XIAOCHEN ZHOU, Purdue University, USA BOSHENG LI, Purdue University, USA BEDRICH BENES, Purdue University, USA SONGLIN FEI, Purdue University, USA SÖREN PIRK, Adobe Research, USA



Fig. 1. Illustration of our DeepTree (*DeepTree*) generation approach: we use a common procedural model to generate a dataset of different species (a). As part of our algorithm, we train a neural network pipeline (b) to predict the topological and geometric properties of branching patterns based on a situated latent space. This means that branch features are predicted locally and by only considering a single node (c). Given a node, our network pipeline is able to predict branchlets, the branching structures of the immediate successors of a node – thereby mimicking the developmental process of a tree model (d, e). Eventually, our algorithm will automatically terminate when a complex tree model has been generated (f). Branchlets are shown with colors in (c), (d), and (e).

In this paper, we propose DeepTree, a novel method for modeling trees based on learning developmental rules for branching structures instead of manually defining them. We call our deep neural model "situated latent" because its behavior is determined by the intrinsic state -encoded as a latent space of a deep neural model- and by the extrinsic (environmental) data that is "situated" as the location in the 3D space and on the tree structure. We use a neural network pipeline to train a situated latent space that allows us to locally predict branch growth only based on a single node in the branch graph of a tree model. We use this representation to progressively develop new branch nodes, thereby mimicking the growth process of trees. Starting from a root node, a tree is generated by iteratively querying the neural network on the newly added nodes resulting in the branching structure of the whole tree. Our method enables generating a wide variety of tree shapes without the need to define intricate parameters that control their growth and behavior. Furthermore, we show that the situated latents can also be used to encode the environmental response of tree models, e.g., when trees grow next to obstacles. We validate the effectiveness of our method by measuring the similarity of our tree models and by procedurally generated ones based on a number of established metrics for tree form.

CCS Concepts: • Computing methodologies \rightarrow Generative and developmental approaches; Shape analysis; Computer vision problems.

Additional Key Words and Phrases: Botanical Tree Models, Deep Learning, Shape Modeling, Generative Methods

1 INTRODUCTION

Vegetation is ubiquitous in almost all environments, ranging from vast outdoor landscapes to indoor spaces. Consequently, models of trees and plants serve as essential assets in applications such as games and movies, mixed reality, architecture and urban planning, agriculture, and forestry, or even the training of autonomous agents. For many applications, it is of utmost importance to define plant models with a high degree of geometric detail to produce highquality renderings or model interactions with plant models, e.g., to plan a path around them. Procedural models have proven to effectively capture the wide variety of plant forms found in Nature and efficiently generate large sets of models.

Procedural and developmental modeling of vegetation in CG has been used in diverse ways, often combined with other approaches, such as modeling the environmental response of trees [32], the reconstruction of tree models from images or point sets [8, 19, 23], inverse procedural modeling [45], by focusing on interactive modeling with user-defined sketches [24], or by leveraging neural networks that compose trees of procedurally generated branching patterns [22]. The breadth of these approaches is a testament to the complexity of realistically modeling plants. However, despite the progress, generating trees with procedural models remains an open and challenging research problem. Our key inspiration comes from Nature's ability to compress and encode structural and behavioral properties into a highly compact form, such as DNA. Bud development is intrinsically determined by the DNA, and extrinsically by the environment (light, gravity, nutrients), and the position in the tree hierarchy [1, 49]. The bud development is a response to these factors; it can grow to a new location, create lateral buds, become dormant etc. The ability of the DNA to encode complex interactions that lead to shape as an emergent 3D structure (gene expression) is one of the key open problems in science [17]. Inspired by this problem, we seek to develop a generative model that would capture the resulting shape of the behavior. While engineering such behavior is challenging because each tree node experiences varying conditions, deep learning is suitable for learning and encoding the bud's response to such conditions.

We propose a novel deep neural model that we call "situated latent" because its behavior is determined by the intrinsic state encoded as a latent space of a deep neural model- and by the extrinsic (environmental) data that is "situated" as the location in the 3D space and on the tree structure. We use a large dataset of 3D tree models to train a pipeline of neural networks on branchlets - atomic branching structures consisting of a node and its immediate children. We learn a situated latent space, a representation that makes predictions for branch growth based on where a branch node is located in the growth space of a tree model. Once trained, the neural networks can predict the topology (number of children) and geometry (position and the branch width of children) for a single input node. The deep neural model is a compact representation that encodes the tree response based on the environment (extrinsic) and the positional information within the tree structure (intrinsic). Similar to the DNA, we encode the entire tree and its environmental response as a single compact deep neural representation.

A new tree is generated iteratively by executing the network on individual nodes of the branching structure. Starting from the root node, the single network is executed for each node, and it predicts the number and positions of its children. The predicted child nodes are added to the branch graph, and in the next iteration, the network predicts their immediate children again. Similar to procedural modeling algorithms that are also implemented recursively or iteratively, this allows producing complex branching structures only with a few iterations. While *DeepTree* learns the response of branchlets, it is *not* a reconstruction method as the generated trees are non-deterministic and only share low-level common properties of the entire training set.

Our results show that the situated latent space can encode and generate diverse branching patterns of multiple tree species by successfully predicting branch growth's primary and secondary attributes (e.g., branching angles, internode length, branch width, etc.). The network also predicts the termination of branches (no immediate children) for twigs toward the outer regions of the tree crown, which then terminates the development of a model. Moreover, we show that the situated latent space can also be used to encode the environmental response of trees, i.e., the adaptation of branching structures according to tropisms and obstacles in the model's environment.

We are not aware of any deep neural models capable of generating 3D biological tree shape. Therefore, we validate the effectiveness of our method in encoding tree form by comparing tree models generated with a state-of-the-art procedural model, and our qualitative results show that our method generates branching structures with almost identical geometric properties. We measure the similarity of branching structures based on histograms over the geometric properties at different levels of fully developed tree models. We also use a recently proposed perceptual metric to assess the quality of generated tree models further.

Examples of a procedurally generated tree and a *DeepTree* model are shown in Fig. 1. The *DeepTree* model is generated by iteratively querying a neural network on the current set of terminal nodes. The network predicts new branch nodes – thereby mimicking the growth process – that are then added to the branch graph. A complex tree model can be generated with a few iterations.

In summary, our contributions are: (1) we propose a novel method for learning to predict the topological and geometric properties of branchlets based on a novel neural network pipeline; (2) we show that the learned situated latent space can be used to iteratively generate tree models, which provides a novel way to encode and generate tree form; (3) we validate our method through state-of-theart qualitative and quantitative metrics.

2 RELATED WORK

Due to the importance of vegetation in several application domains, modeling trees and plants has received significant research attention over the past decades (see review [35]). The goal of the early approaches was to focus on faithfully modeling the morphology of branching structures. Approaches include fractals [2], repetitive patterns [31], environmentally sensitive automata [10], and particlesystems [43]. L-systems [41] - as one of the most fundamental procedural approaches - enable to define the tree growth as a set of parallel production rules and has been shown to allow for capturing a wide variety of plant forms, even when considering the environmental response of plants [26]. However, while L-systems are a powerful mechanism for generating plant form, defining production rules to generate complex trees is an elaborate task that requires a substantial amount of knowledge, which often renders their use for applications impractical. Early procedural models [51] aim to address these shortcomings by combining rule-based approaches with geometric modeling and user interaction [21]. However, they are more specialized and do not support modeling the same range of branching patterns. Finally, the goal of inverse procedural modeling [45, 46] is to automatically find the parameter values of procedural models of plants through optimization schemes or to find the production rules [11].

Recent procedural methods for plants emphasize more principled representations [18, 57] or the underlying biological processes of plant development. Stava et al. [45], for example, add biological priors to their procedural model to provide more nuanced control for shape-defining parameters. Other methods focus on the selforganization of plants [32], their environmental response [26, 39], the competition for resources [4, 10], or modeling features of specific types of plants [53]. The availability of more powerful graphics hardware-enabled approaches that use procedural models of plants to simulate the animation of growth [15, 37], their physical response to wind [12, 38] and fire [36], advanced material properties



Fig. 2. Overview: We use branchlets of tree models generated with a common procedural model to train our neural networks (a). To generate a tree, we iteratively use our neural network pipeline to predict the topology and geometry of new nodes. A classification network predicts the number of new children for a node that triggers a cascade of regression networks predicting the signatures of each child by using the output of the previous regression network (e.g., Regressor 2 uses the parent node signature as well as the output signature of Regressor 1). Starting with the root node, we maintain a list of active nodes. We pop an active node from the list in each iteration and trigger the neural networks to predict its children. The children are then added to the branch graph and the list of active nodes (b, c). The generation of a tree model stops when the classification network predicts zero children for an active node, which happens individually for different branches (d).

of wood [50, 59], or to find efficient representations of tree models [28]. More recently, module-based tree representations -similar to our branchlets- have gained popularity to enable simulating complex vegetation-related phenomena, such as wildfires [14] or climatic gradients [33].

