BRIDGEUQ: BAYESIAN ALEATORIC UNCERTAINTY QUANTIFICATION FOR TIME-SERIES IMAGE REGISTRATION VIA BROWNIAN BRIDGE PRIOR

Anonymous authors
Paper under double-blind review

ABSTRACT

This paper presents *BridgeUQ*, a novel Bayesian framework for aleatoric uncertainty quantification in pretrained deterministic registration models. In contrast to previous approaches that mainly focus on pairwise registration, our work is the first to model posterior uncertainty over spatiotemporal deformation trajectories that captures the intrinsic ambiguity of registration solutions across both space and time. To achieve this, we first introduce a novel hierarchical Gaussian process prior in the temporal deformation space, based on a Brownian bridge formulation with an inverse-Gamma hyper prior on the diffusion variance, to govern stochastic sampling around the predicted registration solution. We then define an energy-based likelihood function derived from the pretrained registration objective to ensure the sampled deformations remain faithful to the observed data. To facilitate efficient inference, we further develop an amortized maximum-a-posterior (MAP) scheme combined with conditional diffusion sampling to approximate this resulting posterior. It is important to note that BridgeUQ is a plug-and-play module that can be seamlessly integrated with arbitrary pretrained registration models without retraining. We validate BridgeUQ on widely used registration networks using comprehensive cardiac and brain magnetic resonance imaging datasets. Experimental results show that our approach is able to produce anatomically meaningful uncertainty estimates for a variety of learning-based registration models. The proposed BridgeUQ provides a practical and principled pathway toward uncertainty-aware and safety-critical medical image registration. Our code is publicly available at *https://anonymous.4open.science/r/BridgeUQ-8F33/*.

1 INTRODUCTION

Deep learning-based time-series image registration plays an important role in medical imaging analysis, including but not limited to template-based segmentation Nekrasov et al. (2020); Hossain & Zhang (2025); Aksoy et al. (2013), image reconstruction Jayakumar et al. (2023); Li et al. (2020), and object motion analysis Becker et al. (2006); Pan et al. (2021). While recent work has made significant strides in improving registration accuracy Wu et al. (2024b;a); Lu et al. (2023), relatively little attention has been paid to quantifying the models uncertainty in its predicted deformation fields. One major source of the uncertainty in the networks predictions is *aleatoric uncertainty*, which arises from ill-posed nature of registration, where multiple plausible deformation solutions can satisfy the same network objective (especially under large spatial and temporal changes) Pennec et al. (2019); Younes (2010). This typically occurs in regions with low texture or homogeneous intensity, or in the presence of noise or low contrast Risholm et al. (2013); Hu et al. (2025); Gong et al. (2022). Capturing this aleatoric uncertainty is critical in time-series medical imaging, where trajectories such as cardiac motion or neurodegenerative progression evolve over structured temporal domains (Davatzikos et al., 2001; Fiot et al., 2014). It enables identification of *when* and *where* registration becomes unreliable; hence providing confidence estimates for assessing whether learning-based registration results can be trusted in high-stakes medical imaging applications Chen et al. (2025).

Related Works. Aleatoric uncertainty quantification in medical image registration has traditionally relied on probabilistic models, such as Bayesian inference or bootstrap sampling to approximate posterior distributions over deformation parameters (Simpson et al., 2013; Risholm et al., 2013;

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Kybic, 2009; Le Folgoc et al., 2016; 2017). Beyond this, registration uncertainty has also been modeled through stochastic differential equation (SDE)-based formulations, which allow Langevin sampling, as well as Gaussian process (GP) for landmark-based deformations (Janoos et al., 2012; Hub et al., 2009; Marsland & Shardlow, 2017; Wassermann et al., 2014). While effective, these methods are computationally expensive due to intensive sampling or iterative inference procedures; hence do not scale well. To address this issue, later approaches focused on improving inference tractability via low-dimensional probabilistic representations of deformations (Zhang et al., 2017; Marsland & Shardlow, 2017; Peter et al., 2021; Wang et al., 2018). With the blooming advances in deep learning, learning-based methods have increasingly replaced traditional registration approaches due to their superior performance in fast prediction. As a result, much recent work has shifted to modeling registration uncertainty within deep networks. Exemplary research includes probabilistic registration frameworks, where neural networks directly learn distributions over latent deformation fields Dalca et al. (2019b); Krebs et al. (2019); Siegert et al. (2023); Gong et al. (2022); Sedghi et al. (2019); Gan et al. (2022). Building on this idea, hierarchical variational methods Krebs et al. (2019); Siegert et al. (2023) further extend uncertainty modeling to multi-resolution deformation representations. However, these approaches primarily focus on pairwise image registration and do not explicitly model temporal correlations, making them less suitable for time-sequential imaging data. As a result, uncertainty is typically estimated independently at each time point, which fails to capture the temporally coherent structure of object motion.

In this paper, we present BridgeUQ, a novel Bayesian framework with an efficient inference scheme for quantifying aleatoric uncertainty of any pre-trained registration network on time-series data. In contrast to previous methods Dalca et al. (2019b); Wang et al. (2019); Siegert et al. (2023); Grzech et al. (2021), we are the first to develop a spatiotemporal posterior over deformation trajectories, which enables coherent uncertainty characterization that is both spatially and temporally coherent. This is particularly critical in high-stake domains, where respecting anatomical and physiological constraints are desired properties Mansilla et al. (2020). Our BridgeUQ provides information not only on *where* but also on *when* registration is reliable or uncertain along the time sequence. To achieve this, we introduce a new GP prior via a learnable Brownian bridge to naturally model the stochasticity around predicted spatiotemporal deformation trajectories. A data likelihood is then designed to appropriately balance data fidelity and stochastic variability; hence providing principled uncertainty quantification without sacrificing registration accuracy. To enable efficient and tractable inference within this Bayesian framework, we employ an amortized inference. Our main contributions are:

(i) Develop a Bayesian uncertainty quantification framework for spatiotemporal deformation fields by introducing a hierarchical GP Brownian bridge prior with a learnable variance schedule, which is treated as a random variable governed by an inverse-Gamma prior.

(ii) Introduce an energy-based data likelihood that explicitly governs the stochastic process of deformation trajectories. An amortized MAP scheme is designed for tractable and efficient approximation of the posterior over temporal deformations.

(iii) Provide a flexible plug-and-play module that can be integrated into a wide range of pretrained deterministic registration networks without requiring retraining from scratch.

Experimental results demonstrate that BridgeUQ produces anatomically consistent uncertainty estimates across diverse registration backbones, substantially outperforming state-of-the-art UQ baselines Hu et al. (2025); Siegert et al. (2023); Grzech et al. (2021); Wang et al. (2018) without compromising registration accuracy. Notably, BridgeUQ highlights its potential for safety- and reliability-aware deployment of learning-based registration in real-world clinical applications.

