



**HAL**  
open science

## Privacy Preserving Image Registration

Riccardo Taiello, Melek Önen, Olivier Humbert, Marco Lorenzi

► **To cite this version:**

Riccardo Taiello, Melek Önen, Olivier Humbert, Marco Lorenzi. Privacy Preserving Image Registration. MICCAI 2022 - Medical Image Computing and Computer Assisted Intervention, Sep 2022, Singapore, Singapore. hal-03697446v3

**HAL Id: hal-03697446**

**<https://inria.hal.science/hal-03697446v3>**

Submitted on 18 Sep 2022

**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.

# Privacy Preserving Image Registration

Riccardo Taiello<sup>1,2,3</sup>, Melek Önen<sup>2</sup>, Olivier Humbert<sup>3</sup>, and Marco Lorenzi<sup>1,3</sup>

<sup>1</sup> Epione Research Project, Inria, Sophia Antipolis, France

<sup>2</sup> EURECOM, Sophia Antipolis, France

<sup>3</sup> Université Côte d’Azur, Nice, France

{riccardo.taiello,marco.lorenzi}@inria.fr

melek.onen@eurecom.fr

olivier.humbert@univ-cotedazur.fr

**Abstract.** Image registration is a key task in medical imaging applications, allowing to represent medical images in a common spatial reference frame. Current literature on image registration is generally based on the assumption that images are usually accessible to the researcher, from which the spatial transformation is subsequently estimated. This common assumption may not be met in current practical applications, since the sensitive nature of medical images may ultimately require their analysis under privacy constraints, preventing to share the image content in clear form. 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 first propose to optimize the underlying image registration operations using gradient approximations. We further revisit the use of homomorphic encryption and use a packing method to allow the encryption and multiplication of large matrices more efficiently. We demonstrate our privacy preserving framework in linear and non-linear registration problems, evaluating its accuracy and scalability with respect to standard image registration. Our results show that privacy preserving image registration is feasible and can be adopted in sensitive medical imaging applications.

**Keywords:** Image Registration · Privacy enhancing technologies · Trustworthiness

## 1 Introduction

Image Registration is a crucial task in medical imaging applications, allowing to spatially align imaging features between two or multiple scans. Image registration is a key component of state-of-the-art methods for atlas-based segmentation [9,31], morphological and functional analysis [3,11], multi-modal data integration [17], and longitudinal analysis [4,26].

Overall, 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 underlying the spatial transformations, similarity metric, image dimensionality and optimization strategy [30].

Image registration is the workhorse of many real-life medical imaging software and applications, including public web-based services for automated segmentation and labeling 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. There are also emerging data analysis paradigms, such as Federated Learning (FL) [22], where medical images can be jointly analysed in multi-centric scenarios to perform group analysis [14]. In these setting, the creation of registration-based image templates [3] is currently not possible without disclosing the image information. Due to the evolving juridical landscape on data protection, these applications of image registration are no longer compliant with regulations currently existing in many countries, such as the European General Data Protection Regulation (GDPR) [2], or the US Health Insurance Portability and Accountability Act (HIPAA) [1]. Medical imaging information falls within the realm of personal health data [20] 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 enabling data processing without disclosing it in clear hold great potential in sensitive data analysis problems (e.g., [19]). Examples of such approaches are Secure-Multi-Party-Computation (MPC) [33] and Homomorphic Encryption (HE) [27]. 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 neither the input data nor the result of the computation.

This work presents a new methodological framework allowing image registration under privacy constraints. To this end, we reformulate the typical 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 privacy preserving techniques, we investigate strategies for the practical use of privacy preserving image registration (PPIR) through gradient approximations, array packing and matrix partitioning. In our experiments we evaluate the effectiveness of PPIR in linear and non-linear registration problems. Our results demonstrate the feasibility of PPIR, and pave the way to the application of image registration in sensitive medical imaging applications.

