Supplementary Information

Incorporating sliding and local rigidity constraints in neural-based image registration

Ziad Kheil, Soleakhena Ken, Laurent Risser

S1 Physics-informed losses reduce sensitivity to regularization hyperparameters

We first assessed the impact of regularization loss weight described in the methods section for $\mathcal{D}_{rigid}$ and $\mathcal{D}_{shearing}$. Figure S1 summarizes the results we obtained.

As expected, increasing regularization improves field smoothness $\mathbb{L}_{rigid}$ Fig. S1(b,d) and $SD \log |J_\Phi|$ (see Fig. S1(a,c)) at the cost of registration accuracy. More interestingly, a secondary trade-off emerges between overall smoothness and biomechanical plausibility as highlighted by the alternating optimal fronts in Figures S1(a,b) and S1(c,d). Penalties which include our proposed $\mathbb{L}_{rigid}$ far better respect local rigidity while only marginally affecting $SD \log |J_\Phi|$, an effect roughly one order of magnitude smaller. Furthermore, the Jacobian-only loss shows strong sensitivity at high α values (minimal regularization), leading to abrupt variations in deformation behavior.

These results confirm that our biomechanical regularizers produce smoother and more physically plausible deformation fields across unseen test data from the same distribution. We note that this sweep varies the overall regularization strength α with the three loss weights fixed at $\beta_r = \beta_s = \beta_j = 1$ and supports this equal-weight choice only in that sense; a decoupled grid search varying β_r,β_s and β_j independently has not been performed.

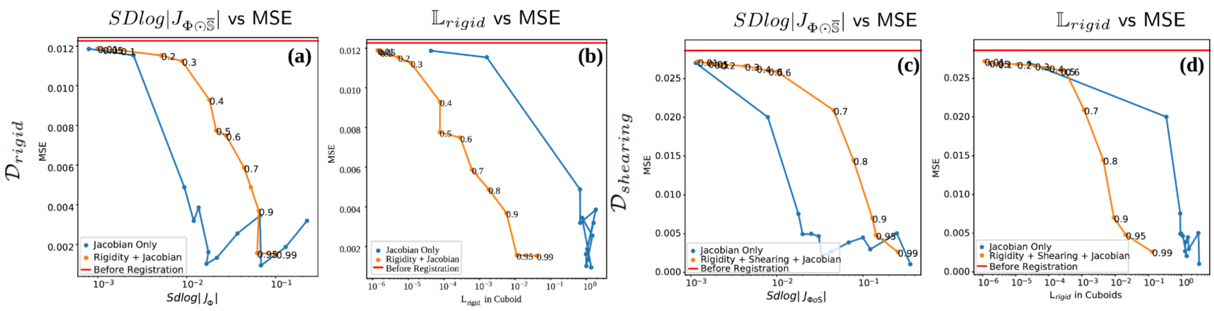

Figure S1: Sensitivity analysis to loss coefficients on the test sets of simulated datasets. Figures **(a-b)** correspond to $\mathcal{D}_{rigid}$ while **(c-d)** represent $\mathcal{D}_{shearing}$. Furthermore, in **(a-c)**, $SDlog|J_\Phi|$ is plotted against MSE, whereas **(b-d)** illustrates the relationship between measured rigidities inside the cuboids and the corresponding MSE measures. The plotted values reflect the weighting of the similarity term in the loss function. Blue curves represent models trained exclusively with the Jacobian term, while orange curves correspond to models incorporating our proposed physical regularization term.

S2 Computational cost

Table S1: Computational cost of the proposed regularization compared with the smooth baseline. Runtimes are medians over the 99 measurable epoch intervals of each run, averaged over the three weight configurations trained per strategy, on identical hardware (NVIDIA GeForce RTX 3090 GPU) with identical architecture, volume size ($96 \times 80 \times 128$) and batch size. Total training time is the median epoch time extrapolated to 100 epochs. Additional memory is the per-sample annotation (1-channel region mask, 3-channel normal field) required at training time only.

