



HAL
open science

Learn2Reg: Comprehensive Multi-Task Medical Image Registration Challenge, Dataset and Evaluation in the Era of Deep Learning

Alessa Hering, Lasse Hansen, Tony Mok, Albert Chung, Hanna Siebert, Stephanie Hager, Annkristin Lange, Sven Kuckertz, Stefan Heldmann, Wei Shao, et al.

► To cite this version:

Alessa Hering, Lasse Hansen, Tony Mok, Albert Chung, Hanna Siebert, et al.. Learn2Reg: Comprehensive Multi-Task Medical Image Registration Challenge, Dataset and Evaluation in the Era of Deep Learning. IEEE Transactions on Medical Imaging, 2023, 42 (3), pp.697-712. 10.1109/TMI.2022.3213983 . hal-04525072

HAL Id: hal-04525072

<https://hal.science/hal-04525072v1>

Submitted on 30 May 2024

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 4.0 International License

Learn2Reg: Comprehensive Multi-Task Medical Image Registration Challenge, Dataset and Evaluation in the Era of Deep Learning

Alessa Hering¹, Lasse Hansen¹, Tony C. W. Mok², Albert C. S. Chung³, Hanna Siebert, Stephanie Häger, Annkristin Lange⁴, Sven Kuckertz, Stefan Heldmann⁵, Wei Shao⁶, Sulaiman Vesal⁷, Mirabela Rusu, Geoffrey Sonn⁸, Théo Estienne⁹, Maria Vakalopoulou, Luyi Han¹⁰, Yunzhi Huang¹¹, Pew-Thian Yap¹², *Senior Member, IEEE*, Mikael Brudfors¹³, Yaël Balbastre, Samuel Joutard, Marc Modat, Gal Lifshitz¹⁴, Dan Raviv, Jinxin Lv¹⁵, Qiang Li, Vincent Jaouen¹⁶, Dimitris Visvikis, Constance Fourcade, Mathieu Rubeaux, Wentao Pan¹⁷, Zhe Xu¹⁸, Bailiang Jian, Francesca De Benetti¹⁹, Marek Wodzinski²⁰, *Associate Member, IEEE*, Niklas Gunnarsson²¹, Jens Sjölund²², Daniel Grzech²³, Huaqi Qiu²⁴, Zeju Li²⁵, *Graduate Student Member, IEEE*, Alexander Thorley, Jinming Duan²⁶, Christoph Großbröhm²⁷, *Graduate Student Member, IEEE*, Andrew Hoopes, Ingerid Reinertsen²⁸, Yiming Xiao, Bennett Landman²⁹, *Senior Member, IEEE*, Yuankai Huo³⁰, *Senior Member, IEEE*, Keelin Murphy³¹, Nikolas Lessmann³², Bram van Ginneken, Adrian V. Dalca, and Mattias P. Heinrich

Abstract—Image registration is a fundamental medical image analysis task, and a wide variety of approaches have been proposed. However, only a few studies have comprehensively compared medical image registration approaches on a wide range of clinically relevant tasks. This limits the development of registration methods, the adoption of research advances into practice, and a fair benchmark across competing approaches. The Learn2Reg challenge addresses these limitations by providing a multi-task medical image registration data set for comprehensive characterisation of deformable registration algorithms. A continuous evaluation will be possible at <https://learn2reg.grand-challenge.org>. Learn2Reg covers a wide range of anatomies (brain, abdomen, and thorax), modalities (ultrasound, CT, MR), availability of annotations, as well as intra- and inter-patient registration evaluation. We established an easily accessible framework for training and validation of 3D registration methods, which enabled the compilation of results of over 65 individual method submissions from more than 20 unique teams. We used a complementary set of metrics, including robustness, accuracy, plausibility, and runtime, enabling unique insight into the current state-of-the-art of medical image registration. This paper describes datasets, tasks, evaluation methods and results of the challenge, as well as results of further analysis of transferability to new datasets, the importance of label supervision, and resulting bias. While no single approach worked best across all tasks, many methodological aspects could be identified that push the performance of medical

image registration to new state-of-the-art performance. Furthermore, we demystified the common belief that conventional registration methods have to be much slower than deep-learning-based methods.

Index Terms—Medical image registration, challenge, evaluation.

I. INTRODUCTION

IMAGE registration is a fundamental task in medical image analysis and has been an active field of research for decades [1], [2], [3], [4]. Most studies that compared registration methods were focused on specific tasks or algorithmic aspects, and did not comprehensively characterise current approaches. With the recent success of deep learning strategies in image analysis, the degree and dependency of algorithms on (partially) labelled training data is often a crucial aspect in current research. The Learn2Reg challenge aims to gain insight into which methodological components and supervision strategies are best suited for a wide range of clinically useful 3D image registration tasks, and sets a new benchmark to evaluate and distinguish strengths and weaknesses of task-tailored solutions. Learn2Reg covers a wide range of anatomies (brain, abdomen and thorax), modalities (ultrasound, CT, MR) and auxiliary annotations (e.g. segmentation, keypoints). The challenge also includes both intra- and inter-patient registration tasks. Due to this broad range, it serves as a unique benchmark to evaluate the current state-of-the-art with respect to various qualities of registration algorithms: accuracy, robustness, plausibility and speed. Furthermore, no other medical image registration challenge has thoroughly analysed the benefits and shortcomings of learning- and optimisation-based strategies. To engage a wider participation from new research groups, Learn2Reg removes entry barriers by providing pre-processed

Manuscript received 27 July 2022; revised 30 September 2022; accepted 8 October 2022. Date of publication 20 October 2022; date of current version 2 March 2023. This work was supported in part by the German Federal Ministry of Education and Research under Grant 031L0202B. (Alessa Hering and Lasse Hansen contributed equally to this work.) (Corresponding author: Lasse Hansen.)

Please see the Acknowledgment section of this article for the author affiliations.

Digital Object Identifier 10.1109/TMI.2022.3213983

and pre-aligned images with additional annotations, as well as evaluation scripts and code for all evaluation metrics.

This overview ranks and scores results from over 65 entries from more than 20 teams throughout 2020 and 2021. We perform additional experiments to analyse the robustness towards cross-dataset transfer, the influence of the bias induced by only labelling certain anatomical regions, and direct comparisons of the supervision level of selected methods.

A. Related Work

In the following a brief overview of important related work on comparing (bio)-medical image registration, and its fundamental methodological choices that differentiate the wide range of metrics, optimisation, and supervision is given. General guidelines for setting up a fair and unbiased challenge have been recently thoroughly discussed in literature [5]. These criteria were adhered to in Learn2Reg and externally reviewed and confirmed by the MICCAI challenge team.

1) *Challenges*: There have previously been four prominent challenges for medical image registration. Three challenges focused on a single task: EMPIRE10 (lung CT) [6], CuRIOUS (intra-operative US and MR) [7], and ANHIR (histology) [8]. Each attracted at least ten participating teams and used various metrics for quantifying the performance. The EMPIRE10 challenge provided the most comprehensive evaluation including distances of manual landmark pairs, fissure segmentations, and Jacobian determinant values of the deformation field. This challenge also required (original) participants to perform live registrations during the MICCAI workshop in Beijing and therefore employed a time constraint on the computations. The Continuous Registration Challenge [9] co-organised with WBIR 2018 aimed at combining multiple tasks from previous benchmarks (lung CT and inter-patient brain MR). It addressed assessing registration quality as a service but is limited to algorithms that can be incorporated into the SuperElastix framework and therefore had limited participation.

2) *Benchmark Papers*: Several papers have compared multiple registration algorithms for a given dataset. In contrast to challenges, these benchmark papers did not have an open workshop format that enabled wide-spread participation. Nevertheless, their findings provided meaningful insights. Starting from RIRE [10], which compared rigid-body alignment of head MR (T1, T2), PET and CT, there have been several brain registration benchmarks - most notably the evaluation of 14 nonlinear iterative registration algorithms [11]. Fewer studies analysed abdominal registration, and included the evaluation of six affine and non-linear algorithms on inter-patient registration of the “beyond the cranial vault” dataset [12]. This study revealed large performance gaps and motivated our inclusion of this dataset to study the potential benefit of supervised (learning-based) algorithms. The DIR-Lab datasets [13] have been widely used to benchmark intra-patient CT lung motion estimation and provide a leaderboard for state-of-the-art comparison. All landmarks are publicly available, which makes the dataset prone to overfitting on the test data.

3) *Survey Papers and Baseline Methods*: Surveys on conventional medical image registration [2], [3] have comprehensively reviewed typical categories of approaches including

similarity metric, regulariser, and optimiser criteria. Due to the strong increase in the number of deep-learning-based registration paper in the last few years, several new surveys have been published (e.g. [4]) extending the typical categories with deep-learning specific categories like supervision-type and network architecture. Moreover, the training data and thus the registered body region and image modality are more important for deep-learning-based methods and get more into the focus of those survey papers. While few papers have evaluated their proposed registration method on more than two different registration tasks, there is a variety of public methods SyN [14], Elastix [15], NiftyReg [16] and deeds [17], and Voxelmorph [18] that are commonly used as baseline or comparison methods. When comparing only among deep-learning based methods simply re-training specific architectures on new data may be insufficient. Hence the use of a challenge benchmark that incorporates several generally applicable baselines is essential for a fair evaluation.

B. Contributions

Learn2Reg provides both datasets and an easily accessible benchmark for the first comprehensive evaluation of a wide-range of methods for inter- and intra-patient, mono- and multimodal medical registration. We introduce a complementary set of metrics, including robustness, accuracy, plausibility and speed, that follows the principles defined by the BIAS group [5] and could become an important data set collection for comparing new algorithms. Further analysis of label bias (for supervised methods), domain transfer and statistical testing of significant differences across algorithms and types of methods highlight the complementary strength and weaknesses of learning vs. non-learning-based approaches.

II. MATERIAL AND METHODS

A. Challenge Organisation

The Learn2Reg challenge is organised by Alessa Hering, Lasse Hansen, Adrian Dalca and Mattias Heinrich and is associated with MICCAI 2020 and 2021. The following tasks were included in 2020: CuRIOUS, Hippocampus MR, Abdomen CT-CT and Lung CT. In 2021, Abdomen MR-CT and OASIS were newly introduced and the Lung CT task was continued. The Learn2Reg challenge consisted of two phases (mainly organised on grand-challenge.org).

- Phase 1 - Validation Phase: The participants downloaded the training and validation datasets and trained a registration network or tuned hyperparameters on them. The calculated displacement fields on the validation dataset were submitted and evaluated using grand-challenge.org. Challenge participants were allowed to create five submissions per day to this phase.
- Phase 2 - Test phase: Within one week after the test data release, the participants had to send either the generated displacement fields to the organisers or a Docker container containing the algorithm. A Docker submission was preferred and made more attractive by evaluating the runtime of the algorithm.

Members of the organisers' institutes could participate in the challenge having the same data access as any other participant. However, they were not eligible for awards. A continuous evaluation for test data will be possible at grand-challenge.org.¹ All methods that solve at least four of the six tasks are included into the overall ranking of this paper. To remove entry barriers for new participants with expertise in deep learning but not necessarily registration, the organisers provided pre-processed data. A detailed description of the used pre-processing is given in section II-B. Furthermore, the evaluation code for voxel displacement fields as well as an example Docker container submissions were provided. All additional resources can be found at the Learn2Reg repository.²

B. Tasks

Learn2Reg consists of six clinically relevant complementary tasks (datasets). Table I summarises the dataset details, and we discuss them in detail below.

