



HAL
open science

Extended Modality Propagation: Image Synthesis of Pathological Cases

Nicolas Cordier, Hervé Delingette, Matthieu Lê, Nicholas Ayache

► **To cite this version:**

Nicolas Cordier, Hervé Delingette, Matthieu Lê, Nicholas Ayache. Extended Modality Propagation: Image Synthesis of Pathological Cases. *IEEE Transactions on Medical Imaging*, 2016, PP (99), 10.1109/TMI.2016.2589760 . hal-01343233

HAL Id: hal-01343233

<https://inria.hal.science/hal-01343233>

Submitted on 7 Jul 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Extended Modality Propagation: Image Synthesis of Pathological Cases

Nicolas Cordier, Hervé Delingette, Matthieu Lê, Nicholas Ayache

Abstract—This paper describes a novel generative model for the synthesis of multi-modal medical images of pathological cases based on a single label map. Our model builds upon *i*) a generative model commonly used for label fusion and multi-atlas patch-based segmentation of healthy anatomical structures, *ii*) the Modality Propagation iterative strategy used for a spatially-coherent synthesis of subject-specific scans of desired image modalities. The expression **Extended Modality Propagation** is coined to refer to the extension of Modality Propagation to the synthesis of images of pathological cases. Moreover, image synthesis uncertainty is estimated. An application to Magnetic Resonance Imaging synthesis of glioma-bearing brains is *i*) validated on the training dataset of a Multimodal Brain Tumor Image Segmentation challenge, *ii*) compared to the state-of-the-art in glioma image synthesis, and *iii*) illustrated using the output of two different tumor growth models. Such a generative model allows the generation of a large dataset of synthetic cases, which could prove useful for the training, validation, or benchmarking of image processing algorithms.

Index Terms—patch-based, multi-atlas, glioma, generative model, medical image simulation, modality synthesis.

I. INTRODUCTION

A. Motivation

The availability of public datasets [1]–[3] of annotated medical images is a key factor in the development of medical image computing. For instance, the organization of the Multimodal Brain Tumor Image Segmentation (BraTS) benchmark challenge [3] has led to an important and fruitful research activity in glioma segmentation. However, the creation of a benchmark dataset is costly for obvious reasons: *i*) a large number of cases is required in the training and testing datasets to capture the variability of structures or pathologies, *ii*) images should be annotated by experts, which requires a complex and time-consuming manual work, and can still lead to debatable results (e.g. inter-rater variability in the range 74–85% for glioblastoma [3], datasets can contain incorrect segmentations [4]), *iii*) the distribution of medical images leads to ethical concerns, since the identification of patients may be possible despite anonymization steps such as skull-stripping [4], *iv*) and finally, ensuring the continued quality of a very large dataset of anonymized images is complex, due to the presence of longitudinal data and duplicates [5]. The development of image synthesis could allow *i*) the augmentation of a dataset by including new realistic synthetic cases, and *ii*) the




Fig. 1: The generative model aims at synthesizing subject-specific MRI, provided a label map of a pathological brain.

creation of large, annotated, unbiased, anonymized, and easy-to-maintain datasets of synthetic medical images of virtual patients (namely atlases completely generated by a computational biophysical model). The availability of a large dataset is especially important for pathologies such as glioblastoma which exhibit a high variability of shape and appearance. For reference, the 2013 BraTS challenge only consisted of 20 real high-grade cases for training, and 10 real high-grade cases for benchmarking [3]). The 2014 BraTS challenge included about 250 additional cases obtained from The Cancer Imaging Archive [6], however other problems arose: the ground truth was unavailable and had to be obtained through a consensus of automatic glioma segmentation algorithms, which resulted in a ground truth of lower reliability [7].

B. Related work

The problem which is tackled in this paper is the synthesis of multi-modal medical images of pathological cases, based on a single label map, as illustrated in Figure 1 for brain tumors imaged with Magnetic Resonance Imaging (MRI). In the following, the focus is on related work regarding MRI synthesis of healthy and tumor-bearing brains.

1) *MRI synthesis of healthy brains*: Original attempts at generating synthetic MRI of healthy brains relied on a numerical simulation of MR acquisition physics. Given MR scanner parameters such as echo time and relaxation time, a discrete-event simulation model [8]–[10] describes the dynamics of magnetization vectors, at each spatial position, according to Bloch equations [11]. Such numerical simulations are computationally expensive, except for specific cases for which there exists a closed-form solution to Bloch equations, such as for spin-echo or gradient-echo MR sequences [8], [12], [13].

To describe the imaged object, two strategies are possible. The first strategy consists in describing the geometry of the imaged tissues with tissue-specific templates, and providing the tissue-specific biophysical properties (spin density and

relaxation times). This allows the synthesis of MRI acquired on virtual geometrical templates. However, it requires the knowledge of tissue-specific biophysical properties [8], which are poorly referenced in the literature for glioblastoma compartments. The second strategy does not rely on the definition of tissues: biophysical properties are specified voxel-wise, after an estimation from several MR scans obtained in a short time-frame with a very strict acquisition protocol (quantitative MRI or relaxometry), or by optimization methods [12], [13]. This strategy does not allow to generate synthetic images on virtual patients. Moreover, it cannot be applied to tumor scans when a proton density map is not available.

Recently, an iterative patch-based heuristic [14] was used for Modality Propagation, i.e. the synthesis of a realistic subject-specific scan of one modality given a scan acquired with another modality. Roots of the Modality Propagation framework include image reconstruction works, e.g. PET image reconstruction incorporating MR data as prior information for anatomical boundary [15] and example-based approaches using non-local patch similarity constraints, known as brain image hallucination [16]–[19]. The core of the algorithm consists in modeling the covariation of local intensity patterns across modalities, in a fashion inspired by multi-atlas patch-based segmentation of healthy brain structures [20], [21]. Successful applications of Modality Propagation or related methods include synthesis of a patient-specific attenuation map for hybrid MR-PET [22], more accurate registration [23], super-resolution [24], and outlier detection [14], [24], [25]. However, Modality Propagation shows two limitations for the synthesis of MRI for virtual pathological cases: *i*) the local search window framework is adopted, which does not take into account the variability of lesion position, and might restrict the method to the image synthesis of healthy cases, *ii*) the input is a real image, which is assumed not to be available for virtual cases.

2) *MRI synthesis of tumor-bearing brains*: Previous works regarding MRI synthesis of tumor-bearing brains [26]–[28] build upon tissue-specific templates, as in the framework introduced for healthy brains [8]. Typically, a tumor seed is artificially positioned in a healthy brain atlas, then a tumor growth model simulates the tumor extension over time and its mass effect, i.e. the displacement of neighboring healthy tissues, which results in templates of healthy tissues and tumor compartments. However, in contrast to [8], tissue-specific biophysical properties are not considered in favor of direct specifications of tissue-specific average MR intensities [26], [27]. In [28], textures are learnt for each healthy tissue and tumor compartment, which results in a set of 3D texture images. The only difference between different synthetic MRI lies in tissue-specific templates: tissue-specific average MR intensities and 3D texture images are therefore always identical.

This last approach is applicable to a wide variety of cases, but also bears some limitations: *i*) the complex distribution of image intensities for tumor compartments is summarized by its expectation, which is oblivious of multi-modal intensity distributions, *ii*) the inter-patient MR normalization procedure is not specified, which makes it difficult to standardize real MRI so that they look like synthetic MRI, typically for the

training of machine learning algorithms [29], *iii*) simulated images do not show the variability of intensity of realistic MR scans, and the addition of a very high Gaussian noise only limits this effect.

The simulator of synthetic pathological MRI¹ described in [28] has been used in a number of research articles mostly for prototyping and validation, in the context of glioma segmentation [3], [29], [30], outlier detection algorithm [31], registration of a healthy brain atlas to a tumor-bearing patient image [32], and construction of a brain atlas [33]. Other applications include the training of machine-learning algorithms for glioma segmentation. [29] trained random forest on a large dataset consisting of 740 synthetic cases and showed good segmentation results on a few real cases for testing. [31] performed a thorough validation of an outlier detection algorithm, based on 100 MR scans synthesized with different tumor volumes to test the robustness of the algorithm to the amount of outliers. Image synthesis could also be useful in the context of tumor growth modeling: in [34], PET images are synthesized for tumor-bearing brains, so as to allow clinically relevant interpretations of tumor growth model outputs. Although the applications of MRI synthesis of pathological cases are numerous, the competitors in the BraTS benchmark challenge tend to completely ignore the provided synthetic cases [3], which is likely due to their lack of realism.