2 Background: Time-series Image Registration

This section briefly reviews time-series image registration, which estimates spatiotemporal deformations from a given image sequence Miller et al. (2003); Wu & Zhang (2023); Arnold (1966). Let $\Omega = \mathbb{R}^d/\mathbb{Z}^d$ be a d-dimensional torus domain with periodic boundary conditions. Consider a given dataset of N time-sequential images, $\{\mathcal{I}^i\}_{i=1}^N, i \in \{1, \cdots, N\}$, where each sequence $\mathcal{I}^i$ contains $T+1$ temporal frames, i.e., $\mathcal{I}^i = \{\mathcal{I}_t^i\}_{t=0}^T, t \in [0, T]$. Each image is definded on the torus domain, Ω, i.e., $\mathcal{I}_t^i : \Omega \to \mathbb{R}$. The goal of deformable time-series image registration is to estimate a family of transformations $\{\phi_t^i\}_{t=1}^T$ with $\phi_t^i : \Omega \to \Omega$, such that the reference frame $\mathcal{I}_0^i$ warped by ϕ_t^i aligns with

the corresponding target frame $\mathcal{I}_t^i$. In general, the time-series registration problem can be formulated as the minimization of an energy functional

$$E(\{\phi_t^i\}) = \sum_{t,i=1}^{T,N} \text{Dist}\left(\mathcal{I}_0^i \circ \phi_t^i, \mathcal{I}_t^i\right) + \lambda \,\text{Reg}(\phi_t^i), \tag{1}$$

where $\circ$ denotes an interpolation operator. The $\text{Dist}(\cdot,\cdot)$ measures image dissimilarity, with common choices including normalized cross-correlation (NCC) Avants et al. (2008), sum of squared differences (SSD) Rueckert et al. (2006), and mutual information (MI) Wells III et al. (1996). The $\text{Reg}(\cdot)$ is a regularization term that enforces the spatiotemporal smoothness of the transformation field. The parameter $\lambda > 0$ controls the trade-off between image matching fidelity and deformation regularity.

Various regularity schemes. The regularization term, $\text{Reg}(\cdot)$, typically includes complementary spatial and temporal regularizer. The temporal regularization encourages temporal consistency by penalizing changes in the transformation between consecutive frames Krebs et al. (2021); Ye et al. (2021); hence to ensure the estimated motion evolves smoothly over time.

In contrast, the spatial regularizer is relatively more complex, as it encourages voxel-wise smoothness of the estimated transformation on dense image grid; thereby suppressing spatial artifacts and promoting physically plausible deformation fields, such as diffeomorphic transformations. Common spatial regularizations, such as stationary velocity fields (SVF) Arsigny et al. (2006) and large deformation diffeomorphic metric mapping (LDDMM) Arnold (1966); Beg et al. (2005), are widely adopted to enforce diffeomorphic constraints (i.e., one-to-one smooth, and invertible smooth mappings) to ensure anatomically plausible and topology-preserving transformations. Mathematically, they penalize the tangent vector of the transformation, i.e., the velocity field v_t^i, defined through the evolution equation Beg et al. (2005); Arsigny et al. (2006); Arnold (1966); Miller et al. (2003)

$$\frac{\partial \phi_t^i(\tau)}{\partial \tau} = v_t^i(\tau) \circ \phi_t^i(\tau), \;\; \tau \in [0,1], \tag{2}$$

where τ is an integration variable of the deformation flow.

As a result, the objective function in Eq. 1 is often reparameterized to optimize over the velocity fields instead and can be formulated as

$$E(\{v_t^i\}) = \sum_{t,i=1}^{T,N} \text{Dist}\left(\mathcal{I}_0^i(\phi_t^i(v_t^i)), \mathcal{I}_t^i\right) + \lambda \,\text{Reg}(v_t^i), \;\; \text{s.t. Eq. 2.} \tag{3}$$

3 Our Method

This section introduces BridgeUQ, a novel Bayesian framework for quantifying uncertainty in pretrained registration networks (as shown on the top panel of Fig. 1). The bottom panel of Fig. 1 illustrates an overview of two main components of BridgeUQ derived from the Bayesian model: (i) a hierarchical GP prior via Brownian bridge to model stochastic perturbations around the deterministic registration solution while enforcing fixed endpoints; and (ii) a diffusion variance learning module that adaptively controls the magnitude of the stochastic process.

3.1 Hierarchical GP In Spatiotemporal Velocity Space

Given an image sequence $\mathcal{I}^i$, a pretrained registration network predicts a sequence of velocity fields, $\hat{v}^i \triangleq \{\hat{v}_t^i\}_{t=0}^T \in \mathbb{R}^{|\Omega| \times d \times T}$, where $\mid \Omega \mid$ denotes the number of voxels in the image domain. We model the registration uncertainty by introducing a stochastic perturbation, $\eta^i \triangleq \{\eta_t^i\}_{t=1}^T$. The perturbed velocity field at each time frame, v_t^i, is then defined as $v_t^i = \hat{v}_t^i + \eta_t^i$.

Instead of treating η^i as independent noise, we impose a hierarchical GP via Brownian bridge to capture temporal correlations while enforcing fixed physical boundary conditions. Moreover, in contrast to conventional approaches that treat bridge diffusion variance σ^i as a fixed parameter Wang et al. (2024); Särkkä & Solin (2019), we consider it as a spatially-varying and time-invariant random variable governed by an inverse-Gamma hyper prior to adaptively control the magnitude of stochastic perturbations in the bridge process. This effectively accounts for the heterogeneous nature of

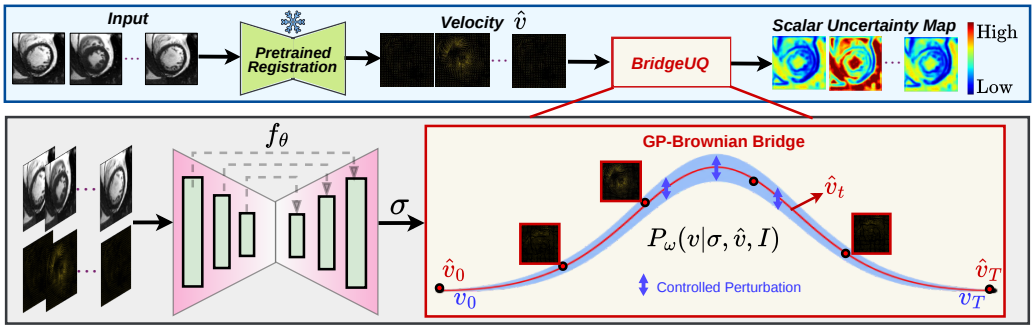


Figure 1: An overview of BridgeUQ integrated with pretrained registration networks (Top); Detailed design of BridgeUQ (Bottom). Once trained, BridgeUQ can be used as a plug-and-play module.

deformable registration, where anatomical regions exhibit varying levels of geometric deformability and registration uncertainty Wang et al. (2019); Luo et al. (2019). The temporal uncertainty evolution is governed by the Brownian bridge process, whereas spatial heterogeneity is encoded by the diffusion variance. Mathematically, we formulate the hierarchical GP on the i-th perturbation sequence as

$$\begin{aligned} \sigma^i &\sim \text{InverseGamma}(\alpha, \beta), \\ \eta_t^i \,|\, \sigma^i &\sim \mathbb{GP}\Big(0,\, \sigma^i\, \Sigma_{\text{BB}}^i(t)\Big), \quad \text{with boundary conditions } \eta_0^i = \eta_T^i = 0, \end{aligned} \tag{4}$$

where, α, β are the shape and scale parameter respectively, and $\Sigma_{\text{BB}}^i(\cdot)$ is the Brownian Bridge kernel. Here, we define the GP prior in continuous time t for notational convenience, while in practice all observations are available only at discrete time points corresponding to the acquired imaging frames.

Periodic Boundary Conditions. Physiologically, we assume the boundary condition defined in Eq. 4 generalizes to a broad class of time-evolved anatomical changes that exhibit progressive, and often approximately cyclic or pseudo-cyclic deformation patterns. In such settings, temporal evolution can be interpreted as a trajectory that departs from an initial anatomical configuration, reaches a peak or final transformation, and then returns (either intrinsically or by construction) to a reference state (see a cardiac example on the bottom panel in Fig. 1). For processes that are not inherently periodic (e.g., brain atrophy), we induce a pseudo-cyclic structure by temporally extending the sequence, for instance, by reversing the observed frames to form a symmetric second half-cycle. This perspective can also be naturally incorporated into pairwise or sequential image registration frameworks by treating the observed evolution between the source and target images as the first half-cycle and augmenting it with a mirrored second phase. Such a formulation establishes well-defined boundary conditions at the temporal endpoints, thereby making the Brownian bridge prior a principled and flexible choice for various applications.

