

A Riemannian Approach for Spatiotemporal Analysis and Generation of 4D Tree-shaped Structures

Tahmina Khanam¹, Hamid Laga¹, Mohammed Bennamoun², Guanjin Wang¹, Ferdous Sohel¹, Farid Boussaid², Guan Wang³, and Anuj Srivastava⁴

¹ Murdoch University, WA, Australia 34719017@student.murdoch.edu.au, {H.Laga, Guanjin.Wang, F.Sohel}@murdoch.edu.au

² University of Western Australia, WA, Australia {mohammed.bennamoun, farid.boussaid}@uwa.edu.au

³ Yangtze Delta Region Institute (Huzhou), University of Electronic Science and Technology of China, China wangguan12621@gmail.com

⁴ Florida State University, USA anuj@stat.fsu.edu

Abstract. We propose the first comprehensive approach for modeling and analyzing the spatiotemporal shape variability in tree-like 4D objects, *i.e.*, 3D objects whose shapes bend, stretch and change in their branching structure over time as they deform, grow, and interact with their environment. Our key contribution is the representation of tree-like 3D shapes using Square Root Velocity Function Trees (SRVFT) [21]. By solving the spatial registration in the SRVFT space, which is equipped with an \mathbb{L}^2 metric, 4D tree-shaped structures become time-parameterized trajectories in this space. This reduces the problem of modeling and analyzing 4D tree-like shapes to that of modeling and analyzing elastic trajectories in the SRVFT space, where elasticity refers to time warping. In this paper, we propose a novel mathematical representation of the shape space of such trajectories, a Riemannian metric on that space, and computational tools for fast and accurate spatiotemporal registration and geodesics computation between 4D tree-shaped structures. Leveraging these building blocks, we develop a full framework for modelling the spatiotemporal variability using statistical models and generating novel 4D tree-like structures from a set of exemplars. We demonstrate and validate the proposed framework using real 4D plant data.

Keywords: 4D registration · Statistical analysis · 4D tree generation

1 Introduction

We propose a novel framework for the statistical analysis of the spatiotemporal shape variability in tree-like 3D objects that deform over time, hereinafter referred to as 4D (or 3D + time) tree-like structures. Such objects, which undergo complex non-rigid and topological deformations, are abundant in nature. Understanding and modelling the deformation patterns of their shapes can benefit

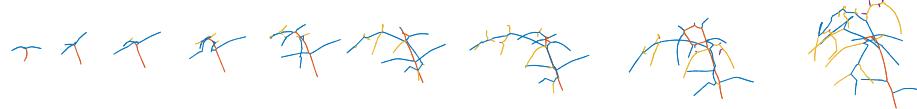


Fig. 1: The 4D skeleton of Fig. 1 in the Supplementary Material, structured it into layers of branches. The red branch is the main branch, blue corresponds to the 2nd layer of branches, yellow to the 3rd layer, and purple to the 4th layer.

many applications, from plant biology and medicine to 4D content generation in computer graphics and virtual reality. For example, plants and blood vessels deform over time due to either normal growth or the progression of diseases. Modeling the spatiotemporal variability can help distinguish shape deformations that are due to normal growth from those due to disease progression.

This problem has received growing attention. However, most of the work focused on the 4D reconstruction of dynamic objects [9, 10] and on the statistical analysis of 4D shapes that bend and stretch while preserving their topology [12]. In contrast, this paper focuses, for the first time, on 4D shapes that bend, stretch, and change in topology. Given a set of such 4D objects, our goal is to learn statistical summaries, such as the mean and modes of variation, characterize the spatiotemporal shape variability in the set using statistical models, and learn a generative model capable of synthesizing and generating novel 4D tree-shaped structures. Achieving these goals requires accurate spatiotemporal registration of the 4D tree-like shapes. Spatial registration refers to the problem of finding one-to-one branchwise and pointwise correspondences across 3D shapes. This is very challenging in our setup since tree-like 3D shapes such as plants and botanical trees exhibit not only large elastic deformations, *i.e.*, the bending and stretching of their branches, but also topological differences, *i.e.*, differences in the branching structures, within and across shape classes. Temporal registration refers to the problem of temporally aligning 4D shapes. For example, two botanical plants can grow at different rates even if they are of the same species. This adds another level of complexity, which is the high temporal variability due to differences in the execution rates (speeds) between 4D tree-shaped objects.

