Zhiwei ZHU, Xiang GAO, Lu YU, Yiyi LIAO, 2025. Neural mesh refinement. *Frontiers of Information Technology & Electronic Engineering*, 26(5):695-712. https://doi.org/10.1631/FITEE.2400344

Neural mesh refinement

Project page: https://zhuzhiwei99.github.io/NeuralMeshRefinement/

Key words: Geometry processing; Mesh refinement; Mesh subdivision; Disentangled representation learning; Neural network; Graph attention

Corresponding authors: Lu YU, Yiyi LIAO
E-mail: yul@zju.edu.cn; yiyi.liao@zju.edu.cn
ORCID: Lu YU, https://orcid.org/0000-0002-0550-7754;
Yiyi LIAO, https://orcid.org/0000-0001-6662-3022

Motivation

Detailed 3D meshes can model complex details but put a strain on storage and transmission. Coarse 3D meshes are favored for efficient storage and quick transmission but struggle to accurately represent fine geometric details. To tackle this trade-off, a practical solution is to store and transmit a coarse mesh and then utilize a **refinement** method at the user end to restore the details.

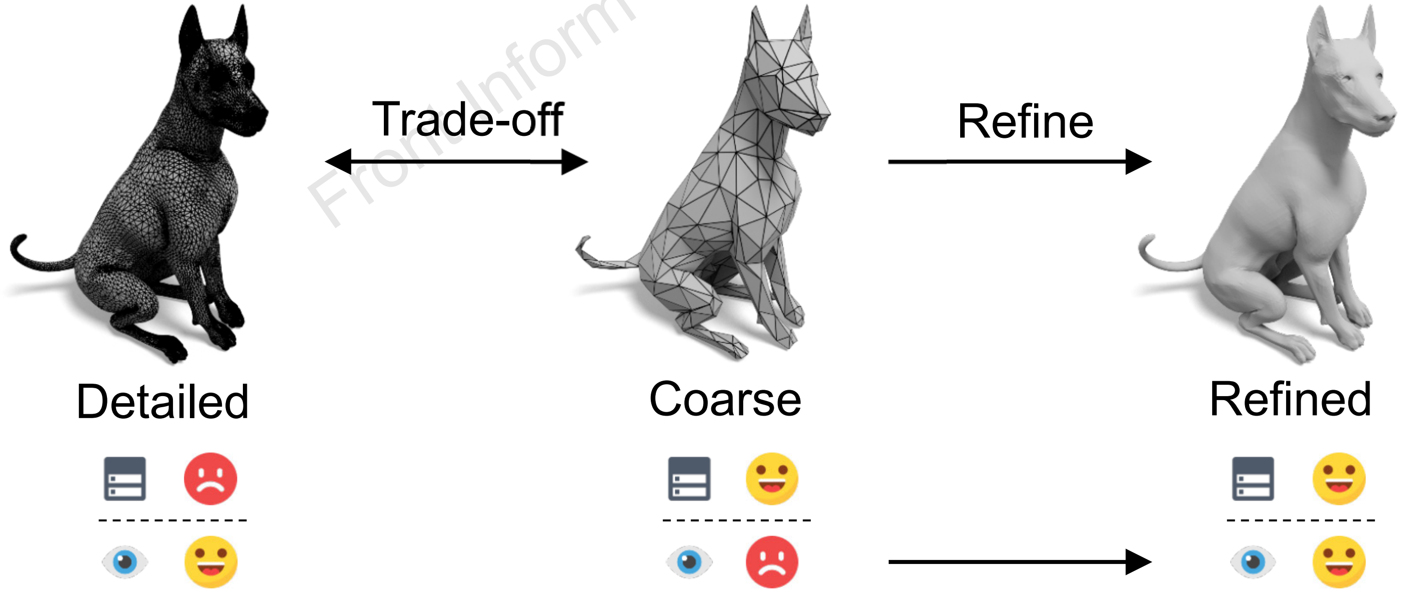

Motivation

Subdivision is a widely used technique for mesh refinement. Classic methods rely on fixed manually-defined weighting rules and struggle to generate a finer mesh with appropriate details, while advanced neural subdivision methods achieve data-driven nonlinear subdivision but lack robustness, suffering from limited subdivision levels and artifacts on novel shapes.

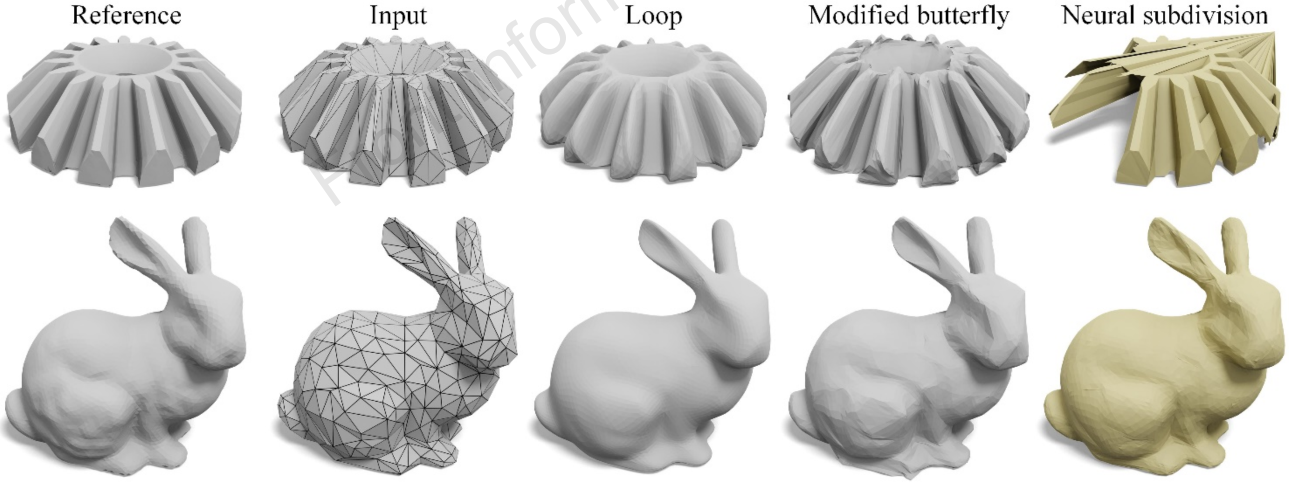

Existing subdivision methods

Main idea

1. To address these issues, we introduce a neural mesh refinement (NMR) method that uses the geometric structural priors learned from fine meshes to adaptively refine coarse meshes through subdivision, resulting in refined meshes with more reasonable details and exhibiting robust generalization.

