

W N R S · 2 0 2 6

From Pixels to Prognosis

What next-generation AI can already do for carotid plaque


● VW-MR

● CTA

● US

Abhi Jain, DO, CIIP


Clinical Assistant Professor · Neuroradiology · Stanford University



WNRS · 2026

Disclosures

I have no disclosures



We know what matters. We can't measure it at scale.



Composition adds what
stenosis misses

+2.3%

absolute rise in plaque burden with IPH.
Plaques without IPH showed minimal
progression over ~6 years.

On CTA, once lipid-core burden enters the
model, stenosis is no longer significant (HR
1.01, $p = 0.28$).



Measuring it is slow

20–30 min

to hand-quantify plaque components on a
single carotid CTA.

Head-and-neck vessel wall MRI: ~32 min to
read and analyze, ~105 mouse clicks per case.



Readers disagree

κ 0.49

inter-reader agreement on MRI plaque
vulnerability, unassisted (external cohort).

Six experienced sonographers vs. multimodal
truth: AUC only 0.72–0.80.


The bottleneck isn't knowledge — it's measurement.

What we're really asking AI to do



THE EVIDENCE LADDER

In 2026, multicenter AI on **MRI**, **CTA**, and **ultrasound** each reached the prognostic step.




Framework synthesized across the studies in this talk

MRI: a fully automated read of multi-sequence VWI

CFAPS — 1,315 patients · 1,610 arteries · 4 centers · 3 vendors · prospective PET/MRI arm

- Fully automated localization**
 ResNet50 finds the bifurcation; a U-Net centers a 128×128×30 box. Nothing is drawn by hand.
- One encoder per sequence**
 T1, T1C+, and TOF are encoded separately, then fused. Forcing one shared encoder dropped AUC 0.94 → 0.87.
- CNN + Transformer in every encoder**
 Conceptually, the CNN captures local texture; the Transformer integrates context across the plaque.
- Segment first, then classify**
 The plaque mask becomes the anatomical prior for the vulnerability head.



Gai et al., BMC Med 2026 — Fig. 2 (panels A–C) · CFAPS pipeline & architecture

CONVERGING EVIDENCE ON VESSEL WALL MRI

Head-and-neck automation Centerlines, CPR, and wall indices from carotid to intracranial — total read 32 → 13 min, clicks 105 → 12.

Anatomy as a prior Feeding the vessel wall in as a second channel lifts plaque Dice 78% → 83% (automatic) or 88% (manual).

CFAPS climbs every rung — all the way to stroke

Internal, pooled external (3 centers), and prospective test sets

TECHNICAL

AUC 0.92

Prospective (internal 0.94, external 0.88). Dice 0.86 ≈ human inter-observer 0.84.

BIOLOGICAL

κ 0.73

vs. endarterectomy histology (42 CEA patients). “Vulnerable” plaques are hotter on FDG-PET.

CLINICAL


72 s → 0.5 s

Inter-reader κ 0.49–0.57 → 0.72–0.75 with AI; junior-reader specificity +13 points.

PROGNOSTIC

HR 4.3–8.9

Ipsilateral stroke-free survival at two centers, 2–4 years of follow-up.



Gai et al., BMC Med 2026 — Fig. 7d–e · Kaplan–Meier stroke-free survival by CFAPS risk

Gai et al., BMC Med 2026 (CFAPS)


The catch

- The label is binary — any 1 of 4 features = “vulnerable.” No component volumes.
- Single-country cohort; histology in only 42 CEA specimens.
- Access: contrast, dedicated coil, long table time, expert reader.

→ *In the stroke bay, we get CTA.*

CTA: the workhorse becomes a composition scanner

2,164 CTAs · 29,113 annotated plaques · 9 centers



COMPONENTS ARE HU-DEFINED

Lipid core < 30 HU	Fibro-lipid 31–130	Fibrotic 130–350	Calcium > 350
------------------------------	------------------------------	----------------------------	-------------------------

76 s

per case, fully automated — vs. 20–30 min by hand

ICC 0.86–0.96

vs. radiologist measurements (fibrotic burden weaker, 0.67)

12×


higher AI-measured lipid-core volume in MRI-defined high-risk plaques; same direction vs. OCT

