Curve-Aware Gaussian Splatting for 3D Parametric Curve Reconstruction

Supplementary Material

This supplementary material first provides detailed experimental settings, including data processing procedures, implementation details for comparison baseline methods, and evaluation metrics.

In Fig. 1, we provide additional visual comparisons of our method against state-of-the-art baselines on ABC-NEF [10]. Our code and data are available at https://github.com/zhirui-gao/Curve-Gaussian.

1. Datasets

The proposed method is evaluated on three publicly available datasets: ABC-NEF [10], Mv2Cyl's Real Objects [2], and the Replica Dataset [7]. Detailed descriptions of each dataset and the experimental setups are provided below.

ABC-NEF Dataset. The ABC-NEF dataset is a widely adopted benchmark for evaluating 3D curve reconstruction quality. It contains precise CAD models with diverse curve types, comprising 115 objects in total. Following the protocols of EMAP [3] and EdgeGaussians [1], objects with indistinct sharp curve features were filtered out, resulting in 82 objects for evaluation. Each object includes 50 images with a resolution of 800×800. For computational efficiency, all images were resized to 400×400.

Mv2Cyl's Real Objects Dataset. The Mv2Cyl's Real Objects dataset consists of multiple 3D-printed objects captured using an iPhone 12. Multi-view images were extracted from continuous videos, with camera poses computed using COLMAP. Ground truth CAD models are also provided for evaluation. This dataset presents a significant challenge for 3D edge detection due to the top-down perspective of the captured images. To generate edge maps, the following pipeline is employed:

- SAM2 [6] is utilized to segment objects from the background.
- A monocular normal estimation network [9] is applied to identify high-curvature regions as object edges. This approach demonstrated superior performance compared to edge detection methods [5, 8], which often extracted irrelevant edges due to lighting interference.
- Edge maps are resized to 480×480 for faster processing.

Since the camera poses and CAD models are not aligned in the same coordinate system, the Iterative Closest Point (ICP) was used to register the reconstructed edge points with the ground truth CAD curves for quantitative evaluation.

The lack of real-world benchmarks for 3D curve reconstruction is a notable gap in the field. The Mv2Cyl's Real Objects dataset addresses this limitation. After getting permission from the authors of Mv2Cyl, a standardized benchmark for evaluating 3D edge reconstruction will be proposed based on their multi-view images. We believe it will be an important contribution to the field.

Replica Dataset. We follow the experimental setup in EMAP [3] on this dataset, focusing on three scenes: Room 0, Room 1, and Room 2.

2. Baselines

Our method is compared against four state-of-the-art 3D line and curve reconstruction baselines. These include three learning-based methods—NEF [10], EdgeGaussians [1], and EMAP [3]—and one line-based Structure-from-Motion (SfM) method, LIMAP [4]. For a fair comparison, the default parameters and settings provided by the authors are adopted for all baselines. In the case of NEF and EMAP, their pre-trained models were directly applied to generate visual results.

3. Evaluation Metrics

To quantitatively assess the performance of our method, we adopt a set of evaluation metrics that align with established protocols in this field. Points are uniformly sampled along both the reconstructed parametric curves and the corresponding ground-truth edges, enabling a direct comparison between them. The metrics are introduced as follows:

- **Accuracy**: This metric calculates the average distance from each predicted point to its closest counterpart on the ground-truth curve. Smaller values correspond to higher accuracy.
- **Completeness**: This measures the average distance from each ground-truth point to the nearest predicted point. Improved performance is indicated by lower values.
- **Precision at Threshold** τ **(**$P(\tau)$**)**: This quantifies the proportion of predicted points that lie within a distance τ of any ground-truth point. Higher precision values reflect better alignment with the ground truth.
- **Recall at Threshold** τ **(**$R(\tau)$**)**: This evaluates the proportion of ground-truth points that have at least one predicted point within a distance τ. Higher recall values signify better coverage of the ground truth.

In addition to these conventional metrics, we introduce a new metric, **Curve Count**, which evaluates the compactness of the reconstructed curves, which can be found in Table 2 of the main paper. This metric counts the total number of curves generated by the method, providing insight into the efficiency of the representation. A smaller curve count indicates a more concise and compact reconstruction, which

is advantageous for downstream tasks that require efficient curve processing.

For consistency with prior evaluations, precision and recall are computed at distance thresholds of $\tau = 5$, 10, and 20 millimeters (mm). For other parameters in the evaluation, such as the number of sampling points, we adhere to the same settings as those used in EMAP [3].