Note that given the Dice loss term, segmentations are required regardless of regularization method

	Smooth reg.	Biomechanical reg.
Training time (s/epoch)	47.3 ± 0.1	55.7 ± 0.1 (+18%)
Total training, 100 epochs	1h19	1h33
Additional memory per sample	–	12.2 MB (4 channels)
Annotation, per volume (one-off)	segmentation	segmentation, mask, normals
Inference, per image pair	**identical**	

S3 Detailed deformations

Random samples are shown in Fig. S2. The regular grid structure is preserved within the bones for the model trained with our loss function. In shearing-prone regions, Fig. S3 shows that the *Biomechanical reg.* model sometimes produces deformations with opposing directions, which respects the local discontinuities expected in movements, whereas the *Smooth reg.* model, lacking the shearing term, defaults to a uniform flow and maintains a consistent flow direction.

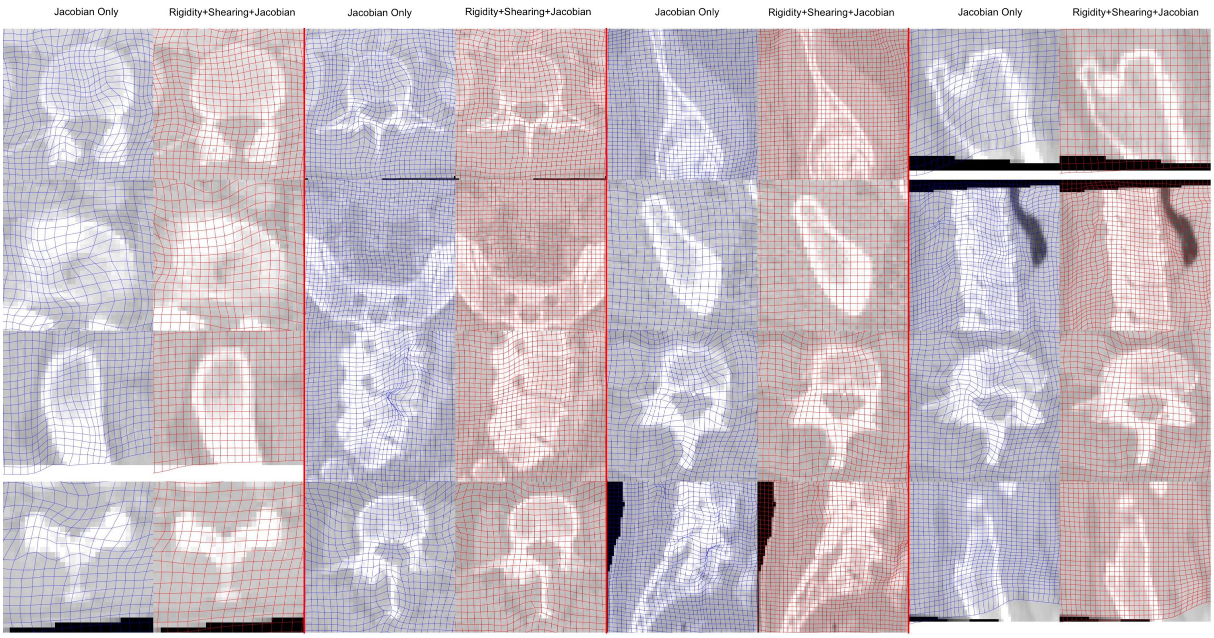

Figure S2: Grid display of deformation fields produced by both models, randomly sampled on structures marked as rigid, with the moving images overlaid on the deformation fields. The columns alternate between *Smooth reg.* (blue grids) and *Biomechanical reg.* (red grids). Rigidity is improved in the second columns of each pair.

