

# Probabilistic Plant Modeling via Multi-View Image-to-Image Translation

Takahiro Isokane<sup>1,\*</sup> Fumio Okura<sup>1,2,\*</sup> Ayaka Ide<sup>1</sup> Yasuyuki Matsushita<sup>1</sup> Yasushi Yagi<sup>1</sup>  
<sup>1</sup>Osaka University <sup>2</sup>JST PRESTO

{isokane, okura, ide, yagi}@am.sanken.osaka-u.ac.jp yasumat@ist.osaka-u.ac.jp



Figure 1: Result of branch structure estimation from simulated and real leafy plant images. From multi-view plant images, our method infers the branch structure in a probabilistic framework. An explicit graph structure (red lines) can be derived from the probabilistic plant model.

## Abstract

*This paper describes a method for inferring three-dimensional (3D) plant branch structures that are hidden under leaves from multi-view observations. Unlike previous geometric approaches that heavily rely on the visibility of the branches or use parametric branching models, our method makes statistical inferences of branch structures in a probabilistic framework. By inferring the probability of branch existence using a Bayesian extension of image-to-image translation applied to each of multi-view images, our method generates a probabilistic plant 3D model, which represents the 3D branching pattern that cannot be directly observed. Experiments demonstrate the usefulness of the proposed approach in generating convincing branch structures in comparison to prior approaches.*

## 1. Introduction

We propose an approach to estimate branch (skeleton) structures of plants from multi-view images that are severely occluded by leaves. Unveiling hidden skeleton structures is one of the most challenging tasks in computer vision, because it naturally involves inference of the unobserved structures. Skeleton estimation is actively studied for human pose estimation [6]. A recent trend in image-based human skeleton estimation relies on the prior knowledge of

the relationship among joints for extracting and interpreting joints from images. Unlike human bodies, the branch structures of plants are less organized, namely, the number of joints and their connections are generally unknown. In addition, heavy occlusions due to leaves makes the problem harder. These aspects pose a unique challenge in plant’s structure estimation. Although difficult, branch structure estimation has a number of applications, such as synthetic plant generation in computer graphics and computational agriculture.

Toward this goal, this paper presents a multi-view image-to-image translation approach to 3D branch structure estimation. To estimate the branching paths hidden under leaves, we cast the estimation problem to an *image-to-image translation* problem [12], which converts an image from one domain to another domain. In our context, we convert an input image of a leafy plant to a map that represents the branch structure, *i.e.*, each pixel containing the prediction of “branch” or “not-branch”. To deal with the uncertainty of the prediction, we develop a Bayesian extension of image-to-image translation applied to each of the multi-view images. It yields the prediction together with its reliability, resulting in *probabilistic* estimates in a similar manner to a Bayesian semantic segmentation approach [13].

The probabilistic representation is advantageous in a few important aspects; not only that we have access to the credibility of the estimates, but also that it allows us to consolidate the view-dependent inferences in a 3D space in a well-defined probabilistic framework. Thus, in our method,

\* Authors contributed equally.

instead of directly predicting 3D branch structures, it makes inference in individual views of the multi-view input and subsequently aggregate them in a 3D voxel space to obtain a probabilistic 3D plant model. We also develop a method for generating explicit branch structures from the probabilistic model based on particle flows so that the probabilistic representation can be converted to a form that could be used in intended applications. A few examples of the output are shown in Fig. 1.

The primary contributions of this study are twofold. First, we propose a method of recovering 3D structures of a plant based on Bayesian image-to-image (leafy- to branch-image) translation applied in a multi-view manner. By representing the *branchness* in a probabilistic framework, we show that both faithful estimates and comprehensive aggregation over multi-view predictions can be achieved. Second, we show that the branch structure estimation is made possible by the proposed approach, together with a method for extracting explicit branch structures from the probabilistic representation. Experimental results show that the proposed method can generate convincing branch structures even with severe occlusions by leaves and other branches. In comparison to a traditional tree reconstruction approach, we found notable improvement in the resultant 3D branching structure.

## 2. Related Works

Our goal is to reconstruct a 3D branch structure of a plant from multi-view images that exhibit severe occlusions due to plant leaves. At the heart of the proposed method, we develop a Bayesian extension of image-to-image translation for making statistical predictions of branch existence even for the parts that are not observed at all. In what follows, we discuss the prior arts for 3D reconstruction of plants and trees, image-to-image translation and Bayesian neural networks that are related to our work.

**3D reconstruction of plants and trees.** (Semi-) automatic 3D modeling of plants and trees is actively studied in the graphics community [32] because of their importance as a rendering subject and that their manual modeling is notably time-consuming. Interactive 3D modeling methods of trees using manually provided hints, *e.g.*, lines, are proposed in the early 00's [4, 21]. Growth models of trees, which include branching rules, are occasionally utilized [10, 29, 22], some of them are augmented by realistic textures [16]. These approaches generally heavily rely on the branch models (or rules), and the resulting structure cannot deviate much from the presumed models.

On the other hands, an approach based on observations of real-world trees that uses photographs [27] or 3D scans [34, 17] is shown promising for automatic tree modeling. Several tree modeling approaches using multi-view images have been proposed [25, 31, 20]. These approaches

have been further extended to single-image based methods [30, 2] for better applicability. The major focus of the most image-based modeling approaches has been to generate 3D tree models that well fit the silhouette or volumes, not necessarily aiming at recovering branch structures. Using multi-view images of bare trees, previous approaches achieve geometric reconstruction of a branch structure [18, 28, 35]; however, leaves of plants or trees make 3D reconstruction considerably difficult due to occlusions. The scope of plant modeling now goes beyond computer graphics and is becoming an important application of computer vision, *i.e.*, analysis of plant shapes and growth for vision-assisted cultivation and plant phenotyping. Our method aims to recover branch structures via a multi-view approach with an emphasis on dealing with uncertainty due to severe occlusions by leaves.

