SparseVM: Fast Learning-based Registration of Sparse Clinical Images

Kathleen Lewis[1,2], Guha Balakrishnan[1,2], Natalia S. Rost[3], John Guttag[1.2], Adrian V. Dalca[1,2,3]

EECS[1], CSAIL[2], MGH/HMS[3]

Motivation for Clinical Scan Registration

code available at voxelmorph.mit.edu

Motivation: Existing registration methods fail on clinical images because of sparsity

Challenges

Missing up to 85% of slices
Inconsistent anatomy across acquired slices
Traditional methods: slow
Learning-based methods: low accuracy

Our Contribution

Faster and more accurate clinical scan registration
> 100x faster on a CPU
More accurate on **>86%** of the test images

SparseVM

Current Methods

All methods linearly interpolate acquired scans
Patch-Based Approach: **Slow**, High Accuracy
Learning-Based Approach: **Fast**, Low Accuracy

Our Method (SparseVM)

New loss function combines best of both methods
Evaluate loss only on acquired voxels
Unsupervised learning method using CNN

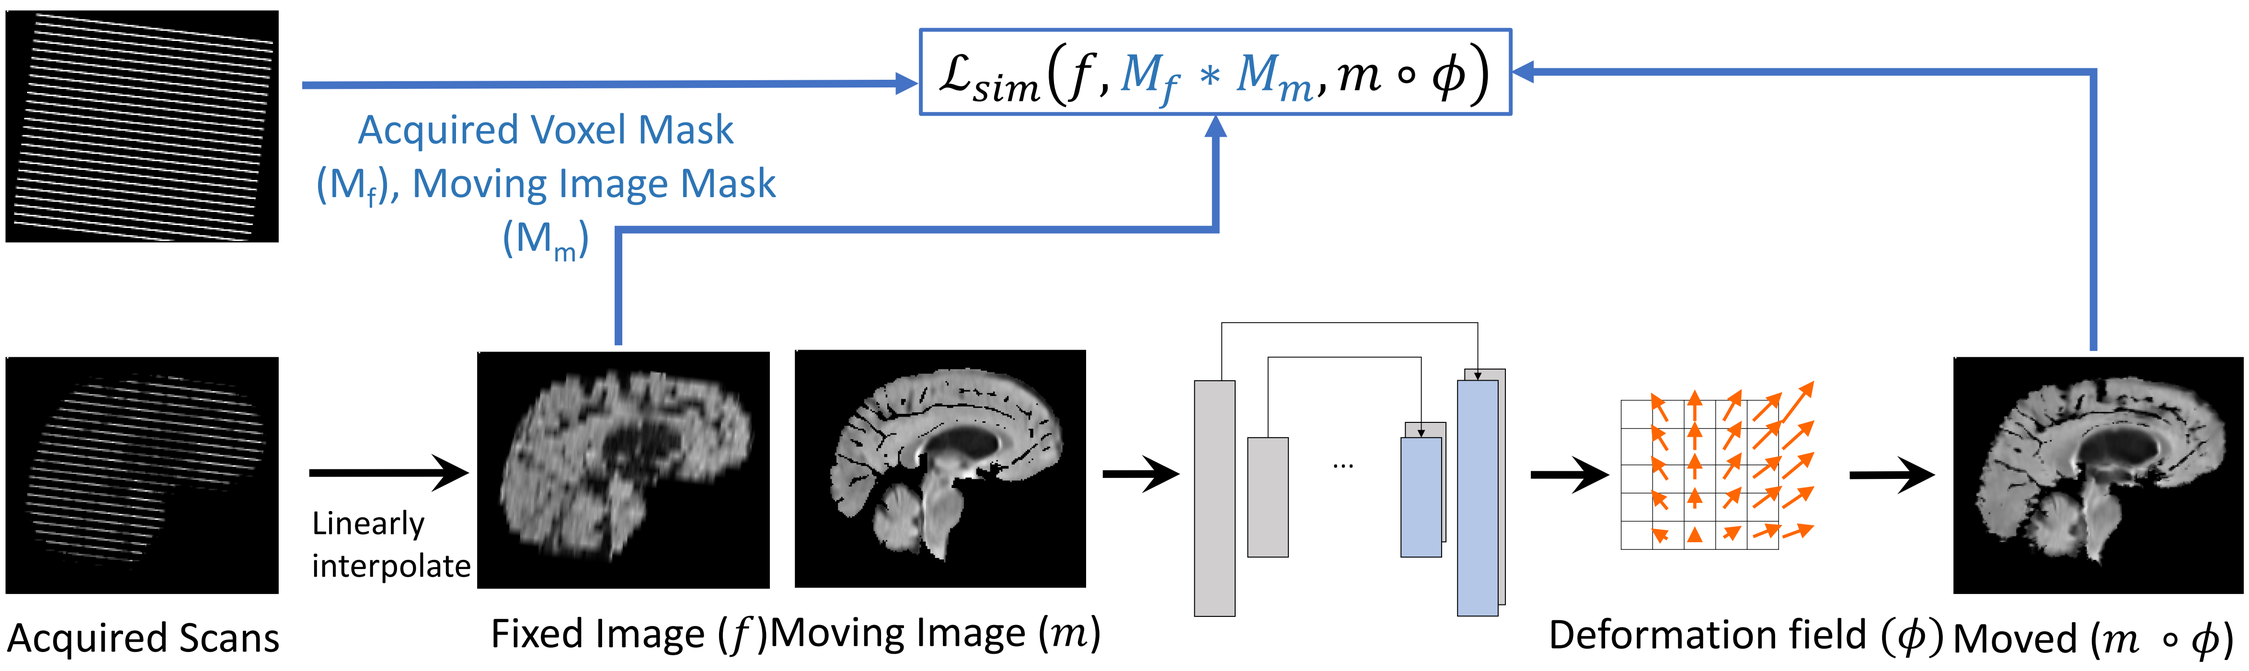

Experiments and Results

Baselines

ANTs: commonly used
Patch-Based Registration (PBR) : most consistently accurate
VoxelMorph (VM) with CC loss: fastest