2. Our key insight is that it is necessary to **disentangle the network from non-structural information** such as scale, rotation, and translation, enabling the network to focus on learning and applying the structural priors of local patches for adaptive refinement. For this purpose, we propose an intrinsic structure descriptor and a locally adaptive neural filter.

Overview

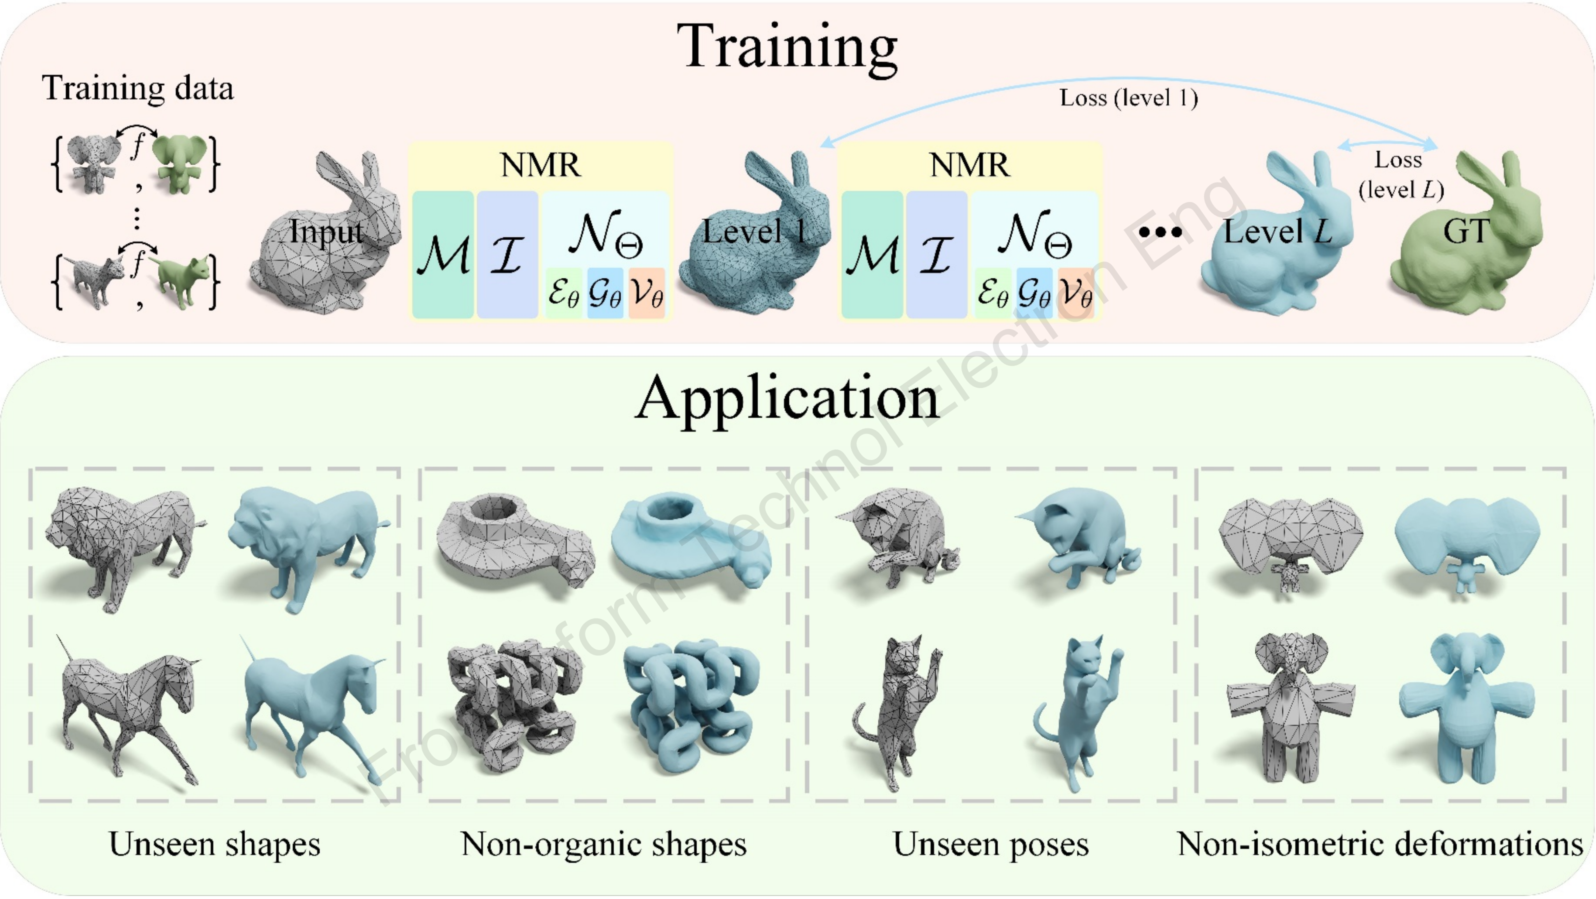

Neural mesh refinement (NMR) performs data-driven nonlinear refinement and demonstrates robust generalization to unseen shapes, poses, and non-isometric deformation. It can also refine coarse non-organic shapes into finer ones with appropriate geometric details, even when trained on organic shapes.

