

ICML 2026


Geometrically Constrained Stenosis Editing in Coronary Angiography via Entropic Optimal Transport

OT-Bridge Editor — pixel-level geometric control for synthetic coronary stenosis

Jialin Li · Zhuo Zhang · Yue Cao · Guipeng Lan · Jiabao Wen · Xiao Shuai · Jiachen Yang

School of Electrical and Information Engineering, Tianjin University

github.com/LiJialin001/OTBridge-Editor



Stenosis = boundary geometry

the diagnostic label lives in the pixels


MOTIVATION

Why pixel-level geometry — not just “where”

In angiography, reading a stenosis is a geometric measurement.

- Percent diameter narrowing (%DS) is read off the vessel boundary, pixel by pixel.
- Detectors learn the label from that exact geometry — a wrong boundary is a wrong label.
- Most controllable generation today controls semantics & position: “is there a lesion, and roughly where.”

Our goal: pixel-level control of lesion shape & location, while the surrounding anatomy stays untouched.



The red target box is what we must reproduce exactly — boundary, not just region.

Why conditional diffusion isn't enough

1 Soft guidance

Conditions enter as guidance terms during denoising — they nudge, they don't lock geometry. Boundary alignment is not guaranteed.

2 Starts from noise

The reverse process is initialized from pure noise, even though the input already holds reliable anatomy outside the edit region — so it struggles to preserve non-target content and to localize the change.

The operation we want is not to rebuild the image from noise — it's to transport the source toward a target that differs only locally.

KEY IDEA

Reframe editing as transport, not reconstruction

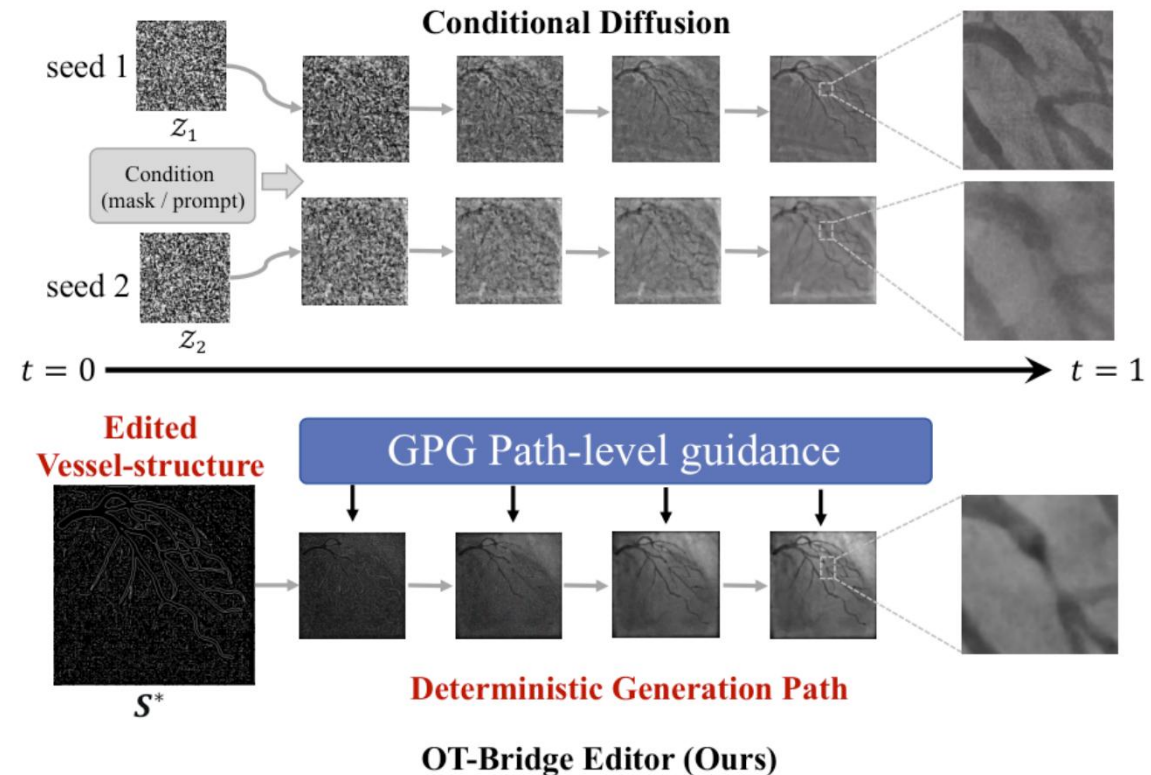
Move probability mass from the source — anchored at the input — to a target restricted by a **geometric feasibility set**, minimizing unnecessary change.