Figure: Zhao et al., Acad Radiol 2026, Fig. 3
See the published article

Zhao et al., Acad Radiol 2026 — Fig. 3 · CTA, AI component map, and matched OCT

Lipid-core burden predicts recurrence

610 symptomatic patients · median 19 months · 59 recurrent strokes / TIAs



1.08

Adjusted HR per percentage point of lipid-core burden — remained independently predictive

1.01

HR for stenosis after lipid core is added ($p = 0.28$)

13%

cut-off: KM HR 4.6 · sensitivity 0.68 · specificity 0.71

Figure: Zhao et al., Acad Radiol 2026, Fig. 4
See the published article

Zhao et al., Acad Radiol 2026 — Fig. 4 · Kaplan–Meier by lipid-core burden $\geq 13\%$


The catch

- IPH isn't reliably separable — its HU overlaps fibrotic tissue and contrast
- Retrospective; AUC 0.74 can't decide revascularization

→ **Prevention happens on ultrasound.**

Ultrasound: an AI that reasons like a sonographer

6,618 high-risk participants · 38,090 B-mode images · 4 centers



Multimodal ground truth

The “vulnerable” label was a multimodal consensus: B-mode findings corroborated by CEUS, CTA, or MR within 2 weeks. The model learns on routine B-mode to approximate that label.




Reasons arrive with the call

Vulnerability is only scored when plaque is present and feature evidence accumulates — so every call comes with named features a sonographer can check.

Outscores six sonographers and predicts stroke or TIA over 3 years

Vulnerability vs. multimodal consensus (internal test)



External AI AUC 0.88. Also flags vulnerability seen only on advanced imaging: AUC 0.84

The catch

- Incorporation bias — B-mode reads were part of the reference
- Events were overall stroke/TIA, not adjudicated as ipsilateral
- Still long-axis images; one country; fixed 0.5 cut-off

→ If it beats us, how do we know why?

Prognosis in held-out cohorts

2.09


adjusted HR, AI high-risk

0.795

C-index with the AI score
(0.769 without)

277

strokes / TIAs in 2,174
people over 37 mo



Xu et al., npj Digit Med 2026 — Fig. 4a · Event-free survival by AI risk group

Opening the black box: four ways these models earn trust



Saliency

where it looked

MRI Grad-CAM lands on IPH, ulcer, lipid core — and stays diffuse on stable fibrous plaque. But a heat map is not a reason.



Glass box

how it reasoned

The US cascade must name features first. 98.6% of vulnerable calls carried irregular morphology — yet the full model beat any single feature.



Measurement

what it measured

CTA outputs mm³ of lipid core, not just a probability. You can inspect it, quantify it, and audit it.




Biology

whether it's real


Model-vulnerable plaques match histology (κ 0.73) and are more inflamed on FDG-PET (TBR 2.67 vs 2.12).

weaker evidence

stronger evidence



Gai et al., BMC Med 2026 — Fig. 6 · Grad-CAM heatmaps with matched histology



Xu et al., npj Digit Med 2026 — Fig. 5a · Sankey: feature pathways to the vulnerability call

Three modalities, one pattern — and shared blind spots

	Built on	AI output	Taught by	Discrimination	Outcome signal	Biggest gap
MRI Gai 2026	1,315 pts · 4 centers + prospective arm	Plaque mask + vulnerable yes/no	Expert VWI read; checked vs histology & PET	AUC 0.88–0.94	HR 4.3–8.9 · ipsilateral stroke	Binary label; access
CTA Zhao 2026	2,164 CTAs · 9 centers; 610 followed	Component volumes & burdens	HU-defined labels; checked vs OCT & MRI	AUC 0.74 (with clinical)	HR 1.08 per point LCB · recurrent stroke/TIA	No IPH; retrospective
US Xu 2026	6,618 pts · 4 centers; 2,174 followed	5 named features → vulnerability	Multimodal consensus: CEUS, CTA, MR	AUC 0.88–0.90	Adj. HR 2.09 · overall stroke/TIA	Incorporation bias; laterality adjudicated

Notice the teacher: HR-VWI validated the CTA model; CEUS, CTA, and MR helped adjudicate the ultrasound labels. Advanced imaging is becoming the teacher for routine imaging.



One country

All three anchor cohorts come from China. We need diverse, Western validation.