Sketch-based and data-driven approaches often use procedural models by guiding the generation according to the defined input data to generate convincing branching structures. User-defined sketches provide the intriguing advantage that plants can be modeled with direct control to meet artistic requirements [6, 16, 30, 52]. Longay et al. [24] propose an advanced framework to efficiently produce complex tree models that can even run on mobile hardware. Datadriven approaches, on the other hand, focus on solving a reconstruction task to generate plant models to faithfully match a captured tree that can either be represented as one or several images [5, 27, 47, 48], videos [20], or point clouds [23, 55]. More recently, it has also been recognized that neural networks are a powerful means to guide procedural modeling by learning bounding volumes to aid the reconstruction from single images [19] by decomposing point clouds [22] into semantically meaningful patterns, or by learning parameters for the placement of plants [29]. Moreover, synthesizing large data collections from a few examples appears to be an increasingly important pursuit [3]. Finally, some approaches also focus on defining trees based on partially defined branching structures - procedurally generated - to enable a level of detail schemes [58] or the processing of large collections of plants [25].

The work of Estrada et al. [9] follows a similar objective to ours. Their approach also uses a parametric tree-growth model to regularize topology estimation for tree-like structures. However, despite these advances, none of the existing approaches for trees and plants uses neural networks to encode the rules and parameters for locally encoding branching patterns for botanical tree models.

3 OVERVIEW

Our goal is to learn all possible combinations of branching patterns as they occur from the perspective of a single node in a branch graph. Instead of defining these patterns manually based on parameters and rules, we employ a set of neural networks that learn to predict the child nodes based on the developmental (growth stage, trunk vs. twig, etc.) and environmental context (obstacles, light, gravity, etc.) of a parent node in the branch graph. We train the networks on *branchlets*, a tree node, and its immediate children (Fig.2, a). A node is defined by all its attributes that we refer to as *node signature*. The node signature serves as input to our network pipeline, and the output is the signatures of up to three child nodes that are then used for the generation of the next nodes.

The *DeepTree* neural network pipeline consists of a classification network that predicts the topology of a *branchlet*, i.e., how many children a node needs to have. We train individual regression networks for each child node that predict the attribute values of the child node's signature, i.e., their thickness, length, and growth direction. The input of the first regression network is the signature of the active parent node, and it predicts the signature of the first child. We then query the second regression network to predict the signature of the second child. This network receives the signature of the active parent node and the predicted signature of the first child from the previous step as input. This process is continued for the number of children the classifier predicted.

The key idea behind our method is that the networks learn to predict meaningful attributes of child nodes only based on where the parent node – along with all its signature attributes – is located within the tree model. Training the network with branchlets enables the network to learn a latent space that encodes which topological and geometric configurations are reasonable for certain locations in the tree's growth space. As an example, for the illustrated tree graph (Fig. 2, a), the networks first need to predict the trunk – each node only has one child – for the first few iterations. The further the generation of a tree model progresses, the more the networks



Fig. 3. Network architecture: our neural network pipeline consists of a classification network to predict 0-3 children for a given node (a). Depending on the number of children, we use a cascade of regression networks (c) to predict output signatures for each child. Each regression network is defined as an MLP (b). For each network, we use the current node's signature, the local voxel space for the node, and the global voxel space as inputs. We add a point encoder network to obtain a feature from local point sets and to enable predicting branchlets for reconstruction tasks.

need to predict more diverse and finer branching patterns up to the point where only small twigs are generated (Fig. 2, b-d). Eventually, the network also has to predict the termination of branches, which stops the generation of a tree model (Fig. 2, d).

The objective of our pipeline is not to reconstruct a given tree model. Instead, our goal is to propose an alternative method for defining the rules and parameters commonly used for generating branch graphs with procedural models. A procedural model generates a distribution of tree models based on value ranges defined for the various parameters. A parameter value is obtained at each growth step by sampling from the range. For complex procedural models, the ranges of parameter values can also vary over time. For example, to capture the impact of gravitropism, branching angles may vary from a more narrow parameter range for the trunk to a wider one for the thinner branches in the tree crown. Our method aims to learn these distributions of branching patterns and how they vary across a tree model. We use neural networks to imitate the developmental process locally.

4 DEEP MODELING OF BRANCHING STRUCTURES

Here we introduce our framework, including a formal definition of the used representation of trees, node signatures, the used neural networks, and the implementation of the environmental response of tree models.

4.1 Tree Representation

We represent tree models as a directed acyclic graph $G = \{N, E\}$, where N are the nodes and E the edges. Given an oriented edge $e_i = (n_s, n_t)$ with a starting node n_s and the end node n_t , a hierarchical relationship is generated. The node with no predecessor is the root node n_{root} . Branches are defined as chains of a varying number of edges $C = \{e_1, e_2, \ldots, e_m\}$; the length of a chain is denoted as m. We assign to each node the Hack [13] ordering number o (also called Gravellius ordering) that assigns o = 0 to the trunk and one order higher for each branch recursively (hierarchical ordering). Formally,



Fig. 4. Illustration of our tree representation: a branchlet is defined as a node with its immediate successors (0-3 children). Branchlets can have different geometries, depending on where they are located within a tree model.

a node is defined as a tuple of attributes

$$n = (p, t, d_r, d_s, q_p, l_p, o, m_p, a, v, s),$$
(1)

where *p* represents the 3D position of a node in Euclidean space, *t* is the branch thickness, d_r is the distance of a node to the root node in the node hierarchy, d_s is the distance to the start of the branch (chain), q_p is the normalized quaternion representing the orientation of the node w.r.t. the parent node, l_p is the length of the edge to the parent node, m_p is the number of children of the parent node. Additionally, we add the global attributes: age *a*, gravitropism *v*, and species *s*, where *a* and *s* are integers. We refer to the tuple *n* as the *input signature* of a node, which is used as the input of our neural networks.

A branchlet is defined as a set of edges of a node and its children $B = \{(n_i, n_{i+1}), ..., (n_i, n_{i+c})\}$. The number of children of a node is denoted c = 1, ..., 3. A branchlet can also have no children $B = \{(n_i, \emptyset)\}$, which is required for terminal nodes. In Fig. 4, we show renderings of branchlets with different topological and geometric properties and where they are located within a tree model.

4.2 Neural Network Pipeline

Most procedural modeling algorithms aim to produce a graph G to define the properties of a given tree species, which is accomplished by defining the parameters and rules for generating node topology and geometry. For example, one may define the angle at which a branch grows away from its parent branch as a range of values that are then sampled randomly when the procedural model is evaluated to generate a tree model.

Learning Topology: Contrary to common procedural approaches, our goal is to learn the generation of *branchlets* – the branching pattern for a single node $n \in N$ of a branch graph *G*. To successfully generate local branching patterns, we need to simultaneously predict the number of children (topology) and their geometric attributes (e.g., position, thickness, and branching angle). Jointly learning the topological and geometrical properties of a branch graph is challenging as most deep network architectures do not support generating outputs of arbitrary lengths, and the regression of high-dimensional outputs is difficult. Therefore, we learn the number of children and their geometric features separately. To predict the topology of a branchlet, we train a classification network to predict the number of children a given parent node should have as probabilities. More specifically, we aim to learn the mapping:

$$f_{class}(n): N \to \mathcal{X},$$
 (2)

where $x \in X$ denotes a tuple of probabilities $x = (p_1, p_2, p_3)$ for each child node of $n \in N$. The probabilities determine whether or not to grow a child node. The goal of this network is only to predict the topology of a branchlet.

Learning Geometry: To learn the geometric properties of a child node, we define the tuple u = (q, t, l), where q denotes a normalized quaternion representing the direction of a child node w.r.t. the parent node, t is the branch thickness of a node, and l the length of the edge between the parent node and the child node. We refer to the tuple u as the *output signature* of a node and use it as the label for training the regression networks.