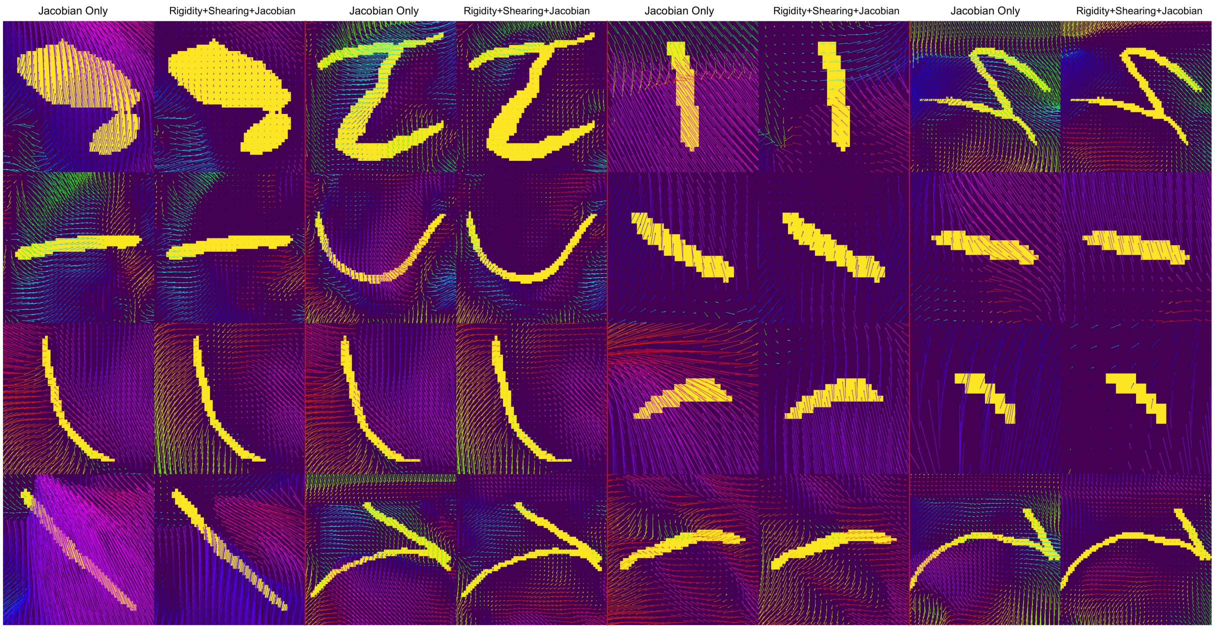

Figure S3: Grid display of deformation fields produced by both models, randomly sampled on regions where the shearing loss was applied in the *Biomechanical reg.* case, with the shearing region overlaid. The columns alternate between *Smooth reg.* and *Biomechanical reg.* In most cases, the model trained with our shearing loss produces sharp transitions inside these regions.

S4 Additional example registrations

Complementary registration examples on further AbdomenCTCT patients, beyond Fig.3 in the main text, are given in Figs. S4–S6, spanning a range of registration difficulty rather than a single best-case illustration.

