

Learning to Simulate Tree-Branch Dynamics for Manipulation

Jayadeep Jacob^{1,2}, Tirthankar Bandyopadhyay², Jason Williams², Paulo Borges², Fabio Ramos^{1,3}

Abstract—We propose to use a simulation driven inverse inference approach to model the dynamics of tree branches under manipulation. Learning branch dynamics and gaining the ability to manipulate deformable vegetation can help with occlusion-prone tasks, such as fruit picking in dense foliage, as well as moving overhanging vines and branches for navigation in dense vegetation. The underlying deformable tree geometry is encapsulated as coarse spring abstractions executed on parallel, non-differentiable simulators. The implicit statistical model defined by the simulator, reference trajectories obtained by actively probing the ground truth, and the Bayesian formalism, together guide the spring parameter posterior density estimation. Our non-parametric inference algorithm, based on Stein Variational Gradient Descent, incorporates biologically motivated assumptions into the inference process as neural network driven learnt joint priors; moreover, it leverages the finite difference scheme for gradient approximations. Real and simulated experiments confirm that our model can predict deformation trajectories, quantify the estimation uncertainty, and it can perform better when base-lined against other inference algorithms, particularly from the Monte Carlo family. The model displays strong robustness properties in the presence of heteroscedastic sensor noise; furthermore, it can generalise to unseen grasp locations.

Index Terms—Probabilistic Inference, Simulation and Animation, Field Robots

I. INTRODUCTION

SIMULATION driven learning approaches have been gaining popularity to address the challenging task of manipulating deformables. In this paper, we study the deformation characteristics of a tree branch under forces exerted by a robotic arm via simulation-based parameter inference. Learning the branch dynamics through simulation has significant practical ramifications, particularly for the robotics community. First, it adds the missing physics link to visual tree reconstruction works; for example [1][2], to create a complete digital twin with both perceptual features and inferred dynamical properties. Second, such a digital model reduces the reality gap, often denoted as the real-to-sim gap in the



Fig. 1: Autonomous navigation task: Our field robot obstructed by vegetation.

robotics context, and facilitates massively parallel policy learning, inexpensive data collection, safe exploration, and most importantly, real-time control estimation. On the application front, the presence of occlusions from branches, other fruit clusters, stems and foliage pose a significant challenge to robotic fruit harvesting tasks [3], which are typically addressed through costly environment modifications to improve visibility and reachability [4]. The harvest success rates of state-of-the-art methods, in the presence of partial and complete occlusions are 50-75% and 5% respectively [3]. Along the same lines, for a mobile robot to perform autonomous navigation in natural environments, such as forests and grasslands; where plant foliage, non-compliant branches and dense vegetation pose obstacles; the conventional approach has been to avoid them altogether, quite unlike human subjects who seamlessly couple locomotion and manipulation to interact with objects, clear the path, and navigate. Having a simulated twin with inferred dynamics can overcome these barriers to improve fruit detection rate and grasp quality in harvesting tasks, and path planning with contact around obstacles for locomotion. Finally, modelling the deformation behaviour across all branch grasp locations allows the arm to effectively trade-off the low force and high reach required at the outer edge vs the high force and low reach at the inner fork, thereby, amplifying the autonomous decision making capacity of the manipulator.

In general, the goal of an inverse inference approach is to estimate the latent parameters of a physical phenomenon based on its observed effects and subsequently to forecast future system behaviour. A subclass of this model, commonly known as simulation-based inference, leverages high fidelity simulations to implicitly define a statistical model and perform the inference through either frequentist or Bayesian paradigms. The parameters themselves may not necessarily have any physical meaning [5]; moreover, the objective of the estimation is to predict a useful outcome rather than to accurately estimate the true parameters [6].

Manuscript received: October, 09, 2023; Accepted December, 08, 2023.

This paper was recommended for publication by Editor Aleksandra Faust upon evaluation of the Associate Editor and Reviewers' comments.

^{1,2}Jayadeep Jacob is with the School of Computer Science, The University of Sydney, NSW, Australia, and with Data61, CSIRO, Pullenvale, QLD, Australia. jjac4485@sydney.edu.au; jay.jacob@data61.csiro.au

²Tirthankar Bandyopadhyay, Jason Williams, and Paulo Borges are with the Robotics and Autonomous System Group, Data61, CSIRO, Pullenvale, QLD, Australia. (tirtha.bandy, jason.williams, paulo.borges)@data61.csiro.au;

^{1,3}Fabio Ramos is with the School of Computer Science, The University of Sydney, NSW, Australia, and with the NVIDIA Corporation, Seattle, USA fabio.ramos@sydney.edu.au

Digital Object Identifier (DOI): see top of this page.

Copyright ©2024 IEEE

A broad outline of our workflow is as follows: to begin with, we capture the branch geometry with crude spring abstractions using purpose-built simulations executed on massively parallel, GPU based, black-box simulators. Next, we actively probe the real branches to measure the observed effects of forces. While individual twigs are rigid by nature, the additional degrees of freedom incurred due to the dynamics of coupled branches makes the problem non-linear and high dimensional. We consider the measured deformation trajectories from the real world as i.i.d samples from a ground truth distribution. Finally, guided by a Bayesian model, we use a probabilistic inference algorithm to infer the posterior density of spring parameters and to predict the branch deformation behaviour.