Similar to Li et al. (2023), we define a SDE on the spatiotemporal perturbation (in Eq. 4) as

$$d\eta_t^i = \underbrace{\frac{\eta_T^i - \eta_t^i}{T - t}}_{\text{Drift}}\, dt + \underbrace{\sqrt{\sigma^i}}_{\text{Diffusion}}\, dW_t, \quad \text{with boundary conditions } \eta_0^i = \eta_T^i = 0, \tag{5}$$

where the drift term guides towards the endpoint constraint, and W_t is a multivariate Wiener process. At each location $\mathbf{x}$, the $\sigma^i(\mathbf{x}) \in \mathbb{R}^{d\times d}$ denotes a diagonal covariance matrix.

Proposition 1 *For a given diffusion variance σ^i, the Brownian bridge SDE in Eq. 5 induces a $\mathbb{GP}$ on the spatiotemporal velocity space, $v_t^i \,|\, \sigma^i \sim \mathbb{GP}\big(\hat{v}_t^i,\, \sigma^i \Sigma_{BB}^i(t)\big)$, with **(i)** mean $\mathbb{E}[v_t^i] = \hat{v}_t^i, \forall t \in [0, T]$; **(ii)** covariance $\sigma^i \Sigma_{BB}^i(s,t) = \sigma_j^i \big[\min(s,t) - \frac{st}{T}\big]\delta_{jk}$; for $0 \le s, t \le T$, and $1 \le j, k \le d$, with σ_j^i the j-th diagonal entry of σ^i and δ_{jk} the Kronecker delta; and **(iii)** Markov transition density $v_t^i \,|\, v_{t-1}^i \sim \mathcal{N}\big(\mu_{t|t-1}^i, \delta_{t|t-1}^i\big)$, $\forall t \ge 1$, with*

$$\mu_{t|t-1}^i = \hat{v}_t^i + \frac{T-t}{T-(t-1)}\big(v_{t-1}^i - \hat{v}_{t-1}^i\big), \qquad \delta_{t|t-1}^i = \sigma^i\, \frac{T-t}{T-(t-1)}. \tag{6}$$

The proof is given in Supplementary Sec.1. Given the pretrained registration results $\hat{v}^i$, the conditional probability of the perturbed velocity trajectory, v^i, conditioned on the diffusion variance, can be factorized according to the Markov property as

$$p(v^i \mid \sigma^i; \hat{v}^i) \;=\; \prod_{t=1}^{T} p\big(v_t^i \mid\; v_{t-1}^i, \sigma^i; \hat{v}^i\big)\,, \tag{7}$$

where each transition $p(v_t^i \mid v_{t-1}^i, \sigma^i; \hat{v}^i)$ is a Gaussian density given in Eq. 6.

3.2 Bayesian Model For Aleatoric Uncertainty Quantification

Energy-based likelihood. To ensure that the stochastic perturbations induced by our hierarchical GP remain consistent with the original registration objective (Eq. 3), we define an energy-based data likelihood derived from the deterministic registration loss as

$$p(\mathcal{I}^i \mid v^i) \propto \exp\!\left(-\frac{1}{\gamma}\, E(v^i)\right), \tag{8}$$

where $\gamma > 0$ is a temperature parameter controlling the concentration of the distribution.

Combining Eq. 4, 7, and 8, we arrive at a joint posterior over the spatiotemporal velocity and diffusion variance using Baye's rule

$$p(\mathbf{v}, \boldsymbol{\sigma} \mid \mathcal{I}, \hat{\mathbf{v}}) \triangleq \prod_{i=1}^{N} p(v^i, \sigma^i \mid \mathcal{I}^i, \hat{v}^i) \propto \prod_{i=1}^{N} p(\mathcal{I}^i \mid v^i)\; p(v^i \mid \sigma^i; \hat{v}^i)\; p(\sigma^i). \tag{9}$$

After taking the logarithm of Eq. 9 and expanding each term, we have

$$\begin{aligned}\log p(\mathbf{v}, \boldsymbol{\sigma} \mid \mathcal{I}, \hat{\mathbf{v}}) \propto -\frac{1}{\gamma} E(\{v_t^i\}) + \sum_{i,t=1}^{N,T-1} \Big[-\frac{1}{2}(v_t^i - \mu_{t|t-1}^i)^\top (\delta_{t|t-1}^i)^{-1} (v_t^i - \mu_{t|t-1}^i) \\ - \frac{1}{2}\log\det(\delta_{t|t-1}^i)\Big] + \sum_{i=1}^{N} \left[-\beta\, \mathbf{1}^\top (\mathrm{vec}(\sigma^i))^{-1} - (\alpha+1)\, \mathbf{1}^\top \log(\mathrm{vec}(\sigma^i))\right], \end{aligned} \tag{10}$$

where $\mathbf{1} \in \mathbb{R}^{|\Omega| \times d}$ denotes an all-ones vector, and $\mathrm{vec}(\cdot)$ denotes a vectorization.

To address the intractable inference of Eq. 10 due to the coupling between $\mathbf{v}$ and $\boldsymbol{\sigma}$, we adopt an iterative procedure Vidal et al. (2020) that decomposes the problem into two alternating subproblems: (i) MAP estimation of the diffusion variance $\boldsymbol{\sigma}$, and (ii) Langevin sampling Grzech et al. (2021) that uses the estimated variance to sample from the posterior distribution of velocities for uncertainty estimation. Specifically, at each iteration r, we perform

- Update $\boldsymbol{\sigma}$: $\boldsymbol{\sigma}^{r+1} = \arg\max\; \log p(\boldsymbol{\sigma} | \mathbf{v}^r, \mathcal{I}, \hat{\mathbf{v}})$;
- Sample $\mathbf{v}^{r+1}$ from the posterior distribution $p(\mathbf{v}^{r+1} \mid \boldsymbol{\sigma}^{r+1}, \mathcal{I}, \hat{\mathbf{v}})$.

For notational simplicity, we omit the index r in the following sections. For MAP inference, we derive the gradient with respect to σ^i for each input image sequence as

$$\nabla_{\sigma^i} \log p(\boldsymbol{\sigma} | \mathbf{v}, \mathcal{I}, \hat{\mathbf{v}}) = -\frac{1}{\gamma} \nabla_{\sigma^i} E(\{v_t^i\}) + \frac{1}{2{\sigma^i}^T \sigma^i} \sum_{t=1}^{T-1} \frac{1}{a_t} (v_t^i - \mu_{t|t-1}^i)^2 + \frac{1}{{\sigma^i}^T \sigma^i} (\beta\, \mathbf{I} - \tilde{\alpha}\, \sigma^i), \tag{11}$$

where $(\cdot)^2$ denotes an outer product, $\mathbf{I}$ is the identity matrix, $a_t \triangleq \frac{T-t}{T-t+1}$, and $\tilde{\alpha} \triangleq \alpha + \frac{T+1}{2}$.