Evaluation Metric

Ventricle Dice: overlap of ventricle segmentations

Dataset

3D T2-FLAIR MR stroke scans from MGH

Bottom Line

Now register clinical images in a reasonable amount of time
Implemented in new pipelines at MGH

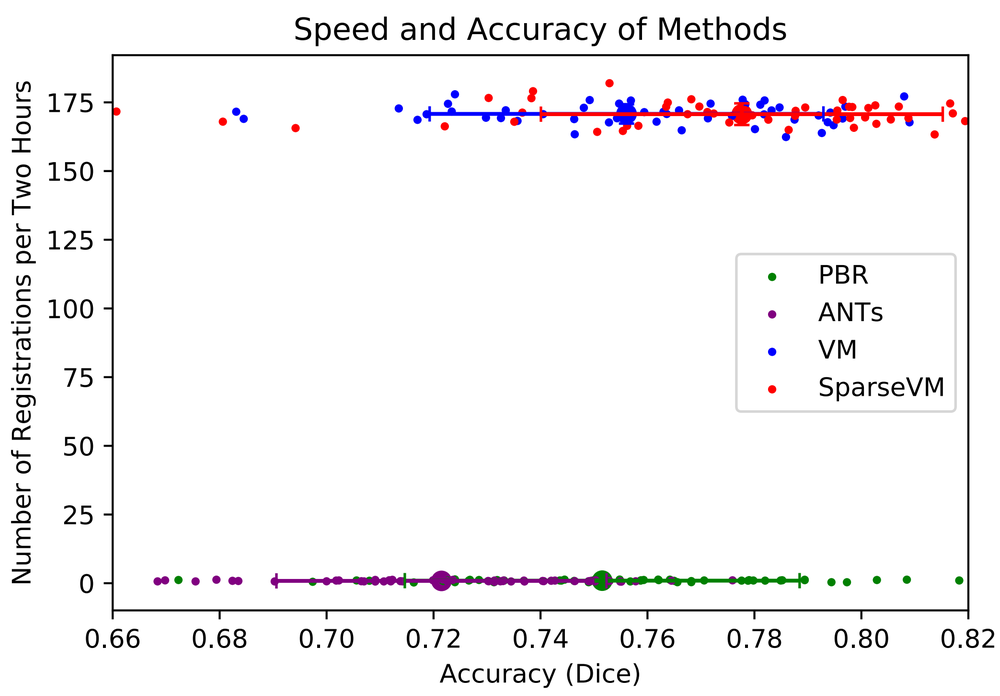

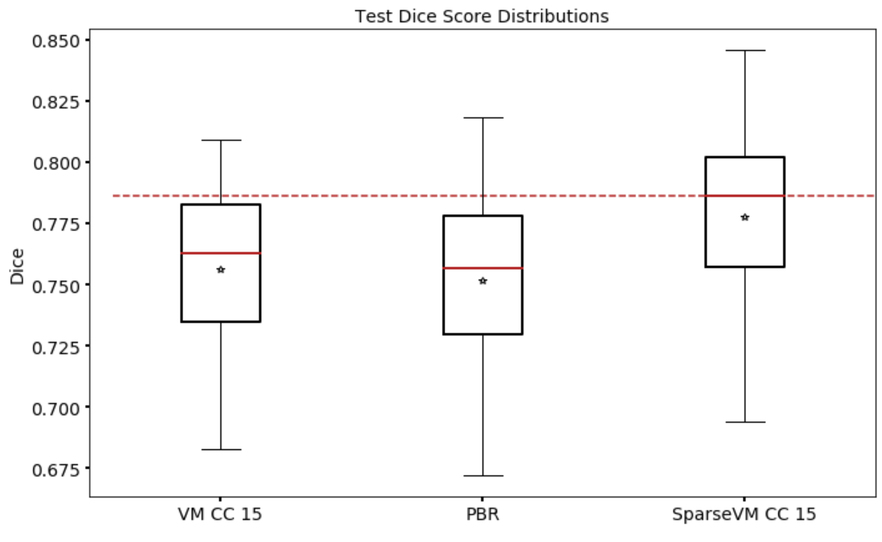

Method	Average Dice (SD)	GPU (sec)	CPU(sec)
ANTs	0.722 (0.031)	-	9059 (2023)
PBR	0.752 (0.037)	-	9269 (5134)
VoxelMorph	0.756 (0.037)	0.313 (0.046)	40 (0.693)
SparseVM CC (ours)	**0.778 (0.038)**	0.303 (0.047)	41 (0.584)