In this paper, we address the image synthesis of pathological cases by using an iterative multi-atlas patch-based algorithm, inspired by recent successful algorithms in MRI synthesis of healthy brains. The expression Extended Modality Propagation (EMP) is coined for two reasons: *i*) to refer to the extension of Modality Propagation for the synthesis of images of pathological cases, *ii*) in contrast to Modality Propagation, EMP can deal with label maps in addition to image intensity. The synthesis process only requires a single label map, or the output of any tumor growth model in terms of the usual tumor compartments. Realistic synthetic images are obtained in the specific application of our algorithm to MRI synthesis of tumor-bearing brains. Our contributions include a novel generative model, a heuristic iterative algorithm to solve for the posterior distribution of multi-channel MR intensities, the estimation of image synthesis uncertainty, the analysis of Modality Propagation iterative feature augmentation, and the public availability of a large dataset of annotated and realistic MRI exhibiting gliomas. In the remainder, the generative model is described and solved for the maximum a posteriori and for uncertainty estimation (Section II). Then synthesized MRI are compared to real MRI from the training dataset of BraTS benchmark challenge, and to the state-of-the-art in glioma image synthesis, and illustrated using the output of two different tumor growth models (Section III). Finally, results and perspectives are discussed (Sections IV and V).

II. EXTENDED MODALITY PROPAGATION

In the following, a label map \mathbf{L} describing the anatomy of a tumor-bearing brain is assumed to be known. The goal of

¹Freely available online at <http://www.nitrc.org/projects/tumorsim>

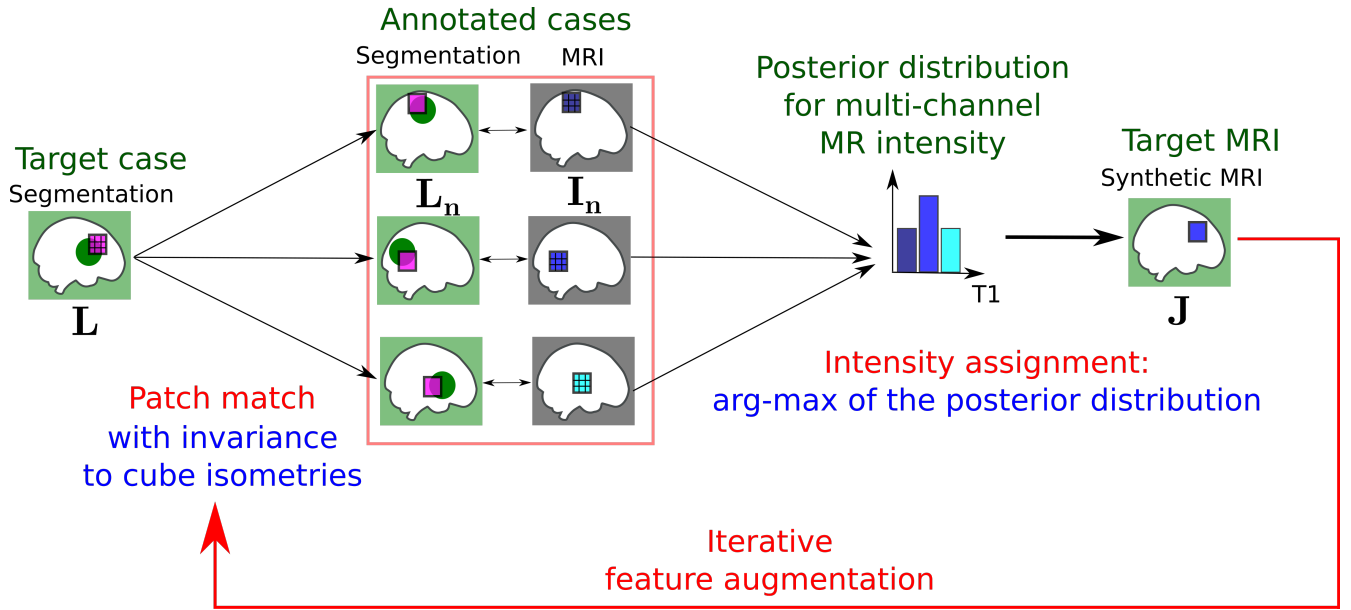


Fig. 2: Pipeline for Extended Modality Propagation. Green disks mark known tumor locations. A fundamental idea is that central voxels of similar label patches should have similar multi-channel image intensity. During the first iteration, multi-channel intensity J is assigned for every test label patch based on comparisons with a database of training label patches. During subsequent iterations, multi-channel intensity J is updated based on comparisons between augmented label patches, which are the concatenation of label patches and multi-channel image patches synthesized during the previous iteration (feedback).

Extended Modality Propagation (EMP) is to synthesize medical images corresponding to the same anatomy. To achieve this goal, EMP relies on a set of N training cases $\{\mathbf{L}_n, \mathbf{I}_n\}_n$ where \mathbf{L}_n is a label map and \mathbf{I}_n denotes multi-channel MRI (T1, T1C, T2, FLAIR). In practice, the set of labels comprises $\mathcal{L} = 7$ classes: cerebrospinal fluid, grey matter, white matter, necrotic core, edema, non-enhancing tumor core, and enhancing core.

The EMP probabilistic generative model is first introduced to describe the synthesis of images. Second, an approximation of the marginalization step is introduced to estimate the maximum a posteriori and the image synthesis model uncertainty. Finally, more consistent patch matching is obtained by augmenting label patches: *i*) multi-scale patches allow to consider larger patches in a computationally efficient manner, *ii*) a patch match invariant to cube isometries allows to consider larger patches without being impacted by a smaller sample size, and *iii*) as for Modality Propagation, an iterative patch match procedure results in smoother and more consistent synthetic images. The EMP pipeline is detailed in Figure 2.

A. Definition of a probabilistic generative model

1) *Graphical model*: The proposed generative model for Extended Modality Propagation builds upon the Bayesian modeling of local weighted voting for multi-atlas segmentation [35], which was then adapted for patch-based segmentation [36], [37]. Conventionally, a patch $S(I, x)$ is a cube of edge length $2r + 1$, centered at spatial position $x \in \mathbb{R}^3$, consisting of image intensity values taken from an image I .

Our model is based on the 5 following assumptions. *i*) Every patch $S(\mathbf{L}, x)$ originates from a single patch $S(\mathbf{L}_n, y)$ from

the training database. This means that the target segmentation L is the result of sampling patches from atlases $(\mathbf{L}_n, \mathbf{I}_n)$ at different positions y in the reference space Ω . The membership index \mathbf{M} encodes, for every position $x \in \Omega$, both the atlas index n and the position y to sample from: $\mathbf{M}(x) = (n, y)$. *ii*) There is a uniform prior on \mathbf{M} : no training case n or patch location y is preferred a priori. This assumption is in sharp contrast with a common assumption of local search window in the patch-based framework. *iii*) The only constraint on \mathbf{M} is that the source patch for a target patch $S(\mathbf{L}, x)$ has the same central label $\mathbf{L}_n(y) = \mathbf{L}(x)$. *iv*) If \mathbf{M} were known, the target image J would be given by intensity, perturbed with white noise, $\mathbf{I}_n(y)$ at central voxels of sampled patches. *v*) Given $\mathbf{M}(x)$, probabilities of observing intensity $J(x)$ and label $\mathbf{L}(x)$ are conditionally independent. This assumption introduces a symmetry between segmentation and synthesis problems. Moreover, this assumption is reasonable: if $\mathbf{M}(x)$ were known, then the information regarding label would be completely defined by \mathbf{L}_n only, and the information regarding intensity would be completely defined by \mathbf{I}_n only.

The corresponding graphical model is shown in Figure 3. This model mirrors the graphical model described in [38] for multi-atlas patch-based segmentation of pathological cases.

