



HAL
open science

Privacy Preserving Image Registration

Riccardo Taiello, Melek Önen, Francesco Capano, Olivier Humbert, Marco Lorenzi

► **To cite this version:**

Riccardo Taiello, Melek Önen, Francesco Capano, Olivier Humbert, Marco Lorenzi. Privacy Preserving Image Registration. *Medical Image Analysis*, 2024, 94, pp.103129. 10.1016/j.media.2024.103129 . hal-04861654

HAL Id: hal-04861654

<https://inria.hal.science/hal-04861654v1>

Submitted on 4 Jan 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.



Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0 International License



Privacy Preserving Image Registration

Riccardo Taiello^{a,b,c,*}, Melek Önen^b, Francesco Capano^b, Olivier Humbert^c, Marco Lorenzi^{a,c}

^aEpione Research Group, Inria, Sophia Antipolis, France

^bEURECOM, France

^cUniversité Côte d'Azur, France

ARTICLE INFO

Article history:

Keywords: Image Registration, Image Registration, Trustworthiness

ABSTRACT

Image registration is a key task in medical imaging applications, allowing to represent medical images in a common spatial reference frame. Current approaches to image registration are generally based on the assumption that the content of the images is usually accessible in clear form, from which the spatial transformation is subsequently estimated. This common assumption may not be met in practical applications, since the sensitive nature of medical images may ultimately require their analysis under privacy constraints, preventing to openly share the image content. In this work, we formulate the problem of image registration under a privacy preserving regime, where images are assumed to be confidential and cannot be disclosed in clear. We derive our privacy preserving image registration framework by extending classical registration paradigms to account for advanced cryptographic tools, such as secure multi-party computation and homomorphic encryption, that enable the execution of operations without leaking the underlying data. To overcome the problem of performance and scalability of cryptographic tools in high dimensions, we propose several techniques to optimize the image registration operations by using gradient approximations, and by revisiting the use of homomorphic encryption through packing, to allow the efficient encryption and multiplication of large matrices. We focus on registration methods of increasing complexity, including rigid, affine, and non-linear registration based on cubic splines or diffeomorphisms parametrized by time-varying velocity fields. In all these settings, we demonstrate how the registration problem can be naturally adapted for accounting to privacy-preserving operations, and illustrate the effectiveness of PPIR on a variety of registration tasks.

© 2025 Elsevier B. V. All rights reserved.

1. Introduction

Image Registration is a crucial task in medical imaging applications, allowing to spatially align imaging features between two or multiple scans. Registration methods are today a central component of state-of-the-art methods for atlas-based segmentation (Shattuck et al., 2009; Cardoso et al., 2013), morphological and functional analysis (Dale et al., 1999; Ashburner and

Friston, 2000), multi-modal data integration (Heinrich et al., 2011), and longitudinal analysis (Reuter et al., 2010; Ashburner and Ridgway, 2013). Typical registration paradigms are based on a given transformation model (e.g. affine or non-linear), a cost function and an associated optimization routine. A large number of image registration approaches have been proposed in the literature over the last decades, covering a variety of assumptions on the spatial transformations, cost functions, image dimensionality and optimization strategy (Schnabel et al., 2016). Image registration is the workhorse of many real-life

*Corresponding author: riccardo.taiello@inria.fr

medical imaging software and applications, including public web-based services for automated segmentation and labelling of medical images. Using these services generally requires uploading and exchanging medical images over the Internet, to subsequently perform image registration with respect to one or multiple (potentially proprietary) atlases. Besides these classical medical imaging use-cases, emerging paradigms for collaborative data analysis, such as Federated Learning (FL) (McMahan *et al.*, 2017), have been proposed to enable analysis of medical images in multicentric scenarios for performing group analysis (Gazula *et al.*, 2021) and distributed machine learning (Kaissis *et al.*, 2021; Zerka *et al.*, 2020). However, in these settings, typical medical imaging tasks such as spatial alignment and downstream operations are generally not possible without disclosing the image information.

Due to the evolving juridical landscape on data protection, medical image analysis tools need to be adapted to guarantee compliance with regulations currently existing in many countries, such as the European General Data Protection Regulation (GDPR)¹, or the US Health Insurance Portability and Accountability Act (HIPAA)². Medical imaging information falls within the realm of personal health data (Lotan *et al.*, 2020) and its sensitive nature should ultimately require the analysis under privacy preserving constraints, for instance by preventing to share the image content in clear form.

Advanced cryptographic tools hold great potential in sensitive data analysis problems (e.g., (Lauter, 2021)). Examples of such approaches are Secure-Multi-Party-Computation (MPC) (Yao, 1982) and Homomorphic Encryption (HE) (Rivest *et al.*, 1978). While MPC allows multiple parties to jointly compute a common function over their private inputs and discover no more than the output of this function, HE enables computation on encrypted data without disclosing either the input data or the result of the computation.

This work presents *privacy-preserving image registration* (PPIR), a new methodological framework allowing image registration under privacy constraints. To this end, we reformulate the image registration problem to integrate cryptographic tools, namely MPC or FHE, thus preserving the privacy of the image data. Due to the well-known scalability issues of such cryptographic techniques, we investigate strategies for the practical use of PPIR. In our experiments, we evaluate the effectiveness of PPIR on a variety of registration tasks and medical imaging modalities. Our results demonstrate the feasibility of PPIR and pave the way for the application of secured image registration in sensitive medical imaging applications.

2. Background

Given images $I, J : \mathbb{R}^d \mapsto \mathbb{R}$, image registration (IR) aims at estimating the parameters θ of a spatial transformation $W_\theta \in \mathbb{R}^d \mapsto \mathbb{R}^d$, either linear or non-linear, maximizing the spatial

Algorithm 1 IR via Gauss-Newton optimization

Input:

- ▷ Moving image I
- ▷ Template image J
- ▷ Distance function f
- ▷ Spatial transformation W_θ , parameterized by θ
- ▷ convergence threshold ϵ

Output:

- ▷ Transformed image $I(W_\theta)$ after convergence is reached

```

1: function IMAGEREGISTRATION( $I, J, W_\theta$ ):
2:    $\theta \leftarrow$  INITIALIZEPARAMETERS()
3:   repeat
4:      $e \leftarrow f[I(W_\theta), J]$ 
5:      $G \leftarrow \frac{\partial f}{\partial \theta}$ 
6:      $H \leftarrow \frac{\partial^2 f}{\partial \theta^2}$ 
7:      $\Delta\theta \leftarrow H^{-1} \cdot G$ 
8:      $\theta \leftarrow \theta + \Delta\theta$ 
9:   until  $\|\Delta\theta\| \leq \epsilon$ 
10:  return  $I(W_\theta), e$ 
11: end function

```

overlap between J and the transformed image $I(W_\theta)$, by minimizing a registration loss function f :

$$\theta^* = \operatorname{argmin}_\theta f(I(W_\theta(x)), J(x)). \quad (1)$$

The loss f can be any similarity measure, e.g., the Sum of Squared Differences (SSD), the negative Mutual Information (MI), or normalized cross correlation (CC). Equation (1) can be typically optimized through gradient-based methods, where the parameters θ are iteratively updated until convergence. In particular, when using a Gauss-Newton optimization scheme (Algorithm 1), the update of the spatial transformation can be computed through Equation (2):

$$\Delta\theta = H^{-1} \cdot G, \quad (2)$$

where $G = \frac{\partial f}{\partial \theta}$ is the Jacobian and $H = \frac{\partial^2 f}{\partial \theta^2}$ the Hessian of f . Besides the Gauss-Newton schemes proposed in the field of IR (Pennec *et al.*, 1999; Modersitzki, 2009), gradient-based techniques are classically adopted to solve the IR task, for example in diffeomorphic image registration problems (Ashburner, 2007; Avants *et al.*, 2011).

In all these cases we consider a scenario with two parties, $party_1$, and $party_2$, whereby $party_1$ owns image I and $party_2$ owns image J . The parties wish to collaboratively optimize the image registration problem without disclosing their respective images to each other. We assume that only $party_1$ has access to the transformation parameters θ and that it is also responsible for computing the update at each optimization step. In what follows, we introduce the basic notation to develop PPIR based on different registration frameworks. We focus on registration methods of increasing complexity, including (i) rigid, (ii) affine, and (iii) non-linear registration based on cubic splines or diffeomorphisms parametrized by time-varying velocity fields (large deformation diffeomorphic metric mapping, LDDMM) (Beg *et al.*, 2005). In all these settings, we demonstrate how the

¹<https://gdpr-info.eu/>

²<https://www.hhs.gov/hipaa/index.html>




Fig. 1: Optimization of SSD loss: proposed framework to compute matrix-vector multiplication $S^T \cdot J$ based on PPIR(MPC) and PPIR(FHE)-v1.

registration problem can be naturally adapted for accounting to privacy-preserving operations, and illustrate the effectiveness of PPIR on a variety of registration tasks.

3. Analysis of classical IR loss functions under a privacy-preserving perspective

In this section we review typical loss functions used in image registration, and analyze the related requirements for privacy-preserving optimization.

3.1. Optimization of SSD loss

A typical loss function to be optimized during the registration process is the sum of squared intensity differences (SSD) evaluated on the set of image coordinates:

$$\text{SSD}(I, J, \theta) = \operatorname{argmin}_{\theta} \sum_{\mathbf{x}} [I(\mathbf{W}_{\theta}(\mathbf{x})) - J(\mathbf{x})]^2 \quad (3)$$

with Jacobian:

$$G = \sum_{\mathbf{x}} S(\mathbf{x}) \cdot (I(\mathbf{W}_{\theta}(\mathbf{x})) - J(\mathbf{x})), \quad (4)$$

where the quantity

$$S(\mathbf{x}) = \nabla I(\mathbf{x}) \frac{\partial \mathbf{W}_{\theta}(\mathbf{x})}{\partial \theta} \quad (5)$$

quantifies image and transformation gradients, and

$$H = \sum_{\mathbf{x}} \left(\nabla I(\mathbf{x}) \frac{\partial \mathbf{W}_{\theta}(\mathbf{x})}{\partial \theta} \right)^T \left(\nabla I(\mathbf{x}) \frac{\partial \mathbf{W}_{\theta}(\mathbf{x})}{\partial \theta} \right) \quad (6)$$

is the second order term obtained from Equation (3) through linearization (Pennec et al., 1999; Baker and Matthews, 2004).

The solution to this problem requires the joint availability of both images I and J , as well as of the gradients of I and of \mathbf{W}_{θ} . In a privacy-preserving setting, this information cannot be disclosed, and the computation of Equation (2) is therefore impossible. We note that to calculate the update of the registration $\Delta \theta$ of Equation (2), the only operation that requires the joint availability of information from both parties is the term $R = \sum_{\mathbf{x}} S(\mathbf{x}) \cdot J(\mathbf{x})$, which can be computed a matrix-vector multiplication of vectorized quantities $R = S^T \cdot J$.