## 2 Problem statement

Given images  $I, J : \mathbb{R}^d \mapsto \mathbb{R}$ , image registration aims at estimating the parameters  $\mathbf{p}$  of a spatial transformation  $\mathbf{W}_{\mathbf{p}} \in \mathbb{R}^d \mapsto \mathbb{R}^d$ , either linear or non-linear, max-

imizing the spatial overlap between  $J$  and the transformed image  $I(\mathbf{W}_{\mathbf{p}})$ . For example, a typical cost function to optimize the registration problem is the sum of squared intensity differences (SSD) evaluated on the set of image coordinates:

$$\text{SSD}(I, J, \mathbf{p}) = \operatorname{argmin}_{\mathbf{p}} \sum_{\mathbf{x}} \left[ I(\mathbf{W}_{\mathbf{p}}(\mathbf{x})) - J(\mathbf{x}) \right]^2 \quad (1)$$

Equation (1) can be typically optimized through gradient-based methods, where the parameters  $\mathbf{p}$  are iteratively updated until convergence. In particular, under a Gauss-Newton optimization scheme, the parameters update of the spatial transformation can be computed through Equation (2):

$$\Delta \mathbf{p} = H^{-1} \cdot \sum_{\mathbf{x}} S(\mathbf{x}) \cdot (I(\mathbf{W}_{\mathbf{p}}(\mathbf{x})) - J(\mathbf{x})), \quad (2)$$

where  $S(\mathbf{x}) = \nabla I(\mathbf{x}) \frac{\partial \mathbf{W}_{\mathbf{p}}(\mathbf{x})}{\partial \mathbf{p}}$  quantifies image and transformation gradients, and  $H = \sum_{\mathbf{x}} \left( \nabla I(\mathbf{x}) \frac{\partial \mathbf{W}_{\mathbf{p}}(\mathbf{x})}{\partial \mathbf{p}} \right)^T \left( \nabla I(\mathbf{x}) \frac{\partial \mathbf{W}_{\mathbf{p}}(\mathbf{x})}{\partial \mathbf{p}} \right)$  is the second order term obtained from Equation (1) through linearization [5, 24].

The solution of this problem requires the joint availability of both images  $I$  and  $J$ , as well as of the gradients of  $I$  and of  $\mathbf{W}_{\mathbf{p}}$ . In a privacy preserving setting, this information may not be available, and the computation of Equation (2) is therefore impossible. We thus consider a scenario with two parties, *party*<sub>1</sub>, and *party*<sub>2</sub>, whereby *party*<sub>1</sub> owns image  $I$  and *party*<sub>2</sub> 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*<sub>1</sub> has access to the transformation parameters  $\mathbf{p}$ , and that is also in charge of computing the update at each optimization step. In particular, to compute the registration update  $\Delta \mathbf{p}$  of Equation (2), the only operation requiring 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 as a matrix-vector multiplication on vectorized quantities,  $R = S^T \cdot J$ .

### 3 Methods

Before presenting PPIR in Section 3.2, we introduce in Section 3.1 the cryptographic tools underlying the proposed framework.

#### 3.1 Secure Computation

**Secure Multi-Party Computation.** Introduced by Yao in [33], 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 [12], 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 [12], SPDZ defines a dedicated MPCMUL operation to compute matrix-vector multiplication within a secure two-party protocol. **Homomorphic Encryption.** Initially introduced by Rivest et al. in [27], HE enables the execution of operations over encrypted data without disclosing neither the input data, nor the result of the computation. Hence, party 1 encrypts the input with her public key and sends this encryption to party 2. Party 2, in turn, evaluates a circuit over this encrypted input and sends the result, which still remains encrypted, back to party 1 which can finally decrypt the result. Among various HE schemes, CKKS [10] supports the execution of all operations over encrypted real values and is considered as a fully homomorphic encryption (FHE). 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 depends on multiple parameters including the degree of the polynomial  $N$ , which is usually sufficiently large (e.g.  $N = 4096$ , or  $N = 8192$ ).

### 3.2 PPIR: Privacy preserving image registration

In order to ensure the privacy of images  $I$  and  $J$  against  $party_2$  and  $party_1$  respectively, we propose to investigate the use of MPC and FHE to develop PPIR. Figure 1 illustrates how these two cryptographic tools are employed to ensure the privacy of the images during registration. As previously mentioned, the only operation that needs to be jointly executed by the parties in a privacy preserving manner is the matrix-vector multiplication:  $R = S^T \cdot J$ , where  $S^T$  is only known to  $party_1$ , and  $J$  to  $party_2$ .