4. Implementation Details

The weight coefficients λ_1, λ_2, λ_3, and λ_4 were set to 0.01, 0.01, 0.01, and 0.0005, respectively. For all Bézier curves and straight lines, a default of 12 Gaussian points was sampled per curve.

For the ABC dataset, the midpoints of Bézier curves were initialized by uniformly sampling $15 \times 15 \times 15$ points in 3D space. In contrast, for COLMAP-based datasets, the midpoints were initialized using the point cloud generated by Structure-from-Motion (SfM). The threshold for curve merging was kept consistent with EMAP [3].

To ensure the quality of the reconstructed curves, curves with an opacity below 0.05 were removed, and curves with bending angles exceeding 20° were split. Additionally, Gaussian components with mask attributes below 0.01 were considered redundant and discarded. Further details on parameter design and implementation can be found in the accompanying code.

References

[1] Kunal Chelani, Assia Benbihi, Torsten Sattler, and Fredrik Kahl. Edgegaussians–3d edge mapping via gaussian splatting. *arXiv preprint arXiv:2409.12886*, 2024. 1

[2] Eunji Hong, Minh Hieu Nguyen, Mikaela Angelina Uy, and Minhyuk Sung. Mv2cyl: Reconstructing 3d extrusion cylinders from multi-view images. *arXiv preprint arXiv:2406.10853*, 2024. 1

[3] Lei Li, Songyou Peng, Zehao Yu, Shaohui Liu, Rémi Pautrat, Xiaochuan Yin, and Marc Pollefeys. 3d neural edge reconstruction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2

[4] Shaohui Liu, Yifan Yu, Rémi Pautrat, Marc Pollefeys, and Viktor Larsson. 3d line mapping revisited. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21445–21455, 2023. 1

[5] Xavier Soria Poma, Edgar Riba, and Angel Sappa. Dense extreme inception network: Towards a robust cnn model for edge detection. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1923–1932, 2020. 1

[6] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 1

[7] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J. Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, Anton Clarkson, Mingfei Yan, Brian Budge, Yajie Yan, Xiaqing Pan, June Yon, Yuyang Zou, Kimberly Leon, Nigel Carter, Jesus Briales, Tyler Gillingham, Elias Mueggler, Luis Pesqueira, Manolis Savva, Dhruv Batra, Hauke M. Strasdat, Renzo De Nardi, Michael Goesele, Steven Lovegrove, and Richard Newcombe. The Replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 1

[8] Zhuo Su, Wenzhe Liu, Zitong Yu, Dewen Hu, Qing Liao, Qi Tian, Matti Pietikainen, and Li Liu. Pixel difference networks for efficient edge detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5117–5127, 2021. 1

[9] Chongjie Ye, Lingteng Qiu, Xiaodong Gu, Qi Zuo, Yushuang Wu, Zilong Dong, Liefeng Bo, Yuliang Xiu, and Xiaoguang Han. Stablenormal: Reducing diffusion variance for stable and sharp normal. *ACM Transactions on Graphics (TOG)*, 2024. 1

[10] Yunfan Ye, Renjiao Yi, Zhirui Gao, Chenyang Zhu, Zhiping Cai, and Kai Xu. Nef: Neural edge fields for 3d parametric curve reconstruction from multi-view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8486–8495, 2023. 1, 3