3.2. Mutual Information

Mutual Information quantifies the joint information content between the intensity distributions of the two images. This is calculated from the joint probability distribution function (PDF):

$$MI(I, J, \theta) = \operatorname{argmin}_{\theta} - \sum_{r,t} p(r, t; \theta) \log \left(\frac{p(r, t; \theta)}{p(r; \theta)p(t)} \right) \quad (7)$$

where, given N_r and N_t the maximum intensity for respectively I and J , we define $r \in [0; N_r - 1] \subseteq \mathbb{N}$ and $t \in [0; N_t - 1] \subseteq \mathbb{N}$ as the range of discretized intensity values of I and J , respectively. A Parzen window (Parzen, 1961) is used to generate continuous estimates of the underlying intensity distributions, thereby reducing the effects of quantization from interpolation, and discretization from binning the data. Let $\psi_I^3 : \mathbb{R} \mapsto [0, 1]$ be a cubic spline Parzen window, and let $\psi_J^0 : \mathbb{R} \mapsto [0, 1]$ be a zero-order spline Parzen window. The smoothed joint histogram of I and J (Viola and Wells III, 1997; Mattes et al., 2003) is given by:

$$p(r, t; \theta) = \frac{1}{N_x} \sum_{\mathbf{x}} \psi_I^3 \left(r - \frac{I(\mathbf{W}_{\theta}(\mathbf{x})) - I(\mathbf{W}_{\theta}(\mathbf{x}))^\circ}{\Delta b_r} \right) \cdot \psi_J^0 \left(t - \frac{J(\mathbf{x}) - J^\circ}{\Delta b_t} \right)$$

In the above formula, the intensity values of I and J are normalized by their respective minimum (denoted by $I(\mathbf{W}_\theta)^\circ$ and J°), and by the bin size (respectively Δb_r and Δb_t), to fit into the specified number of bins (b_r or b_t) of the intensity distribution. The final value for $p(r, t; \theta)$ is computed by normalizing by N_x , the number of sampled voxels. Marginal probabilities are simply obtained by summing along one axis of the PDF, that is, $p(r) = \sum_t p(r, t; \theta)$ and $p(t) = \sum_r p(r, t; \theta)$. Let the matrices $A_I^3 \in \mathbb{R}^{N_x \times N_r}$ and $B_J^0 \in \mathbb{R}^{N_x \times N_t}$ be defined as:

$$A_I^3(\mathbf{x}, r; \theta) = \psi_I^3 \left(r - \frac{I(\mathbf{W}_\theta(\mathbf{x})) - I(\mathbf{W}_\theta)^\circ}{\Delta b_r} \right)$$

and

$$B_J^0(\mathbf{x}, t) = \psi_J^0 \left(t - \frac{J(\mathbf{x}) - J^\circ}{\Delta b_t} \right),$$

the discretized joint PDF can be rewritten in a matrix form via the multiplication:

$$P = \frac{1}{N_x} \cdot (A_I^3)^T \cdot B_J^0. \quad (8)$$

The first derivative of the joint PDF is calculated as follows (Dowson and Bowden, 2006):

$$\frac{\partial p(r, t; \theta)}{\partial \theta} = -\frac{1}{N_x} \cdot \sum_{\mathbf{x}} B_J^0(\mathbf{x}, t) \cdot \frac{\partial \psi_I^3(\epsilon)}{\partial \epsilon} \cdot \frac{\partial \epsilon}{\partial I(\mathbf{W}_\theta(\mathbf{x}))} \cdot \frac{\partial I(\mathbf{W}_\theta(\mathbf{x}))}{\partial \theta}$$

where $\epsilon = \epsilon(\mathbf{x}, r; \theta) = r - \frac{I(\mathbf{W}_\theta(\mathbf{x})) - I(\mathbf{W}_\theta)^\circ}{\Delta b_r}$ is the input of the cubic spline. We also introduce the tensor $C_I^3 \in \mathbb{R}^{N_x \times N_r \times |\theta|}$ defined as $C_I^3(\mathbf{x}, r; \theta) = \frac{\partial \psi_I^3(\epsilon)}{\partial \epsilon} \cdot \frac{\partial \epsilon}{\partial I(\mathbf{W}_\theta(\mathbf{x}))} \cdot \frac{\partial I(\mathbf{W}_\theta(\mathbf{x}))}{\partial \theta}$, to write the discretized first derivative as:

$$P' = -\frac{1}{N_x} \cdot (B_J^0)^T \cdot C_I^3, \quad (9)$$

where $P' \in \mathbb{R}^{N_x \times N_r \times |\theta|}$.

The Jacobian of the MI is obtained from the chain rule and takes the form:

$$G = P' \log \left(\frac{P}{P_I} \right),$$

while the linearized Hessian (Dowson and Bowden, 2007) can be written as:

$$H = P'^T P' \left(\frac{1}{P} - \frac{1}{P_I} \right),$$

where $P_I = \sum_t p(r, t; \theta)$ is a vector which defines the discretized marginal PDF of the moving image.

The derivatives can be easily calculated from the properties of B-splines since we have $\frac{\partial \psi_I^3}{\partial \epsilon} = \psi_I^2(\epsilon + \frac{1}{2}) - \psi_I^2(\epsilon - \frac{1}{2})$. In a privacy-preserving scenario, to calculate the update of the registration $\Delta \theta$ of Equation (7), two operations require the joint availability of information from both parties, which are the matrix P of Equation (8) and the matrix P' of Equation (9).

3.3. Cross Correlation with Advanced Normalization Tools

In the Advanced Normalization Tools (ANTs) introduced by (Avants *et al.*, 2008), the normalized cross-correlation (CC) loss was specified in the context of diffeomorphic image registration. Let ϕ define a diffeomorphism over the domain $\Omega = \mathbb{R}^d$, parameterized by a time-varying velocity field $\mathbf{v}(\mathbf{x}, t)$. In the ANTs setting, inverse consistency is obtained by optimizing the CC loss with respect to both forward and backward (inverse) transformations, here denoted by $\phi_1(\mathbf{x}, t)$ and $\phi_2(\mathbf{x}, t)$, and parameterized by velocity fields $\mathbf{v}_1(\mathbf{x}, t)$ and $\mathbf{v}_2(\mathbf{x}, t)$ respectively. In particular, both images I and J are simultaneously warped towards a ‘‘half-way’’ space, to obtain $I_1 = I(\phi_1(\mathbf{x}, 0.5))$ and $J_2 = J(\phi_2(\mathbf{x}, 0.5))$. The CC loss is thus defined as:

$$CC(\mathbf{x}) = \frac{\sum_{x_i} (\bar{I}(x_i), \bar{J}(x_i))^2}{\sum_{x_i} (\bar{I}(x_i))^2 \sum_{x_i} (\bar{J}(x_i))^2} = \frac{D^2}{EF}$$

where $\bar{I}(x_i) = (I_1(x_i) - \mu_{I_1}(\mathbf{x}))$ and $\bar{J}(x_i) = (J_2(x_i) - \mu_{J_2}(\mathbf{x}))$ quantify the images appearance at location x_i , with respect to the average intensity μ_{I_1} and μ_{J_2} measured in a local window of size M . Coherently with the LDDMM formulation, the variational optimization problem is defined as:

$$E_{CC} = \inf_{\phi_1} \inf_{\phi_2} \int_{t=0}^{0.5} \|\mathbf{v}_1(\mathbf{x}, t)\|_L^2 + \|\mathbf{v}_2(\mathbf{x}, t)\|_L^2 dt + \int_{\Omega} CC(\mathbf{x}) d\Omega.$$

where L is a linear operator prescribing a norm on the velocity fields acting as a regularizer. The equation for the derivative of the forward update is given by:

$$\begin{aligned} \nabla_{\phi_1(\mathbf{x}, 0.5)} CC(\mathbf{x}) &= 2L\mathbf{v}_1(\mathbf{x}, 0.5) + \frac{2D}{EF} \\ &\times \left(\bar{J}(\mathbf{x}) - \frac{D}{E} \bar{I}(\mathbf{x}) \right) |D\phi_1| \nabla \bar{I}(\mathbf{x}), \end{aligned} \quad (10)$$

while the derivative of the backward update is analogously given by:

$$\begin{aligned} \nabla_{\phi_2(\mathbf{x}, 0.5)} CC(\mathbf{x}) &= 2L\mathbf{v}_2(\mathbf{x}, 0.5) + \frac{2D}{EF} \\ &\times \left(\bar{I}(\mathbf{x}) - \frac{D}{F} \bar{J}(\mathbf{x}) \right) |D\phi_2| \nabla \bar{J}(\mathbf{x}) \end{aligned} \quad (11)$$

In privacy-preserving scenario, the sensitive terms carrying private image information are $\frac{2D}{EF}$, $(\bar{J}(\mathbf{x}) - \frac{D}{E} \bar{I}(\mathbf{x}))$ and $(\bar{I}(\mathbf{x}) - \frac{D}{F} \bar{J}(\mathbf{x}))$, which must therefore be computed in a privacy-preserving regime.

4. Building blocks for Secure Computation

After introducing in Section 1 the IR optimization problem and the related functionals addressed in this work, in this section, we review the standard privacy-preserving techniques that will be employed to develop PPIR.

4.1. Secure Multi-Party Computation

Introduced by (Yao, 1982), MPC is a cryptographic tool that allows multiple parties to jointly compute a common function over their private inputs (secrets) and discover no more than the output of this function. Among existing MPC protocols, additive secret sharing consists of first splitting every secret s into additive shares $\langle s \rangle_i$, such that $\sum_{i=1}^n \langle s \rangle_i = s$, where n is the number of collaborating parties. Each party i receives one share $\langle s \rangle_i$ and executes an arithmetic circuit in order to obtain the final output of the function. In this paper, we adopt the two-party computation protocol defined in SPDZ (Damgard et al., 2011), whereby the actual function is mapped into an arithmetic circuit and all computations are performed within a finite ring with modulus Q . Additions consist of locally adding shares of secrets, while multiplications require interaction between parties. Following (Damgard et al., 2011), SPDZ defines: MPC-MUL to compute element-wise multiplication, MPCDOT to compute matrix-vector multiplication and MPCMATMUL to compute matrix-matrix multiplication. These operations are performed within an honest but curious protocol.

4.2. Homomorphic Encryption

Initially introduced by Rivest et al. in (Rivest et al., 1978), Homomorphic Encryption (HE) enables the execution of operations over encrypted data without disclosing either the input data or the result of the computation. Hence, $party_1$ encrypts the input with its public key and sends the encrypted input to $party_2$. In turn, $party_2$ evaluates a circuit over this encrypted input and sends the result, which still remains encrypted, back to $party_1$ which can finally decrypt them. Among various HE schemes, CKKS (Cheon et al., 2017) supports the execution of all operations on encrypted real values and is considered a levelled homomorphic encryption (LHE) scheme. The supported operations are: SUM (+), Element-wise multiplication (*) and DOTPRODUCT. With CKKS, an input vector is mapped to a polynomial and further encrypted with a public key in order to obtain a pair of polynomials $c = (c_0, c_1)$. The original function is further mapped into a set of operations that are supported by CKKS, which are executed over c . The performance and security of CKKS depend on multiple parameters including the degree of the polynomial N , which is usually sufficiently large (e.g. $N = 4096$, or $N = 8192$).