Our key contribution is the representation of tree-like 3D objects using Square Root Velocity Functions (SRVF) [17]. By solving the spatial registration problem in the SRVF space, tree-like 4D shapes become trajectories in this space. This reduces the problem of modeling and analyzing 4D tree-like shapes to that of modeling and analyzing elastic trajectories in the SRVF space, where elasticity refers to time warping. Building on [21], which developed tools for the spatial registration of tree-like 3D shapes in the extended SRVF space, we propose (1) a set of computational tools for the spatiotemporal registration of 4D tree-like shapes, (2) a mechanism for computing geodesics between 4D trees, (3) a framework for the statistical analysis of the spatiotemporal shape variability, and (4) a mechanism for learning, from a set of exemplars, generative models capable

of synthesizing novel 4D tree-like structures. We demonstrate the efficiency of the proposed tools using real 4D shapes of growing tomato and maize plants.

The paper is organized as follows. Sec. 2 summarizes the related work. Sec. 3 presents the proposed SRVF representation of tree-like 3D shapes. Sec. 4 presents the proposed computational tools for the spatiotemporal registration and geodesics computation between 4D tree-shaped structures. Sec. 5 demonstrates how these tools can be used to compute summary statistics and synthesize novel 4D tree-shapes. Sec. 6 presents the results while Sec. 7 concludes the paper.

2 Related Work

We refer to [11, 12, 21] for a detailed survey. Here, we focus on 4D shape analysis and related topics. Studying and modelling the spatio-temporal shape variability in 4D shapes requires defining a proper shape space and metrics on this space, performing spatiotemporal registration under the metric, and computing geodesics between 4D shapes. The bulk of the work is focused on the spatial registration problem, which is traditionally solved using hand-crafted feature matching and the Iterated Closest Point [1, 18, 23]. These methods are well suited for 3D objects that bend and stretch, *e.g.*, human/animal body shapes. Tree-like 3D objects such as plants and blood vessels have complex structures and exhibit topological variability, which makes their spatial registration challenging.

To address these issues, Feragen *et al.* [5–8] treat a tree-shaped object as a point in a tree shape space equipped with the Quotient Euclidean Distance (QED) as a metric. The problem of registering two tree-shaped objects is then reduced to that of finding the optimal deformation path, or geodesic with respect to the metric, that deforms one tree onto the other. These deformations include branch bending, stretching, and topological changes. Under the QED, bending and stretching are measured using the \mathbb{L}^2 metric while topological changes are measured using graph edit distance. This results in a significant shrinkage along the geodesic paths when the deformations are large. Wang *et al.* [20, 22] alleviated this issue by representing the shape of the edges in the tree shapes of Feragen *et al.* [5–8] using the Square Root Velocity Functions (SRVF) [17] and then using the QED in the SRVF space, which is equivalent to the full elastic metric in the original space. This still suffers from the shrinkage that is due to the edge collapses used by the QED metric. Duncan *et al.* [4] addressed these issues by representing trees using sliding branches, instead of collapsing edges. Guan *et al.* [21] extended this representation to tree-like 3D shapes such as complex botanical trees, plant roots, and neuronal structures. It has also been used to build a full generative model of tree-like 3D shapes by fitting Gaussian distributions to tree-like 3D shapes represented as points in the tree-shape space.

This paper builds upon and generalizes the representation of [21] to 4D tree-shaped objects, *i.e.*, sequences of tree-like 3D shapes that deform over time. The key idea is to represent 4D shapes as trajectories in the tree-shape space reducing the problem to that of analyzing elastic trajectories in the tree-shape space. This idea has been introduced by Laga *et al.* [12] for objects that have a manifold structure

and only bend and stretch, but do not change in topology. This paper focuses, for the first time, on 3D shapes that also change in topology. Several recent papers in plant biology focused on the registration of 4D plants [2, 3, 13–15, 19, 24]. They, however, define registration of 4D plants as the spatial registration of 3D plants within plant sequences. This is spatial registration. In this paper, we focus, for the first time, on spatial and temporal registration and develop a comprehensive suite of computational tools that enable the computation of geodesics, summary statistics, and generative models for 4D tree-shaped objects.