Due to space limitations, we present the key computation of Langevin sampling, which is the posterior gradient w.r.t. v_t^i. Note that the gradient of registration loss is provided by the pre-trained registration network. Details of the entire sampling procedure are provided in Supplementary Sec. 2.

$$\nabla_{v_t^i} \log p(\mathbf{v} \mid \boldsymbol{\sigma}, \mathcal{I}, \hat{\mathbf{v}}) = -\frac{1}{\gamma}\, \nabla_{v_t^i} E\big(\{v_t^i\}\big) + \big(\sigma^i\big)^{-1} \Big[\eta_{t+1}^i - 2\,\eta_t^i + \eta_{t-1}^i\Big];\; \eta_t^i = v_t^i - \hat{v}_t^i. \tag{12}$$

3.2.1 Inference Amortization

To speed up the inference, we further amortize the iterative MAP optimization and posterior sampling by adopting the amortization strategy described in Amos (2023). The converged MAP estimates of $\boldsymbol{\sigma}^\star$, along with a set of sampled spatiotemporal velocity fields, $\{\mathbf{v}^{(s)}\}_{s=1}^{S}$, generated by the Langevin sampler (summarized in Alg. 1.1), are used as training targets for our proposed amortization process.

Amortized MAP. We amortize the estimation of diffusion variance by learning a parameterized network, f_θ, to predict the MAP estimate without requiring iterative optimization, i.e., $f_\theta(\mathcal{I}^i, \hat{v}^i) \approx \arg\max_{\boldsymbol{\sigma}} p(\boldsymbol{\sigma} \mid \mathbf{v}, \mathcal{I}, \hat{\mathbf{v}})$. An objective function, $\mathcal{L}(\theta) = \frac{1}{N}\sum_{i=1}^{N} \left\| f_\theta\big(\mathcal{I}^i, \hat{v}^i\big) - \sigma^{i,\star} \right\|_2^2$, is then employed to train the network until convergence. In practice, we implement f_θ using a U-Net architecture Dalca et al. (2019a). To enforce the positiveness of the learned diffusion variance, we apply a softplus activation Dugas et al. (2000) to the output layer.

Amortized Langevin sampling. Similarly, we amortize the iterative Langevin sampling procedure by distilling its posterior samples into a conditional diffusion model Song et al. (2020). Given the set of sampled temporal velocities, $\{\mathbf{v}^{(s)}\}_{s=1}^{S}$, we train a conditional diffusion model that learns to approximate the underlying posterior distribution, $p(\mathbf{v} \mid \boldsymbol{\sigma}, \mathcal{I}, \hat{\mathbf{v}})$, conditioned on the diffusion variance and observed image sequences with predicted registration. Once trained, the diffusion model is used to generate velocity samples directly, rather than performing the computationally more expensive Langevin sampling at inference. We employ a standard diffusion loss $\mathcal{L}(\omega) = \mathbb{E}_{m,\epsilon}\left\| \epsilon - \epsilon_\omega\big(v_m^{i(s)}, m, \sigma^{i,\star}, \mathcal{I}^i, \hat{v}^i\big) \right\|_2^2$, where $v_m^{i(s)}$ is the noisy velocity field at diffusion timestep m.

The overall amortized training and inference are summarized in Alg. 1.2 and Alg. 2. Note that amortization introduces additional training and requires a sufficiently large and representative training set to learn the mapping from the observed inputs to the underlying posterior distribution. Therefore, when training data are limited or substantially different from the target data distribution, our non-amortized BridgeUQ formulation provides an alternative.

4 Experimental Evaluation

Datasets. We validate BridgeUQ on real-world medical imaging datasets. Since ground-truth registration results are not available in these settings, we additionally evaluate our method on synthetic datasets, where ground-truth deformations are known. *(i) Synthetic data.* We construct synthetic deformation trajectories by sampling random velocity fields from a smooth gaussian noise $\mathcal{N}(0, 0.01)$, and then warping a reference template image to simulate the time sequence Wu & Zhang (2023). The spatial dimension of the synthesized dataset is 64×64. Examples are shown in Supplementary. *(ii) Real 2D cardiac cine MRI videos.* We combine two cardiac datasets: a multisite Cardiac MRI Ghadimi et al. (2024); Wu et al. (2024a); Deb et al. (2025) (741 sequences from 284 subjects) and ACDC Bernard et al. (2018)(150 cine MR sequences). All cardiac images were spatially cropped to 128×128, and temporally resampled to 40 frames per cardiac cycle. *(iii) Longitudinal Brain MRI.* We use the ADNI Mueller et al. (2005) dataset comprising 320 subjects with longitudinal follow-up scans at four time points. All brain MRIs are skull-stripped, affine registered, and resampled to a resolution of 128^3 with $1.0\,\text{mm}^3$ isotropic voxel spacing.

Baselines & Registration Backbones. We compare the performance of BridgeUQ against state-of-the-art methods for aleatoric uncertainty in image registration, including Rbr Hu et al. (2025), PULPo Siegert et al. (2023), DIF-VM Dalca et al. (2019b), SGMCMC Grzech et al. (2021), and Wang

Algorithm 1.1 BridgeUQ Optimization

Input: $(\mathcal{I}, \hat{\mathbf{v}})$ pairs; R iterations; S samples
for each $(\mathcal{I}, \hat{\mathbf{v}})$ pair **do**
  $\mathbf{v}^0 \leftarrow \hat{\mathbf{v}}$
  $\boldsymbol{\sigma}^0 \sim \mathcal{N}(0, \mathbf{I})$
  for $r = 0, \ldots, R$ **do**
    Update $\boldsymbol{\sigma}^{r+1}$ with gradient ascent using Eq. 11
    sample $\{\mathbf{v}^{(s)}\}_{s=1}^{S}$ via Langevin sampling
    $\mathbf{v}^{r+1} \leftarrow \{\mathbf{v}^{(s)}\}_{s=1}^{N}$
  end for
end for
return $\boldsymbol{\sigma}^\star \leftarrow \boldsymbol{\sigma}^R$; $\{\mathbf{v}^{(s)}\}_{s=1}^{S}$

Algorithm 1.2 Amortized Training

Input: Training pairs $(\mathcal{I}, \hat{\mathbf{v}})$; $\boldsymbol{\sigma}^\star$, $\{\mathbf{v}^{(s)}\}_{s=1}^{S}$; ϵ_ω; $f_{\boldsymbol{\theta}}$
Train $f_{\boldsymbol{\theta}}(\mathcal{I}, \hat{\mathbf{v}})$ using $\boldsymbol{\sigma}^\star$ as ground truth
Train ϵ_ω on $\{\mathbf{v}^{(s)}\}_{s=1}^{S}$ conditioned on $(\boldsymbol{\sigma}^\star, \mathcal{I}, \hat{\mathbf{v}})$
return $f_{\boldsymbol{\theta}}$, ϵ_ω

Algorithm 2 Amortized Inference

Input: Trained $f_{\boldsymbol{\theta}}, \epsilon_\omega$; test sequence $(\mathcal{I}, \hat{\mathbf{v}})$; S samples
$\boldsymbol{\sigma} \leftarrow f_{\boldsymbol{\theta}}(\mathcal{I}, \hat{\mathbf{v}})$
$\mathbf{v}^{(s)} \leftarrow \text{DDIM-Sampler}\big(\mathbf{z}^{(s)}, \boldsymbol{\sigma}, \mathcal{I}, \hat{\mathbf{v}}\big)$ # $\mathbf{z}^{(s)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
return $\bar{\mathbf{v}}_t \leftarrow \mathbb{E}[\boldsymbol{v}_t^{(n)}]$; $\boldsymbol{u}_t \leftarrow \text{Var}[\boldsymbol{v}_t^{(n)}]^{1/2}, \forall t$

et al. Wang et al. (2019), all of which approximate the posterior over deformation fields or parameters. To demonstrate the flexibility and plug-and-play nature of our approach, we integrate BridgeUQ with four pretrained registration backbones that includes pairwise registration - VoxelMorph (VM) Balakrishnan et al. (2019) and TransMorph (TM) Chen et al. (2022), as well as spatiotemporal registration - TLRN Wu et al. (2024a), and LTMA Deb et al. (2025). For fair comparison, all experiments employ the mean squared error (MSE) as the similarity metric for registration loss.