Method

NMR processes a coarse triangle mesh (gray) to produce a sequence of subdivided meshes (blue) with varying levels of details. During training, we minimize the Charbonnier loss between the ground truth (green) and the output meshes (blue) across levels. Our training data comprise pairs of coarse and fine meshes (left) with a bijective map between each pair. Each refinement involves the processing of three modules: **(1) midpoint subdivision; (2) intrinsic structure descriptor; (3) neural filter.**

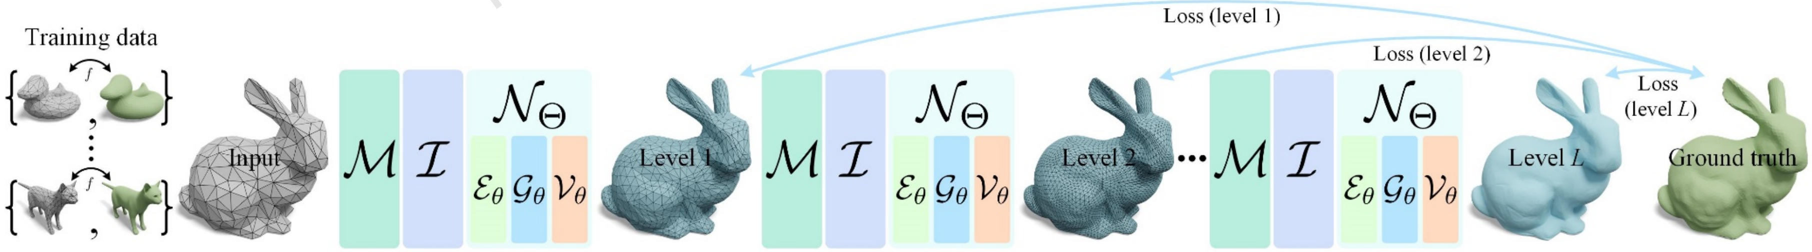

Refinement pipeline of NMR

Method

1. Midpoint subdivision: The original face is subdivided into four faces by inserting new vertices at the midpoint of each edge and connecting new vertices through new edges.

2. Intrinsic structure descriptor: For newly inserted vertices, an intrinsic structure descriptor is created by normalizing their one-ring neighborhood. This normalization makes the features invariant to scale, rotation, and translation, allowing the network to focus on extracting structural priors.

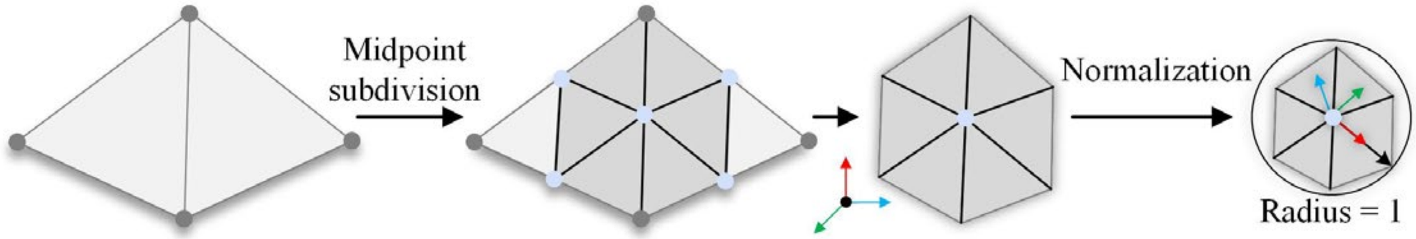

Midpoint subdivision and intrinsic structure descriptor

Method

3. Neural filter: This filter predicts offsets for the new vertices to update their positions. It consists of:

- edge feature embedding: capture features from the outgoing edges of the central vertex;
- graph attention aggregation: aggregate edge features to the central vertex using a graph attention mechanism, adaptively emphasizing relevant local structures;
- vertex repositioning: predict residual from the central vertex feature to refine its position.