5. Methods: from IR to PPIR

In this section, we describe the privacy preserving variants of the three IR methods described in Section 3. We propose two versions of PPIR for SSD according to the underlying cryptographic tool, namely PPIR(MPC), integrating MPC, and PPIR(FHE), integrating FHE. In the case of MI and NCC, we focus on the design of the MPC-variant, due to the non-negligible computational overhead of FHE in these applications. Finally, we also study rigid point cloud registration and describe its privacy-preserving variant in Appendix A 1.

5.1. PPIR based on SSD

As mentioned in Section 3.1, when optimizing the SSD cost, the only sensitive operation that must be jointly executed by the parties is the matrix-vector multiplication: $R = S^T \cdot J$, where S^T is only known to $party_1$ and J to $party_2$. Figure 1 illustrates how cryptographic tools are employed to ensure privacy during registration.

With MPC (Figure 1a), $party_1$ secretly shares the matrix S^T to obtain $(\langle S \rangle_1, \langle S \rangle_2)$, while $party_2$ secretly shares the image J to obtain $(\langle J \rangle_1, \langle J \rangle_2)$. Each party also receives its corresponding share, so that $party_1$ holds $(\langle S \rangle_1, \langle J \rangle_1)$ and $party_2$ holds $(\langle S \rangle_2, \langle J \rangle_2)$. The parties execute a circuit with MPCMUL operations to calculate the 2-party dot product between S^T and J . The parties further synchronize to allow $party_1$ to obtain the product and finally to calculate $\Delta\theta$ (Equations (4) and (6)).

When using FHE (Figure 1b), $party_2$ first uses a FHE key k to encrypt J and obtains $\llbracket J \rrbracket \leftarrow \text{ENC}(k, J)$. This encrypted image is sent to $party_1$, who computes the encrypted matrix-vector multiplication $\llbracket R \rrbracket$. In this framework, only vector J is encrypted, and therefore $party_1$ executes scalar multiplications and additions in the encrypted domain only (which are less costly than multiplications over two encrypted inputs). The encrypted result $\llbracket R \rrbracket$ is sent back to $party_2$, which can obtain the result by decryption: $R = \text{DEC}(k, \llbracket R \rrbracket)$. Finally, $party_1$ receives R in clear form and can therefore compute $\Delta\theta$.

Thanks to the privacy and security guarantees of these cryptographic tools, during the entire registration procedure, the content of the image data S and J is never disclosed to the opposite party.

5.2. PPIR based on MI

In the MPC variant to calculate the joint PDF P in a privacy-preserving manner, the quantity A_j^3 is only known to $party_1$, while the quantity B_j^0 to $party_2$ (Section 3.2). With reference to Supplementary Figure A1a, $party_1$ secretly shares matrix $(A_j^3)^T = (\langle A_j^3 \rangle_1, \langle A_j^3 \rangle_2)$, while $party_2$ does the same with $B_j^0 = (\langle B_j^0 \rangle_1, \langle B_j^0 \rangle_2)$. Each party also receives its corresponding share: $party_1$ now holds $(\langle A_j^3 \rangle_1, \langle B_j^0 \rangle_1)$, and $party_2$ holds $(\langle B_j^0 \rangle_2, \langle A_j^3 \rangle_2)$. The parties execute a circuit with MPCMATMUL operation to calculate the 2-party matrix multiplication between $(A_j^3)^T$ and B_j^0 . The next operation carried out in the privacy-preserving setting is the computation of the first derivative of Equation (9). Supplementary Figure A1b illustrates the MPC variant, where $party_1$ only knows C_j^3 and $party_2$ only knows B_j^0 . Initially, $party_1$ secretly shares the matrix $C_j^3 = (\langle C_j^3 \rangle_1, \langle C_j^3 \rangle_2)$, while $party_2$ does the same with $(B_j^0)^T = (\langle B_j^0 \rangle_1, \langle B_j^0 \rangle_2)$. Each party also receives its corresponding share, namely: $party_1$ holds $(\langle C_j^3 \rangle_1, \langle B_j^0 \rangle_1)$ and $party_2$ holds $(\langle C_j^3 \rangle_2, \langle B_j^0 \rangle_2)$. The parties also execute a circuit with the MPCMATMUL between $(B_j^0)^T$ and C_j^3 .

5.3. PPIR based on ANTS CC loss

According to Section 3.3, $party_1$ has access to \bar{I} and E , whereas $party_2$ has access to \bar{J} and F . The computation begins with the optimization of the CC term, specifically the quantity

Dataset	Dimension	Modality	Registration Type	Loss function	PETs
2D Point Cloud	193 points	Mono	Rigid	SSD	MPC and FHE-v1
2D Whole body PET	1260 × 1090 voxels	Mono	Affine/Cubic splines	SSD	MPC+URS/GMS, FHE-v1 and v2 + URS/GMS
2D Brain MRI	121 × 121 voxels	Mono	Cubic splines	SSD	MPC, FHE-v1 and v2
3D Brain MRI and PET	180 × 256 × 256 and 160 × 160 × 96 voxels	Multi	Affine	MI	MPC
3D Abdomen MR and CT	192 × 160 × 192 voxels	Multi	Diffeomorphic (ANTs)	CC	MPC

Table 1: Overview of the datasets used in the study. PETs: Privacy Enhancing Technologies.

$\frac{2D}{EF}$. In the initial phase, as illustrated in Supplementary Figure A2a, $party_1$ and $party_2$ secretly share $\bar{I} = (\langle \bar{I} \rangle_1, \langle \bar{I} \rangle_2)$, and $\bar{J} = (\langle \bar{J} \rangle_1, \langle \bar{J} \rangle_2)$, respectively. The subsequent multiplication of these shared values results in the computation of the shares of $D = (\langle D \rangle_1, \langle D \rangle_2)$, which is never reconstructed. In the next step of the protocol, $party_1$ secretly shares $\frac{1}{E}$, and $party_2$ secretly shares $\frac{1}{F}$. Through a multiplication of the shares of D , $\frac{1}{E}$, and $\frac{1}{F}$, both parties collectively obtain the final value of $\frac{2D}{EF}$.

The other two terms that need to be jointly computed are the third term of Equation (10) and (11), namely $(\bar{J} - \frac{D}{E}\bar{I})$ and $(\bar{I} - \frac{D}{F}\bar{J})$, reported in Supplementary Figure A2b. In the case of $(\bar{J} - \frac{D}{E}\bar{I})$, $party_1$ secretly shares \bar{I} and $\frac{1}{E}$, while $party_2$ secretly shares \bar{J} . Since both parties already have access to the share of D from the previous protocol, they proceed to multiply the secret shares of D , $\frac{1}{E}$, and \bar{I} . Subsequently, they subtract the result from \bar{J} to obtain the final value of $(\bar{J} - \frac{D}{E}\bar{I})$. The computation of $(\bar{J} - \frac{D}{E}\bar{I})$ is analogous to the one of $(\bar{I} - \frac{D}{F}\bar{J})$, where $party_2$ secretly shares $\frac{1}{E}$, and both parties participate to the multiplication of D , $\frac{1}{E}$, and \bar{J} . The resulting value is finally subtracted from \bar{I} , yielding the final result $(\bar{J} - \frac{D}{E}\bar{I})$.

5.4. Protocols enhancement for SSD loss

Effectively optimizing Equation (1) with MPC or FHE is particularly challenging, due to the computational bottleneck of these techniques when applied to large-dimensional objects (Haralampieva *et al.*, 2020; Benaissa *et al.*, 2021), notably affecting the computation time and the occupation of communication bandwidth between parties. Because cryptographic tools introduce a non-negligible overhead in terms of performance and scalability, in this section we introduce specific techniques to optimize the underlying image registration operations.

5.4.1. Gradient sampling

Since the registration gradient is generally driven mainly by a fraction of the image content, such as the image boundaries in the case of SSD cost, a reasonable approximation of Equations (4) and (6) can be obtained by evaluating the cost only on relevant image locations. This idea has been introduced in medical image registration (Viola and Wells III, 1997; Mattes *et al.*, 2003; Sabuncu and Ramadge, 2004), and here is adopted to optimize Equation (3) by reducing the dimensionality of the arrays on which encryption is performed. We test two different techniques: (i): Uniformly Random Selection (URS), proposed by (Viola and Wells III, 1997; Mattes *et al.*, 2003), in which a random subset of dimension $l \leq d$ of spatial coordinates is sampled at every iteration with uniform probabilities, $p(\mathbf{x}) = \frac{1}{d}$; and (ii): Gradient Magnitude Sampling (GMS) (Sabuncu and Ramadge, 2004), which consists of sampling a subset of coordinates with probability proportional to the norm of the image

gradient, $p(\mathbf{x}) \sim \|\nabla I(\mathbf{x})\|$. We note that gradient sampling is not necessary for computing the MI since in Equation (8) the computation is already performed on a subsample of the image voxels.

5.4.2. Matrix partitioning in FHE

We now describe additional improvements dedicated to PPIR(FHE) and propose two versions of this solution: (i) PPIR(FHE)-v1 implements an optimization of the matrix-vector multiplication by partitioning the vector image into a vector of submatrices, whereas (ii) PPIR(FHE)-v2 enhances the workload of the parties by fairly distributing the computation among them.

PPIR(FHE)-v1. We introduce here a novel optimization dedicated to PPIR with FHE, in particular when the CKKS algorithm is adopted. CKKS allows multiple inputs to be packed into a single ciphertext to decrease the number of homomorphic operations. To optimize matrix-vector multiplication, we propose to partition the image vector J into k sub-arrays of dimension l , and the matrix S^T into k sub-matrices of dimension $|\theta| \times l$. Once all sub-arrays J_i are encrypted, we propose to iteratively apply `DotProduct` as proposed by Benaissa *et al.* (2021), between each sub-matrix and corresponding sub-array; these intermediate results are then summed to obtain the final result, namely: $\llbracket R \rrbracket = \sum_{i=0}^K \text{DotProduct}(\llbracket J_i^T \rrbracket, S_i) = S^T \cdot \llbracket J \rrbracket$.

PPIR(FHE)-v2. In addition to packing multiple inputs into a single ciphertext, the application of FHE to PPIR can be optimized by more equally distributing the workload among the two parties. We note that in PPIR(FHE)-v1, $party_1$ is in charge of computing the matrix-vector multiplication entirely while $party_2$ only encrypts the input and decrypts the result. Following Supplementary Figure A4, PPIR(FHE)-v2 starts by splitting the matrix S of $party_1$ into two sub-matrices S_1 and S_2 using the operation $split_{h,K}$. This operation partitions the matrix into K equally-sized sub-matrices. Next, the operation $flatten$ is applied to S_1 and S_2 , obtaining vectors S'_1 and S'_2 respectively. Then, $party_1$ encrypts S'_2 and sends it to $party_2$, which subsequently applies to its vector J the operation $split_{v,K}$, obtaining J_1 and J_2 . This operation splits J into K equally-sized partitions, and it executes the operation $replicate_d$ to J_1 and J_2 , obtaining J'_1 and J'_2 respectively. $party_2$ encrypts J'_1 and sends it to $party_1$. Both parties then iteratively perform the element-wise multiplication $*$ and sum up the results of the different partitions using the primitive $Sum_{i,k}$. The protocol doesn't rely anymore on `DotProduct` and leads to a significant gain in computational load.