To generate the children of a parent node, we need to predict the geometric properties of up to three tuples u_1, u_2, u_3 for each potential child node. We use a cascade of up to three regression networks to predict these tuples. The first network receives only the *input signature* of a parent node as the input and is defined as:

$$f_{reg}^1(n): N \to \mathcal{U}_1, \tag{3}$$

where $u_1 \in \mathcal{U}_1$ is the output tuple of the first child for a given node *n*. For the subsequent children, the second and third network also receive the predicted output signatures of the previous steps. Consequently, the networks for the second and third child can be defined as:

$$f_{req}^2(n, u_1) \quad : N \to \mathcal{U}_2, \tag{4}$$

$$f_{reg}^3(n, u_1, u_2) \quad : N \to \mathcal{U}_3, \tag{5}$$

where $u_2 \in \mathcal{U}_2$ denotes the output tuple of the second child, and $u_3 \in \mathcal{U}_3$ is the output tuple of the third child for node *n* respectively. Please note that while we did not experiment with predicting more than three child nodes for each parent node, the architecture can be extended as outlined here.

Input: Root node nroot Output: Tree graph G. 1 Procedure: Add n_{root} to G. 2 Add *n_{root}* to list of active nodes *L*. 3 while |L| > 0 do 4 $n_a \leftarrow \text{Pop node from L.}$ 5 $x \leftarrow$ Predict probability of each child $f_{class}(n_a)$. $u_1 \leftarrow$ Predict output signature of child 1 7 with $f_{reg}^1(n_a)$. 8 $u_2 \leftarrow \text{Predict output signature of child } 2$ 9 with $f_{req}^2(n_a, u_1)$. 10 $u_3 \leftarrow$ Predict output signature of child 3 11 with $f_{reg}^3(n_a, u_1, u_2)$. 12 **for** $p_i \in x$ **do**: 13 $\alpha \leftarrow$ Generate a random number in the range [0, 1]. 14 | if $\alpha > p_i$: 15 | Make node n_i with u_i and add to G. 16 | Push n_i to L. 17 end 18 end 19 end 20

ALGORITHM 1: Geometry Generation.

4.3 Iterative DeepTree Generation

We iteratively apply our network for the nodes of the branch graph of a tree model to generate the 3D tree structure. Given the signature of a node n, the classification network predicts the number of children that need to be generated for this node. We then use the cascade of regression networks to predict up to three child nodes. The output of each of the regression networks is an *output signature* containing the quaternion q, the thickness t, and the length l, for each child node of n. Given the position p of a node, we can compute the position of each child node as

$$p_{child} = p + l \cdot d_{front}.$$
 (6)

The growth direction d_{front} and the up direction d_{up} are:

$$d_{front} = \mathbf{R} \times d_{init_front},$$

$$d_{up} = \mathbf{R} \times d_{init_up},$$
 (7)

where *R* is the rotation matrix computed from the quaternion *q*. The vectors of the initial directions d_{init_front} and d_{init_up} are [0, 0, -1] and [0, 1, 0] respectively. The vectors d_{front} and d_{up} are used to generate the branch mesh. The thickness is stored with the child node and later also used to generate the branch mesh.

To generate a new tree model, a user defines the signature of the node n_{root} . The root node is added to a list of active nodes *L*. As long as the list holds nodes, an active node n_a is popped from the list. The classification network then predicts which of the three children should be generated; the result is returned as the tuple $x = (p_1, p_2, p_3)$. If a child node needs to be generated for n_a , the cascade of regression networks predicts the tuple u_i for each child.



Fig. 5. Eight different tree species modeled with the procedural model of Stava et al. [45] (left) and with our *DeepTree* approach (right). The learned situated latent spaces allow us to generate the topological and geometric details to model a wide variety of different tree species.

The child nodes are added to the branch graph, and the list of active nodes and the network pipeline can be queried for each node in the subsequent iterations. The generation of a branch terminates when the classification network predicts zero children for a node. Pseudocode for our algorithm is shown in Alg. 1 and an illustration of the iterative construction of a tree model is shown in Fig. 2.

4.4 Environmental Response

We use a voxel space to learn how trees grow around obstacles in their environment [10]. We represent obstacles as bounding boxes and generate an occupancy grid *Y* that encodes where obstacles are located in the environment. We generate the occupancy grid by testing the bounding boxes against the voxels. We use the computed occupancy grids in two different ways. First, we obtain the $3 \times 3 \times 3$ neighborhood of cells for each node as 27 binary occupancy values denoted as *Z*. We then add them to our node signature $n = n \cup \{Z\}$. As shown in Fig. 3 (a-b), we can use these values as additional input for our classification and regression networks (magenta boxes). Second, to obtain a global feature of the environment, we use an additional encoder consisting of three 3D convolutional layers. The global voxel space is a one-hot 3D vector where obstacle voxels are marked as one and empty voxels as zero; it has a size of 32^3 voxels. The tree's root is set in the middle of the voxel space. We extract the global feature through the additional encoder and added it to the local feature before the dense layers for the output heads of the networks. Adding the global feature of the environment to our neural networks allows us to generate environmentally sensitive tree models that can be placed arbitrarily into the environment. In contrast to more complex models of environmental sensitivity [26, 32, 39], this approach provides a lightweight way to generate realistic branching structures for tree models placed in the vicinity of obstacles.

4.5 Point-based Modeling

To provide additional control for the reconstruction and sketch-based generation of tree models, we have equipped our neural network architecture with an additional point-cloud encoder network (Fig. 3 a-b). We locally sample points from a point cloud

of a procedurally generated tree model to train our architecture for tree reconstruction with point clouds. For each node of a branchlet, we define a sample radius r_s to compute a sphere around a node

Fig. 6. Modeling age and gravitropism: by adding priors for age and gravitropism to our node tuple *n*, our neural network pipeline is able to model realistic branching structures. Here we compare a procedurally generated Oak tree at different developmental stages (a-d) and the corresponding *DeepTree* generated ones (e-h). Our method can also capture the differences in branching patterns when considering different environmental factors, such as gravitropism. The procedurally generated tree models (i-l) show the same structural properties as the *DeepTree* generated ones (m-p).

(inset figure). Given a point cloud of a tree, we select all points within the sphere; the obtained local point set is denoted as *P*. We then associate the local point set to our input signature $n = n \cup \{P\}$ and use it as input to our architecture to obtain a point feature. The point feature is then added to the signature embedding for the classification and regression networks. This training setup aims to condition the prediction of child node attributes based on the captured local point sets. The implementation of our point encoder follows PointNet [42]. For most of our experiments, a radius of $r_s = 2.0$ provides the best results for reconstructing the tree models.

5 IMPLEMENTATION

Our framework is implemented as two components: the first is an interactive framework to efficiently generate large collections of tree models with the procedural model from [45]. Second, we use Pytorch for developing and training our classification and regression neural networks. Below we describe how we used these two components of our framework to generate training data and how we trained our neural network pipeline.

The deep learning model was trained on a single Nvidia RTX A5000, and it took 6 hours for the regressor and 4 hours for the

classifier. Our rendering was performed on a desktop computer equipped with an Intel(R) Core(TM) i9-9900K and Nvidia 3090 RTX.

5.1 Dataset Preparation

We used the 29-dimensional parameter space of the procedural model from Stava et al. [45] to model eight distinct tree species: Beech, Corkscrew, Maple, Oak, Pine, Tulip, Walnut, and Willow (Fig. 5, left). An explanation of the parameters of the procedural model, along with the used parameter values, can be found in the Appx. Tab. 4 and Tab. 6. We then generate 500 unique tree models for each species and normalize their positions (attribute p) into a unit cube. We split these tree models into 400 for training and 100 for validation. The validation data is used for identifying optimal hyperparameters for our neural network architectures (e.g., learning rate, number layers, etc.). The procedural model and our DeepTree algorithm only generate branch graphs. We generate a mesh out of generalized cylinders based on the branch graph and the stored thickness values to render images of trees with high visual quality. The voxel space that we use for representing the environment has $32 \times 32 \times 32$ voxels. We compute the bounding box of each obstacle in the environment (e.g., walls) and check whether they occupy the corresponding voxels. This results in a 3D occupancy grid.

Fig. 7. Ablation study of the classification network: here, we show the impact of different attribute configurations of our node signature as an ablation study. Training only with the attribute position (p) generates a dense branching structure, and the network fails to terminate (a). Even if additional attributes such as the root distance (d_s) and the number of children of the parent node (m_p) are added, the model still fails to terminate (b). c) and d): only after adding additional attributes, such as the thickness (t) or branch order (o), the model generates more convincing branching structures, although some branches are still generated into the ground. Please note that we only use the per-node attributes for this experiment and not the set of global attributes such as species (s), age (a), and gravitropism (v). The tree species we used for this experiment is Beech (Fig.5).

