Intraoperative 3D Reconstruction and Deep Learning based Annotation of Cutaneous Squamous Cell Carcinoma from Mohs Micrographic Surgery

Shritej Mamidala, Abhi Kurapati

Emerging Diagnostic and Investigative Technologies, Department of Pathology, Dartmouth Hitchcock Medical Center

Dartmouth Health

ABSTRACT

- Developed an AI-driven pipeline for intraoperative margin assessment in cutaneous squamous cell carcinoma (cSCC)
- Deep learning classification (CNN–GNN) for automated tumor detection
- Rigid registration (VALIS) for accurate alignment of serial sections
- 3D reconstruction for spatial visualization of tumor depth and continuity
- Takes Mohs surgery from a 2D, labor-intensive process into a semi-automated, 3D workflow
- Improves decision-making, reduces missed margins, and minimizes unnecessary tissue removal
- Entire pipeline executes in ~10 minutes, making it feasible for intraoperative integration.

INTRODUCTION

- Cutaneous squamous cell carcinoma (cSCC): second most common skin cancer worldwide
- Mohs micrographic surgery (MMS): gold standard treatment for high-risk tumors
  - Provides up to 99% margin assessment with maximal healthy tissue preservation
  - Recurrence rates as low as 1–3% vs. higher rates in standard excision
- Limitations:
  - Inherently 2D process, requiring surgeons to mentally reconstruct tumor depth
  - Risk of missed margins or overtreatment
  - Requires coordinated efforts of histotechnologists, dermatologists, and Mohs surgeons, making it labor-intensive and time-sensitive.
- The dataset includes: 95 H&E-stained frozen section slides from MMS, annotated by dermatologists and a Mohs surgeon, which includes both benign and malignant regions
- Existing tools remain 2D, limiting clinical utility for understanding tumor depth and spread

METHODS

- **Tissue preprocessing**:
  - Whole slide images (WSIs) scanned at 20× resolution
  - Tissue masks generated by down sampling and color thresholding (HSV space)
  - Morphological filtering to remove artifacts and fill holes
- **Patch extraction**:
  - Candidate tissue regions identified
  - Sliding window sampling → 256×256 pixel patches at full resolution
- **Patch classification (CNN–GNN)**:
  - CNN: ResNet-101 backbone trained to extract 2048-dimensional embeddings
  - Graph construction: patches treated as nodes, connected by spatial radius
  - GNN: Graph Attention Network (GAT) with multi-head attention, dropout, and normalization for classification
  - Outputs patch-level predictions of benign vs cSCC
- **Registration (VALIS)**:
  - Tissue islands identified and split into separate pyramidal WSIs
  - Rigid registration (translation, rotation, scaling) performed sequentially across sections.
  - Rigid-only approach chosen due to being significantly more stable than non-rigid in frozen section data
- **3D reconstruction**:
  - Tumor predictions converted to binary masks and stacked
  - Registration matrices applied to align excision islands across serial sections
  - Generated volumetric models with voxel dimensions: 0.46 µm (XY), 10 µm (Z)
  - Produced tumor depth maps and continuous 3D renderings of cSCC invasion