Apart from the plant and tree context but somehow related, recently human skeleton (structure) estimation shows a great success [23, 6]. The task has a similarity to what we study in this paper; however, the plant and tree branch structure estimation has its unique difficulty in that the number of joints and their relationships are generally unknown and complex.

**Image-to-image translation.** Image-to-image translation aims at transferring contextual or physical variation between the source and target images. The early works on image-to-image translation include image analogies [11] and texture transfer [8]. Commonly, these approaches divide the image into small patches and transfer the change based on patch-wise correspondences [14, 3, 7]. More recent image-to-image translation largely benefits from deep learning, such as convolutional neural networks (CNNs) with encoder-decoder architectures or using Generative Adversarial Networks (GANs). Pix2Pix [12], which uses conditional GAN, shows impressive performance on a wide variety of translation tasks. Along this context, Cycle GAN [37] has shown the possibility of image-to-image translation without paired training images.

**Bayesian neural networks** To obtain *reliability* of inference, a Bayesian framework has been used together with neural networks in Bayesian neural networks (BNNs) [19]. Recent BNNs use dropout connections in forward passes in the prediction stage [9], which yields variations in inference. The estimate is thus given in the form of distribution rather than a point, effectively modeling the uncertainty of the prediction. This approach can be regarded as an approximation of Monte Carlo integration in traditional Bayesian approaches [15]. An impressive result has been shown in Bayesian SegNet [13] that achieves semantic segmentation augmented by the prediction reliability. Our method also uses the Bayesian approach in image-to-image translation task in order for obtaining the probabilistic representation of the branch structures.

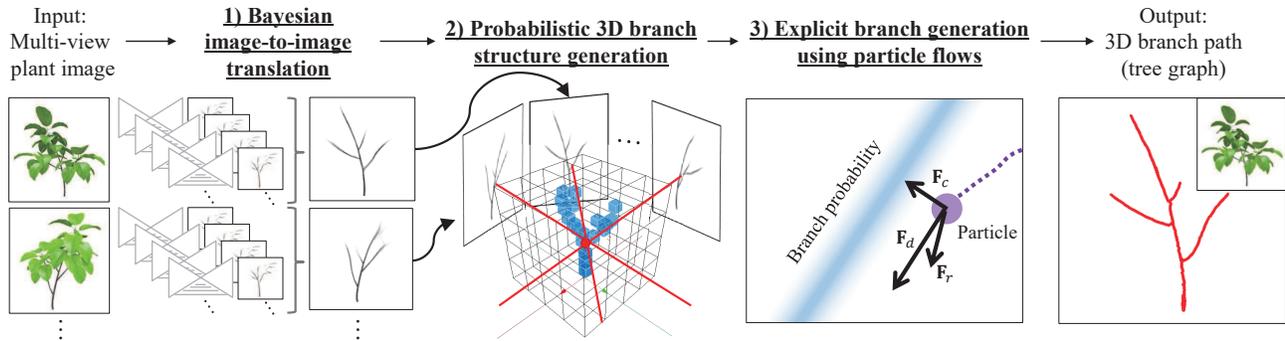


Figure 2: **Overview of our plant’s branch structure modeling.** The proposed method first estimates the 2D branch existence probability on multi-view images using a Bayesian extension of image-to-image translation. The 2D probabilities are then consolidated in a 3D voxel space to form a volume of 3D branch probability. An explicit 3D branch structure is generated by particle flow simulation based on the 3D probability map.

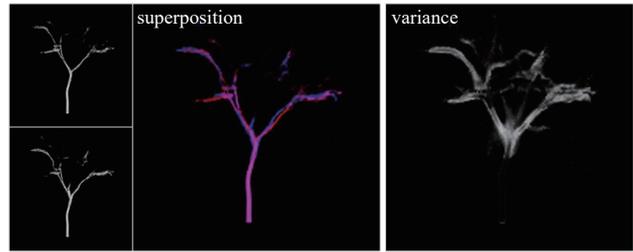
### 3. Probabilistic Branch Structure Modeling

Our method takes as input multi-view images of a plant and generates a probabilistic 3D branch structure in a 3D voxel space. Our method begins with estimating a 2D probabilistic branch existence map in each of the multi-view images based on an altered image-to-image translation method. Once the probabilistic branch existence map is computed for each view, they are merged in a 3D voxel space using the estimated camera poses based on a structure-from-motion method [33] to yield a probabilistic 3D branch structure. Finally, an explicit 3D branch structure is generated by a particle flow simulation, which is inspired by a traditional tree modeling approach [20], using the probabilistic 3D branch structure. Figure 2 illustrates the whole pipeline of the proposed method. In what follows, we explain the individual steps of the proposed method.

#### 3.1. Bayesian image-to-image translation

From a leafy plant image, we first estimate a pixel-wise 2D branch existence probability. The major challenge is to infer branch structure hidden under leaves, which cannot be directly observed by a camera, possibly not from any of the viewpoints. It is here that we adopt a Pix2Pix approach [12] to image-to-image translation as a mean to derive a statistically valid prediction of the existence of branches in the multi-view images. For our context, we train a Pix2Pix network using pairs of a leafy plant and its corresponding label map describing the branch region.

To obtain a meaningful *probabilistic* branch existence as output, we use the original image-to-image translation approach in a Bayesian deep learning framework [9, 19] by Monte Carlo sampling via dropout at the inference stage. By adopting this fluctuation and repeating the prediction multiple times, we obtain a probability of the branch existence for each pixel. This strategy is implemented by inserting dropout layers before the middle four convolution lay-



(a) Two samples generated by Bayesian image-to-image translation (b) Variance in 100 samples

Figure 3: Effect of Bayesian image-to-image translation in branch generation. Two generated samples contain differences in small branches (a). The variance in a larger number of samples (b) shows the uncertainty in the prediction, which cannot be obtained from a single inference.

ers in the encoder-decoder network of Pix2Pix. To further increase variations of inference, dropout is additionally applied to skip connections [26] of Pix2Pix. Figure 3 shows an example of output variations by this Bayesian branch generation. Multiple inferences yield the degree of uncertainty, which cannot be obtained by a single inference.

