

3D Tree Reconstruction from Laser Range Data

Jonathan Binney and Gaurav S. Sukhatme

Abstract— We present a method for reconstructing 3D models of tree branch structure from laser range data. Our approach is probabilistic, and uses a generative model of a tree to guide an iterative reconstruction process. Our goal is to recover parameters such as branch locations, angles, radii, and lengths, as well as connectivity information between branches. These parameters can then be fed into functional-structural plant models to study the relationships between the structure of a plant, its environment, and its internal biology. In this paper we present an algorithm for finding these parameters, and results on both simulated and real datasets.

I. INTRODUCTION

Studies of tree architecture in relation to tree biology and environmental factors often require detailed metric data about the tree structure, such as branch structure, lengths, radii, and angles, as well as leaf areas and positions. [1]. Using this information, computer models of tree function can be created [2], giving valuable information about tree development.

Acquiring this metric information can be extremely time consuming. In one study of an adult walnut tree, [3], measuring the tree architecture by hand using a 3D magnetic measuring device which was moved along the branches took several weeks. We hope to provide a method of acquiring detailed metric information about the branch architecture of a tree which is relatively quick and requires significantly less manual labor. In this paper we address the measurement of the branch structure, but we plan to extend our technique to provide data about the leaves of the tree as well.

This problem is particularly well suited to some of the methods developed in probabilistic robotics. Using sensors to infer metric information about its environment is one of the crucial problems which a robotic system must solve. This task is often broken down into a number of different sub-problems, including localizing the sensors at each step, associating the sensor data with a part of the robot’s environment model, and fitting the model to the data. One example of a method which solves the data association and model fitting problem is [4], which fits a set of planes to laser range data. For the tree problem, the model which we want to fit to sensor data is much more complex.

The method presented in this paper works on laser range data, but can be extended to any sensor for which we can calculate the likelihood, given a specific set of tree model parameters.

Jonathan Binney (binney@usc.edu) and Gaurav S. Sukhatme (gaurav@usc.edu) are with the Robotic Embedded Systems Laboratory and Department of Computer Science at the University of Southern California (USC). This work was funded in part by the National Science Foundation under grants CNS-0540420, CNS-0331481, and CCF-0120778 and a gift from the Okawa Foundation. Jonathan Binney was supported by a doctoral fellowship from USC.

II. RELATED WORK

Previous work on reconstructing tree models from sensor data falls into two main types. The first focuses on methods for creating visually appealing, but not necessarily accurate, models. The second focuses on accurately determining some parameters of the tree, but often requires a large amount of user interaction.

A method for reconstructing realistic looking tree models from laser scan data is described in [5]. The laser data is first converted to a point cloud, and a graph-based technique is used to find the rough branching structure. In order to make the tree more visually appealing, fake branches are added according to the expected structure of the tree.

A graph-based technique is also employed in [6] to find overall branch structure, but images are used as input instead of laser range scans. A structure from motion algorithm is used to create a 3D point cloud from the images. The point cloud and the raw images are used to find the branching structure. Visible branches are copied to fill occluded areas and create a visually appealing tree.

The algorithm in [7] and [8] takes laser data as input, and creates a voxel-based occupancy grid representation of the data. It uses morphological operations to find the underlying branching structure, and fits cylinders for the branches. The algorithm is used to extract metric parameters of the tree (e.g. branch radius).

An approach based on instrumented photographs is described in [9] which generates a model that approximates the appearance of a specific tree, although not its exact structure.

Our work is most similar to [10], which uses a generative statistical model to fit trees to image data. It uses a straight cylinder to model each branch, and generates possible trees using L-systems. In [10] complete trees are generated, and then compared against each image.

We use a generative statistical approach like in [10], because it allows us to use complex sensor models without solving a complex inverse problem. It also allows us to make good use of prior knowledge about tree structure in an intuitive way, by including that information in the generative model of the tree. Unlike [10], we use a detailed model of the tree which allows for curved branches. We use laser range data instead of images because it gives us accurate depth information. By making strong use of priors on the tree structure, we hope to create a method which is tolerant to noise. We use a sequential approach instead of trying to fit a model to the entire tree all at once, because the space of possible trees is huge, especially when allowing curved branches and multiple levels of branching.

