

Tree-D Fusion: Simulation-Ready Tree Dataset from Single Images with Diffusion Priors

Jae Joong Lee¹, Bosheng Li¹, Sara Beery², Jonathan Huang³, Songlin Fei⁴, Raymond A. Yeh¹, and Bedrich Benes¹

¹ Purdue University, Department of Computer Science

² Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science

³ Google

⁴ Purdue University, Department of Forestry & Natural Resources

<https://www.jaejoonglee.com/treedfusion/>




Fig. 1: *Tree-D Fusion* takes a single view image (left) and reconstructs a 3D simulation-ready tree model. The tree model can be used to simulate growth over time with a detailed branching structure with leaves. We provide a dataset of 3D reconstructed tree models from 600,000 Google Street View images.

Abstract. We introduce *Tree-D Fusion*, featuring the first collection of 600,000 environmentally aware, 3D simulation-ready tree models generated through Diffusion priors. Each reconstructed 3D tree model corresponds to an image from Google’s Auto Arborist Dataset, comprising street view images and associated genus labels of trees across North America. Our method distills the scores of two tree-adapted diffusion models by utilizing text prompts to specify a tree genus, thus facilitating shape reconstruction. This process involves reconstructing a 3D tree envelope filled with point markers, which are subsequently utilized to estimate the tree’s branching structure using the space colonization algorithm conditioned on a specified genus.

1 Introduction

Trees provide immense and essential value to human society and underpin diverse ecosystems worldwide [74]. They cool the environment, improve air quality, capture carbon dioxide, produce oxygen, and have a positive effect on human physical and mental health [9, 13, 30, 67, 130, 136]. The complex effect of trees

on the environment has been studied for centuries. Currently, computational models that seek to understand these relationships are hindered by a lack of data. In particular, acquiring phenotypical traits such as branching angles, wood volume, crown size, and trunk width is currently infeasible at a large scale. Simulation-ready 3D reconstructed trees (i.e., data that can be used in algorithms simulating tree development) that fit real-world data would enable the estimation of useful ecological quantities at scale. For example, the overall amount of wood or carbon sequestration, and would even allow for large-scale simulations of “what-if” scenarios, such as estimating shadow cover in both leaf-on and leaf-off scenarios at different times of the day or in different seasons [76, 115, 133].

Existing developmental models of vegetation typically rely on manual parameter tuning [52, 65, 75, 90], or on a model parameter recovery that is fit to the data [73, 81, 103]. Digital datasets are scarce and usually available as unstructured point clouds [57, 91] or images [7]. Providing realistically reconstructed trees that fit the input data at scale is yet to be achieved, as large-scale and geographically diverse data is needed to ensure that the synthetic versions accurately reflect the morphological diversity of real trees. Despite the abundance of tree images, converting them into 3D models has been difficult due to trees’ complex shapes and self-occlusions from branches and leaves, which often hide significant portions of a tree. Thus, reconstruction methods often rely on biologically-based developmental algorithms that approximate missing parts [51, 72].

We introduce *Tree-D Fusion*, the first tree dataset of its kind and scale, containing 600,000 3D simulation-ready tree models grounded to individual real trees via street-level imagery and diffusion priors. The generation process is informed by genus-specific characteristics learned from the Auto Arborist Dataset (AAD) [7] through a diffusion model trained specifically on tree images and captions. The result is a 3D envelope used by a genus-conditioned developmental model [52] that grows into this envelope while respecting the overall shape. The output is a reconstructed simulation-ready 3D tree model with a detailed branching structure and foliage.

We use the two Diffusion priors to reconstruct real trees from AAD, covering 23 cities across North America and over 300 genus categories. *Tree-D Fusion* is validated against [56, 92, 109]. An example in Fig. 1 left demonstrates how, from a single input image, our approach generates a model that can be used to simulate the growth of that specific real-world tree over time, even simulating the growth of leaves. Our dataset, *Tree-D Fusion* contains 600,000 such models grounded on real-world trees, and we believe it will standardize benchmarking and advancements in forestry.

2 Related Work

Remote sensing for trees and forests. Tree and forest monitoring is an important topic in remote sensing, forestry, and ecology. Automated methods to measure canopy cover and forest height have been explored at large scales [5, 11, 21, 29, 49, 62, 63]. Recent work seeks to do, modeling at the level of individual

trees instead of forests, including individual tree detection, counting, and crown delineation [82, 104, 111, 120, 121] and identifying the species of individual trees in both wilderness [22, 32, 33, 39, 40, 66, 82, 99, 119] and urban settings [3, 6, 7, 12, 40, 105, 118, 131]. There is also increasing interest in understanding the structure of trees at an individual level to estimate tree attributes, including height, diameter at breast height (DBH), branching structure, and shadow estimation [45]. In recent years, more sophisticated sensors such as airborne LiDAR have allowed for higher detail representations of individual tree shapes and structures, e.g., [8, 55, 98, 135]. However, these sensor measurements are more expensive and cumbersome to collect at scale, especially at the level of individual trees. Thus, methods that allow for scalable measurement of tree geometry can significantly benefit these fields [110]. Our dataset addresses this concern by offering a vast collection of realistic 3D tree models. Also, our data are grounded in real-world images and exhibit the most realistic reconstruction compared to recent methods, as demonstrated in Sec. 5.1.

Tree modeling has been studied from a very different perspective in the computer graphics literature. The generation of 3D tree geometric and behavioral models dates back to the seminal work of Lindenmayer, who introduced parallel string rewriting systems (L-systems) to simulate cellular subdivision [54]. L-systems were extended by Prusinkiewicz to capture branching patterns [85] and later evolved [90] into a full mathematical formalism that captures continuous development [86, 87], plant signaling [88], environmental interaction [89], and ecosystems [26]. Recent approaches focus on the simulation of interactions of tree models with wind [80], fire [37], climbing trees [38], and ecosystems [65, 73]. One of the most important open problems in tree modeling is reconstructing the developmental model from data [34, 50, 68, 102]. A developmental model that simulates an existing tree would allow for estimating its shape under different conditions and predicting its shape changes and the effect on the environment.

Tree geometry reconstruction methods can be classified according to the input data type, and the most prevailing are RGB images and point clouds. Tree reconstruction from point clouds involves creating a tree skeleton [14] that captures tree topology and basic features (branching angles, branch length) [27, 61, 126]. Other methods use voxels to retrieve the shape from particle flow [72] or fitting geometric proxies to approximate the branch shape [1, 2, 36, 42]. Although tree geometry has been extensively investigated, only a limited number of studies have concentrated on foliage reconstruction [4, 10, 15, 17, 129] or trees incorporating leaves [60, 123]. RGB images are a much more widely accessible modality, so there is significant interest in reconstructing trees from RGB image input. Some methods make use of incomplete information supplied via user input to identify the trunk and branches [19, 59, 93, 107, 108]. Other methods create 3D voxels with approximate information that control tree growth [96] or use multi-image input [35]. Recent deep learning approaches use transformer-based tree reconstruction to L-systems [50], single image-based tree reconstruction using convolutional networks [51], statistical inference [44], or adversarial networks [59]. Our approach is inspired by an algorithm that generates a 3D geometric envelope

completed by a developmental model [51]. Rather than a static proxy, we offer a detailed 3D envelope capturing the tree’s shape more accurately.

Single image 3D model reconstruction has been an active area of research for a long time [18, 23, 28, 64, 69, 71, 77, 79, 114]. Given RGB images, these methods aim to reconstruct common 3D model types such as mesh [114, 122, 127], voxels [20, 78, 112, 124, 128], or point clouds [31]. Recently, diffusion models [97] for 3D generation applications [84, 113, 116] have been used as priors for single image reconstruction [25, 56, 92, 95, 109, 125]. These approaches significantly improve reconstruction quality for occluded parts in a given input image.

Despite these advancements, reconstructing trees from a single image remains challenging, and existing methods fail to create realistic 3D models of trees due to the intricate complexity of a tree’s structure.