1) *CuRIOUS*: EASY-RESECT [19] is a simplified sub-set of the original RESECT dataset [20], previously used in the MICCAI CuRIOUS challenges [21]. The dataset contains 22 training and 10 testing subjects with low-grade brain gliomas, intended to help to develop MR vs. US registration algorithms to correct tissue shift in brain tumour resection. For the Learn2Reg challenge, we included T1w and T2-FLAIR MR scans, and spatially tracked intra-operative ultrasound volumes. All scans were acquired for routine clinical care of brain tumor resection procedures at St Olavs University Hospital (Trondheim, NO). Matching anatomical landmarks were annotated between T2-FLAIR MR and 3D ultrasound volumes [20] to enable evaluation of the registration accuracy. During pre-processing, for each subject, the T1w scan is rigidly registered to the T2-FLAIR scan, and both scans are resampled to the same coordinate space as the 3D ultrasound volume yielding fixed voxel dimensions for all scans ($256 \times 256 \times 288$) at an isotropic resolution of approximately 0.5 mm. The registration to be carried out for this task was difficult for following reasons. First of all, it is a multimodal registration between MR and US images and the US images are typically noisier than the MR images. Furthermore, the pre-operative MR scans show a larger region of the brain whereas the intra-operative US volume was obtained to cover the entire tumor region after craniotomy but before dura opening.

2) *Hippocampus MR*: This dataset consists of 394 MR scans of the hippocampus region acquired in 90 healthy adults and 105 adults with non-affective psychotic disorder taken from the Psychiatric Genotype/Phenotype Project data repository at Vanderbilt University Medical Center (VUMC). The hippocampus head and tail were manually traced in all scans. The ability to establish correspondences for small structures between patients is particularly important for accurate population analysis. Previous to the Learn2Reg challenge, the dataset was used as part of the Medical Segmentation Decathlon [22]. Due to its small volumetric size and large training dataset with two anatomical labels, Hippocampus MR appeared to

be a good entry-level task for learning-based registration approaches.

3) *Abdomen CT-CT*: This task tackles inter-patient registration of abdominal CT scans, which enables statistical modelling of variations of organs for abnormality detection, and can provide a canonical atlas space for further investigations. The dataset contains 50 abdominal CT scans (30/20 for training/testing) with 13 manually labelled anatomical structures: spleen, right/left kidney, gall bladder, esophagus, liver, stomach, aorta, inferior vena cava, portal and splenic vein, pancreas and left/right adrenal gland. Data acquisition and annotation protocols are detailed in [12]. The images were registered affinely in a groupwise manner and resampled to the same voxel resolution and spatial dimensions ($192 \times 160 \times 256$).

4) *Abdomen MR-CT*: The data was compiled from public studies of the cancer imaging archive (TCIA) [23] that contained paired scans of both MR and CT from the same patients. In particular, 16 MR and CT scans from the following studies, TCGA-KIRC [24], TCGA-KIRP [25], and TCGA-LIHC [26], are included in Learn2Reg - that cover routine diagnostic scans and follow-up imaging for kidney surgery. The data has been resampled to an isotropic resolution of 2mm, and cropped and padded to achieve voxel dimensions of $192 \times 160 \times 192$. We have also manually traced 3D segmentation masks for the liver, spleen, left and right kidney. All scans were pre-aligned using a groupwise affine registration based on the deeds-linear algorithm [27]. Additional unpaired and segmented training data from two further challenges - BCV-CT [12] and CHAOS-MR [28], [29] - were provided for pre-training.

5) *OASIS*: The task employed 416 3D whole-brain MR scans from the Open access series of imaging studies (OASIS) [30], a cross-sectional MR data study with a wide range of participants from young, middle-aged, nondemented, and demented older adults. The clinical relevance of this inter-patient registration task lies in quantitative brain analysis, which is of utmost importance for a better understanding of the human brain and for the analysis of various brain diseases. Standard brain MR pre-processing including skull-stripping (optional), normalisation, pre-alignment, and resampling was performed. Semi-automatic labels with manual corrections of 35 cortical and subcortical brain structures were generated using FreeSurfer [31]. For details on data curation, see [32].

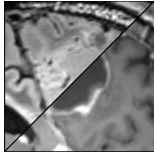
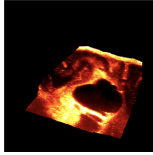







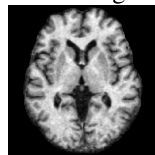
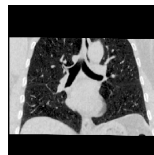

6) *Lung CT*: The aim of the lung CT task was the registration of expiration to inspiration CT scans of the lung. Establishing correspondences between longitudinal lung scans can help to monitor disease progression, estimate motion in radiotherapy planning or enable direct assessment of lung ventilation. The data consists of 20 training [33] and 10 test scan pairs [34]. The scans were acquired at the Dept. of Radiology at the Radboud University Medical Center, Nijmegen, NL. All pairs were affinely pre-registered and resampled to an image size of $192 \times 192 \times 208$. Lung segmentation masks and keypoints were provided as additional training information. The complexity of this registration task is manifold. First, the fields of view of the fixed and moving scan differ largely since the lungs in the expiration scan are not fully visible. Second, the scale of the motion within the lungs can often be larger than the anatomical structures

¹<https://learn2reg.grand-challenge.org>

²<https://github.com/MDL-UzL/L2R>

TABLE I

OVERVIEW OF ALL SIX LEARN2REG TASKS ADDRESSING THE IMMINENT CHALLENGES OF MEDICAL IMAGE REGISTRATION: MULTI-MODAL SCANS ● (TASKS WITH AT LEAST TWO DIFFERENT IMAGE MODALITIES), FEW/NOISY ANNOTATIONS ● (LESS THAN FIVE ANNOTATED ANATOMICAL STRUCTURES FOR TRAINING CASES), PARTIAL VISIBILITY ● (RESTRICTED OR CROPPED FIELD OF VIEW FOR AT LEAST ONE IMAGE OF A REGISTRATION PAIR), SMALL DATASETS ● (LESS THAN 30 TRAINING CASES), LARGE DEFORMATIONS ● (TASKS WITH INITIAL DISPLACEMENTS OF AT LEAST 10CM), SMALL STRUCTURES ● (TASKS CONTAINING CASES WITH TARGET STRUCTURES COMPRISING LESS THAN 100 VOXELS), UNSUPERVISED REGISTRATION ● (NO ANNOTATIONS FOR TRAINING CASES) AND MISSING CORRESPONDENCES ● (E.G. DUE TO REMOVED ORGANS, DIFFERENT FIELD OF VIEWS ETC)

	CuRIOUS		Hippocampus MR		Abdomen CT-CT	
	Fixed	Moving	Fixed	Moving	Fixed	Moving
						
Modalities	MR T1w & FLAIR/US		MR T1w/MR T1w		CT/CT	
Intra-/Inter-patient	Intra-patient		Inter-patient		Inter-patient	
Resolution	256×256×288		64×64×64		192×160×256	
Voxel size	~0.5×0.5×0.5mm		1×1×1mm		2×2×2mm	
Cases (Train/Test)	22/10		263/131		30/20	
Preprocessing	resample		crop/pad/resample		canonical affine pre-align crop/pad/resample	
Annotations (Train/Test)	-/9-18 landmarks/case		2/2 anatomical labels		13/13 anatomical labels	
Additional data						
Metrics	TRE/TRE30 SDlogJ/RT		DSC/DSC30/HD95 SDlogJ/RT		DSC/DSC30/HD95 SDlogJ/RT	
Challenges	● ● ● ● ● ●		● ● ●		● ●	
	Abdomen MR-CT		OASIS		Lung CT	
	Fixed	Moving	Fixed	Moving	Fixed	Moving
						
Modalities	MR T1w / CT		MR T1w / MR T1w		CT / CT	
Intra-/Inter-patient	Intra-patient		Inter-patient		Intra-patient	
Resolution	192×160×192		160×192×224		192×192×208	
Voxel size	2×2×2mm		1×1×1mm		1.75×1.25×1.75mm	
Cases (Train/Test)	8/8		416/39		20/10	
Preprocessing	canonical affine pre-align crop/pad/resample				affine pre-align crop/pad/resample	
Annotations (Train/Test)	4/9 anatomical labels		35/35 anatomical labels		-/100 landmarks/case	
Additional data	90 unpaired MR/CT scans ROI masks				lung masks	
Metrics	DSC/DSC9/HD95 SDlogJ/RT		DSC/DSC30/HD95 SDlogJ/RT		TRE/TRE30 SDlogJ/RT	
Challenges	● ● ● ● ● ●		●		● ● ● ● ● ●	

(vessels and airways) themselves. Therefore, a registration method needs to estimate large displacements that account for substantial breathing motion and also align small structures like individual pulmonary blood vessels precisely. To measure the accuracy manual landmarks are used that are typically located at the boundary or bifurcation of vessels, airways, and parenchyma.

C. Challenge Design

To provide a comprehensive evaluation of the registration performance, we consider a number of complementary metrics (see section II-C.1) that assess the accuracy, robustness, plausibility, and speed of the algorithms. For final task ranks, we further consider the significance of differences in results. The detailed ranking scheme is described in section II-C.2.

1) Metrics:

a) *DSC*: The Dice similarity coefficient (DSC) measures the overlap of two sets of segmentation labels (on the fixed and warped moving scan).

b) *DSC30*: To assess robustness, the DSC30 metric considers the 30th percentile in DSC scores over all anatomical structures and cases.

c) *DSC9*: DSC9 is a special metric introduced for the Abdomen MR-CT task, to assess the effect of label bias. It is evaluated on 9 additional anatomical labels, that were not available during training.

d) *HD95*: The Hausdorff distance (HD) indicates the maximum distance in a metric space (here: Euclidean space, distance specified in millimetres (mm)) between two sets of surfaces (segmentation labels on the fixed and warped moving scan). For a robust score, we consider the 95th percentile instead of the maximum distance (HD95).

e) *TRE*: The target registration error (TRE) is defined as the euclidean distance (in millimetres (mm)) between corresponding landmarks in the warped fixed and moving scan.

f) *TRE30*: Similar to the DSC30 score the TRE30 metric collects the 30th percentile of largest landmark distances.

g) *SDlogJ*: The plausibility (smoothness) of a displacement field is captured using the standard deviation of the logarithm of the Jacobian determinant (SDlogJ) of the displacement field [35], [36]. The Jacobian is calculated by a central differencing approximation.

h) *RT*: In addition, we are able to measure the test-time registration runtime (RT) on the same hardware (CPU: Xeon Silver 4210R, GPU: Quadro RTX 8000), when methods are submitted as a Docker container. Start and stop times are the loading of the first scan and writing of the displacement field to disk, respectively.

2) *Ranking Scheme*: The ranking scheme is based on the ranking scheme of the Medical Decathlon.³ We rank methods using statistically significantly different results. For each metric applied in a task, methods are compared against each other (Wilcoxon signed rank test with $p < 0.05$), ranked based on the number of “won” comparisons and finally mapped to a numerical metric rank score between 0.1 and 1 (with possible score sharing). A task rank score is then obtained as the geometric mean of individual metric rank scores. All methods for which no metric is available (not submitted to the task, no Docker container submitted) share the lowest possible metric rank score of 0.1.

III. CHALLENGE ENTRIES

In 2020, ten teams submitted their solutions. The total number of teams increased to 21 in 2021. Counting the submissions task-wise results in 65 unique challenge entries. Table II provides a summary of important information. Below is a brief description of each of the 21 submissions. For more details, please refer to the respective articles in the proceedings of the MICCAI Learn2Reg workshops.

3Idiots ■: Reference [37] employs deep-learning-based approach using a hybrid similarity loss consisting of intensity (SSD), statistical (MI), and label-based (Dice+L1) penalties. A Voxelmorph [38] model with an increased number of feature channels and halved output resolution is trained in a patch-wise manner and applied to the OASIS task.

