Closed-Form GPU-Accelerated Tensor Voting with Refinement

Helya Goharbavang, Omar Baig, Guoning Chen, David Mayerich

University of Houston, Houston, TX

ISBI 2025
2025 IEEE International Symposium on Biomedical Imaging
April 14-17, 2025 | Houston, TX, USA

Poster No. A02-007

INTRODUCTION

Tensor fields like the structure tensor and Hessian extract local image features but are highly sensitive to noise due to finite difference computation. Smoothing via convolution reduces this noise but also degrades spatial resolution. **Tensor voting (TV)** offers structure-aware smoothing but is computationally intensive and lacks control.

Our work improves TV with:

- ✓ Tunable decay function for better control
- ✓ Normalized energy for iterative refinement.
- ✓ Closed-form solutions for 2D plate field
- ✓ GPU acceleration for large-scale processing
- ✓ Validated on vascular and neural biomedical images

OBJECTIVES

1. **Refine Tensor Voting**
   - Introduce parameters (σ_1, σ_2, p) to control vote spread and directionality.
   - Normalize energy to preserve field integrity.
2. **Derive Closed-Form Solutions**
   - Extend stick tensor decay function
   - Analytically solve plate tensor field
3. **GPU Acceleration**
   - Precompute eigendecompositions for speed.

Refinement and Normalization

We extend the closed-form N-D tensor voting formulation to support negative eigenvalues, enabling voting in Hessian-based fields. The stick tensor field is defined as:

$$\mathbf{V} = \mathrm{sign}(\lambda_1)(|\lambda_1| - |\lambda_0|)\mathbf{S} + \lambda_0 \mathbf{P}$$

We refine the stick vote field **S** using a direction-aware decay function:

$$\boldsymbol{D}(.) = e^{-\frac{t^2}{\sigma_1^2}}\left[1 - \left(\mathbf{q}^T\mathbf{d}\right)^2\right]^{\boldsymbol{p}} + e^{-\frac{t^2}{\sigma_2^2}}\left(\mathbf{q}^T\mathbf{d}\right)^{2\boldsymbol{p}}$$

This includes: (1) two standard deviations $\boldsymbol{\sigma_0}, \boldsymbol{\sigma_1}$ for lateral/axial control, (2) a power **p** to shape angular spread, and (3) a normalization term $\boldsymbol{\eta}$ to conserve total vote energy:

$$\eta = \left[\frac{\pi(2p)!}{2^{2p}(p!)^2}(\sigma_1^2 + \sigma_2^2)\right]^{-1}$$

Our final proposed formulation for the stick tensor is:

$$S(\mathbf{v}, \mathbf{r}, \mathbf{q}, \boldsymbol{\sigma}, p) = \eta(\boldsymbol{\sigma}, p) D(\mathbf{v}, \mathbf{r}, \mathbf{q}, \boldsymbol{\sigma}, p)(\mathbf{Rq})(\mathbf{Rq})^T$$

