



HAL
open science

IMPACT: A Generic Semantic Loss for Multimodal Medical Image Registration

Valentin Bousot, Cédric Hémon, Jean-Claude Nunes, Jason Downling, Simon Rouzé, Caroline Lafond, Anaïs Barateau, Jean-Louis Dillenseger

► **To cite this version:**

Valentin Bousot, Cédric Hémon, Jean-Claude Nunes, Jason Downling, Simon Rouzé, et al.. IMPACT: A Generic Semantic Loss for Multimodal Medical Image Registration. 2025. <hal-05414183>

HAL Id: hal-05414183

<https://hal.science/hal-05414183v1>

Preprint submitted on 12 Dec 2025

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.



HAL Authorization

IMPACT: A Generic Semantic Loss for Multimodal Medical Image Registration

Valentin Boussot, Cédric Hémon, Jean-Claude Nunes, Jason Downling, Simon Rouzé, Caroline Lafond, Anaïs Barateau, Jean-Louis Dillenseger

Abstract—Image registration is fundamental in medical imaging, enabling precise alignment of anatomical structures for diagnosis, treatment planning, image-guided interventions, and longitudinal monitoring. This work introduces IMPACT (Image Metric with Pretrained model-Agnostic Comparison for Transmodality registration), a novel similarity metric designed for robust multimodal image registration. Rather than relying on raw intensities, handcrafted descriptors, or task-specific training, IMPACT defines a semantic similarity measure based on the comparison of deep features extracted from large-scale pretrained segmentation models. By leveraging representations from models such as TotalSegmentator, Segment Anything (SAM), and other foundation networks, IMPACT provides a task-agnostic, training-free solution that generalizes across imaging modalities. These features, originally trained for segmentation, offer strong spatial correspondence and semantic alignment capabilities, making them naturally suited for registration. The method integrates seamlessly into both algorithmic (Elastix) and learning-based (VoxelMorph) frameworks, leveraging the strengths of each. IMPACT was evaluated on five challenging 3D registration tasks involving thoracic CT/CBCT and pelvic MR/CT datasets. Quantitative metrics, including Target Registration Error and Dice Similarity Coefficient, demonstrated consistent improvements in anatomical alignment over baseline methods. Qualitative analyses further highlighted the robustness of the proposed metric in the presence of noise, artifacts, and modality variations. With its versatility, efficiency, and strong performance across diverse tasks, IMPACT offers a powerful solution for advancing multimodal image registration in both clinical and research settings.

Index Terms—Image Registration, Semantic similarity measure, Foundation models

I. INTRODUCTION

Image registration is a fundamental technique in medical imaging that establishes precise anatomical correspondences between multiple scans through spatial transformations. Establishing this correspondence is essential for various clinical and research applications, including contour propagation, disease diagnosis, treatment planning, image-guided treatment and longitudinal monitoring. It enables the integration of information from multiple and complementary imaging modalities, such as Magnetic Resonance Imaging (MRI), Computed Tomography (CT), Cone Beam CT (CBCT), and Positron Emission Tomography (PET), to provide a more complete

understanding of anatomical structures and functional processes. Moreover, accurate registration allows the tracking of anatomical changes over time, for instance, in monitoring tumor progression or assessing treatment response [1], [2].

In unsupervised intensity-based registration, alignment is achieved by optimizing a deformation model through the minimization of a cost function that combines a similarity measure and a regularization term. The similarity measure quantifies the alignment quality between the fixed and moving images, while the regularization term enforces a smooth, anatomically plausible spatial transformation [3], [4], [5].

Similarity measures play an essential role in image registration, directly influencing its effectiveness. However, accurately quantifying similarity is challenging due to the complex, non-linear relationships between multimodal images, which stem from differences in imaging modalities, noise, and acquisition artifacts. Traditional similarity measures such as mean squared error (MSE), normalized cross-correlation (NCC) [6], and mutual information (MI) [7] rely on intensity-based comparisons, assuming that well-aligned images exhibit consistent or statistically related intensities. Although effective in some contexts, these metrics often fail in multimodal scenarios where intensity relationships are not preserved. To overcome these limitations, various handcrafted features have been introduced to incorporate local structural information beyond raw intensity values. Notable examples include the Modality Independent Neighborhood Descriptor (MIND) [8] and the Self-Similarity Context (SSC) [9], both designed to capture spatial patterns that remain stable across modalities. These descriptors demonstrated greater robustness in a variety of registration tasks. However, they remain limited by their sensitivity to noise, their reliance on fixed parameter tuning, and their inability to encode high-level anatomical semantics. Since they rely exclusively on local intensity self-similarity, they can fail in cases of large deformations or anatomical variations, leading to suboptimal alignments [10], [11], [12], [13].

Recent deep learning-based similarity measures address these limitations by replacing handcrafted descriptors with hierarchical representations learned from data [14], [15], [16], [17], [18], [19], [20]. Instead of relying on handcrafted features, these methods leverage deep neural networks to extract multi-scale feature embeddings that capture both structural consistency and semantic context. This shift marks a transition from purely intensity-based registration toward approaches grounded in high-level semantic understanding of image content.

V. Boussot, C. Hémon, J-C. Nunes, C. Lafond, A. Barateau and J-L. Dillenseger are with Univ. Rennes, CLCC Eugene Marquis, INSERM, LTSI - UMR 1099, F-35000 Rennes, France.

J. Downling is with CSIRO Australian e-Health Research Centre, Herston, Queensland, Australia.

S. Rouzé is with CHU Rennes, Department of Cardio-Thoracic and Vascular Surgery, F-35000 Rennes, France.

This transition aligns with recent advances for image synthesis and segmentation. Deep feature-based losses, such as the perceptual loss using VGG embeddings [21], have replaced traditional pixel-wise differences by comparing high-level feature representations rather than raw intensities. By leveraging hierarchical features extracted from pretrained networks, these methods better capture structural and textural correspondences, ensuring that generated images maintain perceptual fidelity while remaining robust to intensity variations and noise. Similarly, in medical image segmentation, comparing deep feature maps has proved more effective than traditional geometrical measures. This feature-based approach allows models to align segmentation outputs with ground truth annotations in a way that better preserves anatomical structures, ensuring sharper boundaries and improved structural coherence [22].

Although these perceptual similarity metrics have shown remarkable success in image synthesis and segmentation, their potential for image registration remains largely unexplored. In particular, no existing approach has leveraged deep feature representations from large-scale models to construct a robust, task-agnostic similarity measure for multimodal registration.

Building on the successes of perceptual loss and handcrafted features, we propose a new feature-based similarity measure for image registration that fully exploits large-scale pretrained models. While previous deep learning-based approaches have explored feature extraction for image registration, they do not explicitly formulate similarity measures as differentiable loss functions [18], [20]. Among these, only one utilizes a generic pretrained model trained on a large-scale dataset without requiring additional task-specific training [19], and only one prior work has proposed a dedicated similarity loss, though it still relies on task-specific training [16].

However, these methods share fundamental limitations: they require task-specific training and operate on feature maps instead of images, leading to suboptimal convergence and spatial inconsistencies, particularly in models with large receptive fields like transformers or deep CNNs. By relying solely on extracted feature maps, these approaches overlook the spatial sensitivity of deep features to transformations. Moreover, they rely on rigid feature extraction frameworks that limit adaptability across various imaging modalities and registration tasks.

In contrast, our method directly leverages the Jacobian of the pretrained model to enable a more principled and spatially coherent optimization process. Furthermore, our approach exploits high-level representations of large-scale generic segmentation models, reassigning them as robust feature extractors for similarity estimation. Moreover, our framework remains highly flexible by formulating similarity estimation as a differentiable loss function, allowing seamless integration into a wide range of registration techniques, including algorithmic and deep learning-based frameworks.

Furthermore, with the increasing availability of generic pretrained models trained on large-scale datasets, leveraging these architectures has become particularly relevant. Models trained on extensive and diverse datasets to segment multiple structures within images learn hierarchical feature representations. These representations capture rich spatial and contextual

information, increasing the robustness of similarity measures while eliminating the need for extensive task-specific training.

In this work, we present IMPACT (Image Metric with Pre-trained model-Agnostic Comparison for Transmodality Registration), a novel generic deep learning-based similarity metric that removes the need for training task-specific models. IMPACT exploits ready-made segmentation models (TotalSegmentator, SAM2.1) as universal 2D/3D feature extractors. By effectively comparing semantic features derived from these large-scale architectures, the proposed method achieves improved alignment in both algorithmic and deep learning-based registration frameworks.

Our contributions are threefold:

- We propose a novel similarity measure that leverages features from large-scale pretrained models such as TotalSegmentator [23], SAM2.1 [24] for multimodal medical image registration
- We establish that these generic features serve as robust and generalizable descriptors for multimodal registration tasks, eliminating the need for large annotated datasets and task-specific training while maintaining high accuracy.
- We evaluate the efficiency and versatility of our method by integrating it into both Elastix [25] (algorithmic registration) and VoxelMorph [26] (deep learning-based registration), leveraging the strengths of each framework, and showing consistent improvements across multiple anatomical regions and imaging modalities.

IMPACT has been extensively evaluated across five challenging multimodal registration tasks, integrating seamlessly into both algorithmic (Elastix) and deep learning-based (VoxelMorph) frameworks. The evaluation encompassed diverse imaging modalities (CBCT, CT, MRI) and anatomical regions (abdomen, pelvis, thorax), and included benchmarks from multiple public challenges. The results demonstrate consistent improvements in alignment accuracy across varying image qualities, acquisition protocols, and anatomical variability, highlighting IMPACT’s robustness and generalizability in real-world clinical settings.

II. BACKGROUND AND RELATED WORK

A. Image Registration: An Optimization Problem

In unsupervised intensity-based image registration, the objective is to estimate a spatial transformation $T_\theta: \Omega_F \rightarrow \Omega_M$ that aligns a moving image $I_M: \Omega_M \rightarrow \mathbb{R}$ with a fixed image $I_F: \Omega_F \rightarrow \mathbb{R}$. Both images are D -dimensional, with each defined over its own spatial domain $\Omega_F, \Omega_M \in \mathbb{R}^D$. The application of the transformation to the moving image I_M is expressed through function composition: $I_M \circ T_\theta$. The goal is to find the optimal transformation by minimizing a cost function \mathcal{C} (2), which depends on a similarity metric \mathcal{S} and a γ -weighted regularization term \mathcal{P} . As the central component of the registration process, \mathcal{S} quantifies the alignment quality between the transformed moving image and the fixed image. Its design is crucial, as it controls the robustness, accuracy, and generalization capability of the registration framework, making it the core focus of this study. To complement this,

the regularization term $\mathcal{P}(T_\theta)$ enforces spatial smoothness by penalizing gradients or curvature (bending energy) of the transformation, ensuring a trade-off between deformation flexibility and anatomical plausibility [27].

This leads to the following optimization problem:

$$\hat{\theta} = \arg \min_{\theta} \mathcal{C}(I_F, I_M \circ T_\theta) \quad (1)$$

$$\text{with } \mathcal{C}(I_F, I_M \circ T_\theta) = -\mathcal{S}(I_F, I_M \circ T_\theta) + \gamma \mathcal{P}(T_\theta) \quad (2)$$

Solving this optimization problem requires either iterative optimization for each image pair or learning-based approaches that generalize over datasets. These lead to two main paradigms in image registration: algorithmic approaches and deep learning-based methods.

B. Registration Framework

Algorithmic approaches iteratively optimize \mathcal{C} for each image pair, typically relying on gradient-based optimization and explicit deformation models. Among these, B-splines [28] provide smooth, regularized transformations by parameterizing the deformation field, while Demons algorithms [29] iteratively estimate displacement fields based on optical flow principles. To enforce topology preservation, diffeomorphic models such as LDDMM [30] and SyN [31] have been introduced. However, optimizing high-dimensional deformation models can be challenging. To improve convergence and capture both global and local deformations, hierarchical multiresolution strategies are commonly employed [28]. By progressively refining the transformation from coarse to fine scales, these methods enhance robustness and prevent local minima in complex deformation spaces. These methods are widely implemented in Elastix [25] and ANTs [32], the two most commonly used open-source frameworks for medical image registration. Both are built on the ITK (Insight Toolkit) framework and provide extensive tools for intensity-based registration. Discrete optimization has emerged as a powerful alternative to continuous methods, offering improved handling of large deformations and reduced sensitivity to local minima. Deeds [33] adopts an MRF framework with densely sampled displacements and efficient regularization via a minimum spanning tree. CorrField [34] extends this idea by introducing sparse keypoints, symmetry constraints, and parts-based regularization to achieve high accuracy even in cases of strong respiratory motion. Recent efforts combine discrete, continuous, and hybrid strategies to further improve registration robustness and accuracy [20].