Bailiang ■: Reference [39] addresses OASIS and is based on the DeepRegNet framework from Project-MONAI. The input of the encoder is the concatenation of fixed and moving images. A dense vector field (DVF) is predicted from summing over different level decoders and integrated using scaling and squaring. The loss function is composed of LNCC, MIND-SCC, Dice, and a diffusion regulariser. https://github.com/BailiangJ/learn2reg2021_task3

ConvexAdam ■: Reference [40] proposes a decoupling of deep learning for semantic feature extraction and the conventional optimisation. They combine a single-level dense discretised displacement correlation with large capture range and convex global optimisation with a local gradient-based instance refinement using the Adam optimiser. The method is applied to all six tasks and uses diffusion regularisation, an inverse-consistency constraint, and MIND similarity. The method extends the input features to learned label-supervised representations for inter-patient tasks: Abdomen CT-CT, Hippocampus MR, and OASIS. <https://github.com/multimodallearning/ConvexAdam>

corrField ■: A faster implementation (from [41]) of the corrField method [42] is introduced as a non-learning based unsupervised baseline. The method estimates sparse correspondences on image-based Förstner keypoints with exact message passing on a minimum spanning tree. MIND-SSC features are used for the similarity term. <https://grand-challenge.org/algorithms/corrfield/>

Driver ■: Reference [43] uses a dual-encoder UNet backbone with separated multi-scale feature extractors that comprises Deformation Field Integration (DFI) and non-rigid feature fusion (NFF) modules. It produces multi-scale sub-fields that progressively align fixed and moving features. The overall framework comprises a rigid transform network and MI or LNCC similarity, weak label-supervision and regularisation.

Epicure ■: Reference [44] addresses the lung CT task using an iterative registration approach based on the Elastix toolbox [15] optimising the object function that is composed of the NCC similarity and a bending energy penalty term.

Estienne ■: References [45] and [46] combines a diffeomorphic symmetric spatial transformer network with an embedding merging step, that eases the learning by subtracting the embeddings of separately encoded fixed and moving scans and thereby leveraging the prior knowledge that swapped inputs should yield negated velocity fields. They extend the label-based pre-training by including additional public datasets with at least partial overlap in segmentation classes, using segmentation masks produced by a CNN. https://github.com/TheoEst/abdominal_registration

Gunnarsson ■: Reference [47] proposes an end-to-end learning-based 3D registration method inspired by the PWC-Net [48]. The method estimates and refines a displacement field from a cost volume at each level of a CNN

³<http://medicaldecathlon.com>

TABLE II

METHODOLOGICAL OVERVIEW OF ALL LEARN2REG METHODS. AN ENTRY IN THE TABLE INDICATES AGREEMENT WITH THE CORRESPONDING HEADING. UNSUPERVISED AND SUPERVISED CHALLENGE ENTRIES ARE MARKED WITH A ■ AND ● SYMBOL IN THE *Tasks* SUBGROUP. IF A CHALLENGE ENTRY USES DIFFERENT APPROACHES FOR DIFFERENT TASKS OR MIXES THEM WITHIN THE METHOD (E.G. DEEP LEARNING + INSTANCE OPTIMISATION) WE MARKED THE PROPERTY WITH A ○ SYMBOL. ALL BASELINE METHODS ARE MARKED WITH AN *. FOR DETAILED DESCRIPTIONS OF THE METHODS SEE SECTION III AND THE ASSOCIATED REFERENCES

	Tasks					Type	Objectives							Reg.	Optimis.	Misc. (architectures, add. objectives, etc.)						
	CuRIOUS	Hippocampus MR	Abdomen CT-CT	Abdomen MR-CT	OASIS	Lung CT	Conventional	Deep Learning	NCC	MIND-SSC	NGF	MI	Dice	Keypoints	Consistency (Inv./Cycl.)		Diffusion	Curvature	Adam	Convex	L-BFGS/Gauss-Newton	Open source
3Diots	■						●									●						Voxelmorph; SSD;
Bailiang	■						●									●						DeepRegNet;
ConvexAdam	■	●	●	■	■	■	○	○	○	○		○	○		○			○	○			UNet; Dense corr.;
corrField*	■	■	■	■	■	■	●			●					●							Dense corr.;
Driver	■						●					●	○	○								PCNet; Cross-entropy loss;
Epicure	■						●															Bending energy regularisation;
Estienne	■	●	●				●									●						UNet; Multi-Task learning;
Gunnarsson	■	●	●				●									●						PWC;
Imperial	■						●				●			●								Structure-guided loss;
Joutard	■		●				●									●						UNet; EDT similarity; Dense corr.;
LapIRN	■	●	●	●	●	■	○	○	●	○		○	○			●						UNet; Conditional NN;
LaTIM	■			■	■	■	●															Directional representations;
Lifshitz	■						●								●							Unrolled L_1 regulariser; Dense corr.;
IWM	■	●					●															2-stream NN;
MEVIS	■	●	■	■	■	■	○	○		●		○	○		●					○		Groupwise; Bayesian modelling;
Multi-brain	■						●									●					●	Bending and Jacobian regularisation;
NiftyReg*	■	■	■	■	■	■	●		○	○												Dense Corr.;
PDD-Net*	■	■	●				●			●			○									UNet; SSTVD similarity; Dense corr.;
PIMed	■						○	○					○						○			UNet; SAD
Thorley	■						●						●								●	Voxelmorph;
Winter	■						○	○													○	

downsampling pyramid and is supervised by a similarity (NCC) and/or segmentation (DICE) loss, as well as a smoothness penalty. The network is trained and evaluated on scan pairs from the four tasks of the 2020 challenge (CuRIOUS, Lung CT, Abdomen CT-CT and Hippocampus MR). <https://github.com/ngunnar/learning-a-deformable-registration-pyramid>

Imperial ■: Reference [49] uses Image-and-Spatial Transformer Networks (ISTN) as the backbone of their method. In the ISTN, the fixed and moving images are first separately processed by the ITN to generate a segmentation mask and a feature map of the input image. Subsequently, both feature maps are used by the STN to predict the displacement field. The loss function consists of a structural-guided and image similarity and a regularisation loss. <https://github.com/biomed-mira/istn>

Joutard ■: Joutard addresses the Abdomen CT-CT task with a weakly supervised deep learning approach. A CNN extracts features from the fixed and moving image, which are concatenated with their spatial image coordinates. The feature distributions for each spatial location are then matched

between the two images which yield a correspondence matrix from which the average displacement can be derived. The network is supervised by a segmentation (Dice) and a regularisation (L2 norm on gradients) loss.

LapIRN ■: References [50] and [51] propose an image registration method based on Laplacian pyramid registration networks to overcome the large inter-and intra-variations of anatomical structures in the input scans. For the 2021 tasks (Abdomen MR-CT, OASIS and Lung CT), [51] extended their initial approach [50] by adding a conditional module that enables the input of the regularisation hyperparameter so that the different solutions for different hyperparameter values can be captured by a single convolutional neural network. https://github.com/cwmok/Conditional_LapIRN

LaTIM ■: Reference [52] addresses the Abdomen CT-CT tasks using an iterative technique exploiting vector-valued directional image representations. The method is implemented within the Elastix framework.

Lifshitz ■: Reference [53] proposes a deep-learning-based solution for the Lung CT task that comprises a 3D extension of ARFlow [54] with multi-resolution warping,

displacement correlation, and flow estimation. To address edge-preservation of sliding motion an unrolling of the total variation (L1) regularisation loss using variable substitution is proposed.

IWM ■: IWM employs a deep-learning-based registration method for the Hippocampus MR and the OASIS task. For the Hippocampus MR task, they use sequential deformation field composition, while the solution for the OASIS task uses an image pyramid separately applied to both input images and a UNet with residual blocks. The objective function includes MIND, Dice, inverse consistency and diffusion losses.

MEVIS ■: The submission of MEVIS [55], [56] solves all tasks besides the Hippocampus MR task using a conventional method and build on cost functions and losses made up from several terms that are selected for the specific task. The method use a coarse-to-fine multi-level iterative registration scheme where a Gaussian image pyramid is generated for both images to obtain downsampled and smoothed images. On each level, a quasi-Newton L-BFGS optimisation is used. For the Hippocampus task, a deep learning approach with a weakly supervised trained UNet is applied using the same cost function as in the conventional approach.

Multi-brain ■: Reference [57] uses groupwise, fully unsupervised registration techniques based on Bayesian modelling and Gauss-Newton optimisation, which learns priors over image intensities and spatial tissue classes. The method requires no pre-processing of the imaging data and does not utilise label information. The method is applied to Abdomen MR-CT, OASIS, and Lung CT. <https://github.com/WTCN-computational-anatomy-group/mb>

NiftyReg ■: [16] is applied as conventional baseline for all tasks without label supervision using NCC for CuRIOUS and otherwise MIND as similarity metric. Both bending and Jacobian regularisation penalties are applied and the number of pyramid levels is restricted to yield competitive run times. <https://github.com/KCL-BMEIS/niftyreg>

PDD-Net ■: The PDD-Net [58], [59] is used as a baseline method. It uses a deformable convolutional network to extract image features and compute a six-dimensional dissimilarity tensor (three spatial + three displacement dimensions). A smooth displacement field is obtained from the dissimilarities by mean field inference over spatial dimensions and approximated min-convolutions over displacement dimensions. The method is adapted to four challenge tasks (CuRIOUS, Hippocampus MR, Abdomen CT-CT, and Lung CT). https://github.com/multimodallelearning/pdd_net

PIMed ■: PIMed uses a multi-slice segmentation network that yields anatomical maps and is employed for Abdomen MR-CT and Abdomen CT-CT in conjunction with a NCC loss and optimised using 1) a translation only and 2) a diffeomorphic deformation model. They adapt a residual VoxelMorph model with weak supervision for OASIS. For lung CT, they apply a conventional method with geodesic density regression and adaptation of intensities to lung tissue density [60].

Thorley ■: The submission from the University of Birmingham (UoB) team tackled the OASIS task using an iterative coarse-to-fine registration scheme, optimizing the classical SAD difference term and a third-order diffusion displacement

regularizer. Additionally, they decomposed the transformation into the composition of a series of small non-stationary velocity fields, and solved the convex optimization using the Nesterov accelerated ADMM [61] with closed-form solutions. An additional post processing step using a UNet supervised with dice and diffusion loss was used to further refine the displacement fields produced by the iterative optimization.

Winter ■: Winter addresses the Abdomen MR-CT, OASIS and Lung CT task by employing a conventional method for Lung CT and a attention-based deep-learning-based registration method for Abdomen MR-CT and OASIS brain. For the Abdomen MR-CT task, a two-step approach is applied that first aligns the provided ROI masks. <https://github.com/WinterPan2017/ADLReg>

IV. ADDITIONAL EXPERIMENTS

A. Label Bias

Previous publications on learning-based registration have already discussed the possibility of bias towards anatomies that are used both for training and evaluation [38]. While this bias is intrinsic to all segmentation approaches, registration is often used as a more generalistic tool in clinical applications that may require accurate alignment of structures that are not defined a priori. To study the effect of adding anatomical labels to the evaluation that were not present during method development and training, we extended both abdomen tasks. For the inter-patient CT-CT registration we included the duodenum with the manual annotations provided by [62], for the intra-patient MR-CT task we extended the predominantly large organs by five smaller ones: gallbladder, stomach, aorta, portal vein, pancreas (semi-automatically generated using a nnUNet trained on the VISCERAL gold corpus [63]).

B. Unsupervised Registration

The top-performing methods are all modular in their use of segmentation labels for supervision. As analysed in the label bias experiment, there is a risk of over-fitting registration performance to the chosen subset of manually annotated anatomies. We, hence, compared the unsupervised counterparts of the following methods: LapIRN and ConvexAdam. ConvexAdam already uses an unsupervised method for all three intra-patient tasks, and LapIRN for CuRIOUS and Lung CT. Therefore the additional comparisons are restricted to the abdomen and brain.