6. Experiments & Results

We demonstrate and assess the different versions of PPIR illustrated in Section 3 on a variety of image registration problem, namely: (i) SSD for rigid transformation of point cloud data, (ii) SSD with linear and non-linear alignment of whole body positron emission tomography (PET) data; (iii) SSD and MI for mono- and multimodal linear alignment of MRI and PET brain scans; (iv) diffeomorphic non-linear registration with CC of multimodal abdomen data from CT and MRI scans. Experiments are carried out on 2D (mainly for the SSD case) and 3D imaging data. In Table 1 is reported an overview of the datasets used specifying their dimensions, modality, registration type, loss functions, and the Privacy Enhancing Technologies (PETs) employed.

6.1. Experimental data

Point Cloud Data. We showcase rigid registration on 2D point cloud data representing the corpus callosum, as presented in Vachet *et al.* (2012), with a set size $n = 193$. The registration loss here considered is SSD between point coordinates (additional details are provided in Appendix A 1).

Whole body PET data. The dataset considered for linear and non-linear registration with SSD consists of 18-Fluoro-Deoxy-Glucose (^{18}FDG) whole body PET scans. The images are a frontal view of the maximum intensity projection reconstruction, obtained by 2D projection of the voxels with the highest intensity across views (1260×1090 pixels).

Brain MRI and PET data. This dataset regroups brain MRI and PET images obtained from the Alzheimer’s Disease Neuroimaging Initiative (Mueller *et al.*, 2005). MRI data were processed via a standard processing pipeline to estimate gray matter density maps (Ashburner and Friston, 2000). Non-linear registration was carried out on the extracted mid-coronal slice, of dimension 121×121 pixels. For 3D multimodal linear registration with MI, we use both MRI images and PET images, with respective dimension of $180 \times 256 \times 256$ and $160 \times 160 \times 96$ voxels.

Abdomen MR and CT data. The multimodal dataset Abdomen-MR-CT (Hering *et al.*, 2022) was used for experiments with ANTs registration based on CC. The data was compiled from public studies of the cancer imaging archive (TCIA) Clark *et al.* (2013) that contains 8 paired scans of MRI and CT from the same patients. The data have an isotropic resolution of $2mm$ and a voxel dimension of $192 \times 160 \times 192$. They also provide 3D segmentation masks for the liver, spleen, and left and right kidney. All scans were pre-aligned by groupwise affine registration.

6.2. Experimental Details

In order to avoid local minima and to decrease computation time, we use a hierarchical multiresolution optimization scheme. The scheme involves M resolution steps, denoted as $r_1 \dots r_M$. At each resolution step r_m , the input data is down-sampled by a scaling factor m , where $m \in [1 \dots M]$. The quality of PPIR is assessed by comparing the registration results with those obtained with standard registration on clear images

(CLEAR). The metrics considered are the difference in image intensity at optimum (for SSD), the total number of iterations required to converge, and the displacement root mean square difference (RMSE) between CLEAR and PPIR. We also evaluate the performance of PPIR in terms of average computation (running time) and communication (bandwidth) across iterations. In the multimodal Abdomen data, the quality of the registration result was assessed by the overlap across the labeled anatomical regions, quantified by the DICE score.

Point Cloud Data. The registration protocol here adopted is detailed in Appendix A 1. For MPC we set as the prime modulus $Q = 2^{32}$. For PPIR(FHE), we define the polynomial degree modulus as $N = 4096$.

Whole body PET data. Whole-body PET image alignment was first performed by optimizing the transformation W_θ in Equation (1) with respect to affine registration parameters. The multiresolution steps used are r_1, r_5, r_{10}, r_{20} . A second whole-body PET image alignment experiment was performed by non-linear registration, without gradient approximation based on a cubic spline model (one control point every four pixels along both dimensions), with multiresolution steps $r_1, r_2, r_5, r_{10}, r_{20}$ and r_{30} . Concerning the PPIR framework, transformations were optimized for both MPC and FHE by using gradient approximation (Section 5) using the same sampling seed for each test. For MPC we set as the prime modulus $Q = 2^{32}$. For PPIR(FHE), we define the polynomial degree modulus as $N = 4096$, and set the resizing parameter D to optimize the trade-off between run-time and bandwidth. Since D needs to be a divisor of the image size image data we set $D = 128$.

Brain MRI and PET data. The registration of brain gray matter density images was performed by non-linear registration based on SSD, without gradient approximation, based on a cubic spline model (one control point every five pixels along both dimensions), with multiresolution steps r_1 and r_2 . For PPIR(MPC) we use the same configuration defined in the previous section, while for PPIR(FHE) we use the same N and we set $D = 121$.

We tested PPIR with MI for multi-modal 3D affine image registration between PET and MRI brain scans where, in addition to varying the multi-resolution steps, a Gaussian blurring filter is applied to the images with a kernel that narrows as multi-resolution proceeds. The kernel size at different resolutions, denoted with $\sigma_1 \dots \sigma_M$, is used to control the amount of blurring applied to the image at each step of the multi-resolution process. The multiresolution steps applied are r_5 and r_{10} , with 10% of the image’s pixels utilized as the number of subsample pixels (N_x). Gaussian image blurring is applied with a degree of $\sigma_5 = 1$ and $\sigma_{10} = 3$. For MPC we set as the prime modulus $Q = 2^{64}$.

Abdomen MRI and CT data. We tested PPIR with CC for multi-modal 3D ANTs image registration between MRI and CT abdomen scans. The multiresolution steps applied are r_3, r_2 and r_1 and the CC window size $M = 5 \times 5 \times 5$.

Implementation Details. The PPIR framework for the SSD is implemented using two state-of-the-art libraries: PySyft (Ryffel *et al.*, 2018), which provides SPDZ two-party computation,

SSD Affine registration and efficiency metrics							
Solution	Intensity Error (SSD)	Num. Iteration	Displacement RMSE (mm)	Time <i>party</i> ₁ (s)	Time <i>party</i> ₂ (s)	Comm. <i>party</i> ₁ (MB)	Comm. <i>party</i> ₂ (MB)
CLEAR	4.34 ± 0.0	118 ± 0.0	-	0.0	0.0	-	-
PPIR(MPC)	4.34 ± 0.0	114.8 ± 4.0	0.13 ± 0.04	0.13	0.13	1.54	1.54
CLEAR + URS	4.38 ± 0.0	61 ± 0.0	-	0.0	0.0	-	-
PPIR(MPC) + URS	4.34 ± 0.0	60.4 ± 6.85	1.75 ± 0.19	0.02	0.02	0.20	0.20
PPIR(FHE)-v1 ($D = 128$) + URS	4.34 ± 0.10	61.80 ± 4.82	13.47 ± 2.87	2.55	0.02	0.06	0.01
PPIR(FHE)-v2 ($D = 128$) + URS	4.34 ± 0.10	61.60 ± 7.21	13.93 ± 4.28	0.01	0.02	0.72	0.88
CLEAR + GMS	4.34 ± 0.0	63 ± 0.0	-	0.0	0.0	-	-
PPIR(MPC) + GMS	4.34 ± 0.0	59.80 ± 6.20	0.93 ± 0.42	0.02	0.02	0.20	0.20
PPIR(FHE)-v1 ($D = 128$) + GMS	4.34 ± 0.05	60.40 ± 5.12	0.59 ± 0.35	2.51	0.02	0.06	0.01
PPIR(FHE)-v2 ($D = 128$) + GMS	4.34 ± 0.05	57.03 ± 4.07	0.50 ± 0.36	0.02	0.02	0.73	0.93

Table 2: Affine SSD registration test, comparison between Clear, PPIR(MPC), PPIR(FHE)-v1 and PPIR(FHE)-v2. Registration metrics are reported as mean and standard deviation. Efficiency metrics in terms of average across iterations. RMSE: root mean square error.

SSD Cubic splines registration metrics							
Solution	Intensity Error (SSD)	Num. Iteration	Displacement RMSE (mm)	Time <i>party</i> ₁ (s)	Time <i>party</i> ₂ (s)	Comm. <i>party</i> ₁ (MB)	Comm. <i>party</i> ₂ (MB)
CLEAR	0.65 ± 0.0	413 ± 0.0	-	0.0	0.0	-	-
PPIR-MPC	0.65 ± 0.0	345.70 ± 91.22	7.31 ± 1.86	0.63	0.63	21.47	28.98
PPIR(FHE)-v1 ($D = 121$)	0.64 ± 0.0	224.7 ± 79.15	9.50 ± 4.34	3.41	0.00	0.06	0.01
PPIR(FHE)-v2 ($D = 256$)	0.64 ± 0.0	379.2 ± 75.82	11.02 ± 4.93	0.98	0.43	40.45	3.56
Clear + URS	0.02 ± 0.0	101 ± 0.0	-	0.0	0.0	-	-
PPIR(MPC) + URS	0.02 ± 0.00	79.3 ± 1.88	5.59 ± 0.39	0.41	0.41	8.00	8.00
PPIR(FHE)-v1 ($D = 128$) + URS	0.02 ± 0.00	105.40 ± 1.71	7.63 ± 0.01	12.23	0.0	0.06	0.01
PPIR(FHE)-v2 ($D = 128$) + URS	0.02 ± 0.00	105.20 ± 2.54	8.74 ± 1.90	0.62	0.26	24.74	3.37
Clear + GMS	0.02 ± 0.0	103.00 ± 0.0	-	0.0	0.0	-	-
PPIR(MPC) + GMS	0.02 ± 0.04	80.20 ± 1.62	6.17 ± 0.37	0.41	0.41	8.00	8.00
PPIR(FHE)-v1 ($D = 128$) + GMS	0.02 ± 0.00	105.70 ± 2.40	5.60 ± 2.22	11.95	0.0	0.06	0.01
PPIR(FHE)-v2 ($D = 128$) + GMS	0.02 ± 0.00	106.32 ± 1.30	9.11 ± 2.34	0.62	0.26	24.91	3.35

Table 3: Non-Linear SSD registration test comparison between Clear, PPIR(MPC), PPIR(FHE)-v1 and PPIR(FHE)-v2. The registration metrics are reported as mean and standard deviation. Efficiency metrics in terms of average across iterations. RMSE: root mean square error.

and TenSeal (Benaissa *et al.*, 2021), which implements the CKKS protocol³.

PPIR based on MI and CC is implemented by extending the Dipy framework of Garyfallidis *et al.* (2014)⁴. Finally, PPIR for point cloud data is released in a separated repository⁵.

All the experiments are executed on a machine with an Intel(R) Core(TM) i7-7800X CPU @ (3.50GHz x 12) using 132GB of RAM. For each registration configuration, the optimization is repeated 10 times to account for the random generation of MPC shares and FHE encryption keys.

6.3. Results

Point Cloud Data. In Supplementary Table A1 we present the registration metrics for PPIR(MPC) and PPIR(FHE)-v1. The registration shows that PPIR(MPC) achieves the best results compared to PPIR(FHE), which exhibits not only a longer computation time but also requires higher bandwidth, thanks to its non-iterative algorithm. However, to carry out MATMUL, a sufficiently large N (4096) is required, and in this scenario, it leads to a significant loss of chipertext slots compared to the dimension of the point set $n = 193$. Finally, the qualitative results reported in Figure A7 show negligible differences between point cloud transformed with CLEAR, PPIR(MPC) and PPIR(FHE)-v1.