Experiments & Evaluation Metrics. We first evaluate BridgeUQ on synthetic datasets, where ground-truth deformations are available to compute voxel-wise registration errors. To assess the quality of aleatoric uncertainty, we adopt ranking-based metrics, such as the normalized Area Under the Sparsification Curve (nAUSC), to measure how registration error changes as voxels are progressively removed according to increasing predicted uncertainty (i.e., from least to most certain) Chen et al. (2025); Wannenwetsch et al. (2017); Ilg et al. (2018). For reference, we compute an oracle curve by sorting voxels using the voxel-wise registration error between the predicted and ground-truth deformation fields, and progressively removing voxels with the largest errors. The nAUSC value quantifies the discrepancy between the predicted-uncertainty and oracle curve.

We further evaluate the nAUSC metric on real-world cardiac and brain datasets. Note that as ground-truth deformation fields are not available in these settings, we adopt a segmentation-based proxy for uncertainty evaluation Chen et al. (2024). Specifically, we propagate manual segmentations at the reference frame through the predicted registration and compare the resulting deformed segmentation sequences at each voxel against ground-truth annotations. This allows us to indirectly assess uncertainty by measuring how segmentation errors correlate with the predicted uncertainty map.

To examine the quality of the sampled velocity fields (deformation trajectories), we report two standard metrics. First, we compute the Dice Similarity Coefficient (DSC) Dice (1945), which measures the overlap between propagated reference segmentations via sampled deformation fields and ground-truth labels. Second, we quantify the physical plausibility of the predicted deformations by reporting the percentage of voxels with negative Jacobian determinants Dalca et al. (2019b); Wu & Zhang (2023), i.e., $|\det(J_\phi) < 0|\,(\%)$. This metric captures non-physical artifacts such as folding, tearing, or topology violations. in the deformation field. Additional implementation details are provided in the Supplementary Sec. 3. Furthermore, to evaluate our BridgeUQ on substantially higher dimensional real-world data, we assess the computational efficiency of the proposed amortized inference (Alg. 2) and compare it with the non-amortized optimization (Alg. 1.1).

4.1 Overall Results

Fig. 2 shows a comparison of sparsification curve across multiple time points between all methods on the synthetic dataset. The sparsification curve of BridgeUQ is monotonically decreasing, which indicates that regions with higher predicted uncertainty consistently correspond to larger registration errors. Additionally, our sparsification curve follows most closely to the oracle curve than all competing methods. This suggests that BridgeUQ provides a more reliable uncertainty ranking, and is more effective at identifying and filtering unreliable predictions.

Tab. 1 compares the computational efficiency of the proposed amortized inference strategy with the non-amortized optimization on real-world data. Results show that amortized inference achieves comparable performance to the non-amortized approach while substantially reducing computational cost. In particular, it preserves the quality of the sampled velocity fields, as reflected by similar Dice scores and determinant of Jacobian, while maintaining comparable uncertainty quantification

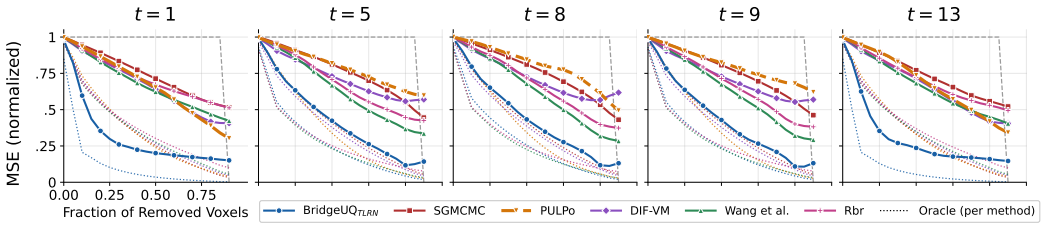


Figure 2: A comparison of sparsification curve between all methods on the synthetic dataset.

performance as measured by nAUSC. For space conciseness, all subsequent BridgeUQ results are obtained by the proposed amortized inference strategy.

Tab. 2 shows that BridgeUQ achieves the lowest temporally averaged nAUSC on both brain and cardiac datasets across all four pretrained registration backbones. In addition, the DSC and the percentage of negative Jacobian determinants further indicate that both registration accuracy and the diffeomorphic properties of deformation trajectories sampled from our model are well preserved.

Fig. 3 visualizes representative cardiac imaging examples of the scalar uncertainty maps across different time frames (additional brain imaging results are provided in the Supplementary). Overall, BridgeUQ produces spatially meaningful uncertainty patterns that are consistent with the underlying image information. Specifically, higher uncertainty is observed in homogeneous regions, where multiple plausible velocity fields can produce similarly good image alignments. In contrast, uncertainty is substantially lower near anatomical boundaries and regions with sharp structural variations, where the observed image structures provide stronger constraints on the underlying deformation.

Ablation Study. We conduct ablation studies to evaluate two key components of BridgeUQ: (i) the proposed Brownian bridge prior for modeling spatiotemporal uncertainty and (ii) learning the diffusion variance. Fig. 4 (a) compares different temporal priors (Brownian motion and Gaussian prior) by measuring the Spearman correlation Wannenwetsch et al. (2017) between registration loss and the per-frame variance induced by each temporal kernel. Although the spatial correlation across each voxels within a frame is similar, they significantly differ along the temporal trajectory. Our BridgeUQ closely tracks the temporal trajectory of registration results, achieving a strong positive Spearman correlation of 0.85. In contrast, both the Brownian motion and Gaussian priors exhibit negative correlations of -0.45 and -0.47, respectively, indicating that their induced temporal variance does not align with the observed registration uncertainty. Fig. 4 (b) compares the learned diffusion variance with a randomly initialized variance field and their corresponding nAUSC values.

Table 1: Non-amortized (w/o amort.) vs. amortized (w/ amort.) inference for BridgeUQ$_{\text{TLRN}}$; time is per subject. nAUSC is averaged across all frames.

	Cardiac Cine MRI (T=15)				Brain MRI – ADNI (T=7)			
	nAUSC(Avg) ↓	DSC ↑	$\lvert J_\varphi\rvert\leq 0$ ↓	Time	nAUSC(Avg) ↓	DSC ↑	$\lvert J_\varphi\rvert\leq 0$ ↓	Time
w/o amort.	0.545	0.847	0.002	6.4 min.	0.310	0.901	0.000	4.5 min.
w/ amort.	0.531	0.843	0.002	53 sec.	0.308	0.895	0.000	46.5 sec.

Table 2: A comparison of nAUSC at selected frames and its average across all the frames (Avg.), DSC (frame-averaged), the percentage of negative Jacobian determinant ($|J_\varphi|\leq 0$, %), and inference time per subject. Best results are in **bold** and the second best are underlined.