When MPC is integrated (Figure 1a),  $party_1$  secretly shares the matrix  $S^T$  to obtain  $(\langle S_1 \rangle, \langle S_2 \rangle)$ , while  $party_2$  secretly shares the image  $J$  to obtain  $(\langle J_1 \rangle, \langle J_2 \rangle)$ . Each party further receives its corresponding share, namely:  $party_1$  holds  $(\langle S_1 \rangle, \langle J_1 \rangle)$  and  $party_2$  holds  $(\langle S_2 \rangle, \langle J_2 \rangle)$ . Parties further execute a circuit with the MPCMUL operations to compute the 2-party dot product between  $S^T$  and  $J$ . The parties further synchronize to let  $party_1$  to obtain the product, and to finally calculate  $\Delta p$  (see Equation (2)).

When using FHE (Figure 1b),  $party_2$  uses a FHE key  $k$  to encrypt  $J$  and obtain:  $\llbracket J \rrbracket \leftarrow \text{ENC}(k, J)$ . This encrypted image is sent to  $party_1$ , who computes the encrypted result  $\llbracket R \rrbracket$  of the matrix-vector multiplication. In this framework, only the 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 through decryption:  $R = \text{DEC}(k, \llbracket R \rrbracket)$ . Finally,  $party_1$  receives  $R$  in clear form and can therefore compute  $\Delta p$ .

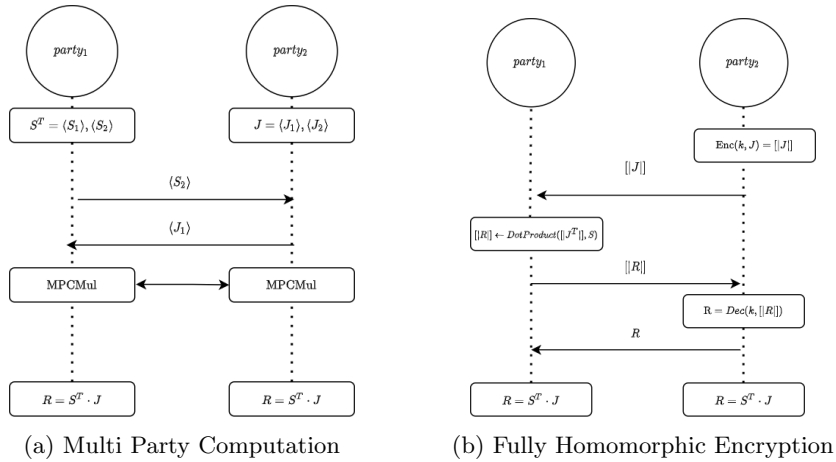


Fig. 1: Proposed framework to compute the matrix-vector multiplication  $S^T \cdot J$  relying on MPC and FHE.

Thanks to the privacy and security guarantees of these cryptographic tools, during the entire registration procedure the content of the image data  $S$  and  $I$  is never disclosed to the opposite party. Nevertheless, 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 [7, 16], notably affecting computation time and communication bandwidth between parties. To tackle this issue, in what follows we introduce in the schemes of Figure 1 computational strategies to effectively reduce the dimensionality of the image information through sampling, and to improve the scalability of the algebraic operations when using these cryptographic tools.

**Gradient sampling.** The update of Equation (2) is computed on the vectorized images, which are large-dimensional arrays representing all the image pixels (or voxels). Since the registration gradient is in general mostly driven by a fraction of the image content, e.g. image boundaries in case of the SSD metric, a reasonable approximation of Equation (2) can be obtained by assessment on relevant image locations only. This idea has been introduced in medical image registration [21, 29, 32], and is here adopted to reduce the dimensionality of the arrays on which encryption is performed. In our works we test two different techniques: (i): Uniformly Random Selection (URS), proposed by [21, 32], in which a random subset of dimension  $l \leq s$  of spatial coordinates is sampled at every iteration with uniform probabilities,  $Pr(\mathbf{x}) = \frac{1}{s}$ ; and (ii): Gradient Magnitude Sampling (GMS) [29], consisting in sampling a subset of coordinates with probability proportional to the norm of the image gradient,  $Pr(\mathbf{x}) \sim \|\nabla I(\mathbf{x})\|$ .