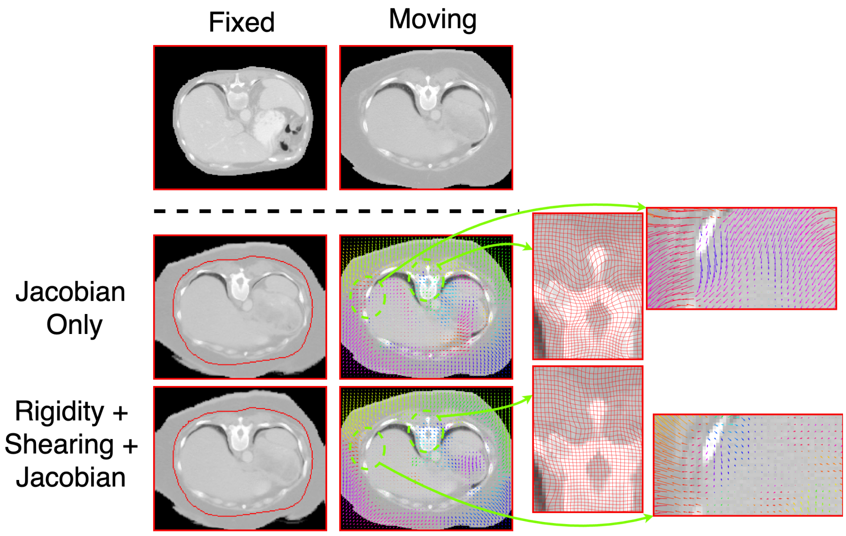

Figure S4: Complementary illustration to Fig.3, obtained on patient 2.

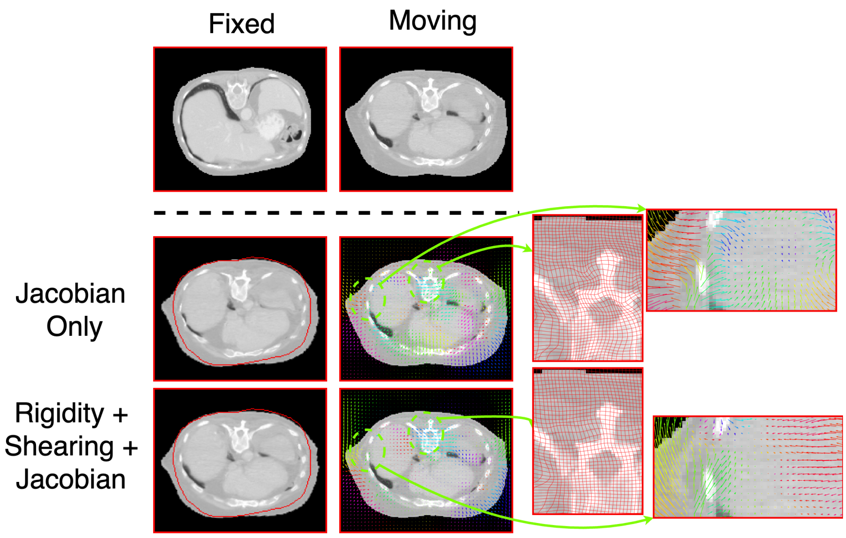

Figure S5: Complementary illustration to Fig.3, obtained on patient 3.

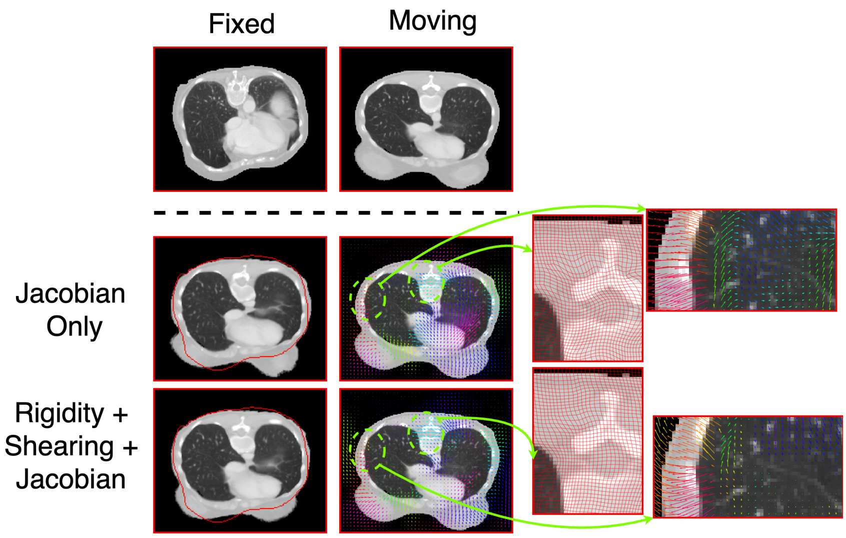

Figure S6: Complementary illustration to Fig.3, obtained on patient 4, among the lower-performing cases in our validation set.