3 The Space of 3D tree-shaped structures

We first present the mathematical framework we use to represent tree-structured 3D shapes (Sec. 3.1), the elastic metric for quantifying their bending, stretching, and topological changes (Sec. 3.2), and the computational tools for their spatial registration and geodesics computation (Sec. 3.3). Throughout the paper, we use the Pheno4D dataset [16], which consists of multiple 4D plants. Each 4D plant is a sequence of 3D plants, in the form of point clouds, captured at different times during its growth. We extract the skeleton of the plants using [15] and partition it into branches; see Fig. 1 and also Fig. 1 in the Supplementary Material.

3.1 Representation

We represent a branch of a 3D tree structure T as an arc length-parameterized curve $\beta : [0, 1] \rightarrow \mathbb{R}^3$ where $\beta(s) = (x(s), y(s), z(s))$. We structure a 3D tree into layers; see Fig. 1. The first layer β^0 is the main branch with k^0 subtrees T_i^1 attached to it at the bifurcation points $s_i^0, i = 1, \dots, k^0$. That is, $T = (\beta^0, \{T_i^1, s_i^0\}_{i=1}^{k^0})$. A 3D subtree T^l at level l of the hierarchy can be represented recursively as $T^l = (\beta^l, \{T_i^{l+1}, s_i^l\}_{i=1}^{k^l})$, for $l = 1, \dots, L - 1$. Here, L is the total number of layers, β^l is the main branch of T^l , k^l is the number of subtrees attached to the branch β^l , and $s_i^l \in [0, 1]$ is a bifurcation point on β^l of T^l . The branches at different levels can have any arbitrary number of subtrees attached to them. A good representation needs to be invariant to similarity-preserving transformations (translation, scale, rotation, reparameterization). To achieve translation invariance, we translate each tree so that the start point of its main branch is located at the origin. Invariance to scale is optional as it may not be required for some applications such as growth analysis. The space of all such normalized trees is referred to as the *pre-tree-shape space* and is denoted by \mathcal{C}_T . The invariance to rotation and reparameterization will be dealt with at the metric level (Sec. 3.2). With this representation, a 3D tree becomes a point in the pre-tree shape space \mathcal{C}_T . A 4D tree H can then be seen as a time-parameterized curve $H : [0, 1] \rightarrow \mathcal{C}_T$ where $[0, 1]$ is the time domain.

3.2 The elastic metric for comparing tree-shaped structures

If two 3D tree structures have the same number of branches, then one can quantify the dissimilarity between the two trees by measuring the amount of bending,

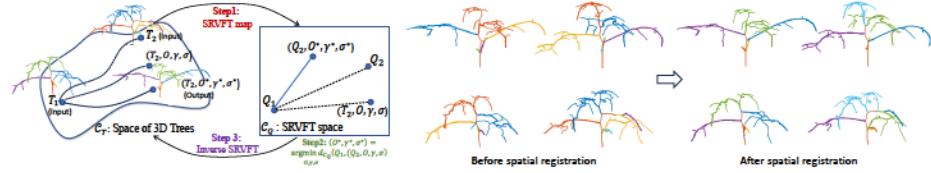


Fig. 2: The proposed spatial registration. We also show four 3D tree shapes before and after their spatial registration using the proposed framework. Our spatial registration took on average 235s to align two tomato and 17s to align two maize 4D plants.

stretching, and branch sliding that one needs to apply to the branches of one tree to align it onto the other. However, instead of explicitly using these metrics, which are non-linear and thus computationally expensive to evaluate and optimize, we use the SRVF representation [17] of 3D curves. Mathematically, the SRVF q of a curve β is defined as:

$$q(\beta)(s) = \frac{\beta'(s)}{\sqrt{\|\beta'(s)\|}}, \text{ if } \|\beta'(s)\| \neq 0, \text{ and } 0 \text{ otherwise.} \quad (1)$$