C. Transferability

A robust registration method should work well for all scan pairs regardless of acquisition parameters and thus on comparable datasets. A limitation of deep-learning-based methods might be that they reach higher accuracy on the dataset they are trained on and show a considerable loss of accuracy on other data. As in [64] and [65], we evaluate the transferability of methods submitted to the lung CT-CT task by registering the DIRLab 4DCT [13] scan pairs. The scans are preprocessed in the same way as the scans of the lung CT-CT task. The evaluation is based on the target registration error of

TABLE III
CuRIOUS

	TRE↓	TRE30↓	SDLogJ↓	RT↓	Rank↑
Initial	6.38	12.00			
corrField	2.84	5.29	0.00	2.70	0.85±0.03
PDD-Net	3.08	6.28	0.00	8.21	0.83±0.03
ConvexAdam	3.31	5.82	0.00	1.33	0.77±0.04
NiftyReg	4.09	7.85	0.00	23.1	0.56±0.06
LapIRN	5.67	11.1	0.00	34.8	0.49±0.06
MEVIS	6.55	10.4	0.00	57.8	0.42±0.04
Gunnarsson	7.1	10.1	0.14	42.2	0.19±0.01

TABLE IV
HIPPOCAMPUS MR

	DSC↑	DSC30↑	HD95↓	SDLogJ↓	RT↓	Rank↑
Initial	0.55	0.36	3.91			
LapIRN	0.88	0.86	1.30	0.05	1.03	0.93±0.01
MEVIS	0.85	0.84	1.55	0.05	0.59	0.78±0.03
ConvexAdam	0.84	0.83	1.85	0.07	0.48	0.75±0.04
IWM	0.79	0.76	2.20	0.08	0.80	0.63±0.04
Estienne	0.85	0.84	1.51	0.09	1.46	0.62±0.04
PDD-Net	0.78	0.76	2.23	0.07	0.35	0.58±0.04
NiftyReg	0.76	0.72	2.72	0.09	4.75	0.37±0.03
corrField	0.72	0.68	2.89	0.05	1.20	0.34±0.02
Gunnarsson	0.74	0.67	2.82	0.16	22.0	0.25±0.01

TABLE V
ABDOMEN CT-CT

	DSC↑	DSC30↑	HD95↓	SDLogJ↓	RT↓	Rank↑
Initial	0.28	0.04	21.78			
ConvexAdam	0.69	0.45	11.03	0.06	2.75	0.94±0.01
LapIRN	0.67	0.47	12.51	0.12	3.80	0.82±0.03
Estienne	0.69	0.51	11.77	0.18	8.23	0.67±0.08
MEVIS	0.51	0.24	18.21	0.14	3.49	0.60±0.04
corrField	0.49	0.24	17.22	0.28	5.40	0.53±0.04
PIMed	0.49	0.23	15.75	0.05		0.49±0.04
PDD-Net	0.49	0.24	17.75	0.41	6.06	0.44±0.02
Joutard	0.40	0.13	17.25	0.05	3.67	0.42±0.01
NiftyReg	0.45	0.20	20.70	0.36	17.1	0.36±0.02
Gunnarsson	0.43	0.17	18.55	0.13	31.5	0.33±0.02

the landmarks and the smoothness of the deformation field. Furthermore, this experiment allows comparison to a variety of other lung registration methods, as the DIRLAB data set is often used as a benchmark (please note that the reduced resolution leads to a general deterioration of TRE of ~ 0.3 mm).

V. RESULTS

A. Challenge Outcome

In this section, we will first present each task separately and subsequently the eight methods that are included in the overall ranking. Tables III to VIII give the numerical results and the

TABLE VI
ABDOMEN MR-CT

	DSC↑	DSC9↑	HD95↓	SDLogJ↓	RT↓	Rank↑
Initial	0.33	0.22	48.65			
ConvexAdam	0.75	0.73	24.92	0.09	1.30	0.82±0.01
corrField	0.76	0.73	23.35	0.10	2.13	0.81±0.02
LapIRN	0.76	0.69	22.81	0.12	1.50	0.77±0.03
PIMed	0.78	0.68	21.99	0.07	59.2	0.75±0.02
MEVIS	0.71	0.65	27.94	0.15	14.7	0.67±0.02
Driver	0.76	0.55	27.02	0.13	1.95	0.63±0.03
NiftyReg	0.65	0.55	33.09	0.12	11.0	0.55±0.02
LaTIM	0.54	0.49	41.17	0.13		0.39±0.03
Winter	0.55	0.41	35.51	0.85	2.79	0.31±0.03
Imperial	0.51	0.41	48.60	0.11	278	0.30±0.02
Multi-brain	0.54	0.44	38.21	0.48		0.30±0.02

TABLE VII
OASIS

	DSC↑	DSC30↑	HD95↓	SDLogJ↓	RT↓	Rank↑
Initial	0.56	0.27	3.86			
LapIRN	0.82	0.66	1.67	0.07	1.21	0.92±0.01
ConvexAdam	0.81	0.64	1.63	0.07	3.10	0.82±0.01
IWM	0.79	0.61	1.84	0.05	2.55	0.79±0.02
Driver	0.80	0.62	1.77	0.08	2.02	0.75±0.02
PIMed	0.78	0.58	1.86	0.06	3.47	0.71±0.02
3Idiots	0.80	0.63	1.82	0.08	1.46	0.70±0.02
Winter	0.77	0.57	2.16	0.08	2.56	0.55±0.02
MEVIS	0.77	0.57	2.09	0.07	10.4	0.51±0.02
Multi-brain	0.78	0.59	1.92	0.57		0.38±0.02
corrField	0.74	0.51	2.36	0.08	5.14	0.37±0.02
Thorley	0.77	0.60	2.21	0.31		0.37±0.02
NiftyReg	0.73	0.51	2.37	0.06	5.00	0.36±0.01
Bailiang	0.67	0.42	2.74	0.04	1.38	0.33±0.00
LaTIM	0.74	0.52	2.31	0.08		0.32±0.01
Imperial	0.76	0.57	2.43	0.19	2610	0.29±0.01

scores for each algorithm for each task averaged over the anatomical structures/landmarks and number of scan pairs that were registered for that task. The algorithms are listed in order of their final placement per task. Standard deviations of final rank scores are calculated using jackknife resampling [66]. Fig. 1 shows boxplots illustrating the distribution of the accuracy (TRE and Dice) of the different methods for each task. Furthermore, for selected task (Abdomen MR-CT, OASIS, and Lung CT), a bubble chart combines the accuracy, smoothness, and runtime metric.

1) *CuRIOUS*: Four methods were submitted to this task in addition to the three baseline methods. For two of these methods, some cases caused negative outliers and the average TRE was worse than the initial TRE (c.f. Table III). Only the registration of the two baseline methods corrField and PDD-Net as well as the ConvexAdam method led to a considerable reduction in TRE from 6.38 mm to 2.84 mm, 3.08 mm, and 3.31 mm, respectively.

2) *Hippocampus MR*: In this task, all algorithms consistently performed very well (median Dice > 0.7). Nevertheless, there