S5 Per-organ registration accuracy

Figure S7 breaks down Dice scores for each of the 13 anatomical structures provided by the official Learn2Reg AbdomenCTCT annotations, complementing the pooled $Dice_{learn2reg}$ figure reported in Table 2, and showing where the proposed regularization helps most and where the two strategies are statistically indistinguishable.

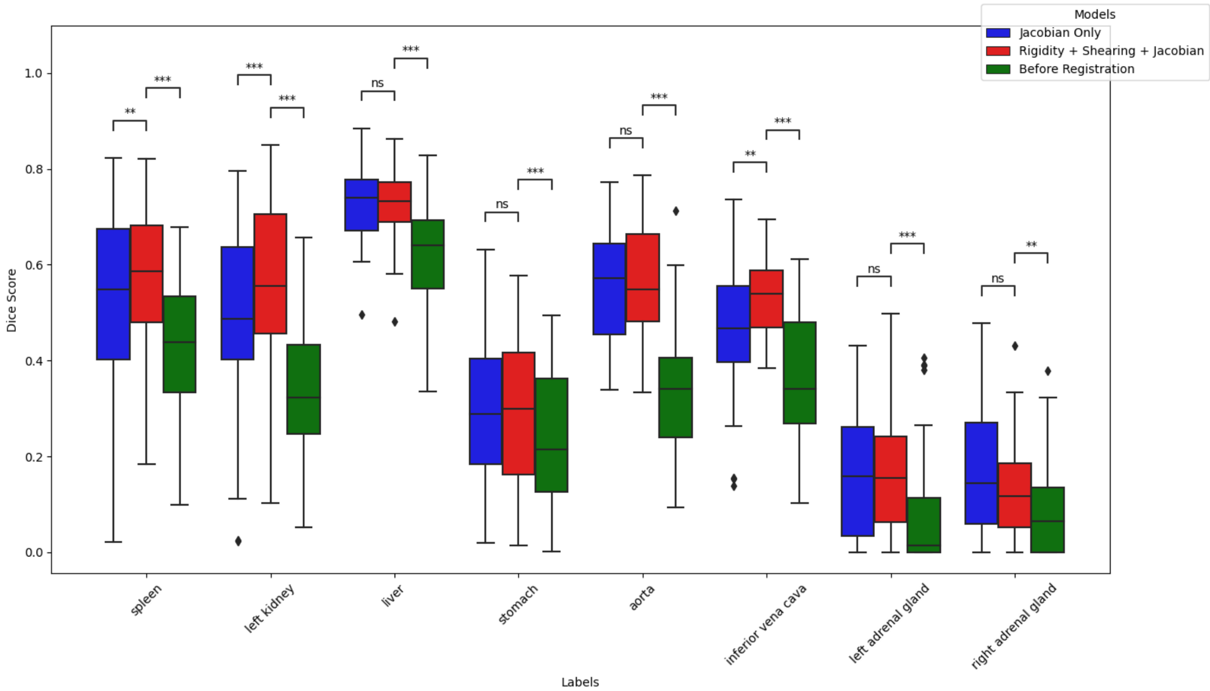

Figure S7: Detailed Dice scores per organ. For each organ and each pair of models, a Wilcoxon signed-rank test is computed on the paired per-validation-pair Dice score differences, under the null hypothesis that the median difference is zero (i.e. no systematic difference in per-organ Dice between the two models); significance tiers are reported a: $n.s \rightarrow \textit{Not Significant}, * \rightarrow p < 10^{-2}, ** \rightarrow p < 10^{-3}, *** \rightarrow p < 10^{-4}$.

