

DistScene: Object-to-Scene Distillation for 3D Scene Generation

Original summary — NOT the original paper or model card.

primary preprint identified · 2026-10-03 [26]

Authors/source: Luo, Kunming, Yan, Hongyu, Deng, Ken, Zhou, Chengcheng, Liu, Tianyu, Li, Haipeng et al. (full author list in manifest)

Coverage: Video topic

What the source reports

Single-image compositional 3D scene generation jointly represents environment and objects, then refines objects using scene context; trains using automatically assembled synthetic scenes distilled from an object generator. [26]

Managing expectations

Indoor results may omit ceilings because training environment assets often omit them. Large outdoor scenes can have holes and lose fine geometry due to finite voxel resolution. Outdoor training corpus is smaller due to compute constraints; project lists code, checkpoints and dataset as coming soon. [26]

Original document and sharing status

The original source remains available at the official links below. Full-document redistribution permission was not established, or this topic has no standalone paper. Only this original summary is reproduced here.

Recorded rights: arXiv perpetual non-exclusive distribution license [26]

Sources and downloads

[Primary source / manuscript details \[26\]](#)

[Official original document / PDF download](#)

[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.