Whole body PET data: affine registration (SSD). Table 2 compares CLEAR, PPIR(MPC), PPIR(FHE)-v1 and v2, showcasing metrics resulting from the affine transformation

of whole-body PET images. Notably, registration through PPIR(MPC) yields negligible differences compared to CLEAR in terms of the number of iterations, intensity, and displacement. In contrast, registering with PPIR(FHE) is not feasible when considering entire images due to computational complexity. Nevertheless, Supplementary Figure A5 shows that neither MPC nor FHE decreases the overall quality of the affine registered images. A comprehensive assessment of the registration results is available in the Appendix. Table 2 (Efficiency metrics) shows that PPIR(MPC) performed on full images requires higher computation time and required communication bandwidth compared to PPIR(MPC)+URS/GMS. These numbers improve sensibly when using URS or GMS (by factors 10× and 20× for time and bandwidth, respectively). Concerning PPIR(FHE)-v1, we note the uneven computation time and bandwidth usage between clients, due to the asymmetry of the encryption operations and communication protocol (Figure 1). PPIR(FHE)-v2, which shares the computational workload between the two parties and avoids DOTPRODUCT, allows obtaining an important speed-up over PPIR(FHE)-v1 (100× faster). Notably, this gain is obtained without affecting the quality of the registration metrics, and improves the execution time of PPIR(MPC). On the other side, although PPIR(FHE)-v1 is able to improve communication with respect to PPIR(MPC), it still suffers from the highest communication among the three proposed solutions. This is due to the fact PPIR(FHE) protocols can find different applications depending on the requirements in terms of computational power or bandwidth.

Brain MRI data and whole body PET data: non-linear registration (SSD). Table 3, comparing CLEAR and PPIR(MPC), PPIR(FHE)-v1 and v2, showcases the metrics resulting from spline-based non-linear registration between grey

³https://github.com/rtaiello/pp_image_registration

⁴https://github.com/rtaiello/pp_dipy/tree/main

⁵https://github.com/rtaiello/ppir_pc

MI Affine registration metrics							
Solution	Intensity Error (MI)	Num. Iteration	Displacement RMSE (<i>mm</i>)	Time <i>party</i> ₁ (s)	Time <i>party</i> ₂ (s)	Comm. <i>party</i> ₁ (MB)	Comm. <i>party</i> ₂ (MB)
Clear	0.22 ± 0.00	213 ± 0.0	-	0.0	0.0	-	-
PPIR(MPC)	0.22 ± 0.00	264 ± 0.0	0.41 ± 0.04	1.02	1.02	15.00	15.00

Table 4: Affine MI registration test, comparison between Clear and PPIR(MPC). Registration metrics are reported as mean and standard deviation. Efficiency metrics in terms of average across iterations. RMSE: root mean square error.

CC ANTs registration registration metrics								
Solution	Initial DICE score	Final DICE score	Num. Iteration	Displacement RMSE (<i>mm</i>)	Time <i>party</i> ₁ (s)	Time <i>party</i> ₂ (s)	Comm. <i>party</i> ₁ (MB)	Comm. <i>party</i> ₂ (MB)
Forward								
Clear	0.54 ± 0.13	0.68 ± 0.19	26 ± 0.0	-	-	-	-	-
PPIR(MPC)		0.67 ± 0.19	26 ± 0.0	0.22 ± 0.04	24.00	24.00	152.25	152.25
Backward								
Clear	0.54 ± 0.13	0.69 ± 0.19	26 ± 0.0	-	-	-	-	-
PPIR(MPC)		0.68 ± 0.19	26 ± 0.0	0.22 ± 0.04	24.00	24.00	152.25	152.25

Table 5: ANTs registration with CC, comparison between Clear and PPIR(MPC). Registration metrics are reported as mean and standard deviation. Efficiency metrics in terms of average across iterations. RMSE: root mean square error.

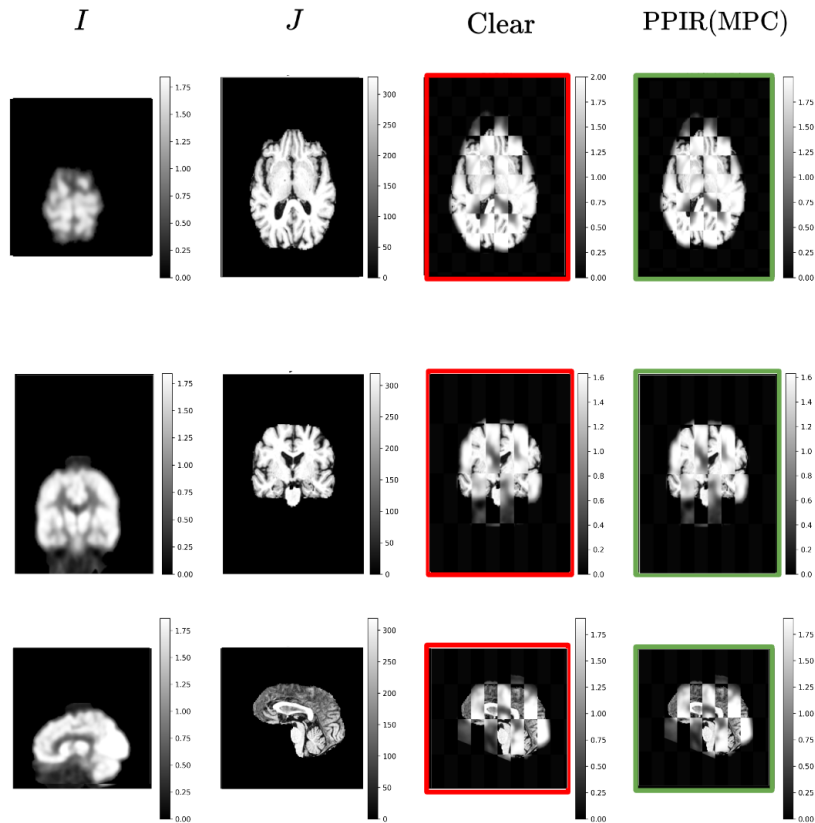


Fig. 2: Qualitative results for affine registration with MI over 3D medical images using ADNI dataset (Mueller *et al.*, 2005). The images are presented in a 3×4 grid, with the first row representing the axial axis, the second row the coronal axis, and the third row the sagittal axis. In the first column of each row, the moving image obtained using PET modality is shown, while in the second column, the fixed image obtained using MRI modality is displayed. The third column shows the checkerboard alignment result using CLEAR, while the fourth column shows the result using PPIR(MPC). The different protocols are highlighted by red and green frames, respectively.

matter density images without the application of gradient approximation. Additionally, the table includes results for the registration between whole-body PET images when the gradient approximation is applied.

Brain MRI data without gradient approximation. Regarding the registration accuracy, we draw conclusions similar to those of the affine case, where PPIR(MPC) leads to minimum differences with respect to CLEAR, while PPIR(FHE)-v1 seems slightly superior. PPIR(MPC) is associated with a lower execution time and a higher computational bandwidth, due to the larger number of parameters of the cubic splines, which affects the size of the matrix S . Although PPIR(FHE)-v1 has a slower execution time, the demanded bandwidth is inferior to the one of PPIR(MPC), since the encrypted image is transmitted only once. PPIR(FHE)-v2, as in the affine case, outperforms PPIR(FHE)-v1 (still about 100× faster) leading to comparable values between registration metrics and is still inferior to PPIR(MPC). Here, the limitations of PPIR(FHE)-v1 on the bandwidth size are even more evident than in the affine case, since the bandwidth increases according to the number of parameters. This result gives a non-negligible burden to the $party_1$, due to the multiple sending of the flattened and encrypted submatrices of updated parameters. Furthermore, in this case, PPIR(FHE)-v1 performs slightly worse than PPIR(MPC) in terms of execution time.

Whole Body PET Data With Gradient Approximation. Incorporating gradient approximation for handling whole-body PET data leads to similar conclusions as for the experiments on brain data. Qualitative results, reported in Supplementary Figure A6, show negligible differences between images transformed with CLEAR+GMS, PPIR(MPC)+GMS, and PPIR(FHE)-v1+GMS.

Brain MRI and PET data: affine registration (MI). Table 4 provides information on the computation cost of the protocol and the registration metrics for both the joint PDF and the First Derivative of the Joint PDF, using both CLEAR and PPIR(MPC). The results demonstrate a reasonable execution time (completed in under 5 minutes for the entire process) and noteworthy data transfer, totaling less than 4GB. Qualitative results for the image registration are shown in Figure 2, indicating that there is no difference between the moving transformed using CLEAR and PPIR(MPC).

Abdomen MRI and CT data: diffeomorphic non-linear registration (CC). Table 5 presents the metrics for ANTs registration using both CLEAR and PPIR(MPC) methods. The initial DICE scores for both forward and backward transformations are consistent between CLEAR and PPIR(MPC), with slight variations in the final DICE scores. The number of iterations and displacement RMSE values also exhibit similar trends. In terms of communication and computation times, PPIR(MPC) demonstrates comparable performance with CLEAR, showcasing its feasibility for secure registration. The computation times and communication bandwidth between the two parties are well within reasonable limits. These qualitative result in Figure 3 indicates that the proposed PPIR(MPC) solution maintains the quality of registration metrics while ensuring secure and private communication between parties.

7. Discussion

This current work builds upon Taiello *et al.* (2022) by extending its application to include MI with linear registration (3D images), CC using the ANTs framework (3D images), enhancing the FHE for the SSD loss function, namely PPIR(FHE)-v2, and also integrating rigid point cloud registration. We recognize specific limitations associated with the use of FHE in our PPIR framework, which limited the effective use of this techniques besides the optimization of the SSD loss. FHE’s computational cost during homomorphic operations poses challenges, limiting the scalability and real-time applicability of PPIR(FHE). This is particularly evident when dealing with large datasets, such as 3D images, or when employing advanced image registration cost functions that demand significant computational resources. To address this gap, researchers are actively exploring the optimization of FHE through the integration of hardware accelerators (Boemer *et al.*, 2021).

For the SSD loss function, we provide comparison experiments with both URS and GMS (Viola and Wells III, 1997; Mattes *et al.*, 2003; Sabuncu and Ramadge, 2004) for sake of completeness and compatibility with subsampling approaches in IR. We recognized that URS doesn’t bring substantial improvements with respect to GRS, and this latter method should be preferred in the considered application or testing scenario.

While PPIR focuses on the privacy-preserving formulation of classical image registration methods based on gradient-based optimization, throughout the past years the research community has been steering the attention towards deep learning (DL)-based image registration (Simonovsky *et al.*, 2016; Krebs *et al.*, 2017; Yang *et al.*, 2017; Balakrishnan *et al.*, 2019). Among the medical imaging application of privacy-preserving methodologies, Kaissis *et al.* (2021) discussed privacy-preserving FL with Secure Aggregation (Bonawitz *et al.*, 2017) and Differential Privacy (Abadi *et al.*, 2016) for 2D medical image classification tasks. However, as highlighted by (Kaissis *et al.*, 2021), deploying DL models for privacy-preserving inference nowadays is predominantly achievable through Multi-Party Computation (MPC). This process necessitates multiple servers and incurs significant overhead, primarily attributed to the size of the DL model, especially when handling 3D image registration tasks within a DL-based framework (Balakrishnan *et al.*, 2019). To the best of our knowledge both DL and non-DL registration methods available in the literature do not satisfy the PPIR requirements investigated in our work, as they always require the disclosure of the target and moving images in clear.