In contrast, deep learning methods take advantage of CNNs [26] or transformer-based architectures [5] to directly predict T_θ from a pair of images, eliminating the need for iterative optimization during inference. Recent advancements in these architectures primarily aim to enhance their generalization ability, ensuring robust performance on unseen test data across diverse anatomical structures and imaging modalities. However, training these models remains an iterative algorithmic process, where the loss function \mathcal{C} is optimized over multiple iterations using gradient-based methods such as Adam. This process updates the network weights to ensure that the learned

deformation field accurately models the transformations observed in the anatomical structures of the training data.

The coexistence of these two paradigms reflects their complementary strengths. Algorithmic methods provide robustness, generalizability, and reliable performance across datasets, making them well-suited for applications where data is scarce. Furthermore, some operate in physical coordinates, making them independent of image size and voxel resolution. This enhances their effectiveness, particularly in scenarios where the fields of view of the images to be registered differ significantly. However, their computational cost poses a challenge for real-time applications. In contrast, deep learning methods excel in speed and accuracy for task-specific applications, particularly when large, high-quality annotated datasets are available. Although these methods can achieve high accuracy, their performance highly depends on the quality and diversity of the training dataset. Furthermore, they are sensitive to domain shifts and may require fine-tuning or retraining when applied to new datasets. However, acquiring such datasets is often challenging due to variations in imaging protocols and the need for manual annotations. The choice between these approaches depends on the application requirements, balancing computational resources, data availability, and the nature of the data [35].

C. Multimodal Similarity Metrics for Image Registration

A major challenge in image registration is handling complex multimodal scenarios, where variations in intensity, noise, and acquisition artifacts introduce significant discrepancies between the images to be aligned. These differences can lead to misalignment when using traditional similarity metrics such as MSE and NCC, making the design of a robust similarity metric \mathcal{S} essential for accurate registration. While these methods assume consistent intensity relationships in monomodal settings, this assumption breaks down in multimodal scenarios, requiring alternative approaches to account for modality-specific variations.

One of the earliest and most widely used approaches to multimodal registration relies on measuring statistical dependence between intensity distributions. MI [7], [36] and its extensions, such as Normalized MI (NMI) [37], estimate the shared information between images, making them effective for aligning images with complex intensity relationships. However, MI-based metrics can be sensitive to noise and prone to local optima due to inaccurate histogram estimation. To mitigate these limitations, gradient-based refinements such as Gradient-Intensity MI [38] integrate spatial gradient information, enhancing alignment consistency and increasing robustness to noise.

Beyond statistical dependencies, multimodal similarity metrics have evolved into three main approaches. Modality gap reduction strategies seek to minimize differences between imaging modalities, either by explicitly learning a shared representation space or by enforcing robustness to modality variations. Handcrafted descriptors extract structural features that remain stable across modalities, improving robustness to intensity variations. Deep-learned representations leverage

neural networks to extract high-level semantic features, enabling robust alignment across diverse imaging modalities.

The following sections provide a detailed analysis of these approaches, highlighting their strengths and limitations in multimodal image registration.

1) *Modality Gap Reduction Strategies for Similarity Learning*: Modality gap reduction strategies aim to bridge differences between imaging modalities by enhancing similarity in image space or feature space. Some methods transform images into a shared representation, while others enforce robustness to modality variations through adversarial learning or contrast-invariant representations.

Han et al. [39] propose JSR, a framework that jointly synthesizes CT images from MR and CBCT while optimizing deformation fields for multimodal alignment. Hémon et al. [40] follow a sequential approach, first synthesizing target-modality images with GANs, then registering them using traditional similarity metrics. While effective, this method requires extensive training data and may introduce artifacts. Yan et al. [41] introduce AIRNet, which uses adversarial learning to distinguish well-aligned from misaligned pairs, guiding the registration network without explicitly transforming images. However, adversarial training is often unstable and requires careful tuning. Hoffmann et al. [42] propose SynthMorph, a contrast-invariant learning method that removes dependency on real training data by generating synthetic images with extreme contrast and deformation variations. Rather than transforming images, SynthMorph trains networks to ignore contrast differences, enabling robust registration across unseen modalities.

While these approaches improve multimodal registration, they introduce challenges. Domain adaptation and adversarial methods require large datasets and may struggle with unseen modalities, whereas contrast-invariant learning reduces dataset dependency but may not fully capture modality-specific structures.

2) *Hand-Crafted Feature-Based Similarity Metrics*: To move beyond raw intensity comparisons, structural descriptors have been introduced to capture local spatial patterns that remain stable across modalities. Normalized Gradient Fields (NGF) [43] align images based on edge structures rather than absolute intensities, making them more robust to modality changes. Similarly, MIND [8] and SSC [9] extract local self-similar features, improving robustness to intensity variations. While these methods outperform statistical metrics in complex cases, they still rely on manually designed descriptors, limiting their ability to capture high-level semantic correspondences.

3) *Learned Feature-Based Similarity Metric*: The evolution of similarity metrics in medical image registration has shifted towards deep learning-based feature extraction, enabling more robust and adaptive similarity measures. Several methods have been proposed to learn feature representations that capture anatomical correspondences across modalities.

Czolbe et al. [16] introduce a semantic similarity metric that aligns images based on learned, task-specific features rather than intensity values, improving robustness to noise and modality variations. SAME (Self-supervised Anatomical Embeddings) [18] is a monomodal registration pipeline that uses

self-supervised embeddings to learn anatomical correspondences without requiring manual annotations. DINO-Reg [19] leverages self-supervised, transformer-based features from DINOv2 for image registration, highlighting the potential of foundation models in medical image alignment. ConvexAdam [20] leverages handcrafted descriptors (MIND) or learned deep features (nnU-Net, when annotations are available) within an adaptive dual-optimization framework.

While these methods improve registration robustness, they have inherent limitations. Many depend on task-specific training, making them highly dataset-dependent and less generalizable. Additionally, because they operate on feature maps rather than raw images, they can suffer from suboptimal convergence and spatial inconsistencies, particularly in models with large receptive fields, such as transformers or deep CNNs, as observed in DINO-Reg [19]. By relying exclusively on extracted feature representations, these methods overlook the spatial sensitivity of deep features to transformations, potentially causing alignment errors. Additionally, they impose rigid feature extraction frameworks, limiting adaptability across modalities, and do not leverage the advantages of algorithmic registration techniques.

Unlike previous approaches that rely on handcrafted descriptors or task-specific deep feature learning, our method introduces a novel paradigm by using large-scale segmentation models as generic feature extractors for multimodal image registration. Instead of designing similarity metrics from scratch, we directly leverage pretrained segmentation networks, such as SAM and TotalSegmentator to provide robust, modality-agnostic feature representations for defining a semantic similarity metric.

Recent foundation segmentation models have transformed medical image analysis by providing modality-agnostic spatial representations that generalize across diverse imaging modalities. Unlike traditional CNNs, these transformer-based architectures capture long-range dependencies and global context, enabling accurate anatomical segmentation and feature-driven alignment without task-specific training.

Models like Segment Anything Model (SAM) [24] and MedSAM [44] extend prompt-based segmentation to medical imaging, leveraging large-scale heterogeneous datasets to enhance generalization. Similarly, TotalSegmentator [23] and STU-Net [45] unify segmentation across multiple modalities, reinforcing the shift from intensity-based registration to deep feature-driven alignment. These networks offer a promising avenue for improving multimodal image registration by leveraging robust anatomical representations rather than intensity similarities.

By leveraging these pretrained “foundation” segmentation models (SAM, MedSAM, TotalSegmentator, STU-Net, etc.), image registration can be guided by structural similarity rather than raw intensity values. The challenge remains in efficiently integrating these extracted feature maps into optimization frameworks, ensuring alignment is driven by semantic correspondences while maintaining computational efficiency and adaptability.

III. METHOD

A similarity metric based on feature-driven alignment between images is proposed. Features from fixed and moving images are extracted using a selected pretrained model and compared using a selected feature distance measure. The next section introduces this feature-based metric in the context of both registration paradigms: deep learning-based methods, such as VoxelMorph [26], and algorithmic approaches implemented in the Elastix framework [25].

A. Semantic Similarity Metric Based on Pretrained Model for Algorithmic Registration

The fundamental principles of algorithmic registration are outlined within the *Elastix* framework, a widely used open-source tool in biomedical imaging that is built upon the ITK library [46]. *Elastix* offers a comprehensive set of transformation models and robust optimization techniques, supporting various similarity metrics, including MSE, NCC, and NMI that are well-suited for both mono and multimodal applications. However, it does not support the MIND or any semantic similarity measures, motivating the development of a novel metric. This paper presents the full architecture of the registration framework, encompassing not only the proposed semantic similarity metric, but also the iterative optimization process, deformation models, and multi-resolution schemes. An advanced strategy was employed using weakly supervised registration with anatomical masks, allowing the optimization to focus on clinically relevant regions. Lastly, this work provides a comprehensive description of the similarity measure implementation, supporting two distinct operational modes (Jacobian and Static) and four distance metrics \mathcal{D} : L1, L2, NCC, and cosine similarity, to quantify similarity between extracted feature representations.

1) *Adaptive Stochastic Gradient Descent (ASGD)*: A fundamental aspect of the framework is its optimization strategy, which iteratively refines the transformation parameters. The optimization problem (Eq.2) is typically addressed using a first-order derivative-based iterative approach :

$$\theta_{i+1} = \theta_i - a_i \frac{\partial \mathcal{C}(I_F, I_M \circ T_\theta)}{\partial \theta} \quad (3)$$

where a_i denotes the step size or gain factor. Among the available strategies, Adaptive Stochastic Gradient Descent (ASGD) [47] is one of the most commonly used due to its efficiency and robustness in medical image registration. ASGD improves computational efficiency by estimating gradients on a randomly selected subset of voxels rather than the full image domain [48], [49]. This stochastic sampling significantly reduces computation time while maintaining a representative cost function. Additionally, ASGD employs an adaptive stepsize, dynamically adjusting the learning rate to stabilize updates. Unlike deterministic optimization methods, ASGD does not include a built-in convergence criterion, as the stochastic nature of gradient estimation leads to fluctuations in the cost function, making standard stopping conditions unreliable. As a result, our experiments are conducted for a fixed number of iterations to ensure consistent comparisons across different configurations.

2) *Deformation Models*: The experiments conducted with *Elastix* use a free-form deformation model based on B-spline interpolation [28]. The transformation T_θ is represented as a sparse function with a limited number of parameters. B-splines are particularly well-suited for capturing smooth and continuous deformations while maintaining computational efficiency, making them a robust choice for image registration tasks [28].

3) *Multiresolution Strategy*: To address the high computational complexity of optimizing transformations across the entire image domain, a multiresolution strategy is employed [25]. The registration begins at a coarse resolution, where large-scale deformations are estimated, and progressively refines the alignment at higher resolutions as finer details are introduced. This hierarchical approach simplifies the optimization landscape, mitigates the risk of local minima, and enables efficient handling of both global and local transformations.

4) *Weakly Supervised Approach using Masks*: In many medical imaging applications, partial annotations or region-specific masks are available. These can be effectively used within a weakly supervised registration framework to guide the alignment toward clinically relevant structures. Instead of using the entire image domain for optimization, the *Elastix* framework allows restricting the registration process to regions defined by these masks. This approach focuses the optimization on clinically relevant areas, such as organs or pathological regions, while ignoring irrelevant or noisy regions in the image. Additionally, this strategy is particularly useful for capturing complex anatomical interactions, such as sliding motions between adjacent organs [50], [51], by limiting the optimization to the ROI where such motions are most relevant. This ensures more accurate and anatomically meaningful registration results.

Technically, the cost function is computed using only voxels within the mask domain $\Omega_F^{\text{mask}} \subset \Omega_F$ of the fixed image and the mask domain $\Omega_M^{\text{mask}} \subset \Omega_M$ of the moving image, reducing the complexity and improving the robustness of the registration. The transformation T_θ is then optimized to align the fixed and moving images only within the masked regions.