**Matrix partitioning in FHE.** In addition to gradient sampling, we propose a specific optimization dedicated to PPIR with FHE, in particular when the CKKS algorithm is adopted. CKKS allows packing multiple inputs into a

single ciphertext to decrease the number of homomorphic operations. In order to optimize the matrix-vector multiplication, we propose to partition the image vector  $J$  into  $K$  sub-arrays of dimension  $D$ , and the matrix  $S^T$  into  $K$  submatrices of dimension  $|\mathbf{p}| \times D$ . Once all sub-arrays  $J_i$  are encrypted, we propose to iteratively apply the (DOTPRODUCT) proposed by [7] between each sub-matrix and corresponding sub-array; these intermediate results are then summed up to obtain the final result, namely:  $\llbracket R \rrbracket = \sum_{i=0}^K \text{DOTPRODUCT}(\llbracket J_i^T \rrbracket, S_i) = S^T \cdot \llbracket J \rrbracket$ .

## 4 Experimental Results

We demonstrate and assess PPIR in two examples based on linear and non-linear alignment of respectively positron emission tomography (PET) and anatomical magnetic resonance (MR) images.

**Dataset.** PET data consists of 18-Fluoro-Deoxy-Glucose ( $^{18}FDG$ ) whole body Positron Emission Tomography (PET). The images here considered are a frontal view of the maximum intensity projection reconstruction, obtained by 2D projection of the voxels with highest intensity across views ( $1260 \times 1090$  pixels).

MR images are obtained from brain scans of the Alzheimer’s Disease Neuroimaging Initiative [23]. Images underwent a standard processing pipeline to estimate grey matter density maps [3]. The subsequent registration experiments are performed on the extracted mid-coronal slice, of dimension  $121 \times 121$  pixels.

**Implementation.** PET image alignment was performed by optimizing the transformation  $\mathbf{W}_p$  of Equation (1) with respect to affine registration parameters. The registration of brain grey matter density images was performed by non-linear registration based on a cubic spline model (one control point every five pixels along both dimensions). For both affine and non-linear cases, the registration was performed between two randomly selected patients’ images.

Concerning the PPIR framework with the affine transformation, tests are carried for both MPC and FHE by considering the entire images, and by using gradient approximation techniques (Section 3). The sampling seed for gradient approximation is the same for each test. Due to the smaller dimension of the brain grey matter images, non-linear PPIR with cubic-splines is applied directly to the full data. For MPC we set as prime modulus  $Q = 2^{32}$ . For FHE, we define the polynomial degree modulus as  $N = 4096$ , and set the resizing parameter  $D$  to optimize the trade-off between runtime and bandwidth. Since  $D$  needs to be a divider of the image size, for PET image data we set  $D = 128$ , while for the grey matter images we set  $D = 121$ . The PPIR framework is implemented using two state-of-the-art libraries: PySyft [28] supporting SPDZ’s two-party computation, and TenSeal [7] for CKKS. 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 the MPC shares and FHE encryption keys. The code is released in a GitHub repository<sup>4</sup>. We used Weights & Biases [8] for

<sup>4</sup> [https://github.com/rtaiello/pp\\_image\\_registration](https://github.com/rtaiello/pp_image_registration)

experiment tracking, and the links of our tracked results are available in the [GitHub repository](#).

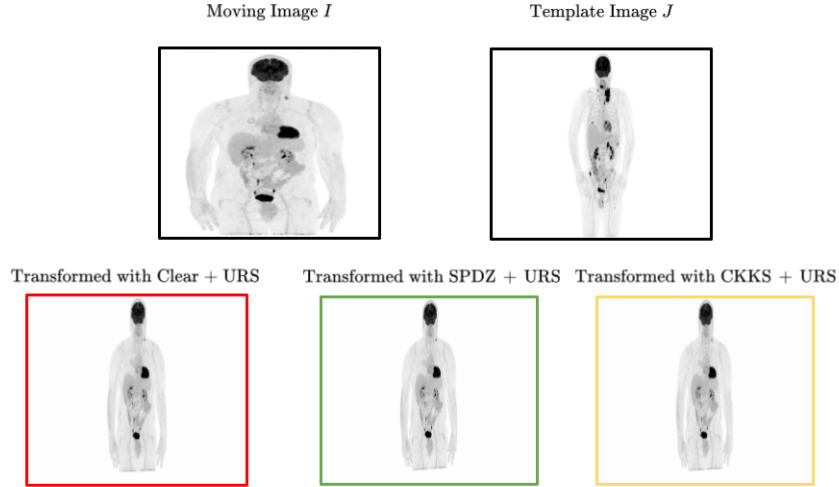


Fig. 2: Qualitative results for affine registration. The red frame is the transformed moving image using `CLEAR+URS` registration. Green and Yellow frames are the transformed images using respectively `MPC+URS` and `FHE+URS` PPIR. The yellow image is the transformed moving using PPIR with `CKKS`.