Conditional diffusion (top)

samples from noise — different seeds drift apart.

OT-Bridge (bottom)

fixes the start state S^* and follows a deterministic generation path to the target.



Conceptual comparison (paper Fig. 1).

From optimal transport to a Schrödinger Bridge

Classical OT gives a coupling, not a generative path — and under strong geometric constraints in high dimensions it is numerically fragile.


Schrödinger Bridge: lift to path space; find the trajectory closest in KL to a reference diffusion, matching both endpoints → a deterministic, steerable path.

Outside the mask
strict preservation

Inside the mask
boundary matches spec

feasibility set \mathcal{F}

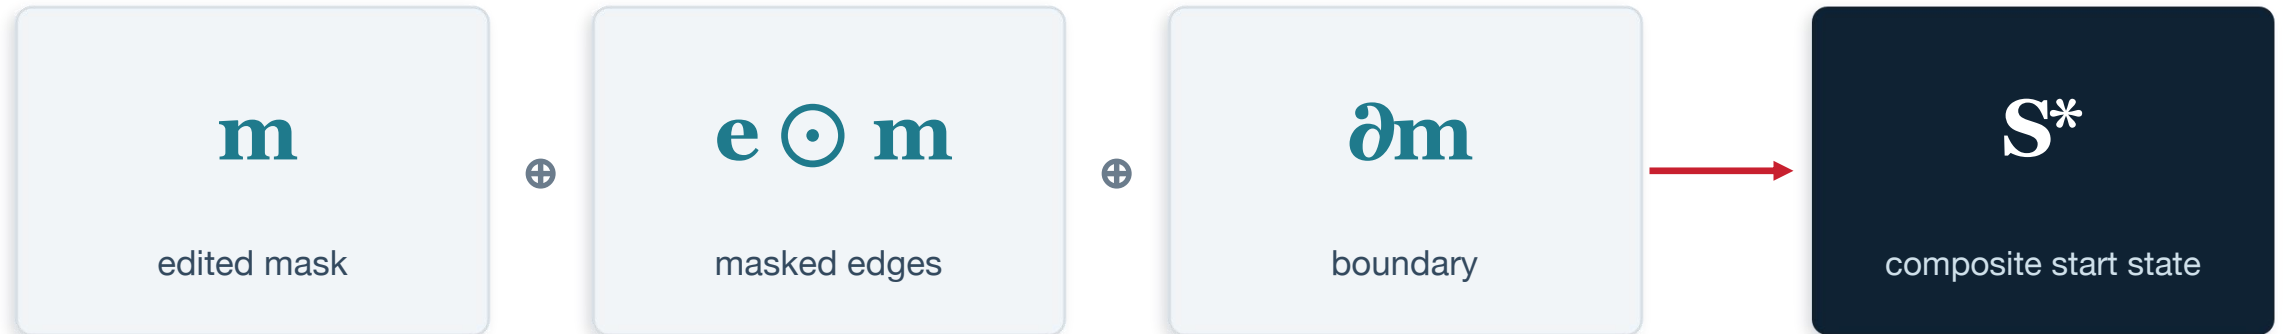
Deterministic generation path (paper Fig. 11): structure → angiogram, topology preserved.



Editing in a Vessel-structure Composite Domain

We don't edit raw pixels. The start state stacks the edited mask with its geometric cues:


$$S^* = \text{Stack}(m, e \odot m, \partial m)$$



Editing in a representation aligned with vascular morphology makes the geometry far easier to enforce.