5) *IMPACT : Jacobian mode*: Fig. 1 provides a schematic overview of the Jacobian and static mode workflow, illustrating how patches are sampled from the fixed and moving images, passed through the shared feature extractor, and then compared to produce gradients for updating the transformation parameters.

For each iteration, N patches of dimension D are extracted uniformly at random. Each patch has a size $P = (P_1, P_2, \dots, P_D)$, where $P_i \in \mathbb{N}^D$, corresponding to the receptive field at the deeper layer, and is sampled from both the fixed and moving images. Based on these patches, the similarity measure is defined as follows:

$$\mathcal{S}(I_F, I_M \circ T_\theta) = \frac{1}{NC} \sum_{l=1}^L \gamma_l \sum_{i=1}^N \mathcal{D} \left(\phi_l(\widehat{I}_F(x_i)), \phi_l(\widehat{I}_M(T_\theta(x_i))) \right) \quad (4)$$

where $\widehat{I}_F(x_i) : \Omega_F \rightarrow \mathbb{R}^{P_1 \times P_2 \times \dots \times P_D}$ and $\widehat{I}_M(x_i) : \Omega_M \rightarrow \mathbb{R}^{P_1 \times P_2 \times \dots \times P_D}$ denote the patches extracted from the fixed and moving images, centered at the coordinate x_i .

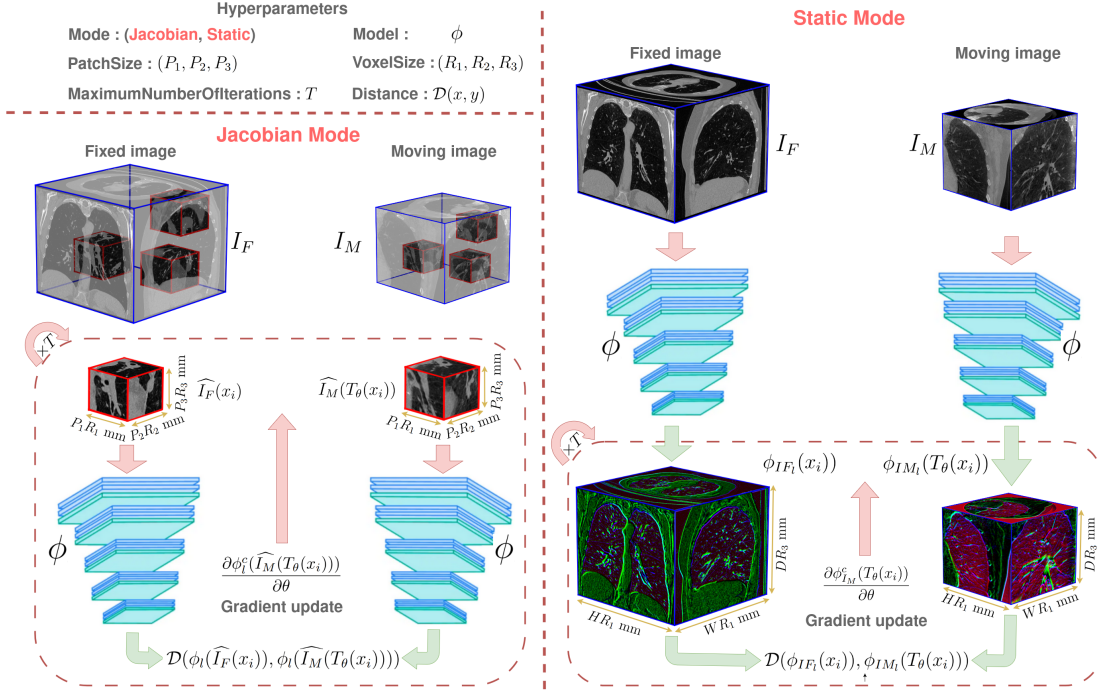


Fig. 1. Schematic overview of the proposed feature-based similarity metric the algorithmic registration framework in the 3D case. Features are extracted from the fixed and moving images using a shared pretrained encoder ϕ . The similarity is computed between features of the fixed image $\phi(I_F(x_i))$ and the transformed moving image $\phi(I_M(T_\theta(x_i)))$ using a distance function \mathcal{D} . Two modes are supported: Jacobian Mode, where local patches are sampled and gradients are backpropagated through the feature extractor, and Static Mode, where features are precomputed over the entire image and reused throughout the optimization. Gradients are used to iteratively update the transformation T , aligning the moving image to the fixed image.

These patches are uniformly sampled such that all positions within each patch lie entirely within their respective mask regions. The extracted patches are then resampled to a resolution of $R \in \mathbb{R}^D$ mm using third-order B-spline interpolation over the image function I .

The function $\phi_l : \mathbb{R}^{P_1 \times P_2 \times \dots \times P_D} \rightarrow \mathbb{R}^C$ extracts C feature representations from the patch at layer l , while γ_l represents the weight assigned to that layer. The function $\mathcal{D}(x, y)$ quantifies the similarity between two feature representations.

The focus is placed exclusively on first-order gradient-based optimization methods, specifically employing gradient descent to iteratively minimize the loss function by adjusting the transformation parameters, θ . The gradient of the loss function with respect to θ is expressed as:

$$\begin{aligned} \frac{\partial \mathcal{S}(I_F, I_M \circ T_\theta)}{\partial \theta} = & \quad (5) \\ \frac{2}{NC} \sum_{l=1}^L \gamma_l \sum_{i=1}^N \sum_{c=1}^C \left. \frac{\partial \mathcal{D}(f, m)}{\partial m} \right|_{\substack{f = \phi_l^c(\hat{I}_F(x_i)) \\ m = \phi_l^c(\hat{I}_M(T_\theta(x_i)))}} & \\ \times \frac{\partial \phi_l^c(\hat{I}_M(T_\theta(x_i)))}{\partial \theta} & \end{aligned}$$

Using the chain rule, the derivative of the c -th feature with respect to the transformation parameters θ can be expressed as:

$$\frac{\partial \phi_l^c(\hat{I}_M(T_\theta(x_i)))}{\partial \theta} = \nabla_{\hat{I}_M} \phi_l^c \cdot \nabla_x \hat{I}_M \cdot \frac{\partial T_\theta(x_i)}{\partial \theta} \quad (6)$$

The equation (Eq.6) highlights the interplay between three components:

- **Gradient of the feature extractor** ($\nabla_{\hat{I}_M} \phi_l^c \in \mathbb{R}^{P_1 \times P_2 \times \dots \times P_D}$): This term captures how the c -th feature, extracted by ϕ_l , responds to changes in the intensity of the image patch \hat{I}_M .
- **Spatial gradient of the moving image** ($\nabla_x \hat{I}_M \in \mathbb{R}^{(P_1 \times P_2 \times \dots \times P_D) \times D}$): This term measures how \hat{I}_M varies with respect to spatial coordinates x . For smoother results, $\nabla_x \hat{I}_M$ is computed using third-order B-spline interpolation, which ensures stability and continuity.
- **Local Jacobian of the transformation** ($\frac{\partial T_\theta(x_i)}{\partial \theta} \in \mathbb{R}^{D \times |\theta|}$): This represents how the transformation T_θ deforms spatial coordinates in response to changes in its parameters θ .

By integrating these components, the resulting gradient $\frac{\partial \phi_l^c(\hat{I}_M(T_\theta(x_i)))}{\partial \theta} \in \mathbb{R}^{|\theta|}$ reflects the contribution of the feature extractor ϕ , whose gradients with respect to the moving image I_M are efficiently computed via backpropagation. Notably, this computation is independent of the specific deformation model considered, as the Jacobian $\frac{\partial T_\theta(x_i)}{\partial \theta}$ encapsulates the transformation influence on the spatial coordinates. This decoupling ensures that the approach is flexible and compatible with a wide range of transformation models.

6) *IMPACT : Static mode*: The time required to compute the Jacobian of the feature extraction function ϕ increases with the model's receptive field, making backpropagation through deep networks computationally intensive. While these gradients are very useful for efficiently minimizing the loss function

in a few iterations, they become impractical for models with large receptive fields. However, they are not strictly necessary when the extracted features maintain spatial alignment with the input images and when the spatial sensitivity of deep features to geometric transformations is negligible. In such cases, full backpropagation can be avoided without significantly affecting alignment performance. For this reason, a slightly simplified version, referred to as Static mode, was developed for loss function computation. This approach bypasses gradient calculation by precomputing the features and directly comparing them.

In static mode, the feature maps for the fixed and moving images are computed once per resolution level or after a few iterations. Once generated, these feature maps are treated as multi-channel images.

Formally, at each iteration, N spatial coordinates are randomly selected from the fixed image domain Ω_F (or its corresponding mask). The similarity measure is expressed as follows:

$$\begin{aligned} \mathcal{S}(I_F, I_M \circ T_\theta) = & \quad (7) \\ \frac{1}{NC} \sum_{l=1}^L \gamma_l \sum_{i=1}^N \mathcal{D}(\phi_{IF_l}(x_i), \phi_{IM_l}(T_\theta(x_i))) \end{aligned}$$

where $\phi_{IF_l} : \Omega_F \rightarrow \mathbb{R}^C$ and $\phi_{IM_l} : \Omega_M \rightarrow \mathbb{R}^C$ represents respectively, from the fixed and moving images, the C features maps extracted at the layer l and γ_l represent the weight assigned to that layer. The gradient expression of the similarity measure (Eq.8) is simplified as there is no need to compute the Jacobian of ϕ .

$$\begin{aligned} \frac{\partial \mathcal{S}(I_F, I_M \circ T_\theta)}{\partial \theta} = & \quad (8) \\ \frac{2}{NC} \sum_{l=1}^L \gamma_l \sum_{i=1}^N \sum_{c=1}^C \frac{\partial \mathcal{D}(f, m)}{\partial m} \Big|_{\substack{f = \phi_{IF_l}^c(x_i) \\ m = \phi_{IM_l}^c(T_\theta(x_i))}} & \times \frac{\partial \phi_{IM_l}^c(T_\theta(x_i))}{\partial \theta} \end{aligned}$$

The extracted feature maps are recorded as matrices, however, ϕ_I can be interpreted as a continuous function through third-order B-spline interpolation.

B. Integration of IMPACT in Elastix Framework

The implementation of the IMPACT within the Elastix framework is highly flexible, allowing for advanced customization. Users can fine-tune the similarity metric parameters to achieve the desired level of precision.

The deep learning-based feature extraction process is defined by seven parameters: the model path, which specifies where to load the model from; the dimensionality and number of channels of the input patch images fed into the model; the patch size P ; the processing resolution R ; a mask for selecting the desired layers within the model to be included in the similarity comparison; and the choice of the distance metric \mathcal{D} for feature comparison. Multiple feature extraction processes from

different models can be applied simultaneously. Additionally, users can define a distinct feature extraction setup for each resolution level.

The calculation of $\phi(\cdot)$ (4) on each patch and the computation of Jacobians (5) are performed using the Torch library [52]. The use of TorchScript ensures compatibility with a wide range of pretrained networks, including 2D and 3D DL-based architectures. The system supports parallel computation on the CPU, with tasks divided into batches to accelerate execution across multiple cores. Additionally, GPU support is fully integrated for the computation of ϕ and its Jacobian.

To further reduce the computational burden, a mechanism is introduced to randomly sample a subset of features uniformly at each iteration. Instead of comparing all features extracted from the pair of patches, only a randomly selected subset of features is used to compute the similarity metric. In static mode, a Principal Component Analysis (PCA) of all features can also be used to compare a reduced and more compact set of features.

If the number of channels required by the model input differs from that of the image, the missing channels are either duplicated or replaced by the mean value.

When the model's dimensions are smaller than the image, such as when using a pretrained model like SAM2.1 originally designed for natural images in a 3D image registration task, 2D patches are randomly sampled from various planes within the 3D volume. These planes are not restricted to the principal axes (sagittal, coronal, and axial) but are instead randomly defined in the 3D space, allowing for the extraction of slices with diverse orientations. This approach captures the spatial complexity of the image while adapting the data to the 2D model.

In static mode, the extracted feature maps can be obtained via patch inference, where the input image is divided into smaller overlapping patches that are processed independently. This approach is particularly advantageous for handling high resolution 3D images. By dividing each image into patches, the system avoids memory limitations and enables the processing of large datasets. By default, this inference is performed once before each resolution but can be repeated every k iterations.