Overall, our work lies within the broader objective of translating the success of manipulation schemes within structured environments, guided by simulation driven inverse inference [7][8], out to the natural world. However, simulation of outdoor environment is beset by both epistemic (e.g., unaccounted factors, such as wind; model simplicity) and aleatoric (e.g., blurred vision due to limited lighting, noisy sensors) uncertainties. Therefore, optimisation models that capture point estimates of simulation parameters, from the frequentist paradigm, are not sufficient in place of an uncertainty aware Bayesian model, that we use. Although not the focus of this paper, we postulate that, on a temporal scale, each inferred posterior can feed into the next iteration as a prior, leading to faster convergence over time and that the posteriors can generalise to different branches, trees and perhaps even across species.

Our experiments demonstrate that the learnt model can estimate deformation trajectories given the end-effector forces, quantify the model uncertainty, display robustness in presence of perturbations and can perform the inference in near real time. Overall, our contributions are summarised as below:

- 1) We propose to model the dynamics of deformable branches as mass-spring abstractions on a parallel non-differentiable simulator, obtain reference trajectories by actively probing real trees, and infer the spring parameters via a probabilistic inverse inference algorithm.
- 2) We incorporate biologically motivated assumptions into the Bayesian inference workflow by formulating inequality constraints as differentiable and learnt joint prior distributions.
- 3) We demonstrate the real-to-sim transfer with two different real-world robots, operating on distinct physical trees, thereby additionally validating model independence from both arm kinematics and branch topology.
- 4) Finally, we show that the learnt model and the predicted deformation trajectories are robust to noise perturbations from the sensors and variations in branch grasp locations where the manipulation forces are applied.

II. RELATED WORK

Manipulation of deformable objects, such as clothes, elastic bands, and ropes has gained much traction over the years, owing to the broad range of ensuing applications. The comprehensive surveys [9] and [10] detail established techniques to represent shapes, manage the dynamics, and apply learning

strategies, all in order to perform control and planning with deformables. While mass-spring models have been used to build coarse-grained, but computationally efficient models for suturing [11] and fabric management [12], its more accurate, but slower counterpart, the finite element method (FEM) models have been used for object tracking [13] and manipulating rubber tubes [14]. Perhaps more in line with our own work, both mass-spring and FEM models are used in conjunction, the latter as a reference, to acquire deformation trajectories for motion planning around soft toys [15].

In contrast, dynamic behaviour modelling of tree branches for active manipulation is rather rare in literature. Spring based models have been used to aid visual reconstruction of trees [16] or to animate the interaction between rain drops and branches [17]; however, the spring parameters are specified in advance. On the other hand, applying spring models to estimate branch behaviour under wind loading has been extensively studied in forestry [18]; for example, to measure the overall sway of tree [19]. Such works depend on accurate modelling of the wind forces, can only compute parameters along the wind direction, and therefore they are unsuitable for active robotic manipulation. To the best of our knowledge, no works exist in the literature that actively manipulate tree branches to estimate the spring parameters through simulation.

On the inference front, probabilistic methods have been proposed to estimate parameters to bridge the gap between simulation and reality, the real-to-sim problem, for both rigid bodies [6] and deformables [8]. While these approaches estimate parameter posteriors through the Bayesian framework, they almost always start with an uninformative prior, a uniform prior to enforce limits, or other well known distributions to represent the subjective belief about the parameters. Weighing the inference technique alone, our approach appears closest to [6], which introduces a constrained version of the SVGD algorithm inspired by the Lagrange multiplier optimization strategy. In comparison, we treat the simulator as a black-box, do not rely on accurate gradients for the simulation rollouts, and our inference formulation can embody inequality assumptions as prior beliefs.

III. APPROACH

A. Mass-Spring-Damper Model

Our work utilises the well established Mass-Spring-Damper System (MSS) abstraction to simplify the intricate tree geometry. MSS models are computationally efficient and simple to implement; however, the spring constants are hard to obtain from material properties without reference models. In this work, we describe the branch structures as cylindrical links coupled with a torsional spring-damper system (Fig. 2a).

The dynamics of a single branch, and consequently that of the entire branch system, can be derived from Newton's second law of motion. Given an external torque T_{ext} , branch stiffness P , damping factor D , moment of inertia J , and angular displacement ψ , this relationship can be obtained as:

$$T_{ext} = J\ddot{\psi} + D\dot{\psi} + P\psi \quad (1)$$

We propose to model the relation $f : T_{ext} \mapsto (\psi, \dot{\psi})$, by comparing the model displacement trajectories to reference




Fig. 2: (a) Mass-Spring-Damper representation of a branch system (b) Parallel instances of our coarse-grain tree geometry simulations, executed on NVIDIA Isaac Gym. (c) Physical manipulator setups. Left: Kinova Jaco-2 arm on a synthetic plant. Right: Franka-Emika Panda arm on a farm tree.

trajectories from the real world. We use physics simulators to model this MSS abstraction, which additionally helps to account for the latent attributes that affect the branch motion, such as gravity and the branch fork angle.

B. Tree Simulation

Generating a digital twin for a physical entity enables cost effective data generation and robust control policy learning; moreover, GPU-accelerated simulators provide the advantage of faster training and efficient run time inference. This section describes our approach to generating crude digital tree replicas that run in parallel to perform simulator driven learning.

To begin with, we purpose-built a tool for generating tree geometry. The geometric structures can be generated from scratch, or alternatively, it can be integrated with a perception system to create replicas of real world trees. The focus of this paper is the former, where the geometric attributes are specified in advance. The resulting coarse tree model (Fig. 2b) takes cues from the natural world to replicate traits of real trees, such as having a fractal pattern and adhering to the cross sectional area-preservation principle [20], also known as the Leonardo da Vinci’s rule. Furthermore, our approach is not bound by branch topology, and it can accommodate larger plant models built from Lindenmayer systems (L-system).