Fig. 8. Ablation study of regression networks: the selection of node attributes for training the regression networks can profoundly impact the generated branch graphs. Only using a few attributes leads to slim and sparse branching structures (a, b). Adding more attributes to the input signature, such as the position or the order (c, d), will lead to more realistic branching patterns. For this experiment, we used a classification network with all parameters. The model with all node attributes is shown in Fig.7 (e).

To train the point cloud network, we generate view-dependent (partial) point clouds for our procedurally generated tree meshes. To generate the point scans, we mimic a LiDAR scanner that we place on a $10m \times 10m$ ground plane outside the bounding box of a tree that is located at the origin of the plane. We then cast rays toward the tree by sampling points on a sphere located at the scanner position. Each ray is tested against the surface mesh of a procedural tree model of our dataset. This procedure generates partial point scans of our tree models with 8k-35k points. We then sample this point cloud for each branchlet to obtain a local point cloud with a maximum of 512 points as described in Sec. 4.5.

5.2 Training

We train the classification and regression neural networks of our pipeline based on *branchlets*. We generate the branchlets by decomposing the generated branch graph *G*. Any node in the tree graph (with 0-3 children) can serve as a branchlet. We traverse each graph in our training dataset and select each node along with its immediate children to generate a set of branchlets \mathcal{B} . A parent node of a branchlet with all its attributes becomes the *input signature*. We then compute the quaternion *q* and the length *l* based on how child

nodes are connected to the parent node of the branchlet and store these attributes along with the thickness value *t* as the *output signature* that is used as a label for training the regression networks. For training the classification network, we simply obtain the number of children attached to the parent node as the label.

Classification Network Training: we train our classification network for 150 epochs jointly on the 3,200 tree models of all species. As illustrated in Fig. 3, our classification network has three output heads. The last layer uses a sigmoid activation function. For each head, we aim to simultaneously predict the probability of whether a particular child needs to be generated (0=child is omitted, 1=child is added). We train for this objective with an MSE loss, where the label for adding a child is 1 and 0 to omit a child node. Depending on the species the network is trained on we obtain a child classification accuracy of 83 - 93%.

To train this network more robustly, we use data augmentation. We add random noise to the position with a mean of zero and std of 0.0001. Additionally, we ensure a balanced sampling of nodes of the same properties. Considering that nodes with a different number of children appear with different frequencies (e.g., most nodes may have one child), we balance the number of nodes of each type during training to avoid network overfitting. For example, nodes at the trunk are rarely present in our dataset compared to the many twig nodes in the tree crown. Balancing the nodes helps to avoid overfitting. Training the classification network takes about four hours.

Regression Network Training: each regression network is trained individually for 200-300 epochs on the 3,200 tree models of all species. We trained the regression networks for up to 12 hours. Each regression network has two output heads: the first generates outputs for thickness and length. We train this head with an MSE loss. The second head predicts a quaternion and is trained with a cosine similarity loss. To avoid overfitting, we also add random noise to the position for training the regression networks with a mean of zero and std of 1e-4. Depending on the species the network is trained on, we obtain an MSE of 1.2e-4 - 3e-4 (branch length and thickness) and a cosine similarity of 0.04 - 0.20 (quaternion). The networks are trained with the Adam optimizer, a learning rate of

Fig. 9. Regrowth: our method can be used to regrow tree models to thereby adapt them to changing environmental conditions. A fully developed tree model (a) generated with our method is pruned and can be developed anew (b) to a fully grown tree model (c) by iteratively applying our *DeepTree* pipeline to the outermost nodes in the branch graph. If the tree model is placed in the vicinity of an obstacle (d) and then regrown (e), our method supports generating environmentally sensitive tree models.

Fig. 10. Modeling with environmental sensitivity: when provided with the environment as occupancy grids, our situated latent space is able to generate convincing branching structures to mimic the growth response of trees to obstacles in their vicinity (a-d). Similar to existing approaches, such as Pirk et al. [39] (e), our method is also able to generate environmentally-aware branching structures.

Fig. 11. Reconstruction of tree models from point cloud data: we sample our procedural trees to generate point clouds. We then train our neural networks with local point clouds and use them to reconstruct unseen tree models from the validation dataset.

1e-4, and a batch of 512. An illustration of our network architectures – also including the used number of units in each layer – is shown in Fig. 3. Please note that each of the used regression networks in the cascade of networks (Fig. 3, c) is trained individually.

5.3 Rendering

We use the generated branch graph from the procedural model and *DeepTree* to compute a surface mesh of branches from generalized cylinders. Leaves are represented as textured quads placed along with the outermost branches with additional procedural parameters

for their placement (e.g., phyllotaxis) [7]. All models were rendered with a path tracer written with the Nvidia OptiX 7.4 API [34] tracing 512 rays per pixel.

6 RESULTS, EVALUATION, AND APPLICATIONS

Here, we discuss experiments on how our framework can be used to generate tree models of different species, ages, or tropisms. We also discuss applications for our method and show how we validate our framework based on geometric and perceptual metrics.

6.1 Results

Fig. 5 shows a qualitative comparison of eight procedurally generated tree species and the same eight species generated with *DeepTree*. Our method can capture the wide variety of distinct features across the eight tree species, including the fine nuances in shape and geometric detail. Generating these species with a procedural model is a significant modeling effort that requires manually specifying the procedural model and carefully fine-tuning all the parameters for each species. In contrast, our method learns the branching patterns of all species from data by training our network pipeline.

The results shown in Fig. 6 indicate that our neural network pipeline and the training of situated latent spaces enable capturing the growth response of tree models and the impact of gravitropism. By adding global priors for age a and gravitropism v to the node tuple n, our pipeline can faithfully generate the structural details for modeling different age stages (e-h) and varying degrees of gravitropism (m-p). Compared to the procedurally generated branching structures (a-h) and (j-l), our method can generate branching patterns with almost identical topological and geometric features.

We conducted ablation studies to analyze the impact of different attributes in our node tuple *n* during training of the classification and regression networks. Fig. 7 demonstrates the impact of using more or fewer parameters during the training of the classifier. For the result shown in Fig. 7 (a), we only used the position p attribute for our classification network. Training the network only on this attribute prevents the network from terminating, which results in very dense and unrealistic branching patterns. Adding the distance to the branch root d_s and the number of child nodes of the parent node m_p by itself also does not enable training of more realistic branching structures (Fig. 7 b) as the model still fails to predict the correct number of children. However, by adding the branch thickness t, the network starts to successfully predict the termination of branches, which leads to more realistic, still dense branching structures (Fig. 7 c). When also using the other attributes, the model starts predicting branchlet topologies and geometries that lead to realistic branching patterns (Fig. 7 d, e). Please note that for (a) and (b), we manually stopped the iterative generation of the models as the neural networks would not automatically terminate. In Fig. 8, we show a similar ablation study for the regression networks. For this experiment, we used a classification network with all node attributes (Fig 7, e). Only using a few attributes for training the regression networks leads to slim and sparse branching structures (a, b). Adding more attributes generates more realistic branching patterns (c, d).

The result in Fig. 9 shows the usefulness of our method for content creation. After we generate a fully developed tree model with our method (a), we can prune and regenerate the branch graph (b) to another fully developed tree model (c). If the model is placed next to an obstacle and then regenerated, it adapts to its new environment (d) and develops an adapted and realistic branch graph (e). Generating a new tree model by pruning a branch graph can be repeated indefinitely. However, please note that each generated model is unique – our method cannot regenerate precisely the same model.

Fig. 12. Reconstruction of a real tree from captured point cloud data: *Deep-Tree* was trained on synthetic data and generated plausible tree models from real point clouds.

Fig. 13. Comparison to AdTree [8, Fig 11 e]: we use an existing point cloud from their paper (a-b) and show their reconstruction (c). *DeepTree* generates a similar tree model (d-f).

Fig. 10 shows that our situated latent spaces can encode the environmental response of a tree model. By adding the local and global occupancy features, the neural network pipeline can be trained to model the growth response of tree models to obstacles in their environment. For this experiment, we also generated a dataset of 500 scenes (400 training, 100 validation) where obstacles were randomly placed in the vicinity of the procedurally generated trees. The procedural model accounts for the obstacles by computing the availability of light in its environment - branches then grow into regions where more light is available. By training with tree models and the occupancy grid from the scene, the neural network pipeline is able to mimic the modeling of the environmental response. Fig. 10 (a)-(d) compares our results to the method of Pirk et al. [39] (e), where environmental sensitivity is modeled based on an inverse algorithm. Our DeepTree learning-based approach can generate a similar response of adapted branching structures for trees grown in the vicinity of obstacles.