Please refer to Appendix B for a complete list of metric parameters.

C. Semantic Similarity Metric Using Pretrained Models for Deep Learning-Based Registration

IMPACT was integrated into the VoxelMorph framework [26] to optimize the learning of a deformable transformation model. VoxelMorph is an unsupervised UNet-based architecture designed to compute a deformation vector field (DVF) for aligning two input images. Its encoder-decoder with skip connections structure allows the encoder to extract multiscale features, while the decoder predicts the DVF. The semantic loss function used in this integration provided two distinct types of features: encoder features, offering general and abstract image representations, and decoder features, optimized for segmentation relevant tasks. Regularization of the DVF is achieved through a diffusion regularizer applied to the spatial

gradients, ensuring smooth and diffeomorphic transformations that preserve anatomical topology and continuity. Compared to the Elastix integration, VoxelMorph offers greater flexibility, enabling access to both encoder and decoder features in Jacobian mode while maintaining computational efficiency. This adaptability facilitates efficient exploration of feature representations. This approach can incorporate embeddings from the same broad set of pretrained models (TotalSegmentator, SAM/MedSAM, STU-Net) without modifying VoxelMorph main architecture.

IV. EXPERIMENTAL SETUP

This section outlines the experimental protocol used to evaluate the IMPACT metric. Five multimodal registration tasks, involving different imaging modalities (CT, MRI, CBCT) and anatomical regions (thorax, abdomen, and pelvis), are used as benchmarks. Tasks 1, 2, 3, and 5 are addressed using the Elastix framework, while tasks 3 and 4 are addressed using the VoxelMorph framework. The clinical tasks are first presented, followed by the configuration details of both registration frameworks. Finally, the evaluation protocol and the metrics used for quantitative and qualitative assessment are described.

A. Clinical tasks

- 1) Thorax CT/CBCT registration (VATSop [53]): The task focuses on lung nodule localization during video-assisted thoracoscopic surgery (VATS), a minimally invasive technique used for the surgical treatment of early-stage lung cancer [53], [54], [55]. This dataset is particularly challenging due to inherent CBCT artifacts (scatter, noise, truncation artifacts) [56] and its limited field of view compared to CT. Additionally, the change in patient positioning from supine to lateral decubitus causes significant lung deformations, including displacements exceeding 40 mm, sliding motions of up to 30 mm between lung lobes and the thoracic wall, and localized volume changes [57]. For validation, precise lobe segmentations and approximately forty paired anatomical landmarks (vessel and airway bifurcations) were manually selected per CT/CBCT pair by the VATS surgeon at Rennes University Hospital. This retrospective study received ethical approval (2016-A01353-48 35RC169838), and informed consent was obtained from all patients.
- 2) Thorax FBCT/CBCT registration (ThoraxCBCT Learn2Reg 2023 MICCAI Challenge [58], [12]): This task focuses on thoracic image registration between pre-therapeutic fan beam CT (FBCT) and interventional CBCT scans [58]. The dataset includes three patients, each with longitudinal imaging acquired at different stages of image-guided radiation therapy. For each patient, one FBCT scan is obtained before treatment during maximum inspiration, and two CBCT scans are acquired during maximum expiration. The CBCT images contain substantial artifacts such as scatter, noise, and truncation, which reflect real-world clinical conditions and increase the complexity of the registration task. Evaluation is based on ground truth

- segmentation masks for lung lobes, tumors, and several organs at risk, including the heart, spinal cord, and esophagus, these annotations are not publicly accessible.
- 3) Abdomen MR/CT registration (AbdomenMRCT Learn2Reg 2021 MICCAI Challenge [59], [12]): This task involves multimodal registration between abdominal MRI and CT scans. The dataset includes three patients, each with paired volumetric T1-weighted MRI and contrast-enhanced CT scans. Significant modality differences and anatomical variability make the task particularly challenging. Evaluation is based on ground truth segmentation masks for key abdominal structures, including the liver, spleen, kidneys, and pancreas, these annotations are not publicly accessible.
- 4) Pelvis MR/CT registration (PelvisMRCT [60]): The task consists of 39 prostate cancer patients, each having both CT and MRI scans. The MRI scans were acquired using 3D T2-weighted SPACE sequences on a 3T Siemens Skyra scanner. Routine planning CT scans were performed using either a GE LightSpeedRT large-bore scanner or a Toshiba Aquilion. Each patient CT scan was coregistered to their whole-pelvis T2 with a robust symmetric rigid registration. For evaluation purposes, manual delineations of the prostate, rectum, bladder, and bones were independently performed on both CT and MRI scans. Ethics approval was obtained from the local area health ethics committee, and informed consent was obtained from all patients.
- 5) Thorax CT/CT registration (LungCT Learn2Reg 2021 MICCAI Challenge [61], [12]): This task focuses on registering lung CT scans acquired at two respiratory phases—full inspiration and full expiration. The dataset includes eleven patients in total: eight for testing and three for validation. It poses significant challenges due to the large deformations caused by respiratory motion, resulting in pronounced lung displacements and localized volume changes. Additionally, expiration scans have a reduced field of view, further complicating the registration process. Quantitative evaluation is performed based on anatomical landmarks, which are not publicly accessible.

B. Common configuration

The IMPACT metric was integrated as a custom similarity measure and configured with task-specific hyperparameters. Depending on the task, features were extracted from different layers of various pretrained models, with the specific combinations detailed in Appendix A, and compared using different feature distances. In the Elastix framework, two modes of operation were used: Jacobian mode, which enables back-propagation through the feature extractor, and Static mode, which relies on precomputed features and is better suited for extracting deeper semantic representations. The detailed configurations applied to each task are summarized in Table I. All input images were intensity normalized and resampled into a canonical orientation to ensure compatibility with the pretrained models used for semantic feature extraction, following the specifications detailed in Appendix A.

TABLE I
HYPERPARAMETERS USED FOR IMPACT WITHIN THE ELASTIX AND VOXELMORPH FRAMEWORKS ACROSS DIFFERENT TASKS. THE PARAMETERS LISTED CORRESPOND TO THE OFFICIAL CHALLENGE SUBMISSIONS. ADDITIONAL CONFIGURATIONS USED IN ABLATION STUDIES ARE ALSO REPORTED, WHERE "VARIABLE" INDICATES PARAMETERS THAT WERE SYSTEMATICALLY MODIFIED FOR EXPERIMENTAL COMPARISON.

Framework	Task	Model	PatchSize	Layers	Distance	Mode
Elastix	1	M258	11*11*11	2	L2	Jacobian
	1(Ablation)	Variable	Variable	Variable	Variable	Jacobian
	2	M730-731	11*11*11	2	L1	Jacobian
	3	M730	-	8	L1	Static
	5	M258	11*11*11	2	L2	Jacobian
VoxelMorph	5(Ablation)	Variable	Variable	Variable	Variable	Jacobian
	3	M730	-	7	L1	Jacobian
	4	M730	-	Variable	L1	Jacobian

C. Elastix setup

All Elastix-based experiments shared a common configuration for deformable registration. A 3D B-spline transformation model was optimized using the ASGD optimizer [47]. The similarity measure was computed at each iteration over a set of 2000 points randomly sampled in physical coordinates. The use of 2000 sampled points is sufficient to provide a reliable estimate of image similarity and to ensure stable convergence. This was combined with a four-level multi-resolution strategy, where image resolution was progressively refined through resampling at spacings of 6 mm, 3 mm, 1.5 mm, and 1 mm. Simultaneously, the B-spline control point spacing was halved at each resolution level, reaching a final spacing of 8 mm at the finest resolution. This configuration was chosen as a robust default across tasks and is further justified by the analysis in Appendix C.

A weakly supervised strategy was employed in tasks 1 and 5 by restricting the similarity computation to the lung region of the fixed image. Task 3 included an affine initialization step prior to non-rigid registration, using the same IMPACT configuration but with an affine transformation model. All experiments were conducted on a workstation equipped with an Intel Core i7-12700K CPU, 32 GB of RAM, and an NVIDIA RTX 3080 GPU.

D. VoxelMorph setup

A vanilla VoxelMorph model was implemented in PyTorch and trained using the AdamW optimizer with a learning rate of 0.0002, processing one image pair per iteration. The model was trained for 200 epochs on an NVIDIA RTX A6000 GPU. VoxelMorph was used as an optimization tool rather than a generalizable model, training was therefore performed directly on the image pairs of each task, without the need for a separate training set. As in the Elastix-based setup, task 3 included an initial affine alignment step using *Elastix*.

E. Evaluation Protocol

The evaluation protocol was designed to thoroughly assess the registration performance of the proposed methods across a variety of tasks. Table II provides a quick overview of which figures and metrics correspond to which task and method.

For task 1, the efficiency of IMPACT was demonstrated by comparing its performance in terms of Target Registration Error (TRE), Dice Similarity Coefficient (DSC) and the 95th percentile Hausdorff distance (HD95) against various conventional similarity measures, including MSE, NCC, NMI and MIND. For a fair comparison, an extraction model for the MIND features has been developed, enabling it to be used in algorithmic registration with Elastix. Four MIND configurations were tested with a radius of 1 and 2, as well as a dilation of 1 and 2. Only the best-performing configuration is presented. Moreover, the convergence speed based on the number of iterations was evaluated, exploring several values: [50, 100, 150, 200, 500, 1000, 1500, 2000]. Task 2 and 3 were evaluated primarily using challenge-specific metrics, including the DSC, HD95, and TRE, which was used exclusively for Task 2. For Task 4, registration accuracy was assessed using only the DSC.

For task 4, the model was trained on the full pelvic MR/CT dataset. Different semantic feature configurations were evaluated, including encoder and decoder layers from TotalSegmentator and the features of the second layer of MedSAM3D. As baselines, MI and MIND descriptors were also integrated into the same VoxelMorph framework to enable direct comparisons with conventional similarity metrics.

Qualitative evaluations were conducted across tasks 1 and 4 to complement quantitative results. This involved visual inspections to assess alignment quality, contour accuracy, and the presence of noticeable artifacts, ensuring the robustness and interpretability of the results.

Additionally, an ablation study was conducted on tasks 1 and 5 to assess the impact of different hyperparameter choices on the performance of IMPACT, evaluated using TRE.

Finally, a computational complexity analysis was performed to evaluate execution time. This holistic evaluation framework ensures a well-rounded assessment of both algorithmic and learning-based approaches.

TABLE II
METHOD RESULT PER TASK MAPPING WITH RELEVANT FIGURES

Task	Methodology (Section)	Evaluation Metric	Figure(s)
Task 1: VATSop	Elastix (Section IV-C)	Qualitative, TRE, DSC, HD95	Fig. 2,4, Table III and VIII
Task 2: ThoraxCBCT	Elastix (Section IV-C)	DSC, TRE	Table IV
Task 3: AbdomenMRCT	Elastix & VoxelMorph (Section IV-C-IV-D)	Qualitative, DSC	Fig. 3 and Table V
Task 4: PelvisMRCT	VoxelMorph (Section IV-D)	DSC	Fig. 5
Task 5: LungCT	Elastix (Section IV-C)	TRE	Table VI and IX

V. RESULTS

A. Qualitative results

Qualitative registration outcomes for Task 1 are presented in Fig. 2. Two cropped regions are shown (each row corresponds




Fig. 2. Qualitative comparison of the image registration results. The reference image (first column) is compared with the registered moving image obtained after 2000 iterations (presented in the second and third columns). The registered image obtained using IMPACT is presented in the second column and the one obtained via NCC in the third. Each row presents an example from two distinct patients.

to a different patient): the first column displays the fixed image (CBCT), while the second and third columns show the registered moving image (CT) using IMPACT and NCC, respectively, after 2000 iterations. Notably, when guided by our metric, the lung contours and local structures (circled or indicated with arrows) exhibit visibly improved alignment, even in the presence of CBCT artifacts.

Fig. 3 shows abdomen MR/CT registration results (Task 3) using VoxelMorph with three different similarity metrics: MI, MIND, and our proposed IMPACT. The first column is the reference (fixed) MRI, followed by a rigidly aligned CT, and then the warped CT with MI, MIND, and IMPACT. As the figure shows, MI and MIND-based registrations sometimes fail to align organs with high precision, whereas IMPACT captures more anatomically coherent correspondence.