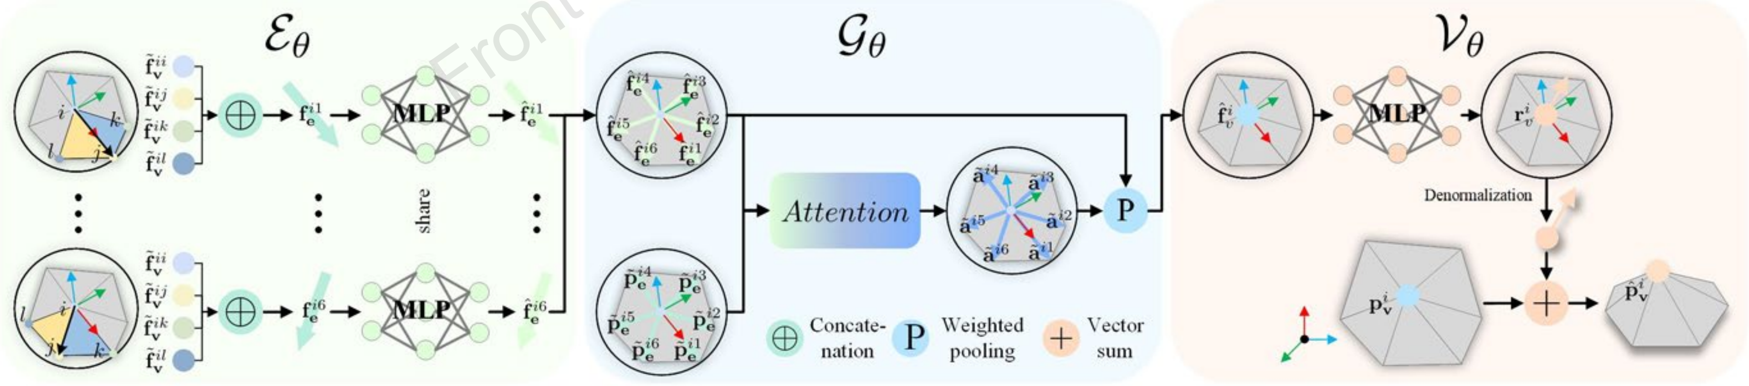

Architecture of the neural filter

Major results

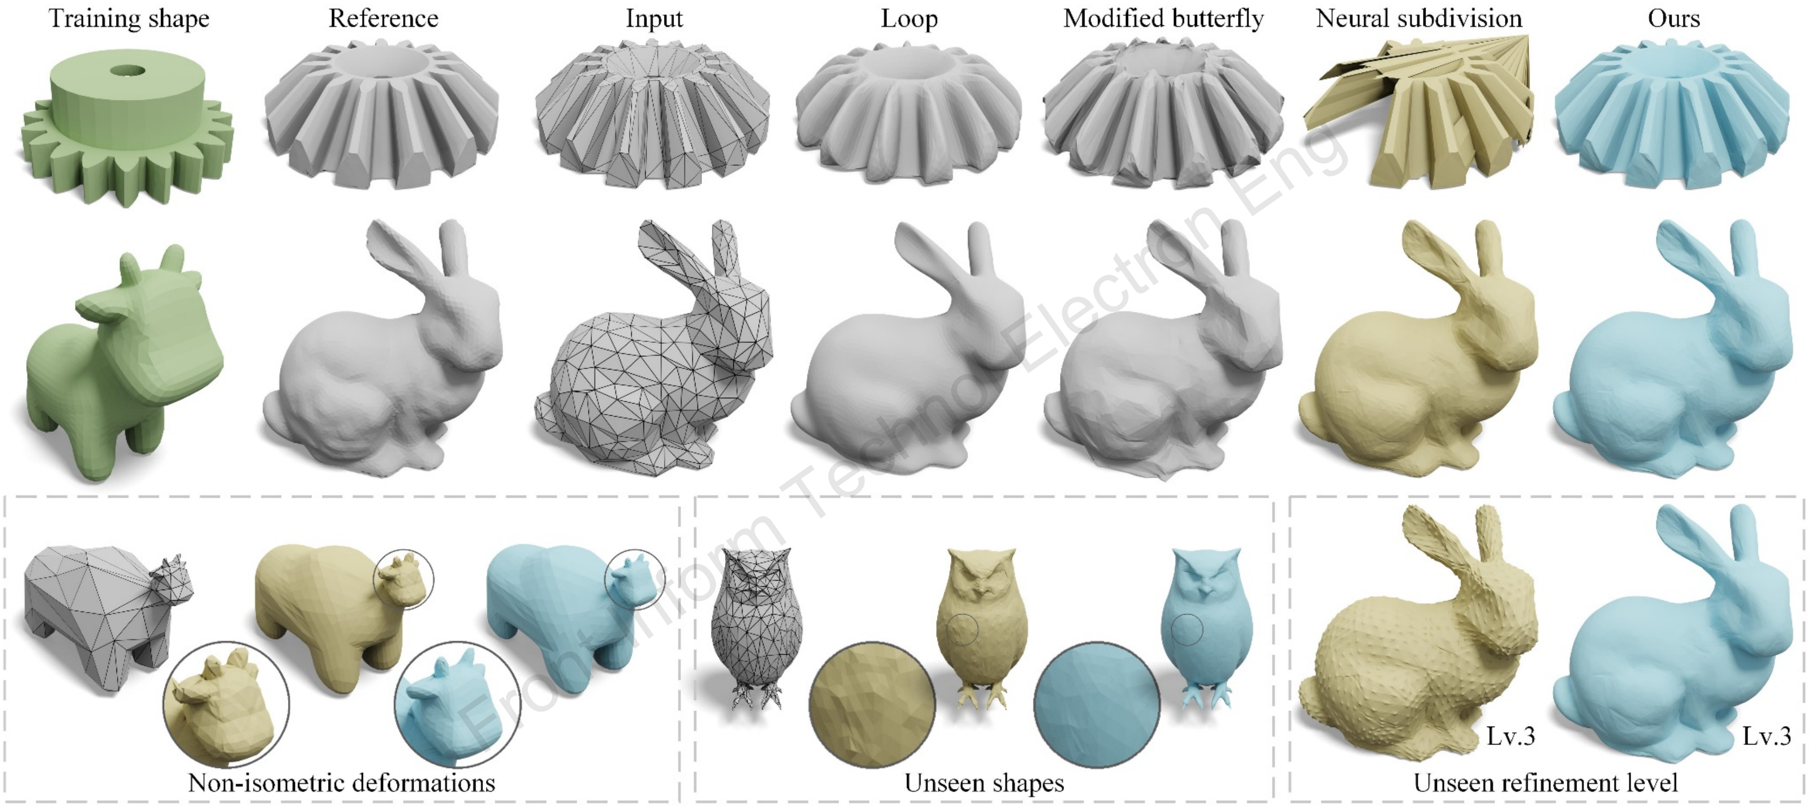

NMR does not suffer from the inherent limitations of existing methods, such as volume shrinkage and over-smoothing (Loop), amplification of tessellation artifacts (modified butterfly), or shape damage (neural subdivision). Moreover, it outperforms neural subdivision in generalization across non-isometric deformations, unseen shapes, and unseen refinement levels.

Major results

NMR outperforms all baseline methods on the **TOSCA** dataset and exhibits significant advantages across all evaluation metrics.