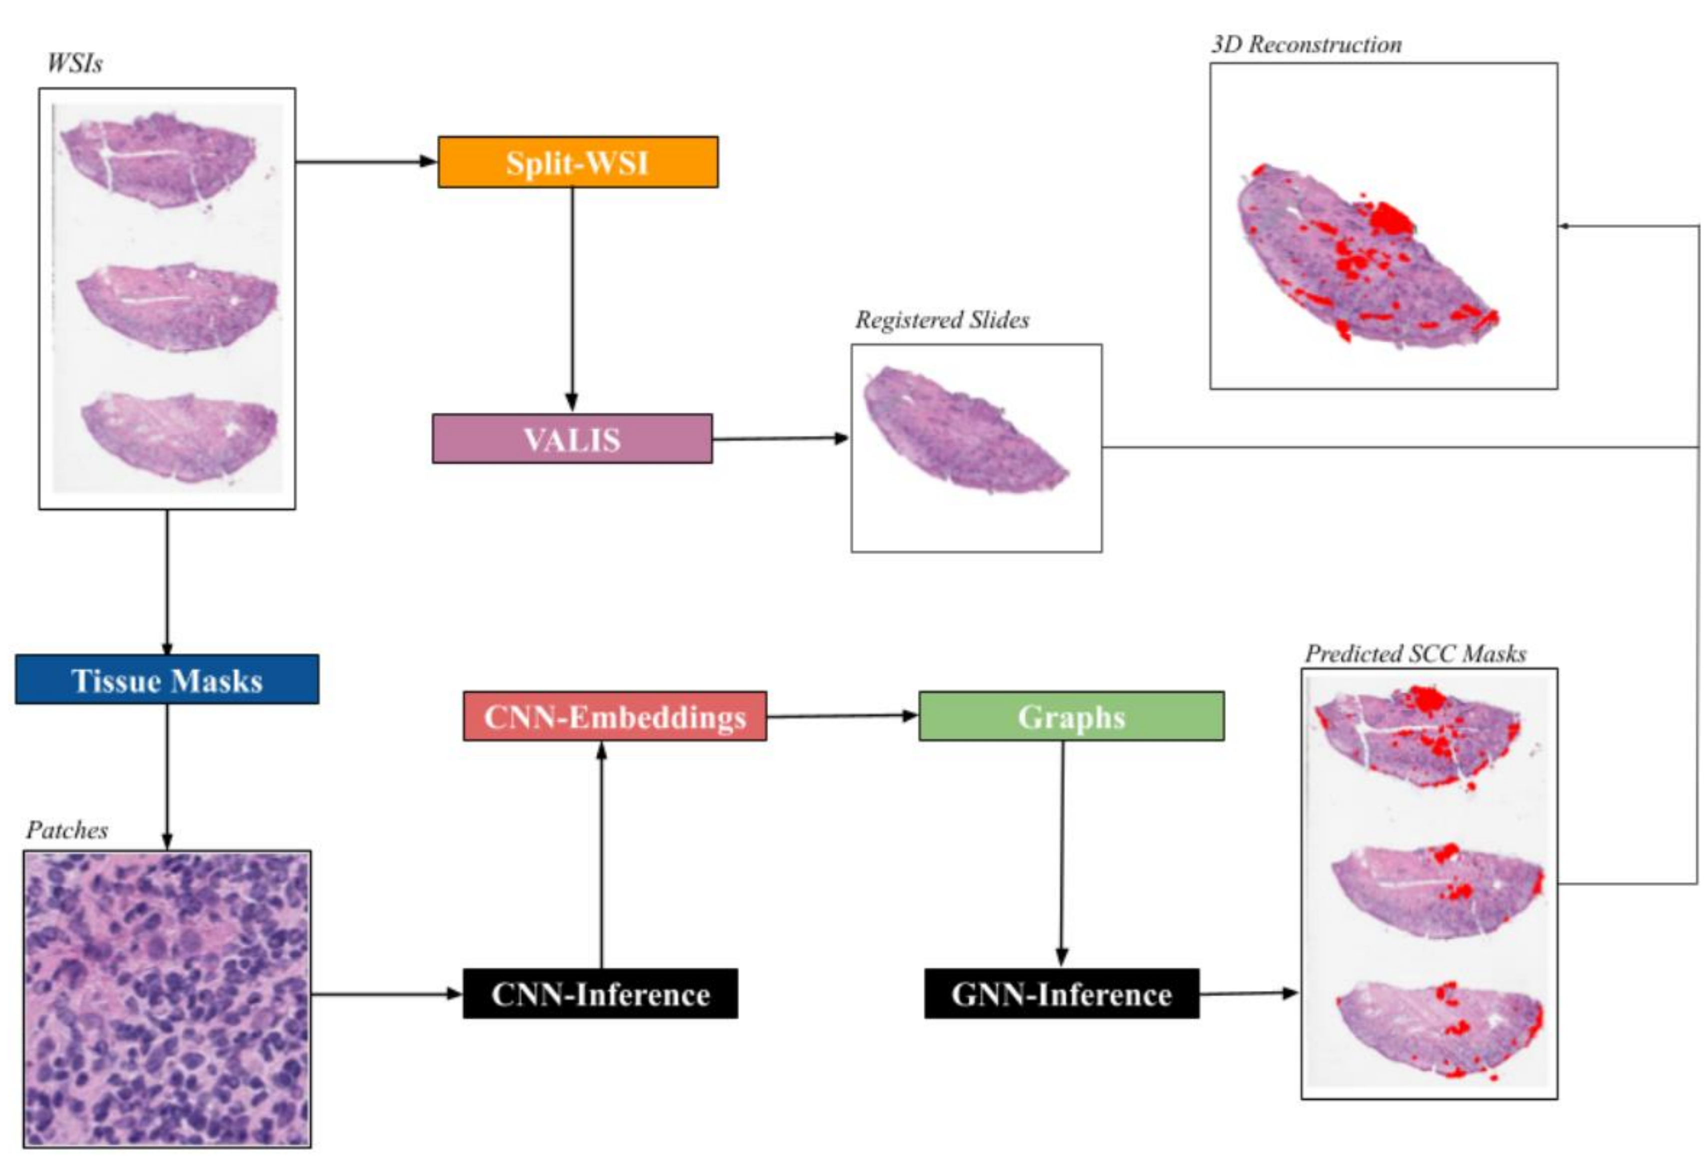

RESULTS

- **Patch-level classification**:
  - CNN baseline AUC = 0.918
  - CNN–GNN achieved **AUC = 0.955**, with higher precision, recall, F1-score across both benign and malignant classes
  - Demonstrates the value of spatial graph context for annotation