Fig. 2 demonstrates reconstructed trees using Magic123 [92] that generates unrealistic tree geometry. These shortcomings of existing methods can be traced back to the inadequate representation of tree geometry in the diffusion models that are used as priors for the 3D shape, and that is the problem we aim to solve. *Tree-D Fusion* addresses the challenge of predicting occluded branches from a single image by outlining the tree’s contours and then using a species-specific procedural model [52] to grow the tree within these boundaries.




Fig. 2: Magic123 [92] fails to capture trees’ complex geometry.

3 Background

We first briefly review diffusion models for 3D generation via score distillation [113]. As a probabilistic model, conditional diffusion aims to learn a model that can efficiently sample from the distribution $P(\mathbf{I}|\mathbf{c})$ of the image \mathbf{I} given some context \mathbf{c} . This is done by learning a family of denoisers $D(\mathbf{I}|\mathbf{c}; \sigma)$ at different Gaussian noise levels σ . As shown in the works [43, 101], the denoiser D approximates the gradient field of the data log-likelihood, also known as the “score function”:

$$\nabla_{\mathbf{I}} \log P_{\sigma}(\mathbf{I}|\mathbf{c}) \approx (D(\mathbf{I}; \sigma) - \mathbf{I})/\sigma^2. \quad (1)$$

With this interpretation of a diffusion model, [113] proposed Score Jacobian Chaining (SJC) that applies the “chain-rule” to *distill* knowledge from a trained diffusion model. Specifically, SJC is applied to optimizing a NeRF [16, 106] for generating 3D assets:

$$\frac{\partial P_{\sigma}(\mathbf{I}|\mathbf{c})}{\partial \theta} = \underbrace{\frac{\partial P_{\sigma}(\mathbf{I}|\mathbf{c})}{\partial \mathbf{I}}}_{\text{score}} \frac{\partial \mathbf{I}}{\partial \theta}, \quad (2)$$

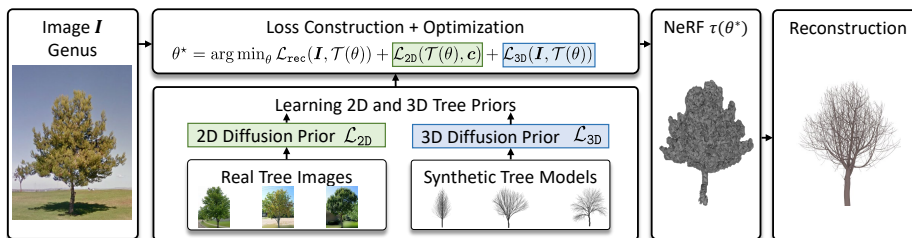


Fig. 3: The input to *Tree-D Fusion* is an RGB image of a tree and its genus. To perform shape reconstruction, we minimize the loss function w.r.t. the NeRF parameter θ . The loss function is constructed from two diffusion models, StableDiffusion with Lora and Zero123, trained on real tree images and synthetic 3D tree models. The output is an optimized NeRF $\tau(\theta^*)$, which is a detailed 3D tree envelope. We then populate the volume of $\tau(\theta^*)$ by markers based on the envelope and reconstruct trees by genus-conditioned space colonization algorithm.

with image I rendered from a NeRF with trainable parameters θ . However, since the rendered image is not noisy and the denoiser is trained on noisy images, naïvely applying the score model leads to an out-of-distribution (OOD) problem.

Perturb-and-Average Scoring (PAAS) [113] is proposed to address this problem

$$\text{PAAS}(\mathbf{x}, \sigma) = \mathbb{E}_{\mathbf{n} \sim \mathcal{N}(0, I)} (D(\mathbf{x} + \sigma \mathbf{n}; \sigma) - \mathbf{x}) / \sigma^2 \quad (3)$$

to be used in place of the score. PAAS adds noise to the rendered image such that the data becomes in-distribution relative to the denoiser’s training data.

We note that Eq. (3) is also known as score distillation in the concurrently proposed DreamFusion [84] under a different mathematical formulation. *Tree-D Fusion* employs SJC/DreamFusion to distill the scores of diffusion models tailored for tree data, enabling the reconstruction of simulation-ready 3D tree models.

4 Simulation Ready Tree Dataset

We construct a large dataset of 600,000 reconstructed 3D trees that faithfully represent the 3D geometry and can simulate their growth, including environmental response to light and obstacles (see Fig. 1).

The release of the AAD [7], featuring extensive tree images from Google Street View, aids the creation of 3D tree models from single images using a developmental model, capturing a specific instance of the tree’s growth. We leverage recent advances in diffusion models to reconstruct 3D trees. Given an input image, we first reconstruct a complete 3D volume of the tree using diffusion models as priors and fit the 3D volume by a developmental model conditioned on the tree genus. In particular, we use the space colonization model [75], also used to reconstruct 3D tree geometry [52] for realistic tree growth.

An overview of our approach is illustrated in Fig. 3. In the following, we describe how we train tree-specific diffusion models to be used as priors in Sec. 4.2, followed by the details of the 3D trees reconstruction in Sec. 4.3.

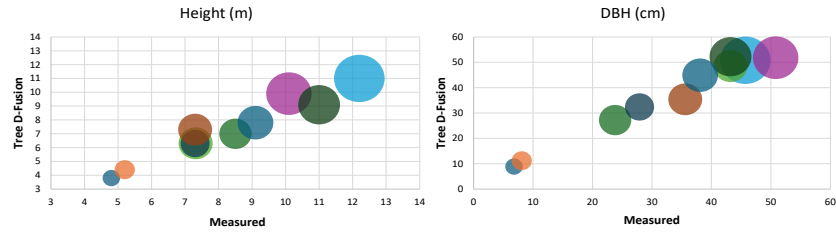


Fig. 4: Models of trees and the phenotypic characteristics derived from them: height and Diameter at Breast Height (DBH). The radius depicted illustrates the projected amount of shade.

4.1 Application

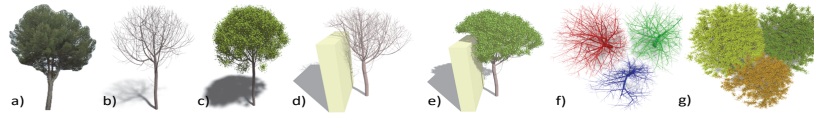


Fig. 5: The input image (a) is reconstructed into a digital twin (b-c) that responds to the environment such as the proximity to a wall (d-e). Growing three copies of the same tree shows their completion for space (f-g).

Our dataset, the first real-world 3D and 2D tree data at a large scale, seeks to standardize benchmarking and propel forestry deep learning, paralleling ImageNet’s role in computer vision. It helps study distant trees, predicts tree diameter, and supports landscape design and forestry management. It also integrates into mapping and AR apps for enhanced 3D tree visualization. Our dataset enables the creation of a 3D scene from a single Google Street View image, without the need for camera parameters, as demonstrated in Fig. 1. The image on the far left serves as the initial input, within which we delineate the three trees using dashed outlines. As time progresses, we model the development of these three trees. The final image depicts the mature trees, complete with foliage, ready for subsequent simulations like predicting urban heat island issues by estimating the total shadow area in the city.

The reconstructed simulation-ready 3D trees are computational models that can be used in many practical applications. An example in Fig. 5 (a) shows a tree from Google Street View images and its reconstructions (b-c). We then put a virtual object and simulate the tree’s growth close to a wall (d-e), and we copy it twice and show its competition for space (f-g).

We quantitatively evaluate other phenotypical traits, such as tree height, the shadow area (as pixel-to-meter



Fig. 6: Shadows cast by a tree without leaves (left) and with foliage (right).



Fig. 7: Low-quality images from the AAD are excluded in *Tree-D Fusion*.

ratio), and the Diameter at Breast Height (DBH) commonly used in the forest industry (see Fig. 4). On average, our estimated DBH is predicted to be 16.4% (4.0 cm) larger, with a standard deviation of 11.6% (2.7 cm). Additionally, our model predicts a height of 12.5% (1.0 m) shorter than the actual, with a standard deviation of 6.7% (0.6 m). An example in Fig. 6 shows the effect of shadow cast with and without leaves.