2) *Class-conditional image intensity likelihood*: Let $j \in \mathbb{R}^d$ denote multi-channel voxel-wise image intensity, with d being the number of MR channels. In our case, $d = 4$. Class-conditional image intensity likelihood $p_n(J(x) = j | \mathbf{I}_n, y, \Sigma_{\mathbf{L}_n(y)}, \mathbf{L}_n(y), \mathbf{L}(x))$ is the probability that the target intensity $J(x) = j$ is sampled from a multivariate Gaussian probability density function \mathcal{N} , with




Fig. 3: Graphical model for Extended Modality Propagation. A membership variable $\mathbf{M} : \Omega \rightarrow \{1, \dots, N\} \times \Omega$ is sampled at every position x in the MNI space Ω to encode the training case \mathbf{n} to sample from, and the spatial offset y : $\mathbf{M}(x) = (n, y)$. The observed Log-Odds map $\tilde{\mathbf{L}}(x)$ results from the sampling of a patch $S(\tilde{\mathbf{L}}, x)$ from a Gaussian distribution with $S(\tilde{\mathbf{L}}_{\mathbf{n}}, y)$ as mean and isotropic scaling $\sigma^2(x)$ as variance. The intensity \mathbf{J} at position x is sampled from a Gaussian distribution with the central value $\mathbf{I}_{\mathbf{n}}(y)$ as mean and $\Sigma_{\mathbf{L}(x)}$ as covariance matrix.

$\mathbf{I}_{\mathbf{n}}(y)$ as mean and $\Sigma_{\mathbf{L}_{\mathbf{n}}(y)} = \Sigma_{\mathbf{L}(x)}$ as covariance matrix:

$$p_n(\mathbf{J}(x) = j \mid \mathbf{I}_{\mathbf{n}}, y, \Sigma_{\mathbf{L}_{\mathbf{n}}(y)}, \mathbf{L}_{\mathbf{n}}(y), \mathbf{L}(x)) \\ \propto \delta_{\mathbf{L}_{\mathbf{n}}(y), \mathbf{L}(x)} \times \mathcal{N}(j \mid \mathbf{I}_{\mathbf{n}}(y), \Sigma_{\mathbf{L}_{\mathbf{n}}(y)})$$

where $\delta_{\cdot, \cdot}$ is Kronecker delta, which translates assumption *iii*. In this case, $\mathbf{L}_{\mathbf{n}}(y) = \mathbf{L}(x)$, and the notation for intensity likelihood is simplified to $p_n(\mathbf{J}(x) = j \mid \mathbf{I}_{\mathbf{n}}, y, \Sigma_{\mathbf{L}(x)})$.

3) *Log-Odds label maps*: The likelihood label model relies on continuous Log-Odds maps $\tilde{\mathbf{L}}$ and $\tilde{\mathbf{L}}_{\mathbf{n}}$ [35], i.e. probability values computed from categorical label maps \mathbf{L} and $\mathbf{L}_{\mathbf{n}}$. Indeed, during the first iteration of image synthesis, patch match only relies on a label map as input. It is crucial to choose the best representation for this single input. A categorical label map could be used, but this representation does not distinguish patches inside a tissue from those at its border. If categorical label maps were directly used, the matching of similar label configurations could be poor for label configurations not present in the training dataset. A better strategy, as recommended in [35], consists in adopting Log-Odds, so as to take into account the distance to borders between tissues and to be able to distinguish different patches among patches completely enclosed in a tissue. Log-Odds distances are defined as proportional to $\exp(\rho \tilde{D}^l(x))$, where $\rho > 0$ is the slope constant (in our experiments, $\rho = 1$ as in [35]), and $\tilde{D}^l(L, x)$ is the signed distance transform of label l for label map L at spatial position x , which is positive inside structures labelled as l .

4) *Augmented Log-Odds label maps*: During the first iteration $t = 1$ of the algorithm, the notation $\tilde{\mathbf{L}}$ corresponds to a multi-label Log-Odds map: a vector in $\mathbb{R}^{\mathcal{L}}$ consisting of Log-Odds label probability values is assigned for each spatial position x . Therefore, a patch of edge length r would result in a vector in $\mathbb{R}^{(2r+1)^3 \times \mathcal{L}}$. However, for subsequent iterations $t \geq 2$, the same notation $\tilde{\mathbf{L}}$ corresponds to augmented Log-Odds maps: for each spatial position x , the vector in $\mathbb{R}^{\mathcal{L}}$ of

Log-Odds label probability values is concatenated with the vector in \mathbb{R}^d of synthetic image intensity of the previous iteration. In this case, a patch of edge length r would result in a vector in $\mathbb{R}^{(2r+1)^3 \times (\mathcal{L}+d)}$.

5) *Label patch likelihood*: Label patch likelihood $p_n(S(\tilde{\mathbf{L}}, x) \mid \tilde{\mathbf{L}}_{\mathbf{n}}, y, \sigma^2(x))$ comprises the similarity between label patches. It is defined as the probability that the target Log-Odds patch $S(\tilde{\mathbf{L}}, x)$ is sampled from a multivariate Gaussian probability density function, with $S(\tilde{\mathbf{L}}_{\mathbf{n}}, y)$ as mean and isotropic scaling $\sigma^2(x)$ as variance:

$$p_n(S(\tilde{\mathbf{L}}, x) \mid \tilde{\mathbf{L}}_{\mathbf{n}}, y, \sigma^2(x)) \\ \propto \mathcal{N}(S(\tilde{\mathbf{L}}, x) \mid S(\tilde{\mathbf{L}}_{\mathbf{n}}, y), \sigma^2(x) \times \mathbb{1})$$

where $\mathbb{1}$ is the identity matrix.

6) *Symmetry*: The symmetry between the segmentation model described in [38] and our synthesis model is emphasized by the definitions of label and intensity likelihoods, as shown in Table I. Indeed, regarding likelihood for observed variables, image intensity likelihood for segmentation [38] and label patch likelihood for synthesis are Gaussian densities on patches, respectively image patches and Log-Odds patches, with an isotropic scaling parameter for variance, respectively $\sigma_n^2(x)$ and $\sigma^2(x)$. Moreover, regarding likelihood for hidden variables, label likelihood for segmentation [38] and class-conditional image intensity likelihood for synthesis are defined voxel-wise instead of patch-wise, and both involve a Kronecker delta.

B. Inference

Since \mathbf{M} is not observed, the inference has to be performed by marginalizing over \mathbf{M} [35]. The marginalization consists in a weighted vote involving patches at every position y in every atlas n , with weights proportional to both label patch and intensity likelihoods.

Variance parameters are first estimated; second, the problem of estimating \mathbf{J} is stated by marginalizing over \mathbf{M} ; third, a strategy to determine the arg-max of the posterior distribution is proposed; finally, a method to estimate the uncertainty of the synthesis procedure is mentioned.

1) *Estimation of the variance parameters*: There are two kinds of variance parameters: *i*) covariance matrices $\{\Sigma_l\}_{l \in \{1, \dots, \mathcal{L}\}}$ for class-conditional image intensity likelihoods, which are pre-computed, and *ii*) isotropic scaling $\sigma^2(x)$ for label patch likelihood, which is computed online.

For each label $l \in \{1, \dots, \mathcal{L}\}$, a covariance matrix Σ_l results from the fit, using Expectation-Maximization [39], of a multivariate Student distribution to average intensity of pure patches, i.e. for which all voxels belong to the same class. Robustness to outliers is achieved with Student distributions to model the heavy-tailed distributions of multi-channel image intensity. Supplementary material, available in the supplementary files /multimedia tab, presents 2D projections of obtained Student distributions.

As with patch-based segmentation, the variance $\sigma^2(x)$ for Log-Odds patch is spatially varying. In contrast to [38], $\sigma^2(x)$

TABLE I: Symmetry between likelihoods for segmentation & synthesis. Observed variables are displayed with grey background.

| Likelihoods | Segmentation model | Synthesis model |
|----------------------|---|---|
| Intensity likelihood | $\mathcal{N}(S(\mathbf{J}, x) S(\mathbf{I}_n, y), \sigma_n^2(x) \times \mathbb{1})$ | $\delta_{\mathbf{L}_n(y), \mathbf{L}(x)} \times \mathcal{N}(\widehat{\mathbf{J}}(x) \mathbf{I}_n(y), \Sigma_{\mathbf{L}_n(y)})$ |
| Label likelihood | $\delta_{\mathbf{L}_n(y), \widehat{\mathbf{L}}(x)}$ | $\mathcal{N}(S(\widehat{\mathbf{L}}, x) S(\widehat{\mathbf{L}}_n, y), \sigma^2(x) \times \mathbb{1})$ |

is a heuristic inspired by [21], [40] and involving the minimal patch distance, defined such that:

$$\sqrt{2}\sigma(x) = \min_{n,y} d\left(S(\widehat{\mathbf{L}}, x), S(\widehat{\mathbf{L}}_n, y)\right)$$

The same formula is applied to augmented Log-Odds patches during subsequent iterations of the proposed algorithm. As proximity between samples is not well characterized by their distance in high dimensional spaces, the use of the minimal distance instead of another measure (such as the average distance) is preferred to limit the number of patches with high probability.