As such, we obtain the mean of variational inferences from the image-to-image translation network. By treating the variational inferences as stochastic samples, each pixel in the mean inference can be regarded as the probability of branch existence spanned in  $[0, 1]$ . For the  $i$ -th image  $I_i$ , the branch probability  $B_{2D_i} : \mathbb{R}^2 \rightarrow [0, 1]$  at a pixel  $\mathbf{x}_{2D} \in \mathbb{R}^2$  is written as:

$$\begin{cases} B_{2D_i}(\mathbf{x}_{2D}) = \frac{1}{n_v} \sum_v \hat{B}_{2D_i,v}(\mathbf{x}_{2D}), \\ \hat{B}_{2D_i,v} = \pi_v(I_i), \end{cases} \quad (1)$$

where  $\pi_v$  denotes the Pix2Pix translation from an image  $I_i$

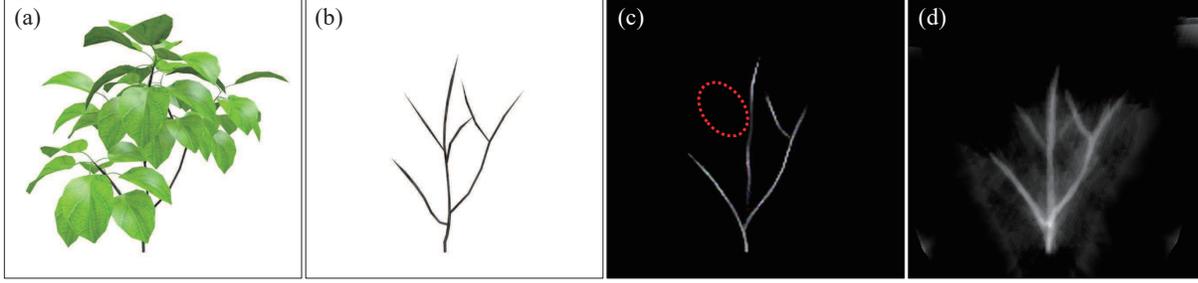


Figure 4: Visualizations of branch probability: a) an input image, b) the ground truth branch, c) 2D branch probability generated by Bayesian image-to-image translation, d) a 2D projection of 3D branch probability via 3D aggregation. Although image-to-image translation (c) infers a few branches as low probability due to heavy occlusions (e.g., in the circle), they are recovered by 3D probability generation (d) due to votes from other views.

to the corresponding branch existence  $\hat{B}_{2D_i,v}$  with the  $v$ -th variation of random dropout patterns. The probability map  $B_{2D_i}$  for each viewpoint  $i$  is then obtained by marginalizing the individual samples  $\hat{B}_{2D_i,v}$  over random trials  $v$ .

### 3.2. Probabilistic 3D branch structure generation

Once the view-wise probability maps  $\{B_{2D_i}\}$  are obtained, our method estimates a 3D probability map  $B_{3D}$  of the branch structure defined in the 3D voxel coordinates. From the multi-view input images, we estimate the camera poses and intrinsic parameters by a structure-from-motion method [33]. It yields a set of projection functions  $\{\theta\}$  that map from the 3D voxel to image coordinates,  $\theta : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ . Using the projections  $\{\theta\}$ , the probability of the branch existence  $B_{3D}$  at voxel  $\mathbf{x}_{3D} \in \mathbb{R}^3$  can be computed as a joint distribution of  $\{B_{2D_i}\}$  by assuming their independence as

$$B_{3D}(\mathbf{x}_{3D}) = \prod_i B_{2D_i}(\theta_i(\mathbf{x}_{3D})), \quad (2)$$

in which  $\theta_i$  represents a projection from the voxel to the  $i$ -th image coordinates. This aggregation process can be regarded as a back projection, which is used in traditional computed tomography [5]. While any of the views may not convey complete information of the branch structure due to heavy occlusions, the aggregation effectively recovers the branch structure in a probabilistic framework as depicted in Fig. 4 (d). To avoid numerical instability, Eq. (2) is computed in the log domain.

### 3.3. Explicit branch generation using particle flows

The probabilistic 3D branch structure can be converted to an explicit representation of 3D branch models that can be used for applications in computer graphics and branch structure analysis. Inspired by a conventional tree modeling approach [20], we develop a branch structure generation method using particle flows. Instead of relying on the attractor graph that is computed directly from images in [20],

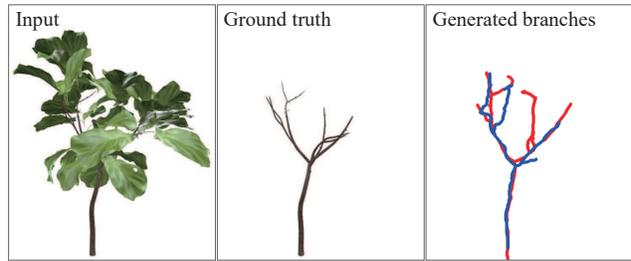


Figure 5: Different branch structures (shown in red and blue) generated from the same branch probability map, by changing position and the number of initial particles.

our method uses the 3D probability map to regulate the particle movement for generating branch path candidates. Finally, the candidates are consolidated via structure refinement that includes smoothing and simplification. The resulting 3D model is represented by a graph that consists of nodes and edges that correspond to joints and branches. Using the particle-flow-based method, our method can yield different plant structures from a single probability map, by changing position and numbers of initial particles as shown in Fig. 5. This is one of the merits of our probabilistic plant modeling for graphics applications.