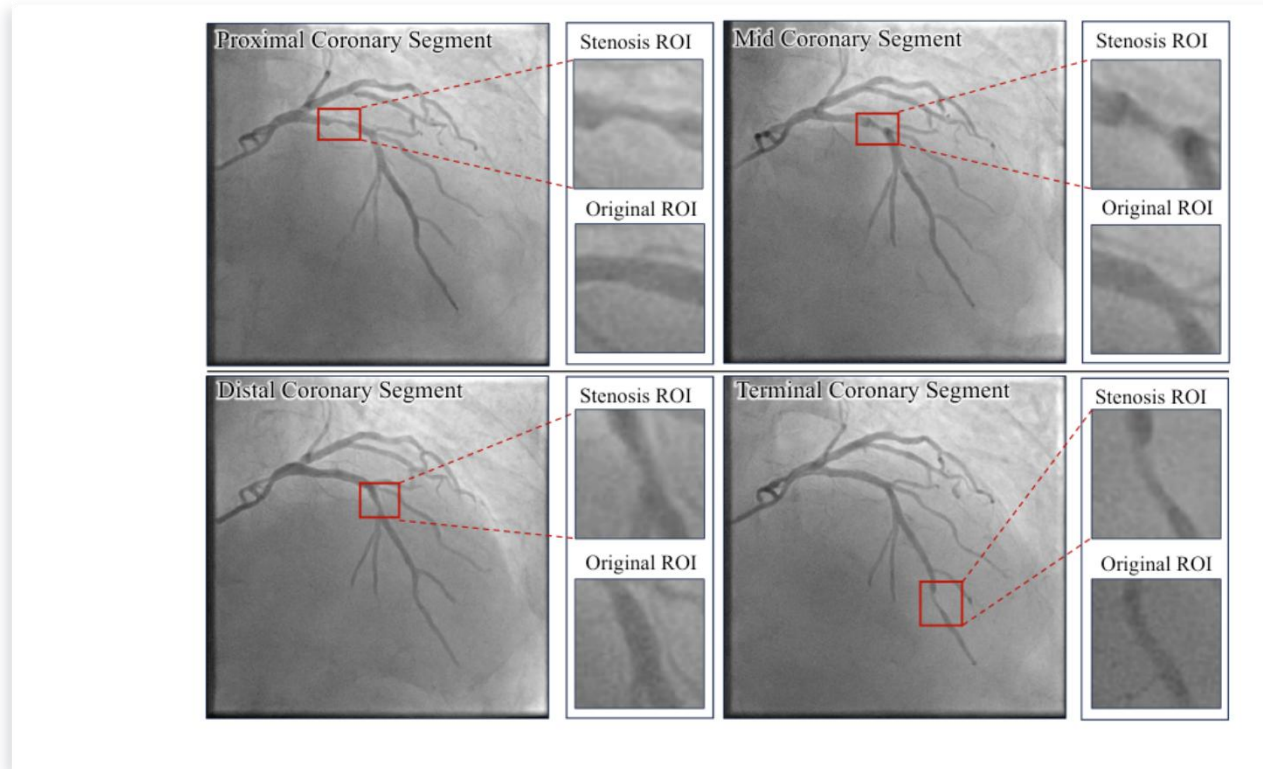
Geometric Generation-Path Guidance (GPG)

Enforce geometry at every bridge step — not only at the endpoint. After each transition we project onto the feasibility set $\Pi_{\mathcal{F}}$: pull toward the target boundary inside the mask, back toward original anatomy outside. Boundary discrepancy via a signed distance transform (stable).



Endpoint-only conditioning → path-level supervision. This is what prevents drift on thin vessels.

Pixel-accurate edits, anatomy preserved



Location-controlled stenosis across coronary segments (paper Fig. 4) — minimal drift outside the target ROI.