2) *Approximate marginalization*: Assuming the membership index $\mathbf{M}(x)$ is independent and identically uniformly distributed $p(\mathbf{M}) = (N \times |\Omega|)^{-|\Omega|}$, the marginalization over \mathbf{M} [35] is given voxel-wise by:

$$\widehat{\mathbf{J}}(x) = \underset{j}{\operatorname{argmax}} \sum_{n=1}^N \sum_{y \in \Omega} p_n \left(S(\widehat{\mathbf{L}}, x) | \widehat{\mathbf{L}}_n, y, \sigma^2(x) \right) \\ \times p_n \left(\mathbf{J}(x) = j | \mathbf{I}_n, y, \Sigma_{\mathbf{L}_n(y)}, \mathbf{L}_n(y), \mathbf{L}(x) \right)$$

The marginalization over \mathbf{M} would require a high number of comparisons between patches, which is infeasible in practice. However, the marginalization step could be approximated, in a similar fashion as [38]. Let $y_n(S(\mathbf{L}, x))$ be the position in reference space of the closest Log-Odds patch, in canonical L_2 distance, of central label $l = \mathbf{L}(x)$ found in the atlas $(\mathbf{L}_n, \mathbf{I}_n)$. For image synthesis of pathological cases, we propose to marginalize over only N patches, where N is the total number of atlases: given a target Log-Odds patch of central label $\mathbf{L}(x)$, the approximate marginalization only includes patches found at the spatial position $y_n(S(\mathbf{L}, x))$ in atlas n , which results in:

$$\widehat{\mathbf{J}}(x) = \underset{j}{\operatorname{argmax}} \sum_{n=1}^N p_n \left(S(\widehat{\mathbf{L}}, x) | \widehat{\mathbf{L}}_n, y_n(S(\mathbf{L}, x)), \sigma^2(x) \right) \\ \times p_n \left(\mathbf{J}(x) = j | \mathbf{I}_n, y_n(S(\mathbf{L}, x)), \Sigma_{\mathbf{L}(x)} \right)$$

3) *Arg-max of the posterior distribution*: With the approximate marginalization, our original belief was that $\widehat{\mathbf{J}}(x)$ could be efficiently computed thanks to Banach Fixed Point Theorem. Indeed, a necessary condition for j^* to be an optimal solution is given by gradient cancelling. Let $F : \mathbb{R}^4 \mapsto \mathbb{R}$ be:

$$F(j) = \sum_{n=1}^N w_n \times p_n \left(\mathbf{J}(x) = j | \mathbf{I}_n, y_n(S(\mathbf{L}, x)), \Sigma_{\mathbf{L}(x)} \right)$$

where:

$$w_n = p_n \left(S(\widehat{\mathbf{L}}, x) | \widehat{\mathbf{L}}_n, y_n(S(\mathbf{L}, x)), \sigma^2(x) \right) > 0$$

Gradient cancelling gives the necessary condition that j^* is

a fixed point of a function $T : \mathbb{R}^4 \mapsto \mathbb{R}^4$:

$$j^* = T(j^*) = \sum_{n=1}^N \lambda_n(j^*) \times \mu_n$$

where:

$$\mu_n = \mathbf{I}_n(y_n(S(\mathbf{L}, x)))$$

$$\lambda_n(j) = \frac{w_n \times p_n \left(\mathbf{J}(x) = j | \mathbf{I}_n, y_n(S(\mathbf{L}, x)), \Sigma_{\mathbf{L}(x)} \right)}{\sum_{n=1}^N w_n \times p_n \left(\mathbf{J}(x) = j | \mathbf{I}_n, y_n(S(\mathbf{L}, x)), \Sigma_{\mathbf{L}(x)} \right)}$$

$$\forall n, \lambda_n(j) \geq 0 \quad \sum_{n=1}^N \lambda_n(j) = 1$$

For voxels x for which T is a contraction mapping, Banach Fixed Point Theorem gives that j^* is unique, and that for any initialization $j_0 \in \mathbb{R}^4$, any sequence $(j_k)_{k \in \mathbb{N}}$ such that $j_{k+1} = T(j_k)$ converges to the fixed point j^* with a geometric rate. An implementation of the EMP algorithm is given in Algorithm 1. In fact, we adopt a slight variant of this approach where $\lambda_n(j)$ is computed with a more robust Student likelihood (Section II-B1) instead of a Gaussian likelihood $p_n(\mathbf{J}(x) = j | \mathbf{I}_n, y_n(S(\mathbf{L}, x)), \Sigma_{\mathbf{L}(x)})$, to cope with the rare case where a few μ_n are far away from j^* . We could not prove that the Lipschitz constant of T is strictly less than 1 for every voxel in the image, hence the assumption of Banach Fixed Point Theorem may not be fulfilled. Nonetheless, in our experiments for all test cases, the algorithm does converge in about $K = 60$ iterations for 99.9% of the voxels of the brain mask. For the remaining 0.01% of the voxels, the intensity value obtained at the last iteration is used.

4) *Estimation of the image synthesis uncertainty*: With the Fixed Point procedure, the estimation of the arg-max of the posterior image intensity distribution involves the iterative estimation of j^* as a weighted average of intensities μ_n found in the atlases. Therefore, the uncertainty of the image synthesis process can be similarly estimated as the weighted sample variance $\widehat{\sigma}^2$, thanks to the optimal weights $\lambda_n(j^*)$. Estimation of image synthesis uncertainty is illustrated in Figure 4. Uncertainty estimation could allow to pin-point regions with unlikely image intensity based on the proposed model.

C. Consistent patch match

Patch match consistency can be improved by relying on augmented Log-Odds patches and by adopting an iterative feature augmentation strategy.

1) *Multi-scale patches*: As shown in [38], multi-scale patches are computationally efficient representations for large patches, which allows more effective patch match. The central

Algorithm 1: Extended Modality Propagation.

```

foreach iteration  $t \in 1, 2, \dots$  do
  foreach voxel  $x \in \Omega$  do
    evaluate label patch  $S(\mathbf{L}, x)$ 
    let  $l = \mathbf{L}(x)$  be the central label in  $S(\mathbf{L}, x)$ 
    evaluate Log-Odds patch  $S(\tilde{\mathbf{L}}, x)$ 
    augment Log-Odds patch with result at iteration  $t - 1$ 
    foreach training case  $n$  do
      find the closest patch of central label  $l$  in  $\tilde{\mathbf{L}}_n$ 
      let  $y_n(S(\mathbf{L}, x))$  be this patch spatial position
      evaluate Log-Odds patch  $S(\tilde{\mathbf{L}}_n, y_n(S(\mathbf{L}, x)))$ 
      evaluate image intensity  $\mathbf{I}_n(y_n(S(\mathbf{L}, x)))$ 
      augment Log-Odds patch with image intensity
    set variance  $\sigma^2(x)$  based on augmented patches
    set the intensity  $\hat{\mathbf{J}}(x)$  by applying a Fixed Point
    procedure as if  $T$  were a contraction mapping:

```

$$j_0 = \frac{\sum_{n=1}^N w_n \times \mu_n}{\sum_{n=1}^N w_n} \in \mathbb{R}^4$$

$$\forall k \in \mathbb{N}, j_{k+1} = T(j_k)$$

$$\hat{\mathbf{J}}(x) = \lim_{k \rightarrow +\infty} j_k \approx j_K$$




Fig. 4: Illustration of image synthesis uncertainty estimation. Along with the model prediction $\hat{\mathbf{J}}(x)$ (synthesis mean), the knowledge of the model uncertainty is illustrated with the standard-deviation $\hat{\sigma}(x)$ (square-root of the synthesis variance). Please note that only the FLAIR MR sequence is shown here, but $\hat{\mathbf{J}}(x)$ consists in fact in four synthesized MR channels, which are taken into account by $\hat{\sigma}(x)$ shown here. Uncertainty is more important at the boundary between cerebrospinal fluid and brain tissues (ventricles, sulci) and in the tumor core.

part of the patch is described voxel-wise, while the peripheral part is described by average intensity values over neighboring patches, by analogy with the foveal vision. Similar ideas regarding foveation and non-uniform sampling are detailed in [41]. In the following, parameters are set as in [38]: a $3 \times 3 \times 3$ patch is used for the central part, and a 26-neighborhood for the peripheral part.