We first generate particles proportionally to the log probability map  $\log B_{3D}$  for avoiding the peaky distribution in the canonical domain. Also, the root position  $\mathbf{r} \in \mathbb{R}^3$  of the plant is set to the bottommost point that has a high probability of being a part of branches. Under these settings, starting from random distribution of particles, the particle positions are iteratively updated in a manner of flow simulation. At each step  $t$ , the  $(t+1)$ -th particle position  $\mathbf{p}_{t+1}$  is computed by the following update rules:

$$\begin{cases} \mathbf{p}_{t+1} \leftarrow \mathbf{p}_t + \mathbf{F}(\mathbf{p}_t), \\ \mathbf{F}(\mathbf{p}_t) = \lambda_c \mathbf{F}_c(\mathbf{p}_t) + \lambda_d \mathbf{F}_d(\mathbf{p}_t) + \lambda_r \mathbf{F}_r(\mathbf{p}_t). \end{cases} \quad (3)$$

As illustrated in Fig. 6,  $\mathbf{F}_c(\mathbf{p}_t)$  and  $\mathbf{F}_d(\mathbf{p}_t)$  represent nor-

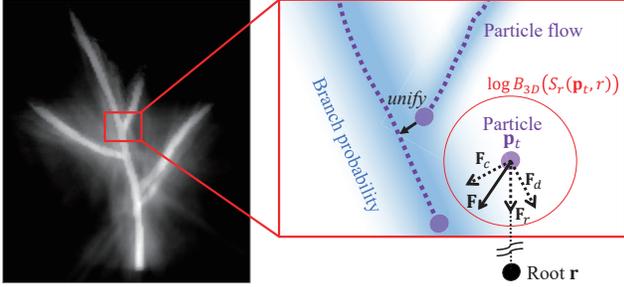


Figure 6: Particle flow simulation. At each step, particles move according to the force computed from the 3D branch probability. The traces of particles are stored as branch path candidates.

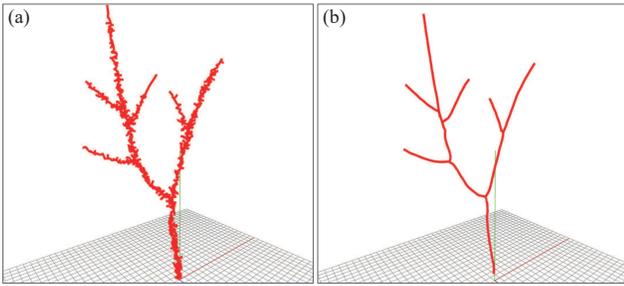


Figure 7: Branch structure refinement: a) Branch structure candidates by flow simulation, and b) the final output after refinement and simplification.

malized vectors toward and parallel to the stream of branch probability.  $\mathbf{F}_r(\mathbf{p}_t)$  represents the unit direction from  $\mathbf{p}_t$  to the root point  $\mathbf{r}$  of the plant. These directions are linearly combined with weight factors  $\lambda_c$ ,  $\lambda_d$ , and  $\lambda_r$  that are determined empirically. The traces of the particles are recorded in a tree graph as vertices and edges, and a unification of particles is treated as a joint.

While the flow simulation generates a lot of branch candidates as shown in Fig. 7(a), they are simplified and refined to yield the final structure. The refinement process involves (1) Smoothing: Apply low-pass filter along branch paths. We simply update the position of each vertex to the mass center of neighboring vertices, (2) Refinement: Move each vertex toward the direction of local probability maximum perpendicular to current branch direction, and (3) Simplification: Unify vertices located close to each other, and delete subtrees that locate in areas with small probability. We iterate the steps (1)–(3) several times and acquire the final branches as a tree graph structure as shown in Fig. 7(b).

### 3.4. Implementation Details

Here we describe some implementation details on the Bayesian image-to-image translation and 3D branch generation.

**Generator training.** For generating the view-wise probability maps  $\{B_{2D_i}\}$ , we train the Pix2Pix network using images rendered using 10 synthetic plants. The plant models are created by changing the parameters of a self-organizing tree model [22], which is implemented in L-studio [24], an L-system-based plant modeler. For each plant, we render images viewed from 72 viewpoints, and each image is additionally flipped for data augmentation. As a result, the total number of images is  $10 \times 72 \times 2 = 1440$ . Because Pix2Pix employs PatchGAN [12] that divides an image into patches for training, we find that this number of images is sufficient for training the generator  $\pi_i$  of Eq. (1).

The 2D branch probability estimation is implemented by modifying a TensorFlow implementation of Pix2Pix [1]. To realize Bayesian image-to-image translation described in Sec. 3.1, we modify the original implementation by adding dropout layers to each of four central layers of the encoder and decoder blocks. Since we intend to model the uncertainty of contextual information in deeper layers (*i.e.*, branch patterns), we set a larger dropout rate in central layers (0.8 for the innermost layer, 0.7 and 0.6 for the next two, and 0.5 for the outermost layer). Although [13] uses a fixed dropout rate (0.5), they also report that applying dropout only for central layers yields better performance for segmentation tasks. We additionally use dropout for all the skip connections with a fixed dropout rate 0.3. The inference is repeated 100 times by randomly changing the dropout patterns. A previous work reports tens of iterations are sufficient to obtain the probability distributions [13].

**Explicit branch generation.** For particle flow simulation described in Sec. 3.3, we generate approximately 10,000 particles in the 3D space. Let  $\log B_{3D}(S_r(\mathbf{p}_t))$  as the subset of the log-probability map around the particle location  $\mathbf{p}_t$  within radius  $r$ , we computed  $\mathbf{F}_c(\mathbf{p}_t)$  and  $\mathbf{F}_d(\mathbf{p}_t)$  as the mass center and a major axis of  $\log B_{3D}(S_r(\mathbf{p}_t))$ . We set  $\lambda_r$  as a constant (0.1 in our experiment) but change  $\lambda_c$  and  $\lambda_d$  according to the distance  $d_c$  between  $\mathbf{p}_t$  and mass center of  $B_{3D}(S_r(\mathbf{p}_t))$  as

$$\lambda_c = \frac{d_c}{r}(1 - \lambda_r), \quad \lambda_d = 1 - \lambda_r - \lambda_c. \quad (4)$$

In this manner,  $\lambda_c$  becomes smaller when the mass center gets closer to  $\mathbf{p}_t$ . The implementation of the weight control is based on the previous approach [20].