III. APPROACH

In order to reconstruct a tree from sensor data, we use a model of the tree to generate likely hypotheses, and then use a sensor model to evaluate the likelihood of each hypothesis. This is done in an iterative fashion starting at the base of the trunk, then working upwards in segments, as show in Figure 1. Before we give details of the algorithm, we will describe the tree and sensor models.

A. Generative Tree Model

Our algorithm is based on a generative model of the tree being reconstructed. This model does not need to be perfect, and a less accurate model simply means that more hypotheses will be needed in order to get a good result. Ideally, the model will be specific to the species of tree which is being reconstructed, and will have values chosen from previous reconstruction results.

For this work, the generative model used for each branch is extremely simple. Each branch is broken up into fixed-length segments, as shown in Figure 2(a). Two angles α and β control the relative direction between a segment and its parent. We select β from a uniform distribution over $[0, 2\pi]$, and α from a normal distribution centered around zero (a branch is most likely to be straight) and with some standard deviation. Additionally, there is a multiplier r_{ratio} which relates the radius of each segment to that of the segment before it. The distribution for r_{ratio} is modeled as a Gaussian distribution centered around a value slightly less than 1.0 (since branches tend to get thinner towards the ends), and has a relatively small standard deviation. To model the end of a branch, at each step there is a small constant probability that the branch will terminate. A zero length segment is added to signify the end of the branch.

After finding the trunk, we use another generative model to provide hypothetical branchpoint locations. We assume a uniform probability of the new branch starting at any distance L along the current branch. After L is chosen, we again choose the angles α and β , and the ratio r_{ratio} , although with slightly different distributions to reflect the fact that this is the start of a new branch, and not just a new segment along the same branch.

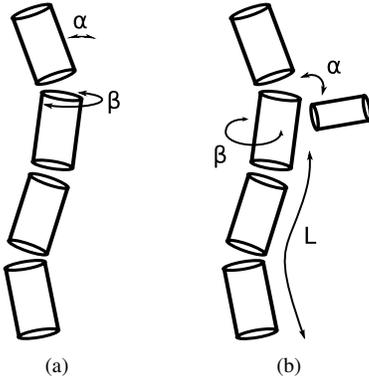


Fig. 2. Branch and Branchpoint models

The values of α , β , and r_{ratio} should ideally be found empirically by looking at previous measurements of the species of tree being reconstructed. During this work we have found that reconstruction of the tree model is possible even with a very general model.

B. Sensor Model

To evaluate hypotheses created by the generative model, we need to be able to calculate $P(Z|hyp_i)$, where hyp_i is a particular segment hypothesis, and Z is the sensor data. This probability comes directly from our model for the sensor. The data we get from the laser is a set of rays, each with a start position, a direction, and a length. Although there may be multiple rays originating at the same sensor pose, for generality we treat each ray separately.

Given the hypothesis segment, z'_j is the ray length which we would expect. We assume a Gaussian distribution over the error $z_j - z'_j$, as in equation 1. In this case we leave out the constant coefficient, because it is irrelevant when we are comparing the relative likelihood of multiple hypothesis segments. We assume that the sensor is well calibrated (that the error in the direction of each ray is negligible.) We also assume that the sensor is well localized (that the error in the start position of each ray is negligible.)

$$P(z_j|hyp_i) \propto \exp(-(z_j - z'_j)^2/2\sigma^2) \quad (1)$$

Equation 1 models the error in the laser readings, but we also need to account for the possibility of other sources of noise. These could be a number of things, such as dust in the air, or as we found, large errors when the beam glances off of the side of small branches. Because we do not have a good model for these errors, we account for them by setting a minimum allowed likelihood for each ray. Any ray with a likelihood below this constant value is considered to be noise and is assigned the minimum allowed likelihood. This gives us:

$$P(z_j|hyp_i) \propto \max(\exp(-(z_j - z'_j)^2/2\sigma^2), P_{min}) \quad (2)$$

Equation 2 gives the likelihood of each ray; to find the likelihood of all of the rays with respect to a given segment we use the usual method of taking the product over the individual likelihoods.

$$P(Z|hyp_i) \propto P(hyp_i) \prod_j P(z_j|hyp_i) \quad (3)$$

z'_j is calculated using ray tracing. Since each segment is a cylinder, there is a closed form solution for finding the intersection with any ray [11]. This method extends easily to much more complicated shapes. If we wanted to allow generalized cylinders, for instance, we would simply need to break them down into a set of quadrics, and then ray trace each one.