The quality of PPIR is assessed by comparing the registration results with respect to the ones obtained with standard registration on clear images (`CLEAR`). The metrics considered are the image intensity difference at the optimum, the overall 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.

**Results.** Table 1 (Registration metrics) shows that affine PPIR through `SPDZ` leads to negligible differences with respect to `CLEAR` in terms of number of iterations, intensity and displacement. Registration with `CKKS` is instead not possible when considering the entire images, due to computational complexity, and is also associated to generally larger approximations when using `URS` and `GMS`. Nevertheless, Figure 2 shows that neither `MPC` nor `FHE` do decrease the overall quality of the affine registered images. Additional registration results are available in Appendix. Table 1 (Efficiency metrics) shows that `SPDZ` performed on the full images has highest computation time and communication bandwidth. These figures sensibly improve when using `URS` or `GMS`, by factors 10x and 20x for respectively time and bandwidth. Concerning `CKKS`, we note the uneven time and bandwidth requirements between clients, due to the asymmetry of



Table 1: Affine registration test. Registration metrics are reported as mean and standard deviation. Efficiency metrics in terms of average across iterations. RMSE: root mean square error.

Affine registration metrics				
Solution	Intensity Error (SSD)	Num. Iteration	Displacement RMSE CLEAR vs PPIR ( <i>mm</i> )	
CLEAR	4.34 ± 0.0	118 ± 0.0	-	
SPDZ	4.34 ± 0.0	114.8 ± 4.0	1.81 ± 0.02	
CKKS	<b>x</b>	<b>x</b>	<b>x</b>	
CLEAR + URS	4.38 ± 0.0	61 ± 0.0	-	
SPDZ + URS	4.34 ± 0.0	60.4 ± 6.85	16.49 ± 3.74	
CKKS ( <i>D</i> = 128) + URS	4.34 ± 0.10	61.80 ± 4.82	23.31 ± 2.72	
CLEAR + GMS	4.34 ± 0.0	63 ± 0.0	-	
SPDZ + GMS	4.34 ± 0.0	59.80 ± 6.20	6.21 ± 1.49	
CKKS ( <i>D</i> = 128) + GMS	4.34 ± 0.05	60.4 ± 5.12	5.17 ± 1.40	
Affine efficiency metrics				
Solution	Time <i>party</i> <sub>1</sub> (s)	Time <i>party</i> <sub>2</sub> (s)	Comm. <i>party</i> <sub>1</sub> (MB)	Comm. <i>party</i> <sub>2</sub> (MB)
CLEAR	0.0	0.0	-	-
SPDZ	0.13	0.13	1.54	1.54
CKKS	<b>x</b>	<b>x</b>	<b>x</b>	<b>x</b>
CLEAR + URS	0.0	0.0	-	-
SPDZ + URS	0.02	0.02	0.20	0.20
CKKS ( <i>D</i> = 128) + URS	2.55	0.02	0.06	0.01
CLEAR + GMS	0.0	0.0	-	-
SPDZ + GMS	0.02	0.02	0.20	0.20
CKKS ( <i>D</i> = 128) + GMS	2.51	0.02	0.06	0.01

Table 2: Non-linear registration test. Registration metrics are reported as mean and standard deviation. Efficiency metrics in terms of average across iterations. RMSE: root mean square error.