**Resolution and processing time.** We used  $128 \times 128$  [px] images for Pix2Pix, and they were upsampled to  $1024 \times 1024$  [px] for the 3D aggregation process. With our unoptimized implementation, it took approximately 10 [sec] for the Pix2Pix with variational inference (100 inferences) per a viewpoint using an NVIDIA GeForce GTX 1080 GPU. Also, it took 1 [min] for 3D aggregation, and 30 [sec] for particle simulation on a CPU (3.70GHz, 6 cores).

		Geometric error (Euclidean distance)				Structure error (difference in joint numbers)			
		72	36	12	6	72	36	12	6
2D branch map $\{B_{2D_i}\}$	Visible branch region	4.20	3.98	3.59	18.05	0.6	<b>0.6</b>	<b>1.0</b>	40.4
	Whole plant region [20]	12.37	18.12	24.75	17.44	305.4	117.6	65.4	<b>35.0</b>
	Image-to-image translation	1.76	2.26	2.40	14.95	0.8	8.2	13.6	88.2
	Bayesian image-to-image translation (proposed)	<b>1.69</b>	<b>1.74</b>	<b>2.14</b>	<b>14.53</b>	<b>0.4</b>	1.4	7.4	60.2

Table 1: The accuracy of 3D branch structure generated using different settings and number of views (averaged over five simulated plants). The geometric error is defined on a relative scale.

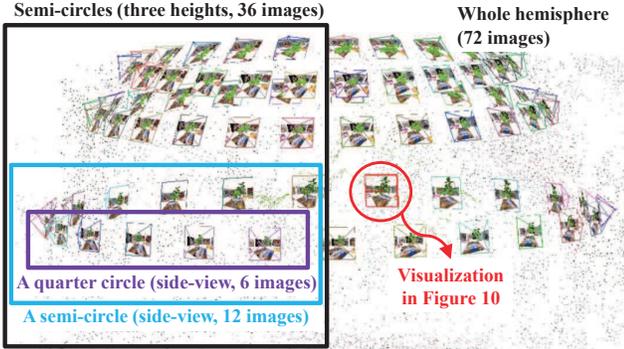


Figure 8: Camera setting employed in the experiments.

## 4. Experiments

We conduct experiments using simulated and real-world plant images and assess the quality of the reconstruction both quantitatively and qualitatively.

### 4.1. Simulated plants

**Quantitative evaluation.** We first describe the experimental result using simulated plant models, which have the ground truth branches for quantitative evaluation. The accuracy of generated plants by the proposed method is assessed using two metrics; *geometric* and *structure* errors.

The geometric error evaluates the Euclidean distances between 3D points in the generated 3D branch structure and points in the ground truth branches. We use a graph representation of the generated 3D branch structure and the ground truth. By sampling 3D points on the graph edges of both generated and the ground truth branches, we assess the geometric error. Let  $\mathbf{g} \in \mathcal{G}$  and  $\mathbf{t} \in \mathcal{T}$  be generated and true 3D branch points, respectively. The geometric error is defined as a bidirectional Euclidean distance  $d$  [36] between the two point sets written as

$$d(\mathcal{G}, \mathcal{T}) = \frac{1}{2} \left( \frac{\sum_{\mathcal{G}} \|\mathbf{g} - N_{\mathcal{T}}(\mathbf{g})\|}{|\mathcal{G}|} + \frac{\sum_{\mathcal{T}} \|\mathbf{t} - N_{\mathcal{G}}(\mathbf{t})\|}{|\mathcal{T}|} \right),$$

where  $N_{\mathcal{G}}(\mathbf{x})$  and  $N_{\mathcal{T}}(\mathbf{x})$  are functions to acquire the nearest neighbor point to  $\mathbf{x}$  from point sets  $\mathcal{G}$  and  $\mathcal{T}$ , respectively, and  $|\mathcal{G}|$  and  $|\mathcal{T}|$  denote the numbers of points in  $\mathcal{G}$  and  $\mathcal{T}$ . The geometric error is defined only up to scale be-

cause our branch recovery method is also up to scale like most multi-view 3D reconstruction methods.

We define the structure error as the difference in joint numbers between the generated and the ground truth structures. From a tree graph generated by the explicit branch generation, we count the number of graph vertices, where the number of connected edges is three (*i.e.*, one incoming and two outgoing) or more, for assessing the accuracy of structure recovery.

Using these error metrics, we assess the method using four different settings described below:

- **Visible branch region:** As the easiest test case, we use the ground truth branch regions in multi-view images that are not occluded by leaves. It naturally skips the image-to-image translation process because the 2D branch structure map  $\{B_{2D_i}\}$  is directly given.
- **Whole plant region:** In this test case, to compare with the image-to-image translation strategy in our problem, a whole plant (*i.e.*, branches and leaves) region is directly used without image-to-image translation for defining the probability map  $\{B_{2D_i}\}$ . This setting is akin to the previous tree modeling method [20], which generates branches that fit the entire volume of a tree by a particle flow simulation.
- **Image-to-image translation:** In this case, the 2D probability map  $\{B_{2D_i}\}$  are generated by the original image-to-image translation [12], where the Bayesian extension is not employed.
- **Bayesian image-to-image translation (Proposed):** In this setting, we use the proposed method, with which branch probability maps  $\{B_{2D_i}\}$  are generated by the Bayesian image-to-image translation.