S6 Regularization regions overview

Through discussions with radiologists and anatomy knowledge, we compiled two lists of organs whose segmentation masks were used to define the different types of expected deformations, hence the different losses, as described in the **Methods** section.

The rigidity loss was applied to bone tissues using the following labels: **Vertebrae** (S1, L5, L4, L3, L2, L1, T12, T11, T10, T9, T8, T7, T6, T5, T4, T3, T2, T1, C7, C6, C5, C4, C3, C2, C1); **Ribs** (left and right: 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12); **Humerus** (left and right); **Scapula** (left and right); **Clavicula** (left and right); **Femurs** (left and right); **Hips** (left and right); **Sacrum**; **Sternum**; **Costal cartilages**; **Skull**. Note that some of these structures lie outside the field of view of the AbdomenCTCT volumes; they are kept in the list for completeness and simply yield empty masks there.

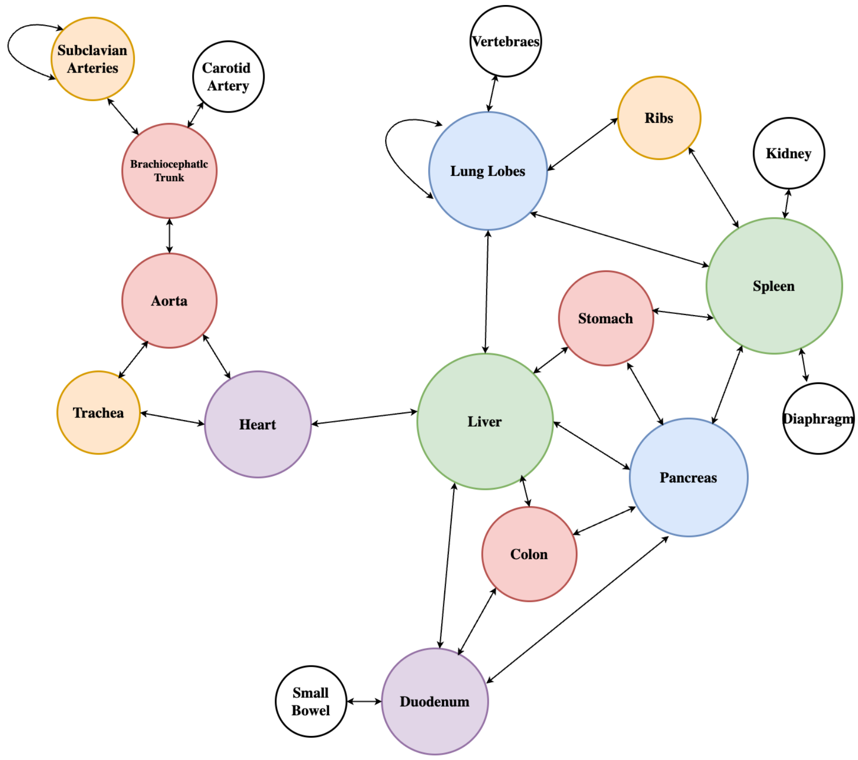

Figure S8: Graph of the tissue interfaces between which sliding is permitted. For readability, different sub-structures of a same tissue are conflated into single nodes (lung lobes, vertebrae, kidneys, ribs, etc.). Nodes with self-edges, e.g. *Lung Lobes*, correspond to interactions between sub-structures of the same tissue, such as the lower and upper lung lobes in contact along the pleura.

Meanwhile, through the process detailed in the **Methods** section, a sliding-prone region was allowed at the interfaces of the tissue pairs identified by radiologists, described in Fig. S8. The dilation $k_{ab} \in \{1, 2, 3\}$ applied to each pair before intersection is specific to that pair, as some interfaces are separated by intervening tissue that is not represented in the segmentation, the diaphragm between liver and lung being the clearest example. The complete pair list with the corresponding tolerances is given by the annotation script of the repository.