Subsequently, the MSS model is implemented as a multi-link structure with branches approximated as cylindrical links and branch forks represented by actuated revolute joints. The branch motion is simulated via a proportional–derivative (PD) controller governed by tunable stiffness and damping gain parameters, $\theta = \{K_p, K_d\}$. Consequently, the branch motion estimated by the PD controller is:

$$T_{ctrl} = K_p\psi + K_d\dot{\psi}, \quad (2)$$

where T_{ctrl} is the controlling torque exerted by the simulated branch to offset the positional ψ and velocity $\dot{\psi}$ errors. The final structure is represented with a simulator agnostic unified robotics description format (URDF). While the simplified mass-spring model is an imprecise approximation of the ground truth, particularly when compared to other more realistic finite element method (FEM) models, we show that it is sufficient to capture the multi-branch dynamics.

Finally, we choose NVIDIA Isaac Gym [21] to simulate tree structures; specifically, its joint space PD controller to apply deformation torques and measure the resulting motion characteristics. Isaac Gym supports tensor based distributed

environments, and it permits parallel execution of the inference without CPU bottlenecks. However, Isaac Gym simulations are non-differentiable; therefore, unlike [6], throughout this paper we use finite difference methods to approximate gradients.

C. Probabilistic Parameter Inference

The key idea of our approach is to model the relationship between the force required to counteract the simulated branch resistance T_{ctrl} and the resulting branch deformation trajectory, through inverse inference of stiffness K_p and damping K_d gain parameters. The moment of inertia is computed by the simulator directly from the branch topology, and therefore is known a priori. The cost of deformation is established by comparing the position and velocity trajectories between the ground truth and the simulation through a loss function. The ground truth trajectories can be generated in two ways; first, as yet another independent articulated body simulation as described in section IV-A, and second using real robotic arms on physical trees, as described in section IV-B.

We define force profile F as a sequence of force vectors applied on the target branch, at a location C , over g time steps, i.e., $F := \{f_0, f_1, \dots, f_g \mid f \in \mathbb{R}^3\}$. The resulting deformation at each time step t is defined as the change in the position of the branch, at the force location, w.r.t its initial position, i.e., the deviation of the point C_t from its rest position C_0 . The position trajectory \mathcal{T} is then the sequence of such deformations, over g time steps, while its derivative w.r.t time gives $\dot{\mathcal{T}}$, the velocity trajectory. \mathcal{T}_{gt} and \mathcal{T}_{sim} represent the branch displacement for ground truth and simulation respectively. However, the force profile is constrained to the gravity axis (downward) as a simplification measure during experiments.

A general parameter inference problem is characterised by three terms: a parameterised simulator, a loss, and an inference framework. First, we define the simulator as a function S that steps forward in time, given simulation parameters θ and force profile, to roll out a trajectory \mathcal{T}_{sim} over time.

$$S : (\theta, s_i, f_i) \mapsto s_{i+1}, \quad \theta \in \mathbb{R}^N \\ \therefore \mathcal{T}_{sim}^\theta := \{s_0^\theta, s_1^\theta, \dots, s_g^\theta\},$$

where N is the simulation parameter count. For each branch, we take $N = 2$ to represent its PD coefficients, K_p and K_d . The relationship to be modelled then is $f : F \mapsto (\mathcal{T}, \dot{\mathcal{T}})$. Second, we define the loss L over simulation parameters as deviations of the branch position and velocity trajectories from

a ground truth.

$$L(\theta) := \|\mathcal{T}_{sim}^\theta - \mathcal{T}_{gt}\|_2 + \|\dot{\mathcal{T}}_{sim}^\theta - \dot{\mathcal{T}}_{gt}\|_2.$$

Finally, similar to [6] and [8], we use Bayesian inference to quantify the uncertainty associated with the simulation parameters. The canonical Bayes' rule relates the parameter posterior probability density $p(\theta|\mathcal{T}_{sim})$ to the corresponding likelihood $p(\mathcal{T}_{sim}|\theta)$ and the prior $p_0(\theta)$,

$$p(\theta|\mathcal{T}_{sim}) \propto p(\mathcal{T}_{sim}|\theta) p_0(\theta). \quad (3)$$

While point estimates of optimal parameters can be obtained as $\theta^* = \operatorname{argmin}_{\theta \in \mathbb{R}^N} L(\theta)$, i.e., by minimising the objective function, we instead attempt to successively sample promising regions of the posterior density $p(\theta|\mathcal{T}_{sim})$ in order to capture its multiple modalities. Furthermore, we compute the likelihood term through an exponential transformation of the loss, along the lines of simulated annealing architectures [22],

$$p(\mathcal{T}_{sim}|\theta) := \frac{e^{-L(\theta)/kT}}{Z_\theta} \quad \because L(\theta) \geq 0, \quad (4)$$

where Z_θ is the constant set of all possible energy states, k the Boltzmann constant, and T the fixed annealing temperature,

$$\therefore p(\theta|\mathcal{T}_{sim}) = \frac{e^{-L(\theta)/kT} p_0(\theta)}{Z_\theta Z_\tau}, \quad (5)$$

where Z_τ is the normalisation constant. We do away with the intractable $Z_\theta Z_\tau$ term by choosing Stein Variational Gradient Descent algorithm (SVGD) as our inference algorithm.