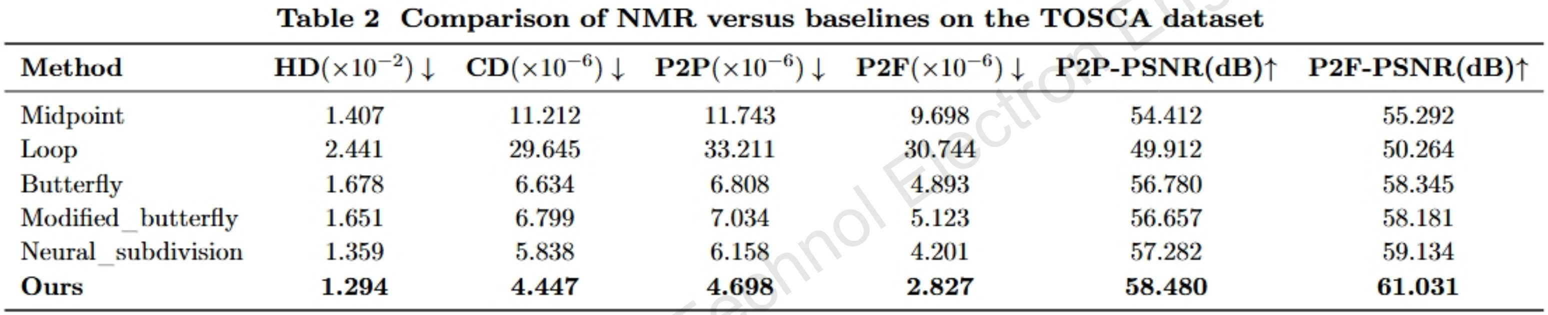

Table 2 Comparison of NMR versus baselines on the TOSCA dataset

Method	HD($\times10^{-2}$)↓	CD($\times10^{-6}$)↓	P2P($\times10^{-6}$)↓	P2F($\times10^{-6}$)↓	P2P-PSNR(dB)↑	P2F-PSNR(dB)↑
Midpoint	1.407	11.212	11.743	9.698	54.412	55.292
Loop	2.441	29.645	33.211	30.744	49.912	50.264
Butterfly	1.678	6.634	6.808	4.893	56.780	58.345
Modified_butterfly	1.651	6.799	7.034	5.123	56.657	58.181
Neural_subdivision	1.359	5.838	6.158	4.201	57.282	59.134
Ours	**1.294**	**4.447**	**4.698**	**2.827**	**58.480**	**61.031**

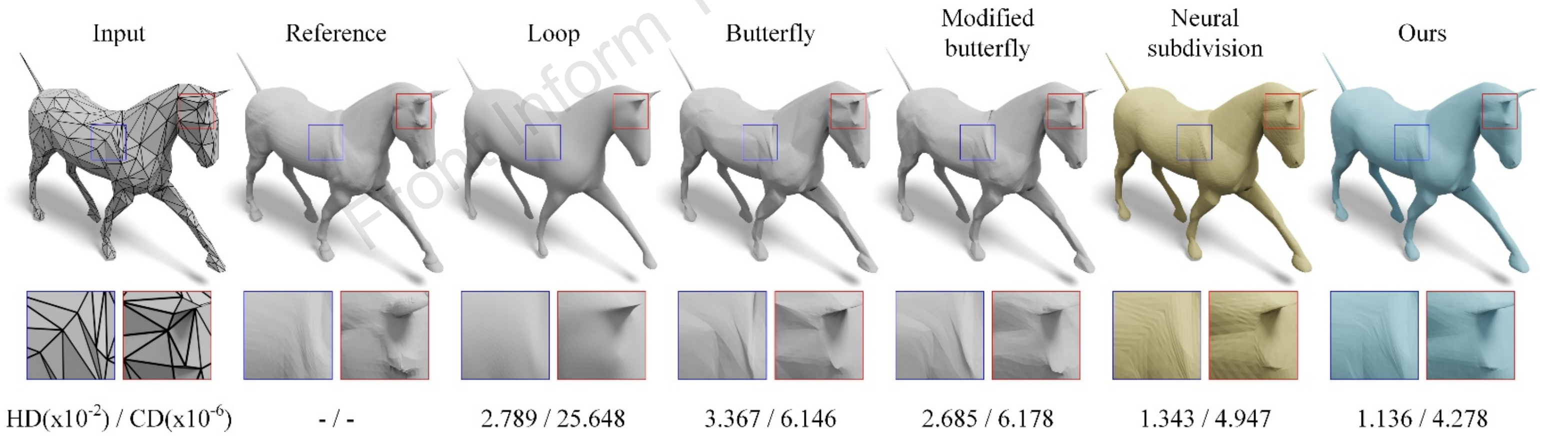

NMR produces finer meshes closer to the ground truth, successfully addressing issues encountered in other methods, such as over-smoothing, volume shrinkage, and artifacts.

Major results

NMR outperforms all baseline methods on the **Thingi10K** dataset and exhibits significant advantages across all evaluation metrics.

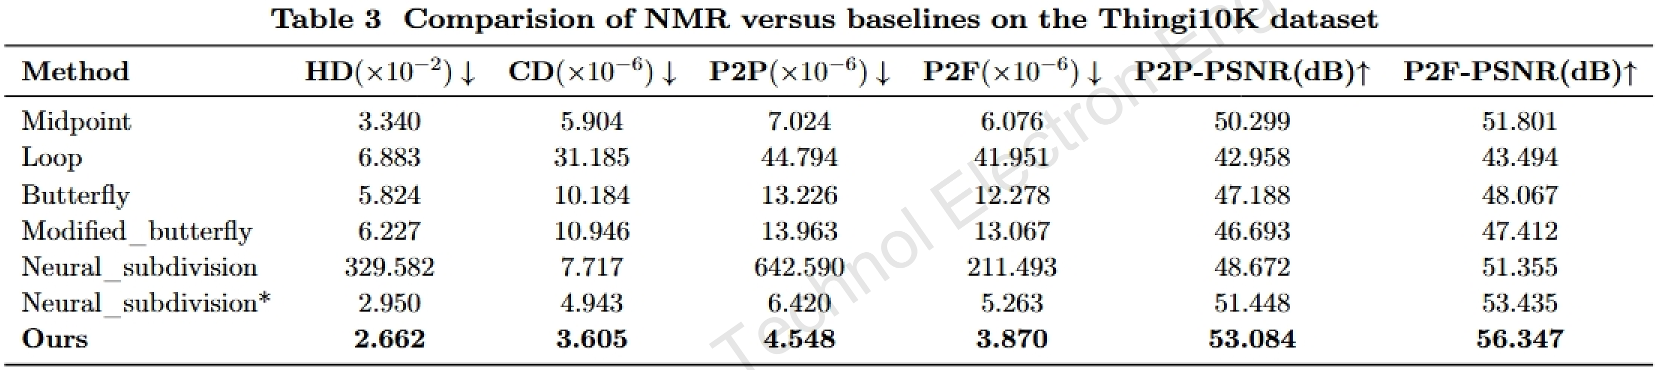

Table 3 Comparision of NMR versus baselines on the Thingi10K dataset

Method	HD($\times10^{-2}$)↓	CD($\times10^{-6}$)↓	P2P($\times10^{-6}$)↓	P2F($\times10^{-6}$)↓	P2P-PSNR(dB)↑	P2F-PSNR(dB)↑
Midpoint	3.340	5.904	7.024	6.076	50.299	51.801
Loop	6.883	31.185	44.794	41.951	42.958	43.494
Butterfly	5.824	10.184	13.226	12.278	47.188	48.067
Modified_butterfly	6.227	10.946	13.963	13.067	46.693	47.412
Neural_subdivision	329.582	7.717	642.590	211.493	48.672	51.355
Neural_subdivision*	2.950	4.943	6.420	5.263	51.448	53.435
Ours	**2.662**	**3.605**	**4.548**	**3.870**	**53.084**	**56.347**