2) *Patch match with invariance properties:* Every atlas is affinely registered to the MNI space [42] so that patch extraction is invariant to the pose of the subject. Moreover, the canonical L_2 distance used for patch matching is sensitive to rotation or symmetry of the patches. As shown in [38],

the combination of multi-scale patches and invariance to cube isometries yields better segmentation results, and thus could be beneficial for patch-based image synthesis as well. The 48 cube isometries consist of certain rotations under which the cube is invariant, plus their composition with central symmetry. For healthy brain structures, rotated patches can be unrealistic, only sagittal plane symmetry is considered. For pathological patches, the similarity measure should be insensitive to rotation and symmetry: the number of atlases is small, and using multi-scale patches further decreases the sample size. Therefore, the addition of plausible configurations of patches, generated by applying sagittal plane symmetry or cube isometries, to observed patches is relevant, and ultimately leads to a patch match invariant to the group of cube isometries. Indeed, if we were to apply cube isometries to the target patch, the minimal patch distance between the transformed target patch and the augmented learning dataset would be equal to the minimal patch distance between the target patch and the augmented learning dataset.

3) *Iterative feature augmentation:* An iterative augmentation procedure guarantees a spatially more coherent patch match, which results in an overall more coherent image synthesis. This is illustrated in Figure 5. Augmented Log-Odds patches are the concatenation of *i*) Log-Odds patches, weighted with $1 - \alpha_t$, and *ii*) multi-channel image patches synthesized during the previous iteration $t - 1$, weighted with α_t . There are different possible definitions for weight α_t . As with Modality Propagation [14], the feedback weight α_t increases with the number of iterations, starting with $\alpha_1 = 0$ for iteration $t = 1$. Unless mentioned, a maximal number of iterations is set to $t_{\max} = 3$, and $\alpha_t = (t - 1) / t_{\max}$, so that the weight of Log-Odds patches is never zero and decreases linearly with the number of iterations.

III. RESULTS

Our goal is to synthesize four MR channels commonly acquired for glioma assessment: pre-contrast T_1 -weighted image (T1), contrast-enhanced T_1 -weighted (TIC), T_2 -weighted (T2), and T_2 -FLAIR MR images. Fast approximate nearest-neighbour search are used for patch match, more precisely multiple randomized k-d trees for high dimensional data [43].

A. Pre-processing pipeline

Every image is affinely registered to an MNI atlas and warped to the same reference space. For experiments which require strictly more than 3 iterations, a resampling to 2 mm isotropic resolution is performed to decrease computation time, otherwise we keep the original 1 mm isotropic resolution. A resampling to 2 mm resolution was used by the top-performing glioma segmentation method at the 2012 BraTS challenge, which indicates that the generation of a synthetic dataset of 2 mm resolution images could already be of interest for applications other than medical image synthesis. Atlas images \mathbf{I}_n are standardized with a global affine intensity transform, applied to each modality independently: non-zero intensity are clipped below 1% and above 99% quantiles, then intensity mean and standard deviation are set to the corresponding




Fig. 5: From left to right: input label map, synthetic contrast-enhanced T_1 -weighted MRI after iterations $t = 1, 2, 3$, and ground truth (real MRI). Image parts near the brain border are better synthesized due to iterative feature augmentation. The presence of dark synthesized regions around the brain, instead of the presence of skull, is due to the fact that atlases are skull-stripped. Supplementary material, available in the supplementary files /multimedia tab, presents maps of absolute differences between synthesized images and ground truth.

average values over the training dataset, respectively 360 as mean and 120 as standard deviation. Except for the clipping, and potentially the resampling if performed, all the proposed pre-processing operations can be reverted, which eases the adaptation of such a synthetic dataset for the analysis of another real dataset.

B. Validation & comparison to state-of-the-art

1) *Dataset and evaluation:* The dataset consists of MRI of 20 high-grade (HG) and 10 low-grade (LG) cases, skull-stripped and made publicly available as training dataset in the 2013 edition of the MICCAI BraTS benchmark [3]. A leave-one-out procedure is always applied to exclude the target image from the training dataset. There is no atlas pre-selection and all training cases, remaining after leave-one-out, are used as atlases for the synthesis of MRI of cases of the same grade. Only the segmentation of the tumor compartments is originally available [3], so healthy tissues are automatically segmented by applying FSL FAST [44] to the T1C image, which consistently has the highest resolution among the different MR channels.

The validation of the experiment consists in comparing real MRI with synthetic MRI generated solely based on a segmentation of the brain and the tumor. Quantitative assessment include Peak Signal to Noise Ratio (PSNR) to quantify the quality of the synthesis, Mean-Absolute-Error (MAE) to quantify bias, and finally mean Structural Similarity (SSIM) [45] (with default values) as an additional image reconstruction criterion. All these measures are classically used for assessment of image reconstruction or synthesis [25], [46]–[48]. Whenever an average is computed, only the voxels which belong to a region-of-interest (ROI) are considered: the ROI is the whole brain mask as defined by the input label map (segmentation of healthy and pathological tissues).

Our image synthesis model is compared to the image synthesis procedure described in [28]. Supplementary material, available in the supplementary files /multimedia tab, includes

the user-specific coefficients chosen in our implementation of this procedure [28].

For a fair comparison, synthetic MRI generated with [28] *i)* use the same input categorical label map as our method, and *ii)* are standardized as described in Section III-A, so that the average image intensity and the variance corresponds to the average and variance of the real image. Supplementary material, available in the supplementary files /multimedia tab, illustrates results obtained with [28] based on tumor growth model probabilistic outputs instead. However, no quantitative evaluation can be performed on these results due to a lack of ground truth.

2) *Results:* A visual inspection of image synthesis results can be performed in axial views for a high-grade glioma in Figure 6. Axial, coronal and sagittal views are also presented in smaller displays in supplementary material, available in the supplementary files /multimedia tab. The model does not replicate artifacts or croppings present on some real cases. This is expected, since the image synthesis model is based on a label map, which was created from a high-resolution T1C image of a whole brain.

Quantitative results are reported for each MR channel in Figure 7 for PSNR, and in supplementary material, available in the supplementary files /multimedia tab, for other assessment measures. A summary of assessment measures is presented in Table II. PSNR, MAE, and mean SSIM are always higher for the proposed image synthesis method. Moreover, with the linearly increasing feedback weight, most results are improved with additional iterations. To confirm this statement, paired two-sample t-tests are performed independently for each MR channel. With 1% significance level, EMP with 2 iterations performs significantly better than the state-of-the-art [28] for every criterion (PSNR/SSIM/MAE) and for every MR channel (T1/T1C/T2/FLAIR), which is not the case when using EMP with a single iteration. Moreover, with 1% significance level, from iteration 1 to iteration 3, EMP improves significantly the image between each iteration.

The influence of the feedback weight α_t was studied. The




Fig. 6: Qualitative evaluation of our synthesis method after 3 iterations on a high-grade case (HG-0009). Axial views. Top to bottom: synthetic MRI obtained with our method, ground truth (real MRI), synthetic MRI obtained with the method described in [28] with the same categorical label map as input. Ground truth FLAIR image is partially dark, corresponding to the brain mask for which the FLAIR image is unknown, most likely for a faster MR acquisition. Interestingly, this does not happen on T1C images, and hence does not affect the label map for healthy brain tissues. Synthetic images are not affected since they rely only on the label map. Supplementary material, available in the supplementary files /multimedia tab, presents maps of absolute differences between synthesized images and ground truth.