Cubic splines registration metrics				
Solution	Intensity Error (SSD)	Num. Iteration	Displacement RMSE CLEAR vs PPIR ( <i>mm</i> )	
CLEAR	0.65 ± 0.0	413 ± 0.0	-	
SPDZ	0.65 ± 0.0	345.70 ± 91.22	7.31 ± 1.86	
CKKS	0.64 ± 0.0	224.7 ± 79.15	9.50 ± 4.34	
Cubic splines efficiency metrics				
Solution	Time <i>party</i> <sub>1</sub> (s)	Time <i>party</i> <sub>2</sub> (s)	Comm. <i>party</i> <sub>1</sub> (MB)	Comm. <i>party</i> <sub>2</sub> (MB)
CLEAR	0.0	0.0	-	-
SPDZ	0.63	0.63	21.47	28.98
CKKS	3.41	0.00	0.06	0.01

encryption operations and communication protocol (Figure 1). Finally, Table 2 reports the metrics for the non-linear registration test. Concerning registration accuracy, we draw similar conclusions to the affine case, where SPDZ leads to minimum differences with respect to CLEAR, while CKKS seems slightly inferior. SPDZ is associated to lower execution time and higher computation bandwidth, due to the larger number of parameters of the cubic splines, affecting the size of the matrix *S*. While CKKS has higher execution time, the demanded bandwidth is inferior to the one of SPDZ, since the encrypted template image is transmitted only once.

## 5 Conclusion

This work introduces privacy preserving image registration (PPIR), a novel framework to allow image registration when images are confidential and cannot

be disclosed in clear. PPIR is developed with MPC and FHE and implements effective strategies to mitigate their known computational and communication overhead. PPIR is demonstrated and evaluated through a series of quantitative benchmarks in both linear and non-linear image registration problems. Our results highlight the existing trade-off between registration performance and efficiency of the different PPIR schemes.

Future extensions of this work will be devoted to the benchmarking of our framework in more general scenarios, involving 3D medical image data and multimodal registration problem. The application to multimodal data will require the extension of our framework to account for different similarity metrics, such as Mutual Information or Normalized Cross-Correlation [25, 32]. Moreover, the effectiveness of sampling through URS and GMS motivates the adoption of sparse image registration frameworks, especially for non-linear registration [13, 15]. Another relevant research direction concerns the development of PPIR in deep-learning based approaches [6, 18].

Overall, this study shows that PPIR is feasible and can therefore be adopted in sensitive medical imaging applications.

## References

1. Health Resources and Services Administration. Health insurance portability and accountability act, 1 (1996), U.S. Dept. of Labor, Employee Benefits Security Administration.
2. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (General Data Protection Regulation) (2016-05-04), European Union.
3. Ashburner, J., Friston, K.J.: Voxel-based morphometry—the methods. *Neuroimage* **11**(6), 805–821 (2000)
4. Ashburner, J., Ridgway, G.R.: Symmetric diffeomorphic modeling of longitudinal structural MRI. *Frontiers in neuroscience* **6**, 197 (2013)
5. Baker, S., Matthews, I.: Lucas-Kanade 20 years on: A unifying framework. *International journal of computer vision* **56**(3), 221–255 (2004)
6. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* **38**(8), 1788–1800 (2019)
7. Benaissa, A., Retiat, B., Cebere, B., Belfedhal, A.E.: Tenseal: A library for encrypted tensor operations using homomorphic encryption. *CoRR* **abs/2104.03152** (2021), <https://arxiv.org/abs/2104.03152>
8. Biewald, L.: Experiment tracking with weights and biases (2020), <https://www.wandb.com/>, software available from wandb.com
9. Cardoso, M.J., Leung, K., Modat, M., Keihaninejad, S., Cash, D., Barnes, J., Fox, N.C., Ourselin, S., Initiative, A.D.N., et al.: STEPs: Similarity and truth estimation for propagated segmentations and its application to hippocampal segmentation and brain parcellation. *Medical image analysis* **17**(6), 671–684 (2013)
10. Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: *International Conference on the Theory and Application of Cryptology and Information Security*. pp. 409–437. Springer (2017)