Cardiac Cine MRI (T=15)

Method	nAUSC at frame ↓					Avg. ↓	Avg. DSC↑	Avg. $\lvert J_\varphi\rvert\leq 0$ ↓	Time
	t=2	t=4	t=8	t=10	t=12				
Learning-Based Probabilistic registration									
DIF-VM (2019b)	0.673±0.022	0.884±0.310	0.820±0.174	0.807±0.101	0.721±0.114	0.807±0.204	0.830±0.077	0.001±0.010	1.4 sec
PULPo (2023)	0.551±0.135	0.573±0.128	0.601±0.146	0.546±0.137	0.580±0.131	0.606±0.104	0.834±0.064	0.005±0.022	3.0 sec
Rbr (2025)	0.573±0.216	**0.440±0.174**	0.516±0.195	0.624±0.203	0.705±0.222	0.599±0.156	0.824±0.085	0.003±0.001	1.4 sec
Classical Bayesian Registration									
SGMCMC (2021)	0.557±0.221	0.768±0.110	0.821±0.170	0.709±0.184	0.586±0.220	0.662±0.110	0.808±0.046	0.000±0.000	8.0 min.
Wang et al. (2019)	0.525±0.520	0.698±0.372	0.581±0.315	0.587±0.193	0.578±0.148	0.592±0.261	0.820±0.081	0.000±0.000	26.2 min
BridgeUQ (Ours, Plug-and-Play on Pretrained Backbones)									
BridgeUQ$_{\text{VM}}$	0.574±0.133	0.585±0.139	0.593±0.144	0.575±0.134	0.562±0.137	0.574±0.121	0.833±0.095	0.001±0.007	53.0 sec.
BridgeUQ$_{\text{TM}}$	0.547±0.120	0.563±0.140	0.551±0.130	0.562±0.110	0.547±0.120	0.562±0.110	0.838±0.087	0.000±0.000	53.0 sec.
BridgeUQ$_{\text{TLRN}}$	**0.522±0.119**	0.563±0.125	**0.549±0.136**	**0.531±0.125**	**0.536±0.129**	**0.531±0.114**	**0.843±0.062**	0.002±0.001	53.0 sec.
BridgeUQ$_{\text{LTMA}}$	0.540±0.122	0.569±0.127	0.557±0.135	0.542±0.127	0.546±0.127	0.554±0.114	0.844±0.080	0.000±0.000	53.0 sec.

Brain MRI – ADNI (T=7)

Method	nAUSC at frame ↓					Avg. ↓	Avg. DSC ↑	Avg. $\lvert J_\varphi\rvert\leq 0$ ↓	Time
	t=1	t=2	t=3	t=4	t=5				
Learning-Based Probabilistic registration									
DIF-VM (2019b)	0.335±0.071	0.398±0.112	0.471±0.162	0.398±0.112	0.335±0.071	0.387±0.088	0.873±0.107	0.000±0.000	0.85 sec
PULPo (2023)	0.360±0.105	0.373±0.170	0.365±0.110	0.367±0.104	0.360±0.105	0.365±0.080	0.858±0.131	0.004±0.001	1.90 sec
Rbr (2025)	0.348±0.060	0.362±0.063	0.364±0.066	0.362±0.063	0.348±0.068	0.352±0.013	0.881±0.124	0.001±0.003	0.85 sec
Classical Bayesian Registration									
SGMCMC (2021)	0.310±0.107	0.377±0.151	0.466±0.195	0.374±0.148	0.306±0.098	0.367±0.155	0.877±0.139	0.000±0.000	2.5 min
Wang et al. (2019)	0.331±0.221	0.341±0.240	0.317±0.201	0.341±0.240	0.331±0.221	0.332±0.228	0.883±0.101	0.000±0.000	15.2 min
BridgeUQ (Ours, Plug-and-Play on Pretrained Backbones)									
BridgeUQ$_{\text{VM}}$	0.327±0.070	0.334±0.068	0.327±0.065	0.334±0.068	0.327±0.070	0.330±0.046	0.878±0.122	0.000±0.000	46.5 sec.
BridgeUQ$_{\text{TM}}$	0.303±0.108	0.317±0.117	0.334±0.098	0.317±0.119	0.303±0.108	0.311±0.103	0.886±0.109	0.001±0.008	46.5 sec.
BridgeUQ$_{\text{TLRN}}$	**0.297±0.073**	**0.308±0.069**	0.332±0.065	**0.308±0.069**	**0.297±0.073**	**0.308±0.049**	**0.895±0.106**	0.000±0.000	46.5 sec.
BridgeUQ$_{\text{LTMA}}$	0.312±0.071	0.318±0.069	**0.320±0.066**	0.318±0.069	0.312±0.071	0.314±0.047	0.893±0.108	0.000±0.000	46.5 sec.

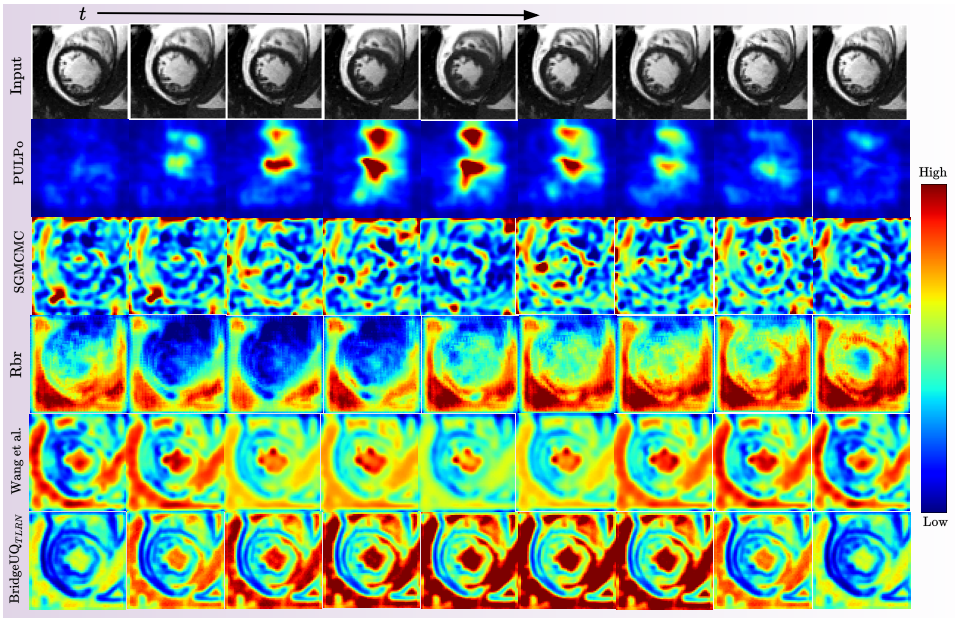


Figure 3: Examples of the predicted uncertainty maps across the cardiac cycle, showing lower uncertainty along anatomical boundaries and higher uncertainty along homogeneous regions.

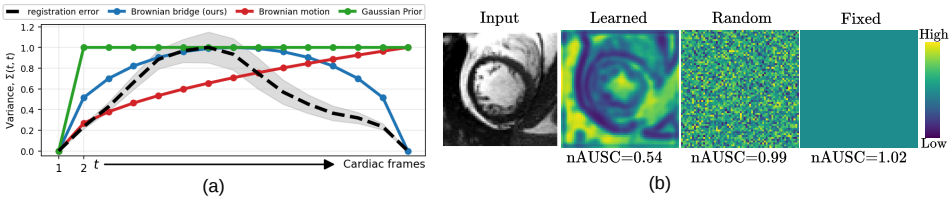


Figure 4: Ablation on (a) Brownian bridge prior for spatiotemporal modeling (with shaded line represents standard deviation); and (b) learnable diffusion variance visualized by its standard deviation.

The learned variance produces spatially informative uncertainty estimates, whereas independent random voxel-wise perturbations fail to capture the spatial structure of registration uncertainty, resulting in substantially less informative uncertainty estimates.