While our approach can be generalized to a full 3D system, in this paper, we only consider deformations in the vertical plane, implemented by restricting the force profile to the gravity-axis, i.e., perpendicular to the ground. Deformation along other axes, such as horizontal twists can be estimated by using higher-dimensional hinges at each joint in the sequence, i.e., by increasing the θ dimensions. However, we restrict the forces to tolerable limits and leave the discontinuous dynamics of branch rupture for future work. Finally, we treat each \mathcal{T}_{sim} trajectory as an i.i.d sample from the underlying distribution, compute the likelihood term (4) over multiple trajectories, and use the log form to avoid vanishing floating points,

$$\log p(\mathcal{T}_{sim}|\theta) := \sum_\tau \frac{-L(\theta)}{kT} - \log Z_\theta. \quad (6)$$

D. Stein Variational Gradient Descent

This section provides an overview of the SVGD [23] and its utility in estimating branch simulation parameters. SVGD belongs to a family of Variational Inference methods that takes an optimisation approach to simulate intractable probability distributions. SVGD has been shown to converge faster than alternatives such as Markov Chain Monte Carlo (MCMC) predominantly due to its deterministic and iterative progression, akin to the traditional Gradient Descent optimisation.

The algorithm starts with n randomly initialised points in \mathbb{R}^N , termed as particles, representing draws from the target density. At each step of the iteration, the particles are pushed closer by a small step size $\epsilon \in \mathbb{R}$ as below:

$$\theta_i^{t+1} \leftarrow \theta_i^t + \epsilon \phi^*(\theta_i^t) \quad \forall i = 1, 2, \dots, n, \quad (7)$$

where θ_i^t represents the i^{th} particle at time step t and ϕ^* represents the optimal function that minimises the Kullback–Leibler (KL) divergence between the distribution represented by the current particle state and the expected distribution, $p(\theta|\mathcal{T}_{sim})$ in our case. However, in contrast to SGD, this minimisation is performed in a space of functions, named the Reproducing Kernel Hilbert Space (RKHS), uniquely identified by a kernel function $k : \mathbb{R}^N \times \mathbb{R}^N \mapsto \mathbb{R}$. We use the popular Radial Basis Function (RBF) kernel: $k(\theta, \theta') = \exp(-\frac{\|\theta - \theta'\|^2}{2\sigma^2})$, where the bandwidth σ is computed using the median heuristic. Using the Stein approximation from [23] to compute the steepest direction of descent in RKHS, i.e, the optimal function ϕ^* :

$$\hat{\phi}^*(\theta) = \frac{1}{n} \sum_{j=1}^n k(\theta_j, \theta) \nabla_{\theta_j} \log p(\theta_j|\mathcal{T}_{sim}) + \nabla_{\theta_j} k(\theta_j, \theta).$$

While SVGD allows for placing the particles on a GPU and performing the operations independently, the traditional performance bottleneck has been in executing the simulations to derive the gradients. In our case, the term $\nabla_{\theta_j} \log p(\theta_j|\mathcal{T}_{sim})$ is approximated via finite differences instead. Our simulator choice, Isaac Gym, is non-differentiable but supports efficient gradient approximations through GPU based simulations; therefore, the entire inference process runs in parallel. On substituting the prior and likelihood terms from (5), we obtain:

$$\hat{\phi}^*(\theta) = \frac{1}{n} \sum_{j=1}^n -k(\theta_j, \theta) \nabla_{\theta_j} \frac{L(\theta)}{kT} + k(\theta_j, \theta) \nabla_{\theta_j} \log p_0(\theta_j) + \nabla_{\theta_j} k(\theta_j, \theta). \quad (8)$$

Notice that the intractable constant terms from Boltzmann distribution and from the Bayes' rule, Z_θ and Z_τ vanish due to the gradient, which is a quintessential feature of SVGD.

E. Joint Parameter Inference

A multi-branch joint parameter estimation problem is an extension of (3) to R coupled branches that actively take part in the collective dynamics.

$$p(\theta_1, \dots, \theta_R|\mathcal{T}_{sim}) \propto p(\mathcal{T}_{sim}|\theta_1, \dots, \theta_R) p_0(\theta_1, \dots, \theta_R) \quad (9)$$

This includes the child branch which is actively probed and all its main parent branches up to the fork where any deformation is observed due to the probing. While we limit our experiments with $R = 3$, the model itself has no such limitation, except for the marginal increase in computational cost per increment due to the increased dimensionality of the search space.

F. Smooth Box Prior

We use the prior $p_0(\theta)$ from (5) to define a search range for the inference algorithm in order to ensure convergence within a reasonable time. Using a uniform prior results in non-differentiable regions around the prior limits. For that reason, we use the Smoothed Box prior from GPyTorch library [24] to generate a smoothed version of the uniform prior to ensure differentiability at all points. Given a box defined by

the lower (θ_l) and upper (θ_u) bound of simulation parameters $B := \{\theta \mid \theta \in [\theta_l, \theta_u], \theta \in \mathbb{R}^N\}$, and a variance σ^2 , the fully differentiable smoothed box prior is approximated as:

$$p_0(\theta) \sim \exp\left(-\frac{d(\theta, B)^2}{\sqrt{2\sigma^2}}\right)$$

where the distance d is defined as $d(\theta, B) = \min_{\theta' \in B} |\theta - \theta'|$.

G. Neural Network based Inequality Prior

Further to constraining the search range with a smooth box prior, we propose to incorporate biological assumptions on the relationship between the branch parameters into the prior $p_0(\theta)$ when more than one branch is involved.