Fig. 14. Different viewpoints of tree models show that *DeepTree* generates tree models without dominant artifacts (e.g., half-reconstructed model) and captures tree phyllotaxis.

Table 1. Tree model compression: our framework can regenerate the same tree model from a root node by setting the same random seed. We show the memory footprint of the root node (Seed Size = 52B) along with the weights required for our neural networks (NN Size) and compare it with the average memory footprint of a single tree model, which can either be stored as a skeletal graph (Graph Size) or surfaces mesh (Mesh Size). We show the average and standard deviation of the size comparison for the graph and mesh size, having the seed and the NN as the base. The graph size is, on average, $1.58 \times$ larger, and the mesh size is about $20.81 \times$.

Species	Seed + NN Size	Graph Size	# Nodes	Mesh Size	# Vertices
Beech	913 kB	1.85 MB	5.7k	24.08 MB	115k
Corkscrew	913 kB	1.20 MB	3.7k	15.66 MB	75k
Maple	913 kB	1.36 MB	4.2k	17.68 MB	85k
Oak	913 kB	0.78 MB	2.4k	10.18 MB	49k
Pine	913 kB	1.81 MB	5.5k	23.86 MB	114k
Tulip	913 kB	1.84 MB	5.6k	24.27 MB	115k
Walnut	913 kB	0.74 MB	2.3k	10.97 MB	53k
Willow	913 kB	1.60 MB	4.9k	20.96 MB	101k
Size (avg,std)	[1,0]	[1.58,0.46]	-	[20.81,5.78]	-

6.2 Applications

Fig. 11 shows a tree reconstruction result. We use our training data set to generate point clouds of procedurally generated tree models. We then train our neural network pipeline on local point clouds to provide the network with a point feature (Sec. 4.5). Training with the point feature allows us to generate branching structures that follow the scanned points, which also works for point clouds of real trees, as shown in Figs. 12 and 13.

We show multiple views of the same tree models in Fig. 14 to demonstrate that our method generates plausible tree models from all view directions. Once trained, our framework enables the generation of tree models of a specific tree species. Moreover, *Deep-Tree* generalizes the branchlet and not the entire tree. Thus, each generated model is different. This makes our approach suitable for generating large datasets of new models, as shown in the selection of randomly sampled tree models in Fig. 18. Fig. 15 shows that our method can replicate the growth of more constrained tree shapes. For this result, we generated a dataset of 500 tree models for each of the shown shapes, including cone (a), cube (b), and ring (c).

Finally, in Tab. 1 we show that our method can also be used for compressing generated tree models. We compare the memory footprint of a tree model generated by our method to common representations for tree models, such as skeletal graphs and surface meshes. *DeepTree* can deterministically re-generate the same tree model by using the same random seed and replacing the stochastic probability with a given number. As we only need to store the root node and the weights for our neural networks, our method allows us to compress tree models with a lightweight memory footprint that is even smaller than most skeletal graphs.

6.3 Evaluation

We are not aware of any deep neural generative model for 3D tree geometry. Therefore we compare our algorithm to the state-of-theart procedural model of [19, 45]. We used 500 procedurally generated tree models (P) for each species, and we generated another 500 trees using *DeepTree* (DT). We then compare their geometry and perceived level of realism.

Geometry: We validate the geometric structure of the trees generated by our *DeepTree* approach by comparing their geometric properties to the ground truth, and the results are shown in Fig. 17 and Tab. 6 (Appx.). The overall branch length varies by 11% for all trees, the number of generated branches by 18%, branching distance by 12%, the number of generated nodes by 11%, and the angles (in order from the trunk) by 6%, 4%, and 6%. The variations are minor, and it is essential to note that there is a great difference in visual importance for the presented features, e.g., the angle of the branches coming from the trunk has a strong effect on the overall tree shape, the branching angle of the small branchlets is not so important. This observation is further supported by the perceptual metrics.

Perceptual Metrics: While the geometric comparison shows that our trees have geometric traits comparable to the ground truth, we also want to answer whether the generated trees are perceived as realistic by humans. While a standard and tedious approach to this validation is to run a user study, we validate the generated trees by recently introduced deep neural perceptual metrics ICTree [40] that estimate the visual realism of the 3D tree models. ICTree has been trained by a response of 4,000 human subjects, and it provides a value between zero to one that corresponds to the perceived visual realism of the input tree. The authors introduced view-independent ICTreeF and image-based ICTreeI metrics. We use the ICTreeF that uses geometric tree properties.

Tab. 2 shows the average and the standard deviation of the perceived realism of the ground truth and the generated trees. The overall perceived realism metrics of all generated trees is 0.47 ± 0.11 that is 4.3% different from the ground truth of 0.49 ± 0.11 .

Runtime Performance: Tab. 3 shows the comparison of the runtime performance of procedurally generated tree models and *DeepTree.* The reported numbers represent the average measurements of each method for ten models for each species. we report the measured time that the two algorithms take to generate the branching structure of a tree model. As shown, our method outperforms the procedural

Fig. 15. Our method can also replicate the growth process of form-controlled tree models. For this result, we trained on a dataset of tree models that were grown into meshes of cones (a), cubes (b), and rings (c).

Fig. 16. Three failure cases: if the used neural networks are not properly trained our method tends to generate branching structures with too thin and too long branches. Here we show examples of Oak (a), Tulip (b), and Maple (c) trees.

algorithm for tree models of the same species and similar complexity. While the procedural model is implemented with a recursive algorithm to construct the branch graph, our method queries the neural network pipeline of classification and a cascade of regression networks to generate the branch graph iteratively.

Network Comparison. To validate our cascaded neural network architecture (Fig. 3c), we compare it with a simpler architecture that jointly trains the classification and regression networks. The architecture for this experiment is the same as shown in Fig.3a, except that we add two additional heads on top of the joined embedding to generate thickness and length as well as the quaternion. We trained both architectures on all species with the same setup for loss functions and data processing. With the cascaded network architecture, we are able to obtain a classification result of 83 - 93%, an MSE of 1.2e-4 - 3e-4, and a cosine similarity of 0.04 - 0.20, whereas, for the joint model, we obtain a classification result of 89% - 90%, an MSE of 0.28 - 0.51, and a cosine similarity of 1.03 - 0.51. While the classification head is trained with similar accuracy, the regression results are inferior. The jointly trained network immediately predicted erroneous branching that caused the termination of the tree generation.

trees (mean and std for the entire dataset).

Table 2. ICTree perceptual validation of the ground truth and generated

	All trees	[0.49, 0.11]	[0.47, 0.11]
	Willow	[0.49, 0.11]	[0.46, 0.10]
	Walnut	[0.49, 0.11]	[0.47, 0.10]
	Tulip	[0.49, 0.11]	[0.45, 0.10]
	Pine	[0.50, 0.11]	[0.48, 0.11]
	Oak	[0.50, 0.11]	[0.48, 0.12]
	Maple	[0.49, 0.11]	[0.46, 0.12]
(Corkscrew	[0.48, 0.11]	[0.48, 0.12]
	Beech	[0.47, 0.10]	[0.49, 0.09]
	Species	Ground Truth	DeepTree

Table 3. Runtime performance comparison of procedurally generated trees (P) obtained with the method of [45] and *DeepTree* generated trees (DT). The runtime is reported in seconds and only includes the time required to generate the branch graph.

Species	Method	Runtime	# Nodes	# Branches
Decel	Р	23	3376	777
beech	DT	2.7	5672	1634
Carlinaria	Р	26	2872	751
Corkscrew	DT	2.1	3041	566
Manla	Р	25	2928	861
Maple	DT	1.9	2951	817
Oala	Р	27	4261	1084
Oak	DT	2.1	6238	1738
Dina	Р	26	3506	1000
rine	DT	1.9	3168	839
Tulin	Р	26	3050	726
Tulip	DT	1.8	3360	773
W-1+	Р	27	5638	589
wamut	DT	2.0	4975	677
W/:11	Р	30	4796	1304
w mow	DT	2.0	4975	1691

7 DISCUSSION AND LIMITATIONS

Our focus was on exploring the capabilities of neural networks to predict branching structures by only training them on local branching patterns. Training a neural network this way learns a situated latent space - a representation that can encode the necessary topological and geometric information to mimic branch growth. As we have shown, a trained situated latent space can serve as a powerful representation capable of generating a wide range of complex branching patterns and encoding the environmental response of trees to obstacles in their environment.