B. Quantitative evaluation and convergence analysis

Figure 4 illustrates an experiment designed to evaluate the convergence behavior and accuracy of different similarity metrics during registration. We compare MSE, NMI, NCC, and IMPACT (in both Jacobian and Static modes) on Task 1. The plot shows the TRE distribution as a function of the number of iterations. IMPACT not only converges faster but also achieves the lowest median TRE with a notably tighter interquartile range, indicating more consistent performance across patients.

Table III summarizes the registration performance on Task 1 after 2000 iterations, comparing different similarity measures using TRE, DSC, and HD95. IMPACT in Jacobian mode achieves the best overall performance. The Static variant performs similarly, with slightly reduced boundary accuracy. Notably, IMPACT outperforms the strongest baseline, MIND, across all evaluation metrics. Traditional metrics MSE, NMI,

and NCC perform substantially worse. The rigid case, used as initialization, reflects the absence of deformation and highlights the large displacements to be estimated by the deformable models.

TABLE III
EVALUATION OF DIFFERENT SIMILARITY METRICS ON TASK 1 AFTER 2000 ITERATIONS, USING TRE, DSC, AND HD95.

Method	TRE (25%)	TRE (50%)	TRE (75%)	DSC	HD95
Rigid (Before Registration)	9.44	13.59	17.58	0.79	14.75
MSE	1.09	1.68	4.08	0.91	9.91
NMI	1.09	1.62	3.03	0.85	14.54
NCC	1.06	1.58	2.92	0.88	12.01
MIND (R1D2)	0.93	1.25	1.85	0.96	3.70
IMPACT Jacobian	0.85	1.2	1.71	0.97	3.19
IMPACT Static	0.85	1.22	1.71	0.95	6.06

Elastix with IMPACT was used in the Learn2Reg Thorax-CBCT challenge, which involves highly artifact-prone CBCT images acquired across different respiratory phases. As shown in Table IV, our method (team “BreizhReg”) ranked 6th among all participants.

The Abdomen MRCT task of the Learn2Reg 2021 MIC-CAI Challenge, characterized by strong intensity and contrast differences between T1-weighted MRI and CT, was used to evaluate both Elastix and VoxelMorph frameworks with IMPACT. As shown in Table V, Elastix ranked 7th overall, while VoxelMorph achieved 2nd place among 26 participants (average DSC: 0.8973, HD95: 3.32 mm). Although both use the same similarity metric, they rely on different deformation models. VoxelMorph predicts DVFs through a learning-based approach trained on the dataset. In contrast, Elastix uses a B-spline transformation with strong regularization, estimated independently for each image pair.

Fig. 5 compares TotalSegmentator encoder and decoder layers, MedSAM3D features, and traditional similarity metrics (MI, MIND) within the VoxelMorph registration framework. Overall, the encoder layers from TotalSegmentator yield results roughly equivalent to MI and MIND, except for layer 3, which surpasses these metrics. Minor variations do appear among different decoder depths, features based on decoder layers outperform all other metrics. Meanwhile, MedSAM3D features produce Dice scores better than MI and MIND but still lower than those achieved by TotalSegmentator decoder layers.

We further evaluated Elastix with IMPACT on the LungCT task of the Learn2Reg 2021 Challenge, which focuses on aligning lung CT scans using anatomical landmarks as evaluation targets. As reported in Table VI, our method (team “BreizhReg”) achieved a 3th place ranking out of 32 participants, with a mean TRE of 1.90 mm.

C. Computational Complexity Analysis

The computational cost of the method depends linearly on the number of iterations N , the number of sampled points per iteration S , and the cost $f(\cdot)$ of computing the similarity on a single sample. The overall complexity can be expressed as:

$$\mathcal{O}(N.S.f(\cdot)) \quad (9)$$




Fig. 3. Qualitative results obtained using VoxelMorph on Task 3 with different loss functions. The first column presents the reference image, followed by results from rigid alignment (Rigid) and the registered CT images estimated using three loss functions: MI, MIND, and the proposed IMPACT loss. Each row illustrates a different example, highlighting the influence of the loss function on registration quality.

TABLE IV
RANKING RESULTS OF THE 21 PARTICIPANTS IN **LEARN2REG THORAXCBCT CHALLENGE**. THE TABLE PRESENTS RESULTS BASED ON DSC. THE NON-RIGID REGISTRATION BASED ON ELASTIX WITH IMPACT RANKED 6TH.

#	User (Team)	DSC
1st	mysterious_man	0.732 ± 0.0026
2nd	anonymous2024	0.689 ± 0.0053
3rd	SSKJLBW	0.669 ± 0.0021
6th	Our team (BreizhReg) (Elastix)	0.652 ± 0.0032
8th	Challenge Organizers(<i>deedsBCV</i>)	0.648 ± 0.0030
12th	Challenge Organizers(<i>ConvexAdam 30 label</i>)	0.586 ± 0.0027
15th	Challenge Organizers(<i>NiftyReg</i>)	0.568 ± 0.0036
17th	Challenge Organizers(<i>Voxelmorph++</i>)	0.503 ± 0.113
21th	Challenge Organizers(<i>Initial Displacements</i>)	0.313 ± 0.093




Fig. 4. Registration performance measured in terms of TRE. The box plots represent the distribution of errors across varying numbers of iterations, comparing different cost functions: MSE, NMI, NCC, and IMPACT in both Jacobian and Static modes.

An estimate of the total runtime can be directly inferred from this relation for a given configuration. Table VII provides empirical runtimes for a representative setting with $N = 500$ and $S = 2000$, covering both standard similarity metrics and IMPACT using different pretrained models, including MIND, TotalSegmentator (TS), and SAM2.1, in both static and Jacobian modes for the first and second feature layer.




Fig. 5. Boxplot of DSC per organ for rigid registration and non-rigid registration using VoxelMorph with different loss functions, including MI, MIND, and IMPACT. The figure also presents DSC values for IMPACT when applied with feature representations from layers 2, 3, and 4 of the decoder (Dec) and encoder (Enc) of the TotalSegmentator model M730, and the second layer of the MedSAM3D network.

VI. ABLATION STUDIES

To analyze the impact of different hyperparameter choices on the performance of IMPACT, an ablation study was con-

TABLE V

RANKING RESULTS OF THE 26 PARTICIPANTS IN [TASK 1 \(ABDOMEN MRCT\) OF THE LEARN2REG 2021 CHALLENGE](#). THE TABLE PRESENTS RESULTS BASED ON DSC, AND HD95 METRICS WHERE HIGHER DSC VALUES AND LOWER HD95 VALUES INDICATE BETTER PERFORMANCE IN MEDICAL IMAGE REGISTRATION. THE NON-RIGID REGISTRATION BASED ON ELASTIX WITH IMPACT RANKED 7TH AND NON-RIGID REGISTRATION BASED ON VOXELMORPH WITH IMPACT RANKED 2ND.

#	User (Team)	DSC	HD95
1st	cwmokab (Orange)	0.9148 ± 0.0018	2.7269
2nd	Our team (BreizhReg) (Voxelmorph)	0.8973 ± 0.0152	3.3169
3rd	honkamj	0.8956 ± 0.0137	2.8451
6th	Challenge Organizers (<i>corrField</i>)	0.8757 ± 0.0168	4.4036
7th	Our team (BreizhReg) (Elastix)	0.8748 ± 0.0066	3.1051
13th	Challenge Organizers (<i>AdamReg Square MIND</i>)	0.8040 ± 0.0253	7.0
19th	Challenge Organizers (<i>NiftyReg MIND</i>)	0.5395 ± 0.1593	15.0666
24th	Challenge Organizers (<i>Initial Displacements</i>)	0.3096 ± 0.1556	22.8887

TABLE VI

RANKING RESULTS OF THE 32 PARTICIPANTS IN [TASK 2 \(LUNGCT\) OF THE LEARN2REG 2021 CHALLENGE](#). THE TABLE PRESENTS RESULTS BASED ON TRE. ELASTIX WITH IMPACT RANKED 3TH IN THIS LEADERBOARD.

#	User (Team)	TRE
1st	(Challenge Organizers) (<i>Fraunhofer MEVIS</i>)	1.8356 ± 0.5685
2nd	(Challenge Organizers) (<i>corrField</i>)	1.8497 ± 0.7033
3th	Our team (BreizhReg) (Elastix)	1.8986 ± 0.7207
11th	(Challenge Organizers) <i>SLIC-Reg++</i>)	2.6359 ± 1.0163
21th	(Challenge Organizers) (<i>Initial Displacements</i>)	14.6407 ± 6.0765

TABLE VII

EMPIRICAL COMPUTATIONAL COST FOR DIFFERENT SIMILARITY METRICS, EVALUATED WITH $N = 500$ ITERATIONS AND $S = 2000$ SAMPLES. FOR EACH CONFIGURATION, THE RUNTIME IN SECONDS IS REPORTED FOR THE FIRST AND SECOND FEATURE LAYERS.

Intensity-based Metrics			
MSE (s)		NCC (s)	NMI (s)
15.10		16.45	114.62
Per-Model Evaluation			
Model	Mode	First Layer (s)	Second Layer (s)
MIND	Jacobian	230.98	–
	Static	105.29	–
TS	Jacobian	131.55	586.72
	Static	160.83	170.56
SAM2.1	Jacobian	393.58	408.87

ducted on tasks 1 and 5. This experiment aims to evaluate three key parameters in terms of TRE : (1) the pretrained model used for feature extraction, (2) the distance measure between extracted features, and (3) the feature extraction level within the model.

A. Impact of the Pretrained Model

Several pretrained models were evaluated to assess their influence on the quality and relevance of extracted features for the registration tasks. The models included ConvNeXt, SAM2.1, MedSAM2, and the TotalSegmentator variants M258, M291, and M730 (see Appendix A for details on training data and target tasks).

Tables VIII and IX indicate that M258 consistently achieves the lowest TRE across Tasks 1 and 5, particularly when using second-layer features combined with the L2 distance, yielding TREs of 1.20 ± 0.82 and 1.32 ± 1.83 , respectively. Other segmentation-based models, including M291, M730, SAM2.1, and MedSAM2, also deliver strong performance, though with slightly higher variability. In contrast, the classification-based model ConvNeXt consistently shows the weakest results, with TREs reaching 1.25 ± 0.86 on Task 1 and 7.14 ± 5.78 on Task 5 under its best configuration. Finally, the handcrafted descriptor MIND, which we reimplemented as a pretrained model within the IMPACT optimization scheme, shows competitive performance in Task 1, especially with L2 distance. However, it is consistently outperformed by learned features on Task 5, even in its best configuration (R1D2).

B. Impact of the Distance Measure

We analyzed how different distance metrics affect registration performance by evaluating their ability to emphasize orientation, magnitude, or a combination of both within the feature space. The compared metrics included cosine similarity, NCC, and L1/L2 distances.

Cosine similarity captures only the angular relationship between feature vectors and is invariant to scale. While this may be beneficial in tasks where direction alone is informative, it disregards magnitude differences that often encode essential anatomical variations. As a result, cosine consistently underperformed across tasks, particularly in Task 5 (Table IX).

NCC incorporates both orientation and relative magnitude and generally performs well when features are linearly related.

TABLE VIII

QUANTITATIVE COMPARISON OF PRETRAINED MODELS, FEATURE DISTANCE METRICS, AND FEATURE EXTRACTION LEVELS USED WITHIN IMPACT ON TASK 1, REPORTING THE MEAN AND STANDARD DEVIATION OF THE TRE. INTENSITY-BASED METRICS ARE ALSO REPORTED FOR REFERENCE.