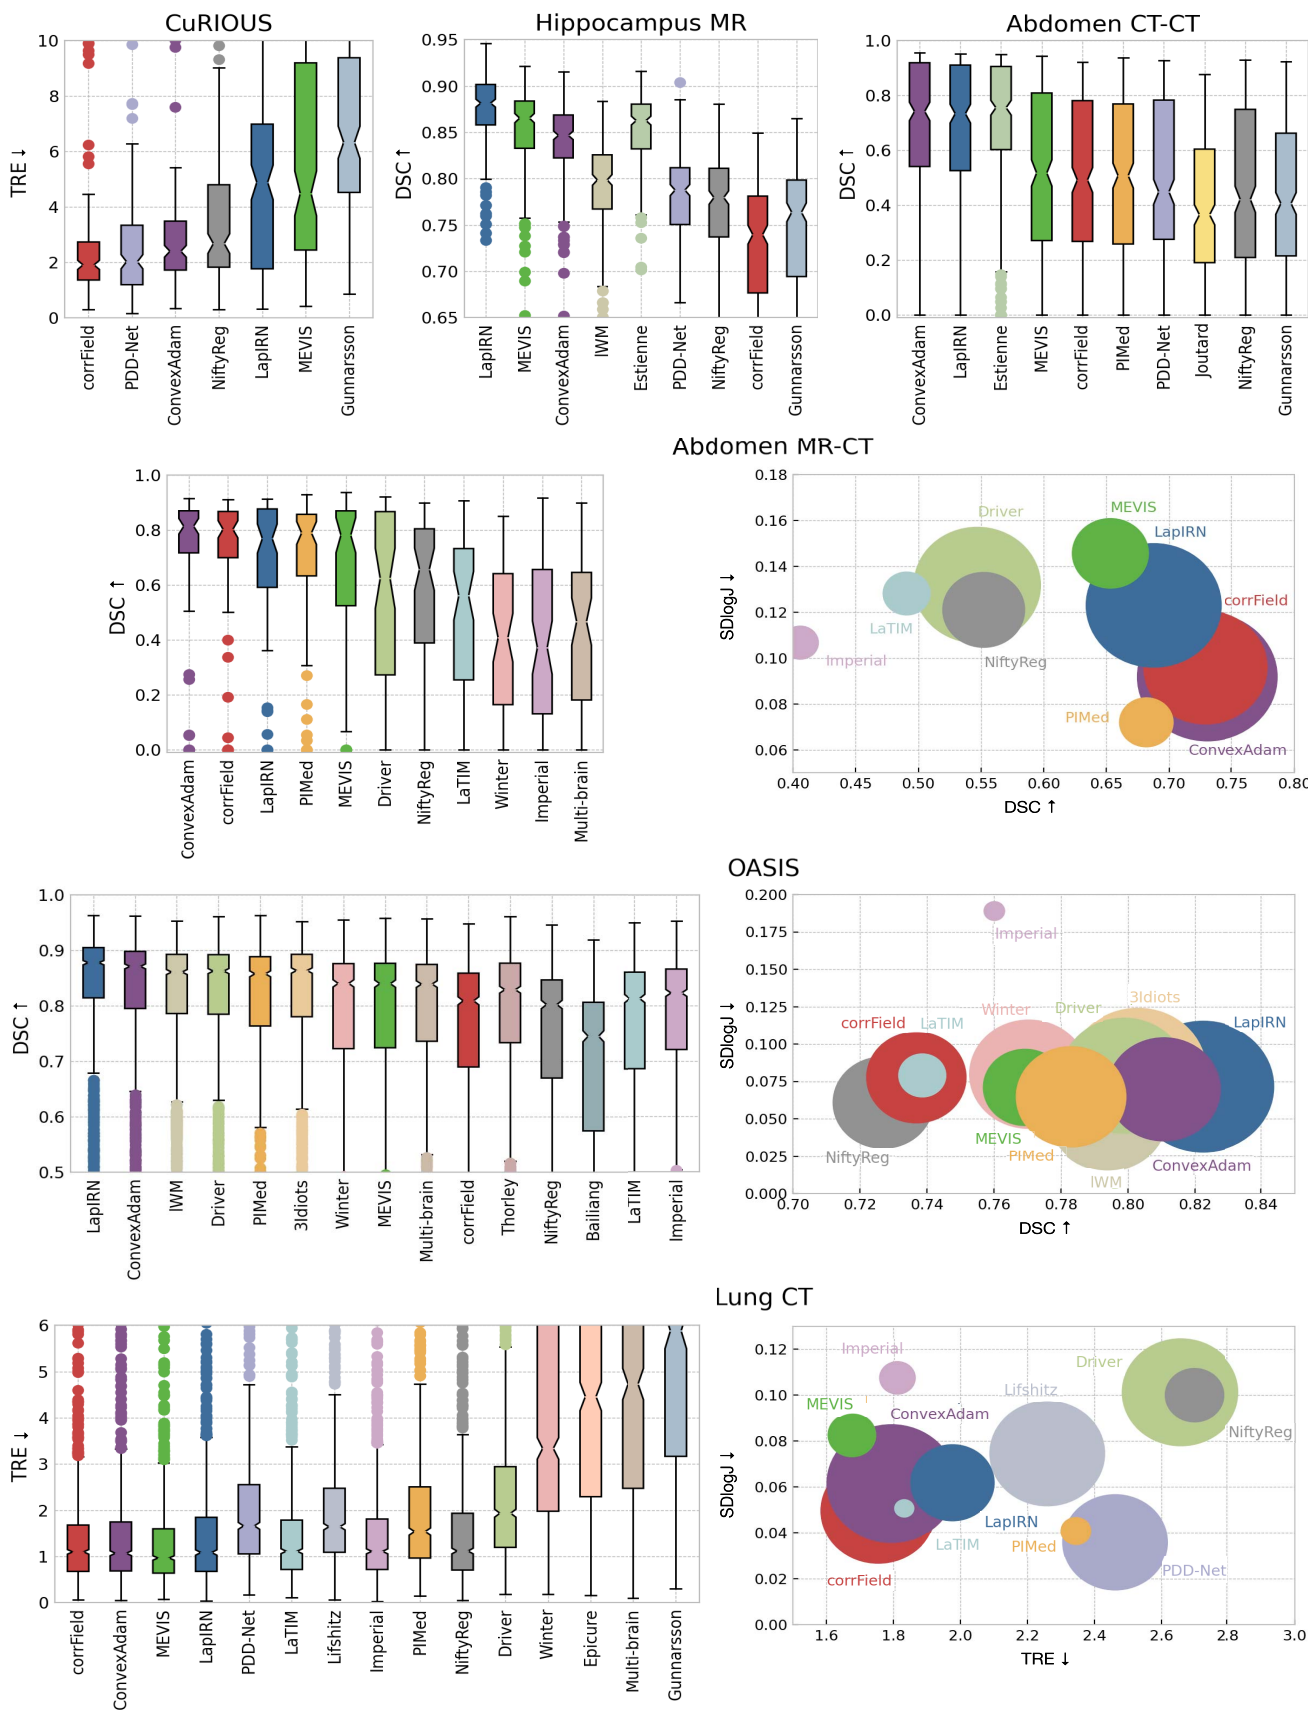


Fig. 1. Boxplots and (selected) bubble charts visualising the results for the six challenge tasks. While the boxplots show the main accuracy metric (DSC and TRE, respectively), the bubble charts combine the accuracy, smoothness and runtime metric (a larger bubble means a faster runtime). Arrows (↑, ↓) indicate the favourable direction of metrics. Comparison methods are color coded: ConvexAdam (purple), LapIRN (blue), MEVIS (green), corrField (red), NiftyReg (grey), PDD-Net (light blue), PIMed (orange), Gunnarsson (light grey), IWM (olive), Estienne (light green), Joutard (yellow), Driver (light green), LaTIM (teal), Winter (pink), Imperial (purple), Multi-brain (brown), 3ldiots (orange), Thorley (brown), Bailliang (teal), Epicure (orange), and Lifshitz (grey). Methods are sorted according to final rank scores.

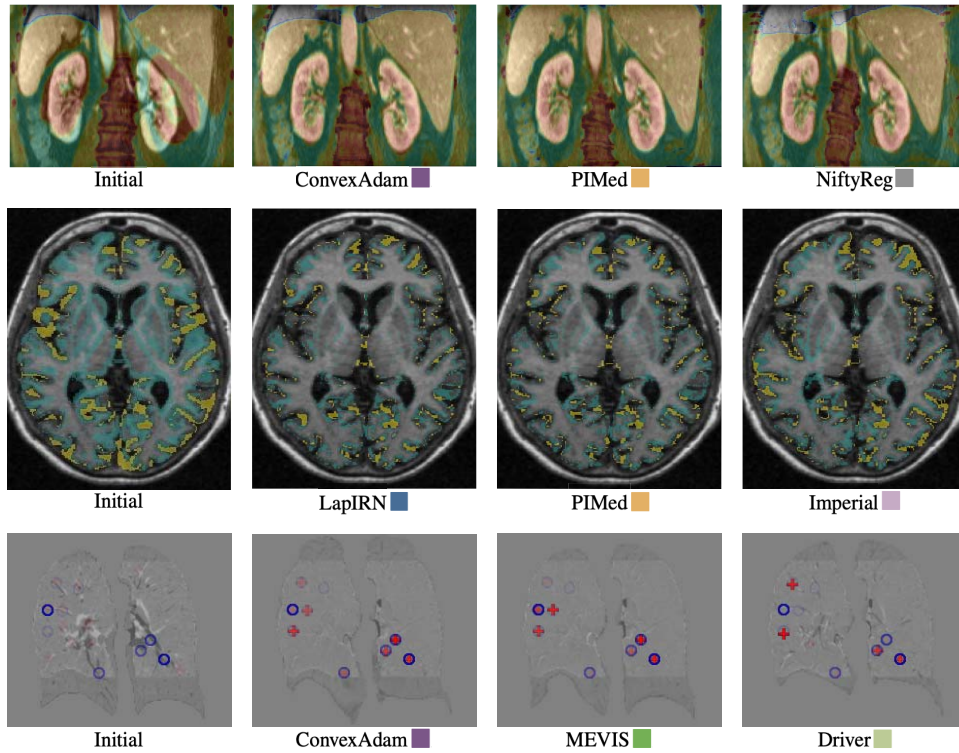


Fig. 2. Exemplary qualitative results for selected methods and tasks. Top row: Overlay of coronal abdominal MR (gray) and warped CT (color) slices. Middle row: False-negative (green) and false-positive (yellow) voxels of propagated segmentation labels on transversal slices of the OASIS dataset. Bottom row: Coronal slices of difference images between exhale and warped inhale lung CT scans (including exhale (blue circle) and warped inhale (red cross) landmarks).

TABLE VIII
LUNG CT

	TRE↓	TRE30↓	SDLogJ↓	RT↓	Rank↑
Initial	10.24	16.80			
corrField	1.75	2.48	0.05	2.91	0.87±0.01
ConvexAdam	1.79	2.70	0.06	1.82	0.81±0.01
MEVIS	1.68	2.37	0.08	95.4	0.78±0.01
LapIRN	1.98	2.95	0.06	10.3	0.73±0.02
PDD-Net	2.46	3.81	0.04	4.22	0.62±0.02
LaTIM	1.83	2.50	0.05		0.62±0.01
Lifshitz	2.26	3.01	0.07	2.90	0.61±0.02
Imperial	1.81	2.54	0.11	300	0.57±0.01
PIMed	2.34	3.27	0.04	623	0.55±0.02
NiftyReg	2.70	5.28	0.10	42.2	0.51±0.02
Driver	2.66	3.50	0.10	2.66	0.44±0.02
Winter	7.41	10.11	0.09	12.0	0.40±0.02
Epicure	6.55	10.29	0.07		0.29±0.02
Multi-brain	6.61	8.75	0.08		0.27±0.01
Gunnarsson	9.00	11.27	0.12	30.9	0.21±0.00

is a performance gap between algorithms using label supervision (LapIRN, MEVIS, ConvexAdam, and Estienne) and unsupervised methods (NiftyReg, PDD-Net and corrField). However, despite label-supervision, the methods of IWM and Gunnarsson perform comparably to unsupervised methods. This is the only task that enabled sub-second runtimes.

3) Abdomen CT-CT: In this task, a clear three-way partition of the algorithms appears. The methods of Estienne, LapIRN,

and ConvexAdam achieved a Dice Score of 0.67-0.69 across the eight individual organs and thus at least a 0.2 higher Dice Score than all other participants. The midfield includes the unsupervised methods MEVIS, corrField, and PDD-Net and the supervised method PIMed which achieve a Dice Score of 0.49-0.51. The final group is formed by the methods Joutard, NiftyReg and Gunnarsson with a Dice Score of 0.40-0.45. This structure can also be found in the other accuracy measures DSC30 and HD95. All methods, apart from NiftyReg and Gunnarsson, have a runtime of fewer than 10 seconds.

4) Abdomen MR-CT: In the abdominal MR-CT task, the algorithms can also be divided into three groups based on the median Dice Score (c.f. Fig. 1). The leading group can be further divided into the algorithms that achieve a similar Dice Score on the segmentations provided in the training as on the nine unknown organ segmentations (ConvexAdam and corrField) and those that show a performance loss on the nine unknown organs (LapIRN, PIMed, MEVIS). This division is also reflected in the variance of the achieved Dice Scores. In respect of runtime, PIMed stands out in this task with a runtime of approximately one minute. In Fig. 2, exemplary qualitative registration results are shown.

5) OASIS: The OASIS inter-subject brain task attracted the most learning-based solutions. The results are summarised in Table VII and visualised in Fig. 1 showing that most of these methods achieve very similar results in terms of Dice Score for the cases with the highest scores (Dice of 80-90%). The differences are primarily in the more difficult cases and thus in the DSC30 score, where the LapIRN,

TABLE IX
OVERALL RANK SCORES OF METHODS SUBMITTED TO FOUR OR MORE TASKS

	CuRIOUS	Hippocampus MR	Abdomen CT-CT	Abdomen MR-CT	OASIS	Lung CT	Overall	Intra-Patient	Inter-Patient
ConvexAdam	0.77	0.75	0.94	0.82	0.82	0.81	0.82	0.80	0.83
LapIRN	0.49	0.93	0.82	0.77	0.92	0.73	0.76	0.65	0.89
MEVIS	0.42	0.78	0.60	0.67	0.51	0.78	0.61	0.61	0.62
corrField	0.85	0.34	0.53	0.81	0.37	0.87	0.59	0.84	0.41
NiftyReg	0.56	0.37	0.36	0.55	0.36	0.51	0.44	0.54	0.36
PIMed			0.49	0.75	0.71	0.55	0.35	0.39	0.33
PDD-Net	0.83	0.58	0.44			0.62	0.34	0.37	0.32
Gunnarsson	0.19	0.25	0.33			0.21	0.19	0.16	0.22

ConvexAdam, and the methods of Driver and 3Idiots methods perform slightly better than for example PIMed and Winter. The conventional methods of MEVIS and corrField achieve mid-ranked accuracies and have a higher runtime. Fig. 2 shows an example transversal slice of the fixed image overlaid with the false-negative segmented voxels (green) and false-positive segmented voxels (yellow) for initial moving segmentation and the propagated segmentations by the methods of Imperial, PIMed, and LapIRN. All methods were able to align the small structures of the brain with only very small visible differences.

6) *Lung CT*: This task was carried out in both years because in 2020 only the MEVIS, which uses automatically computed keypoints as additional metric, achieved a TRE of less than 2mm (1.72mm), while other teams performed considerably worse (e.g. LapIRN 3.24mm and PDD-Net 2.46mm). In 2021, keypoint correspondences were provided for training and the submissions improved, with six teams (corrField, ConvexAdam, MEVIS, LapIRN, LaTIM, Liftschitz) achieving a TRE of less than 2mm. Compared to the other tasks, the runtime in the lung CT task is considerably longer for several algorithms due to the additional time needed to compute keypoints or perform instance optimisation. Fig. 2 visualises the difference images of an example coronal slices for the methods of Driver, ConvexAdam, and MEVIS overlaid with manual landmarks.

7) *Overall Ranking*: Table IX gives the overall rank scores of the eight methods submitted to four or more tasks. Additionally, we separately listed the scores for inter- and intra-patient registration tasks. ConvexAdam was among the top three on each task (winning Abdomen CT-CT and Abdomen MR-CT) and ranked first overall. The GPU-acceleration brings down computation cost of this optimisation-based method to a few seconds for 3D registration and that is why it consistently achieves high scores for the run time in addition to the very good quality scores. LapIRN reached the overall second rank and yielded the best result for Hippocampus MR and OASIS. This demonstrates that a well-designed convolutional feed-forward network (instance optimisation was used only for CuRIOUS and Lung CT) can outperform conventional approaches in particular for inter-patient tasks. MEVIS achieved the third place overall, with top ranks in particular for Lung CT and Hippocampus MR based on a combination of NGF metric, curvature regularisation, and L-BFGS optimisation with additional learning components only employed for the brain task. CorrField uses no label supervision at all, but

relies on highly optimised graph-based registration, and comes fourth overall winning two individual tasks: CuRIOUS and LungCT. It is the best method for intra-patient registration. PIMed’s method achieves strong performance on Abdomen MR-CT and OASIS and generalises well to Abdomen CT-CT.