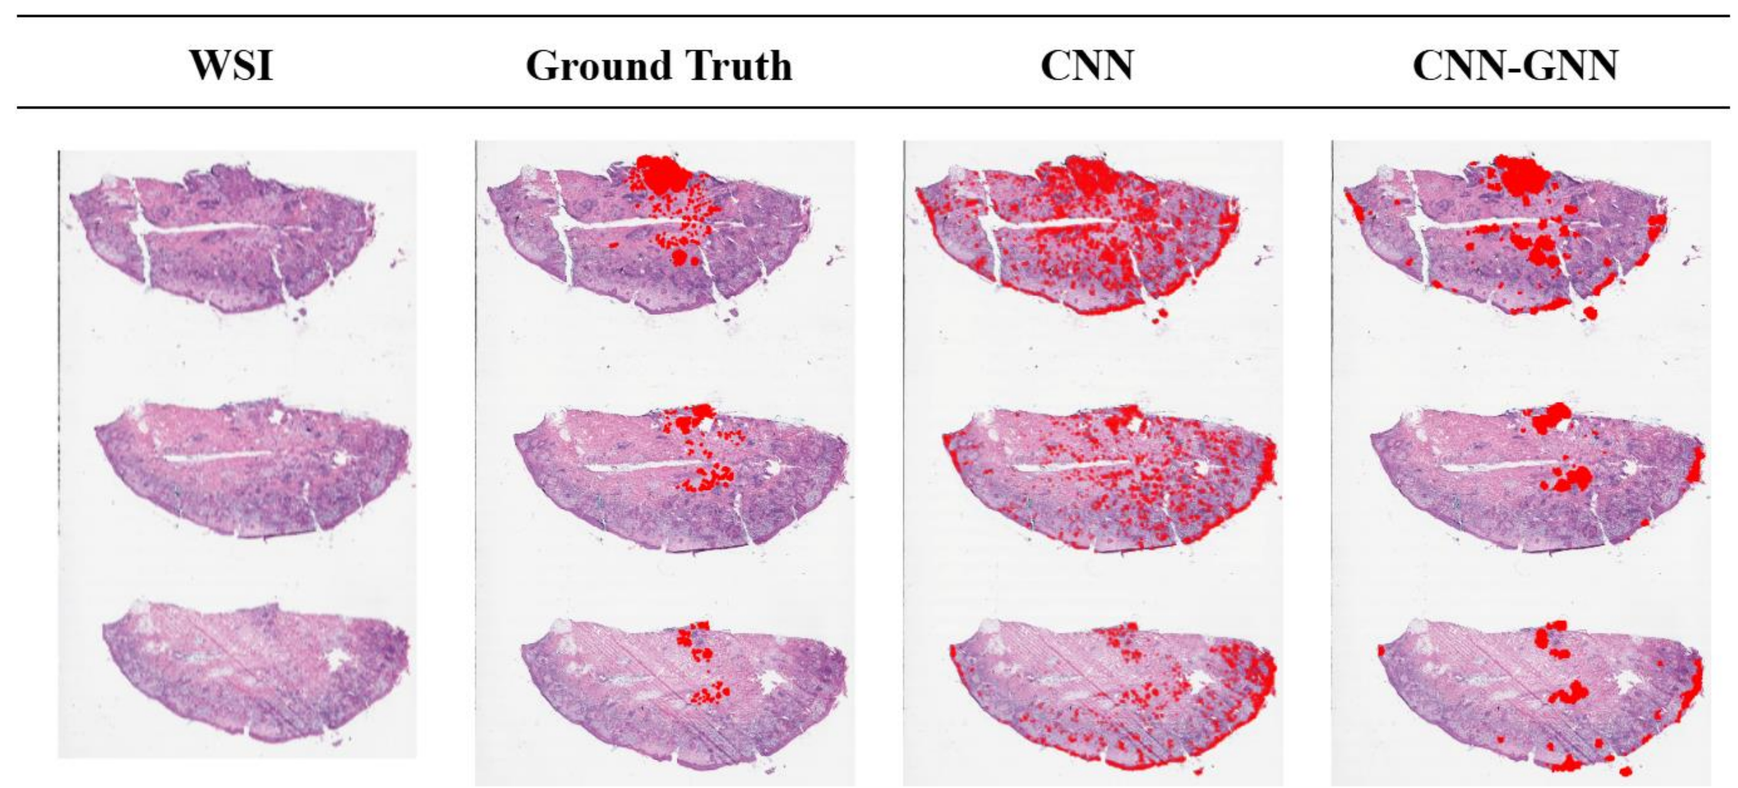

- Rigid registration reduced mean alignment error by **>95%** compared with baseline and non-rigid methods
- Rigid-only registration shown to be both more accurate and biologically faithful for Mohs SCC slides

Stack	*Baseline rTRE*	*Rigid rTRE*	*Improvement (%)*
61_A1a	0.73	0.002	99.69%
61_B1a	1.05	0.01	99.47%
345_a	1.08	0.01	99.32%
341_a	1.69	0.01	99.55%
345_b	0.97	0.01	99.03%
362_A1a	0.37	0.01	97.42%
341_b	2.71	0.01	99.63%
362_A1c	2.35	0.01	99.56%
362_A1b	1.79	0.01	99.38%
344_b	1.73	0.01	99.22%
344_a	2.57	0.01	99.46%
12_A1c	0.51	0.02	95.99%