Using situated latent spaces, our approach is orthogonal to approaches that aim to learn graph structures end-to-end. In contrast to these approaches, our method is driven by the idea of only encoding local distributions, e.g., local variations of branching patterns. This is inspired by Nature's ability to encode structural and behavioral properties as DNA, where the DNA serves as a blueprint for generating complex structures. *DeepTree* is a step in this direction that intriguingly shows that locally learned representations can provide an interesting modeling alternative.

Our work is also similar to graph neural networks (GNNs) in that we also focus on predicting the topological and geometric properties of graphs. Research toward GNNs has recently gained a lot

Fig. 17. Geometric and perceptual comparison of the ground truth (GT) (left bar) and the trees generated by DeepTree (DT) (the right bar).

of momentum, and many approaches aim to predict graphs and their properties in different ways [44, 54, 56, 60]. GNNs leverage the connectivity of vertices to their neighbors in a graph by simultaneously predicting vertex properties and their connectivity information. Most of these methods focus on making predictions for local properties in graphs to solve classification tasks. Our work aims to simultaneously estimate the topology and the geometric properties of graph nodes. Regressing multiple geometric attributes remains a challenging problem that we addressed by training a cascade of individually trained regression networks.

Our method is currently limited in three ways: first, we rely on using data generated by a procedural model. We decided to showcase our method on procedurally generated tree models as this is the only way to generate a large dataset to train our neural network pipeline successfully. To the best of our knowledge, large collections of reconstructed models from scans or images do not exist. However, as long as the tree models are generated from branching graphs, which is also common practice for reconstruction algorithms and manual artistic modeling, our method can still work with the same degree of sophistication. Moreover, DeepTree is a data-driven method, and the training data determines the set of generated trees. While we can generate variations via environmental changes, modifying the parameters of the intrinsic model is not possible. The second limitation is that our method is currently unable to model tree-tree or branch-branch interactions. Procedural models can distinguish between the intrinsic plant environment (the plant itself) and the extrinsic (obstacles). However, as we rely on a coarse voxel

space to encode the environment, we cannot solve intricate treetree or branch-branch collisions during tree generation. Finally, we observed that our method has difficulties correctly predicting the termination of branches for some node configurations, which generates artificial-looking branching structures (Fig. 16). While this leads to unrealistic branching structures for some tree models, we observed that this is highly correlated with the size of the dataset and the training duration.

8 CONCLUSIONS AND FUTURE WORK

We have advanced tree modeling in computer graphics by introducing *DeepTree*, a deep-learning-based method for automatic and adaptive tree form generation. *DeepTree* is inspired by Nature's ability to encode the development of structural and behavioral features as DNA and decides the shape generation on a local level based on the signature of a single node and the environment. Instead of manually defining parameters and rules – as is commonly done in procedural modeling – our method learns branching patterns locally as they can be observed for a single node in a branch graph. We have shown that a situated latent space that is evaluated for a single node in a tree graph, somewhere in the growth space of a tree model, can generate complex tree models with similar topological and geometric features as contemporary procedural and developmental modeling techniques.

We have shown that our method can generate a diverse set of branching structures from different species. Adding additional information to each node's signature (e.g., a species identifier) shows that our method can even be trained in a generalized manner, i.e., we

Fig. 18. Random sampling: a set of tree models of the same species generated by our neural network pipeline shown from the front (top row) and top (bottom).

can jointly train our networks on multiple species and successfully generate tree models. Furthermore, by encoding the environment, we have shown that our situated latent space can even support mimicking the environmental response of tree models when they grow next to walls or other trees.

A novel way of encoding tree form opens multiple avenues for future work. First, it would be interesting to explore how a situated latent space can be used to support the authoring of branching structures. In our current framework, we only automatically use our network pipeline. However, we could also control our network pipeline by user-defined sketches or adapt it to dynamically changing scenes. Second, our current architecture does not avoid inter-branch collisions. To this end, it would be interesting to explore additional global or semi-global encodings for the generated branches when querying the networks for producing new nodes. For example, encoding the neighboring nodes of a node could provide an additional feature for producing branching structures that do not collide. It also seems promising to use our proposed algorithm of iteratively querying a neural network pipeline to generate other objects, e.g., road networks or buildings.

ACKNOWLEDGMENTS

We want to thank Liangliang Nan for providing a point cloud for the tree model from Fig. 12.

REFERENCES

- [1] G. Acquaah. 2009. Principles of plant genetics and breeding. John Wiley & Sons.
- [2] M. Aono and T.L. Kunii. 1984. Botanical Tree Image Generation. IEEE Comput.
- Graph. Appl. 4(5) (1984), 10–34.
 [3] O. Argudo, C. Andújar, and A. Chica. 2020. Image-Based Tree Variations. Computer Graphics Forum 39, 1 (2020), 174–184.
- [4] B. Benes and E. U. Millán. 2002. Virtual Climbing Plants Competing for Space. In Proceedings of the Computer Animation (CA '02). IEEE Computer Society, USA, 33.
- [5] D. Bradley, D. Nowrouzezahrai, and P. Beardsley. 2013. Image-based Reconstruction and Synthesis of Dense Foliage. ACM Transactions on Graphics 32, 4, Article 74 (2013), 74:1–74:10 pages.
- [6] X. Chen, B. Neubert, Y.-Q. Xu, O. Deussen, and S. B. Kang. 2008. Sketch-Based Tree Modeling Using Markov Random Field. ACM Transactions on Graphics 27, 5, Article 109 (Dec. 2008), 9 pages.
- [7] O. Deussen, A. Dowden-Williams, and B. Lintermann. 2005. Digital Design of Nature: Computer Generated Plants and Organics. Springer Berlin Heidelberg.
- [8] S. Du, R. Lindenbergh, H. Ledoux, J. Stoter, and L. Nan. 2019. AdTree: Accurate, detailed, and automatic modelling of laser-scanned trees. *Remote Sensing* 11, 18 (2019), 2074.
- [9] R. Estrada, C. Tomasi, S. C. Schmidler, and S. Farsiu. 2015. Tree topology estimation. IEEE Trans. Pattern Anal. Mach. Intell. 37, 8 (Aug. 2015), 1688–1701.