4.2 Training Tree-Specific 2D and 3D Diffusion Prior

Tree specific 2D prior. In order to obtain an effective 2D prior, we need a diffusion model $P_{2D}(\mathbf{I}|\mathbf{c};\sigma)$ that is trained to generate tree images \mathbf{I} given a corresponding text caption \mathbf{c} . We utilize the AAD [7], which includes around one million real tree images.

When training with these noisy labels, we observe that the diffusion model generalizes poorly and does not generate high-quality tree images. The dataset also contains images with heavy light exposure, extensive blurry images resulting from privacy filters, and low-image quality of distant trees (see Fig. 7). We semi-manually removed noisy images from the AAD to train this diffusion model. First, we divide the image into patches and measure the variance of the Laplacian in each patch. This approach measures image sharpness, averaging these values to exclude very blurry images for manual review. We will release this data along with this manuscript.

We train a tree-genus-conditioned diffusion model using Low-Rank Adaptation (LoRA) [41] on top of the pre-trained Stable Diffusion 1.5 version [97] as LoRA [41] is more data efficient. We create image captions using BLIP [53] for genus conditioning and prefix the word “tree” within the caption with the specific genus annotation from AAD. For example, BLIP outputs the caption of “A tree next to a house.”. The caption will be A {genus} tree next to a house, with {genus} replaced by the name of the genus.

Tree specific 3D prior. To train the 3D prior, $P_{3D}(\mathbf{I}_\pi|\mathbf{I},\pi;\sigma)$, we need data that contains paired images $(\mathbf{I}, \mathbf{I}_\pi)$ of the same tree at two different viewing angles, and their corresponding viewing angle π .

To obtain such data, we generated synthetic 3D tree models following Lee et al. [50] over 12 different types of trees: Pine, Maple, Oak, Acacia, Birch, Cabbage, Corkscrew, Elm, Hazel, Tulip, Walnut, and Willow (some shown in Fig. 8).

We produce 2,000 synthetic trees for each of the 12 tree categories. We train the diffusion model of Zero123 [56] from scratch for 1,000 iterations. We train this model to be genera-agnostic, i.e., not conditioned on the genera, because

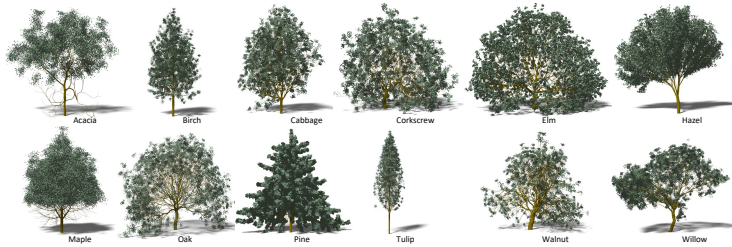


Fig. 8: Examples of synthetic trees used for training the 3D prior.

our synthetic trees do not cover all the genera in the AAD dataset. For example, there are more than 100 genera of trees in San Jose. For training this 3D prior, we use tree geometry information, which includes 12 rendered views, each 30 degrees apart, per tree 3D object.

4.3 Single Image Tree Shape Reconstruction

We propose a specialized approach for reconstructing trees based on the generic 3D shape reconstruction method of Magic123 [92]. We use our trained tree-specific diffusion models (Sec. 4.2) and eliminate the depth and normal smoothness loss terms from Magic123 that we found not applicable to tree reconstruction.

Problem formulation. Given an tree image \mathbf{I} with a text prompt \mathbf{c} , we formulate shape reconstruction as an optimization problem:

$$\min_{\theta} \mathcal{L}_{\text{rec}}(\mathbf{I}, \mathcal{T}(\theta)) + \alpha \mathcal{L}_{2\text{D}}(\mathcal{T}(\theta), \mathbf{c}) + \beta \mathcal{L}_{3\text{D}}(\mathbf{I}, \mathcal{T}(\theta)), \quad (4)$$

where $\mathcal{T}(\theta)$ denotes the underlying 3D model, Instant-NGP [70] in our case, with parameters θ , \mathcal{L}_{rec} denotes a reconstruct loss, $\mathcal{L}_{2\text{D}}$ denotes a loss based on the 2D diffusion prior, $\mathcal{L}_{3\text{D}}$ denotes a loss based on the 3D diffusion prior, and α and β are weights for each of the priors.

Loss definitions. *Reconstruction loss \mathcal{L}_{rec} :* As in Magic123, \mathcal{L}_{rec} encourages the rendered image $\mathcal{T}_{\pi}(\theta)$ from the view π to match the input image \mathbf{I} . The loss consists of two terms:

$$\mathcal{L}_{\text{rec}} = \lambda_{\text{rgb}} \|\mathcal{M}(\mathbf{I}) \odot (\mathbf{I} - \mathcal{T}_{\pi}(\theta))\|_2^2 + \lambda_{\text{mask}} \|\mathcal{M}(\mathbf{I}) - \mathbf{M}(\mathcal{T}_{\pi}(\theta))\|_2^2. \quad (5)$$

The first term calculates the difference between the pixels of the image within the segmentation mask of the tree. The symbol \odot stands for element-wise multiplication. The second term encourages the 3D occupancy to be consistent with the 2D tree mask obtained from the segmentation model. To do this, the algorithm extracts a mask \mathbf{M} that indicates the occupancy at each pixel by integrating the volume density from the 3D model $\mathcal{T}(\theta)$. We use the Segment Anything Model (SAM) for 2D instance segmentation masks [48].

2D diffusion prior loss \mathcal{L}_{2D} : A pre-trained diffusion model provides a prior for the non-observed parts of the tree. Specifically, we minimize the expected negative log probability of the diffusion model from different viewpoints π' , i.e.,

$$\mathcal{L}_{2D}(\mathcal{T}(\theta), \mathbf{c}) = \mathbb{E}_{\pi'} [-\log P_{2D}(\mathcal{T}_{\pi'}(\theta) | \mathbf{c}; \sigma)]. \quad (6)$$

Here, P_{2D} corresponds to the adapted diffusion model that specializes in tree image generation (Sec. 4.2).

3D diffusion prior loss \mathcal{L}_{3D} utilizes another pre-trained diffusion model to provide a prior for tree generation. In contrast to our 2D prior, the 3D prior’s diffusion model is trained to be 3D aware. Instead of conditioning on a text prompt, this diffusion model is trained to generate the image at a particular viewpoint π' of a reference image \mathbf{I} . Once again, we minimize the expected negative log probability of the diffusion model P_{3D} over different viewpoints π' , i.e.,

$$\mathcal{L}_{3D}(\mathbf{I}, \mathcal{T}(\theta)) = \mathbb{E}_{\pi'} [-\log P_{3D}(\mathcal{T}_{\pi'}(\theta) | \mathbf{I}, \pi'; \sigma)]. \quad (7)$$

With the loss functions defined, we then solve the optimization problem in Eq. (4) using Adam [47], using PAAS (Eq. 3) to compute the gradient of \mathcal{L}_{2D} and \mathcal{L}_{3D} w.r.t. to θ . We use a learning rate of 0.001 and optimize for 2,000 iterations without weight decay. We set $\lambda_{\text{rgb}} = 5.0$, $\lambda_{\text{mask}} = 20.0$, $\alpha = 1.0$ and $\beta = 8.0$.

4.4 3D Simulation-Ready Tree Reconstruction

After computing the tree’s 3D geometry, we extract its 3D volume and use existing methods to create a simulation-ready 3D tree. This algorithm is parameterized to generate the tree model. While various developmental models of trees (e.g., [90, 103]) could be used, these are not suited for fitting the tree into a particular shape. Instead, we use the realistic tree growth algorithm [52] that extends the previous space colonization algorithm [75] for fitting the tree shape. The algorithm [52] is conditioned on tree genus, which prescribes branching angles, internode lengths, branch radii, and other parameters [24, 117]. Environmental sensitivity is simulated by following existing works [81, 89]. Fig. 5 shows a reconstructed tree under varying conditions and Fig. 9 shows the simulation of the tree over time.