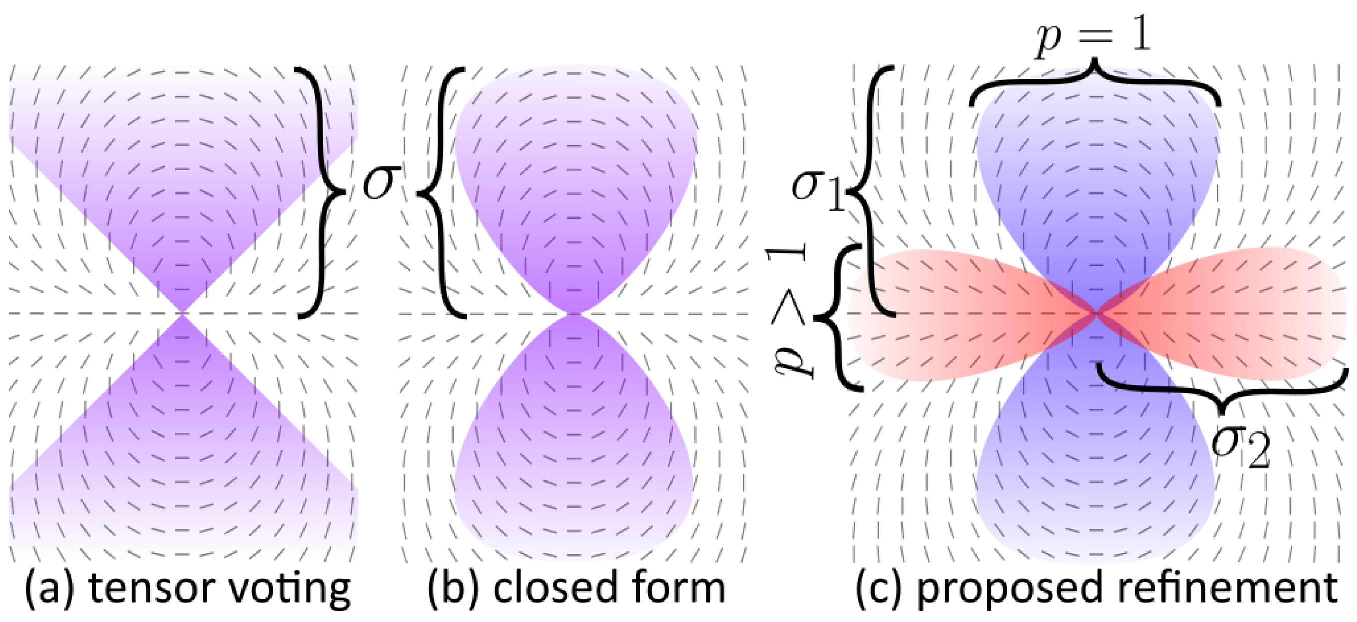

Fig. 1. Previous tensor voting fields compared to the pro-posed refined field

The plate tensor is calculated by integrating the stick tensor across all possible directions of **q**:

$$\mathrm{P}(\mathrm{v}, \mathrm{r}, \sigma) = \int_0^{\pi} \mathrm{S}(\boldsymbol{v}, \boldsymbol{r}, \boldsymbol{q}, \boldsymbol{\sigma}, \boldsymbol{p})\, d\beta$$

RESULTS

Qualitative Improvements:

1. **EM Images (Zebrafish Tectum):**
   - Original TV misaligns with fiber tracts (Fig. 3d)
   - Refined TV (p=6) preserves fiber continuity without resolution loss (Fig. 3f)
2. **Retinal Fundus Images:**
   - Original TV creates "halo“ effect near vessels (Fig. 4d)
   - Refined TV (p=10) aligns with vessel directions (Fig. 4e)

Quantitative Validation: $E = \min_{\tau} \sum_{\boldsymbol{x}} ||\tau \mathbf{G}(\mathbf{x}) - \mathbf{R}(\mathbf{x})||_2$

Using the normalized error metric (E), the proposed method reduces error vs. Gaussian blur and original TV (Fig. 2).

GPU Acceleration: Parallelized gather implementation achieves 300x speedup (Table 1)