- [10] N. Greene. 1989. Voxel Space Automata: Modeling with Stochastic Growth Processes in Voxel Space. SIGGRAPH Comp. Graph. 23, 3 (1989), 175–184.
- [11] J. Guo, H. Jiang, B. Benes, O. Deussen, X. Zhang, D. Lischinski, and H. Huang. 2020. Inverse Procedural Modeling of Branching Structures by Inferring L-Systems. ACM Transactions on Graphics 39, 5 (sep 2020), 1–13.
- [12] R. Habel, A. Kusternig, and M. Wimmer. 2009. Physically Guided Animation of Trees. 28. 2 (2009), 523–532.
- [13] J. T. Hack. 1957. Studies of longitudinal stream profiles in Virginia and Maryland. Vol. 294. US Government Printing Office.
- [14] T. Hädrich, D. T. Banuti, W. Pałubicki, S. Pirk, and D. L. Michels. 2021. Fire in Paradise: Mesoscale Simulation of Wildfires. ACM Trans. Graph. 40, 4, Article 163 (July 2021).
- [15] T. Hädrich, B. Benes, O. Deussen, and S. Pirk. 2017. Interactive Modeling and Authoring of Climbing Plants. *Comput. Graph. Forum* 36, 2 (2017), 49–61.
- [16] T. Ijiri, S. Owada, and T. Igarashi. 2006. Seamless Integration of Initial Sketching and Subsequent Detail Editing in Flower Modeling. 25, 3 (2006), 617–624.
- [17] P. Khatri and S. Drăghici. 2005. Ontological analysis of gene expression data: current tools, limitations, and open problems. *Bioinformatics* 21, 18 (2005), 3587– 3595.
- [18] Š. Kohek and D. Strnad. 2015. Interactive Synthesis of Self-Organizing Tree Models on the GPU. Computing 97, 2 (feb 2015), 145–169.
- [19] B. Li, J. Kałużny, J. Klein, D. L. Michels, W. Pałubicki, B. Benes, and S. Pirk. 2021. Learning to Reconstruct Botanical Trees from Single Images. ACM Transaction on Graphics 40, 6, Article 231 (2021), 15 pages.
- [20] C. Li, O. Deussen, Y.-Z. Song, P. Willis, and P. Hall. 2011. Modeling and Generating Moving Trees from Video. ACM Transactions on Graphics 30, 6, Article 127 (2011), 127:1–127:12 pages.
- [21] B. Lintermann and O. Deussen. 1999. Interactive Modeling of Plants. IEEE Comput. Graph. Appl. 19, 1 (1999), 56–65.
- [22] Y. Liu, J. Guo, B. Benes, O. Deussen, X. Zhang, and H. Huang. 2021. TreePart-Net: Neural Decomposition of Point Clouds for 3D Tree Reconstruction. ACM Transaction on Graphics 40, 6, Article 232 (2021), 16 pages.
- [23] Y. Livny, S. Pirk, Z. Cheng, F. Yan, O. Deussen, D. Cohen-Or, and B. Chen. 2011. Texture-lobes for Tree Modelling. ACM Transactions on Graphics 30, 4, Article 53 (2011), 53:1–53:10 pages.
- [24] S. Longay, A. Runions, F. Boudon, and P. Prusinkiewicz. 2012. TreeSketch: interactive procedural modeling of trees on a tablet. In Proc. of the Intl. Symp. on SBIM. 107–120.
- [25] M. Makowski, T. Hädrich, J. Scheffczyk, D. L. Michels, S. Pirk, and W. Pałubicki. 2019. Synthetic Silviculture: Multi-Scale Modeling of Plant Ecosystems. ACM Trans. Graph. 38, 4, Article 131 (July 2019), 14 pages.
- [26] R. Měch and P. Prusinkiewicz. 1996. Visual models of plants interacting with their environment. In Proc. of SIGGRAPH. ACM, 397-410.
- [27] B. Neubert, T. Franken, and O. Deussen. 2007. Approximate Image-based Treemodeling Using Particle Flows. ACM Transactions on Graphics 26, 3, Article 88 (2007).
- [28] B. Neubert, S. Pirk, O. Deussen, and C. Dachsbacher. 2011. Improved Model- and View-Dependent Pruning of Large Botanical Scenes. *Computer Graphics Forum* 30, 6 (2011), 1708–1718.
- [29] T. Niese, S. Pirk, M. Albrecht, B. Benes, and O. Deussen. 2022. Procedural Urban Forestry. ACM Trans. Graph. 41, 2, Article 20 (mar 2022), 18 pages.
- [30] M. Okabe, S. Owada, and T. Igarashi. 2007. Interactive Design of Botanical Trees Using Freehand Sketches and Example-Based Editing. In ACM SIGGRAPH 2007 Courses (San Diego, California) (SIGGRAPH '07). 26-es.
- [31] P. E. Oppenheimer. 1986. Real time design and animation of fractal plants and trees. Proc. of SIGGRAPH 20, 4 (1986), 55-64.

- [32] W. Palubicki, K. Horel, S. Longay, A. Runions, B. Lane, R. Měch, and P. Prusinkiewicz. 2009. Self-organizing Tree Models for Image Synthesis. ACM Trans. Graph. 28, 3, Article 58 (2009), 10 pages.
- [33] W. Pałubicki, M. Makowski, W. Gajda, T. Hädrich, D. L. Michels, and S. Pirk. 2022. Ecoclimates: Climate-Response Modeling of Vegetation. ACM Trans. Graph. 41, 4, Article 155 (2022), 19 pages.
- [34] S. G. Parker, J. Bigler, A. Dietrich, H. Friedrich, J. Hoberock, D. Luebke, D. McAllister, M. McGuire, K. Morley, A. Robison, and M. Stich. 2010. OptiX: A General Purpose Ray Tracing Engine. ACM Trans. Graph. 29, 4, Article 66 (jul 2010), 13 pages.
- [35] S. Pirk, B. Benes, T. Ijiri, Y. Li, O. Deussen, B. Chen, and R. Měch. 2016. Modeling Plant Life in Computer Graphics. In ACM SIGGRAPH 2016 Courses (Anaheim, California) (SIGGRAPH '16). ACM, New York, NY, USA, Article 18, 180 pages.
- [36] S. Pirk, M. Jarząbek, T. Hädrich, D. L. Michels, and W. Palubicki. 2017. Interactive Wood Combustion for Botanical Tree Models. ACM Transactions on Graphics 36, 6, Article 197 (Nov. 2017), 12 pages.
- [37] S. Pirk, T. Niese, O. Deussen, and B. Neubert. 2012. Capturing and animating the morphogenesis of polygonal tree models. ACM Transactions on Graphics 31, 6, Article 169 (2012), 10 pages.
- [38] S. Pirk, T. Niese, T. Hädrich, B. Benes, and O. Deussen. 2014. Windy Trees: Computing Stress Response for Developmental Tree Models. ACM Transactions on Graphics 33, 6, Article 204 (2014), 11 pages.
- [39] S. Pirk, O. Stava, J. Kratt, M. A. M. Said, B. Neubert, R. Mčch, B. Benes, and O. Deussen. 2012. Plastic Trees: Interactive Self-adapting Botanical Tree Models. ACM Trans. Graph. 31, 4, Article 50 (July 2012), 10 pages.
- [40] T. Polasek, D. Hrusa, B. Benes, and M. Cadik. 2021. ICTree: Automatic Perceptual Metrics for Tree Models. ACM Transaction on Graphics 40, 6, Article 230 (Dec. 2021), 15 pages. https://doi.org/10.1145/3478513.3480519
- [41] P. Prusinkiewicz and A. Lindenmayer. 1990. The Algorithmic Beauty of Plants. Springer-Verlag New York, Inc.
- [42] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. 2017. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In CVPR. 77–85.
- [43] W. T. Reeves and R. Blau. 1985. Approximate and Probabilistic Algorithms for Shading and Rendering Structured Particle Systems. SIGGRAPH Comput. Graph. 19, 3 (July 1985), 313–322.
- [44] H. Shao, T. Kugelstadt, T. Hädrich, W. Pałubicki, J. Bender, S. Pirk, and D. L. Michels. 2021. Accurately Solving Rod Dynamics with Graph Learning. In *NeurIPS*.
- [45] O. Stava, S. Pirk, J. Kratt, B. Chen, R. Měch, O. Deussen, and B. Benes. 2014. Inverse Procedural Modelling of Trees. 33, 6 (2014), 118–131.
- [46] J. O. Talton, Y. Lou, S. Lesser, J. Duke, R. Měch, and V. Koltun. 2011. Metropolis procedural modeling. ACM Transactions on Graphics 30, Article 11 (2011), 14 pages.
- [47] P. Tan, T. Fang, J. Xiao, P. Zhao, and L. Quan. 2008. Single Image Tree Modeling. ACM Transactions on Graphics 27, 5, Article 108 (2008), 7 pages.
- [48] P. Tan, G. Zeng, J. Wang, S. B. Kang, and L. Quan. 2007. Image-based Tree Modeling. ACM Transactions on Graphics 26, 3, Article 87 (2007).
- [49] D. W Thompson. 1942. On growth and form. Vol. 2. Cambridge university press.
 [50] B. Wang, Y. Zhao, and J. Barbič. 2017. Botanical Materials Based on Biomechanics.
- ACM Transactions on Graphics 36, 4, Article 135 (July 2017), 13 pages. [51] J. Weber and J. Penn. 1995. Creation and rendering of realistic trees. In Proc. of
- [31] J. Weber and J. Temi. 1993. Creation and related in the internet internet in the cost of SIGGRAPH '95. New York, NY, USA, 119–128.
 [52] J. Wither, F. Boudon, M.-P. Cani, and C. Godin. 2009. Structure from silhouettes:
- a new paradigm for fast sketch-based design of trees. 28, 2 (2009), 541–550.
- [53] S.-K. Wong and K.-C. Chen. 2015. A Procedural Approach to Modelling Virtual Climbing Plants With Tendrils. Comput. Graph. Forum (2015).
- [54] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. 2021. A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems* 32, 1 (2021), 4–24. https://doi.org/10.1109/TNNLS.2020.2978386
- [55] H. Xu, N. Gossett, and B. Chen. 2007. Knowledge and heuristic-based modeling of laser-scanned trees. ACM Transactions on Graphics 26, 4 (2007), Article 19, 13 pages.
- [56] J. Xu, J. Chen, S. You, Z. Xiao, Y. Yang, and J. Lu. 2021. Robustness of deep learning models on graphs: A survey. AI Open 2 (2021), 69–78.
- [57] L. Xu and D. Mould. 2015. Procedural Tree Modeling with Guiding Vectors. Computer Graphics Forum 34, 7 (2015), 47–56.
- [58] X. Zhang, G. Bao, W. Meng, M. Jaeger, H. Li, O. Deussen, and B. Chen. 2017. Tree Branch Level of Detail Models for Forest Navigation. *Computer Graphics Forum* 36, 8 (2017), 402–417.
- [59] Y. Zhao and J. Barbič. 2013. Interactive Authoring of Simulation-ready Plants. ACM Transactions on Graphics 32, 4, Article 84 (2013), 12 pages.
- [60] J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun. 2020. Graph neural networks: A review of methods and applications. AI Open 1 (2020), 57–81.