For the simplest multi-branch case, where $R = 2$, i.e., one parent child pair each with a single parameter $\theta_1 \in \mathbb{R}$ and $\theta_2 \in \mathbb{R}$ respectively, an inequality assumption that defines the joint parameter relationship could be:

$$y(\theta_1, \theta_2) := \begin{cases} 0 & \theta_1 \geq \theta_2 \\ -\eta & \theta_1 < \theta_2 \end{cases}, \quad (10)$$

where η is a scaling factor and $\eta \gg 0$. We propose to train a simple one-hidden layer, feed forward, neural network regressor to model y , and then use the predicted y' as an approximation for log of prior, i.e., $\log p_0(\theta_1, \theta_2) \approx y(\theta_1, \theta_2)$.

$$\therefore p_0(\theta_1, \theta_2) \approx \begin{cases} 1 & \theta_1 \geq \theta_2 \\ e^{-\eta} & \theta_1 < \theta_2 \end{cases}, \quad (11)$$

implying that the child branch parameter θ_2 (K_p or K_d) is assumed to be lower than that of the parent θ_1 . Intuitively, from the area-preservation principle [20], the child branch cross-sectional area is less than that of the parent, while elasticity theory dictates that the bending stiffness is proportional to the area for a cylinder. The neural network is modelled as:

$$y'(\theta_1, \theta_2) = w_0 + \sum_k w_k \sigma(w_{k0} + w_{k1}\theta_1 + w_{k2}\theta_2), \quad (12)$$

where σ is any activation function, k the number of hidden units, and w the network weights.

We justify our design choices with the following arguments: First, the regressor ensures that the predicted y' is continuous and fully differentiable at all points in the domain, subject to the network activation function choice. Second, equation (11) guarantees that $p_0(\theta)$ is non-negative. Third, the use of SVGD as the inference algorithm renders irrelevant the need to normalise $p_0(\theta)$. Therefore, y' can indeed be considered as an un-normalised probability density that satisfies the differentiability requirements of the term $\nabla_{\theta_j} \log p_0(\theta_j)$ from (8). Fourth, training the model on a log scale prevents prior values from getting extremely small to machine limits, keeps it comparable to the likelihood, and therefore adequately constrains the SVGD convergence to the region of interest. Finally, using a trained model to predict joint prior opens up the possibility of learning directly from the tree geometry instead of hand coding an assumption *a priori*. Most importantly, we assert that this approach is consistent with the Bayesian paradigm, where assumed knowledge is expressed with priors,

unlike other methods that penalise constraint violations at the likelihood - CSVGD [6], for example.

However, parameter estimation with (10) and (12) generates non-trivial flat regions in the target density, thereby slowing down the inference process. A better alternative is to penalise the difference in parameters as per the prior belief.

$$y(\theta_1, \theta_2) := \begin{cases} 0 & \theta_1 \geq \theta_2 \\ -\eta(\theta_2 - \theta_1) & \theta_1 < \theta_2 \end{cases} \quad (13)$$

$$\therefore p_0(\theta_1, \theta_2) \approx \begin{cases} 1 & \theta_1 \geq \theta_2 \\ e^{-\eta(\theta_2 - \theta_1)} & \theta_1 < \theta_2 \end{cases} \quad (14)$$

where η is a scaling factor that determines the sharpness of the gradient and $\eta > 0$. Note that we do not make any assumptions about the relationship between the parameters when the relationship is valid, and the penalty applies only when the constraint is violated. For the simple inequality case we implemented using (14), the reduction in search space is at least exponential in the number of priors, indicating that the benefits of adding the prior grows exponentially for every increment in R .

IV. EXPERIMENTAL SETUP

The following subsections describe our experimental strategies and the evaluation metrics used.

A. Sim-to-sim

In the most basic set up, we capture both ground truth and the replica measurements from the Isaac Gym simulator. While the geometric representation of the branch structure is same in both (as described in III-B), the execution process remains independent, i.e., no information is leaked between ground truth and inference workflows except for the force-deformation profiles. As opposed to the hardware experiments, in the sim-to-sim case, we use the simulator API to apply the deformation forces rather than using simulated arms.

B. Real-to-sim

This section describes the scenario where the ground truth is captured from real world using a physical robot arm. To show that our framework is agnostic to arm kinematics, we use two independent hardware platforms (Fig. 2c) for the measurements: 1) a Kinova Jaco-2 6 DoF curved wrist manipulator, and 2) a Franka-Emika Panda 7 DoF manipulator. In the first setup, we use the Kinova manipulator to probe and deform an artificial olive tree. In contrast, we operate the mobile base mounted Franka arm on a larger, real farm tree.

A few general principles applicable to both executions are examined below. The end-effector grasp position is taken as synonymous to the branch force location C ; therefore, the Cartesian trajectory of the end-effector acts as a proxy for the branch deformation trajectory \mathcal{T} , at the grasp location. We command the arm using position control to deform the branch downward, in the gravity axis, while ignoring force and deformation in all other axes. To infer the joint parameters of coupled branches, we only measure the trajectories of the child

branch to which the arm is attached, obviating the need for any visual tracking. We utilize the end-effector force readings exposed by the arm, which are typically inferred using forward kinematics, instead of the more accurate external sensor attachments at the end-effector itself, and consequently the readings are rather noisy. Additionally, the robot model could have position-sensitive bias in force readings, due to gravity compensation issues, sensor misalignment or improper arm calibration. However, we show that our framework is robust to these distortions. Both force and deformation measurements are independently retrieved from the arms using their respective Robot Operating System (ROS) APIs. Additionally, we neither use a perception mechanism to capture the tree characteristics nor do we attempt automated grasping.