B. Additional Experiments

1) *Label Bias and Unsupervised Registration*: When evaluating the influence of supervision with anatomical labels, we found a clear distinction between intra- (Abdomen MR-CT) and inter-patient registration (Abdomen CT-CT, Hippocampus and OASIS), see Table IX. The former shows nearly no advantage of including such information and it is therefore possible to avoid a risk of overfitting towards certain anatomies. The latter, however, shows a clear deterioration in accuracy when excluding structures from training that are used for evaluation. CorrField (unsupervised) achieves the highest scores for intra-patient registration trails nearly all learning-based methods on the remaining inter-patient tasks. LapIRN trained without Dice loss (i.e. without anatomical knowledge) improves upon those results and achieves very strong results for OASIS and Abdomen CT-CT. This indicates that a large training database and an advanced deep learning architecture may narrow the gap between supervised and unsupervised approaches. We evaluated ConvexAdam for Abdomen CT-CT in three settings, each time evaluating on 8 test labels: 1) all 13 labels in training (DSC=69%), 2) 4 labels in training (DSC=55%) and 3) no labels in training (DSC=45%). This shows that partial supervision clearly leads to improvement of those identical anatomies but can also help to align nearby structures: aligning the esophagus which was excluded as label from training improved by 16% points (likely through the guidance of liver and aorta) and pancreas overlap was increased by 12% points (possibly by including portal vein and adrenal gland). As mentioned in Sec. V-A training on 4 and evaluating on 9 abdominal organs for MR-CT fusion results in a moderate performance gap between supervised and unsupervised methods.

2) *Transferability*: We were able to show that the three best methods of the lung registration task also perform very well on the DIRLab dataset (MEVIS 1.22 mm, ConvexAdam 1.31 mm, and corrField 1.34 mm) without further hyperparameter adaptations. Since the inspiration and expiration images of the DIRLab dataset are extracted from a 4DCT dataset with

shallow breathing, the registration task is probably easier than the Learn2Reg lung CT task. This might explain the lower TRE values on the DIRLab dataset compared to the Learn2Reg lung task (improved TRE of 0.46 mm, 0.48 mm, and 0.41 mm for MEVIS, ConvexAdam, and corrField, respectively). Due to the preprocessing and the reduced resolutions, the Learn2Reg methods achieve slightly worse results than state-of-the-art methods evaluated on the DIRLab dataset. For example, the method of MEVIS as part of their complete registration pipeline and applied to the original images reaches a TRE of 0.94 mm [67]. LapIRN achieves similar results on both datasets (Learn2Reg lung CT 1.98 mm and DIRLab 1.98 mm) showing that the best deep-learning-based methods can also be successfully applied to other datasets without retraining.

VI. DISCUSSION

A. Reducing Entry Barriers

By pre-processing each dataset to the same dimensions and isotropic resolution and providing anatomical annotations for training data wide participation was achieved from research groups across the world. The OASIS inter-subject brain task attracted the most learning-based solutions, which highlights the importance of large, labelled training datasets for deep-learning registration and mirrors the focus of recent research. Lung CT intra-patient registration was addressed by the same number but more diverse set of methods, including conventional, fully deep-learning-based, and hybrid approaches. Some aspects of medical image registration, including affine or rigid pre-alignment, dealing with differences in field-of-view of voxel resolutions, and the processing of very high-resolution scans have been omitted due to our challenge design and could be addressed in future.

B. Task Specific Results

In general, it is difficult to find the exact reasons why one or the other method performed better or worse in the various tasks. Nevertheless, there are some relevant patterns that can be identified. In the CuRIOUS task, the three methods using a dense discretised displacement correlation (ConvexAdam, corrField and PDD-Net) cope best with the difference in the field of view of the input images. In the case of the Hippocampus MR task, the learning-based methods perform considerably better. This can be explained by the fact that the structures used for the evaluation were also available in the training data set, so that the learning-based methods were specialised in the alignment of these structures during the training. A similar result was observed on the OASIS and Abdomen CT-CT task. The OASIS dataset has already been used in the past in various training-test splits by several groups to develop and test the registration algorithms, so that consistently good results were to be expected and which became true for both deep-learning based and conventional methods. On the Abdomen CT-CT dataset, it is difficult to explain the large performance difference of a nearly 0.2 higher Dice. A successful strategy for inter patient registration can be identified in the ConvexAdam method. Instead of using the segmentations directly in the training of a registration network,

a segmentation network is trained. This is used to first generate the segmentations on new data and then to utilise them in the cost function of the optimisation-based registration. In the Abdomen MR-CT task, we found that using a Dice loss for certain structures can lead to overfitting on these structures and therefore the registration network might not registering other structures as well. Furthermore, it has been shown that a multimodal distance metric, as used by most participants, is essential. Successful strategies for lung registration seem to be the use of keypoints and the combination of deep learning registration + instance optimisation. Gunnarsson's learning-based method performs worse in comparison, this is most likely due to the fact that a common network was trained for the Lung CT, Abdomen CT-CT and Hippocampus MR tasks showing that task-specific solutions might be beneficial. Nevertheless, this result shows that a registration network is capable of solving very different tasks at the same time.

C. Comparison of Learning- vs Optimisation-based Registration

We argue that Learn2Reg has helped to demystify common beliefs of fundamental differences between learning- and optimisation registration. First and foremost, there is virtually no difference in computational speed. GPU-acceleration brings down computation cost of optimisation-based methods to a few seconds for 3D registration, i.e. the extraction of features using CNNs often outweighs optimisation times. Furthermore, we see a clear trend that learning on segmentation labels is primarily beneficial for inter-subject registration. For Abdomen CT-CT for instance large improvements of 20%points in Dice overlap compared to previous work [12] could be achieved using Dice losses. All three highest ranked approaches employ a combination of DL and optimisation: LapIRN primarily uses a deep network, but add instance optimisation for Lung CT, MEVIS mainly use conventional optimisation but a DL network for Hippocampus MR, and ConvexAdam combines discrete optimisation with UNet-based semantic features for inter-patient tasks. Our current challenge design did not consider any computational constraints (GPU memory, runtime on CPU), which might limit the practical impact for some applications and should be considered in future studies.

D. Algorithmic Design Choices

There are no direct ablation studies possible for the used architectures and loss functions since each method differs in multiple aspects (see Table II), but some general trends are visible nonetheless. Most approaches use a combination of contrast-invariant intensity metrics (LNCC, NGF and MIND) as well as a Dice loss for tasks where anatomical labels are available. To address larger motion (all tasks expect brain) DL registration methods employ multi-scale (and residual) architectures, multiple warps or often dense correlation layers. Two-stream approaches that process both input scans independently are commonplace to deal with multimodality or contrast variations.

E. Comparison to Baselines

We evaluated two conventional methods, NiftyReg [16] and corrField [42] (using the GPU implementation of [41]), and two learning-based approaches, PDD-Net [58] and the original VoxelMorph [38] as baselines. The latter two were only applied to a subset of tasks. NiftyReg achieves reasonable accuracies but falls behind supervised methods on inter-patient tasks. The original VoxelMorph variant reaches an average Dice overlap of $76.88\% \pm 2.17\%$ for OASIS (7th-10th place based on DSC alone) and a TRE of 7.51 ± 3.43 mm for lung CT (13th place). When trained on a large additional lung dataset [64] a TRE of 1.71 ± 2.86 mm was achieved for the additional DIRLAB lung experiment for which the best performing methods in this challenge achieved 1.3 mm. PDD-Net achieved a second rank for CuRIOUS and fifth place for Lung CT. CorrField achieved the best scores overall for CuRIOUS and LungCT and second place for Abdomen MR-CT, making it stand out as the best performing intra-patient approach (without supervision). This demonstrates that conventional methods are still very competitive for datasets without strong label supervision.

F. Plausibility of Transformations

We analysed the smoothness of transformations with respect to the log-standard deviation of Jacobian determinants for all experiments. While this measure is far from perfect, it enabled a ranking of different solutions to the inherently ambiguous nonlinear registration task that may achieve similar accuracy with large differences in complexity (the common assumption being: the smoother transform is then preferable). As visualised in Fig. 1 there is a tendency that more accurate solutions are also smoother, which indicates that enforcing regularity is an effective means of avoiding overfitting and improves robustness. Some notable exceptions can be found for lung CT, where Imperial appears to suffer from too low regularisation while PDD-Net and PIMed may have reduced accuracy in exchange for overly smooth fields. A potential explanation for the positive correlation of smoothness and accuracy could be the hypothesis that accurate methods are able to establish strong (correct) correspondences at relevant anatomies and extrapolate as smooth as possible in uncertain areas. That means putting emphasis on either surfaces (e.g. based on segmentation estimates) or geometric keypoints (for lung scans) can be beneficial.

G. Limitations of the Challenge Design

We have identified a number of limitations that should be addressed in future studies. First, for computational reasons the training of algorithms was performed offline by participants. This could introduce a bias when additional data is used by certain teams that is not accessible to others and prevents the use of larger datasets that cannot be made public due to privacy concerns. Enabling docker-based training or fine-tuning of models directly at grand-challenge.org would be desirable. Second, the amount of available annotated training data varied across tasks and made in particular intra-patient tasks harder for learning-based approaches. Unfortunately, the problem is

that large datasets are often not publicly available and therefore cannot be used in this type of challenge. Decoupling anatomical feature learning from patient-wise optimisation could be a next step, e.g. by providing training data for airway and fissure segmentation for lung CT. The registration accuracy cannot be measured directly but must be evaluated via auxiliary metrics such as the overlap of segmentation masks which disregards the plausibility of correspondences along the surface or within the structure. While this is an inherent problem in evaluating image registration, this issue can be mitigated by generating further manual annotations for certain structures. The provision of all segmentation classes for training that were used for testing is in our opinion the most problematic limitation of this challenge. This was due to the fact, that for 3 out of 4 tasks with segmentation labels these annotations were already publicly available prior to Learn2Reg and we considered it in-transparent to simply not point participants to their availability. We aimed to mitigate the influence of overfitting towards labelled anatomies by performing additional experiments for partial supervision. And finally, statements about the quality of the registration algorithms can only be generalised to a limited extent, but apply mainly to the selected tasks.

H. Impact and Clinical Adoption

With regard to the five-year-old survey on medical image registration by [3], we can reflect that the shift from surface- to intensity-based registration has somewhat been reverted with a majority of approaches employing segmentation-based overlap or keypoints as driving force. The establishing of learning-based strategies, including hybrid approaches that decouple semantic feature extraction from optimisation or combine feed-forward networks with instance optimisation, can be seen as an important new trend. To assess the likelihood of adopting registration in clinical practice, we are encouraged to see that a number of previous obstacles have been successfully addressed by the participants. First, robustness against variations in scanner protocol and patient characteristics was shown to be very high for top-ranking methods that tackled both multi-centric MR studies (OASIS) as well as the transferability issue for lung CT. Second, run times have been considerably reduced to a few seconds, which will enable clinicians to interact with algorithmic solutions by adjusting hyper-parameters, e.g. the strength of regularisation in near realtime (this holds only true for DL-based methods if they are either decoupled or trained with conditioning cf. [51]). Third, it became clear that highly nonrigid transformations are as well solved as rigid alignment, opening up the promise for clinical applications in image-guided surgery/radiotherapy. In fact, it appears as if pre-alignment remains an active problem in particular for DL solutions.