For the evaluation, we use five plant 3D models, where the branching parameters are different from the training dataset described in Sec. 3.4. To assess the effect of the varying number of cameras, we generate branch structures from multi-view images (a) covering a whole hemisphere (72 cameras), (b) camera paths on semi-circles at three heights (36 cameras), (c) a semi-circle (12 cameras), and (d) a quarter circle (6 cameras), as illustrated in Fig. 8.

**Results.** Table 1 summarizes the accuracy evaluation of the four settings with varying number of views, and Fig. 9

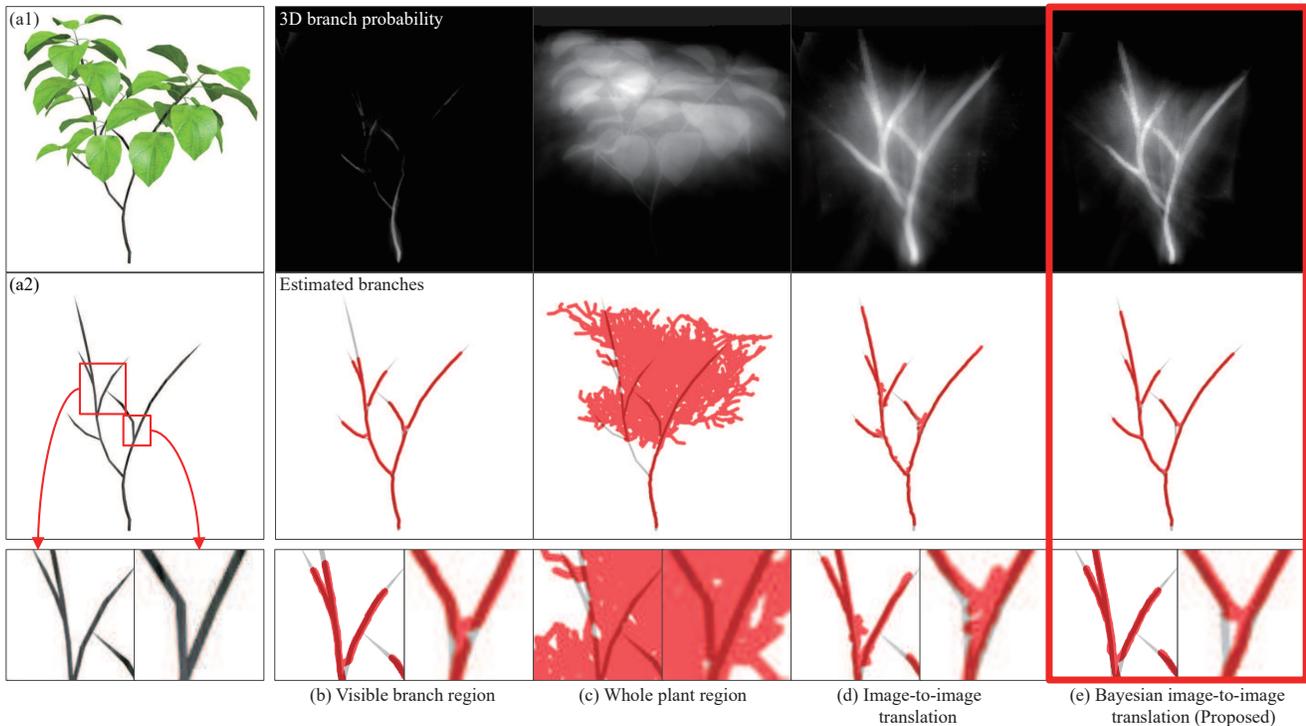


Figure 9: Results using a simulated plant. The leftmost column shows (a1) an input image and (a2) the ground truth branch structure. The other columns show the 3D branch probability and the resultant branch structures using different settings. In comparison to other settings, the proposed approach (e) generates accurate and stable branch structure.

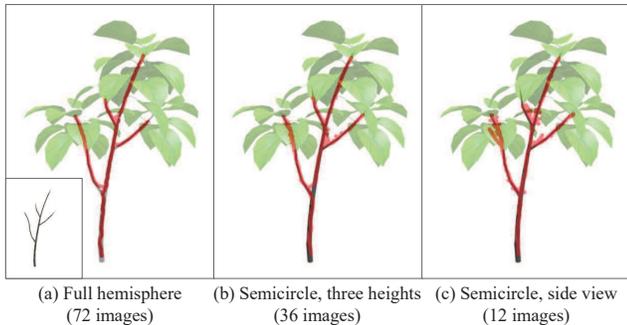


Figure 10: Results when reducing the cameras. The ground truth branch is shown in bottom left corner.

shows some of the results when using 72 images. The proposed approach generates accurate branch structures across all the view settings in geometric error, while the accuracy degrades in the smaller number of views. For the case of “visible branch regions,” although we extracted the ground truth branch labels seen from each camera, branches are occluded by leaves and there are parts that are unobserved from most of the viewpoints, resulting in greater geometric errors. The case of “image-to-image translation,” the output includes branch paths that are not supported by the reliability of inference, instead generated from one sample, which results in small sub-branches (see Fig. 9(d)). Since the pro-

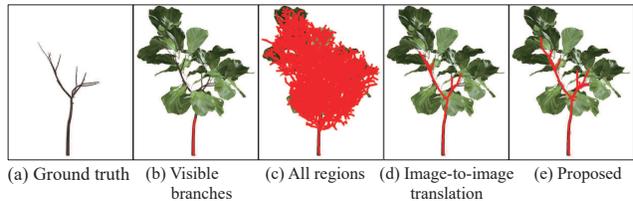


Figure 11: Results using a plant of a different species than the ones used for training.

posed approach averages multiple inferences, the estimated branches become more stable and precise compared to the non-Bayesian approach, thus shows better agreement between the estimated probability map and the branch paths.

Figure 10 compares the results of the proposed method together with the Bayesian image-to-image translation along the varying number of views. It can be seen that the more views make the estimation more faithful, while the method is still able to recover the overall branch structures.