5 Conclusion & Discussion

This paper presents BridgeUQ, a Bayesian framework for quantifying aleatoric uncertainty in pretrained time-series medical image registration. We are the first to introduce a hierarchical GP prior defined through a Brownian bridge in spatiotemporal velocity space with a learnable diffusion variance. To ensure that the inferred uncertainty remains consistent with the observed image sequence, we incorporate an energy-based data likelihood that couples the stochastic velocity field to the pretrained registration model. We further develop an amortized inference strategy to substantially reduce the computational cost of Bayesian inference, enabling efficient uncertainty estimation for high-dimensional real-world medical images.

Limitations. Our current framework focuses exclusively on aleatoric uncertainty coming from ambiguity in the pretrained registration process. It therefore does not explicitly capture epistemic uncertainty associated with uncertainty in the network parameters or model specification. A natural direction for future work is to extend BridgeUQ to jointly characterize aleatoric and epistemic uncertainty, for example, by incorporating Bayesian approximations such as MC dropout or model ensembles into the registration network.

REFERENCES

Timur Aksoy, Gozde Unal, Stefanie Demirci, Nassir Navab, and Muzaffer Degertekin. Template-based cta to x-ray angio rigid registration of coronary arteries in frequency domain with automatic x-ray segmentation. *Medical physics*, 40(10):101903, 2013.

Brandon Amos. Tutorial on amortized optimization. *Foundations and Trends in Machine Learning*, 16(5):592–732, 2023.

Vladimir I Arnold. Sur un principe variationnel pour les écoulements stationnaires des liquides parfaits et ses applications aux problemes de stabilité non linéaires. *Journal de mécanique*, 5(1): 29, 1966.

Vincent Arsigny, Olivier Commowick, Xavier Pennec, and Nicholas Ayache. A log-euclidean framework for statistics on diffeomorphisms. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 924–931. Springer, 2006.

Brian B Avants, Charles L Epstein, Murray Grossman, and James C Gee. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis*, 12(1):26–41, 2008.

Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019.

Michael Becker, Rainer Hoffmann, Harald P Kühl, Helena Grawe, Markus Katoh, Rafael Kramann, Arno Bücker, Peter Hanrath, and Nicole Heussen. Analysis of myocardial deformation based on ultrasonic pixel tracking to determine transmurality in chronic myocardial infarction. *European heart journal*, 27(21):2560–2566, 2006.

M Faisal Beg, Michael I Miller, Alain Trouvé, and Laurent Younes. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision*, 61(2):139–157, 2005.

Olivier Bernard, Alain Lalande, Clement Zotti, Frederick Cervenansky, Xin Yang, Pheng-Ann Heng, Irem Cetin, Karim Lekadir, Oscar Camara, Miguel Angel Gonzalez Ballester, et al. Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging*, 37(11):2514–2525, 2018.

Junyu Chen, Eric C Frey, Yufan He, William P Segars, Ye Li, and Yong Du. Transmorph: Transformer for unsupervised medical image registration. *Medical image analysis*, 82:102615, 2022.

Junyu Chen, Yihao Liu, Shuwen Wei, Zhangxing Bian, Aaron Carass, and Yong Du. From registration uncertainty to segmentation uncertainty. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1–5. IEEE, 2024.

Junyu Chen, Yihao Liu, Shuwen Wei, Zhangxing Bian, Shalini Subramanian, Aaron Carass, Jerry L Prince, and Yong Du. A survey on deep learning in medical image registration: New technologies, uncertainty, evaluation metrics, and beyond. *Medical Image Analysis*, 100:103385, 2025.

Adrian Dalca, Marianne Rakic, John Guttag, and Mert Sabuncu. Learning conditional deformable templates with convolutional networks. *Advances in neural information processing systems*, 32, 2019a.

Adrian V Dalca, Guha Balakrishnan, John Guttag, and Mert R Sabuncu. Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. *Medical image analysis*, 57: 226–236, 2019b.

Christos Davatzikos, Dinggang Shen, Ashraf Mohamed, and Stelios K Kyriacou. A framework for predictive modeling of anatomical deformations. *IEEE transactions on medical imaging*, 20(8): 836–843, 2001.

Swakshar Deb, Nian Wu, Frederick H Epstein, and Miaomiao Zhang. Unsupervised cardiac video translation via motion feature guided diffusion model. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 648–658, 2025.

Lee R Dice. Measures of the amount of ecologic association between species. *Ecology*, 26(3): 297–302, 1945.

Charles Dugas, Yoshua Bengio, François Bélisle, Claude Nadeau, and René Garcia. Incorporating second-order functional knowledge for better option pricing. *Advances in neural information processing systems*, 13, 2000.

Jean-Baptiste Fiot, Hugo Raguet, Laurent Risser, Laurent D Cohen, Jurgen Fripp, François-Xavier Vialard, Alzheimer's Disease Neuroimaging Initiative, et al. Longitudinal deformation models, spatial regularizations and learning strategies to quantify alzheimer's disease progression. *NeuroImage: Clinical*, 4:718–729, 2014.

Ziyu Gan, Wei Sun, Kaimin Liao, and Xuan Yang. Probabilistic modeling for image registration using radial basis functions: Application to cardiac motion estimation. *IEEE Transactions on Neural Networks and Learning Systems*, 34(10):7324–7338, 2022.

SONA Ghadimi, Derek Bivona, Kenneth Bilchick, and Frederick Epstein. Deep learning-based prognostic model using cine dense mri for outcome prediction after cardiac resynchronization therapy. *Journal of Cardiovascular Magnetic Resonance*, 26, 2024.

X Gong, L Khaidem, W Zhu, et al. Uncertainty learning towards unsupervised deformable medical image registration 2022 ieee. In *CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 1555–64, 2022.

Daniel Grzech, Mohammad Farid Azampour, Huaqi Qiu, Ben Glocker, Bernhard Kainz, Loïc Le Folgoc, et al. Uncertainty quantification in non-rigid image registration via stochastic gradient markov chain monte carlo. *Machine Learning for Biomedical Imaging*, 1(UNSURE2020 special issue):1–25, 2021.

Tonmoy Hossain and Miaomiao Zhang. Corld: Contrastive representation learning of deformable shapes in images. In *International Conference on Information Processing in Medical Imaging*, pp. 342–357. Springer, 2025.

Xiaoling Hu, Karthik Gopinath, Peirong Liu, Malte Hoffmann, Koen Van Leemput, Oula Puonti, and Juan Eugenio Iglesias. Hierarchical uncertainty estimation for learning-based registration in neuroimaging. In *... International Conference on Learning Representations*, volume 2025, pp. 91200, 2025.

Martina Hub, Marc L Kessler, and Christian P Karger. A stochastic approach to estimate the uncertaintyinvolved in b-spline image registration. *IEEE Transactions on Medical Imaging*, 28 (11):1708–1716, 2009.

Eddy Ilg, Ozgun Cicek, Silvio Galesso, Aaron Klein, Osama Makansi, Frank Hutter, and Thomas Brox. Uncertainty estimates and multi-hypotheses networks for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 652–667, 2018.

Firdaus Janoos, Petter Risholm, and William Wells. Robust non-rigid registration and characterization of uncertainty. In *2012 IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 4321–4326. IEEE, 2012.

Nivetha Jayakumar, Jiarui Xing, Tonmoy Hossain, Fred Epstein, Kenneth Bilchick, and Miaomiao Zhang. Activation from sparse 2d cardiac mris. In *Machine Learning for Health (ML4H)*, pp. 190–200. PMLR, 2023.

Julian Krebs, Hervé Delingette, Boris Mailhé, Nicholas Ayache, and Tommaso Mansi. Learning a probabilistic model for diffeomorphic registration. *IEEE transactions on medical imaging*, 38(9): 2165–2176, 2019.