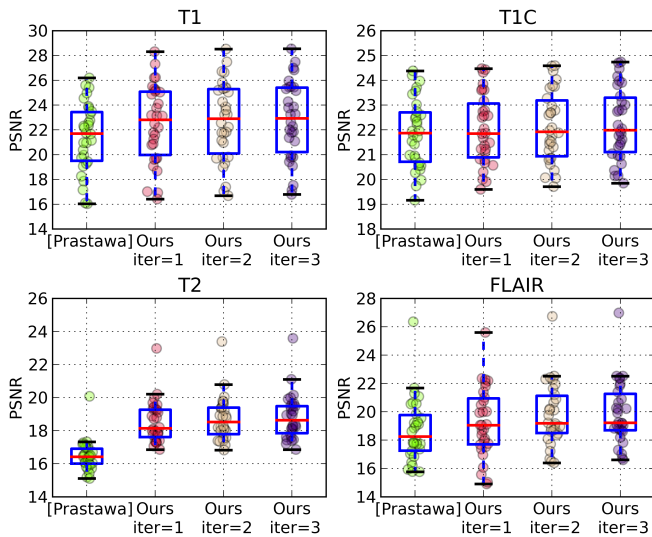


Fig. 7: Comparison to ground truth MRI in terms of PSNR for each MR sequence. From left to right: state-of-the-art [28], iterations 1, 2 and 3 of our method.

protocol is the following: different values are chosen for feedback weights, then Extended Modality Propagation is run for 10 iterations with $\alpha_1 = 0$ and subsequent α_t fixed at the chosen value. The highest PSNR (18.6) and highest mean-SSIM (0.63) are obtained for the smallest non-zero feedback weight (1/9). The variance of the PSNR is lower when a non-zero feedback weight is chosen, however the median PSNR is nearly the same for all feedback weight values. MAE increases when the weight increases. In the end, any small non-zero feedback weight is reasonable for image synthesis.

The influence of the number of iterations of EMP was studied. A high constant feedback weight, fixed at $2/3$, is chosen to emphasize the effect of the number of iterations. The highest PSNR (18.7) and highest mean SSIM (0.64) are obtained at the end of the second iteration. MAE decreases notably with the number of iterations. In the end, a small number of iterations should be preferred, typically 2 to 3 iterations. A similar trend is observed with a linearly increasing feedback weight α_t .

TABLE II: Average assessment measures for image synthesis, based on the 2013 BRATS Evaluation dataset.

| T1 | PSNR | mean SSIM | MAE |
|-----------------------|-------------|-------------|-------------|
| State-of-the-art [28] | 21.4 | 0.70 | 80.7 |
| Ours (iteration 1) | 22.5 | 0.76 | 65.6 |
| Ours (iteration 2) | 22.7 | 0.81 | 63.0 |
| Ours (iteration 3) | 22.8 | 0.82 | 62.6 |
| T1C | PSNR | mean SSIM | MAE |
| State-of-the-art [28] | 21.9 | 0.77 | 75.2 |
| Ours (iteration 1) | 22.0 | 0.76 | 70.1 |
| Ours (iteration 2) | 22.1 | 0.81 | 68.6 |
| Ours (iteration 3) | 22.2 | 0.82 | 68.0 |
| T2 | PSNR | mean SSIM | MAE |
| State-of-the-art [28] | 16.5 | 0.67 | 97.2 |
| Ours (iteration 1) | 18.5 | 0.69 | 76.0 |
| Ours (iteration 2) | 18.7 | 0.74 | 73.3 |
| Ours (iteration 3) | 18.8 | 0.75 | 73.0 |
| FLAIR | PSNR | mean SSIM | MAE |
| State-of-the-art [28] | 18.7 | 0.63 | 95.4 |
| Ours (iteration 1) | 19.3 | 0.66 | 78.2 |
| Ours (iteration 2) | 19.8 | 0.73 | 73.6 |
| Ours (iteration 3) | 19.9 | 0.75 | 73.1 |

C. Illustration using an advanced tumor growth model

In this section, we use an advanced tumor growth model to generate probabilistic label maps, based on which we create a hard label map. The proposed image synthesis model is then used to generate synthetic MRI of virtual patients. The tumor growth model is a multi-population cell model driven by angiogenesis and vascularization. It is a variation of the Proliferation Invasion Hypoxia Necrosis Angiogenesis (PIHNA) model described in 1D in [49]. It was also used in 3D to simulate the impact of anti-angiogenic therapy in an atlas of a healthy brain [50]. Supplementary material, available in the supplementary files /multimedia tab, details the model and its implementation. Synthetic MRI based on the PIHNA model are shown in Figure 8.

IV. DISCUSSION

In the experiments presented in this paper, the iterative nature of the proposed image synthesis method allows to reach higher PSNR and synthesize qualitatively more realistic images. The optimal number of iterations seems to be about 2 or 3. A progressive increase of the feedback weight with the number of iterations, as suggested in [14] leads to better synthesis results.

Automatically setting the value for the feedback weight would require to consider the weight as an additional random variable together with an hyper prior which depends on the amount of spatial correlation of \mathbf{J} .

Several MR channels are simultaneously synthesized, which leads to a more constrained patch match (for iterations $t > 1$), and ultimately to more consistent synthetic MRI. However, using a higher number of MR sequences implies that patches lie in a space of higher dimension, which tends to homogenize patch distances. Indeed, for high dimensional data, every patch is more or less a close neighbor of others, and the notion of nearest-neighborhood becomes less relevant. This could result in fuzzier synthetic images. Conversely, less consistent, yet potentially more textured, results could arise when trying to synthesize a single image modality.

EMP relies on a marginalization over several training cases, which has an averaging effect. Therefore, if labeling errors are assumed to be incoherent in the training cases, an incorrect labeling should have little impact, theoretically proportional to 1 over the number of atlases. However, if the target label map \mathbf{L} is close to one of the incorrect segmentations, then the impact would be greater in practice.

One limitation of the proposed method is the running time: with 20 atlases, one iteration to simultaneously synthesize the 4 MR channels (for the whole brain) requires about 2 days of computation (16 cores, 100 GB of RAM) for 1 mm isotropic MRI, and about 3 hours of computation (3 cores, 20 GB of RAM) for 2 mm isotropic MRI. Typically, sulci are less visible on 2 mm synthetic MRI. Atlas selection could result in a decrease of the number of atlases, which could ultimately lead to finer details on 2 mm synthetic MRI. However, this would require a strategy to pre-select relevant atlases based on some similarity criterion between label maps.

Regarding running time in the perspective of generating a large database of synthetic cases, a region-of-interest enclosing the tumor could be defined such that the rest of the synthetic image is supplied by a template. Indeed, tumor growth models are based on a healthy atlas in which a tumor seed is placed, and a T1 template is available for the healthy atlas [42]. In this case, image intensity does not have to be optimized for voxels outside the ROI, since these voxels are assumed not to present intensity abnormalities.

Without resorting to tumor growth models, by only performing the patch-based synthesis on pathological regions, the extended modality propagation could be applied to any brain MR image, provided that a realistic label map of tumor compartments exists for that image. The generation of those tumor label maps could be either borrowed and transported from real pathological cases in the BRATS database or could be generated procedurally through an additional label map model. This would allow to create a high number of complex and detailed cases, with different locations of the same tumor label maps in the brain. The impact of this perspective would have to be carefully considered since some tumor patterns might be more or less likely in different brain regions.

All experiments were conducted with a label map as single input. Due to the iterative process described in this paper, it would be straightforward to provide as input an MR channel on top of the label map. Typically, if a contrast-enhanced T1 MR image is available on top of the label map, we could synthesize MR channels such as FLAIR or T2. Additional inputs are merely additional constraints for the patch match.