This ends our description of the models used; next we describe the algorithm shown in Figure 1.

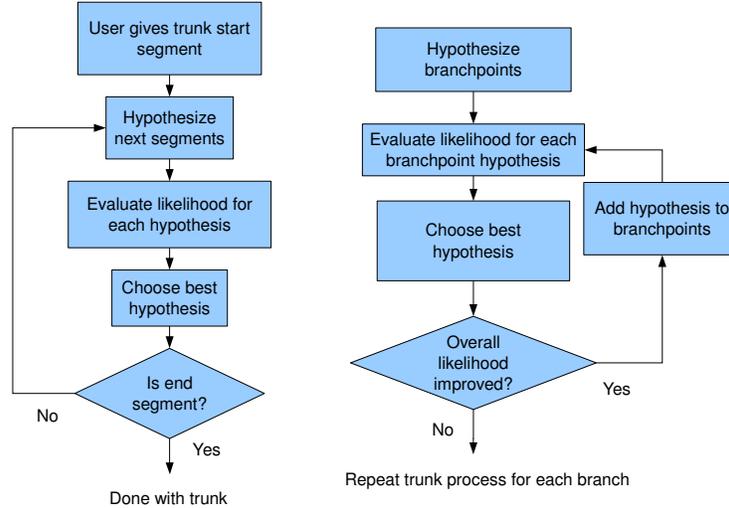


Fig. 1. Branch and Branchpoint models

C. Finding the Trunk

The algorithm for finding the trunk is shown on the left in Figure 1 and described in more detail in Algorithm 1. We explain it below.

Algorithm 1 Find the trunk

```

trunk  $\leftarrow$  [seg0]
while Trunk not terminated do
  Hypothesize  $N_b$  next segments
  for  $i = 1$  to  $N_b$  do
     $w_i = \prod_j P(z_j | hyp_i)$ 
  end for
   $best \leftarrow \operatorname{argmax}_i \{w_i\}$ 
  Append segment  $hyp_{best}$  to trunk
  if  $hyp_i$  is an end segment then
    Done
  end if
end while

```

The user initializes the algorithm by providing the approximate start location, direction, and radius of the trunk. From this known start segment, N_b hypotheses for the next segment are created using the generative model of the tree branch. Next, the likelihood of each of these segments with respect to the sensor data is calculated up to a constant, using equation 4.

$$P(hyp_i | Z) \propto P(hyp_i)P(Z | hyp_i) \quad (4)$$

Here, $P(hyp_i)$ is the prior likelihood of the hypothesis hyp_i . We do not explicitly calculate this, because it is used in our generative model to create the hypotheses. For example, a value of $\alpha_i = \pi/2$ would indicate a right angle in a branch, and such a hypothesis is extremely unlikely to be generated by the model. This leaves us to evaluate $P(Z | hyp_i)$, which is done using the sensor model.

After calculating the likelihood of each hypothesized segment, the one with the highest likelihood is added to the branch, and the process is repeated. Because the generative model creates some hypotheses which are endpoints, termination of branches is automatically taken care of. At the end of branches, the sensor data will give much higher likelihood to endpoints than to other hypothesized segments, and the branch will be terminated.