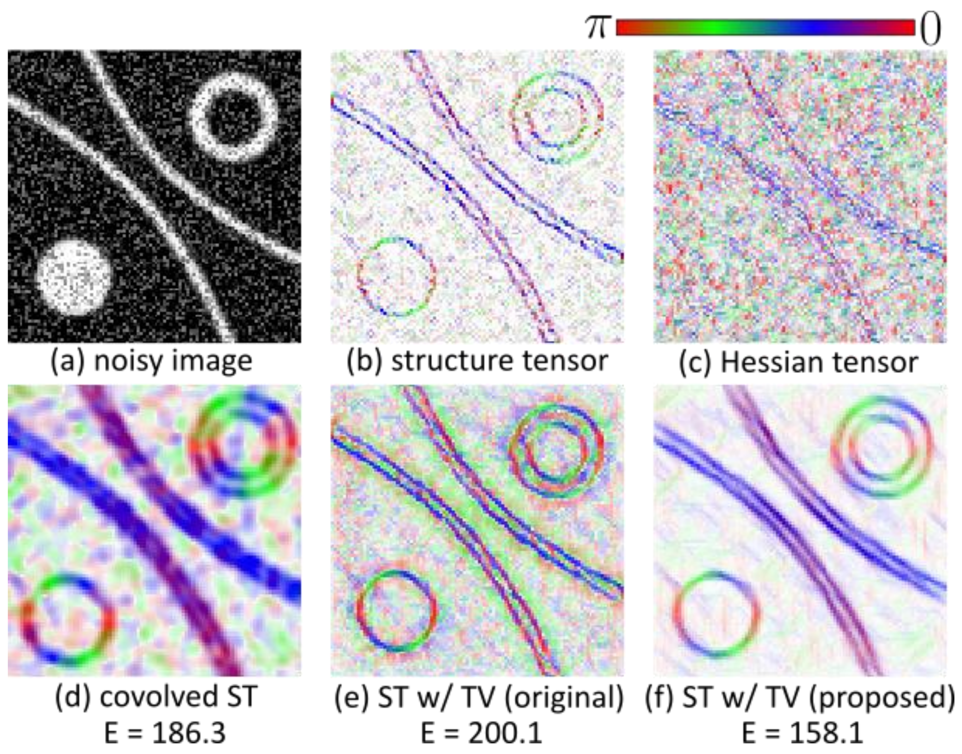

Fig. 2. Synthetic images and their corresponding tensor fields

GPU Speedup	Image Size	1000x1000	2000x2000
	σ_1	5	10
	GPU Time (ms)	218	2,837
	CPU Time (ms)	49,092	875,071
	Speedup	230x	310x

Table 1. Timing results using CPU and GPU ($\sigma_2 = 0$)