8. Conclusion and future works

This study introduces the novel paradigm of Privacy Preserving Image Registration, designed for allowing image registration in privacy-preserving scenarios where images are confidential and cannot be shared in clear. Leveraging both secure multi-party computation (MPC) and Fully Homomorphic Encryption (FHE), we propose in PPIR effective strategies integrating cryptographic techniques into a variety of state-of-the-art registration frameworks, encompassing different parameterization and loss functions. We evaluate the framework’s per-

formance across various registration benchmarks, conducting quantitative and qualitative assessment for all the considered image registration problems. Our future direction involve extending PPIR to encompass additional cost functions commonly used in image registration, aiming to enhance the framework's versatility and applicability.

Acknowledgments

This work has been supported by the French government, through the 3IA Côte d'Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002, and by the ANR JCJC project Fed-BioMed 19-CE45-0006-01.

References

- Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., Zhang, L., 2016. Deep learning with differential privacy, in: Proceedings of the 2016 ACM SIGSAC conference on computer and communications security, pp. 308–318.
- Arun, K.S., Huang, T.S., Blostein, S.D., 1987. Least-squares fitting of two 3-d point sets. *IEEE Transactions on pattern analysis and machine intelligence*, 698–700.
- Ashburner, J., 2007. A fast diffeomorphic image registration algorithm. *Neuroimage* 38, 95–113.
- Ashburner, J., Friston, K.J., 2000. Voxel-based morphometry—the methods. *Neuroimage* 11, 805–821.
- Ashburner, J., Ridgway, G.R., 2013. Symmetric diffeomorphic modeling of longitudinal structural MRI. *Frontiers in neuroscience* 6, 197.
- Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C., 2008. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis* 12, 26–41.
- Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C., 2011. A reproducible evaluation of ants similarity metric performance in brain image registration. *Neuroimage* 54, 2033–2044.
- Baker, S., Matthews, I., 2004. Lucas-Kanade 20 years on: A unifying framework. *International journal of computer vision* 56, 221–255.
- Balakrishnan, G., Zhao, A., Sabuncu, M.R., Gutttag, J., Dalca, A.V., 2019. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* 38, 1788–1800.
- Beg, M.F., Miller, M.I., Trounev, A., Younes, L., 2005. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision* 61, 139–157.
- Benaissa, A., Retiat, B., Cebere, B., Belfedhal, A.E., 2021. Tenseal: A library for encrypted tensor operations using homomorphic encryption. *arXiv preprint arXiv:2104.03152*.
- Boemer, F., Kim, S., Seifu, G., DM de Souza, F., Gopal, V., 2021. Intel hexl: Accelerating homomorphic encryption with intel avx512-ifma52, in: Proceedings of the 9th on Workshop on Encrypted Computing & Applied Homomorphic Cryptography, pp. 57–62.
- Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H.B., Patel, S., Ramage, D., Segal, A., Seth, K., 2017. Practical secure aggregation for privacy-preserving machine learning, in: proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 1175–1191.
- Cardoso, M.J., Leung, K., Modat, M., Keihaninejad, S., Cash, D., Barnes, J., Fox, N.C., Ourselin, S., Initiative, A.D.N., et al., 2013. STEPs: Similarity and truth estimation for propagated segmentations and its application to hippocampal segmentation and brain parcellation. *Medical image analysis* 17, 671–684.
- Cheon, J.H., Kim, A., Kim, M., Song, Y., 2017. Homomorphic encryption for arithmetic of approximate numbers, in: International Conference on the Theory and Application of Cryptology and Information Security, Springer, pp. 409–437.
- Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P., Moore, S., Phillips, S., Maffitt, D., Pringle, M., et al., 2013. The cancer imaging archive (tcia): maintaining and operating a public information repository. *Journal of digital imaging* 26, 1045–1057.
- Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis: I. segmentation and surface reconstruction. *Neuroimage* 9, 179–194.
- Damgard, I., Pastro, V., Smart, N., Zakarias, S., 2011. Multiparty computation from somewhat homomorphic encryption. *Cryptology, ePrint Archive, Report 2011/535*. <https://ia.cr/2011/535>.
- Dowson, N., Bowden, R., 2006. A unifying framework for mutual information methods for use in non-linear optimisation, in: European Conference on Computer Vision, Springer, pp. 365–378.
- Dowson, N., Bowden, R., 2007. Mutual information for lucas-kanade tracking (milk): An inverse compositional formulation. *IEEE transactions on pattern analysis and machine intelligence* 30, 180–185.
- Garyfallidis, E., Brett, M., Amirbekian, B., Rokem, A., Van Der Walt, S., Descoteaux, M., Nimmo-Smith, I., Contributors, D., 2014. Dipy, a library for the analysis of diffusion mri data. *Frontiers in neuroinformatics* 8, 8.
- Gazula, H., Holla, B., Zhang, Z., Xu, J., Verner, E., Kelly, R., Jain, S., Bharath, R.D., Barker, G.J., Basu, D., et al., 2021. Decentralized multisite vbm analysis during adolescence shows structural changes linked to age, body mass index, and smoking: a coinstac analysis. *Neuroinformatics* 19, 553–566.
- Haralampieva, V., Rueckert, D., Passerat-Palmbach, J., 2020. A systematic comparison of encrypted machine learning solutions for image classification, in: Proceedings of the 2020 workshop on privacy-preserving machine learning in practice, pp. 55–59.
- Heinrich, M.P., Jenkinson, M., Bhushan, M., Matin, T., Gleeson, F.V., Brady, J.M., Schnabel, J.A., 2011. Non-local shape descriptor: A new similarity metric for deformable multi-modal registration, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, pp. 541–548.
- Hering, A., Hansen, L., Mok, T.C., Chung, A.C., Siebert, H., Häger, S., Lange, A., Kuckertz, S., Heldmann, S., Shao, W., et al., 2022. Learn2reg: comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning. *IEEE Transactions on Medical Imaging* 42, 697–712.
- Kaissis, G., Ziller, A., Passerat-Palmbach, J., Ryffel, T., Usynin, D., Trask, A., Lima Jr, I., Mancuso, J., Jungmann, F., Steinborn, M.M., et al., 2021. End-to-end privacy preserving deep learning on multi-institutional medical imaging. *Nature Machine Intelligence* 3, 473–484.
- Krebs, J., Mansi, T., Delingette, H., Zhang, L., Ghesu, F.C., Miao, S., Maier, A.K., Ayache, N., Liao, R., Kamen, A., 2017. Robust non-rigid registration through agent-based action learning, in: Medical Image Computing and Computer Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part I 20, Springer, pp. 344–352.
- Lauter, K., 2021. Private AI: Machine Learning on Encrypted Data. Technical Report. Eprint report <https://eprint.iacr.org/2021/324.pdf>.
- Lotan, E., Tschider, C., Sodickson, D.K., Caplan, A.L., Bruno, M., Zhang, B., Lui, Y.W., 2020. Medical imaging and privacy in the era of artificial intelligence: myth, fallacy, and the future. *Journal of the American College of Radiology* 17, 1159–1162.
- Mattes, D., Haynor, D.R., Vesselle, H., Lewellen, T.K., Eubank, W., 2003. Pet-ct image registration in the chest using free-form deformations. *IEEE transactions on medical imaging* 22, 120–128.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A., 2017. Communication-efficient learning of deep networks from decentralized data, in: Artificial intelligence and statistics, PMLR, pp. 1273–1282.
- Modersitzki, J., 2009. FAIR: Flexible Algorithms for Image Registration. SIAM, Philadelphia.
- Mueller, S.G., Weiner, M.W., Thal, L.J., Petersen, R.C., Jack, C., Jagust, W., Trojanowski, J.Q., Toga, A.W., Beckett, L., 2005. The alzheimer's disease neuroimaging initiative. *Neuroimaging Clinics* 15, 869–877.
- Parzen, E., 1961. Mathematical considerations in the estimation of spectra. *Technometrics* 3, 167–190.
- Pennec, X., Cachier, P., Ayache, N., 1999. Understanding the “Demon's algorithm”: 3D non-rigid registration by gradient descent, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, pp. 597–605.
- Reuter, M., Rosas, H.D., Fischl, B., 2010. Highly accurate inverse consistent registration: a robust approach. *Neuroimage* 53, 1181–1196.
- Rivest, R.L., Adleman, L., Dertouzos, M.L., et al., 1978. On data banks and

privacy homomorphisms. *Foundations of secure computation* 4, 169–180.

Ryffel, T., Trask, A., Dahl, M., Wagner, B., Mancuso, J., Rueckert, D., Passerat-Palmbach, J., 2018. A generic framework for privacy preserving deep learning. arXiv preprint arXiv:1811.04017.

Sabuncu, M.R., Ramadge, P.J., 2004. Gradient based nonuniform subsampling for information-theoretic alignment methods, in: *The 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, IEEE. pp. 1683–1686.

Schnabel, J.A., Heinrich, M.P., Papież, B.W., Brady, J.M., 2016. Advances and challenges in deformable image registration: from image fusion to complex motion modelling. *Medical Image Analysis* 33, 145–148.

Shattuck, D.W., Prasad, G., Mirza, M., Narr, K.L., Toga, A.W., 2009. Online resource for validation of brain segmentation methods. *NeuroImage* 45, 431–439.

Simonovsky, M., Gutiérrez-Becker, B., Mateus, D., Navab, N., Komodakis, N., 2016. A deep metric for multimodal registration, in: *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2016: 19th International Conference, Athens, Greece, October 17-21, 2016, Proceedings, Part III* 19, Springer. pp. 10–18.

Taiello, R., Önen, M., Humbert, O., Lorenzi, M., 2022. Privacy preserving image registration, in: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (Eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, Springer Nature Switzerland, Cham. pp. 130–140.

Vachet, C., Yvernault, B., Bhatt, K., Smith, R.G., Gerig, G., Hazlett, H.C., Styner, M., 2012. Automatic corpus callosum segmentation using a deformable active fourier contour model, in: *Medical Imaging 2012: Biomedical Applications in Molecular, Structural, and Functional Imaging*, SPIE. pp. 79–85.

Viola, P., Wells III, W.M., 1997. Alignment by maximization of mutual information. *International journal of computer vision* 24, 137–154.

Yang, X., Kwitt, R., Styner, M., Niethammer, M., 2017. Quicksilver: Fast predictive image registration—a deep learning approach. *NeuroImage* 158, 378–396.

Yao, A.C., 1982. Protocols for secure computations, in: *23rd annual symposium on foundations of computer science (sfcs 1982)*, IEEE. pp. 160–164.

Zerka, F., Urovi, V., Vaidyanathan, A., Barakat, S., Leijenaar, R.T., Walsh, S., Gabrani-Juma, H., Miraglio, B., Woodruff, H.C., Dumontier, M., et al., 2020. Blockchain for privacy preserving and trustworthy distributed machine learning in multicentric medical imaging (c-distrim). *Ieee Access* 8, 183939–183951.