D. Finding New Branches

Algorithm 2 Find branches

```

 $hyp_1$  to  $hyp_{N_{bp}}$  generated by branchpoint model
loop
  branchpoints  $\leftarrow$  []
  for  $i = 1$  to  $N_b$  do
     $n_{bp} \leftarrow \operatorname{length}(\text{branchpoints})$ 
     $w_{prev} \leftarrow P(n_{bp}) \times P(Z | \text{trunk} \cup \text{branchpoints})$ 
     $w_i \leftarrow P(n_{bp} + 1) \times P(Z | \text{trunk} \cup \text{branchpoints} \cup hyp_i)$ 
  end for
   $best \leftarrow \operatorname{argmax}_i \{w_i\}$ 
  if  $w_{best} > w_{prev}$  then
    Append  $hyp_{best}$  to branchpoints
  else
    Done
  end if
end loop

```

After the trunk is reconstructed, a similar process is applied to reconstruct each branch. Because there may be many branches, it is not reasonable to require a user to provide the start point and direction for every branch. Instead, we have an automated way of finding likely branchpoints. This process is shown on the right side of Figure 1, and described in more detail in Algorithm 2. We explain it below.

First, using the generative model we hypothesize N_{bp} branchpoints. We do not know beforehand how many branchpoints there are, but as part of the model we have a distribution over that number. For a set of n_{bp} specific branchpoints, the likelihood is $P(n_{bp})P(\text{branchpoints})$, where $P(n_{bp})$ is the prior probability that there are n_{bp} branchpoints. Unfortunately, there are $\binom{N_{bp}}{n_{bp}}$ ways to select n_{bp} branchpoints, and we have found that N_{bp} must be on the order of 1000. This gives far too many possibilities to evaluate all of them.

Instead, we use a greedy approximation. First, we calculate the log-likelihood of the sensor data based on the trunk. Next, we re-evaluate the log-likelihood for the trunk plus each branch hypothesis. Good hypotheses increase the log-likelihood by better explaining some rays which hit branches near the trunk. Bad hypotheses decrease the log-likelihood, because they occupy space near the trunk which sensor data suggests is empty. We choose the branchpoint which maximizes $P(Z|\text{trunk} \cup \text{hyp}_i)$, but only if $P(n_{bp} = 1)P(Z|\text{trunk} \cup \text{hyp}_i) > P(0 \text{ branches})P(Z|\text{trunk})$. Then we repeat the process to choose a second branch, third branch, and so forth until we reach a point at which the likelihood decreases.

After this process, we are left with a set of branchpoints, each with a start position, direction and radius. Beginning at each branchpoint, we repeat the process used to reconstruct the trunk. After each branch is reconstructed, the same branchpoint finding process can be repeated recursively to find sub-branches, although we do not present results for that here.

IV. RESULTS WITH SIMULATED DATA

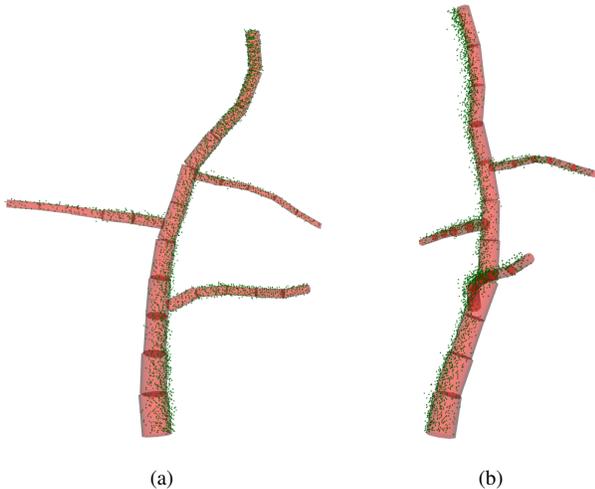


Fig. 3. Simulated data overlaid on ground truth model

We first tested our algorithm on simulated data, for which we had ground truth about the tree structure, sensor poses and sensor noise. To create a number of datasets with similar characteristics, the number of branches in the generative model was fixed, as was the length in segments of the branches. Then random models were generated using our model. After the generative model was used to create the

simulated trees, data was generated for a simulated laser range scanner moving in the direction of the trunk, scanning cross-sections of the branch. All data was taken from one side of the branch, since this is usually enough to reconstruct the tree model, and is easier to acquire when working with real trees. The laser was translated in 4 mm increments, and a scan with 180 rays each spaced one degree was taken at each position.