Its main property is that the \mathbb{L}^2 metric in the SRVF space is equivalent to a weighted sum of bending and stretching in the original space. It is also invertible, up to translation. In other words, given an SRVF, one can retrieve, up to translation, its corresponding original 3D curve. This significantly simplifies the analysis tasks: instead of measuring the similarity between two 3D curves with a complex elastic metric, one can map them to the SRVF space, perform the analysis there using the \mathbb{L}^2 metric and then map the results back to the original space. Let Q be the SRVF Tree (SRVFT) of an entire 3D tree T defined by computing the SRVF of each of its branches and appending their location with respect to their parent branch, and C_Q the pre-shape space of SRVFTs. A proper metric on that space needs to be invariant to the global rotation $O \in SO(3)$, reparameterization γ of the branches, and permutations σ of the orders of the lateral subtrees attached to a branch. We define γ and σ recursively, *i.e.*, $\gamma = (\gamma^0, \{\gamma^i\}_{i=1}^{k^0})$ is the reparameterization of the main branch and its subtrees. Specifically, $\gamma^0 \in \Gamma$ is a diffeomorphism that applies to the main branch of q and γ^i is the reparameterization, defined recursively, of the i -th subtree attached to the main branch. $\sigma = (\sigma^0, \{\sigma^i\}_{i=1}^{k^0})$ where $\sigma^0 \in \sigma$ is the permutation of the orders of the lateral subtrees on their corresponding main branch q^0 , and σ^i defines recursively these permutations for the i -th subtree. Following [4, 21], we define a rotation, reparameterization, and index permutation-invariant distance between two 3D trees, represented by their SRVFTs Q_1 and Q_2 , as the infimum over all possible rotations, branch reparameterizations, and branch index permutations:

$$d(Q_1, Q_2) = \inf_{O, \gamma, \sigma} d_{C_Q}(Q_1, (Q_2, O, \gamma, \sigma)), \text{ and} \quad (2)$$

$$d_{C_Q}(Q_1, (Q_2, O, \gamma, \sigma)) = \lambda_m \| q_1^0 - O(q_2^0, \gamma_0) \|^2 + \lambda_p \sum_{i=1}^n \left(s_1^i - s_2^{\sigma(i)} \right)^2 + \lambda_s \sum_{i=1}^n d(Q_1^i, Q_2^{\sigma(i)}). \quad (3)$$

Here, (Q, O, γ, σ) denotes the result of applying these transformations to Q .

3.3 Spatial registration and geodesics

Fig. 2 summarizes the spatial registration process. With this representation, the optimal rotation, diffeomorphism, and permutations that align Q_2 onto Q_1 can be found by solving the following optimization problem; see Section 2 in the Supplementary Material for the detailed algorithm:

$$(O^*, \gamma^*, \sigma^*) = \underset{O, \gamma, \sigma}{\operatorname{argmin}} d_{\mathcal{C}_Q}(Q_1, (Q_2, O, \gamma, \sigma)). \quad (4)$$

In practice, the input trees have different numbers of branches, thus they are elements of different subspaces. We address this by adding null branches at the different levels of the trees. This way, trees become elements of the same pre-tree shape space. The location of each additional null branch, *i.e.*, the value of its parameter s , is initialized to the s value of its initial corresponding branch. These will get updated during the branchwise correspondences.

Eq. (4) registers onto each other a pair of tree-like 3D shapes. To analyze 4D trees, we need to jointly register and align all the 3D instances within a 4D tree. We perform this sequentially: each 3D tree-shape within a sequence is aligned to the next 3D tree-shape in the sequence, using Eq. (4). With this formulation, computing a geodesic, or the optimal deformation that aligns Q_2 onto Q_1 is straightforward. Let $Q_2^* = (Q_2, O^*, \gamma^*, \sigma^*)$. Since the metric is a weighted norm of \mathbb{L}^2 distances, the geodesic α^* between Q_1 and Q_2 is the straight line that connects Q_1 to Q_2^* , *i.e.*, $\alpha^*(t) = (1-t)Q_1 + tQ_2^*, t \in [0, 1]$. For visualization, we map $\alpha^*(t)$ back to the non-linear space of 3D tree-shaped structures \mathcal{C}_T using the inverse SRVF mapping, which has a closed analytical form; see [17].