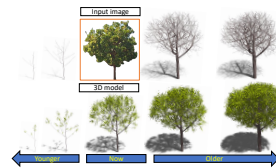


Fig. 9: The developmental model captures the real-tree shape (red frame) as a single instance in time. It allows us to look back and forward in time.

With the reconstructed 3D tree models from the Auto Arboist Dataset [7], we aim to share the first real-world large 3D and 2D tree dataset. Our dataset can standardize benchmarking in forestry deep learning, similar to ImageNet’s impact on computer vision. This will advance research, aid in studying inaccessible trees, manage urban heat islands, and support landscape design and forestry

Table 1: ICTree is the perceived realism scores of generated trees. *Tree-D Fusion* shows an average improvement of $44.83\% \pm 25.9\%$.

ICTree [83]↑	DGauss [109]	Magic123 [92]	Zero123 [56]	RBV [51]	Ours
Cupressus	0.46±0.03	0.45±0.05	0.46±0.04	0.75 ± 0.09	0.65 ± 0.07
Magnolia	0.46±0.04	0.45±0.03	0.46±0.03	0.616 ± 0.16	0.71 ± 0.08
Pinus	0.44±0.04	0.46±0.03	0.44±0.04	0.627 ± 0.16	0.69 ± 0.08
Ligustrum	0.45±0.04	0.45±0.03	0.45±0.04	0.665 ± 0.16	0.71 ± 0.07
Cinnamomum	0.45±0.02	0.47±0.02	0.46±0.02	0.658 ± 0.16	0.71 ± 0.08
Total	0.45±0.03	0.45±0.03	0.45±0.03	0.67 ± 0.12	0.71 ± 0.05

Table 2: LPIPS [132] between a frontal view of a tree envelope from RBV [51] and *Tree-D Fusion* shows $20.21\% \pm 8.89\%$ improvement.

LPIPS [132]↓	DGauss [109]	Magic123 [92]	Zero123 [56]	RBV [51]	Ours
Cupressus	0.68± 0.02	0.68± 0.02	0.69±0.02	0.53 ± 0.01	0.52 ± 0.02
Magnolia	0.73± 0.03	0.72± 0.03	0.73±0.03	0.59 ± 0.05	0.55 ± 0.06
Pinus	0.70± 0.01	0.69± 0.02	0.70±0.02	0.55 ± 0.04	0.50 ± 0.05
Ligustrum	0.74± 0.03	0.74± 0.02	0.74±0.03	0.59 ± 0.04	0.56 ± 0.05
Cinnamomum	0.74± 0.02	0.73± 0.02	0.74±0.02	0.59 ± 0.05	0.56 ± 0.05
Total	0.72 ± 0.03	0.71± 0.04	0.72 ± 0.03	0.58 ± 0.05	0.54 ± 0.06

management. Additionally, it can enhance map and augmented reality applications by displaying 3D trees.

5 Experiments and Applications

Implementation. We implemented our interactive procedural vegetation framework with C++ and Vulkan. We also added a path tracer based on Nvidia OptiX 8.0. The 2,000 synthetic trees (tree specific 3D priors $P_{3D}(\mathbf{I}_\pi | \mathbf{I}, \pi; \sigma)$ from Sec. 4.2) were generated for 113 hours on a laptop with NVIDIA RTX 3070. Using the 3D synthetic tree models, training the 3D prior model using Zero123 [56] took 89 hours using 6 NVIDIA RTX5000 GPUs, and training the 2D prior model using LoRA [41] took 9 hours using 4 NVIDIA RTX5000 GPU devices on the San Jose dataset. Reconstructing a 3D tree model from an RGB image takes 20 minutes using a single NVIDIA RTX 3090 GPU, and it would require ≈ 23 GPU-years on the single GPU to complete our 600,000 3D tree dataset.

Experiment setup. We conduct experiments using the AAD dataset [7] that contains 23 cities; among these, we use the city of San Jose to benchmark our system. We select the five most common tree genera, Cupressus, Magnolia, Pinus, Ligustrum, and Cinnamomum, from the city.

Baselines. We compare with the Radial Bounding Volume (RBV) [51], a tree-specific reconstruction using a single image with Deep Learning. We also provide comparisons with recent state-of-the-art single-image 3D object reconstruction

Table 3: CLIP-Similarity [94] between four views of a tree envelope shows an improvement of $45.34\% \pm 23.86\%$.

CLIP-Sim. [94]↑	DGauss. [109]	Magic123 [92]	Zero123 [56]	RBV [51]	Ours
Cupressus	0.47±0.05	0.49±0.03	0.38±0.04	0.54 ± 0.04	0.67 ± 0.06
Magnolia	0.46±0.04	0.46±0.03	0.39±0.04	0.54 ± 0.04	0.65 ± 0.07
Pinus	0.38±0.04	0.38±0.04	0.34±0.03	0.54 ± 0.04	0.64 ± 0.08
Ligustrum	0.45±0.03	0.45±0.03	0.35±0.06	0.54 ± 0.04	0.65 ± 0.06
Cinnamomum	0.40±0.08	0.42±0.05	0.34±0.05	0.54 ± 0.04	0.65 ± 0.06
Total	0.43±0.06	0.44±0.06	0.36±0.05	0.54 ± 0.04	0.63 ± 0.07

Table 4: A Chamfer distance (CD) from real LiDAR scanned trees shows an improvement of $32.62\% \pm 6.44\%$.

CD(10^{-2})↓	DGauss. [109]	Magic123 [92]	Zero123 [56]	RBV [51]	Ours
Tree 1	5.09	6.05	5.31	6.44	4.62
Tree 2	5.13	4.53	6.09	4.71	2.81
Tree 3	2.52	3.87	4.13	4.36	2.31
Total	4.25	4.82	5.18	5.17	3.25

methods based on diffusion models, including Zero123 [56], Magic123 [92], Dream-Gaussian [109].

Evaluation metrics. We evaluate our generated 3D tree models from the San Jose subset of AAD, which contains 4,400 trees with 133 genus-level categories. We use four evaluation metrics:

1. **ICTree** [83] evaluates tree plausibility, focusing on their unique traits. It assesses various geometric aspects, including branch straightness and segment ratios, using 29 parameters to analyze branching structures, as discussed in works [50, 134].
2. **LPIPS** [132] measures image similarity between the rendered front view of the reconstructed tree model (3D tree envelope) and the input image.
3. **CLIP-Similarity** [94] measures the similarity between images at a semantic level. We compute the average CLIP similarity between the input image and four renderings of the reconstructed tree at orientations 0° , 90° , 180° , and 270° . Higher values indicate the renderings are more similar to the input.
4. **Chamfer distance** measures the differences in an overall tree structure using lidar scanned tree point cloud data [58] and sampled point clouds from our generated 3D tree models.

5.1 Quantitative Results

Tab. 1 reports the ICTree [83] score to compare the perceptual realism of tree branching structures generated using *Tree-D Fusion*. We use 133 tree genera from the San Jose dataset, and on average, *Tree-D Fusion* scores 0.71 with a standard deviation of 0.05, which is an improvement in the average of $20.21\% \pm 8.89\%$ compared to the recent works [51, 56, 92, 109]. *Tree-D Fusion* achieves unparal-




Fig. 10: Single tree image reconstruction front-view results.

leled realism in generating tree structures by confirming 29 geometric branching attributes.

In Tab. 2, we use the pre-trained VGG [100] network to calculate a perceptual similarity between our generated and RBV tree front-views. We show the LPIPS score of five common tree genera from the San Jose dataset, including all the tree genera in Tab. 2. Ours shows $20.21\% \pm 8.89\%$ greater perceptual similarity to a conditioned image than the generated trees from the previous works [51, 56, 92, 109].