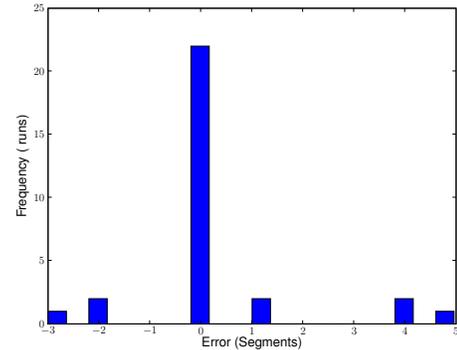


Fig. 4. Trunk length errors on simulated data

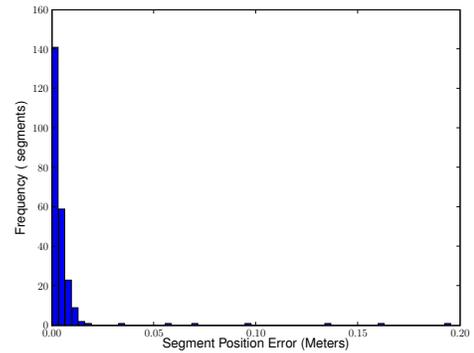


Fig. 5. Segment position errors on simulated data

As described in the sensor model, ray-tracing was used to create the simulated data. After finding the expected length of each ray, zero-mean Gaussian noise with a standard deviation of 5mm is added to the length of each ray to better simulate real data. Figure 3 shows the endpoints of the rays overlaid on the ground truth model for one of the simulated models.

Because it is easier to express the metric error for the branch fitting than for the sub-branch finding, 30 trials were first run on simulated branches with no sub-branches.

Figure 4 shows the length errors in segments for the 30 runs. Most of the reconstructed models had exactly the right number of segments. In several trials, the trunk following algorithm failed, usually because it failed to follow a turn in the branch. These show up as outliers which are much longer or shorter than the ground truth. This could be remedied somewhat by using more hypotheses, at the cost of running time.

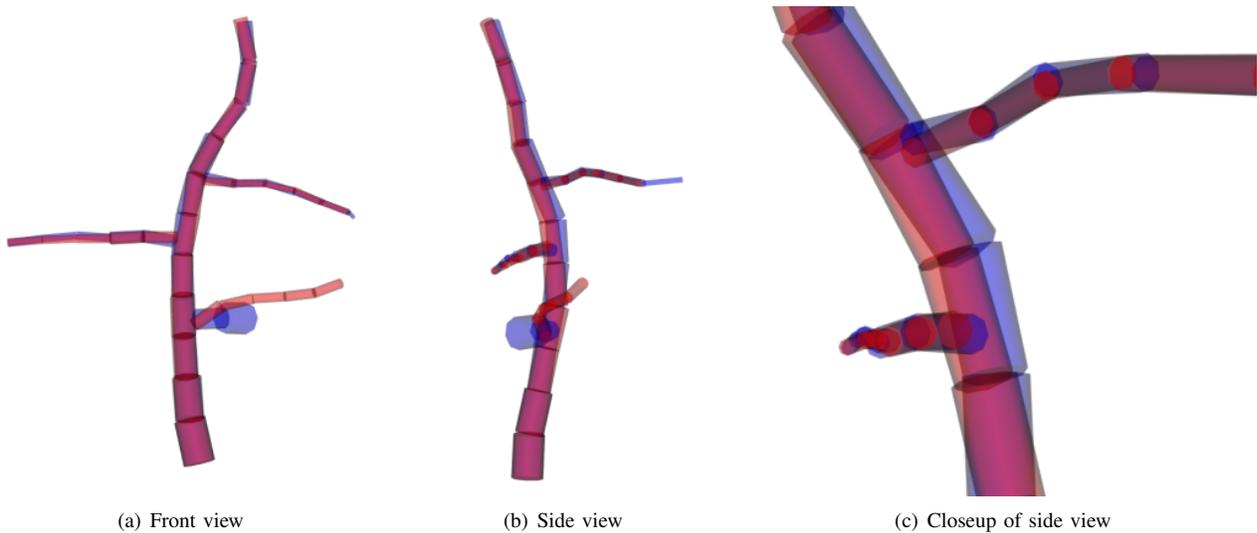


Fig. 7. Reconstruction results on simulated data (blue), overlaid on the ground truth (red)