Solution	Rigid Point Clouds Registration				
	Displacement RMSE (mm)	Time <i>party</i> ₁ (s)	Time <i>party</i> ₂ (s)	Comm. <i>party</i> ₁ (MB)	Comm. <i>party</i> ₂ (MB)
PPIR(MPC)	1.11 ± 0.31	0.02	0.02	0.03	0.03
PPIR(FHE)	1.26 ± 0.37	1.10	0.18	1.42	35.93

Table A1: Rigid Point Clouds registration test, comparison between PPIR(MPC) and PPIR(FHE). Registration metrics are reported as mean and standard deviation. Efficiency metrics in terms of average across iterations. RMSE: root mean square error.

Appendix

A 1. Rigid Point Cloud Registration

Let $\{z_i\}$, $\{z'_i\}$ two finite n size point sets where $z_i, z'_i \in \mathbb{R}^d$, to continue ... A non-iterative least-squares approach to match two sets of points, was proposed by Arun et al. (1987). The method uses singular value decomposition (SVD) and is trying to minimize the following cost function:

$$\Sigma^2 = \sum_{i=1}^n \|z'_i - (Rz_i + \mathbf{t})\|^2, \quad (\text{A1})$$

where R is the rotation matrix and \mathbf{t} is the translation vector. Let define \bar{z} and \bar{z}' to represent the centroids of $\{z_i\}$ and $\{z'_i\}$ respectively. Let \hat{R} being the estimated rotation matrix, and $\hat{\mathbf{t}}$ being the estimated translation. The method lies in the algorithm for finding \hat{R} detailed below:

1. Calculate the following quantities:

$$\begin{aligned} q_i &= p_i - \bar{z} \\ q'_i &= p'_i - \bar{z}' \end{aligned}$$

for all $0 \leq i \leq n$, which are distances from each point to its centroid.

2. Calculate $L \in \mathbb{R}^{d \times d}$, $L = \sum_{i=1}^n q_i q_i^T$, in vectorized form:

$$L = Q \cdot Q^T, \quad (\text{A2})$$

where $Q, Q' \in \mathbb{R}^{d \times n}$.

3. Find SVD of L , namely $L = U \Lambda V^T$;
4. Calculate $X = V U^T$;
5. Check the determinant of X . If it equals to $+1$, then $\hat{R} = X$ and $\hat{\mathbf{t}} = q' - \hat{R}q$. If the determinant equals to -1 , the algorithm has failed.

We note that the only operation that requires the joint availability of information from both parties is Equation A2 which can be computed with a matrix multiplication with PPIR(MPC) and PPIR(FHE) as reported in Figure A7.

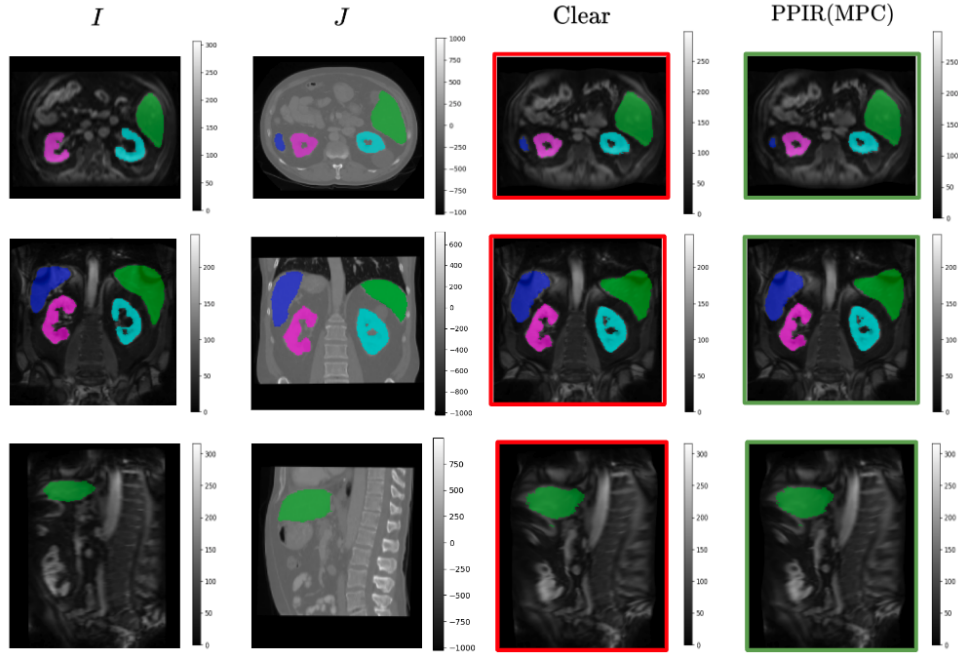


Fig. 3: Qualitative results for diffeomorphic registration with CC between 3D medical images from the AbdomenMRCT dataset (Hering *et al.*, 2022). The images are presented in a 3×4 grid, with the first row representing the axial axis, the second row the coronal axis, and the third row the sagittal axis. First and second column show respectively MRI and CT images. The third column shows the MRI transformed using CLEAR, while the fourth column shows the MRI transformed using PPIR(MPC). The transformed images are highlighted by red and green frames, respectively.

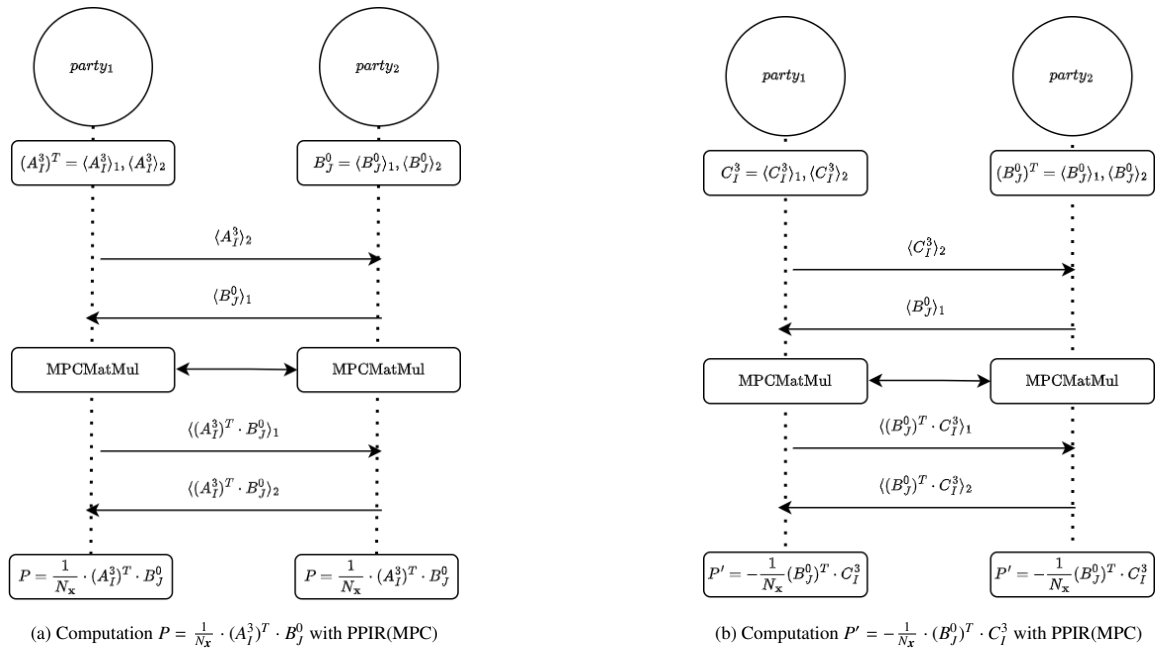

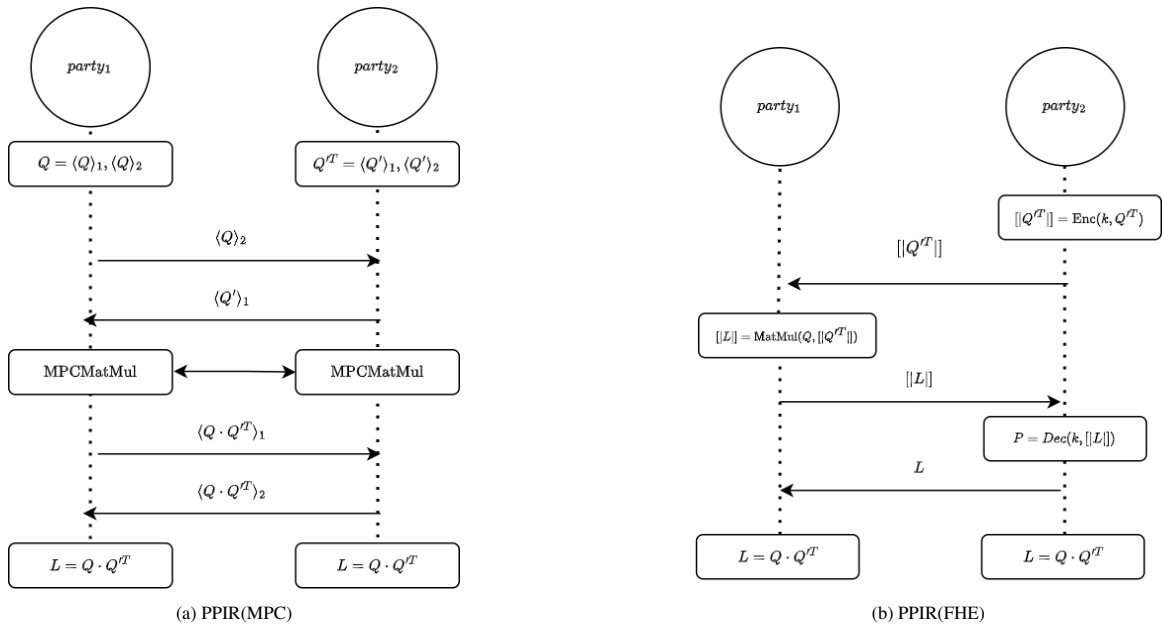


Fig. A1: Optimization of MI loss: proposed framework to calculate matrix multiplication $P = \frac{1}{N_x} \cdot (A_I^3)^T \cdot B_J^0$ and $P' = -\frac{1}{N_x} \cdot (B_J^0)^T \cdot C_I^3$ based on PPIR(MPC).


 Fig. A2: Optimization of ANTS NCC loss: proposed framework to calculate $\frac{2D}{EF}$ and $(\bar{J} - \frac{D}{E} \bar{I})$ based on PPIR(MPC).

 Fig. A3: Optimization of rigid point cloud: proposed framework to compute matrix multiplication $L = Q \cdot Q^T$ based on PPIR(MPC) and PPIR(FHE).