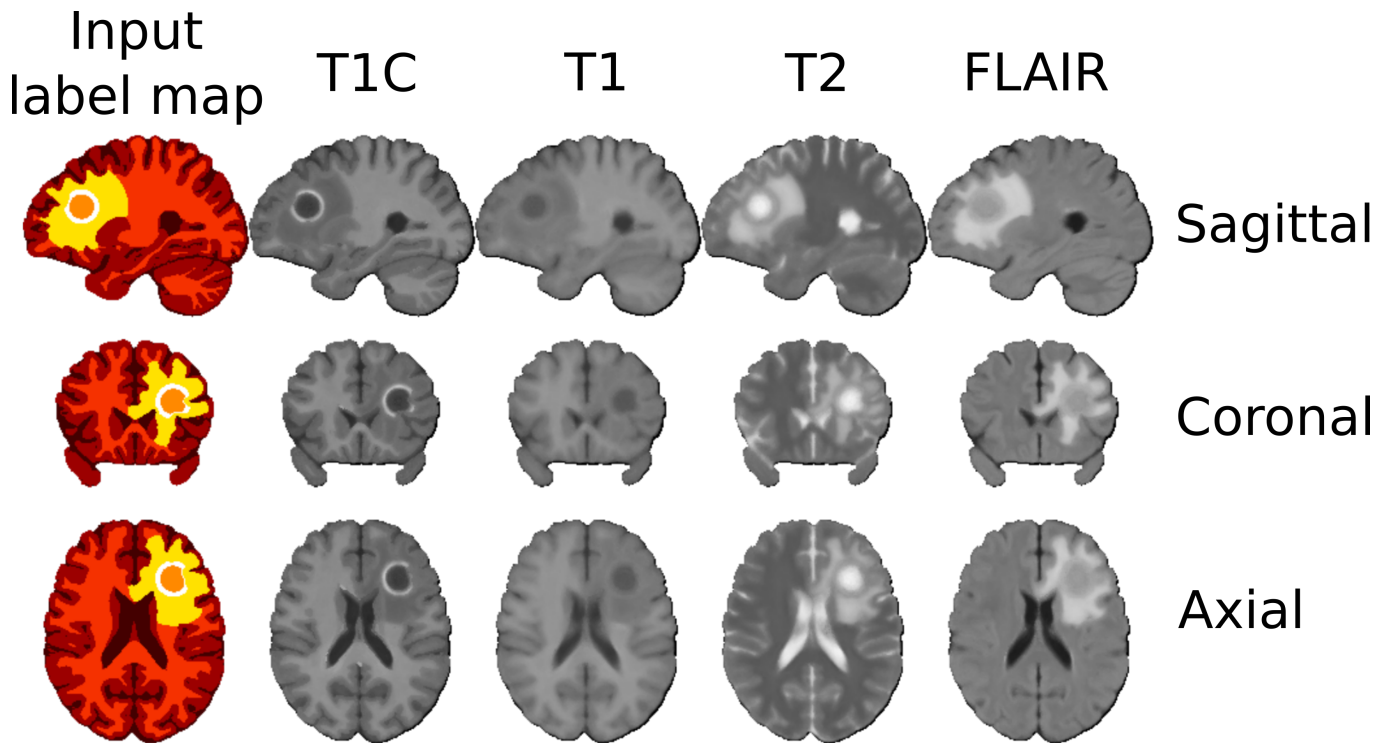


Fig. 8: Illustration of the modularity of our image synthesis model. Synthetic 1 mm isotropic MRI, obtained after 3 iterations, based on a categorical label map generated by the PIHNA tumor growth 3D model.

V. SUMMARY AND CONCLUSIONS

We presented a generative model for the synthesis of multi-modal medical images of pathological cases by using an iterative multi-atlas patch-based algorithm. The method was illustrated on MRI synthesis of glioma-bearing brains. A single label maps allows to synthesize realistic images, with clear improvements with regard to the state-of-the-art. The approach is flexible and can be applied *i)* either to label maps of real cases to synthesize artifact-free MRI, *ii)* or to label maps simulated by any tumor growth model. The standardization of images is made transparent and reversible, so that synthetic data can be used for any application with any other dataset, including intensity inhomogeneity or different levels of Rician noise. Moreover, due to the strategy used to solve for the maximum a posteriori, estimation of image synthesis uncertainty is straightforward.

Such a generative model would allow the availability of large, public and annotated datasets of synthetic medical images obtained with virtual patients, which could prove useful for the validation or benchmarking of image processing algorithms. Promising results have been obtained on real data after training machine learning algorithms on synthetic data [29], which could be further improved by using a more realistic synthetic dataset, generated with the proposed image synthesis model. Finally, image synthesis of pathological cases bridges the gap between outputs of tumor growth model and medical images, which are the most clinically relevant observations. The development of image synthesis could benefit to interpretability and personalization of tumor growth models.

ACKNOWLEDGEMENTS

Part of this work was funded by the European Research Council through the ERC Advanced Grant MedYMA 2011-291080 (on Biophysical Modeling and Analysis of Dynamic Medical Images).

REFERENCES

- [1] C. R. Jack, M. A. Bernstein, N. C. Fox, P. Thompson, G. Alexander, D. Harvey, B. Borowski, P. J. Britson, J. L. Whitwell, C. Ward et al., “The Alzheimer’s disease neuroimaging initiative (ADNI): MRI methods,” *Journal of Magnetic Resonance Imaging*, vol. 27, no. 4, pp. 685–691, 2008.
- [2] E. A. Regan, J. E. Hokanson, J. R. Murphy, B. Make, D. A. Lynch, T. H. Beaty, D. Curran-Everett, E. K. Silverman, and J. D. Crapo, “Genetic epidemiology of copd (copdgene) study design,” *COPD: Journal of Chronic Obstructive Pulmonary Disease*, vol. 7, no. 1, pp. 32–43, 2011.
- [3] B. Menze, M. Reyes, K. Van Leemput et al., “The Multimodal Brain Tumor Image Segmentation Benchmark (BraTS),” *Medical Imaging, IEEE Transactions on*, vol. 34, no. 10, pp. 1993–2024, October 2015.
- [4] C. Wachinger, P. Golland, W. Kremen, B. Fischl, M. Reuter, A. D. N. Initiative et al., “Brainprint: A discriminative characterization of brain morphology,” *NeuroImage*, vol. 109, pp. 232–248, 2015.
- [5] M. Toews, C. Wachinger, R. S. J. Estepar, and W. M. W. III, “A feature-based approach to big data analysis of medical images,” in *Information Processing in Medical Imaging*. Springer Berlin Heidelberg, 2015.
- [6] K. Clark, B. Vendt, K. Smith, J. Freymann, J. Kirby, P. Koppel, S. Moore, S. Phillips, D. Maffitt, M. Pringle et al., “The cancer imaging archive (tcia): maintaining and operating a public information repository,” *Journal of digital imaging*, vol. 26, no. 6, pp. 1045–1057, 2013.
- [7] M. Havaei, A. Davy, D. Warde-Farley, A. Biard, A. Courville, Y. Bengio, C. Pal, P.-M. Jodoin, and H. Larochelle, “Brain tumor segmentation with deep neural networks,” *arXiv preprint arXiv:1505.03540*, 2015.
- [8] R. K. Kwan, A. C. Evans, and G. B. Pike, “MRI simulation-based evaluation of image-processing and classification methods,” *Medical Imaging, IEEE Transactions on*, vol. 18, no. 11, pp. 1085–1097, 1999.