**Different species.** To study the generalization ability, we apply the proposed method to plant models where their species (*i.e.*, texture and shape of leaves) are different from the simulated plants used for training. Figure 11 shows that the proposed method still generates accurate branch structures compared to the other approaches even for different leaf textures.

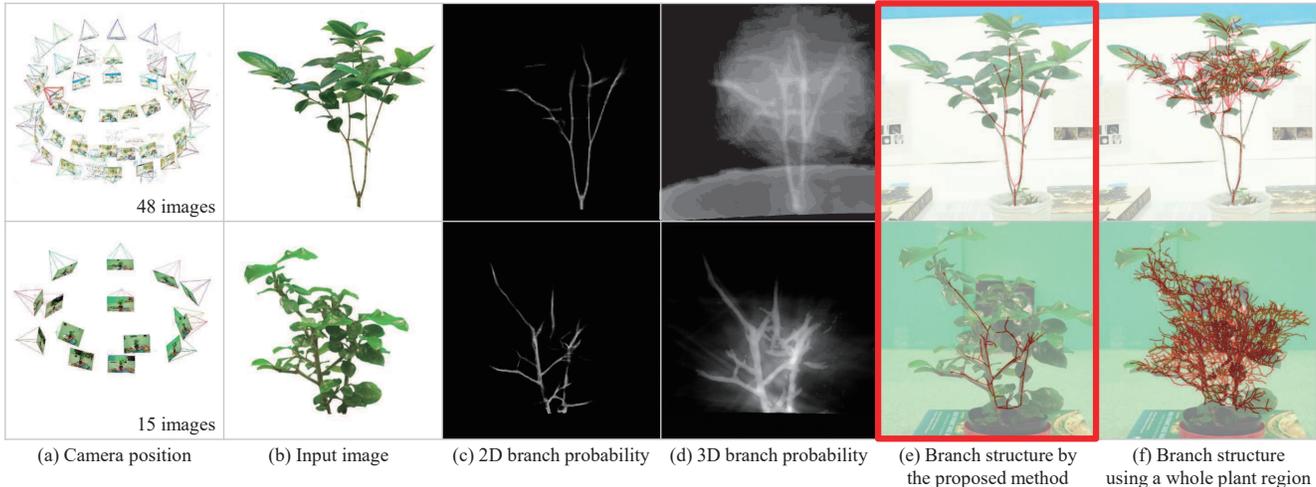


Figure 12: Results using real plants. Using input images where the foreground plants are extracted (b), the proposed approach generates convincing branch structures (e), in comparison to branch structures generated by a previous tree reconstruction approach [20] (f).

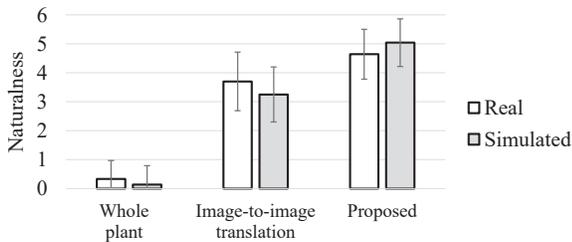


Figure 13: Subjective evaluation of the branch naturalness.

## 4.2. Real plants

Now we show the result of the proposed method using the images of real-world plants. In this experiment, we use the same trained model for the image-to-image translation that is used in the simulation experiment in Sec. 4.1. To avoid the unmodelled factors in the experiment, we first manually extract the plant regions from images, which can be alternatively achieved by chroma keying. Figure 12 shows the results including the intermediate outputs in comparison to the method [20]. The proposed method qualitatively yields convincing branch structures, even though the image-to-image translation network is trained using simulated plants. Compared with a previous tree reconstruction approach [20], our approach shows its effectiveness in generating the branch structure under leaves.

Since for this experiment we did not have access to the ground truth, we conducted a small subjective evaluation by 10 participants to assess the perceptual naturalness of output plant skeletons. The participants were asked to watch the skeletons generated by each method overlaid on the plant image and assessed naturalness of the skeletons in a 7-step Likert scale. Figure 13 summarizes the result. The pro-

posed approach yields the best score among the comparisons, and the trend is consistent with the evaluation of simulated plants.

## 5. Discussion

We presented a plant modeling approach via image-to-image translation to estimate branch structures of 3D plants from multi-view images, even if the branches are occluded under leaves. The combination of Bayesian image-to-image (leafy- to branch-image) translation and 3D aggregation generates the branch existence probability in a 3D voxel space, resulting in a probabilistic model of a 3D plant structure. We have shown that explicit branch structures can be generated from the probabilistic representation via particle flow simulation.

The experimental results using simulated plants showed that the proposed approach is able to generate the accurate recovery of the branch structure of a plant compared with a previous tree modeling approach [20]. It has also been shown that a Bayesian extension of image-to-image translation is effective in obtaining a stable estimate of the branch structure in the form of probability compared with a non-Bayesian one. Qualitatively, the result also applies to the real-world plants as demonstrated in the experiment.

One of the advantages of our probabilistic approach is that it can generate a variety of plant branch structures from a single branch probability map, by changing the parameters of particle flow simulation, as shown in Fig. 5, which may benefit computer graphics applications.

## Acknowledgement

This work was supported in part by JST PRESTO Grant Number JPMJPR17O3.