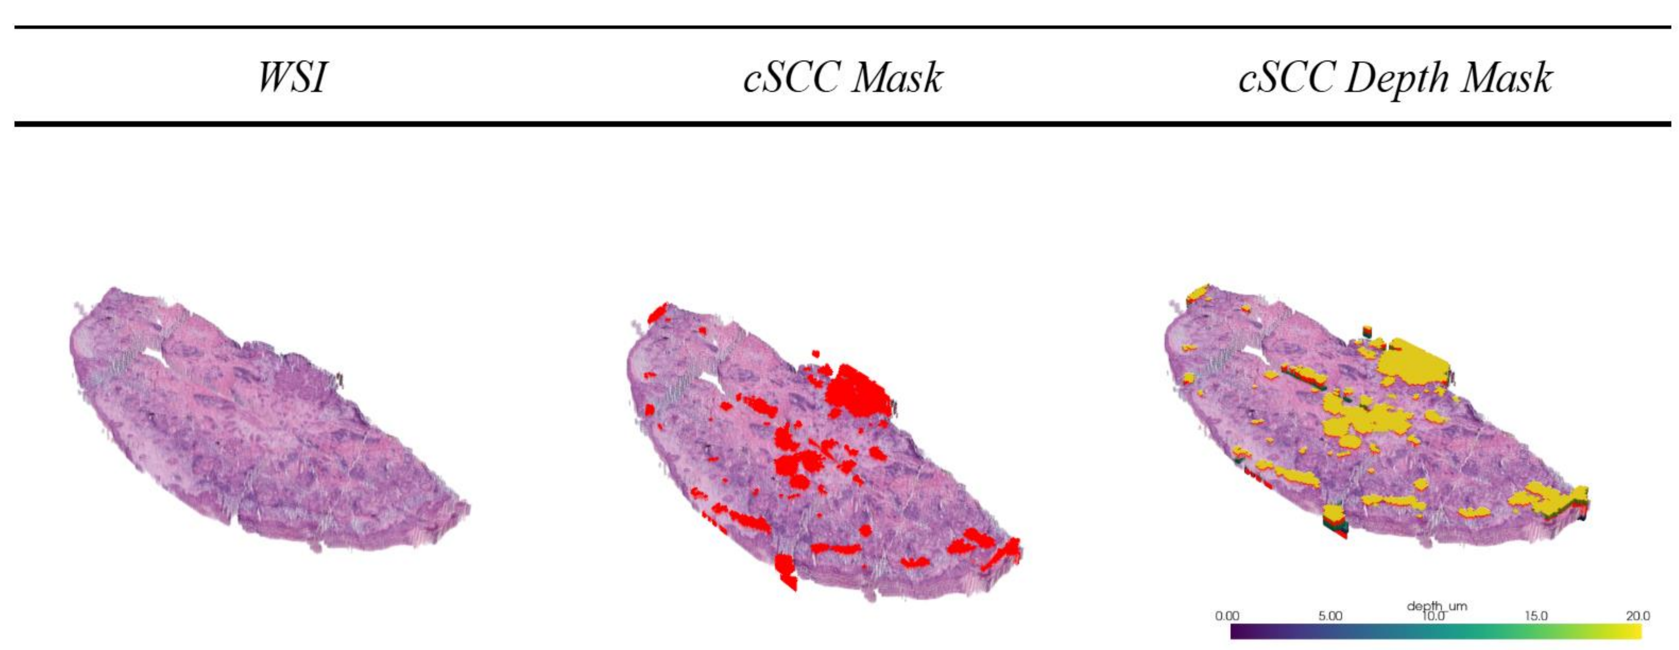

CONCLUSION

- Developed and validated an **AI-driven pipeline** for intraoperative annotation and 3D reconstruction of cSCC in Mohs specimens
- **CNN–GNN model** outperformed CNN baseline across all metrics (AUC = 0.955), highlighting the benefit of incorporating spatial graph context
- **Rigid registration** reduced error by **>95%** compared with baseline and non-rigid methods, enabling accurate alignment of serial frozen sections
- **Runtime ~10 minutes**, making the workflow **feasible for real-time intraoperative support**
- Provides **3D reconstructions** of tumor invasion depth and continuity, addressing the inherent limitations of traditional 2D frozen section review
- Potential to **reduce reliance on multiple intraoperative specialists** and to improve surgical precision by reducing both missed margins and overtreatment
- **Limitations**: single-institution dataset, short time window, and loss of true anatomical orientation in frozen specimens
- **Future directions**:
  - Incorporate orientation-preserving methods
  - Extend graph learning to **3D neighborhood structures** by leveraging registration, allowing edges to span across serial sections and enforce spatial continuity of tumor invasion across depth

edit.