- [9] H. Benoit-Cattin, G. Collewet, B. Belaroussi, H. Saint-Jalmes, and C. Odet, "The SIMRI project: a versatile and interactive MRI simulator," *Journal of Magnetic Resonance*, vol. 173, no. 1, pp. 97–115, 2005.
- [10] T. Glatard, C. Lartizien, B. Gibaud, R. F. da Silva, G. Forestier, F. Cerve-nansky, M. Alessandrini, H. Benoit-Cattin, O. Bernard, S. Camarasu-Pop et al., "A Virtual Imaging Platform for multi-modality medical image simulation," *Medical Imaging, IEEE Transactions on*, vol. 32, no. 1, pp. 110–118, 2013.
- [11] F. Bloch, W. Hansen, and M. Packard, "Nuclear induction," *Physical review*, vol. 70, no. 7-8, pp. 460–474, 1946.
- [12] R. Maitra and J. J. Riddles, "Synthetic magnetic resonance imaging revisited," *Medical Imaging, IEEE Transactions on*, vol. 29, no. 3, pp. 895–902, 2010.
- [13] J. Iglesias, I. Dinov, J. Singh, G. Tong, and Z. Tu, "Synthetic MRI signal standardization: application to multi-atlas analysis," *MICCAI 2010*, pp. 81–88, 2010.
- [14] D. H. Ye, D. Zikic, B. Glocker, A. Criminisi, and E. Konukoglu, "Modality Propagation: coherent synthesis of subject-specific scans with data-driven regularization," in *MICCAI 2013*. Springer, 2013, pp. 606–613.
- [15] R. Leahy and X. Yan, "Incorporation of anatomical MR data for improved functional imaging with PET," in *Information Processing in Medical Imaging*. Springer Berlin Heidelberg, 1991, pp. 105–120.
- [16] F. Rousseau, "Brain hallucination," *Computer Vision—ECCV 2008*, pp. 497–508, 2008.
- [17] F. Rousseau, A. D. N. Initiative et al., "A non-local approach for image super-resolution using intermodality priors," *Medical image analysis*, vol. 14, no. 4, pp. 594–605, 2010.
- [18] S. Roy, A. Carass, and J. L. Prince, "Synthesizing MR contrast and resolution through a patch matching technique," *Proceedings-Society of Photo-Optical Instrumentation Engineers*, vol. 7623, p. 76230j, 2010.
- [19] S. Roy, A. Carass, N. Shiee, D. L. Pham, and J. L. Prince, "MR contrast synthesis for lesion segmentation," in *Biomedical Imaging: From Nano to Macro, 2010 IEEE International Symposium on*. IEEE, 2010, pp. 932–935.
- [20] F. Rousseau, P. A. Habas, and C. Studholme, "A supervised patch-based approach for human brain labeling," *Medical Imaging, IEEE Transactions on*, vol. 30, no. 10, pp. 1852–1862, 2011.
- [21] P. Coupé, J. V. Manjón, V. Fonov, J. Pruessner, M. Robles, and D. L. Collins, "Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation," *NeuroImage*, vol. 54, no. 2, pp. 940–954, 2011.
- [22] N. Burgos, M. J. Cardoso, K. Thielemans, M. Modat, S. Pedemonte, J. Dickson, A. Barnes, R. Ahmed, C. J. Mahoney, J. M. Schott et al., "Attenuation correction synthesis for hybrid PET-MR scanners: application to brain studies," *Medical Imaging, IEEE Transactions on*, vol. 33, no. 12, pp. 2332–2341, 2014.
- [23] J. E. Iglesias, E. Konukoglu, D. Zikic, B. Glocker, K. Van Leemput, and B. Fischl, "Is synthesizing MRI contrast useful for inter-modality analysis?" in *MICCAI 2013*. Springer, 2013, pp. 631–638.
- [24] S. Roy, A. Carass, and J. L. Prince, "Magnetic resonance image example-based contrast synthesis," *Medical Imaging, IEEE Transactions on*, vol. 32, no. 12, pp. 2348–2363, 2013.
- [25] M. J. Cardoso, C. H. Sudre, M. Modat, and S. Ourselin, "Template-Based Multimodal Joint Generative Model of Brain Data," in *Information Processing in Medical Imaging*. Springer International Publishing, 2015, pp. 17–29.
- [26] J. Rexilius, H. K. Hahn, M. Schlüter, S. Kohle, H. Bourquain, J. Böttcher, and H.-O. Peitgen, "A framework for the generation of realistic brain tumor phantoms and applications," in *MICCAI 2004*. Springer, 2004, pp. 243–250.
- [27] M. Prastawa, E. Bullitt, and G. Gerig, "Synthetic ground truth for validation of brain tumor MRI segmentation," in *MICCAI 2005*. Springer, 2005, pp. 26–33.
- [28] —, "Simulation of brain tumors in MR images for evaluation of segmentation efficacy," *Medical Image Analysis*, vol. 13, no. 2, p. 297, 2009.
- [29] E. Geremia, B. H. Menze, M. Prastawa, M.-A. Weber, A. Criminisi, and N. Ayache, "Brain tumor cell density estimation from multi-modal MR images based on a synthetic tumor growth model," in *Medical Computer Vision. Recognition Techniques and Applications in Medical Imaging*. Springer, 2013, pp. 273–282.
- [30] A. Hamamci, N. Kucuk, K. Karaman, K. Engin, and G. Unal, "Tumor-cut: segmentation of brain tumors on contrast enhanced MR images for radiosurgery applications," *Medical Imaging, IEEE Transactions on*, vol. 31, no. 3, pp. 790–804, 2012.
- [31] A. Galimzianova, F. Pernus, B. Likar, and Z. Spiclin, "Robust estimation of unbalanced mixture models on samples with outliers," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. PP, no. 99, 2015.
- [32] S. Bauer, C. May, D. Dionysiou, G. Stamatakos, P. Büchler, and M. Reyes, "Multiscale modeling for image analysis of brain tumor studies," *Biomedical Engineering, IEEE Transactions on*, vol. 59, no. 1, pp. 25–29, 2012.
- [33] X. Liu, M. Niethammer, R. Kwitt, N. Singh, M. McCormick, and S. Aylward, "Low-rank Atlas Image Analyses in the Presence of Pathologies," *Medical Imaging, IEEE Transactions on*, vol. PP, no. 99, 2015.
- [34] S. Gu, G. Chakraborty, K. Champley, A. M. Alessio, J. Claridge, R. Rockne, M. Muzi, K. A. Krohn, A. M. Spence, E. C. Alvord et al., "Applying a patient-specific bio-mathematical model of glioma growth to develop virtual [18F]-FMISO-PET images," *Mathematical Medicine and Biology*, vol. 29, no. 1, pp. 31–48, 2012.
- [35] M. R. Sabuncu, B. T. Yeo, K. Van Leemput, B. Fischl, and P. Golland, "A generative model for image segmentation based on label fusion," *Medical Imaging, IEEE Transactions on*, vol. 29, no. 10, pp. 1714–1729, 2010.
- [36] W. Bai, W. Shi, D. P. O'Regan, T. Tong, H. Wang, S. Jamil-Copley, N. S. Peters, and D. Rueckert, "A probabilistic patch-based label fusion model for multi-atlas segmentation with registration refinement: application to cardiac MR images," *Medical Imaging, IEEE Transactions on*, vol. 32, no. 7, pp. 1302–1315, 2013.
- [37] J. E. Iglesias and M. R. Sabuncu, "Multi-Atlas Segmentation of Biomedical Images: A Survey," *Medical Image Analysis*, 2015.
- [38] N. Cordier, H. Delingette, and N. Ayache, "A patch-based approach for the segmentation of pathologies: Application to glioma labelling," *Medical Imaging, IEEE Transactions on*, vol. PP, 2015.
- [39] C. Liu and D. B. Rubin, "ML estimation of the t distribution using EM and its extensions, ECM and ECME," *Statistica Sinica*, vol. 5, no. 1, pp. 19–39, 1995.
- [40] Z. Wang, K. K. Bhatia, B. Glocker, A. Marvao, T. Dawes, K. Misawa, K. Mori, and D. Rueckert, "Geodesic patch-based segmentation," in *MICCAI 2014*. Springer International Publishing, 2014, pp. 666–673.
- [41] D. Ciresan, A. Giusti, L. M. Gambardella, and J. Schmidhuber, "Deep neural networks segment neuronal membranes in electron microscopy images," in *Advances in neural information processing systems*, 2012, pp. 2843–2851.
- [42] J. Mazziotta, A. Toga, A. Evans, P. Fox, J. Lancaster, K. Zilles, R. Woods, T. Paus, G. Simpson, B. Pike et al., "A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM)," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 356, no. 1412, pp. 1293–1322, 2001.
- [43] M. Muja and D. G. Lowe, "Scalable Nearest Neighbor Algorithms for High Dimensional Data," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 36, 2014.
- [44] Y. Zhang, M. Brady, and S. Smith, "Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm," *Medical Imaging, IEEE Transactions on*, vol. 20, no. 1, pp. 45–57, 2001.
- [45] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *Image Processing, IEEE Transactions on*, vol. 13, no. 4, pp. 600–612, 2004.
- [46] J. V. Manjón, P. Coupé, A. Buades, V. Fonov, D. L. Collins, and M. Robles, "Non-local MRI upsampling," *Medical image analysis*, vol. 14, no. 6, pp. 784–792, 2010.
- [47] E. Konukoglu, A. van der Kouwe, M. R. Sabuncu, and B. Fischl, "Example-based restoration of high-resolution magnetic resonance image acquisitions," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2013*. Springer, 2013, pp. 131–138.
- [48] P. A. Gómez, J. I. Sperl, T. Sprenger, C. Metzler-Baddeley, D. K. Jones, and P. Saemann, "Joint Reconstruction of Multi-Contrast MRI for Multiple Sclerosis Lesion Segmentation," in *Bildverarbeitung für die Medizin 2015*. Springer, 2015, pp. 155–160.
- [49] K. R. Swanson, R. C. Rockne, J. Claridge, M. A. Chaplain, E. C. Alvord, and A. R. Anderson, "Quantifying the role of angiogenesis in malignant progression of gliomas: in silico modeling integrates imaging and histology," *Cancer Research*, vol. 71, no. 24, pp. 7366–7375, 2011.
- [50] O. Saut, J.-B. Lagaert, T. Colin, and H. M. Fathallah-Shaykh, "A multilayer grow-or-go model for gbm: effects of invasive cells and anti-angiogenesis on growth," *Bulletin of mathematical biology*, vol. 76, no. 9, pp. 2306–2333, 2014.