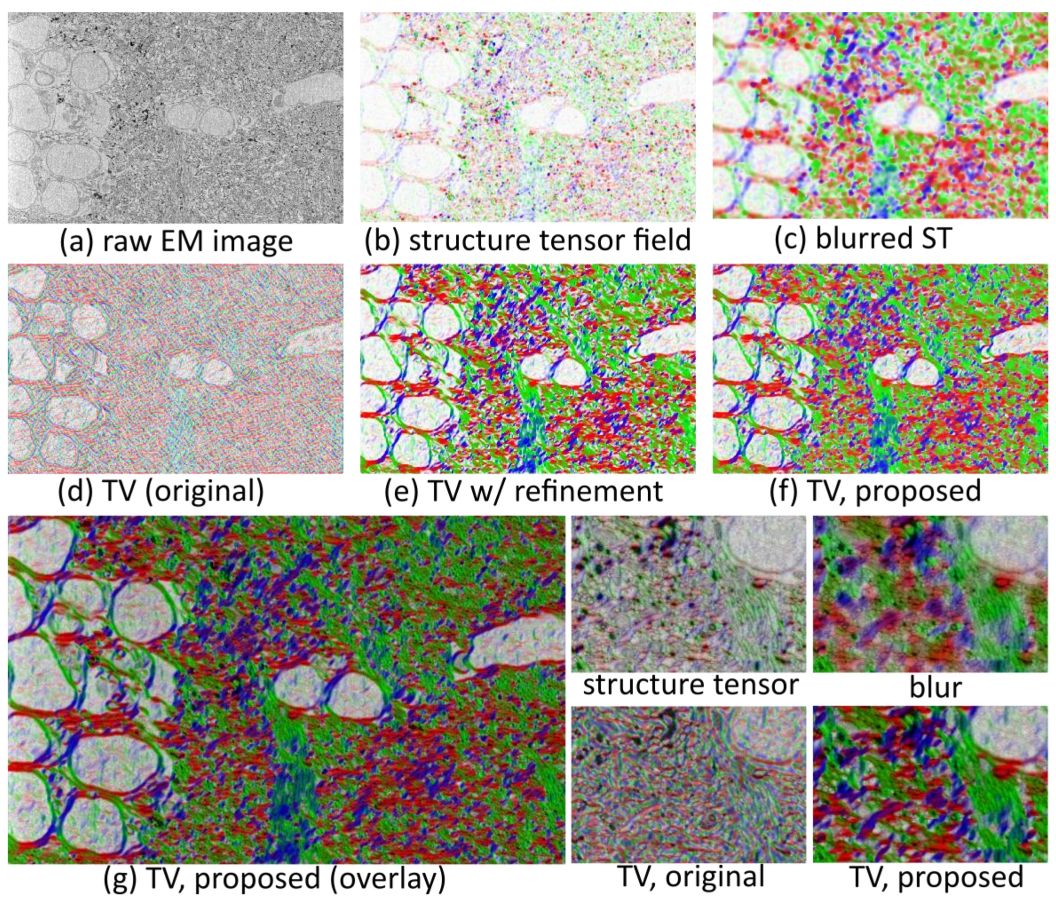

Fig. 3. Structure tensor field labeling fiber orientation in an EM image of zebrafish tectum

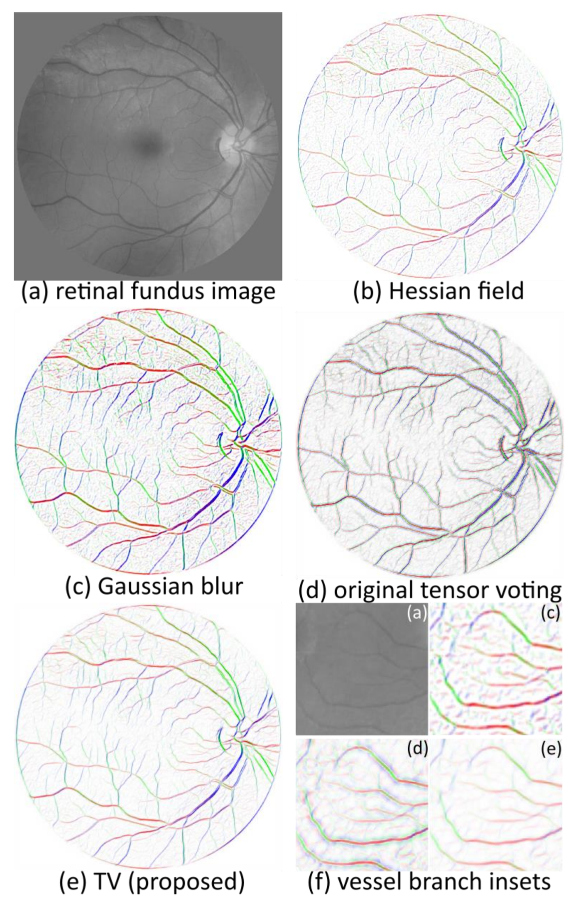

Fig. 4. Vessel orientations in a retinal fundus image

CONCLUSION

Key Contributions:

- Tunable, normalized TV for biomedical images
- Closed-form solutions for stick/plate tensors
- GPU implementation (300x faster)

Future Work:

- Extend to 3D for volumetric fiber tracing.
- 3D stick normalization and plate tensor derivation
- Integration of ball tensors to enhance 3D segmentation

REFERENCES

1. Mordohai & Medioni, Tensor Voting, 2006.
2. Wu et al., Closed-form tensor voting, IEEE TPAMI 2012.
3. Frangi et al., Multiscale vessel filtering, MICCAI 1998.
4. Lin et al., Analytical tensor voting, IEEE TPAMI 2023.

This work is funded in part by: National Institutes of Health / National Heart, Lung, and Blood Institute (NHLBI) #R01HL146745, National Science Foundation CAREER Award #1943455. DM is a stakeholder in SwiftFront, LLC.

UNIVERSITY OF HOUSTON