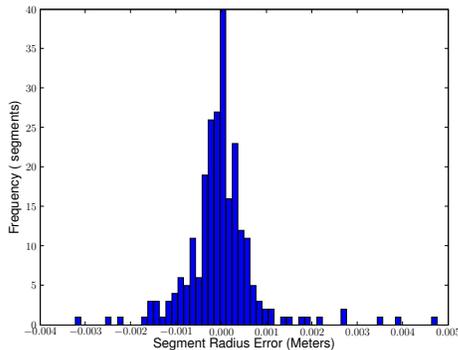


Fig. 6. Segment radius errors on simulated data

For the reconstructed trunks which were the correct length, Figure 5 shows distribution of the position errors between each correspond segment in the actual and reconstructed models. Most of the segments are within about a centimeter of where they should be. There are several outliers, in cases where the trunk finding algorithm got off track at some point.

Figure 6 shows the distribution of mean radius errors for corresponding actual and reconstructed segments. Most errors here are less than a millimeter, largely because over-estimating the radius would cause a segment to overlap rays which are known to continue, giving it low likelihood, and underestimating the radius would fail to explain rays which hit the branch, also giving it low likelihood.

Figure 7 shows the ground truth model and the reconstructed model overlaid. As found in the trials, the trunk following works quite well. Two of the branches are found correctly, but one fails badly (most likely because the initial guess for the branchpoint was bad). Because of the length of the trunk and number of possible angles, the space of possible branchpoints is very large. Because of this, even with a fairly large number of branchpoint hypotheses, in

some cases none may be close enough. In this particular case 10,000 hypotheses were used.

V. RESULTS WITH REAL DATA

To test the method in more realistic conditions, we scanned a branch with a SICK LMS200 laser scanner. The scanner takes 180 range measurements at 1 degree intervals, scanning back and forth at 75 Hz. Instead of an entire tree, one (leafless) branch was used, to allow collection of data in a laboratory setting. The scanner was started at the base of the branch, and translated along the length of the branch. A second laser was used to localize the main laser along the translational axis. The laser scan data consists of 41,876 rays, each with a start point, direction, and measured length.

The radii of the branch and its subbranches vary from 2 centimeters down to about 0.5 centimeters, and the length of the main branch is 78 centimeters. The error listed in the laser's specification is 0.5 centimeters, but we saw very noisy measurements along the smaller branches, especially in places where the laser hit the edge of the branch, as can be seen in figure 9. In these places the error was as much as 5 to 10 centimeters.