4 The proposed space of 4D trees-shaped structures

With this setup, a 4D tree-like shape becomes a time-parameterized 1D curve, or trajectory, $\alpha : [0, 1] \rightarrow \mathcal{C}_Q$. Thus, analyzing 4D tree-shaped structures becomes the problem of analyzing 1D curves in \mathcal{C}_Q . Here, we mathematically define the shape space \mathcal{C}_α of such curves and equip it with a proper metric (Sec. 4.1) that will allow us to temporally register and compare such trajectories (Sec. 4.2).

4.1 The shape space of 4D tree-shaped structures

Let $\alpha : [0, 1] \rightarrow \mathcal{C}_Q$ be a curve in \mathcal{C}_Q representing the SRVFT of a 4D tree-shaped object. Let \mathcal{C}_α be the space of all such curves. To temporally register, compare, and summarize samples of such curves, we need to define an appropriate metric on \mathcal{C}_α that is invariant to the execution rate of the curves. Mathematically, variations in the execution rate correspond to time-warping and thus can be represented using diffeomorphisms $\xi : [0, 1] \rightarrow [0, 1]$ that map the temporal domain to itself. Two curves α and $\alpha \circ \xi$ only differ in their execution rates and thus should be deemed equivalent. With this, the temporal registration of two 4D trees α_1 and α_2 becomes the problem of finding the optimal time

warping ξ^* that brings the two curves as close as possible to each other, with closeness measured using a metric $d(\cdot, \cdot)$: $\xi^* = \operatorname{argmin}_{\xi \in \Xi} d(\alpha_1, \alpha_2 \circ \xi)$, where Ξ is the space of all diffeomorphisms $\xi : [0, 1] \rightarrow [0, 1]$. We are left with the problem of defining the metric, or measure of closeness, $d(\cdot, \cdot)$. As time-warping of trajectories corresponds to the elasticity of curves, we follow the same approach used to compare branches (Sec. 3.2 and Eq. (1)), *i.e.*, instead of using a complex non-linear metric to measure the dissimilarity between two trajectories α_1 and α_2 , we first map them to their SRVF space, denoted by q_1 and q_2 , and perform the analysis there. Working in the SRVF space has many benefits. **(First)**, the elastic metric in the original space reduces to the \mathbb{L}^2 metric in the SRVF space. **Second**, under the \mathbb{L}^2 metric, the action of the diffeomorphism group Ξ is by isometries, *i.e.*, $\|q_1, q_2\| = \|q_1 \circ \xi, q_2 \circ \xi\|$. **Third**, the SRVF is invertible, up to translation. Thus, one can perform all the analysis tasks in the SRVF space, which has an \mathbb{L}^2 structure, and then map the results back to the original space for visualization without loss of information. Thus, temporal registration becomes:

$$\xi^* = \operatorname{argmin}_{\xi \in \Xi} \|q_1, q_2 \circ \xi\|^2 \text{ and } d(q_1, q_2) = \inf_{\xi \in \Xi} \|q_1, q_2 \circ \xi\|^2. \quad (5)$$

Here, $d(\cdot, \cdot)$ defines the rate-invariant distance between two trajectories (and thus two 4D shapes). Directly working with Eq. (5) to solve the temporal registration problem is computationally very expensive since α , which is a function, is of infinite dimension and when discretized, it will be of a very high dimension. Thus, we propose to learn a low dimensional subspace \mathcal{C}_{PCA} of 3D trees, similar to PCA but on \mathcal{C}_Q . Then, we can treat 4D tree-shapes as trajectories in \mathcal{C}_{PCA} , instead of the original space \mathcal{C}_Q , and perform their registration (Sec. 4.2) and statistical analysis (Sec. 5) using the elastic metric of Eq. (5).