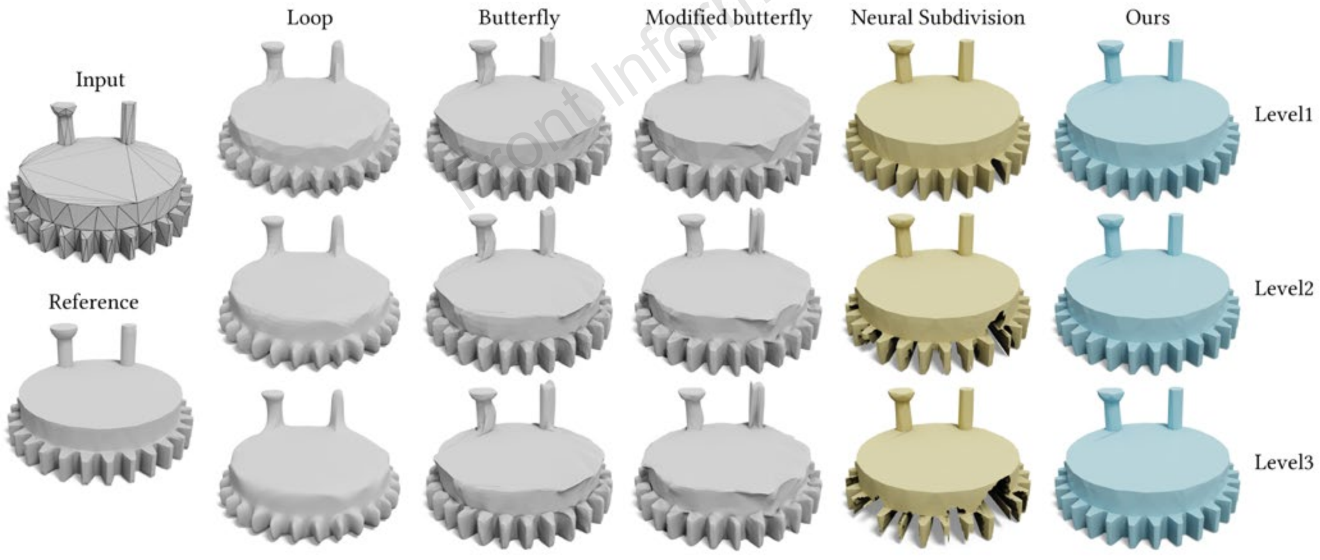

Classical methods encounter issues of volume shrinkage and amplification of tessellation artifacts. Neural subdivision disrupts the mesh shape. In contrast, our method robustly preserves the original shape.

Applications

We mimic the modeling scenario by applying isometric deformations to the coarse cage (gray), resulting in various poses. NMR can generalize to refining **unseen poses** (blue).

Applications

We mimic the modeling scenario by applying non-isometric deformations to the coarse cage (gray). NMR generalizes well to **unseen non-isometric deformations** compared to neural subdivision.

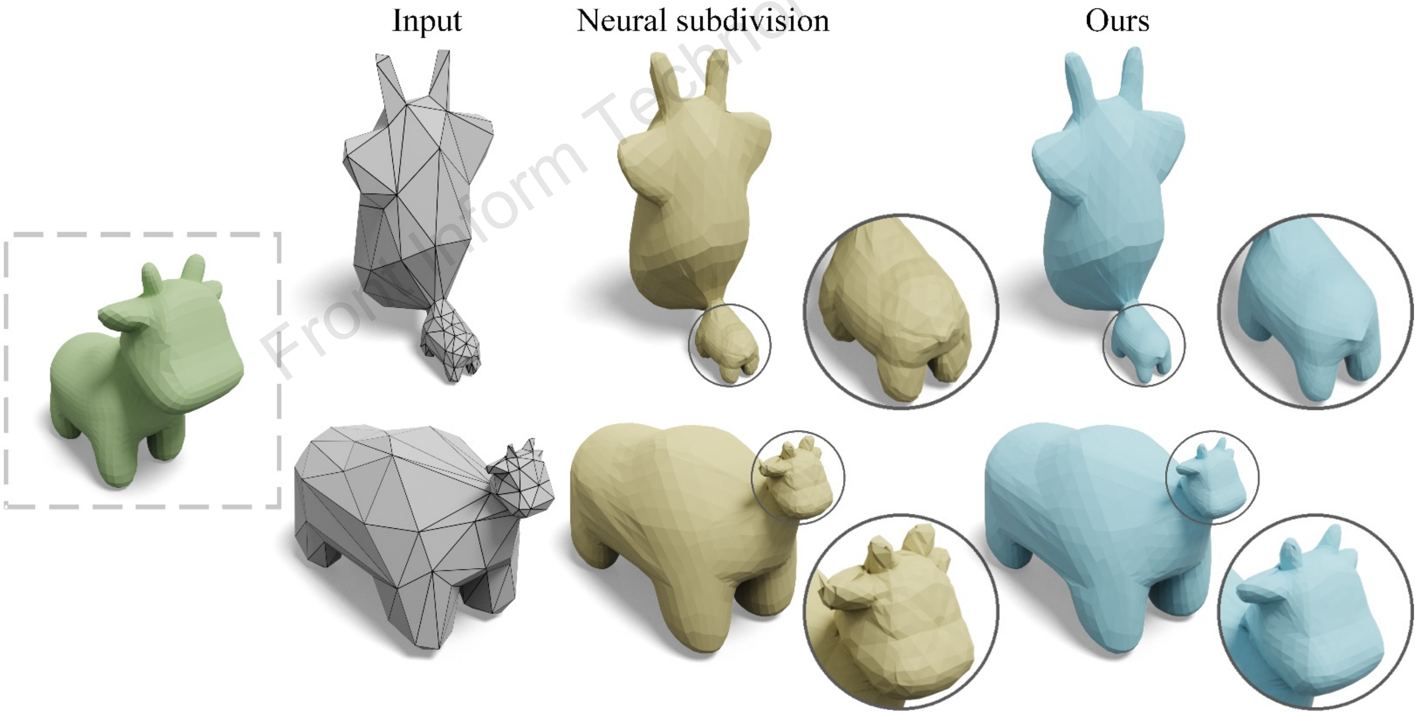

Applications

Even when trained on a single shape (green fish), NMR can generalize to refining **unseen shapes** and outperform neural subdivision.