Tab. 3 shows the CLIP-Similarity between RBV [51] and *Tree-D Fusion*. We used four evenly distributed views around the object’s 360-degree to validate the consistency of the 3D-generated tree model. Using the five common tree genera from the San Jose dataset, *Tree-D Fusion* outperforms other works [51, 56, 92, 109] by an average of $45.34\% \pm 23.86\%$.

Tab. 4 compares the Chamfer Distance between LiDAR-scanned tree point clouds and our results using captured images [58]. We show that our generated tree models are reconstructed in the average of $32.62\% \pm 6.44\%$ closer to the real trees than other works [51, 56, 92, 109].

Based on the four reconstruction metrics, *Tree-D Fusion* shows it becomes a new state-of-the-art method to reconstruct a 3D tree model from a single unposed RGB image. Thus, our dataset contains 3D tree models that are perpetually, visually, and geometrically realistic.

5.2 Qualitative Results

We show the existing methods generate unsatisfying tree structures from a single image due to a lack of tree geometry information and visualize its front (Fig. 10) and top-views (Fig. 11). RBV [51] capture the approximated contour estimation, Zero123 [56] also captures contour information but includes holes in the Ginkgo































Genus	Input	RBV [51]	Zero123 [56]	Magic123 [92]	DGauss. [109]	Ours
Cupressus						
Ginkgo						
Ligustrum						
Magnolia						
Tristaniopsis						

Fig. 11: Single tree image reconstruction top-view results.

and Ligustrum genus, and the front-view is flat. Also, it generates non-tree-like shapes near the top area in the Tristaniopsis tree genus. Magic123 [92] fails to estimate tree depth, concentrating on the center and creating a gap between the inner and outer sides. Moreover, it does not generate overall balanced tree volumes and tends to look overly simplistic. DreamGaussian [109] captures the overall tree contour but it creates numerous holes around the center of 3D tree models with bubble-like artifacts visible from the top view. *Tree-D Fusion* captures the tree contours without any holes and generates tree-like balanced volumetric trees from the front and the top view.

5.3 Ablation Study

We conduct ablation studies using 3D Gaussian Splatting [46] on a tree and different α/β ratios. Synthesizing novel views for trees, as seen in Fig. 12, has proven challenging, hindering the extraction of high-quality 3D boundaries. Successful synthesis relies on pre-trained data, as shown by our use of 24,000 synthetic 3D tree models. This highlights the need for large-scale real image-based 3D tree reconstruction datasets in forestry to utilize deep learning efficiently. We evaluate the tree reconstruction metrics, CLIP-Similarity [94], LPIPS [132] and ICTree [83] using four different α/β ratios. Empirically, from Tab. 5, we found that the ratio around 0.1 achieves a good performance.




Fig. 12: Novel view failed using 3DGS.

6 Conclusion

We introduce *Tree-D Fusion*, the first large dataset of 3D reconstructed tree models enabled by a diffusion-based approach to construct 3D tree models from

Table 5: A reconstruction quality comparison between different α/β ratios.

α/β	CLIP [94]↑	LPIPS [132]↓	ICTree [83]↑
0.01	0.53 ± 0.06	0.58 ± 0.07	0.52 ± 0.07
0.1	0.60 ± 0.07	0.59 ± 0.06	0.55 ± 0.06
1	0.59 ± 0.07	0.58 ± 0.07	0.50 ± 0.05
10	0.57 ± 0.07	0.59 ± 0.06	0.47 ± 0.03

single images. We use images from the AAD [7] and synthetic trees to train diffusion models that reconstruct 3D tree volumes that are then completed by a developmental model to provide functional-structural tree model for individual real trees. Our generated dataset has the potential to enable diverse and impactful applications in 3D tree phenotyping at a previously unprecedented scale.

This is the first large dataset of its kind, but it comes with a few **limitations**. First, the 3D model shapes are constrained by the capabilities of the simulator. Second, single-image reconstruction, limited by a single viewpoint and missing information, is best suited for symmetric trees. An example in Fig. 13 epitomizes this case. The same asymmetric tree from two different angles yields significantly varied 3D models. High branch density can also hinder precise reconstruction due to the complexity of tree segmentation. Finally, predicting a branch structure obscured by leaves with 100% fidelity to nature is an ongoing challenge due to its complex geometry.

**Fig. 13:** Single-image reconstruction fails with asymmetric trees, producing different 3D models from varying views.

There are many possible avenues for future work. First, a comprehensive 3D tree database could enable the study of light and environmental effects on reconstruction. Additionally, incorporating real tree models with algorithms for partially occluded parts during training could improve detection significantly, as single-image segmentations struggle to disentangle nearby tree crowns.

Acknowledgements

This work is based upon efforts supported by the EFFICACI grant under USDA NRCS #NR233A750004G044. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government or NRCS. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. This project was also sponsored by USDA NIFA, Award #2023-68012-38992 grant “Promoting Economic Resilience and Sustainability of the Eastern U.S. Forests” and by NIFA-AFRI award #2024-67021-42879, “Improving Forest Management Through Automated Measurement of Tree Geometry”.

References

1. Aiteanu, F., Klein, R.: Hybrid tree reconstruction from inhomogeneous point clouds. *The Visual Computer* (2014) [3](#)
2. Aiteanu, F., Klein, R.: Exploring shape spaces of 3d tree point clouds. *Comput. Graph.* (2021) [3](#)
3. Alonzo, M., Bookhagen, B., Roberts, D.A.: Urban tree species mapping using hyperspectral and lidar data fusion. *Remote Sensing of Environment* (2014) [3](#)
4. Ando, R., Ozasa, Y., Guo, W.: Robust surface reconstruction of plant leaves from 3d point clouds. *Plant Phenomics* (2021) [3](#)
5. Aubry-Kientz, M., Laybros, A., Weinstein, B., Ball, J.G., Jackson, T., Coomes, D., Vincent, G.: Multisensor data fusion for improved segmentation of individual tree crowns in dense tropical forests. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2021) [2](#)
6. Aygunes, B., Cinbis, R.G., Aksoy, S.: Weakly supervised instance attention for multisource fine-grained object recognition with an application to tree species classification. *ISPRS Journal of Photogrammetry and Remote Sensing* (2021) [3](#)
7. Beery, S., Wu, G., Edwards, T., Pavetic, F., Majewski, B., Mukherjee, S., Chan, S., Morgan, J., Rathod, V., Huang, J.: The auto arborist dataset: A large-scale benchmark for multiview urban forest monitoring under domain shift. In: *CVPR* (2022) [2](#), [3](#), [5](#), [7](#), [9](#), [10](#), [14](#)
8. Bergseng, E., Ørka, H.O., Næsset, E., Gobakken, T.: Assessing forest inventory information obtained from different inventory approaches and remote sensing data sources. *Annals of forest science* (2015) [3](#)
9. Blum, J.: Contribution of ecosystem services to air quality and climate change mitigation policies: the case of urban forests in barcelona, spain. In: *Urban Forests*, Apple Academic Press (2017) [1](#)
10. Bradley, D., Nowrouzezahrai, D., Beardsley, P.: Image-based reconstruction and synthesis of dense foliage. *ACM TOG* (2013) [3](#)
11. Brandt, M., Tucker, C.J., Kariryaa, A., Rasmussen, K., Abel, C., Small, J., Chave, J., Rasmussen, L.V., Hiernaux, P., Diouf, A.A., et al.: An unexpectedly large count of trees in the west african sahara and sahel. *Nature* (2020) [2](#)
12. Branson, S., Wegner, J.D., Hall, D., Lang, N., Schindler, K., Perona, P.: From google maps to a fine-grained catalog of street trees. *ISPRS Journal of Photogrammetry and Remote Sensing* (2018), URL <http://dx.doi.org/10.1016/j.isprsjprs.2017.11.008> [3](#)
13. Bratman, G.N., Anderson, C.B., Berman, M.G., Cochran, B., De Vries, S., Flanders, J., Folke, C., Frumkin, H., Gross, J.J., Hartig, T., et al.: Nature and mental health: An ecosystem service perspective. *Science advances* (2019) [1](#)
14. Cárdenas-Donoso, J.L., Ogayar, C.J., Feito, F.R., Jurado, J.M.: Modeling of the 3d tree skeleton using real-world data: a survey. *IEEE TVCG* (2022) [3](#)

