NeuraLeaf: Neural Parametric Leaf Models with Shape and Deformation Disentanglement

Yang Yang[1] Dongni Mao[1] Hiroaki Santo[1] Yasuyuki Matsushita[1,2] Fumio Okura[1]
[1]The University of Osaka [2]Microsoft Research Asia – Tokyo
{yang.yang,mao.dongni,santo.hiroaki,yasumat,okura}@ist.osaka-u.ac.jp

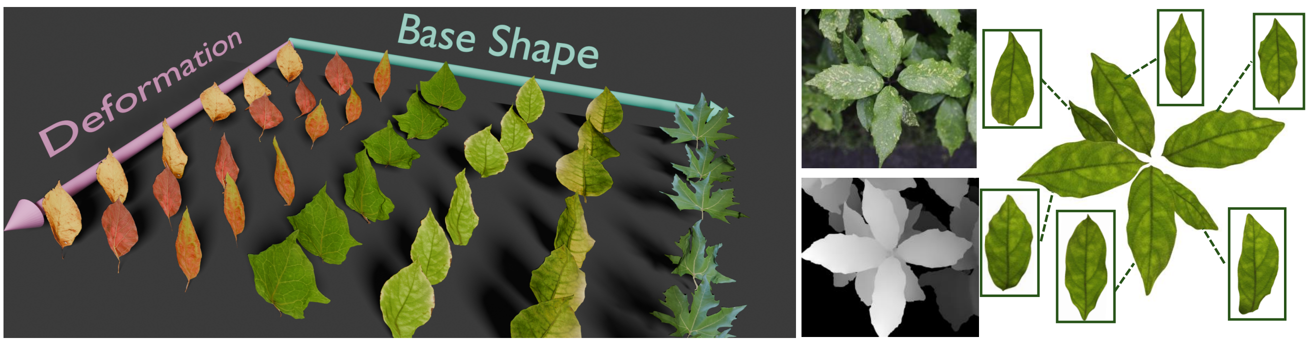

Figure 1. Our neural parametric model for leaves, NeuraLeaf, represents **shapes of various leaf species** and **natural 3D deformation**. Our model represents the leaves' flattened shape and their 3D deformation in disentangled latent spaces (left). Our method enables the instance-wise reconstruction of leaves via fitting to real-world observations, besides pure CG modeling (right).

1. Introduction

3D leaf reconstruction is important as it provides insights into plant growth in agricultural applications and is essential for creating virtual computer graphics (CG) assets. Parametric 3D models have been built on plant leaves [1, 2]. Although for modeling static leaves, some methods take bio-inspired approaches [12, 14] using venation patterns or the nature of flattened shapes.

However, the requirement of large datasets poses challenges for leaf modeling due to the lack of a 3D dataset containing leaves' deformations.

To these issues, we integrate leaves' bio-inspired insights with a data-driven framework. We propose a leaf-specific NPM representation, NeuraLeaf, leveraging the nature of leaves, *i.e.*, the flattened leaves can be approximated as a plane. As illustrated in Fig. 1, we disentangle the leaves' 3D geometry into their 2D base (*i.e.*, flattened) shape and 3D deformation represented by different latent codes. 2D base shapes can be learned from 2D-scanned image datasets of leaves and for the 3D deformation modeling, we newly construct a dataset, DeformLeaf, by acquiring real-world 3D leaf shapes in natural deformation.

Beyond data limitations, leaf deformation also presents challenges due to its high flexibility, and shape variance across species further complicates defining a common skeleton structure, making conventional skinned vertex-based models [3, 8, 9, 11, 13, 15–17] unsuitable. We, therefore, newly introduce a skeleton-free skinning model that learns a parametric space for skinning parameters.

We showcase applications of our NeuraLeaf for reconstruction and modeling purposes, where our method enables the *leaf-wise* reconstruction from the observation as shown in Fig. 1.

2. NeuraLeaf: Neural Parametric Leaf Model

The overview of our disentangled representations and training strategy is illustrated in Fig. 2. We first represent the base leaf shapes in a flattened state using a neural SDF conditioned by latent codes, which are learned from a large-scale 2D dataset [5]. The base shape and texture spaces are modeled by a shape decoder f_{θ_s} and texture generator f_{θ_t} conditioned on shape & texture latent codes $\{\mathbf{z}_s, \mathbf{z}_t\}$. We then convert the SDF into a mask in a differentiable manner.

Based on the learned shape space, we then train two deformation decoders to predict skinning parameters, called skinning weight decoder f_{θ_w} and transformation decoder f_{θ_d}, which infer skinning weight of each point and transformation of each control point. We use linear blend skinning (LBS) [7] to deform the points. We learn the deformation space from our **DeformLeaf** dataset containing natural deformation.