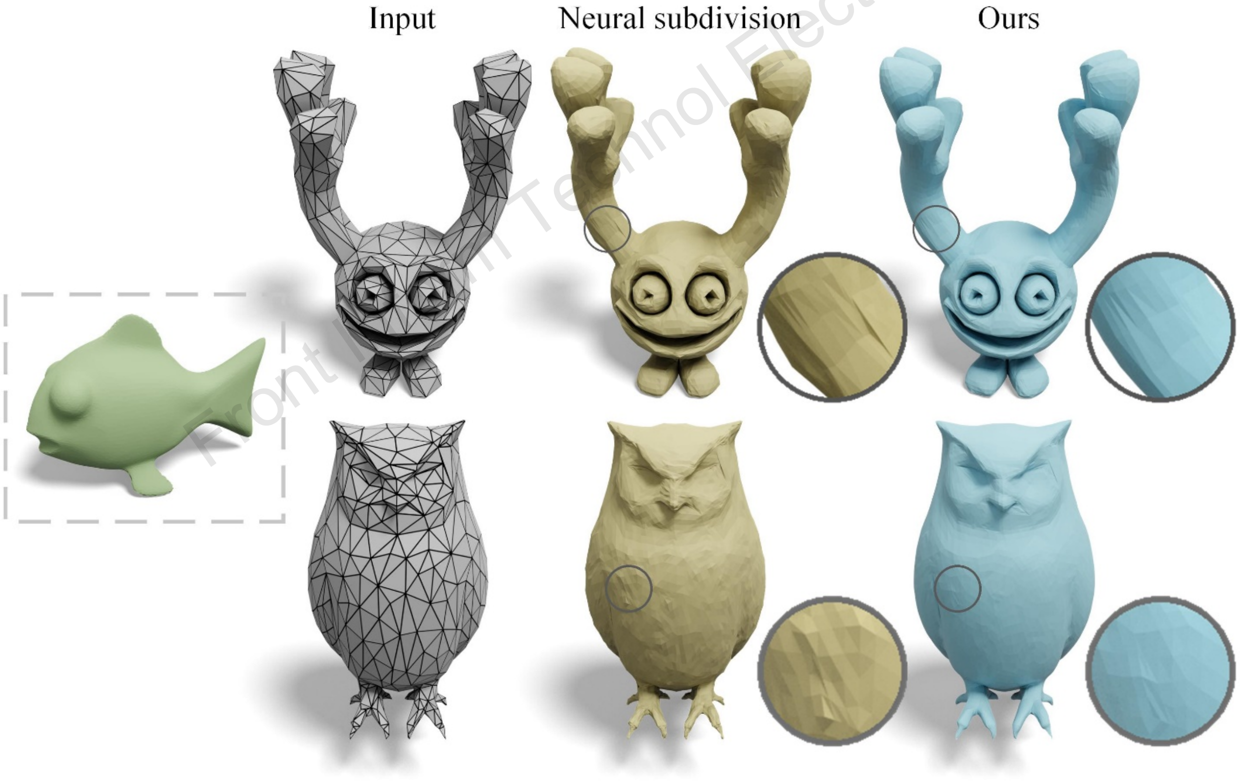

Applications

Even when the **test refinement level exceeds the training level**, NMR can still output a reasonably smooth surface, while neural subdivision cannot.