Intensity-based metrics								
MSE			NCC			NMI		
1.86 ± 4.97			1.79 ± 3.68			1.78 ± 4.61		
Per-Model Evaluation								
Model	First Layer				Second Layers			
	L1	L2	NCC	Cosine	L1	L2	NCC	Cosine
MIND (R1D2)	1.28 ± 0.97	1.27 ± 0.89	1.26 ± 0.95	1.32 ± 0.97				
M258	1.42 ± 1.53	1.58 ± 5.08	1.29 ± 0.86	1.7 ± 8.56	1.23 ± 0.87	1.2 ± 0.82	1.25 ± 0.86	1.26 ± 0.89
M291	1.23 ± 0.86	1.25 ± 0.87	1.3 ± 0.89	1.24 ± 0.99	1.24 ± 0.94	1.22 ± 0.92	1.32 ± 1.02	1.23 ± 0.88
M730	1.23 ± 0.93	1.2 ± 0.84	1.22 ± 0.86	1.27 ± 0.91	1.22 ± 0.84	1.2 ± 0.84	1.3 ± 1.02	1.22 ± 0.91
ConvNeXt	1.25 ± 0.86	1.44 ± 1.25	1.51 ± 2.51	1.35 ± 1.14	1.31 ± 1.01	2.19 ± 4.67	2.29 ± 6.02	2.44 ± 4.86
SAM2.1	1.23 ± 0.97	1.34 ± 1.17	1.48 ± 1.74	1.39 ± 1.48	1.22 ± 0.82	1.27 ± 0.84	1.25 ± 0.89	1.26 ± 0.84
MedSAM2	1.27 ± 0.97	1.34 ± 1.15	1.49 ± 1.89	1.56 ± 1.97	1.25 ± 0.83	1.23 ± 0.84	1.23 ± 0.84	1.23 ± 0.83

TABLE IX

QUANTITATIVE COMPARISON OF PRETRAINED MODELS, FEATURE DISTANCE METRICS, AND FEATURE EXTRACTION LEVELS USED WITHIN IMPACT ON TASK 5, REPORTING THE MEAN AND STANDARD DEVIATION OF THE TRE. INTENSITY-BASED METRICS ARE ALSO REPORTED FOR REFERENCE.

Intensity-based metrics								
MSE			NCC			NMI		
9.63 ± 10.37			5.15 ± 8.2			2.91 ± 3.77		
Per-Model Evaluation								
Model	First Layer				Second Layer			
	L1	L2	NCC	Cosine	L1	L2	NCC	Cosine
MIND (R1D2)	1.75 ± 3.18	1.88 ± 3.73	1.9 ± 3.82	25.63 ± 32.67				
M258	1.95 ± 2.91	1.94 ± 2.82	2.11 ± 2.84	2.35 ± 3.35	1.36 ± 1.87	1.32 ± 1.83	1.42 ± 1.97	1.34 ± 1.92
M291	2.07 ± 2.96	2.36 ± 3.85	2.9 ± 3.67	3.31 ± 3.86	1.66 ± 2.6	1.62 ± 2.68	1.69 ± 2.44	3.41 ± 5.69
M730	1.5 ± 2.25	1.4 ± 2.07	1.73 ± 2.44	1.81 ± 2.6	1.43 ± 2.01	1.47 ± 2.32	1.56 ± 2.41	1.69 ± 3.18
ConvNeXt	9.12 ± 5.7	8.07 ± 7.78	7.45 ± 6.78	9.12 ± 5.7	7.14 ± 5.78	7.71 ± 6.25	7.42 ± 5.84	8.19 ± 5.66
SAM2.1	2.52 ± 4.45	2.59 ± 4.95	2.5 ± 4.56	6.19 ± 9.28	1.48 ± 2.37	1.56 ± 2.58	1.49 ± 2.29	1.74 ± 2.85
MedSAM2	2.49 ± 4.38	2.62 ± 5.18	2.53 ± 4.52	6.16 ± 9.0	1.48 ± 2.36	1.57 ± 2.59	1.5 ± 2.3	1.72 ± 2.83

However, in our experiments, it showed inconsistent improvements and higher variability across models and tasks, limiting its robustness.

In contrast, L1 and L2 distances, which measure absolute differences in feature intensity, systematically achieved the best results across all configurations (Tables VIII and IX). This was especially pronounced when using deeper semantic features.

C. Impact of the Feature Extraction Level

Figure 6 illustrates the hierarchical nature of feature extraction within the nnU-Net encoder. The feature maps, visualized as RGB images using PCA, show the progressive abstraction from fine-grained spatial details in early layers to more semantic representations in deeper layers. To evaluate the impact of this hierarchy on registration performance, features from the first and second encoder layers were tested independently. As reported in Tables VIII and IX, the extraction level significantly influences TRE. In most cases, features from the second layer lead to lower errors and more stable results compared to those from the first layer.

VII. DISCUSSION

This work introduced IMPACT, a semantic similarity metric for multimodal medical image registration that leverages semantic features extracted from large-scale pretrained segmentation models as a generic, modality-agnostic feature space. The metric was integrated into both algorithmic and deep learning-based registration frameworks, and evaluated across five diverse tasks involving different anatomical regions (thorax, abdomen, pelvis) and imaging modalities (CT, CBCT, MRI).

Experimental results demonstrate that semantic features significantly improve registration performance, particularly in the presence of modality-specific artifacts and nonlinear intensity relationships. IMPACT was systematically compared to standard similarity metrics (MSE, NCC, MI, MIND) and several state-of-the-art methods from public challenges. Consistent improvements in TRE, DSC, and HD95 were achieved across all tasks without task-specific training or tuning. By leveraging features from pretrained models like TotalSegmentator and SAM2.1, IMPACT provides a robust alternative to conventional intensity-based metrics and handcrafted descrip-




Fig. 6. Visualization of the aligned fixed and moving images (first column), feature maps extracted at each encoder stage of the nnU-Net encoder (columns 1–6), and the Softmax output (column 7) for the pretrained models M258 and M291. The feature maps are displayed as RGB images, where the three components are obtained using PCA, retaining the top three principal components for visualization. The resolution (in mm), the number of features, and the receptive field size are indicated below each column in the format resolution|receptive field (as 32 | 5). This representation highlights the hierarchical feature extraction process, with increasing spatial context and abstraction at successive encoder stages.

tors. The plug-and-play integration of these features within both algorithmic and deep learning frameworks ensures broad applicability across diverse datasets and clinical contexts.

In both thoracic CBCT registration tasks (Tasks 1 and 2), strong modality-specific artifacts, such as scatter, truncation, and beam hardening, compromise the reliability of intensity-based metrics, particularly in peripheral and low-contrast regions, limiting the effectiveness of measures like NCC and MI. On Task 1, IMPACT (Jacobian) achieved a median TRE of 1.20 mm [0.85–1.71], outperforming NCC (1.58 mm [1.06–2.92]) and MIND (1.25 mm [0.93–1.85]) (Table III). Qualitative results (Fig. 2) confirm improved anatomical alignment. On Task 2 (Learn2Reg ThoraxCBCT), IMPACT achieved a DSC of 0.652, comparable to Deeds (0.648), which combines MIND with a task-specific discrete optimization framework, and higher than ConvexAdam (0.586), which also uses MIND but with a generic hybrid strategy. This performance secured 6th place out of 21 submissions in the MICCAI-ThoraxCBCT challenge (Table IV). These results underscore that, while MIND provides a strong handcrafted baseline, its performance is sensitive to the optimization framework, whereas IMPACT offers generalizable accuracy without task-specific engineering.

In the Lung CT registration task (Task 5), characterized by substantial anatomical variability and large respiratory deformations, IMPACT combined with Elastix achieved competitive performance, ranking third on the final leaderboard of the MICCAI-LungCT challenge with a mean TRE of 1.89 mm (Table VI). It ranked ahead of the SAME method, which is specifically designed for monomodal registration. A detailed test set analysis (Table IX) showed that the best configuration achieved a TRE of 1.32 ± 1.93 mm, outperforming MIND (1.75 ± 3.18 mm) and all intensity-based baselines.

The observed improvements on these 3 tasks are largely attributable to the early layers of pretrained segmentation networks, which act as learned feature denoisers. These lay-

ers effectively suppress high-frequency noise and modality-specific artifacts, while preserving and enhancing fine anatomical details, such as pulmonary fissures, airway bifurcations, and vascular trees, key landmarks that remain stable despite significant deformations, as shown in Fig. 6. In contrast to intensity-based metrics or handcrafted descriptors, these learned features produce denoised and meaningful anatomical representations, suitable for guiding registration in highly deformable settings.

In the MR/CT tasks (Tasks 3 and 4), IMPACT again demonstrated strong performance, despite substantial intensity differences and anatomical variability between modalities. On the abdomen (Task 3), VoxelMorph combined with IMPACT ranked second out of 26 submissions, outperforming several established methods in both DSC and HD95 (Table V). On the pelvis (Task 4), IMPACT consistently outperformed MI and MIND across all structures (Fig. 5). The improvement was particularly pronounced for the prostate, where intensity and structural differences between MR and CT are most significant, indicating the method’s ability to align semantically corresponding anatomy despite weak appearance similarity.

These results are consistent with results from the Thorax CBCT/CT tasks and confirm that IMPACT’s structure-alignment capability is not modality-specific. In MR/CT scenarios, performance gains are primarily driven by deep semantic features extracted from the final layers of pretrained segmentation models, which encode high-level anatomical concepts such as organ identity and shape rather than local appearance. These features remain effective despite pronounced differences in intensity and texture between modalities, as anatomical structures remain consistent. By projecting both MR and CT volumes into a shared semantic space, IMPACT enables anatomically meaningful correspondence even in the absence of intensity similarity, something traditional metrics MI and MIND fail to achieve.

The results demonstrate that IMPACT consistently matches

or surpasses learned and handcrafted methods such as Convex-Adam, SAME, and DEEDS on challenging public registration benchmarks, all while using the widely adopted Elastix framework and requiring no training or task-specific tuning.

As shown in Fig. 4, IMPACT in Jacobian mode achieves the highest final accuracy and significantly faster convergence than both static mode and traditional intensity-based metrics. By 200 iterations, Jacobian mode reaches a low median TRE with reduced variability, while intensity-based methods converge more slowly and remain less accurate. Although static mode improves over traditional metrics, its convergence is slower and more variable. These results show that Jacobian optimization provides more informative gradients, accelerating convergence and stabilizing alignment early. However, Jacobian mode is more computationally demanding, especially with deep feature extractors like SAM (Table VII). Static mode is more efficient but depends strongly on the spatial resolution of the feature maps, becoming ineffective with deeper layers or with transformer-based models.

The effectiveness of semantic similarity-based registration depends strongly on the choice of pretrained model used for feature extraction. Models trained for segmentation tasks, as opposed to generic classification, produced more anatomically structured and spatially coherent features. This distinction was consistently reflected in Tasks 1 and 5, where segmentation-based models outperformed classification networks such as ConvNeXt.

Within segmentation models, performance was further enhanced when the pretraining dataset included anatomically relevant structures. For example, M258, trained specifically on CT thoracic anatomy, outperformed the more generic M730 in corresponding tasks, suggesting that anatomical specificity during pretraining improves the relevance of extracted features for alignment.

In contrast, both the architectural design (CNN versus transformer) and the domain on which the models were pretrained had a comparatively limited influence. Despite being trained on natural images, the transformer-based SAM2.1 achieved comparable performance to the CNN-based TotalSegmentator, which was trained on medical images. This finding is significant as it demonstrates that sufficiently expressive architectures can learn semantic representations that are general enough to transfer across domains. Furthermore, ablation studies comparing SAM2.1 and MedSAM fine-tuned on medical images, showed identical performance, indicating that domain adaptation had little impact on feature quality. This underscores that the semantic richness introduced by segmentation-based pretraining plays a more decisive role than the pretraining domain itself. These results suggest that models like SAM2.1 can act as truly general-purpose feature extractors, even when applied beyond their original domain.

The ablation studies showed that registration performance is sensitive to the choice of distance metric. Simple magnitude-based measures (L1, L2) consistently outperformed cosine similarity and NCC, especially with deeper features (Tables VIII, IX). These results indicate that the semantic features bring both images into a common representation space, where the absolute differences provide a significant alignment

indicator.

The ablation studies further highlighted the impact of feature extraction depth on registration performance. Across tasks, features from the second encoder layer generally yielded superior results (Tables VIII, IX). However, the optimal layer was task-dependent: early encoder layers proved most effective in CT/CBCT registration by suppressing modality-specific artifacts while preserving spatial detail, whereas decoder layers performed better in MR/CT tasks by capturing high-level anatomical semantics necessary to bridge large appearance discrepancies. This behavior is illustrated in Fig. 5, particularly for complex structures such as the prostate.

This study shows that features from large-scale pretrained segmentation models are effective for multimodal image registration, offering strong generalization across tasks without requiring task-specific training. However, a natural trade-off exists between generality and task specificity, as some anatomically relevant information may not be fully leveraged when features are used as-is, especially since the choice of the pretrained model influences the nature and granularity of the extracted features, potentially limiting performance in highly specialized scenarios. To address this limitation, future work could enhance semantic exploitation by developing adaptive feature-weighting methods or training lightweight embedding networks to model task-specific relationships.