0.892

mIoU vs target mask

0.774

Dice

16.7

FID (best realism)

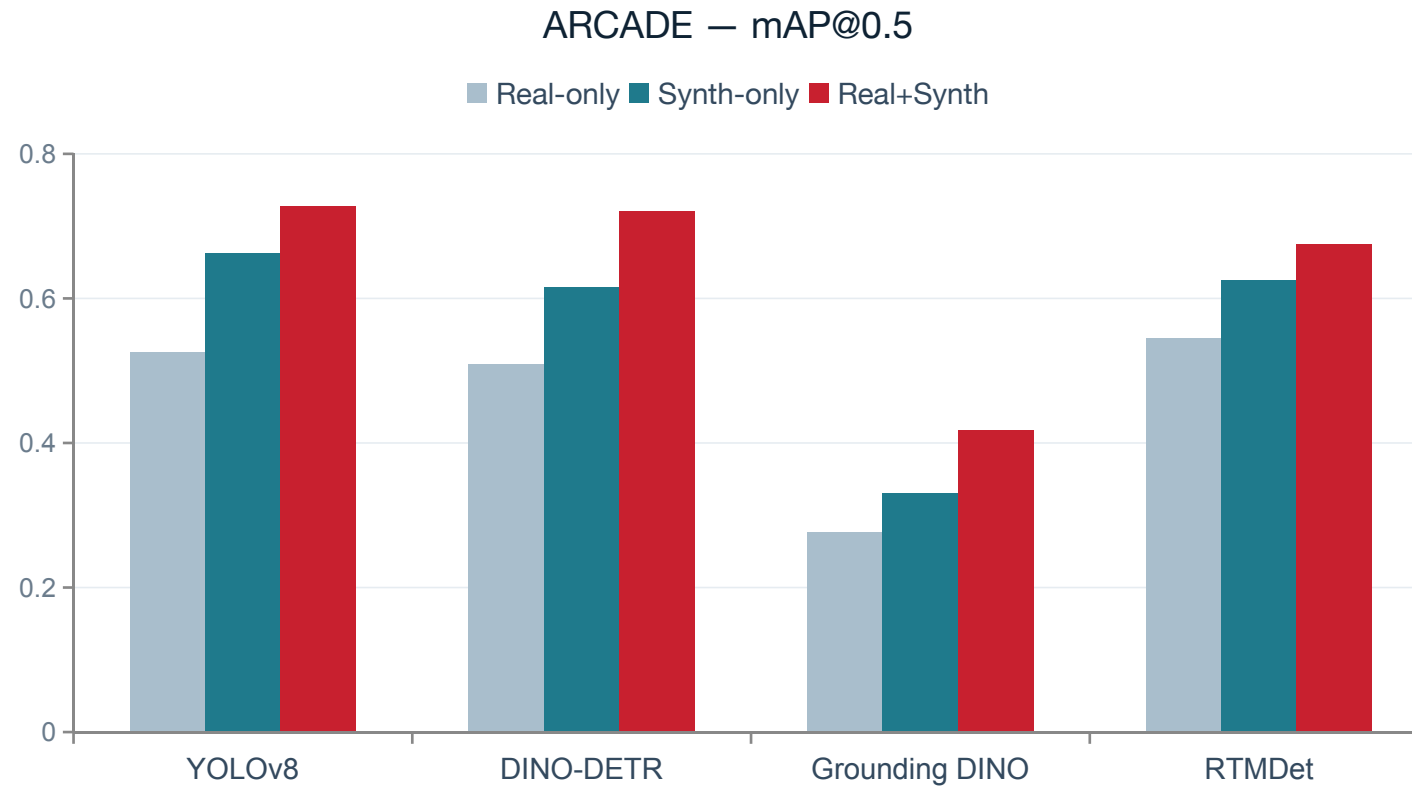
0.878

SSIM

Highest geometric compliance and best realism

vs GAN/diffusion baselines (best baseline SiameseDiff: mIoU 0.837, FID 34.2).

Synthetic data lifts every detector



+27.8%

relative mAP on public ARCADE

+23.0%

relative on multi-center dataset

Real+Synth is best across all four detectors.

Ablations: composite domain + path-level GPG each matter; GPG gives the lowest, smoothest boundary error.

CONCLUSION

OT-Bridge Editor

*Reframe localized editing as constrained optimal transport, and guide the entire generation path — **pixel-level geometric control that soft-conditioned diffusion can't guarantee.***

Constrained entropic OT solved via a Schrödinger Bridge

Path-level guidance (GPG) for pixel-accurate boundaries

Clinically realistic data → measurable detection gains

Code: github.com/LiJialin001/OTBridge-Editor

Thank you.