In the inference stage, to obtain a reasonable initialization of latent codes, we use an inversion-based approach

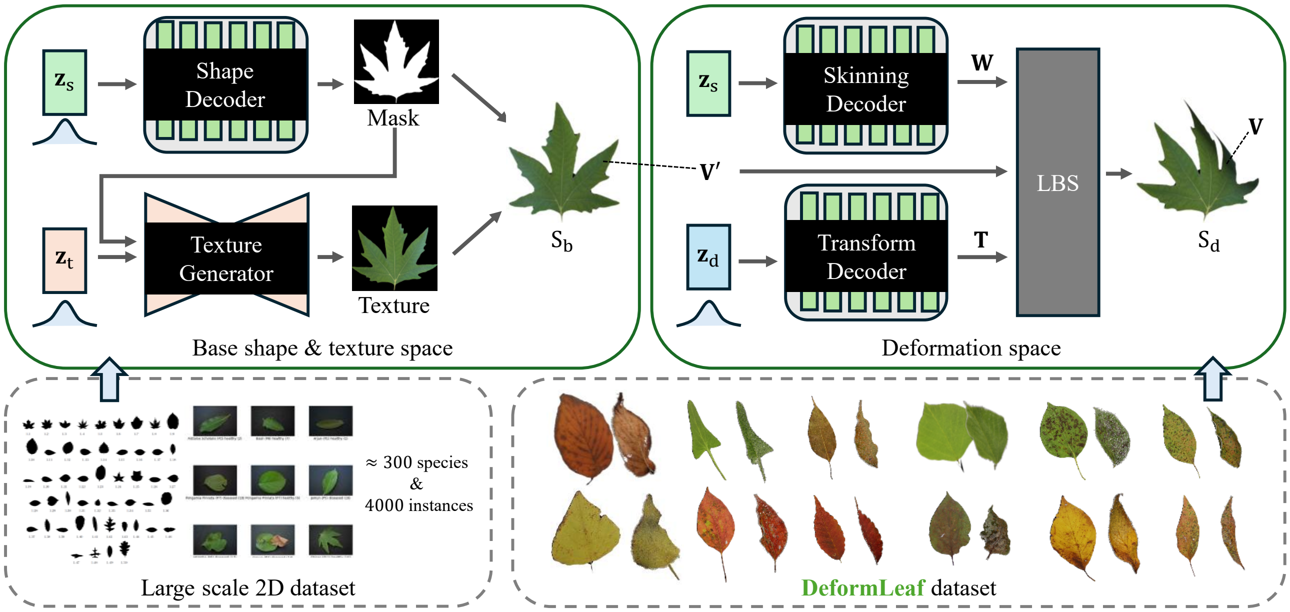

Figure 2. **Overview of NeuraLeaf.**

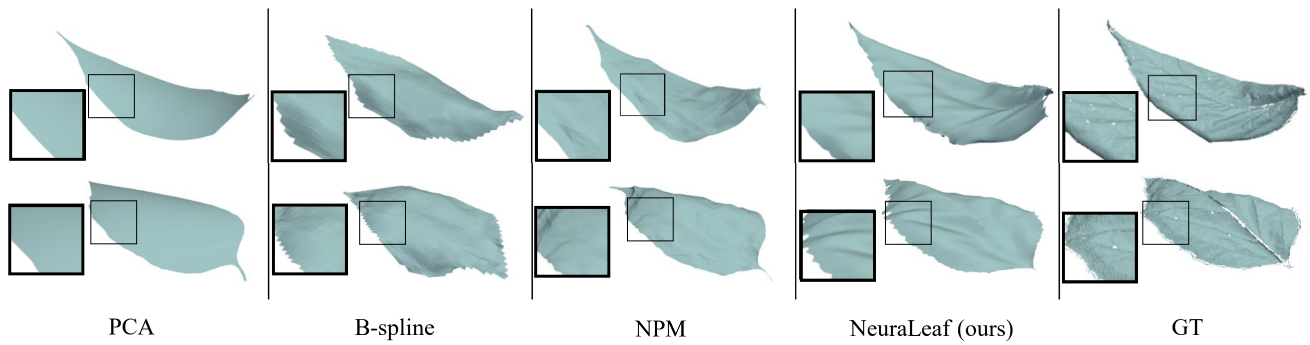

Figure 3. **Visual comparisons of reconstructed leaf samples.** We compare the ground truth with predicted shapes. Our method shows accurate reconstruction compared to baseline methods and has significantly more surface detail, such as leaf veins.

Table 1. **Quantitative comparison** between baselines on the test set of DeformLeaf dataset for single leaf reconstruction. The best result is highlighted in **bold**.