Beyond its performance, IMPACT is also a modular platform that enables systematic experimentation with pretrained models, feature layers, and distance functions. It allows the exploration of feature extraction methods in both algorithmic and deep learning-based registration approaches. This flexibility is especially valuable for challenging modalities like PET or functional MRI, where semantic guidance can help overcome low anatomical contrast or complex intensity variations.

VIII. CONCLUSION

This study introduces IMPACT, a novel similarity metric for multimodal medical image registration that leverages pretrained segmentation models to guide alignment via high-level anatomical features. By integrating feature representations from TotalSegmentator and SAM-based models, IMPACT addresses key limitations of traditional similarity measures, including sensitivity to noise, artifacts, and modality-specific intensity variations. Integrated into both algorithmic (Elastix) and deep learning-based (VoxelMorph) frameworks, IMPACT consistently improved alignment accuracy over traditional intensity-based and handcrafted metrics across multiple datasets.

While IMPACT demonstrated robust performance, particularly in tasks involving substantial modality variation or artifacts, its effectiveness depends on the choice of feature extractor and distance function. Ablation studies highlighted that pretrained segmentation models, such as TotalSegmentator and SAM2.1, provide superior features for guiding registration, with L1 and L2 losses yielding the most consistent results.

Moreover, based on our conclusions, we propose the pretrained model M730 with 2 layers and a L2 features distance measure as the recommended generic feature extractor for

IMPACT, balancing robustness and efficiency across different registration tasks. However, if the user wishes to further optimize performance for a specific task, the feature extractors in IMPACT can be replaced with a custom model trained on their own data. This allows for more task-specific features while maintaining the versatility of IMPACT, ensuring seamless integration with frameworks like Elastix and VoxelMorph. This adaptability reinforces IMPACT’s role as a scalable and efficient tool, suitable for both research and clinical applications.

The integration of semantic feature alignment into image registration holds considerable promise for clinical applications, including tumor monitoring, surgical planning, and image-guided therapy. The plug-and-play nature of IMPACT, combined with its scalability across multimodal datasets, ensures broad applicability and ease of adoption. Future work should focus on extending generalizability to additional modalities and developing strategies for adaptive feature selection tailored to specific clinical contexts.

IX. CODE AVAILABILITY AND INTEGRATION DETAILS

The resources for implementing the IMPACT similarity metric, including the PyTorch-based loss function and a version of Elastix with IMPACT, are publicly available for reproducibility and community use at: <https://github.com/vbousso/ImpactLoss.git>. Both resources are maintained under open licenses to support research applications. Detailed explanations of the methodology and its integration into registration pipelines are provided in this article and the associated repository documentation. These tools are designed to enable users to reproduce and adapt the framework to their specific use cases, ensuring both accessibility and extensibility.

ACKNOWLEDGMENT

The work presented in this article was supported by the Brittany Region through its Allocations de Recherche Doctorale framework and by the French National Research Agency as part of the VATSop project (ANR-20-CE19-0015). Additionally, it was supported by a PhD scholarship Grant from Elekta AB (C.Hémond). The authors have no relevant financial or non-financial interests to disclose. While preparing this work, the authors used ChatGPT to enhance the writing structure and refine grammar. After using these tools, the authors reviewed and edited the content as needed and took full responsibility for the publication’s content.

REFERENCES

- [1] X. Chen, X. Wang, K. Zhang, K.-M. Fung, T. C. Thai, K. Moore, R. S. Mannel, H. Liu, B. Zheng, Y. Qiu, Recent advances and clinical applications of deep learning in medical image analysis, *Medical Image Analysis* 79 (2022) 102444.
- [2] M. A. Viergever, J. B. A. Maintz, S. Klein, K. Murphy, M. Staring, J. P. W. Pluim, A survey of medical image registration – under review, *Medical Image Analysis* 33 (2016) 140–144.
- [3] A. V. Dalca, G. Balakrishnan, J. Guttag, M. R. Sabuncu, Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces, *Medical Image Analysis* 57 (2019) 226–236.
- [4] B. D. De Vos, F. F. Berendsen, M. A. Viergever, H. Sokootti, M. Staring, I. Išgum, A deep learning framework for unsupervised affine and deformable image registration, *Medical Image Analysis* 52 (2019) 128–143.
- [5] J. Chen, E. C. Frey, Y. He, W. P. Segars, Y. Li, Y. Du, Transmorph: Transformer for unsupervised medical image registration, *Medical image analysis* 82 (2022) 102615.
- [6] C. Studholme, D. L. Hill, D. J. Hawkes, Automated 3-d registration of mr and ct images of the head, *Medical image analysis* 1 (2) (1996) 163–175.
- [7] P. Viola, W. M. Wells III, Alignment by maximization of mutual information, *International journal of computer vision* 24 (2) (1997) 137–154.
- [8] M. P. Heinrich, M. Jenkinson, M. Bhushan, T. Martin, F. V. Gleeson, S. M. Brady, J. A. Schnabel, MIND: Modality independent neighbourhood descriptor for multi-modal deformable registration, *Medical Image Analysis* 16 (7) (2012) 1423–1435.
- [9] M. P. Heinrich, M. Jenkinson, B. W. Papież, S. M. Brady, J. A. Schnabel, Towards realtime multimodal fusion for image-guided interventions using self-similarities (2013).
- [10] Z. Gao, B. Gu, J. Lin, Monomodal image registration using mutual information based methods, *Image and Vision Computing* 26 (2) (2008) 164–173.
- [11] G. P. Penney, J. Weese, J. A. Little, P. Desmedt, D. L. Hill, et al., A comparison of similarity measures for use in 2-d-3-d medical image registration, *IEEE transactions on medical imaging* 17 (4) (1998) 586–595.
- [12] A. Hering, L. Hansen, T. C. Mok, A. C. Chung, H. Siebert, S. Häger, A. Lange, S. Kuckertz, S. Heldmann, W. Shao, et al., Learn2reg: comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning, *IEEE Transactions on Medical Imaging* 42 (3) (2022) 697–712.
- [13] S. Shirraishi, M. P. Grams, L. E. Fong de los Santos, Image-guided radiotherapy quality control: statistical process control using image similarity metrics, *Medical Physics* 45 (5) (2018) 1811–1821.
- [14] G. Haskins, J. Kruecker, U. Kruger, S. Xu, P. A. Pinto, B. J. Wood, P. Yan, Learning deep similarity metric for 3D MR–TRUS image registration, *International Journal of Computer Assisted Radiology and Surgery* 14 (3) (2018) 417–425.
- [15] M. Ronchetti, W. Wein, N. Navab, O. Zettinig, R. Prevost, Disa: Differentiable similarity approximation for universal multimodal registration (2023).
- [16] S. Czolbe, P. Pegios, O. Krause, A. Feragen, Semantic similarity metrics for image registration, *Medical Image Analysis* 87 (2023) 102830.
- [17] J. Chen, Y. Liu, S. Wei, Z. Bian, S. Subramanian, A. Carass, J. L. Prince, Y. Du, A survey on deep learning in medical image registration: New technologies, uncertainty, evaluation metrics, and beyond, *Medical Image Analysis* 100 (2025) 103385.
- [18] F. Liu, K. Yan, A. P. Harrison, D. Guo, L. Lu, A. L. Yuille, L. Huang, G. Xie, J. Xiao, X. Ye, et al., Same: Deformable image registration based on self-supervised anatomical embeddings (2021).
- [19] X. Song, X. Xu, P. Yan, Dino-reg: General purpose image encoder for training-free multi-modal deformable medical image registration (2024).
- [20] H. Siebert, C. Großbröhmer, L. Hansen, M. P. Heinrich, Convexadam: Self-configuring dual-optimization-based 3d multitask medical image registration, *IEEE Transactions on Medical Imaging* (2025).
- [21] J. Johnson, A. Alahi, L. Fei-Fei, Perceptual losses for real-time style transfer and super-resolution (2016).
- [22] G. G. Pihlgren, K. Nikolaidou, P. C. Chhipa, N. Abid, R. Saini, F. Sandin, M. Liwicki, A systematic performance analysis of deep perceptual loss networks: Breaking transfer learning conventions, *arXiv preprint arXiv:2302.04032* (2023).
- [23] J. Wasserthal, H.-C. Breit, M. T. Meyer, M. Pradella, D. Hinck, A. W. Sauter, T. Heye, D. T. Boll, J. Cyriac, S. Yang, et al., Totalsegmentator: Robust segmentation of 104 anatomic structures in ct images, *Radiology: Artificial Intelligence* 5 (5) (2023) e230024.
- [24] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, et al., Segment anything (2023).
- [25] S. Klein, M. Staring, K. Murphy, M. A. Viergever, J. P. W. Pluim, Elastix: A toolbox for intensity-based medical image registration, *IEEE Transactions on Medical Imaging* 29 (1) (2010) 196–205.
- [26] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, A. V. Dalca, VoxelMorph: A learning framework for deformable medical image registration, *IEEE Transactions on Medical Imaging* 38 (8) (2019) 1788–1800.
- [27] A. Reithmeir, V. Spieker, V. Sideri-Lampretsa, D. Rueckert, J. A. Schnabel, V. A. Zimmer, From model based to learned regularization in medical image registration: A comprehensive review, *arXiv preprint arXiv:2412.15740* (2024).