11. Dale, A.M., Fischl, B., Sereno, M.I.: Cortical surface-based analysis: I. segmentation and surface reconstruction. *Neuroimage* **9**(2), 179–194 (1999)
12. Damgard, I., Pastro, V., Smart, N., Zakarias, S.: Multiparty computation from somewhat homomorphic encryption. *Cryptology*, ePrint Archive, Report 2011/535 (2011), <https://ia.cr/2011/535>
13. Fawzi, A., Frossard, P.: Image registration with sparse approximations in parametric dictionaries. *SIAM Journal on Imaging Sciences* **6**(4), 2370–2403 (2013)
14. Gazula, H., Holla, B., Zhang, Z., Xu, J., Verner, E., Kelly, R., Jain, S., Bharath, R.D., Barker, G.J., Basu, D., et al.: Decentralized multisite vbm analysis during adolescence shows structural changes linked to age, body mass index, and smoking: a coinstac analysis. *Neuroinformatics* **19**(4), 553–566 (2021)
15. Ha, I.Y., Wilms, M., Handels, H., Heinrich, M.P.: Model-based sparse-to-dense image registration for realtime respiratory motion estimation in image-guided interventions. *IEEE Transactions on Biomedical Engineering* **66**(2), 302–310 (2018)
16. Haralampieva, V., Rueckert, D., Passerat-Palmbach, J.: 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 (2020)
17. Heinrich, M.P., Jenkinson, M., Bhushan, M., Matin, T., Gleeson, F.V., Brady, J.M., Schnabel, J.A.: Non-local shape descriptor: A new similarity metric for deformable multi-modal registration. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 541–548. Springer (2011)
18. Krebs, J., Delingette, H., Mailhé, B., Ayache, N., Mansi, T.: Learning a probabilistic model for diffeomorphic registration. *IEEE transactions on medical imaging* **38**(9), 2165–2176 (2019)
19. Lauter, K.: Private AI: Machine Learning on Encrypted Data. Tech. rep. (2021), eprint report <https://eprint.iacr.org/2021/324.pdf>
20. Lotan, E., Tschider, C., Sodickson, D.K., Caplan, A.L., Bruno, M., Zhang, B., Lui, Y.W.: Medical imaging and privacy in the era of artificial intelligence: myth, fallacy, and the future. *Journal of the American College of Radiology* **17**(9), 1159–1162 (2020)
21. Mattes, D., Haynor, D.R., Vesselle, H., Lewellen, T.K., Eubank, W.: Pet-ct image registration in the chest using free-form deformations. *IEEE transactions on medical imaging* **22**(1), 120–128 (2003)
22. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: *Artificial intelligence and statistics*. pp. 1273–1282. PMLR (2017)
23. Mueller, S.G., Weiner, M.W., Thal, L.J., Petersen, R.C., Jack, C., Jagust, W., Trojanowski, J.Q., Toga, A.W., Beckett, L.: The alzheimer’s disease neuroimaging initiative. *Neuroimaging Clinics* **15**(4), 869–877 (2005)
24. Pennec, X., Cachier, P., Ayache, N.: Understanding the “Demon’s algorithm”: 3D non-rigid registration by gradient descent. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 597–605. Springer (1999)
25. Pilu, M.: A direct method for stereo correspondence based on singular value decomposition. In: *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. pp. 261–266. IEEE (1997)
26. Reuter, M., Rosas, H.D., Fischl, B.: Highly accurate inverse consistent registration: a robust approach. *Neuroimage* **53**(4), 1181–1196 (2010)
27. Rivest, R.L., Adleman, L., Dertouzos, M.L., et al.: On data banks and privacy homomorphisms. *Foundations of secure computation* **4**(11), 169–180 (1978)

28. Ryffel, T., Trask, A., Dahl, M., Wagner, B., Mancuso, J., Rueckert, D., Passerat-Palmbach, J.: A generic framework for privacy preserving deep learning. arXiv preprint arXiv:1811.04017 (2018)
29. Sabuncu, M.R., Ramadge, P.J.: Gradient based nonuniform subsampling for information-theoretic alignment methods. In: The 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. vol. 1, pp. 1683–1686. IEEE (2004)
30. Schnabel, J.A., Heinrich, M.P., Papież, B.W., Brady, J.M.: Advances and challenges in deformable image registration: from image fusion to complex motion modelling. *Medical Image Analysis* **33**, 145–148 (2016)
31. Shattuck, D.W., Prasad, G., Mirza, M., Narr, K.L., Toga, A.W.: Online resource for validation of brain segmentation methods. *NeuroImage* **45**(2), 431–439 (2009)
32. Viola, P., Wells III, W.M.: Alignment by maximization of mutual information. *International journal of computer vision* **24**(2), 137–154 (1997)
33. Yao, A.C.: Protocols for secure computations. In: 23rd annual symposium on foundations of computer science (sfcs 1982). pp. 160–164. IEEE (1982)