VII. CONCLUSION

The Learn2Reg challenge was the first to evaluate a wide-range of methods for various inter- and intra-patient as well as mono- and multimodal medical image registration tasks. The main goal was to provide a standardised benchmark on complementary tasks with clinical impact and a platform for comparison of conventional and learning-based medical image

registration methods. We established a low entry barrier for training and validation of 3D registration, which helped us compile results of over 65 individual method submissions from more than 20 unique teams. Although registration is highly dependent on the task, two methods (ConvexAdam and LapIRN) and a baseline method (corrField) were shown to work robustly on all tasks with only minor adjustments to the hyperparameters. The submission of MEVIS also works robustly for all tasks. It should be noted, however, that they use a deep-learning-based method for the hippocampal tasks. Furthermore, several teams (Estienne, PIMed, Driver, 3idiots, Multi-brain LaTIM, Lifshitz and Imperial) have submitted tailored solutions to individual tasks and achieve very good results with it. Our additional *Transferability experiment* (c.f. section V-B) gives a tentative indication that the conventional methods ConvexAdam, MEVIS, and corrField can be directly applied to new data sets without much loss of accuracy. Furthermore, we demystified the common belief that conventional registration methods have to be much slower than deep-learning-based methods. Nevertheless, with LapIRN a deep-learning-based registration method achieves state-of-the-art registration results within seconds. We could not identify any architecture that was advantageous over others. In our experiments, it was found that for deep-learning-based methods using a Dice loss for inter-patient registration is particularly useful and instance optimisation helped increasing the accuracy for intra-patient registration. The results presented in this paper initially apply to the submitted methods on the six data sets used in this challenge. However, they may provide a reference for further research on additional data sets. With the Learn2Reg challenge, we have created a dataset for comparing future registration papers. Furthermore, the dataset has the potential to allow the development of dataset-independent and self-configuring registration methods.

ACKNOWLEDGMENT

The authors would like to thank Yipeng Hu and Tom Vercauteren who co-organised the first Learn2Reg tutorial at MICCAI 2019 which initiated the challenge. The results published here are in part based upon data generated by the TCGA Research Network: <https://www.cancer.gov/tcga>.

Alessa Hering and Bram van Ginneken are with the Department of Radiology and Nuclear Medicine, Radboud University Medical Center, 6525 GA, The Netherlands, and also with Fraunhofer MEVIS, Institute for Digital Medicine, 28359 Bremen, Germany (e-mail: alessa.hering@mevis.fraunhofer.de; Bram.vanGinneken@radboudumc.nl).

Lasse Hansen, Hanna Siebert, Christoph Großbrömer, and Matthias P. Heinrich are with the Institute of Medical Informatics, Universität zu Lübeck, 23562 Lübeck, Germany (e-mail: lassehansen@posteo.de; h.siebert@uni-luebeck.de; c.grossbroehmer@uni-luebeck.de; matthias.heinrich@uni-luebeck.de).

Tony C. W. Mok and Albert C. S. Chung are with the Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Hong Kong (e-mail: cwmokab@connect.ust.hk; achung@cse.ust.hk).

Stephanie Häger, Annkristin Lange, Sven Kuckertz, and Stefan Heldmann are with Fraunhofer MEVIS, Institute for Digital Medicine, 23562 Lübeck, Germany (e-mail: stephanie.haeger@mevis.fraunhofer.de; annkristin.lange@mevis.fraunhofer.de; sven.kuckertz@mevis.fraunhofer.de; stefan.heldmann@mevis.fraunhofer.de).

Wei Shao and Mirabela Rusu are with the Department of Radiology, Stanford University, Stanford, CA 94305 USA (e-mail: weishao@stanford.edu; mirabela.rusu@stanford.edu).

Sulaiman Vesal and Geoffrey Sonn are with the Department of Urology, Stanford University, Stanford, CA 94305 USA (e-mail: svesal@stanford.edu; gsonn@stanford.edu).

Théo Estienne is with the CentraleSupélec, Mathématiques et Informatique pour la Complexité et les Systèmes, Inria Saclay, Institut Gustave Roussy, Inserm, Radiothérapie Moléculaire et Innovation Thérapeutique, Université Paris-Saclay, 91190 Gif-sur-Yvette, France (e-mail: theo.estienne@centralesupelec.fr).

Maria Vakalopoulou is with the CentraleSupélec, Mathématiques et Informatique pour la Complexité et les Systèmes, Inria Saclay, Université Paris-Saclay, 91190 Gif-sur-Yvette, France (e-mail: maria.vakalopoulou@ecp.fr).

Luyi Han, Keelin Murphy, and Nikolas Lessmann are with the Department of Radiology and Nuclear Medicine, Radboud University Medical Center, 6525 GA Nijmegen, The Netherlands (e-mail: hanluyi4869@gmail.com; keelin.murphy@radboudumc.nl; nikolas.lessmann@radboudumc.nl).

Yunzhi Huang is with the School of Automation, Nanjing University of Information Science and Technology, Nanjing 211544, China (e-mail: yunzhi.huang.scu@gmail.com).

Pew-Thian Yap is with the Department of Radiology and Biomedical Research Imaging Center, University of North Carolina, Chapel Hill, NC 28804 USA (e-mail: ptyap@med.unc.edu).

Mikael Brudfors is with NVIDIA, RG2 6UJ Berkshire, U.K. (e-mail: mbrudfors@nvidia.com).

Yaël Balbastre is with the Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Boston, MA 02114 USA, and also with the Harvard Medical School, Boston, MA 02115 USA (e-mail: y.balbastre@ucl.ac.uk).

Samuel Joutard and Marc Modat are with the Department of Biomedical Engineering, King's College London, WC2R 2LS London, U.K. (e-mail: samuel.joutard@kcl.ac.uk; marc.modat@kcl.ac.uk).

Gal Lifshitz and Dan Raviv are with the Faculty of Engineering, the School of Electrical Engineering, Tel Aviv University, Tel Aviv 69978, Israel (e-mail: lifshitz@mail.tau.ac.il; darav@tauex.tau.ac.il).

Jinxin Lv and Qiang Li are with the Wuhan National Laboratory for Optoelectronics-Huazhong University of Science and Technology, Collaborative Innovation Center for Biomedical Engineering, Huazhong University of Science and Technology, Wuhan 430074, China (e-mail: jinx@hust.edu.cn; liqiang@hust.edu.cn).

Vincent Jaouen and Dimitris Visvikis are with UMR 1101 LaTIM, IMT Atlantique, Inserm, 29200 Brest, France (e-mail: vjaouen@gmail.com; visvikis@univ-brest.fr).

Constance Fourcade is with the Ecole Centrale de Nantes, LS2N, UMR CNRS, 44321 Nantes, France and also with Keosys Medical Imaging, 44800 Saint Herblain, France (e-mail: constance.fourcade@ec-nantes.fr).

Mathieu Rubeaux is with Keosys Medical Imaging, 44800 Saint Herblain, France (e-mail: mathieu.rubeaux@keosys.com).

Wentao Pan is with the Shenzhen International Graduate School, Tsinghua University, Beijing 100190, China (e-mail: panwentao1301@gmail.com).

Zhe Xu is with the Department of Biomedical Engineering, The Chinese University of Hong Kong, Hong Kong (e-mail: jackxz@link.cuhk.edu.hk).

Bailiang Jian and Francesca De Benetti are with the Chair for Computer Aided Medical Procedures and Augmented Reality, TUM, 85748 Garching, Germany (e-mail: bailiang.jian@tum.de; francesca.de-benetti@tum.de).

Marek Wodzinski is with the Department of Measurement and Electronics, AGH University of Science and Technology, 30-059 Krakow, Poland, and also with the Information Systems Institute, University of Applied Sciences Western Switzerland, 3960 Sierre, Switzerland (e-mail: marek.m.wodzinski@gmail.com).

Niklas Gunnarsson and Jens Sjölund are with the Department of Information Technology, Uppsala University, 75236 Uppsala, Sweden, and also with Elekta Instrument AB, 113 57 Stockholm, Sweden (e-mail: niklas.gunnarsson@it.uu.se; jens.sjolund@it.uu.se).

Daniel Grzech, Huaqi Qiu, and Zeju Li are with the Department of Computing, Imperial College London, SW7 2BX London, U.K. (e-mail: d.grzech17@imperial.ac.uk; huaqi.qiu15@imperial.ac.uk; zeju.li18@imperial.ac.uk).

Alexander Thorley is with the School of Computer Science, University of Birmingham, B15 2SQ Birmingham, U.K. (e-mail: AJT973@student.bham.ac.uk).

Jinming Duan is with the School of Computer Science, University of Birmingham, B15 2SQ Birmingham, U.K., and also with the Alan Turing Institute, NW1 2DB London, U.K. (e-mail: jduan@bham.ac.uk).

Andrew Hoopes is with the Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Boston, MA 02114 USA (e-mail: ahoopes@mgh.harvard.edu).

Ingerid Reinertsen is with the Department of Health Research, SINTEF Digital, 7037 Trondheim, Norway (e-mail: Ingerid.Reinertsen@sintef.no).

Yiming Xiao is with the Department of Computer Science and Software Engineering, Concordia University, Montréal, QC H3G 1M8, Canada (e-mail: yiming.xiao@concordia.ca).

Bennett Landman and Yuankai Huo are with the Department of Electrical and Computer Engineering, Vanderbilt University, Nashville, TN 37235 USA (e-mail: bennett.landman@vanderbilt.edu; yuankai.huo@vanderbilt.edu).

Adrian V. Dalca is with the Computer Science and Artificial Intelligence Laboratory, Martinos Center for Biomedical Imaging, Massachusetts General Hospital, MIT, Cambridge, MA 02139 USA, and also with the Harvard Medical School, Boston, MA 02115 USA (e-mail: adalca@mit.edu).