Finally, we measure the branch attributes essential to create the digital replica, which are: the radius, length and density of the branch; and the angle at which the branch forks from its parent. While it is not necessary to simulate the entire tree, it is imperative that all upper level branches that participate in the collective dynamics of the target branch are accounted for. The density of wood material is assumed to be constant for a tree species (or the synthetic material), which can be derived indirectly from the mass and volumetric measurements of a single cut out fragment.

C. Evaluation Metrics

To assess the estimation reliability, first we compare the inferred parameter distribution with the pre-specified, parameter values from the ground truth, using the sim-to-sim setup. In both sim-to-sim and real-to-sim experiments, we compare the unseen ground truth (GT) deformation trajectory to the predicted samples drawn from the estimated trajectory distribution $p(\mathcal{T}_{sim}|\theta^*)$ corresponding to the converged particle set θ^* . Additionally, in the latter case, we compute the 95% confidence interval (CI) over the estimated trajectories and assess the fraction of true trajectory points that fall outside the interval. The distribution itself is constructed from the samples of trajectories corresponding to θ^* , leveraging Kernel Density Estimation with Gaussian kernels and Scott's rule of thumb for bandwidth selection. In the sim-to-sim experiments, we use a hold-out dataset for evaluation, while in real-to-sim, we use k-fold cross validation due to the smaller sample size. In either case, the evaluation trajectories are unseen during the parameter inference workflow. Samples of GT force-deformation measurements, along with all experiment details, such as the dataset volume and the hyper-parameters used are available in the supplementary material.

V. EXPERIMENTS AND RESULTS

We assess the quality of our approach using the following experimental scenarios.

A. Baselines

In the first scenario, illustrated in Fig. 3, we baseline our inference approach, SVGD with a neural network prior

(NNSVGD), against four alternatives, which are: a) the classical Metropolis–Hastings (MCMH) algorithm b) Stochastic gradient Hamiltonian Monte Carlo (SGHMC) method c) Stochastic Gradient Langevin dynamics (SGLD) and a basic d) Stein Variational Gradient Descent (SVGD). In each of the four baseline algorithms, we use the smooth box prior described in section III-F to set the search limits, and wherever a gradient is required, we use the finite difference based approximation. This scenario is executed using the sim-to-sim setup with a multi-branch simulated tree geometry ($R = 3$), imitating a joint grandparent-parent-child motion. We keep all hyper-parameters, such as the number of iterations and the particle count, similar across the algorithms. In the gradient based methods, we use an Adam optimizer with a small learning rate to guide the descent.

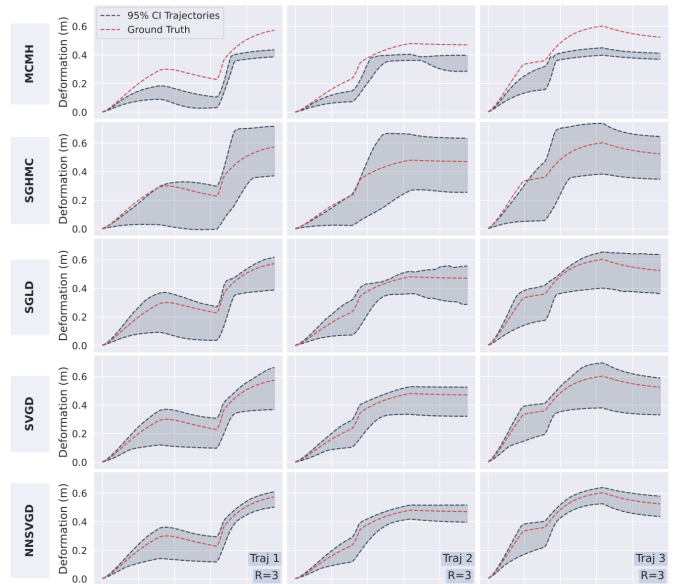


Fig. 3: Comparison of NNSVGD (ours) against baseline inference algorithms. 95% CI of the predicted distribution of test trajectories $p(\mathcal{T}_{sim}|\theta^*)$, computed with $R = 3$, is shown in grey against the preset ground truth \mathcal{T}_{gt} in red

The results in Fig. 3 indicate all gradient based methods outperform the random-walk based MCMH. Furthermore, samples from the Stein methods, i.e., SVGD and NNSVGD, are better distributed around the ground truth \mathcal{T}_{gt} trajectory, while additionally displaying sharper convergence and lower prediction uncertainty, indicated by the CI width. For identical iteration counts, Stein methods outperform chain based SGHMC and SGLD chiefly because the particles are updated in parallel, a feature exploited by our parallel simulator.

B. NNSVGD vs SVGD

With the same setup, the benefit of having a neural network prior, described in section III-G, is examined both visually and quantitatively. First, the intuitive visual result in Fig. 4, shows that the search space restriction from the inequality prior forces the spring parameters to have a sharper convergence, for all three branches. Consequently, the sample trajectories

predicted by NNSVGD show lower CI widths in comparison to SVGD, from Fig. 3. Furthermore, the multi-modal nature of the branch1 posterior indicates the presence of multiple solutions, a quintessential feature of inverse inference.




Fig. 4: Joint posteriors with and without the neural network prior, computed with $R = 3$. The red dots indicate the preset ground truth for stiffness K_p and damping K_d parameters.