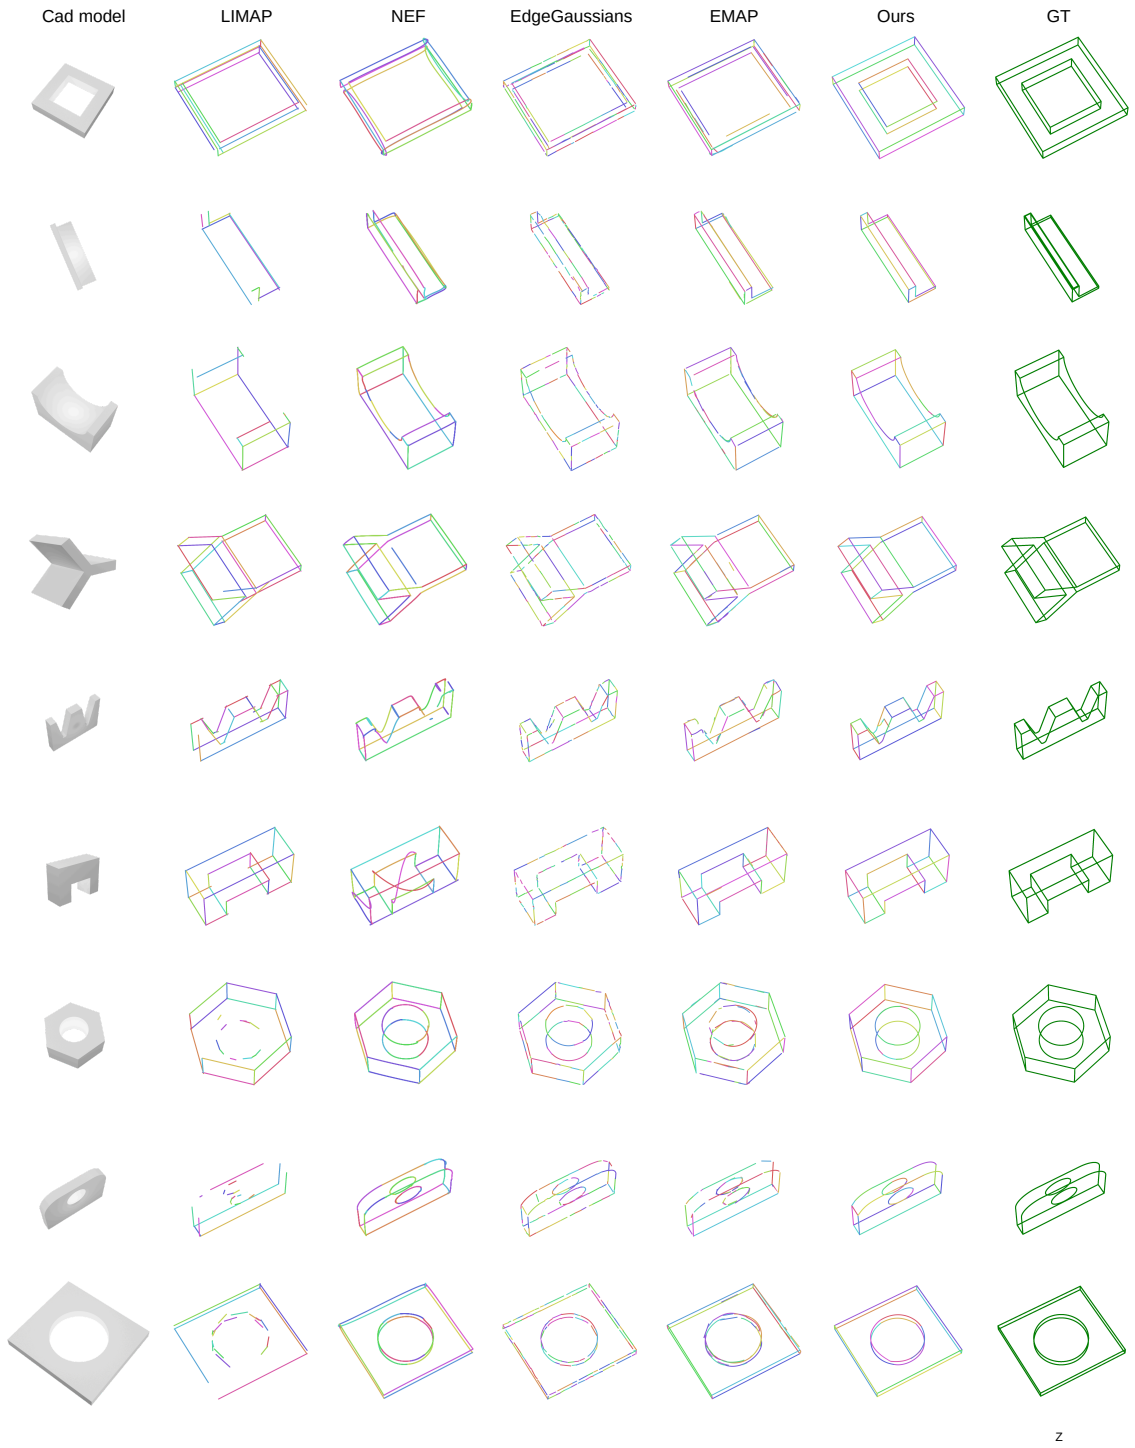

Figure 1. **Qualitative comparisons on ABC-NEF [10].** Distinct colors represent different curves. Our method achieves more complete and accurate edge reconstruction of objects while maintaining parametric compactness.