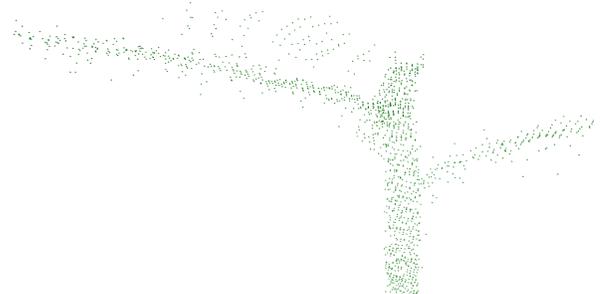


Fig. 9. Closeup of laser data; note the noisy measurements above the left branch

The trunk and all three branchpoints are correctly iden-

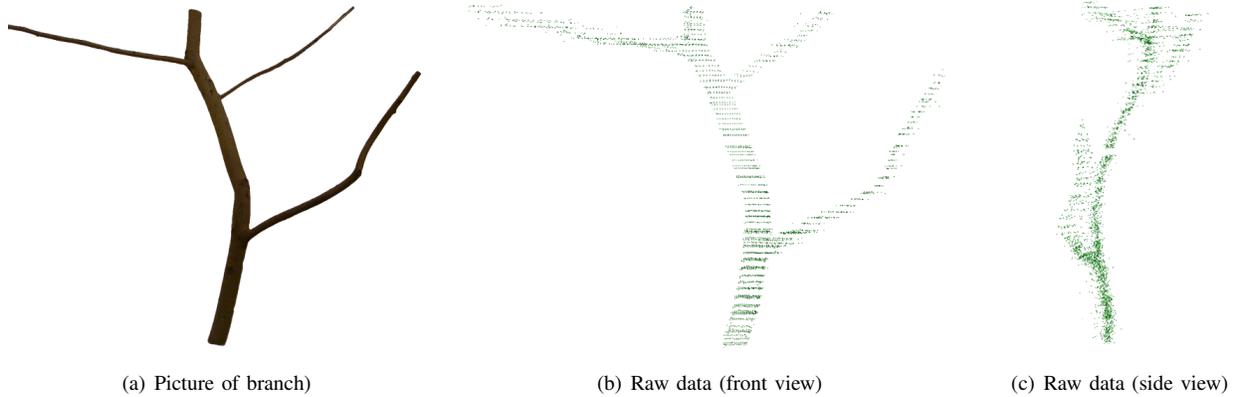


Fig. 8. Collected data

tified. The trunk radius is overestimated by approximately 0.4 cm, most likely because of the large errors in the measurements at the sides of the branch. Because of the fixed segment length, there is some error in the reconstruction of branches at sharp angles. If a branch bends in the middle of a segment, the reconstruction process is unable to account for it, as can be seen in the bottom left branch of figure 10(a). To remedy this, the segment length would need to be varied along with other parameters, or the cylinders would need to be replaced with more complicated shapes. The same issue causes some error in the estimate of trunk length, because the branches can only be terminated between segments.

VI. CONCLUSIONS, DISCUSSION, AND FUTURE WORK

We have developed a technique which can recover tree branch structure from one or more laser scans, and shown that it works reasonably well on both real and simulated data sets.

Using the raw rays from the laser instead of first converting to a point cloud preserves valuable information about what areas of space appear to be empty.

Using a top down, generative model based approach allows us to incorporate prior knowledge about tree structure in an intuitive way to guide the reconstruction process. This has three main advantages.

First, it allows us to easily ignore noise in the sensor data when it does not match any reasonably likely tree structure. Although we do not explicitly identify outlier points, noise is easily dealt with as long as there are more good points than noise points. Fitting entire segments at a time incorporates a large enough number of data measurements at each step to smooth out large amounts of noise.

Second, it means that the parameters which must be set for our algorithm almost all map directly to tangible properties of the tree. For example, we specify the distribution over possible trunk shapes and branching angles. We do not need to specify parameters such as the minimum distance between two data points for them to be considered part of the same branch, as might be required by a bottom up type approach.

Third, we do not need to invert the sensor model function to directly solve for the most likely model parameters. All we

need to be able to do is evaluate the sensor model function for a given instantiation of parameters. Although we got fairly good results with a simple sensor model and using cylinders for hypothesis segments, much more complicated (and accurate) models and segment types could be used with little additional effort.

One disadvantage of our current method is that it requires all sensor data to be explainable by some part of the model. A better approach would simultaneously reconstruct the tree and decide which points belonged to the tree, and which were unimportant. Another disadvantage is that because of the iterative nature of our approach, mistakes early on could lead to extremely bad results. We are looking for ways to include more global information at each step, to avoid local errors leading to poor solutions.

In the future we would like to combine data from both imaging and ranging sensors, since they each have their own strengths. We also plan to incorporate actuation in the data collection process, so that additional data can be collected automatically to improve the model. This actuation could be done by mobile robots, or by some simpler mechanism. There are also some interesting possibilities for combining sensor localization with the model fitting process.

There are some interesting ways in which the reconstruction process could be sped up. Most of the computation time is taken up by raytracing, an operation which is repeated millions of times each time the algorithm is run. At each step tens of thousands of rays must be traced, but the results do not depend on each other, and could be parallelized. For best results, this could be done in several processes on a multi-core CPU, or even in hundreds of processes on a modern GPU.

The Python/C implementation of this algorithm used to create the results in this paper can be found online at [12]. The simulated and real datasets used are also available for download at that site.

REFERENCES

- [1] R. W. Pearcy and W. Yang, "A three-dimensional crown architecture model for assessment of light capture and carbon gain by understory plants," *Oecologia*, vol. 108, pp. 1–12, 1996.