The fact that \mathcal{C}_Q is equipped with **(1)** a proper metric that measures bending, stretching, and topological changes, and **(2)** a mechanism for computing geodesics (Sec. 3) allows us to learn a subspace, of finite dimension, that captures variability in collections of tree-shaped objects, similar to PCA in standard Euclidean spaces. Let $\{T_i\}_{i=1}^m$ be a set of tree-shaped 3D objects and $\{Q_i\}$ their corresponding SRVFTs. To compute the mean tree $\boldsymbol{\mu}$, we first compute the mean SRVFT $\boldsymbol{\mu}_q$ and then map it back to the space of tree shapes. Mathematically, $\boldsymbol{\mu}_q$ is the point in \mathcal{C}_Q that is as close as possible to all tree shapes in $\{Q_i\}_{i=1}^m$. The closeness is defined with respect to the metric of Eq. (2). In other words,

$$\boldsymbol{\mu}_q = \operatorname{argmin}_{(O_i, \gamma_i, \sigma_i)_{i=1}^m} \sum_{i=1}^m d_{\mathcal{C}_Q}^2(Q_i, (O_i, \gamma_i, \sigma_i)). \quad (6)$$

Solving Eq. (6) involves finding $\boldsymbol{\mu}_q$, known as the Karcher mean, while registering every Q_i onto $\boldsymbol{\mu}_q$. This can be efficiently done via a gradient descent approach:

1. Set $\boldsymbol{\mu}_q = Q_1$.
2. For $i = 1 : m$, Find $(\tilde{O}_i, \tilde{\gamma}_i, \tilde{\sigma}_i)$ that optimally register Q_i onto $\boldsymbol{\mu}_q$ (Eq. (4)).
3. Set $\boldsymbol{\mu}_q = \frac{1}{m} \sum_{i=1}^m (Q_i, \tilde{O}_i, \tilde{\gamma}_i, \tilde{\sigma}_i)$.
4. Repeat steps 2 and 3 until convergence, and return $\boldsymbol{\mu}_q$ and $(\tilde{O}_i, \tilde{\gamma}_i, \tilde{\sigma}_i)_{i=1}^m$.

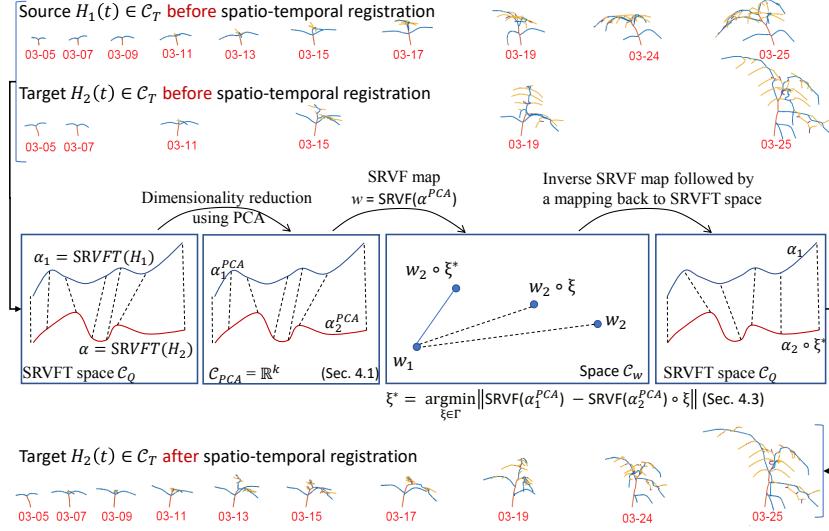


Fig. 3: The proposed framework for the analysis of 4D tree-shaped structures. The key idea is to represent 4D trees H_1 and H_2 as curves in the SRVFT space \mathcal{C}_Q , which is Euclidean but of infinite dimension. By learning a PCA subspace \mathcal{C}_{PCA} , 4D tree-shaped structures become curves in \mathbb{R}^k . However, instead of using the nonlinear elastic metric in \mathcal{C}_T to model temporal variability, we further map the curves to the SRVF space where the \mathbb{L}^2 metric is equivalent to the full elastic metric. All the analysis can be performed in this space using standard vector calculus and mapped back to the original space for visualization. The computation time of the proposed temporal registration between two 4D plants is on average 0.037s for tomato and 0.006s for maize 4D plants.