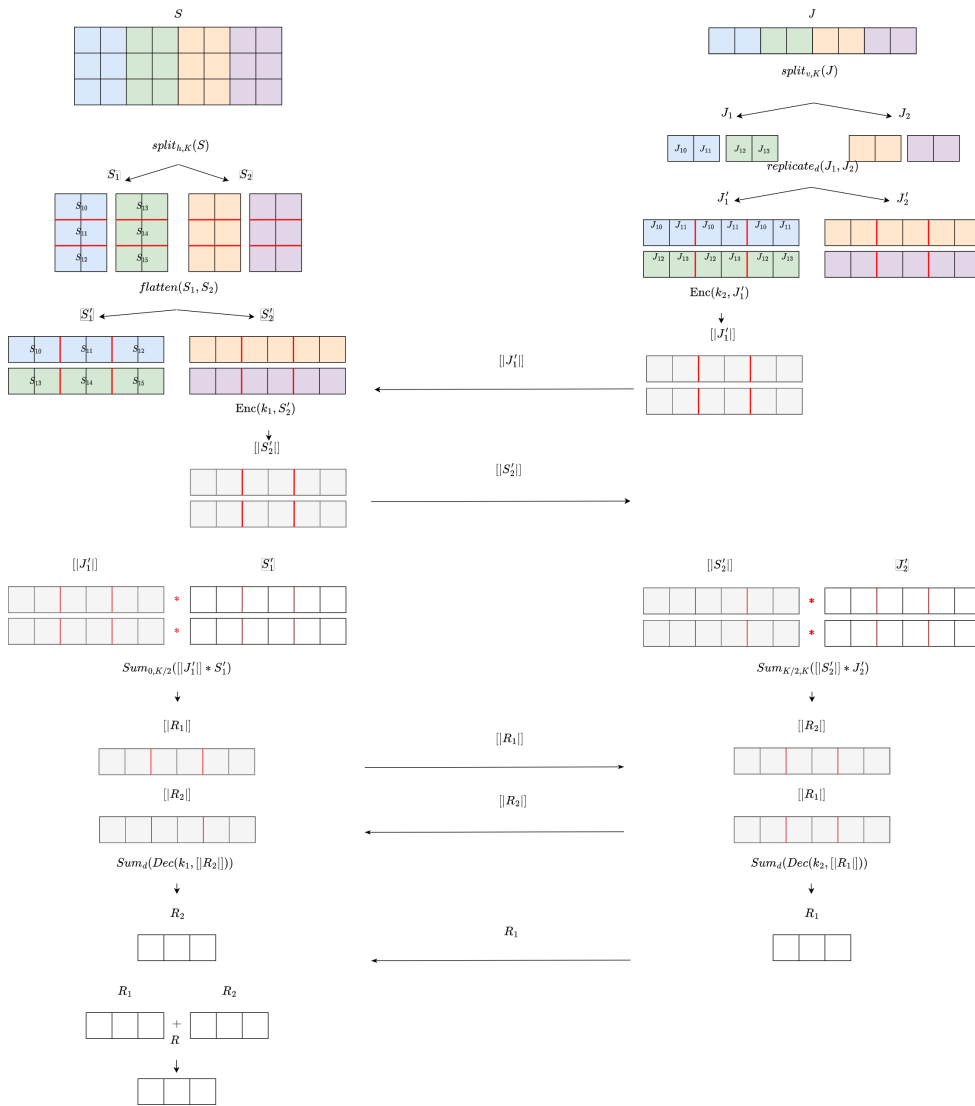


Fig. A4: Proposed framework to compute matrix-vector multiplication $S^T \cdot J$ based on PPIR(FHE)-v2.

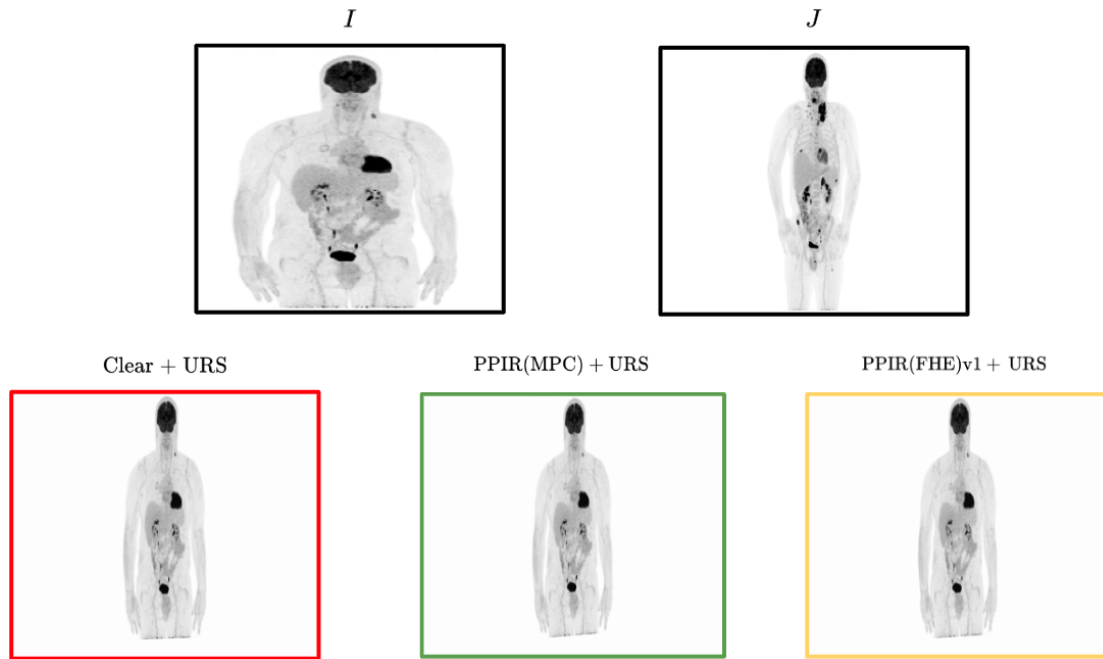


Fig. A5: Qualitative results for affine registration with SSD between 2D medical images. The red frame is the transformed moving image using CLEAR+URS registration. Green and Yellow frames are the transformed images using respectively PPIR(MPC)+URS and PPIR(FHE)v1+URS.

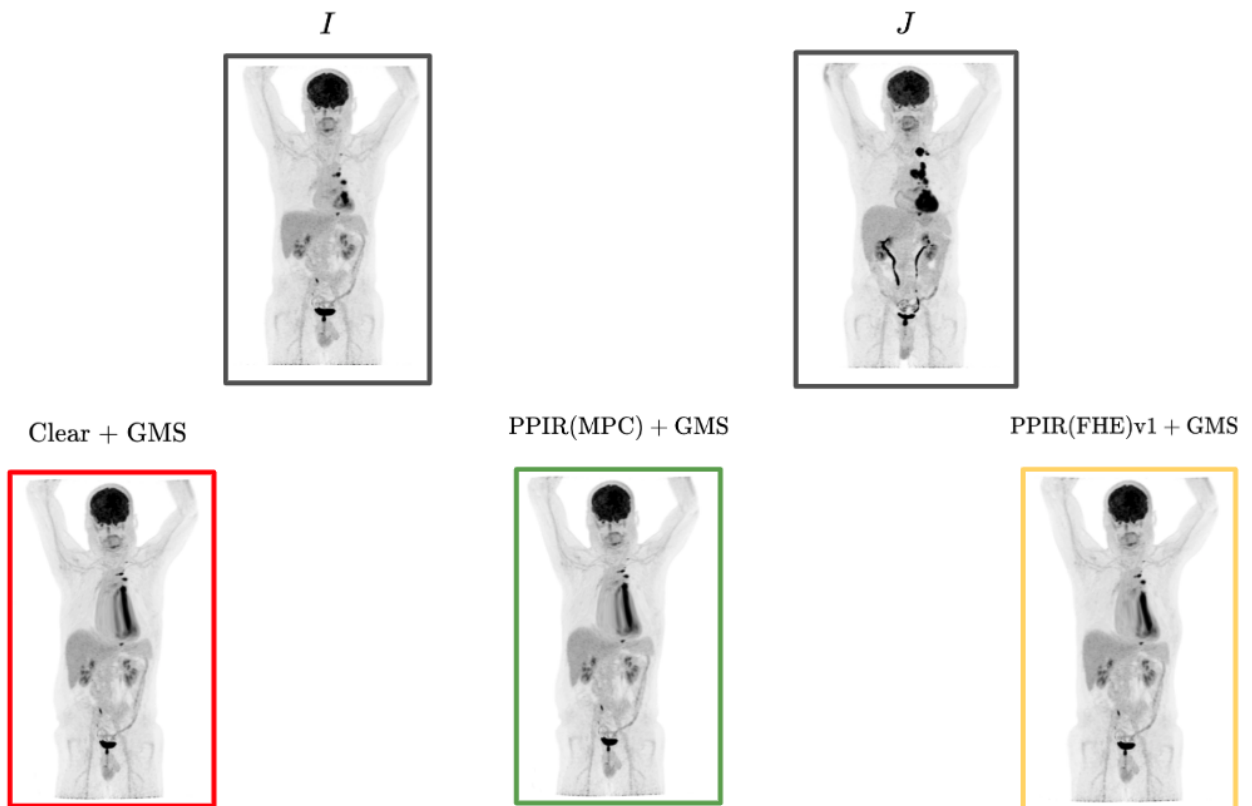


Fig. A6: Qualitative results for Cubic splines registration with SSD between 2D medical images. The red frame is the transformed moving image using CLEAR+GMS registration. Green and Yellow frames are the transformed images using respectively PPIR(MPC)+GMS and PPIR(FHE)v1+GMS.

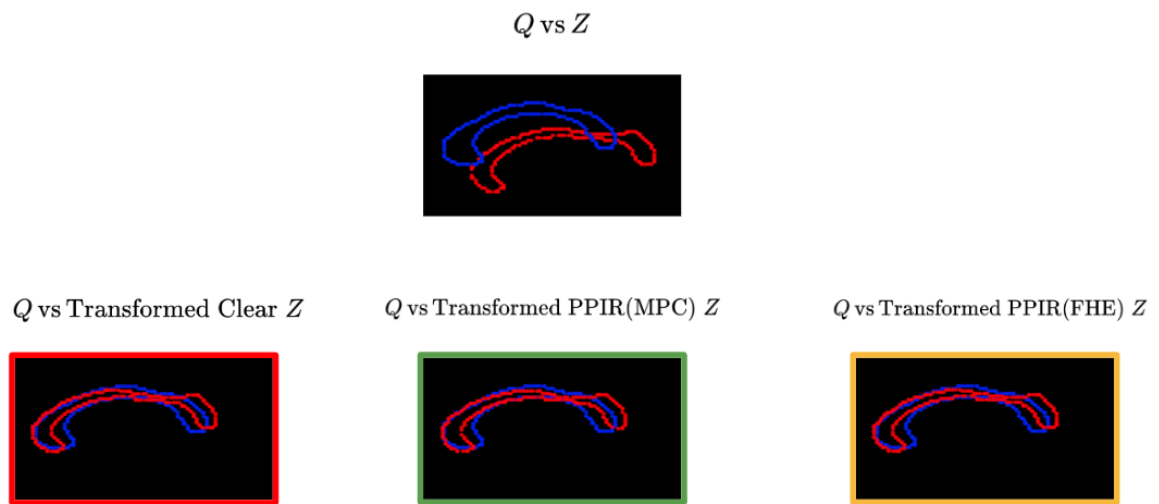


Fig. A7: Qualitative results for rigid point cloud registration between 2D corpus callosum point sets. The red frame is the transformed moving image using CLEAR registration. Green and Yellow frames are the transformed images using respectively PPIR(MPC) and PPIR(FHE)v1.