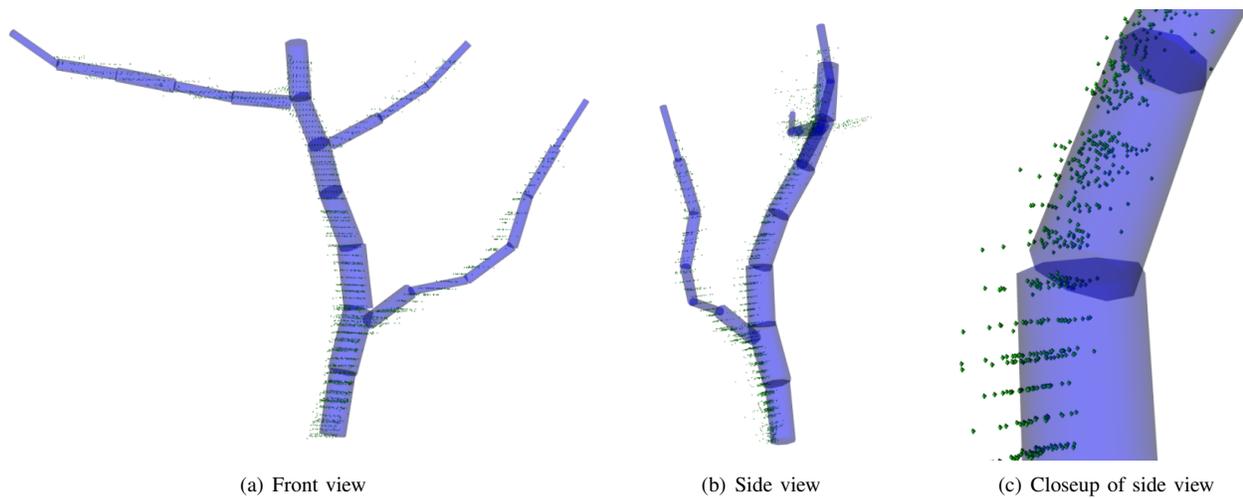


Fig. 10. Reconstructed branch

- [2] C. Godin and H. Sinoquet, "Functional-structural plant modelling," *New Phytologist*, vol. 166, pp. 705–708, 2005. Very good overview of the methods and purpose of functional-structural plant modeling. Briefly mentions some methods for obtaining the models.
- [3] H. Sinoquet and P. Rivet, "Measurement and visualization of the architecture of an adult tree based on a three-dimensional digitising device," *Trees - Structure and Function*, vol. 11, pp. 265–270, April 1997.
- [4] Y. Liu, R. Emery, D. Chakrabarti, W. Burgard, and S. Thrun, "Using em to learn 3d models of indoor environments with mobile robots," in *Proceedings of the International Conference on Machine Learning (ICML)*, 2001.
- [5] H. Xu, N. Gossett, and B. Chen, "Knowledge and heuristic-based modeling of laser-scanned trees," *ACM Transactions on Graphics*, vol. 26, no. 4, p. 19, 2007.
- [6] P. Tan, G. Zeng, J. Wang, S. B. Kang, and L. Quan, "Image-based tree modeling," *ACM Transaction on Graphics (TOG) and Proc. of SIGGRAPH*, 2007.
- [7] B. Gorte and N. Pfeifer, "Structuring laser-scanned trees using 3d mathematical morphology," in *Conference proceedings ISPRS conference*, 2004.
- [8] N. Pfeifer, B. Gorte, and D. Winterhalder, "Automatic reconstruction of single trees from terrestrial laser scanner data," in *Conference proceedings ISPRS conference*, vol. XXXV, pp. 114–119, 2004.
- [9] I. Shlyakhter, M. Rozenoer, J. Dorsey, and S. Teller, "Reconstructing 3d tree models from instrumented photographs," *Computer Graphics and Applications, IEEE*, vol. 21, pp. 53–61, May/June 2001.
- [10] H. Huang and H. Mayer, "Extraction of 3d unfoliated trees from image sequences via a generative statistical approach," *Lecture Notes in Computer Science*, vol. 4713/2007, pp. 385–394, 2007.
- [11] C. A. Lindley, *Practical Ray Tracing in C*. John Wiley & Sons, 1992.
- [12] "<http://robotics.usc.edu/~binney/icra2009/>".