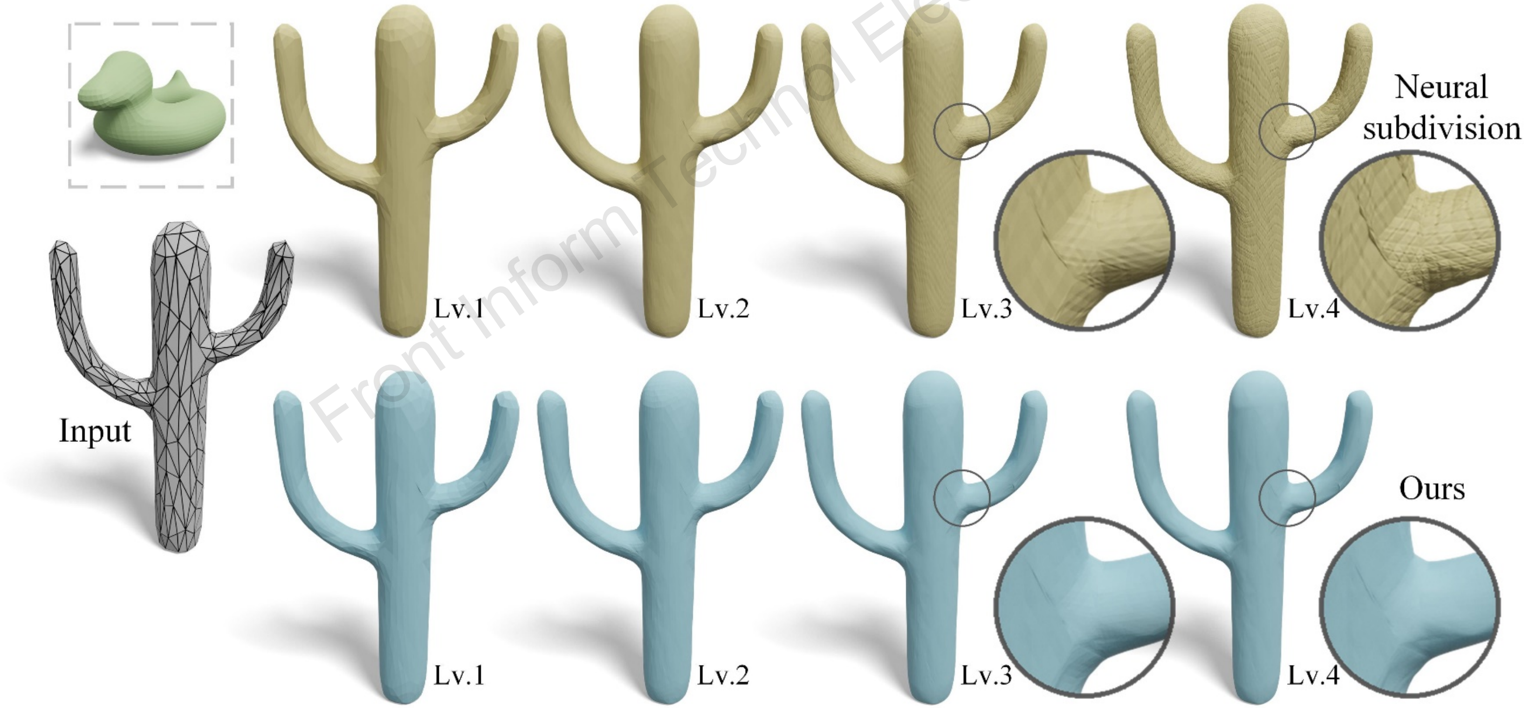

Applications

When trained on a single shape, NMR exhibits **style transfer** characteristics. Using different shapes in training leads to stylized refinement results (blue) biased towards the training shapes (green).

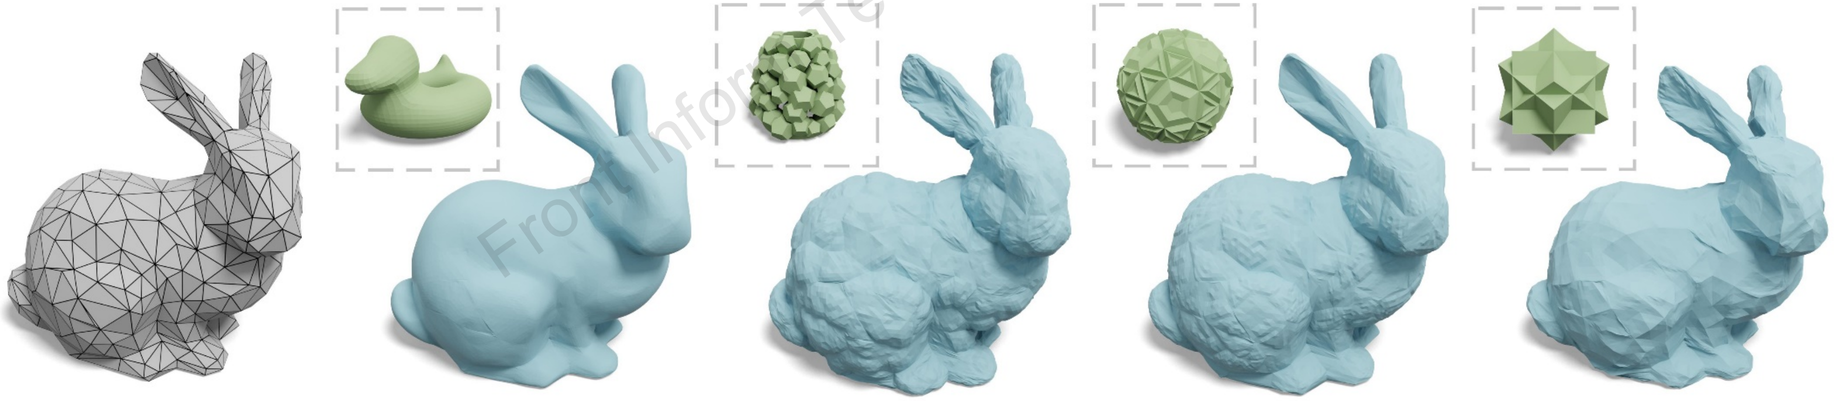

Conclusions

1. We propose a data-driven neural mesh refinement (NMR) method that performs adaptive refinement and exhibits robust generalization.
2. Our main contribution is that we propose and demonstrate that **disentangling the network from non-structural information** (e.g., scale, rotation, and translation) allows the network to focus on learning and applying structural priors for adaptive refinement.
3. Extensive experiments demonstrate that the proposed method outperforms state-of-the-art mesh subdivision methods.

Zhiwei ZHU received the BE degree in electronic information engineering from Tianjin University, Tianjin, China, in 2021. He is currently a PhD candidate in information and communication engineering at the College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, China. His research interests include 3D computer vision, machine learning, and visual media coding.

Xiang GAO received the BE degree in communication engineering from Wuhan University of Technology, Wuhan, China, in 2022, and the MS degree in electronics and information engineering from Zhejiang University, Hangzhou, China, in 2025. His research interests include image/point cloud/Gaussian splatting compression and signal processing.

Yiyi LIAO is an assistant professor at Zhejiang University, leading the X-Dimensional Representations Lab. She received the BS degree from Xi'an Jiaotong University in 2013 and the PhD degree in control science and engineering from Zhejiang University in 2018. Her research interests include 3D computer vision and machine learning.

Lu YU received the BE degree in radio engineering and the PhD degree in communication and electronic systems from Zhejiang University. She is currently a Distinguished Professor with Zhejiang University. She authored more than 150 peer-reviewed papers on image and video coding, quality assessment, and hardware architecture design, and contributed more than 100 adopted technical proposals to define a series of IEEE, ISO/IEC, as well as ISO/IEC and ITU-T joint standards. She is the convenor of ISO/IEC JTC 1/SC 29/WG 4 MPEG Video Coding, one of the voting members of the Standard Activities Subdivision of IEEE Circuits and Systems Society, and on the Editorial Board of *IEEE Trans Circ Syst Video Technol*.