A APPENDIX

Table 4. Table of parameters for the procedural model.

Param	Name	Description
G _{NLB}	Number of lateral buds	The number of lateral buds created per internode
		during growth.
G_{AAM}	Apical angle Mean	The mean of the angle between the direction of
		parent shoot and the direction of apical bud.
G_{AAV}	Apical angle variance	The variance of the angle between the direction of
		parent shoot and the direction of apical bud.
G_{BAM}	Branching angle mean	The mean of angle between the direction of parent
		shoot and the direction of lateral bud.
G_{BAV}	Branching angle variance	The variance of angle between the direction of par-
		ent shoot and the direction of lateral bud.
G_{RAM}	Roll angle mean	The mean of orientation angle between two lateral
		buds created with the same internode.
G_{RAV}	Roll angle variance	The variance of orientation angle between two lat-
		eral buds created with the same internode.
F_{AKP}	Apical bud kill probability	The probability that a given apical bud will die dur-
		ing a growth cycle.
F_{LKP}	Lateral bud kill probability	The probability that a given lateral bud will die
		during a growth cycle.
F_{ALF}	Apical bud lighting factor	The influence of the lighting condition on the
		growth probility of a apical bud.
F_{LLF}	Lateral bud lighting factor	The influence of the lighting condition on the
		growth probility of a lateral bud.
F_{ADB}	Apical dominance base	The base level of auxin produced to inhibit parent
		shoots from growing.
F_{ADF}	Apical dominance distance	The reduction of auxin due to the transimission
_	factor	along parent shoots.
F_{AAF}	Apical dominance age fac-	The reduction of auxin due to increasing age of the
	tor	tree.
FGR	Growth rate	The expected number of internodes generated along
		the branch during a growth cycle.
FILB	Internode length base	The base distance between two adjacent internodes
F		on the same shoot.
FACB	Apical control base	The impact of the branch level on the growth rate.
<i>EPHO</i>	Phototropism	The impact of the average growth direction of in-
F	0	coming light.
LGGB	Gravitropism base	The impact of the average growth direction of the
F	D	gravity.
LPPF	Pruning factor	The impact of the amount of incoming light on the
F	I and have also amontian of a star	The being of branches.
LLPF	Low branch pruning factor	The neight below which all lateral branches are
Fana	Cravity handing strop -th-	pruneu. The impect of gravity on branch structured bar direction
EGBS	Gravity bending strength	The impact of gravity on branch structural bending.
LGAF	tor	ne relation of gravity bending related to the thick-
+	Desired age	The avpected are of the tree
	Desireu age	The expected age of the tree.

Table 5. Parameter values for each species used in our framework.

Params.	Beech	Corkscrew	Maple	Oak	Pine	Tulip	Walnut	Willow
G_{NLB}	2	2	2	2	2	2	2	2
G_{AAM}	0	45	0	20	0	0	20	25
G_{AAV}	2	3	0	2	0	1	2	3
G_{BAM}	45	45	50	30	80	30	30	15
G_{BAV}	1	2	1	3	1	1	3	1
G_{RAM}	120	120	110	130	120	125	120	120
G_{RAV}	3	3	2	3	3	3	1	2
F_{AKP}	0.0	0.0	0.0	0.01	0.0	0.05	0.1	0.0
F_{LKP}	0.01	0.03	0.0	0.03	0.03	0.21	0.15	0.12
F_{ALF}	0.1	0.23	0.12	0.15	0.05	0.09	0.04	0.3
F_{LLF}	0.54	1.13	0.94	0.2	0.2	1.0	0.1	0
F_{ADB}	0.2	0.12	0.2	0.12	0.5	0.5	0.12	0.11
F_{ADF}	1.0	1.0	1.0	0.98	1.0	1.0	0.1	1.0
F_{AAF}	0.5	0.3	0.5	0.3	0.3	0.3	0.3	0.3
F_{GR}	1.5	1.45	1.45	1.35	1.2	1.1	1.75	1.12
F_{ILB}	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
F_{ACB}	1.0	1.0	1.5	1.0	1.6	2.0	6.2	1.0
E_{PHO}	0.01	0.1	0.0	0.05	0.0	0.1	0.05	0.1
E_{GGB}	-0.25	-0.21	-0.1	-0.1	-0.09	-0.15	-0.1	0.07
E_{PPF}	0.1	0.2	0.05	0.01	0.01	0.7	0.12	0
E_{LPF}	0.1	0.15	0.1	0.15	0.15	0.15	0.15	0.2
E_{GBS}	1.0	6.0	2.0	6.0	4.0	2.0	6.0	0.0
E_{GAF}	2.0	3.0	2.0	3.0	2.0	3.0	3.0	3.0
t	35	35	40	35	45	60	40	40

Table 6. Evaluation metrics values for each species. P indicates procedurally generated models and DT *DeepTree* generated ones. The top row in each pair of rows shows the mean value and the second row the standard deviation.

Params.	Beech		Beech Corkscrew		Maple Oak		Pine 7		Tu	lip	Wa	Walnut		Willow		Total			
	Р	DT	Р	DT	Р	DT	Р	DT	Р	DT	Р	DT	Р	DT	Р	DT	Р	DT	
Branch length	3.30	3.75	4.47	5.92	3.34	3.69	4.16	4.61	2.92	3.41	4.44	4.80	4.84	5.13	4.43	5.19	4.50	4.04	11%
(BL)	3.16	3.31	3.48	4.29	2.70	2.75	3.44	3.97	2.67	2.88	3.56	3.40	3.81	4.67	3.46	3.97	3.77	3.36	
Branch number	777	1634	751	566	816	817	1084	1738	1000	839	726	773	589	677	1304	1691	894	1092	18%
(BN)	222	308	224	95	122	42	354	194	129	65	49	33	175	259	440	495	337	527	
Branching distance	1.98	2.10	2.43	3.16	1.91	2.11	2.28	2.50	1.65	1.96	2.56	2.79	2.60	2.79	2.47	2.80	2.47	2.80	12%
(BD)	1.51	1.55	1.79	2.57	1.27	1.38	1.76	1.92	0.98	1.30	2.06	2.08	1.82	2.26	2.03	2.14	2.03	2.14	
Node number	1.19	1.06	1.18	1.09	1.20	1.29	1.18	1.06	1.16	1.03	1.10	1.05	1.11	1.07	1.15	1.01	1.15	1.08	11%
(NN)	0.62	0.68	0.63	0.67	0.75	0.68	0.60	0.66	0.69	0.65	0.51	0.51	0.45	0.37	0.58	0.52	0.59	0.59	
Angle: parent vs child1	0.15	0.15	0.70	0.69	0.07	0.08	0.33	0.33	0.06	0.06	0.06	0.06	0.34	0.32	0.42	0.41	0.31	0.29	6%
	0.10	0.11	0.17	0.09	0.03	0.04	0.08	0.05	0.04	0.04	0.04	0.04	0.10	0.06	0.15	0.06	0.24	0.19	
Angle: parent vs child2	0.79	0.79	0.81	0.81	0.89	0.89	0.54	0.53	1.40	1.40	0.53	0.53	0.56	0.55	0.26	0.27	0.70	0.67	4%
	0.02	0.06	0.18	0.07	0.03	0.04	0.17	0.04	0.05	0.05	0.02	0.06	0.20	0.10	0.02	0.01	0.38	0.31	
Angle: parent vs child3	0.79	0.80	0.81	0.80	0.88	0.88	0.54	0.56	1.40	1.40	0.53	0.53	0.56	0.57	0.26	0.27	0.70	0.66	6%
	0.02	0.07	0.18	0.10	0.03	0.05	0.17	0.09	0.05	0.07	0.02	0.04	0.20	0.18	0.02	0.06	0.38	0.32	