

Iris-3B: Going Beyond the Latent with Pixel-Space Diffusion Training, Conversion and Fine-Tuning

Original summary — NOT the original paper or model card.

Research preprint · 2026-10-07 [28]

Authors/source: Hanqiu Li Cai, Chema Garabito

Coverage: Video topic

What the source reports

A three-billion-parameter pixel-space image generator, with experiments in depth estimation and restoration. [28]

Managing expectations

The paper reports no significant downstream improvement from a pixel-space generative prior; neither pixel model beats the compared latent fine-tune on the reported restoration experiment. This negative result matters more than gallery impressions. [28]

Original document and sharing status

An unmodified original paper is included separately under CC0 1.0. This brief does not replace it.

Recorded rights: CC0 1.0 [30]

Sources and downloads

[Primary source / manuscript details \[28\]](#)

[Official original document / PDF download](#)

[Hosted, unmodified original PDF](#)

[Complete document library](#) · [Read the research roundup](#)

Source briefs are Managing Expectations editorial summaries, not independent replications, peer review, or medical advice. Source claims and release details may change.