



Fast Localization of Topological Features in fMRI Voxels with Union-Find



Locating components and Voids with Union-Find – Padma Tanikella and Nicole Lazar



Motivation: Many fMRI studies reduce each brain volume to a connectivity matrix or a vector before comparing different groups using topological summaries like persistence diagrams. Very few studies interpret which brain regions drive the differences.





Goal: Efficiently map every dim 0 and 2 feature back to the voxels that produce it, as a base for interpretation, using union-find algorithm.



Dim 0: connected components: time per brain volume			
	gudhi-based	union-find	speed-up
Top-6 components with their voxels	2.29 s	0.39 s	6×
All components with size and centroid	2.72 s	0.35 s	8×
Barcode only	2.22 s	0.35 s	6×

gudhi-based: gudhi cubical persistence, then one scipy label call for every component whose voxels are needed. Union-find: one sweep of the sorted voxel graph. Median time per brain volume over 18 volumes (3 per age group, TR 50, 79x95x68 voxels, trim 150, T construction), one CPU thread. The persistence values are identical in every volume. 'All components' for gudhi was timed on 3 volumes.

Estimated for all 155 subjects (26,040 volumes), all components with size and centroid: gudhi-based 19.7 h vs. union-find 2.6 h (median time per volume $\times$ number of volumes, one CPU thread).



Dim 2: voids: time per brain volume			
	gudhi-based	union-find	speed-up
Top-6 voids with their voxels	2.17 s	0.16 s	14×
All voids with size and centroid	3.13 s	0.13 s	24×
Barcode only	2.16 s	0.11 s	19×



gudhi-based: gudhi cubical persistence of the complement (its H_0 is the volume's H_2), then one scipy label call per void. 'Barcode only' for gudhi is H_2 of the volume itself, which gives no voxel locations. Union-find: one sweep on the complement, with sizes and centroids from running sums. Median time per brain volume over 18 volumes (3 per age group, TR 50, 79x95x68 voxels, trim 150, T construction), one CPU thread. The persistence values are identical in every volume. 'All voids' and the H_2 barcode for gudhi were timed on 3 volumes.

Estimated for all 155 subjects (26,040 volumes), all voids with size and centroid: gudhi-based 22.7 h vs. union-find 0.9 h (median time per volume $\times$ number of volumes, one CPU thread).

Future Work:

- **Connect the topology to function.** Do the features track brain function? Subject-level and condition-specific tests against the ToM and pain networks.
- Cubical representative cycles are not unique, and the reduction picks one of many equivalent choices. Can we define or visualize what an “average” 1-cycle looks like?

Finding a loop by matrix reduction: pairing $(e03, sq)$ in $R = DV$ gives a 1-cycle whose representative is V_{e03}



1. Fast Cubical Persistent Homology on 2D and 3D Images via Union-Find, Pruning, and Lookup Tables, Le Breton et al., 2026, arXiv
2. Development of the social brain from age three to twelve years, Richardson et al., Nat Commun 2018, OpenNeuro dataset ds000221