid registration of real time 3D ultrasound to cardiovascular MR images," in *Proc. 20th Int. Conf. Information processing in medical imaging*, ser. IPMI'07. Berlin, Heidelberg: Springer-Verlag, 2007, pp. 50–61. 1
- [3] Y. Sakai, Y. Matsuyama, H. Yoshihara, H. Nakamura, S. Nakashima, and N. Ishiguro, "Simultaneous registration with CT-fluoro matching for spinal navigation surgery. A case report," *Nagoya J Med Sci*, vol. 68, no. 1-2, pp. 45–52, 2006. 1, 10
- [4] G. J. Webster, J. E. Kilgallon, K. F. Ho, C. G. Rowbottom, N. J. Slevin, and R. I. Mackay, "A novel imaging technique for fusion of high-quality immobilised MR images of the head and neck with CT scans for radiotherapy target delineation," *The British Journal of Radiology*, vol. 82, no. 978, pp. 497–503, 2009. 1
- [5] M. M. Chakravarty, A. F. Sadikot, J. Germann, P. Hellier, G. Bertrand, and D. L. Collins, "Comparison of piece-wise linear, linear, and non-linear atlas-to-patient warping techniques: Analysis of the labeling of subcortical nuclei for functional neurosurgical applications," *Human Brain Mapping*, vol. 30, no. 11, pp. 3574–3595, 2009. 1
- [6] S. Ji, Z. Wu, A. Hartov, D. W. Roberts, and K. D. Paulsen, "Mutual-information-based image to patient re-registration using intraoperative ultrasound in image-guided neurosurgery," *Medical Physics*, vol. 35, no. 10, pp. 4612–4624, Oct. 2008. 1
- [7] A. Roche, X. Pennec, G. Malandain, and N. Ayache, "Rigid registration of 3-D ultrasound with MR images: A new approach combining intensity and gradient information," *IEEE Transactions on Medical Imaging*, vol. 20, no. 10, pp. 1038–49, 2001. 1, 10
- [8] D. D. Nigris, D. Collins, and T. Arbel, "Multi-modal image registration based on gradient orientations of minimal uncertainty," *IEEE Transactions on Medical Imaging*, vol. 31, no. 12, pp. 2343–54, Dec. 2012. 1
- [9] D. Mattes, D. Haynor, and et al., "PET-CT image registration in the chest using free-form deformations," *IEEE TMI*, vol. 22, no. 1, pp. 120–128, 2003. 1, 7
- [10] P. Viola and W. M. Wells, "Alignment by maximization of mutual information," *Int. J. Computer Vision*, vol. 24, no. 2, pp. 137–154, 1997. 1, 2, 3, 5, 7
- [11] D. I. Barnea and H. F. Silverman, "A class of algorithms for fast digital image registration," *IEEE Trans. Computers*, vol. C-21, no. 2, pp. 179–186, Feb. 1972. 1
- [12] R. N. Nagel and A. Rosenfeld, "Ordered search techniques in template matching," *Proceedings of the IEEE*, vol. 60, no. 2, pp. 242–244, Feb. 1972. 1
- [13] S. J. Reeves and R. Hezar, "Selection of observations in magnetic resonance spectroscopic imaging," in *Proc. ICIP*, vol. 1, Oct. 1995, pp. 641–644. 1, 2, 5
- [14] F. Dellaert and R. Collins, "Fast image-based tracking by selective pixel integration," in *Proc. ICCV-99*, Kerkira, Greece, Sept. 1999. 1, 2
- [15] R. Brooks and T. Arbel, "The importance of scale when selecting pixels for image registration," in *Proc. CRV*, Paris, France, May 2007, pp. 235–242. 1
- [16] M. R. Sabuncu and P. J. Ramadge, "Gradient based nonuniform subsampling for information-theoretic alignment methods," in *Proc. IEEE IEMBS '04*, vol. 1, sept. 2004, pp. 1683–1686. 2, 7
- [17] B. N. Oreshkin and T. Arbel, "Optimization over random and gradient probabilistic pixel sampling for fast, robust multi-resolution image registration," in *WBIR*, ser. Lecture Notes in Computer Science, B. M. Dawant, G. E. Christensen, J. M. Fitzpatrick, and D. Rueckert, Eds., vol. 7359. Springer, 2012, pp. 150–159. 2, 7, 8
- [18] M. Freiman, M. Werman, and L. Joskowicz, "A curvelet-based patient-specific prior for accurate multi-modal brain image rigid registration," *Med. Imag. Anal.*, vol. 15, no. 1, pp. 125–132, Feb. 2010. 2
- [19] J. Fitzpatrick, J. West, and C. Maurer, "Predicting error in rigid-body point-based registration," *IEEE TMI*, vol. 17, no. 5, pp. 694–702, Oct. 1998. 2, 7
- [20] C. Studholme, D. L. G. Hill, and D. J. Hawkes, "An overlap invariant entropy measure of 3D medical image alignment," *Pattern Recognition*, vol. 32, no. 1, pp. 71–86, 1999. 2, 7
- [21] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. Dover Books on Engineering, 2005. 4
- [22] H. Everett III, "Generalized Lagrange multiplier method for solving problems of optimum allocation of resources," *Operations Research*, vol. 11, pp. 399–417, 1963. 4, 12
- [23] M. Taron, N. Paragios, and M.-P. Jolly, "Registration with uncertainties and statistical modeling of shapes with variable metric kernels," *IEEE PAMI*, vol. 31, no. 1, pp. 99–113, 2009. 6
- [24] R. K.-S. Kwan, A. C. Evans, and G. B. Pike, "MRI simulation-based evaluation of image-processing and classification methods," *IEEE TMI*, vol. 18, no. 11, pp. 1085–1097, Nov. 1999. 6
- [25] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proc. IEEE Int. Conf. Neural Networks*, vol. 4, nov 1995, pp. 1942–1948. 7
- [26] X. Lu, S. Zhang, H. Su, and Y. Chen, "Mutual information-based multimodal image registration using a novel joint histogram estimation," *Computerized Medical Imaging and Graphics*, vol. 32, no. 3, pp. 202–209, 2008. 7
- [27] R. Brooks, "Efficient and reliable methods for direct parameterized image registration," Ph.D. dissertation, McGill University, 2008. 7
- [28] T. Arbel, X. Morandi, R. M. Comeau, and D. L. Collins, "Automatic non-linear MRI-ultrasound registration for the correction of intra-operative brain deformations," in *Proceedings of the 4th International Conference on Medical Image Computing and Computer-Assisted Intervention*, ser. MICCAI '01. London, UK: Springer-Verlag, 2001, pp. 913–922. 10
- [29] N. Hansen, S. Niederberger, L. Guzzella, and P. Koumoutsakos, "A method for handling uncertainty in evolutionary optimization with an application to feedback control of combustion," *IEEE Transactions on Evolutionary Computation*, vol. 13, no. 1, pp. 180–197, 2009. 11
- [30] B. T. Polyak and A. B. Juditsky, "Acceleration of stochastic approximation by averaging," *Siam Journal on Control and Optimization*, vol. 30, no. 4, pp. 838–855, 1992. 11
- [31] H. Kahn and A. Marshall, "Methods of reducing sample size in Monte Carlo computations," *Journal of the Operations Research Society of America*, vol. 1, no. 5, pp. 263–271, 1953. 11