REFERENCES

- [1] J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration," *Med. Image Anal.*, vol. 2, no. 1, pp. 1–36, Mar. 1998.
- [2] A. Sotiras, C. Davatzikos, and N. Paragios, "Deformable medical image registration: A survey," *IEEE Trans. Med. Imag.*, vol. 32, no. 7, pp. 1153–1190, Jul. 2013.
- [3] M. A. Viergever, J. B. A. Maintz, S. Klein, K. Murphy, M. Staring, and J. P. W. Pluim, "A survey of medical image registration—under review," *Med. Image Anal.*, vol. 33, pp. 140–144, Oct. 2016.
- [4] G. Haskins, U. Kruger, and P. Yan, "Deep learning in medical image registration: A survey," *Mach. Vis. Appl.*, vol. 31, nos. 1–2, pp. 1–18, Feb. 2020.
- [5] L. Maier-Hein et al., "BIAS: Transparent reporting of biomedical image analysis challenges," *Med. Image Anal.*, vol. 66, Dec. 2020, Art. no. 101796.
- [6] K. Murphy et al., "Evaluation of registration methods on thoracic CT: The EMPIRE10 challenge," *IEEE Trans. Med. Imag.*, vol. 30, no. 11, pp. 1901–1920, Nov. 2011.
- [7] Y. Xiao et al., "Evaluation of MRI to ultrasound registration methods for brain shift correction: The CuRIOUS2018 challenge," *IEEE Trans. Med. Imag.*, vol. 39, no. 3, pp. 777–786, Mar. 2020.
- [8] J. Borovec et al., "ANHIR: Automatic non-rigid histological image registration challenge," *IEEE Trans. Med. Imag.*, vol. 39, no. 10, pp. 3042–3052, Oct. 2020.
- [9] K. Marstal, F. Berendsen, N. Dekker, M. Staring, and S. Klein, "The continuous registration challenge: Evaluation-as-a-Service for medical image registration algorithms," in *Proc. IEEE 16th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2019, pp. 1399–1402.
- [10] J. West et al., "Comparison and evaluation of retrospective intermodality brain image registration techniques," *J. Comput. Assist. Tomogr.*, vol. 21, no. 4, pp. 554–568, 1997.
- [11] A. Klein et al., "Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration," *NeuroImage*, vol. 46, no. 3, pp. 786–802, Jul. 2009.
- [12] Z. Xu et al., "Evaluation of six registration methods for the human abdomen on clinically acquired CT," *IEEE Trans. Biomed. Eng.*, vol. 63, no. 8, pp. 1563–1572, Aug. 2016.
- [13] R. Castillo et al., "A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets," *Phys. Med. Biol.*, vol. 54, no. 7, p. 1849, 2009.
- [14] B. B. Avants, C. L. Epstein, M. Grossman, and J. C. Gee, "Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain," *Med. Image Anal.*, vol. 12, no. 1, pp. 26–41, 2008.
- [15] S. Klein, M. Staring, K. Murphy, M. A. Viergever, and J. P. Pluim, "Elastix: A toolbox for intensity-based medical image registration," *IEEE Trans. Med. Imag.*, vol. 29, no. 1, pp. 196–205, Jan. 2010.
- [16] M. Modat et al., "Fast free-form deformation using graphics processing units," *Comput. Methods Programs Biomed.*, vol. 98, no. 3, pp. 278–284, 2010.
- [17] H. P. Heinrich, M. Jenkinson, M. Brady, and J. A. Schnabel, "MRF-based deformable registration and ventilation estimation of lung CT," *IEEE Trans. Med. Imag.*, vol. 32, no. 7, pp. 1239–1248, Jul. 2013.
- [18] A. V. Dalca, G. Balakrishnan, J. Guttag, and M. Sabuncu, "Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces," *Med. Image Anal.*, vol. 57, pp. 226–236, Oct. 2019.
- [19] Y. Xiao, M. Fortin, G. Unsgard, H. Rivaz, and I. Reinertsen. (2020). *Easy-ResectEASY-RESECT*. [Online]. Available: <https://archive.sigma2.no/pages/public/datasetDetail.jsf?id=10.11582/2020.00025>
- [20] Y. Xiao, M. Fortin, G. Unsgård, H. Rivaz, and I. Reinertsen, "REtrospective evaluation of cerebral tumors (RESECT): A clinical database of pre-operative MRI and intra-operative ultrasound in low-grade glioma surgeries," *Med. Phys.*, vol. 44, no. 7, pp. 3875–3882, Jul. 2017.
- [21] Y. Xiao et al., "Evaluation of MRI to ultrasound registration methods for brain shift correction: The CuRIOUS2018 challenge," *IEEE Trans. Med. Imag.*, vol. 39, no. 3, pp. 777–786, Mar. 2020.
- [22] M. Antonelli et al., "The medical segmentation decathlon," 2021, *arXiv:2106.05735*.
- [23] K. Clark et al., "The cancer imaging archive (TCIA): Maintaining and operating a public information repository," *J. Digit. Imag.*, vol. 26, no. 6, pp. 1045–1057, Dec. 2013.
- [24] O. Akin et al., "The cancer genome atlas kidney renal clear cell carcinoma collection (TCGA-KIRC), version 3," The Cancer Imaging Archive, 2016. [Online]. Available: <https://wiki.cancerimagingarchive.net/pages/viewpage.action?pageId=5800386#5800386a4a7b137e03049b2b4865d5b23cba736>, doi: [10.7937/K9/TCIA.2016.V6PBVTDR](https://doi.org/10.7937/K9/TCIA.2016.V6PBVTDR).
- [25] M. Linehan et al., "The cancer genome atlas cervical kidney renal papillary cell carcinoma collection (TCGA-KIRP), version 4," The Cancer Imaging Archive, 2016, doi: [10.7937/K9/TCIA.2016.ACWOGBEF](https://doi.org/10.7937/K9/TCIA.2016.ACWOGBEF).
- [26] B. J. Erickson et al., "The cancer genome atlas liver hepatocellular carcinoma collection (TCGA-LIHC), version 5," The Cancer Imaging Archive, 2016, doi: [10.7937/K9/TCIA.2016.IMMQW8UQ](https://doi.org/10.7937/K9/TCIA.2016.IMMQW8UQ).
- [27] M. P. Heinrich, M. Jenkinson, M. Brady, and J. A. Schnabel, "Globally optimal deformable registration on a minimum spanning tree using dense displacement sampling," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2012, pp. 115–122.
- [28] A. E. Kavur, M. A. Selver, O. Dicle, M. Barış, and N. S. Gezer, "CHAOS—Combined (CT-MR) healthy abdominal organ segmentation challenge data, v1.03," Zenodo, Apr. 2019. [Online]. Available: https://zenodo.org/record/3431873#_Y1G8Ti-21pQ, doi: [10.5281/zenodo.3362844](https://doi.org/10.5281/zenodo.3362844).
- [29] A. E. Kavur et al., "CHAOS challenge-combined (CT-MR) healthy abdominal organ segmentation," *Med. Image Anal.*, vol. 69, Apr. 2021, Art. no. 101950.
- [30] D. S. Marcus, T. H. Wang, J. Parker, J. G. Csernansky, J. C. Morris, and R. L. Buckner, "Open access series of imaging studies (OASIS): Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults," *J. Cognit. Neurosci.*, vol. 19, no. 9, pp. 1498–1507, Sep. 2007.
- [31] B. Fischl, "Freesurfer," *NeuroImage*, vol. 62, no. 2, pp. 774–781, Aug. 2012.
- [32] A. Hoopes, M. Hoffmann, B. Fischl, J. Guttag, and A. V. Dalca, "Hypermorph: Amortized hyperparameter learning for image registration," in *Information Processing in Medical Imaging*. Cham, Switzerland: Springer, 2021, pp. 3–17.
- [33] A. Hering, K. Murphy, and B. van Ginneken, "Learn2Reg challenge: CT lung registration—Training data," May 2020, doi: [10.5281/zenodo.3835682](https://doi.org/10.5281/zenodo.3835682).
- [34] A. Hering, K. Murphy, and B. van Ginneken, "Learn2Reg challenge: CT lung registration—Test data," Zenodo, Sep. 2020, doi: [10.5281/zenodo.4048761](https://doi.org/10.5281/zenodo.4048761).
- [35] A. D. Leow et al., "Statistical properties of Jacobian maps and the realization of unbiased large-deformation nonlinear image registration," *IEEE Trans. Med. Imag.*, vol. 26, no. 6, pp. 822–832, Jun. 2007.
- [36] S. Kabus, T. Klinder, K. Murphy, B. van Ginneken, C. Lorenz, and J. P. Pluim, "Evaluation of 4D-CT lung registration," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2009, pp. 747–754.
- [37] L. Han, H. Dou, Y. Huang, and P.-T. Yap, "Deformable registration of brain MR images via a hybrid loss," 2021, *arXiv:2110.15027*.
- [38] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, "VoxelMorph: A learning framework for deformable medical image registration," *IEEE Trans. Med. Imag.*, vol. 38, no. 8, pp. 1788–1800, Aug. 2019.
- [39] DeepRegNet. (2021). *DeepregNet*. [Online]. Available: <https://github.com/Project-MONAI/MONAI/blob/dev/monai/networks/nets/regunet.py>
- [40] H. Siebert, L. Hansen, and M. P. Heinrich, "Fast 3D registration with accurate optimisation and little learning for Learn2Reg 2021," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2021, pp. 174–179.
- [41] L. Hansen and M. P. Heinrich, "GraphRegNet: Deep graph regularisation networks on sparse keypoints for dense registration of 3D lung CTs," *IEEE Trans. Med. Imag.*, vol. 40, no. 9, pp. 2246–2257, Sep. 2021.

- [42] M. P. Heinrich, H. Handels, and I. J. Simpson, "Estimating large lung motion in COPD patients by symmetric regularised correspondence fields," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2015, pp. 338–345.
- [43] J. Lv et al., "Joint progressive and coarse-to-fine registration of brain MRI via deformation field integration and non-rigid feature fusion," *IEEE Trans. Med. Imag.*, vol. 41, no. 10, pp. 2788–2802, Oct. 2022.
- [44] C. Fourcade, M. Rubeaux, and D. Mateus, "Using Elastix to register inhale/exhale intrasubject thorax CT: A unsupervised baseline to the task 2 of the Learn2Reg challenge," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2020, pp. 100–105.
- [45] T. Estienne et al., "Deep learning-based concurrent brain registration and tumor segmentation," *Frontiers Comput. Neurosci.*, vol. 14, p. 17, Mar. 2020.
- [46] T. Estienne et al., "Deep learning based registration using spatial gradients and noisy segmentation labels," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, vol. 12587. Cham, Switzerland: Springer, 2020, p. 87.
- [47] N. Gunnarsson, J. Sjölund, and T. B. Schön, "Learning a deformable registration pyramid," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2020, pp. 80–86.
- [48] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2018, pp. 8934–8943.
- [49] M. C. Lee, O. Oktay, A. Schuh, M. Schaap, and B. Glocker, "Image-and-spatial transformer networks for structure-guided image registration," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2019, pp. 337–345.
- [50] T. C. Mok and A. C. Chung, "Large deformation diffeomorphic image registration with Laplacian pyramid networks," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2020, pp. 211–221.
- [51] T. C. Mok and A. Chung, "Conditional deformable image registration with convolutional neural network," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2021, pp. 35–45.
- [52] V. Jaouen, P.-H. Conze, G. Dardenne, J. Bert, and D. Visvikis, "Regularized directional representations for medical image registration," 2021, *arXiv:2111.15509*.
- [53] G. Lifshitz and D. Raviv, "Cost function unrolling in unsupervised optical flow," 2020, *arXiv:2011.14814*.
- [54] L. Liu et al., "Learning by analogy: Reliable supervision from transformations for unsupervised optical flow estimation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 6489–6498.
- [55] S. Häger, S. Heldmann, A. Hering, S. Kuckertz, and A. Lange, "Variable Fraunhofer MEVIS RegLib comprehensively applied to Learn2Reg challenge," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2020, pp. 74–79.
- [56] A. Hering, A. Lange, S. Heldmann, S. Häger, and S. Kuckertz, "Fraunhofer MEVIS image registration solutions for the Learn2Reg 2021 challenge," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2021, pp. 147–152.
- [57] M. Brudfors, Y. Balbastre, G. Flandin, P. Nachev, and J. Ashburner, "Flexible Bayesian modelling for nonlinear image registration," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2020, pp. 253–263.
- [58] M. P. Heinrich, "Closing the gap between deep and conventional image registration using probabilistic dense displacement networks," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2019, pp. 50–58.
- [59] M. P. Heinrich and L. Hansen, "Highly accurate and memory efficient unsupervised learning-based discrete CT registration using 2.5 D displacement search," in *Medical Image Computing and Computer Assisted Intervention MICCAI 2020*. Cham, Switzerland: Springer, 2020, pp. 190–200.
- [60] W. Shao et al., "Geodesic density regression for correcting 4DCT pulmonary respiratory motion artifacts," *Med. Image Anal.*, vol. 72, Aug. 2021, Art. no. 102140.
- [61] A. Thorley et al., "Nesterov accelerated ADMM for fast diffeomorphic image registration," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2021, pp. 150–160.
- [62] E. Gibson et al., "Automatic multi-organ segmentation on abdominal CT with dense V-networks," *IEEE Trans. Med. Imag.*, vol. 37, no. 8, pp. 1822–1834, Aug. 2018.
- [63] O. Jimenez-del-Toro et al., "Cloud-based evaluation of anatomical structure segmentation and landmark detection algorithms: VISCERAL anatomy benchmarks," *IEEE Trans. Med. Imag.*, vol. 35, no. 11, pp. 2459–2475, Nov. 2016.
- [64] A. Hering, S. Häger, J. Moltz, N. Lessmann, S. Heldmann, and B. van Ginneken, "CNN-based lung CT registration with multiple anatomical constraints," *Med. Image Anal.*, vol. 72, Aug. 2021, Art. no. 102139.
- [65] M. Hoffmann, B. Billot, D. N. Greve, J. E. Iglesias, B. Fischl, and A. V. Dalca, "SynthMorph: Learning contrast-invariant registration without acquired images," *IEEE Trans. Med. Imag.*, vol. 41, no. 3, pp. 543–558, Mar. 2022.
- [66] J. Tukey, "Bias and confidence in not quite large samples," *Ann. Math. Statist.*, vol. 29, no. 2, p. 614, 1958.
- [67] J. Ruhaak et al., "Estimation of large motion in lung CT by integrating regularized keypoint correspondences into dense deformable registration," *IEEE Trans. Med. Imag.*, vol. 36, no. 8, pp. 1746–1757, Aug. 2017.