Julian Krebs, Hervé Delingette, Nicholas Ayache, and Tommaso Mansi. Learning a generative motion model from image sequences based on a latent motion matrix. *IEEE Transactions on Medical Imaging*, 40(5):1405–1416, 2021.

Jan Kybic. Bootstrap resampling for image registration uncertainty estimation without ground truth. *IEEE Transactions on Image Processing*, 19(1):64–73, 2009.

Loic Le Folgoc, Herve Delingette, Antonio Criminisi, and Nicholas Ayache. Quantifying registration uncertainty with sparse bayesian modelling. *IEEE transactions on medical imaging*, 36(2):607–617, 2016.

Loïc Le Folgoc, Hervé Delingette, Antonio Criminisi, and Nicholas Ayache. Sparse bayesian registration of medical images for self-tuning of parameters and spatially adaptive parametrization of displacements. *Medical image analysis*, 36:79–97, 2017.

Bo Li, Kaitao Xue, Bin Liu, and Yu-Kun Lai. Bbdm: Image-to-image translation with brownian bridge diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern Recognition*, pp. 1952–1961, 2023.

Tiantian Li, Mengxi Zhang, Wenyuan Qi, Evren Asma, and Jinyi Qi. Motion correction of respiratory-gated pet images using deep learning based image registration framework. *Physics in Medicine & Biology*, 65(15):155003, 2020.

Jiayi Lu, Renchao Jin, Manyang Wang, Enmin Song, and Guangzhi Ma. A bidirectional registration neural network for cardiac motion tracking using cine mri images. *Computers in Biology and Medicine*, 160:107001, 2023.

Jie Luo, Alireza Sedghi, Karteek Popuri, Dana Cobzas, Miaomiao Zhang, Frank Preiswerk, Matthew Toews, Alexandra Golby, Masashi Sugiyama, William M Wells III, et al. On the applicability of registration uncertainty. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 410–419. Springer, 2019.

Lucas Mansilla, Diego H Milone, and Enzo Ferrante. Learning deformable registration of medical images with anatomical constraints. *Neural Networks*, 124:269–279, 2020.

Stephen Marsland and Tony Shardlow. Langevin equations for landmark image registration with uncertainty. *SIAM Journal on Imaging Sciences*, 10(2):782–807, 2017.

Michael I Miller, Alain Trouvé, and Laurent Younes. The metric spaces, euler equations, and normal geodesic image motions of computational anatomy. In *International Conference on Image Processing*, volume 2, pp. II–635. IEEE, 2003.

Susanne G Mueller, Michael W Weiner, Leon J Thal, Ronald C Petersen, Clifford Jack, William Jagust, John Q Trojanowski, Arthur W Toga, and Laurel Beckett. The alzheimer's disease neuroimaging initiative. *Neuroimaging Clinics*, 15(4):869–877, 2005.

Vladimir Nekrasov, Chunhua Shen, and Ian Reid. Template-based automatic search of compact semantic segmentation architectures. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1980–1989, 2020.

Jiazhen Pan, Daniel Rueckert, Thomas Küstner, and Kerstin Hammernik. Efficient image registration network for non-rigid cardiac motion estimation. In *International Workshop on Machine Learning for Medical Image Reconstruction*, pp. 14–24. Springer, 2021.

Xavier Pennec, Stefan Sommer, and Tom Fletcher. *Riemannian geometric statistics in medical image analysis*. Academic Press, 2019.

Loïc Peter, Daniel C Alexander, Caroline Magnain, and Juan Eugenio Iglesias. Uncertainty-aware annotation protocol to evaluate deformable registration algorithms. *IEEE Transactions on Medical Imaging*, 40(8):2053–2065, 2021.

Petter Risholm, Firdaus Janoos, Isaiah Norton, Alex J Golby, and William M Wells III. Bayesian characterization of uncertainty in intra-subject non-rigid registration. *Medical image analysis*, 17 (5):538–555, 2013.

Daniel Rueckert, Paul Aljabar, Rolf A Heckemann, Joseph V Hajnal, and Alexander Hammers. Diffeomorphic registration using b-splines. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 702–709. Springer, 2006.

Simo Särkkä and Arno Solin. *Applied stochastic differential equations*, volume 10. Cambridge University Press, 2019.

Alireza Sedghi, Tina Kapur, Jie Luo, Parvin Mousavi, and William M Wells III. Probabilistic image registration via deep multi-class classification: characterizing uncertainty. In *Workshop on Clinical Image-Based Procedures*, pp. 12–22. Springer, 2019.

Leonard Siegert, Paul Fischer, Mattias P Heinrich, and Christian F Baumgartner. PULPo: Probabilistic unsupervised Laplacian pyramid registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 717–727. Springer, 2023.

Ivor JA Simpson, Mark W Woolrich, Manuel Jorge Cardoso, David M Cash, Marc Modat, Julia A Schnabel, and Sebastien Ourselin. A bayesian approach for spatially adaptive regularisation in non-rigid registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 10–18. Springer, 2013.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

Ana Fernandez Vidal, Valentin De Bortoli, Marcelo Pereyra, and Alain Durmus. Maximum likelihood estimation of regularization parameters in high-dimensional inverse problems: An empirical bayesian approach part i: Methodology and experiments. *SIAM Journal on Imaging Sciences*, 13 (4):1945–1989, 2020.

Jian Wang, William M Wells III, Polina Golland, and Miaomiao Zhang. Efficient laplace approximation for bayesian registration uncertainty quantification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 880–888. Springer, 2018.

Jian Wang, William M Wells III, Polina Golland, and Miaomiao Zhang. Registration uncertainty quantification via low-dimensional characterization of geometric deformations. *Magnetic resonance imaging*, 64:122–131, 2019.

Peiyong Wang, Bohan Xiao, Qisheng He, Carri Glide-Hurst, and Ming Dong. Score-based image-to-image brownian bridge. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 10765–10773, 2024.

Anne S Wannenwetsch, Margret Keuper, and Stefan Roth. Probflow: Joint optical flow and uncertainty estimation. In *Proceedings of the IEEE international conference on computer vision*, pp. 1173–1182, 2017.

Demian Wassermann, Matthew Toews, Marc Niethammer, and William Wells Iii. Probabilistic diffeomorphic registration: Representing uncertainty. In *International Workshop on Biomedical Image Registration*, pp. 72–82. Springer, 2014.

William M Wells III, Paul Viola, Hideki Atsumi, Shin Nakajima, and Ron Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical image analysis*, 1(1):35–51, 1996.

Nian Wu and Miaomiao Zhang. Neurepdiff: Neural operators to predict geodesics in deformation spaces. In *International Conference on Information Processing in Medical Imaging*, pp. 588–600. Springer, 2023.

Nian Wu, Jiarui Xing, and Miaomiao Zhang. Tlrn: Temporal latent residual networks for large deformation image registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 728–738. Springer, 2024a.

Yifan Wu, Mengjin Dong, Rohit Jena, Chen Qin, and James C Gee. Neural ordinary differential equation based sequential image registration for dynamic characterization. *arXiv preprint arXiv:2404.02106*, 2024b.

Meng Ye, Mikael Kanski, Dong Yang, Qi Chang, Zhennan Yan, Qiaoying Huang, Leon Axel, and Dimitris Metaxas. Deeptag: An unsupervised deep learning method for motion tracking on cardiac tagging magnetic resonance images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7261–7271, June 2021.

Laurent Younes. *Shapes and diffeomorphisms*, volume 171. Springer, 2010.

Miaomiao Zhang, William M Wells III, and Polina Golland. Probabilistic modeling of anatomical variability using a low dimensional parameterization of diffeomorphisms. *Medical image analysis*, 41:55–62, 2017.