One can map μ_q back to the original tree shape to obtain the mean tree μ , which can be used for visualization. Let $Q_i = (Q_i, \tilde{O}_i, \tilde{\gamma}_i, \tilde{\sigma}_i)$ be the SRVFT representation of Q_i but optimally registered onto the mean μ_q , and $\mathbf{v}_i = Q_i - \mu_q$. The leading eigenvectors Λ_i of the covariance matrix $C = \frac{1}{m-1} \sum_{i=1}^m \mathbf{v}_i \mathbf{v}_i^\top$ define the principal directions of variations while the corresponding eigenvalues λ_i define the variance along the i -th eigenvector. Thus, each SRVFT Q can be modeled as a linear combination of the k leading eigenvectors: $Q = \mu_q + \sum_{i=1}^k a_i \sqrt{\lambda_i} \Lambda_i$. The coefficients $a_i \in \mathbb{R}$ are obtained by projecting Q onto eigenvectors Λ_i . Thus, every tree shape Q in the SRVFT space can now be represented as a point $p = (a_1, \dots, a_k)^\top$ in $\mathcal{C}_Q = \mathbb{R}^k$, which is Euclidean and has a finite dimension.

4.2 Temporal registration and geodesics between 4D trees

With this representation, a 4D tree-shaped structure $\alpha : [0, 1] \rightarrow \mathcal{C}_Q$ becomes a trajectory α^{PCA} in $\mathcal{C}_{PCA} = \mathbb{R}^k$. Thus, the analysis framework presented in Sec. 4.1 for the pre-tree shape space can now be performed on \mathbb{R}^k . Fig. 3 summarizes the

entire process. To temporally register two trajectories α_1^{pca} and α_2^{pca} , we first represent them using their SRVFs and then formulate the temporal registration problem as the one of finding an optimal diffeomorphism $\xi : [0, 1] \rightarrow [0, 1]$ that is a solution to Eq. (5). We initialize ξ to be the identity diffeomorphism: $\xi : [0, 1] \rightarrow [0, 1]$ such that $\xi(t) = t$. During the optimization of Eq. (5), only the target domain is updated. In other words, a point t in $[0, 1]$ will be mapped into another point u in $[0, 1]$ such that Eq. (5) is minimized. Since $u \in [0, 1]$, the search for the optimal match is over the entire continuous domain. Also, ξ is enforced to be a diffeomorphism, *i.e.*, $\xi(t_1) = u_1, \xi(t_2) = u_2$, and $t_1 < t_2$, then $u_1 < u_2$. The spatio-temporal registration process is as follows;

- Let $S = \{H_i\}_{i=1}^n$ be a set of n 4D tree-shaped structures where $H_i(t) \in \mathcal{C}_T$, for $t \in [0, 1]$, is a 3D tree-structured shape after normalization for scale.
- **Step 1: Spatial registration.**
 - Map every 4D tree $H_i \in S$ to its SRVFT representation α_i . Thus, we obtain a new set $\{\alpha_i\}_{i=1}^n$ such that $\alpha_i(t) = \text{SRVFT}(H_i(t))$.
 - Within-sequence registration (Stage 1 of the spatial registration, Sec. 3.3):
 - * For $i = 1$ to n , spatially register every 3D tree-shape $\alpha_i(t)$ in the i -th sequence α_i to its next 3D tree-shape in the sequence (Sec. 3.3).
 - * For simplicity of notation, let from now on $\{\alpha_i\}$ denote the new set.
 - Cross sequence registration (Stage 2 of the spatial registration): to spatially align α_1 onto α_2 ,
 - * Map $\alpha_i, i \in \{1, 2\}$ to the PCA space $\mathcal{C}_{\text{PCA}} = \mathbb{R}^k$ to obtain $\{\alpha_i^{\text{pca}}\}$.
 - * Interpolate, linearly, the samples of α_1^{pca} and α_2^{pca} and then discretize them at equidistances.
 - * $\forall t$, spatially register $\alpha_1^{\text{pca}}(t)$ onto $\alpha_2^{\text{pca}}(t)$ (Sec. 3.3).
- **Step 2: Temporal registration.**
 - Map the spatially registered α_i^{pca} to their SRVF representation w_i . Let \mathcal{C}_w be the space of such SRVFs. For any two curves w_1 and w_2 :
 - * Find ξ^* that optimally aligns w_2 onto w_1 by solving $\xi^* = \text{argmin}_{\xi \in \Xi} \|w_1, w_2 \circ \xi\|^2$.
 - * Set $w_2 \leftarrow w_2 \circ \xi^*$ and map it back to \mathcal{C}_T for visualization.