In the first quantitative test, we compute the negative log-likelihood (Nll) of the estimated posterior from 5 random training trajectories. Secondly, we compute the 95% CI of the predicted trajectories averaged over 25 unseen test trajectories and assess the fraction of GT points that fall outside the interval, as detailed in section IV-C.

	95% CI width	% GT outside CI	Nll
NNSVGD	0.11	<1%	2268.1
SVGD	0.19	<1%	3050.7

TABLE I. NN prior benefits: Quantitative metrics.

The results, in Table I demonstrate that the addition of the neural network prior results in a lower Nll and a lower CI width for the same %GT enclosure, suggesting its superiority.

C. Robustness

In the second scenario, we perform two sets of experiments to measure the robustness of the model against erratic sensors and variations in branch grasp locations. In both sets, the metrics are computed as an average over 25 unseen test force-deformation trajectories, 3 of which are shown along with the results. In the first set, we inject Gaussian noise of standard deviation σ into the ground truth deformation and force trajectories, leveraging the sim-to-sim setup. The results in Fig. 5, show that for moderate levels of noise, the predicted distribution is well-aligned to the ground truth; furthermore, the 95% CI attributes indicate that the model is able to capture the injected noise within the estimated uncertainty.

Secondly, we vary the grasp locations, selected randomly along the branch, where the force is applied. This experiment mimics the situation where the arm needs to ascertain whether to grasp towards the outer edge of a branch to reduce the applied forces, or move towards the inner fork due to the arm reach constraints. The obtained results, illustrated in Fig. 6, confirm that our model can successfully infer the deformation on one location using trajectories obtained by grasping another.




Fig. 5: 95% CI of predicted distribution against GT for varying noise σ , measured relative to the force, keeping the grasp location fixed. The bar chart shows the CI width along with %GT points outside the interval (lesser the better for both).




Fig. 6: 95% CI of predicted distribution against GT for varying grasp locations, keeping noise σ fixed. The grasp location is represented as a fraction of the branch length, from the parent joint; for e.g., Loc: $0.75l$ is the mid point of the outer half.

D. Hardware Experiments

Here, we perform real-world experiments and capture multiple pairs of force-deformation trajectories using the two setups described in section IV-B. The converged NNSVGD particles and the corresponding deformation trajectory samples are used to construct a trajectory distribution $p(\mathcal{T}_{sim}|\theta^*)$ and the 95% CI for each time-step, shown in Fig. 7. The CI attributes are tabulated in Table. II.

	Kinova ($R = 1$)	Franka ($R = 1$)	Kinova ($R = 2$)
Avg 95% CI width	0.052	0.066	0.039
Avg %GT outside CI	5.0%	7.0%	23.9%

TABLE II. Real-to-sim CI attributes

The results show that despite the heteroschedastic noise in the ground truth, our model is able to predict the deformation quite well. For instance, across the three farm tree trajectories, on average, 93.0% of true points lie within the 95% confidence interval and the average interval width is 6.6cm. We deem such metrics as acceptable for a majority of practical applications discussed earlier. Remarkably, these high quality real-world predictions are obtained by using as little as three training trajectories. However, from the real-world branch behaviour, we observed that our approach works the best for deformation in regions that are closer to the horizontal ($+45^\circ$ to -45° from the ground axis). Given the force is applied along the gravity axis, this corresponds to a region where its normal component significantly outweighs the shear. Therefore, the proposed method could be improved by using the branch angles captured through perception to apply normal forces to the branch instead, or alternatively, by representing a single real branch with multiple links to increase pliability, both of which we leave aside for future consideration.

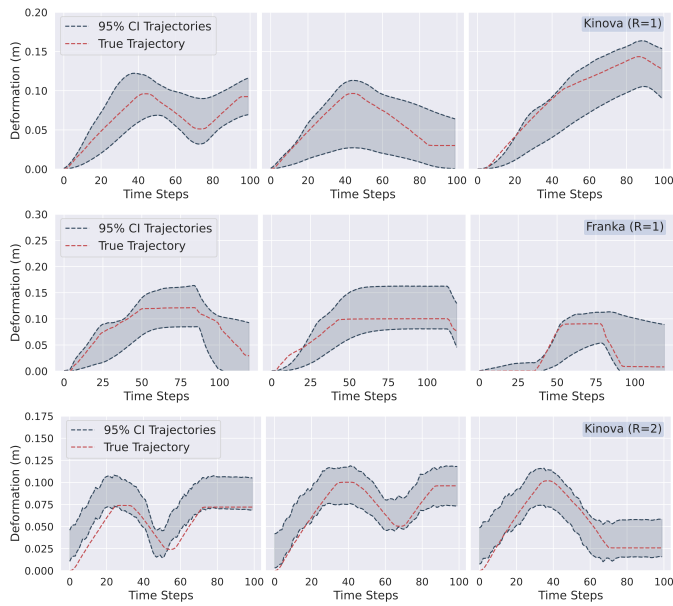


Fig. 7: Real-to-sim results. Top: Kinova ($R = 1$), Middle: Franka ($R = 1$), Bottom: Kinova ($R = 2$). Shaded regions represent the 95% CI of the predicted distribution $p(\mathcal{T}_{sim}|\theta^*)$.