15. Chaurasia, G., Beardsley, P.: Editable parametric dense foliage from 3d capture. In: CVPR (2017) [3](#)
16. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: ECCV (2022) [4](#)
17. Chen, X., Jiang, K., Zhu, Y., Wang, X., Yun, T.: Individual tree crown segmentation directly from uav-borne lidar data using the pointnet of deep learning. *Forests* (2021) [3](#)
18. Chen, Z., Tagliasacchi, A., Zhang, H.: BSP-Net: Generating compact meshes via binary space partitioning. In: CVPR (2020) [4](#)
19. Cheng, Z.L., Zhang, X.P., Chen, B.Q.: Simple reconstruction of tree branches from a single range image. *Journal of computer science and technology* (2007) [3](#)
20. Choy, C.B., Xu, D., Gwak, J., Chen, K., Savarese, S.: 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In: ECCV (2016) [4](#)
21. Crowther, T.W., Glick, H.B., Covey, K.R., Bettigole, C., Maynard, D.S., Thomas, S.M., Smith, J.R., Hintler, G., Duguid, M.C., Amatulli, G., et al.: Mapping tree density at a global scale. *Nature* (2015) [2](#)
22. Culman, M., Rodríguez, A.C., Wegner, J.D., Delalieux, S., Somers, B.: Deep learning for sub-pixel palm tree classification using spaceborne sentinel-2 imagery. In: *Remote Sensing for Agriculture, Ecosystems, and Hydrology XXIII* (2021) [3](#)
23. Dai, A., Ruizhongtai Qi, C., Nießner, M.: Shape completion using 3D-encoder-predictor cnns and shape synthesis. In: CVPR (2017) [4](#)
24. De Reffye, P., Edelin, C., Françon, J., Jaeger, M., Puech, C.: Plant models faithful to botanical structure and development. *ACM Siggraph Computer Graphics* (1988) [9](#)
25. Deng, C., Jiang, C., Qi, C.R., Yan, X., Zhou, Y., Guibas, L., Anguelov, D., et al.: Nerdi: Single-view nerf synthesis with language-guided diffusion as general image priors. In: CVPR (2023) [4](#)
26. Deussen, O., Hanrahan, P., Lintermann, B., Měch, R., Pharr, M., Prusinkiewicz, P.: Realistic modeling and rendering of plant ecosystems. In: *Proc. Computer graphics and interactive techniques* (1998) [3](#)
27. Du, S., Lindenbergh, R., Ledoux, H., Stoter, J., Nan, L.: Adtree: Accurate, detailed, and automatic modelling of laser-scanned trees. *Remote Sensing* (2019) [3](#)
28. Duggal, S., Pathak, D.: Topologically-aware deformation fields for single-view 3d reconstruction. In: CVPR (June 2022) [4](#)
29. Duncanson, L., Dubayah, R.: Monitoring individual tree-based change with airborne lidar. *Ecology and evolution* (2018) [2](#)
30. Eisenman, T.S., Churkina, G., Jariwala, S.P., Kumar, P., Lovasi, G.S., Pataki, D.E., Weinberger, K.R., Whitlow, T.H.: Urban trees, air quality, and asthma: An interdisciplinary review. *Landscape and urban planning* (2019) [1](#)
31. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: CVPR (2017) [4](#)

32. Fassnacht, F.E., Latifi, H., Stereńczak, K., Modzelewska, A., Lefsky, M., Waser, L.T., Straub, C., Ghosh, A.: Review of studies on tree species classification from remotely sensed data. *Remote Sensing of Environment* (2016) [3](#)
33. Graves, S.J., Marconi, S., Stewart, D., Harmon, I., Weinstein, B.G., Kanazawa, Y., Scholl, V.M., Joseph, M.B., McClinchy, J., Browne, L., et al.: Data science competition for cross-site delineation and classification of individual trees from airborne remote sensing data. *bioRxiv* (2021) [3](#)
34. Guo, J., Jiang, H., Benes, B., Deussen, O., Zhang, X., Lischinski, D., Huang, H.: Inverse procedural modeling of branching structures by inferring l-systems. *ACM TOG* (2020) [3](#)
35. Guo, J., Xu, S., Yan, D.M., Cheng, Z., Jaeger, M., Zhang, X.: Realistic procedural plant modeling from multiple view images. *IEEE TVCG* (2020) [3](#)
36. Hackenberg, J., Spiecker, H., Calders, K., Disney, M., Raunonen, P.: Simpletree—an efficient open source tool to build tree models from tls clouds. *Forests* (2015) [3](#)
37. Hädrich, T., Banuti, D.T., Pałubicki, W., Pirk, S., Michels, D.L.: Fire in paradise: Mesoscale simulation of wildfires. *ACM TOG* (2021) [3](#)
38. Hädrich, T., Benes, B., Deussen, O., Pirk, S.: Interactive modeling and authoring of climbing plants. In: *Computer Graphics Forum*, Wiley Online Library (2017) [3](#)
39. Harmon, I., Marconi, S., Weinstein, B., Bai, Y., Wang, D.Z., White, E., Bohlman, S.: Improving rare tree species classification using domain knowledge. *IEEE Geoscience and Remote Sensing Letters* (2023) [3](#)
40. Hartling, S., Sagan, V., Sidike, P., Maimaitijiang, M., Carron, J.: Urban tree species classification using a worldview-2/3 and lidar data fusion approach and deep learning. *Sensors* (2019) [3](#)
41. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models (2021) [7](#), [10](#)
42. Hu, S., Li, Z., Zhang, Z., He, D., Wimmer, M.: Efficient tree modeling from airborne lidar point clouds. *Computers & Graphics* (2017) [3](#)
43. Hyvärinen, A.: Estimation of non-normalized statistical models by score matching. *J. Mach. Learn. Res* (2005) [4](#)
44. Isokane, T., Okura, F., Ide, A., Matsushita, Y., Yagi, Y.: Probabilistic plant modeling via multi-view image-to-image translation. In: *CVPR* (2018) [3](#)
45. Jucker, T., Fischer, F.J., Chave, J., Coomes, D.A., Caspersen, J., Ali, A., Loubota Panzou, G.J., Feldpausch, T.R., Falster, D., Usoltsev, V.A., et al.: Tallo: A global tree allometry and crown architecture database. *Global change biology* (2022) [3](#)
46. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (2023) [13](#)
47. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *Proc. ICLR* (2015) [9](#)

48. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. In: ICCV (2023) [8](#)
49. Lang, N., Jetz, W., Schindler, K., Wegner, J.D.: A high-resolution canopy height model of the earth. *Nature Ecology & Evolution* (2023) [2](#)
50. Lee, J.J., Li, B., Benes, B.: Latent l-systems: Transformer-based tree generator. *ACM Trans. Graph.* (2023) [3](#), [7](#), [11](#)
51. Li, B., Kałużny, J., Klein, J., Michels, D.L., Pałubicki, W., Benes, B., Pirk, S.: Learning to reconstruct botanical trees from single images. *ACM Transaction on Graphics* (2021) [2](#), [3](#), [4](#), [10](#), [11](#), [12](#), [13](#)
52. Li, B., Klein, J., Michels, D.L., Benes, B., Pirk, S., Pałubicki, W.: Rhizomorph: The coordinated function of shoots and roots. *ACM TOG* (2023) [2](#), [4](#), [5](#), [9](#)
53. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: ICML (2022) [7](#)
54. Lindenmayer, A.: Mathematical models for cellular interaction in development. *Journal of Theoretical Biology* **Parts I and II** (1968) [3](#)
55. Lines, E.R., Fischer, F.J., Owen, H.J.F., Jucker, T.: The shape of trees: Reimagining forest ecology in three dimensions with remote sensing. *Journal of Ecology* (2022) [3](#)
56. Liu, R., Wu, R., Hoorick, B.V., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: ICCV (2023) [2](#), [4](#), [7](#), [10](#), [11](#), [12](#), [13](#)
57. Liu, W., Sun, J., Li, W., Hu, T., Wang, P.: Deep learning on point clouds and its application: A survey. *Sensors* (2019) [2](#)
58. Liu, Y., Guo, J., Benes, B., Deussen, O., Zhang, X., Huang, H.: Treepartnet: Neural decomposition of point clouds for 3d tree reconstruction. *ACM Trans. Graph.* (2021) [11](#), [12](#)
59. Liu, Z., Wu, K., Guo, J., Wang, Y., Deussen, O., Cheng, Z.: Single image tree reconstruction via adversarial network. *Graphical Models* (2021) [3](#)
60. Livny, Y., Pirk, S., Cheng, Z., Yan, F., Deussen, O., Cohen-Or, D., Chen, B.: Texture-lobes for tree modelling. *ACM TOG* (2011) [3](#)
61. Livny, Y., Yan, F., Olson, M., Chen, B., Zhang, H., El-Sana, J.: Automatic reconstruction of tree skeletal structures from point clouds. *ACM SIGGRAPH Asia* (2010) [3](#)
62. Locke, D.H., Hall, B., Grove, J.M., Pickett, S.T., Ogden, L.A., Aoki, C., Boone, C.G., O’Neil-Dunne, J.P.: Residential housing segregation and urban tree canopy in 37 us cities. *npj Urban Sustainability* (2021) [2](#)
63. Loving, R., Agarwal, A., Sherstyuk, I., Parra, S.: A network fusion model pipeline for multi-modal, deep learning for tree crown detection (2020) [2](#)
64. Mahmud, J., Price, T., Bapat, A., Frahm, J.M.: Boundary-aware 3D building reconstruction from a single overhead image. In: CVPR (2020) [4](#)
65. Makowski, M., Hädrich, T., Scheffczyk, J., Michels, D.L., Pirk, S., Pałubicki, W.: Synthetic silviculture: Multi-scale modeling of plant ecosystems. *ACM TOG* (2019) [2](#), [3](#)

66. Marconi, S., Weinstein, B.G., Zou, S., Bohlman, S.A., Zare, A., Singh, A., Stewart, D., Harmon, I., Steinkraus, A., White, E.P.: Continental-scale hyperspectral tree species classification in the united states national ecological observatory network. *Remote Sensing of Environment* (2022) [3](#)
67. McDonald, R.I., Kroeger, T., Zhang, P., Hamel, P.: The value of us urban tree cover for reducing heat-related health impacts and electricity consumption. *Ecosystems* (2020) [1](#)
68. McQuillan, I., Bernard, J., Prusinkiewicz, P.: Algorithms for inferring context-sensitive l-systems. In: *Proc. Unconventional Computation and Natural Computation* (2018) [3](#)
69. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3D reconstruction in function space. In: *CVPR* (2019) [4](#)
70. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG* (2022) [8](#)
71. Nash, C., Ganin, Y., Eslami, S.A., Battaglia, P.: PolyGen: An autoregressive generative model of 3D meshes. In: *Proc. ICML* (2020) [4](#)
72. Neubert, B., Franken, T., Deussen, O.: Approximate image-based tree-modeling using particle flows. *ACM SIGGRAPH 2007 papers* (2007) [2](#), [3](#)
73. Niese, T., Pirk, S., Albrecht, M., Benes, B., Deussen, O.: Procedural urban forestry. *ACM TOG* (2022) [2](#), [3](#)
74. Nowak, D.J., Greenfield, E.J.: Us urban forest statistics, values, and projections. *Journal of Forestry* (2018) [1](#)
75. Palubicki, W., Horel, K., Longay, S., Runions, A., Lane, B., Měch, R., Prusinkiewicz, P.: Self-organizing tree models for image synthesis. *ACM TOG* (2009) [2](#), [5](#), [9](#)
76. Pan, L., Liu, L., Condon, A.G., Estavillo, G.M., Coe, R.A., Bull, G., Stone, E.A., Petersson, L., Rolland, V.: Biomass prediction with 3d point clouds from lidar. In: *CVPR* (2022) [2](#)
77. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: *CVPR* (2019) [4](#)
78. Peng, K., Islam, R., Quarles, J., Desai, K.: Tmvnet: Using transformers for multi-view voxel-based 3d reconstruction. In: *CVPR* (2022) [4](#)
79. Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., Geiger, A.: Convolutional occupancy networks. In: *ECCV* (2020) [4](#)
80. Pirk, S., Niese, T., Hädrich, T., Benes, B., Deussen, O.: Windy trees: Computing stress response for developmental tree models. *ACM TOG* (2014) [3](#)
81. Pirk, S., Stava, O., Kratt, J., Said, M.A.M., Neubert, B., Měch, R., Benes, B., Deussen, O.: Plastic trees: interactive self-adapting botanical tree models. *ACM TOG* (2012) [2](#), [9](#)
82. Pleșoianu, A.I., Stupariu, M.S., Șandric, I., Pătru-Stupariu, I., Drăguț, L.: Individual tree-crown detection and species classification in very high-resolution remote sensing imagery using a deep learning ensemble model. *Remote Sensing* (2020) [3](#)

83. Polasek, T., Hrusa, D., Benes, B., Čadík, M.: Ictree: Automatic perceptual metrics for tree models. *ACM Trans. Graph.* (2021) [10](#), [11](#), [13](#), [14](#)
84. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. In: *ICLR* (2023) [4](#), [5](#)
85. Prusinkiewicz, P.: Graphical applications of l-systems. In: *Proc. on Graphics Interface* (1986) [3](#)
86. Prusinkiewicz, P., Cieslak, M., Ferraro, P., Hanan, J.: Modeling plant development with l-systems. *Mathematical modelling in plant biology* (2018) [3](#)
87. Prusinkiewicz, P., Hammel, M.S., Mjolsness, E.: Animation of plant development. In: *ACM Siggraph Computer Graphics*, ACM Press (1993) [3](#)
88. Prusinkiewicz, P., Hanan, J.: Visualization of botanical structures and processes using parametric l-systems. In: *Scientific Visualization and Graphics simulation'90* (1990) [3](#)
89. Prusinkiewicz, P., James, M., Měch, R.: Synthetic topiary. In: *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques* (1994) [3](#), [9](#)
90. Prusinkiewicz, P., Lindenmayer, A.: *The algorithmic beauty of plants*. Springer Science & Business Media (2012) [2](#), [3](#), [9](#)
91. Puliti, S., Pearse, G., Surový, P., Wallace, L., Hollaus, M., Wielgosz, M., Astrup, R.: For-instance: a uav laser scanning benchmark dataset for semantic and instance segmentation of individual trees (2023) [2](#)
92. Qian, G., Mai, J., Hamdi, A., Ren, J., Siarohin, A., Li, B., Lee, H.Y., Skorokhodov, I., Wonka, P., Tulyakov, S., et al.: Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843* (2023) [2](#), [4](#), [8](#), [10](#), [11](#), [12](#), [13](#)
93. Quan, L., Tan, P., Zeng, G., Yuan, L., Wang, J., Kang, S.B.: Image-based plant modeling. *ACM Siggraph* (2006) [3](#)
94. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *ICML* (2021) [11](#), [13](#), [14](#)
95. Raj, A., Kaza, S., Poole, B., Niemeyer, M., Ruiz, N., Mildenhall, B., Zada, S., Aberman, K., Rubinstein, M., Barron, J., et al.: Dreambooth3d: Subject-driven text-to-3d generation. In: *ICCV* (2023) [4](#)
96. Reche-Martinez, A., Martin, I., Drettakis, G.: Volumetric reconstruction and interactive rendering of trees from photographs. *ACM Siggraph Computer Graphics* (2004) [3](#)
97. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *CVPR* (2022) [4](#), [7](#)
98. Roskopf, E., Morhart, C., Nahm, M.: Modelling shadow using 3d tree models in high spatial and temporal resolution. *Remote Sensing* (2017) [3](#)
99. Safonova, A., Tabik, S., Alcaraz-Segura, D., Rubtsov, A., Maglinets, Y., Herrera, F.: Detection of fir trees (*abies sibirica*) damaged by the bark beetle in unmanned aerial vehicle images with deep learning. *Remote sensing* (2019) [3](#)

100. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015) [12](#)
101. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. In: NeurIPS (2019) [4](#)
102. Štáva, O., Benes, B., Měch, R., Aliaga, D.G., Krištof, P.: Inverse procedural modeling by automatic generation of l-systems. *Computer Graphics Forum* (2010) [3](#)
103. Stava, O., Pirk, S., Kratt, J., Chen, B., Měch, R., Deussen, O., Benes, B.: Inverse procedural modelling of trees. In: *Comput. Graph. Forum* (2014) [2](#), [9](#)
104. Stewart, D., Zare, A., Marconi, S., Weinstein, B.G., White, E.P., Graves, S.J., Bohlman, S.A., Singh, A.: Randcrowns: A quantitative metric for imprecisely labeled tree crown delineation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2021) [3](#)
105. Sumbul, G., Cinbis, R.G., Aksoy, S.: Fine-grained object recognition and zero-shot learning in remote sensing imagery. *IEEE Transactions on Geoscience and Remote Sensing* (2017) [3](#)
106. Sun, C., Sun, M., Chen, H.T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In: CVPR (2022) [4](#)
107. Tan, P., Fang, T., Xiao, J., Zhao, P., Quan, L.: Single image tree modeling. *ACM TOG* (2008) [3](#)
108. Tan, P., Zeng, G., Wang, J., Kang, S.B., Quan, L.: Image-based tree modeling. *ACM Siggraph Computer Graphics* (2007) [3](#)
109. Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. arXiv preprint arXiv:2309.16653 (2023) [2](#), [4](#), [10](#), [11](#), [12](#), [13](#)
110. Tompalski, P., Coops, N.C., White, J.C., Goodbody, T.R., Hennigar, C.R., Wulder, M.A., Socha, J., Woods, M.E.: Estimating changes in forest attributes and enhancing growth projections: a review of existing approaches and future directions using airborne 3d point cloud data (feb, 10.1007/s40725-021-00135-w, 2021). *CURRENT FORESTRY REPORTS* (2021) [3](#)
111. Ventura, J., Pawlak, C., Honsberger, M., Gonsalves, C., Rice, J., Love, N.L., Han, S., Nguyen, V., Sugano, K., Doremus, J., et al.: Individual tree detection in large-scale urban environments using high-resolution multispectral imagery. arXiv preprint arXiv:2208.10607 (2022) [3](#)
112. Wang, D., Cui, X., Chen, X., Zou, Z., Shi, T., Salcudean, S., Wang, Z.J., Ward, R.: Multi-view 3d reconstruction with transformers. In: CVPR (2021) [4](#)
113. Wang, H., Du, X., Li, J., Yeh, R.A., Shakhnarovich, G.: Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In: CVPR (2023) [4](#), [5](#)
114. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2Mesh: Generating 3D mesh models from single RGB images. In: ECCV (2018) [4](#)
115. Wang, Y., Pyörälä, J., Liang, X., Lehtomäki, M., Kukko, A., Yu, X., Kaartinen, H., Hyypä, J.: In situ biomass estimation at tree and plot levels: What did data record and what did algorithms derive from terrestrial and aerial point clouds in boreal forest. *Remote Sensing of Environment* (2019) [2](#)

116. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolific-dreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. arXiv preprint arXiv:2305.16213 (2023) [4](#)
117. Weber, J., Penn, J.: Creation and rendering of realistic trees. In: Proceedings Computer Graphics and Interactive Techniques, SIGGRAPH '95 (1995) [9](#)
118. Wegner, J.D., Branson, S., Hall, D., Schindler, K., Perona, P.: Cataloging public objects using aerial and street-level images — urban trees. In: CVPR (2016) [3](#)
119. Weinstein, B., Marconi, S., Graves, S., Zare, A., Singh, A., Bohlman, S., Magee, L., Johnson, D., Townsend, P., White, E.: Capturing long-tailed individual tree diversity using an airborne multi-temporal hierarchical model. bioRxiv pp. 2022–12 (2022) [3](#)
120. Weinstein, B.G., Marconi, S., Bohlman, S., Zare, A., White, E.: Individual tree-crown detection in rgb imagery using semi-supervised deep learning neural networks. Remote Sensing (2019) [3](#)
121. Weinstein, B.G., Marconi, S., Bohlman, S.A., Zare, A., Singh, A., Graves, S.J., White, E.P.: A remote sensing derived data set of 100 million individual tree crowns for the national ecological observatory network. Elife (2021) [3](#)
122. Worchel, M., Diaz, R., Hu, W., Schreer, O., Feldmann, I., Eisert, P.: Multi-view mesh reconstruction with neural deferred shading. In: CVPR (2022) [4](#)
123. Xie, D., Wang, X., Qi, J., Chen, Y., Mu, X., Zhang, W., Yan, G.: Reconstruction of single tree with leaves based on terrestrial lidar point cloud data. Remote sensing (2018) [3](#)
124. Xie, H., Yao, H., Zhang, S., Zhou, S., Sun, W.: Pix2vox++: Multi-scale context-aware 3d object reconstruction from single and multiple images. IJCV (2020) [4](#)
125. Xu, D., Jiang, Y., Wang, P., Fan, Z., Wang, Y., Wang, Z.: Neurallift-360: Lifting an in-the-wild 2d photo to a 3d object with 360deg views. In: CVPR (2023) [4](#)
126. Xu, H., Gossett, N., Chen, B.: Knowledge and heuristic-based modeling of laser-scanned trees. ACM TOG (2007) [3](#)
127. Xu, Q., Wang, W., Ceylan, D., Mech, R., Neumann, U.: Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. In: NeurIPS (2019) [4](#)
128. Yang, B., Wang, S., Markham, A., Trigoni, N.: Robust attentional aggregation of deep feature sets for multi-view 3d reconstruction. IJCV (2020) [4](#)
129. Yin, K., Huang, H., Long, P., Gaissinski, A., Gong, M., Sharf, A.: Full 3d plant reconstruction via intrusive acquisition. In: Computer Graphics Forum (2016) [3](#)
130. Yu, Q., Ji, W., Pu, R., Landry, S., Acheampong, M., O’Neil-Dunne, J., Ren, Z., Tanim, S.H.: A preliminary exploration of the cooling effect of tree shade in urban landscapes. International Journal of Applied Earth Observation and Geoinformation (2020) [1](#)

131. Zhang, K., Hu, B.: Individual urban tree species classification using very high spatial resolution airborne multi-spectral imagery using longitudinal profiles. *Remote Sensing* (2012) [3](#)
132. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *CVPR* (2018) [10](#), [11](#), [13](#), [14](#)
133. Zhang, Y., Onda, Y., Kato, H., Feng, B., Gomi, T.: Understory biomass measurement in a dense plantation forest based on drone-sfm data by a manual low-flying drone under the canopy. *Journal of Environmental Management* (2022) [2](#)
134. Zhou, X., Li, B., Benes, B., Fei, S., Pirk, S.: Deeptree: Modeling trees with situated latents. *IEEE TVCG* (2023) [11](#)
135. Zhou, X., Wang, W., Di, L., Lu, L., Guo, L.: Estimation of tree height by combining low density airborne lidar data and images using the 3d tree model: A case study in a subtropical forest in china. *Forests* (2020) [3](#)
136. Ziter, C.D., Pedersen, E.J., Kucharik, C.J., Turner, M.G.: Scale-dependent interactions between tree canopy cover and impervious surfaces reduce daytime urban heat during summer. *Proceedings of the National Academy of Sciences* (2019) [1](#)