## References

- [1] Image-to-image translation in tensorflow. <https://affinelayer.com/pix2pix/>.
- [2] O. Argudo, A. Chica, and C. Andujar. Single-picture reconstruction and rendering of trees for plausible vegetation synthesis. *Computers & Graphics*, 57:55–67, 2016.
- [3] C. Barnes, E. Shechtman, A. Finkelstein, and D. Goldman. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. on Graphics*, 28(3):24, 2009.
- [4] F. Boudon, P. Prusinkiewicz, P. Federl, C. Godin, and R. Karwowski. Interactive design of bonsai tree models. *Computer Graphics Forum*, 22(3):591–599, 2003.
- [5] R. A. Brooks and G. Di Chiro. Theory of image reconstruction in computed tomography. *Radiology*, 117(3):561–572, 1975.
- [6] Z. Cao, T. Simon, S. Wei, and Y. Sheikh. Realtime multi-person 2D pose estimation using part affinity fields. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'17)*, 2017.
- [7] S. Darabi, E. Shechtman, C. Barnes, D. B. Goldman, and P. Sen. Image melding: Combining inconsistent images using patch-based synthesis. *ACM Trans. on Graphics*, 31(4):82, 2012.
- [8] A. A. Efros and W. T. Freeman. Image quilting for texture synthesis and transfer. In *Proc. SIGGRAPH'01*, pages 341–346, 2001.
- [9] Y. Gal and Z. Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proc. 33rd Int'l Conf. on Machine Learning (ICML'16)*, pages 1050–1059, 2016.
- [10] C. Galbraith, L. Muendermann, and B. Wyvill. Implicit visualization and inverse modeling of growing trees. *Computer Graphics Forum*, 23(3):351–360, 2004.
- [11] A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin. Image analogies. In *Proc. SIGGRAPH'01*, pages 327–340, 2001.
- [12] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'17)*, 2017.
- [13] A. Kendall, V. Badrinarayanan, and R. Cipolla. Bayesian SegNet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. In *Proc. 28th British Machine Vision Conference (BMVC'17)*, 2017.
- [14] S. Lefebvre and H. Hoppe. Parallel controllable texture synthesis. *ACM Trans. on Graphics*, 24(3):777–786, 2005.
- [15] Y. Li and Y. Gal. Dropout inference in Bayesian neural networks with alpha-divergences. In *Proc. 34th Int'l Conf. on Machine Learning (ICML'17)*, 2017.
- [16] Y. Livny, S. Pirk, Z. Cheng, F. Yan, O. Deussen, D. Cohen-Or, and B. Chen. Texture-lobes for tree modelling. *ACM Trans. on Graphics*, 30(4):53, 2011.
- [17] Y. Livny, F. Yan, M. Olson, B. Chen, H. Zhang, and J. El-Sana. Automatic reconstruction of tree skeletal structures from point clouds. *ACM Trans. on Graphics*, 29(6):151, 2010.
- [18] L. D. Lopez, Y. Ding, and J. Yu. Modeling complex unfoliated trees from a sparse set of images. *Computer Graphics Forum*, 29(7):2075–2082, 2010.
- [19] R. M. Neal. *Bayesian Learning for Neural Networks*. Springer Science & Business Media, 2012.
- [20] B. Neubert, T. Franken, and O. Deussen. Approximate image-based tree-modeling using particle flows. *ACM Trans. on Graphics*, 26(3):88, 2007.
- [21] M. Okabe, S. Owada, and T. Igarashi. Interactive design of botanical trees using freehand sketches and example-based editing. *Computer Graphics Forum*, 24(3):487–496, 2005.
- [22] W. Palubicki, K. Horel, S. Longay, A. Runions, B. Lane, R. Měch, and P. Prusinkiewicz. Self-organizing tree models for image synthesis. *ACM Trans. on Graphics*, 28(3):58, 2009.
- [23] L. Pishchulin, E. Insafutdinov, S. Tang, B. Andres, M. Andriluka, P. V. Gehler, and B. Schiele. DeepCut: Joint subset partition and labeling for multi person pose estimation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'16)*, pages 4929–4937, 2016.
- [24] P. Prusinkiewicz, R. Karwowski, R. Měch, and J. Hanan. L-studio/cpfg: A software system for modeling plants. In *Proc. Int'l Workshop on Applications of Graph Transformations with Industrial Relevance*, pages 457–464, 1999.
- [25] A. Reche-Martinez, I. Martin, and G. Drettakis. Volumetric reconstruction and interactive rendering of trees from photographs. *ACM Trans. on Graphics*, 23(3):720–727, 2004.
- [26] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Proc. Int'l Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI'15)*, pages 234–241, 2015.
- [27] I. Shlyakhter, S. Teller, M. Rozenoer, and J. Dorsey. Reconstructing 3D tree models from instrumented photographs. *IEEE Computer Graphics and Applications*, 21(3):53–61, 2001.
- [28] O. Stava, S. Pirk, J. Kratt, B. Chen, R. Měch, O. Deussen, and B. Benes. Inverse procedural modelling of trees. *Computer Graphics Forum*, 33(6):118–131, 2014.
- [29] L. Streit, P. Federl, and M. C. Sousa. Modelling plant variation through growth. *Computer Graphics Forum*, 24(3):497–506, 2005.
- [30] P. Tan, T. Fang, J. Xiao, P. Zhao, and L. Quan. Single image tree modeling. *ACM Trans. on Graphics*, 27(5):108, 2008.
- [31] P. Tan, G. Zeng, J. Wang, S. B. Kang, and L. Quan. Image-based tree modeling. *ACM Trans. on Graphics*, 26(3):87, 2007.
- [32] K. W. Waite. Modelling natural branching structures. *Computer Graphics Forum*, 7(2):105–115, 1988.
- [33] C. Wu. Towards linear-time incremental structure from motion. In *Proc. Int'l Conf. on 3D Vision (3DV'13)*, pages 127–134, 2013.
- [34] H. Xu, N. Gossett, and B. Chen. Knowledge and heuristic-based modeling of laser-scanned trees. *ACM Trans. on Graphics*, 26(4):19, 2007.
- [35] D. Zhang, N. Xie, S. Liang, and J. Jia. 3D tree skeletonization from multiple images based on PyrLK optical flow. *Pattern Recognition Letters*, 2015.

- [36] J. Zhu, S. Du, Z. Yuan, Y. Liu, and L. Ma. Robust affine iterative closest point algorithm with bidirectional distance. *IET Computer Vision*, 6(3):252–261, 2012.
- [37] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proc. Int’l Conf. on Computer Vision (ICCV’17)*, 2017.