We're still the truth

Labels are mostly radiologist reads; histology in 42 CEA specimens, OCT in 16 patients.



Association ≠ intervention

No one has randomized AI-guided vs. standard management.




Endpoints don't match

Ipsilateral stroke vs. recurrent stroke/TIA vs. overall stroke/TIA — no head-to-head yet.

Where this is heading

From snapshot to trajectory


DL segmented wall and IPH on ~5 VWI scans in 28 asymptomatic patients over ~6 years. No IPH: minimal progression. IPH: +2.3% burden, narrowing lumen. New IPH arose in 11 of 39 arteries and persisted.



Guo et al., JCMR 2026 — Fig. 6 · Serial MERGE/SNAP: IPH appears (scan 3), then grows

AI before the read

On intracranial VWI, DL denoising + super-resolution took in-plane resolution from 0.51 to 0.28 mm. IPH sensitivity for culprit plaques rose 24% → 71% at 94% specificity — without an extra acquisition.



Seo et al., Sci Rep 2024 — Fig. 4 · Conventional vs. DL-augmented VWI of culprit plaques



Advanced imaging teaches US

Multimodal truth at ultrasound scale.



Open weights

CFAPS weights + US and VWI code are public — validate on our own patients.



The missing rung

IMPROVE turns MRI IPH into 3-yr stroke risk (C 0.82). AI can standardize that input; then randomize.

TAKE-HOME

Three things to remember

1

Composition is finally measurable

In seconds, reproducibly, on the modalities we already use.

2

2026: plaque AI reached outcomes

MRI, CTA, and ultrasound models now each link their output to future cerebrovascular events.

3

The frontier isn't a better Dice

It's diverse populations, longitudinal tracking, feeding risk models like IMPROVE, and trials of AI-guided care.



References

- Gai Q, Zhang Y, Liu J, et al. A fully automated deep learning pipeline system for carotid plaque segmentation and vulnerability classification on multi-contrast MRI: a multicenter study. *BMC Med.* 2026. doi:10.1186/s12916-026-05199-8
- Zhao MJ, Cheng XQ, Xu PP, et al. Automated analysis of carotid plaque components in CT angiography: performance and prognosis in symptomatic carotid stenosis patients. *Acad Radiol.* 2026. doi:10.1016/j.acra.2026.06.066
- Xu Q, Ding H, Zhou L, et al. AI-derived carotid plaque vulnerability from B-mode ultrasound for cerebrovascular risk stratification: a multicenter cohort study. *npj Digit Med.* 2026. doi:10.1038/s41746-026-03201-4
- Guo Y, Akcicek EY, Hippe DS, et al. Long-term carotid plaque progression and the role of intraplaque hemorrhage: a deep learning-based analysis of longitudinal vessel wall imaging. *J Cardiovasc Magn Reson.* 2026;28:102670. doi:10.1016/j.jocmr.2025.102670
- Fu F, Lin Z, Yang X, et al. Deep learning for high-resolution MR vessel wall imaging: image reconstruction, stenosis diagnosis and plaque calculation. *Eur Radiol.* 2026;36:5169–5181.
- Yang L, Yang X, Gong Z, et al. Deep learning-based automatic segmentation of arterial vessel walls and plaques in MR vessel wall images for quantitative assessment. *Eur Radiol.* 2025;35:7712–7726.
- Seo M, Jung W, Jeong G, et al. Deep learning improves quality of intracranial vessel wall MRI for better characterization of potentially culprit plaques. *Sci Rep.* 2024;14:18983.
- Nies KPH, Smits LJM, van Kuijk SMJ, et al. Individualized MRI-based stroke prediction score using plaque vulnerability for symptomatic carotid artery disease patients (IMPROVE). *Stroke.* 2025;56:2068–2078. doi:10.1161/STROKEAHA.124.050020
- Ma D, Chen J, Li Z, et al. Carotid plaque vulnerability assessment and stroke risk prediction based on radiomics and deep learning: a systematic review and meta-analysis. *BMC Med Imaging.* 2026. doi:10.1186/s12880-026-02824-z
- Ju L, Guo Y, Guo H, et al. Diagnostic performance of deep learning and radiomics in extracranial carotid plaque detection: systematic review and meta-analysis. *J Med Internet Res.* 2026;28:e77092. doi:10.2196/77092