Method	C-ℓ_2 [mm] ↓	NC ↑	Corres-free	Temp-free	Inf. time [s]
PCA [2]	32.7	0.924	×	×	157
B-spline [6]	26.7	0.957	√	√	18
NPM [10]	15.1	0.961	×	√	73
NeuraLeaf (ours)	**2.1**	**0.973**	√	√	55

inspired by [4] to learn the mapping between input observation and the latent codes.

3. Experiments

Table 1 shows a quantitative comparison on fitting to real-world single-view partial leaf point clouds, along with a visual comparison shown in Fig. 3. According to the results, our NeuraLeaf outperforms the baseline methods, achieving higher normal consistency and also reducing chamfer distance.

For multiple leaves cases that contain occlusions, we use top-view RGB-D observations with segmentation masks and obtain point clouds corresponding to each leaf instance. Figure 1 presents a visual example of fitting multiple leaf instances to RGB-D observation. The leaf-wise reconstruction results highlight that, by sharing similar shape latent $\mathbf{z}_s$, our method reasonably recovers leaf-wise shapes even in the portions occluded in the input RGB-D image.

4. Conclusions

This paper has presented NeuraLeaf, a novel leaf-specific NPM that faithfully captures the unique characteristics of leaf geometry. Our method effectively disentangles a leaf's 3D structure into a 2D base shape and 3D deformation, each represented by a distinct latent code. By using large 2D-scanned image datasets and our newly acquired DeformLeaf dataset, NeuraLeaf learns from limited data while managing the flexibility inherent in leaf morphology.

References

[1] Paul Beardsley and Gaurav Chaurasia. Editable parametric dense foliage from 3D capture. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 1

[2] Derek Bradley, Derek Nowrouzezahrai, and Paul Beardsley. Image-based reconstruction and synthesis of dense foliage. *ACM Transactions on Graphics (TOG)*, 32(4):74:1–74:10, 2013. 1, 2

[3] Christoph Bregler, Aaron Hertzmann, and Henning Biermann. Recovering non-rigid 3D shape from image streams. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2000. 1

[4] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3D shape reconstruction and completion. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2

[5] Siddharth Singh Chouhan, Ajay Kaul, Uday Pratap Singh, and Sanjeev Jain. A data repository of leaf images: Practice towards plant conservation with plant pathology. In *Proceedings of International Conference on Information Systems and Computer Networks (ISCON)*, 2019. 1

[6] William J. Gordon and Richard F. Riesenfeld. B-spline curves and surfaces. *Computer Aided Geometric Design*, 1974:95–126, 1974. 2

[7] Alec Jacobson, Zhigang Deng, Ladislav Kavan, and JP Lewis. Skinning: Real-time shape deformation. In *ACM SIGGRAPH 2014 Courses*, 2014. 1

[8] Chen Kong and Simon Lucey. Deep non-rigid structure from motion. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 1

[9] Marko Mihajlovic, Yan Zhang, Michael J Black, and Siyu Tang. LEAP: Learning articulated occupancy of people. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1

[10] Pablo Palafox, Aljaž Božič, Justus Thies, Matthias Nießner, and Angela Dai. NPMs: Neural parametric models for 3D deformable shapes. In *Proceedings of IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2

[11] Pablo Palafox, Nikolaos Sarafianos, Tony Tung, and Angela Dai. SPAMs: Structured implicit parametric models. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1

[12] Adam Runions, Martin Fuhrer, Brendan Lane, Pavol Federl, Anne-Gaëlle Rolland-Lagan, and Przemyslaw Prusinkiewicz. Modeling and visualization of leaf venation patterns. *ACM Transactions on Graphics (TOG)*, 24(3):702–711, 2005. 1

[13] Shubham Tulsiani, Nilesh Kulkarni, and Abhinav Gupta. Implicit mesh reconstruction from unannotated image collections. *arXiv preprint arXiv:2007.08504*, 2020. 1

[14] Weiliang Wen, Baojun Li, Bao-jun Li, and Xinyu Guo. A leaf modeling and multi-scale remeshing method for visual computation via hierarchical parametric vein and margin representation. *Frontiers in Plant Science*, 9(783), 2018. 1

[15] Shangzhe Wu, Ruining Li, Tomas Jakab, Christian Rupprecht, and Andrea Vedaldi. MagicPony: Learning articulated 3D animals in the wild. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1

[16] Yuefan Wu, Zeyuan Chen, Shaowei Liu, Zhongzheng Ren, and Shenlong Wang. CASA: Category-agnostic skeletal animal reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

[17] Chun-Han Yao, Wei-Chih Hung, Yuanzhen Li, Michael Rubinstein, Ming-Hsuan Yang, and Varun Jampani. Hi-lassie: High-fidelity articulated shape and skeleton discovery from sparse image ensemble. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1