High-Throughput Image Alignment for Connectomics

This presentation covers scalable algorithms Quilter and Stacker for efficient image alignment in connectomics.

We explore frugal snap judgments (FSJ) to boost throughput while maintaining accuracy.

S

by Sudha Sree Yerramsetty

CS 455

Introduction to Alignment Pipeline

Multicore Algorithms

Quilter (2D) and Stacker (3D) enable scalable, memory-efficient alignment.

Pipeline Design

Coarse-grained task distribution supports cluster scalability.

Performance-Accuracy Trade-Offs

FSJ balances speed and alignment quality.

2D Image Alignment with Quilter

Algorithm Overview

Processes overlapping image tiles using keypoint matching and spring optimization.

Minimizes I/O and recomputation with line-sweep task ordering.

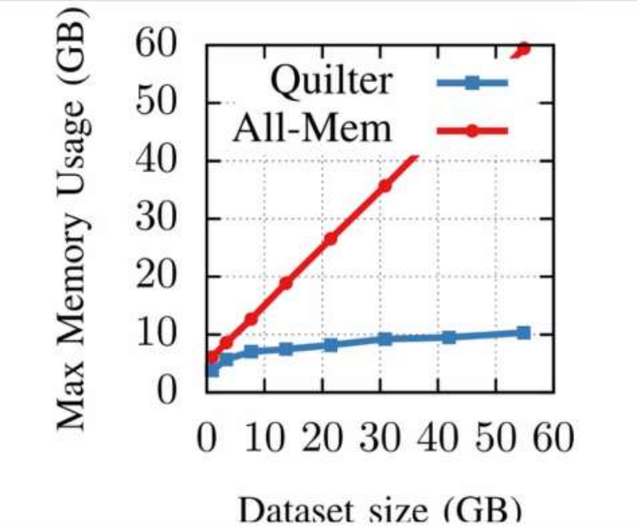

Memory Efficiency

Aligns large brain sections using under 1TB RAM by loading only needed tiles.

Each tile has an approximate location and overlaps slightly with neighboring tiles, enabling the algorithm to identify shared landmarks.

Using these landmarks, the algorithm computes relative positions and applies a spring based optimization model to refine tile placement, minimizing the overall alignment error.

3D Image Alignment with Stacker

(performs 3D alignment on 2D-aligned sections (produced by Quilter) using both affine and non-affine (elastic) transformations to correct distortions from sample preparation and imaging)

Independent Section Alignment

Aligns adjacent sections pairwise independently using a hexagonal triangle mesh overlay.

Memory Usage

Requires ~4TB RAM for a human brain, ~50GB for mouse brain datasets, hence feasible for connectomics.

Scalability

Supports efficient horizontal scaling in cluster pipelines.

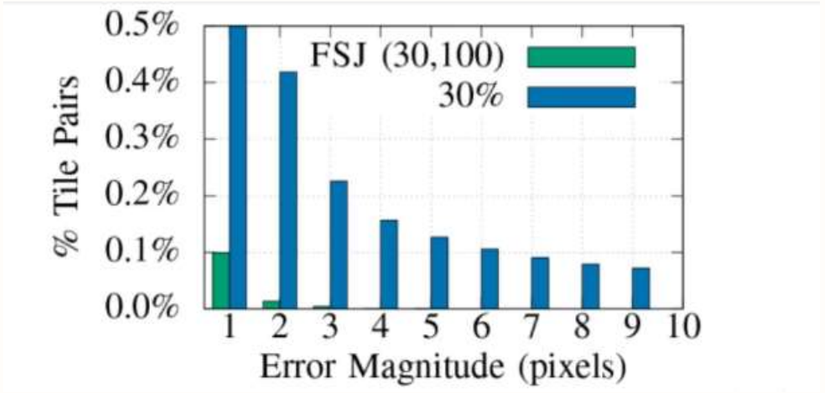

% of overlapping tile pairs misaligned when downsampling to 30% resolution on human100