- [28] D. Rueckert, L. Sonoda, C. Hayes, D. Hill, M. Leach, D. Hawkes, Nonrigid registration using free-form deformations: application to breast mr images, *IEEE Transactions on Medical Imaging* 18 (8) (1999) 712–721.
- [29] J.-P. Thirion, Image matching as a diffusion process: an analogy with maxwell's demons, *Medical image analysis* 2 (3) (1998) 243–260.
- [30] M. F. Beg, M. I. Miller, A. Trounev, L. Younes, Computing large deformation metric mappings via geodesic flows of diffeomorphisms, *International journal of computer vision* 61 (2005) 139–157.
- [31] B. B. Avants, C. L. Epstein, M. Grossman, J. C. Gee, Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain, *Medical image analysis* 12 (1) (2008) 26–41.
- [32] B. B. Avants, N. Tustison, G. Song, et al., Advanced normalization tools (ants), *Insight j* 2 (365) (2009) 1–35.
- [33] M. P. Heinrich, M. Jenkinson, M. Brady, J. A. Schnabel, Mrf-based deformable registration and ventilation estimation of lung ct, *IEEE transactions on medical imaging* 32 (7) (2013) 1239–1248.
- [34] M. P. Heinrich, H. Handels, I. J. Simpson, Estimating large lung motion in copd patients by symmetric regularised correspondence fields (2015).
- [35] R. Jena, D. Sethi, P. Chaudhari, J. Gee, Deep learning in medical image registration: Magic or mirage?, *Advances in Neural Information Processing Systems* 37 (2025) 108331–108353.
- [36] W. M. Wells III, P. Viola, H. Atsumi, S. Nakajima, R. Kikinis, Multimodal volume registration by maximization of mutual information, *Medical image analysis* 1 (1) (1996) 35–51.
- [37] C. Studholme, D. L. Hill, D. J. Hawkes, An overlap invariant entropy measure of 3d medical image alignment, *Pattern recognition* 32 (1) (1999) 71–86.
- [38] J. P. Pluim, J. A. Maintz, M. A. Viergever, Image registration by maximization of combined mutual information and gradient information (2000).
- [39] R. Han, C. Jones, J. Lee, X. Zhang, P. Wu, P. Vagdargi, A. Uneri, P. Helm, M. Luciano, W. Anderson, et al., Joint synthesis and registration network for deformable mr-cbct image registration for neurosurgical guidance, *Physics in Medicine & Biology* 67 (12) (2022) 125008.
- [40] C. Hémon, B. Texier, H. Chourak, A. Simon, I. Bessières, R. de Crevoisier, J. Castelli, C. Lafond, A. Barateau, J.-C. Nunes, Indirect deformable image registration using synthetic image generated by unsupervised deep learning, *Image and Vision Computing* 148 (2024) 105143.
- [41] P. Yan, S. Xu, A. R. Rastinehad, B. J. Wood, Adversarial image registration with application for mr and trus image fusion (2018).
- [42] M. Hoffmann, B. Billot, D. N. Greve, J. E. Iglesias, B. Fischl, A. V. Dalca, Synthmorph: learning contrast-invariant registration without acquired images, *IEEE transactions on medical imaging* 41 (3) (2021) 543–558.
- [43] E. Haber, J. Modersitzki, Intensity gradient based registration and fusion of multi-modal images (2006).
- [44] J. Ma, Y. He, F. Li, L. Han, C. You, B. Wang, Segment anything in medical images, *Nature Communications* 15 (1) (2024) 654.
- [45] Z. Huang, H. Wang, Z. Deng, J. Ye, Y. Su, H. Sun, J. He, Y. Gu, L. Gu, S. Zhang, et al., Stu-net: Scalable and transferable medical image segmentation models empowered by large-scale supervised pre-training, *arXiv preprint arXiv:2304.06716* (2023).
- [46] M. McCormick, X. Liu, J. Jomier, C. Marion, L. Ibanez, Itk: enabling reproducible research and open science, *Frontiers in neuroinformatics* 8 (2014) 13.
- [47] S. Klein, J. P. W. Pluim, M. Staring, M. A. Viergever, Adaptive stochastic gradient descent optimisation for image registration, *International Journal of Computer Vision* 81 (3) (2009) 227–239.
- [48] R. Bollapragada, R. Byrd, J. Nocedal, Adaptive sampling strategies for stochastic optimization, *SIAM Journal on Optimization* 28 (4) (2018) 3312–3343.
- [49] Y. Qiao, B. P. Lelieveldt, M. Staring, An efficient preconditioner for stochastic gradient descent optimization of image registration, *IEEE transactions on medical imaging* 38 (10) (2019) 2314–2325.
- [50] Z. Wu, E. Rietzel, V. Boldea, D. Sarrut, G. C. Sharp, Evaluation of deformable registration of patient lung 4dct with subanatomical region segmentations: Evaluation of deformable registration of 4dct with segmentations, *Medical Physics* 35 (2) (2008) 775–781.
- [51] J. Vandemeulebroucke, O. Bernard, S. Rit, J. Kybic, P. Clarysse, D. Sarrut, Automated segmentation of a motion mask to preserve sliding motion in deformable registration of thoracic ct: Motion mask for deformable registration of thoracic ct, *Medical Physics* 39 (2) (2012) 1006–1015.
- [52] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., Pytorch: An imperative style, high-performance deep learning library, *Advances in neural information processing systems* 32 (2019).
- [53] S. Rouzé, P. A. Alvarez, B. de Latour, E. Flécher, J.-L. Dillenseger, J.-P. Verhoye, Localisation de nodules pulmonaires en réalité augmentée grâce au cone beam computed tomography (cbct) en vidéo-thoroscopie, *Bulletin de l'Académie Nationale de Médecine* 202 (8–9) (2018) 1897–1908.
- [54] P. Alvarez, S. Rouzé, M. I. Miga, Y. Payan, J.-L. Dillenseger, M. Chabanas, A hybrid, image-based and biomechanics-based registration approach to markerless intraoperative nodule localization during video-assisted thoracoscopic surgery, *Medical image analysis* 69 (2021) 101983.
- [55] V. Boussot, J.-L. Dillenseger, Statistical model for the prediction of lung deformation during video-assisted thoracoscopic surgery (2023).
- [56] R. Schulze, U. Heil, D. Gross, D. D. Bruellmann, E. Dranschnikow, U. Schwanecke, E. Schoemer, Artefacts in CBCT: a review, *Dento Maxillo Facial Radiology* 40 (5) (2011) 265–273.
- [57] P. Alvarez, M. Chabanas, S. Sikora, S. Rouzé, Y. Payan, J.-L. Dillenseger, Measurement and analysis of lobar lung deformation after a change of patient position during video-assisted thoracoscopic surgery, *IEEE Transactions on Biomedical Engineering* 70 (3) (2022) 931–940.
- [58] G. D. Hugo, E. Weiss, W. C. Sleeman, S. Balik, P. J. Keall, J. Lu, J. F. Williamson, A longitudinal four-dimensional computed tomography and cone beam computed tomography dataset for image-guided radiation therapy research in lung cancer, *Medical Physics* 44 (2) (2017) 762–771.
- [59] K. Clark, B. Vendt, K. Smith, J. Freymann, J. Kirby, P. Koppel, S. Moore, S. Phillips, D. Maffitt, M. Pringle, L. Tarbox, F. Prior, The cancer imaging archive (TCIA): Maintaining and operating a public information repository, *Journal of Digital Imaging* 26 (6) (2013) 1045–1057.
- [60] J. A. Dowling, J. Sun, P. Pichler, D. Rivest-Hénault, S. Ghose, H. Richardson, C. Wratten, J. Martin, J. Arm, L. Best, et al., Automatic substitute computed tomography generation and contouring for magnetic resonance imaging (mri)-alone external beam radiation therapy from standard mri sequences, *International Journal of Radiation Oncology* Biology* Physics* 93 (5) (2015) 1144–1153.
- [61] A. Hering, K. Murphy, B. van Ginneken, Learn2reg challenge: Ct lung registration - training data (May 2020).
- [62] Z. Liu, H. Mao, C.-Y. Wu, C. Feichtenhofer, T. Darrell, S. Xie, A convnet for the 2020s (2022).
- [63] C. Ryali, Y.-T. Hu, D. Bolya, C. Wei, H. Fan, P.-Y. Huang, V. Aggarwal, A. Chowdhury, O. Poursaeed, J. Hoffman, et al., Hier: A hierarchical vision transformer without the bells-and-whistles (2023).

APPENDIX

A. Pretrained Models Utilized

TABLE X

LIST OF PRETRAINED MODELS USED IN THIS STUDY, COVERING A RANGE OF FEATURE EXTRACTION PARADIGMS, FROM HANDCRAFTED DESCRIPTORS (MIND) AND 2D FOUNDATION MODELS (SAM2.1, MedSAM2) TO CLASSIFICATION NETWORKS (CONVNEXT) AND SEGMENTATION MODELS PRETRAINED ON MEDICAL DATASETS. THE FIELD OF VIEW INDICATES THE SPATIAL RECEPTIVE FIELD AT THE SELECTED EXTRACTION LAYER.

Model	Specialization	Field of View
MIND	A handcrafted model designed to extract MIND descriptors [8]. It is parameterized by a radius r and a dilation d .	$2rd + 1$
SAM2.1	Segment Anything Model (SAM) 2.1 [24], a 2D foundation model designed for general segmentation tasks.	29
MedSAM2	A medical adaptation of SAM [44], specifically fine-tuned for medical image segmentation.	29
ConvNeXt	ConvNeXt Tiny [62], a 2D model trained for a classification task on natural images.	13
M258	Specialized in lung vessel segmentation.	$2^l + 3$ where l is the layer number
M291	Focused on organ segmentation, particularly structures such as lung lobes and the liver.	
M292	Designed for vertebrae segmentation.	
M293	Tailored for cardio-respiratory system segmentation, including structures such as the aorta and myocardium.	
M294	Optimized for muscle segmentation, including brain structures.	
M295	Dedicated to rib segmentation.	
M297	CT model trained at a 3mm resolution	
M298	CT model trained at a 6mm resolution	
M730	Designed for organ segmentation in MRI and CT imaging.	
M731	Specialized in muscle segmentation in MRI and CT.	
M732	MRI and CT model trained at a 3mm resolution.	
M733	MRI and CT model trained at a 6mm resolution.	

1) *TotalSegmentator Models*: All TotalSegmentator models share a common autoencoder architecture with skip connections. The downsampling process is performed using stride convolutions (stride = 2), while upsampling is achieved through transposed convolutions. Each resolution level consists of two convolutional blocks, with feature dimensions progressively increasing as [32, 64, 128, 320, 320]. Each block comprises a convolutional layer, followed by instance normalization and a Leaky ReLU activation function. Layers are extracted at each level.

Preprocessing for TotalSegmentator models:

- CT images: Standardized with a canonical orientation, intensities clipped to the range $[-1024, 276]$, and normalized to a mean of -370 and a variance of 436 .
- MRI images: Standardized with canonical orientation and intensity standardization.

2) *SAM2.1 and MedSAM2*: The encoders of SAM2.1 and MedSAM2 are built upon Hiera, a hierarchical vision transformer architecture designed for efficient multi-scale feature extraction [63].

Hiera is a hierarchical transformer that refines self-attention across multiple scales. It processes images as non-overlapping patches, extracts multi-scale features using transformer blocks, and progressively reduces spatial resolution while enriching feature representations. Unlike standard vision transformers, Hiera applies local attention at lower levels and maintains global context at higher levels, improving efficiency. Skip connections help merge features across scales, preserving spatial details while enabling hierarchical abstraction. Features are extracted at each downsampling stage.

Preprocessing for models pretrained on natural images.

- SAM2.1, MedSAM2 and ConvNeXt utilize ImageNet normalization, where image intensities are standardized using a mean and variance derived from ImageNet statistics to ensure consistency in feature representation.

B. Hyperparameters of the IMPACT Similarity Metric

Table XI lists the hyperparameters available for configuring the IMPACT metric. These settings govern the sampling strategy, multi-resolution scheme, feature extraction, and distance computation. Unless stated otherwise, the same parameters were used for both fixed and moving images.

TABLE XI
HYPERPARAMETERS FOR THE IMPACT METRIC USED IN ELASTIX-BASED REGISTRATION. PARAMETERS COVER SPATIAL SAMPLING, FEATURE EXTRACTION, AND SIMILARITY COMPUTATION.

Parameter	Description	Default value
MaximumNumberOfIterations	Number of iterations for the optimization process at each resolution.	500
NumberOfSpatialSamples	Number of spatial samples used during optimization.	2000
NumberOfResolutions	Number of multi-resolution levels used in the registration process.	3
FinalGridSpacingInPhysicalUnits	Grid spacing for the final resolution in physical units.	8
ModelsPath	Path to the pretrained models used for feature extraction.	"Path"
Dimension	Dimensionality of the input images used by the model (2D or 3D).	3
NumberOfChannels	Number of input channels in the model (like grayscale = 1, RGB = 3).	1
PatchSize	Size of the patches extracted for processing.	5*5*5
VoxelSize	Resampled voxel size for input images.	1.5*1.5*1.5
LayersMask	Binary mask indicating which feature extractor layers are selected.	1
SubsetFeatures	Number of selected feature channels.	32
LayersWeight	Weight assigned to the selected layers during feature extraction.	1
Mode	Mode selection for feature extraction (Static, Jacobian).	Jacobian
GPU	GPU configuration (CPU mode (-1) or specific GPU device selection).	-1
FeaturesMapUpdateInterval	Frequency of feature map updates during optimization.	-1
PCA	Number of principal components retained for the PCA.	0
Loss	Choose loss comparison for each layer (L1, L2, NCC, Cosine, L1Cosine).	L2

TABLE XII
COMPARISON OF MULTI-RESOLUTION STRATEGIES AND THEIR IMPACT ON REGISTRATION ACCURACY (MEASURED WITH TRE) ON TASK 1 AFTER 500 ITERATIONS.

Multi-resolution strategy	TRE (mm)	
	Jacobian	Static
Gaussian smoothing + downsampling (standard multi-resolution)	1.20	1.22
Downsampling only (no smoothing)	1.24	1.24
Gaussian smoothing only (no downsampling)	1.43	1.46
No image pyramid (full-resolution only)	1.94	3.27

C. Choice of Multi-Resolution Strategy with Extracted Features

We conducted an ablation study on Task 1 (Thorax CT/CBCT) to assess the impact of various multi-resolution strategies on registration accuracy using deep features for similarity computation. Table XII shows the median TRE after 500 iterations with the second-layer model M258, employing L2 distance for feature similarity computation, across four configurations: downsampling only, Gaussian smoothing only, a combination of both (standard multi-resolution), and no pyramid (full resolution only).

The results show that combining Gaussian smoothing and downsampling yields the most accurate performance, especially in Jacobian mode. Additionally, downsampling alone, without smoothing, still produces competitive results, suggesting that spatial rescaling is sufficient on its own. In contrast, operating at full resolution (without the pyramid) results in significantly poorer performance. These findings highlight the advantage of using a four-level multi-resolution strategy, which combines both smoothing and downsampling, as implemented in all Elastix-based experiments.