VI. CONCLUSION

In this work, we have demonstrated a method to learn the dynamic behaviour of tree branches by utilising the simulation-based parameter inference approach. In addition, we described a method for injecting biological assumptions into the Bayesian model via a neural network prior. Previous works on tree dynamics are either focused on perception or on the biological aspects of the tree; therefore, they are not designed for active branch manipulation, a key requirement for robotic harvesting and locomotion tasks. The comprehensive results indicate a robust model that captures the parameter uncertainty quite well even under noise and grasp location perturbations. Future work in this direction are as follows: First, the spring model can be alternatively built from a perception acquired branch topology [1]. Second, a generalised version of the model could be used to learn branch manipulation policies. Third, severe discrepancies between model estimates and true deformation behaviour could be used as an indicator of poor tree health or branch rupture. Overall, we foresee works that entail combining our uncertainty aware model with visual tree reconstruction works, generalising it to different branches, and eventually generating robust control policies for manipulation.

ACKNOWLEDGMENT

This work was (partially) funded by CSIRO’s Reimagine Farming project part of CSIRO’s Reinvent Science program. The authors are grateful for technical and systems support from Tea Molnar, Brendan Tidd, Lauren Hanson and Peyman Moghadam on CSIRO’s mobile manipulation hardware and software integration for field testing.

REFERENCES

1 T. Lowe and J. Pinskiel, “Tree reconstruction using topology optimisation,” *Remote Sensing*, vol. 15, no. 1, p. 172, 2022.

2 E. Quigley, W. Lin, Y. Zhu, and R. Fedkiw, “Three dimensional reconstruction of botanical trees with simulatable geometry,” *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, vol. 4, no. 3, pp. 1–16, 2021.

3 H. Zhou, X. Wang, W. Au, H. Kang, and C. Chen, “Intelligent robots for fruit harvesting: Recent developments and future challenges,” *Precision Agriculture*, vol. 23, no. 5, pp. 1856–1907, 2022.

4 L. van Herck, P. Kurtser, L. Wittemans, and Y. Edan, “Crop design for improved robotic harvesting: A case study of sweet pepper harvesting,” *Biosystems Engineering*, vol. 192, pp. 294–308, 2020.

5 K. Cranmer, J. Brehmer, and G. Louppe, “The frontier of simulation-based inference,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 48, pp. 30055–30062, 2020.

6 E. Heiden, C. E. Denniston, D. Millard, F. Ramos, and G. S. Sukhatme, “Probabilistic inference of simulation parameters via parallel differentiable simulation,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 3638–3645.

7 Y. Chebotar, A. Handa, V. Makovychuk, M. Macklin, J. Issac, N. Ratliff, and D. Fox, “Closing the sim-to-real loop: Adapting simulation randomization with real world experience,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8973–8979.

8 R. Antonova, J. Yang, P. Sundaresan, D. Fox, F. Ramos, and J. Bohg, “A bayesian treatment of real-to-sim for deformable object manipulation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 5819–5826, 2022.

9 H. Yin, A. Varava, and D. Kragic, “Modeling, learning, perception, and control methods for deformable object manipulation,” *Science Robotics*, vol. 6, no. 54, p. eabd8803, 2021.

10 V. E. Arriola-Rios, P. Guler, F. Ficuciello, D. Kragic, B. Siciliano, and J. L. Wyatt, “Modeling of deformable objects for robotic manipulation: A tutorial and review,” *Frontiers in Robotics and AI*, p. 82, 2020.

11 J. Schulman, A. Gupta, S. Venkatesan, M. Tayson-Frederick, and P. Abbeel, “A case study of trajectory transfer through non-rigid registration for a simplified suturing scenario,” in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 4111–4117.

12 S. Makris, E. Kampourakis, and D. Andronas, “On deformable object handling: Model-based motion planning for human-robot co-manipulation,” *CIRP Annals*, 2022.

13 A. Bozic, M. Zollhofer, C. Theobalt, and M. Nießner, “Deepdeform: Learning non-rigid rgb-d reconstruction with semi-supervised data,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 7002–7012.

14 M. Rambow, T. Schaub, M. Buss, and S. Hirche, “Autonomous manipulation of deformable objects based on teleoperated demonstrations,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 2809–2814.

15 B. Frank, C. Stachniss, R. Schmedding, M. Teschner, and W. Burgard, “Learning object deformation models for robot motion planning,” *Robotics and Autonomous Systems*, vol. 62, no. 8, pp. 1153–1174, 2014.

16 F. Yandun, A. Silwal, and G. Kantor, “Visual 3d reconstruction and dynamic simulation of fruit trees for robotic manipulation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 54–55.

17 M. Yang, M.-C. Huang, G. Yang, and E.-H. Wu, “Physically-based animation for realistic interactions between tree branches and raindrops,” in *Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology*, 2010, pp. 83–86.

18 E. De Langre, “Effects of wind on plants,” *Annu. Rev. Fluid Mech.*, vol. 40, pp. 141–168, 2008.

19 K. R. James *et al.*, “A study of branch dynamics on an open-grown tree,” *Arbiculture & Urban Forestry*, vol. 40, no. 3, pp. 125–34, 2014.

20 R. Minamino and M. Tateno, “Tree branching: Leonardo da vinci’s rule versus biomechanical models,” *PLoS one*, vol. 9, no. 4, p. e93535, 2014.

21 V. Makovychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and G. State, “Isaac gym: High performance gpu-based physics simulation for robot learning,” 2021.

22 S. Kirkpatrick, C. D. Gelatt Jr, and M. P. Vecchi, “Optimization by simulated annealing,” *science*, vol. 220, no. 4598, pp. 671–680, 1983.

23 Q. Liu and D. Wang, “Stein variational gradient descent: A general purpose bayesian inference algorithm,” *Advances in neural information processing systems*, vol. 29, 2016.

24 J. Gardner, G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson, “Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration,” *Advances in neural information processing systems*, vol. 31, 2018.