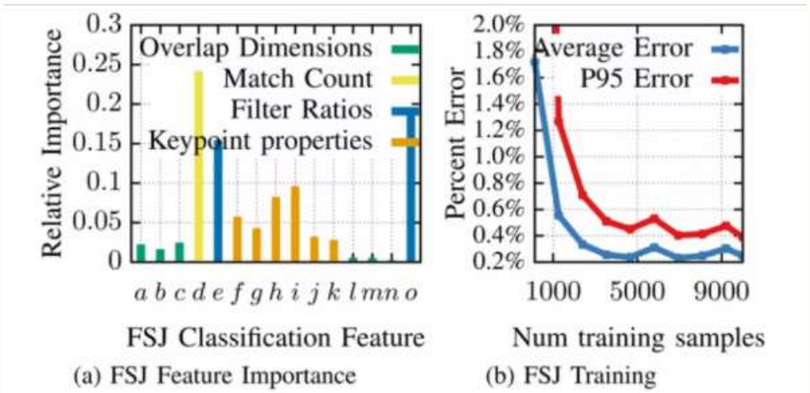

relative importance of variables used by FSJ on human 100 dataset & FSJ classification error during training

Frugal Snap Judgments (FSJ)

Fast Path

Uses 30% downsampled images for quick alignment.

Speedup

Achieves 3–5× faster alignment with minimal accuracy loss.

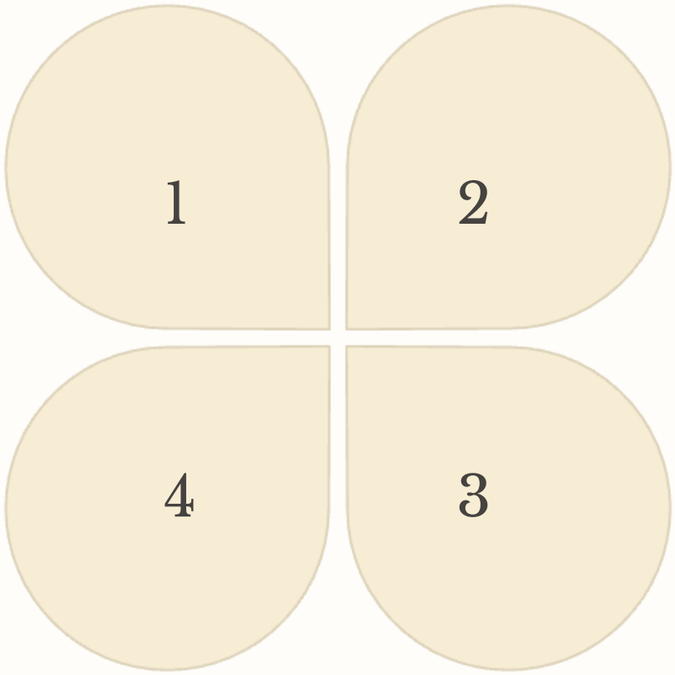

Random Forest Classifier

Predicts when full-resolution processing is needed.

Accuracy

Nearly no errors >1 pixel; no errors >5 pixels.

System Evaluation Setup

Software

C++ pipeline with Cilk Plus, LLVM Tapir, OpenCV, OpenJPEG, and Google Protocol Buffers.

Datasets

- Mouse50: 550GB, 65K tiles
- Mouse200: 2TB, 200K tiles
- Human100: 38TB, very high resolution

Hardware

Tested on Common Multicore, Large Multicore, LLSC and AWS Clusters.

Multicore Performance Results

Dataset	Speedup vs FijiBento
Mouse50	11.3× (Quilter with FSJ)
Mouse200	Up to 39× on 112 cores
Human100	Near-linear scaling (similar to Mouse200)

Stacker completes 3D alignment of mouse50 in 8 minutes.

Cluster Scaling Performance

Strong Scaling

Near-linear strong scaling on AWS Cluster with Mouse200 dataset.

Weak Scaling

Maintains efficiency from 1 to 128 nodes as the dataset size and cluster size grow proportionally.

End-to-End Throughput

The pipeline achieves 21.4 TB/hour throughput on AWS Cluster with 1600 cores.

Full results for human100 dataset obtained on Large Multicore due to data transfer limits.

Conclusion and Future Work

1

Efficient Pipeline

Quilter and Stacker enable scalable, memory-conscious alignment.

2

FSJ Technique

Balances speed and accuracy with fast-path judgments.

3

Challenges Ahead

Address knife replacement artifacts and tissue folds for robustness.