4D geodesics: The advantage of the proposed representation is that the \mathbb{L}^2 metric in the SRVF space \mathcal{C}_w of curves is equivalent to the full elastic metric. Thus, geodesics under the complex metric become straight lines in this space, *i.e.*, the geodesic Λ_w between w and \tilde{w}_2 is the linear interpolation: $\Lambda_w(\tau) = (1 - \tau)w + \tau\tilde{w}_2$, $\tau \in [0, 1]$. To visualize it, we map all 4D trees in Λ_w to the original space of trees. Fig. 4 shows an example of a geodesic between two 4D trees where each row corresponds to one 4D tree along the geodesic. The first row corresponds to $\Lambda_w(0) = w_1$ and the last row corresponds to $\Lambda_w(1) = \tilde{w}_2$. The middle row, *i.e.*, $\tau = 0.5$, is the mean 4D tree shape between w_1 and \tilde{w}_2 .

5 Statistical Analysis of 4D Tree Shapes

The ability to compute spatio-temporal correspondences and geodesics between 4D tree-shaped structures enables a wide range of 4D shape analysis tasks. In

this section, we show how these fundamental tools can be used to compute a 4D atlas, which includes the mean and modes of variability, of a collection of tree-shaped 4D structures. Let $\{H_1, \dots, H_n\}$ be a set of n 4D tree shapes and $\{w_1, \dots, w_n\}$ their corresponding trajectories in \mathcal{C}_w . To simplify the notation, we assume that these 4D shapes are all spatially registered using the procedure described in Sec. 3.3 and Sec. 4.2. Thus, one can compute the mean, the modes of variations, and fit generative models for 4D tree-structured shape generation.

Mean of 4D tree-like shapes. It is the 4D tree that is as close as possible to all of the 4D trees in the dataset under the specified metric. This is the Karcher mean μ_q and is mathematically defined as $\bar{w} = \operatorname{argmin}_w \sum_{i=1}^n \min_{\xi_i \in \Xi} \|w - q_i \circ w_i\|^2$, which we solve via a gradient descent approach:

1. Set $\bar{w} = w_1$.
2. For $i = 1 : n$
 - Optimally register w_i onto \bar{w} using Step 2 of the algorithm of Sec. 4.2.
 - Let ξ_i be the reparameterization that temporally aligns w_i onto w .
3. Set $w = \frac{1}{n} \sum_{i=1}^n w_i \circ \xi_i$.
4. Repeat steps 2 and 3 until convergence, and return \bar{w} , $\{\xi_i, \tilde{w}_i = w_i \circ \xi_i\}_{i=1}^n$.

Finally, the mean curve \bar{w} can be mapped back to a 4D tree, \bar{H} , which represents the mean of the 4D tree-shaped structures in the dataset. The inverse mapping is performed following the procedure presented in Sec. 4.2.

Modes of variations of 4D tree-shaped structures. Since the SRVF space is Euclidean, we compute the principal directions of variation using linear PCA. Let $\{\tilde{w}_i\}_{i=1}^n$ be the spatio-temporally registered 4D tree structures in \mathcal{C}_w . We first compute their covariance matrix and then take its k -leading eigenvalues $u_i, i = 1, \dots, k$ and their corresponding eigenvectors $\delta_i, i = 1, \dots, k$. This allows us to represent any curve (*i.e.*, 4D trees) in the i^{th} principal direction as $w_\tau = \bar{w} + \tau \sqrt{u_i} \delta_i$, $\tau \in \mathbb{R}$. For visualization, we map w_τ back to the original representation of 4D tree shapes, using the one-to-one inversion discussed in Sec. 4.2.

4D tree-like shape generation: Now that we have the mean \bar{w} and k principal directions of variation, $\delta_i, i = 1, \dots, k$, we can generate a 4D tree-like shape by sampling k real values $\tau_i \in \mathbb{R}$ and then computing $w = \bar{w} + \sum_{i=1}^k \tau_i \sqrt{u_i} \delta_i$, $\tau_i \in \mathbb{R}$. Next, w can be mapped back to the original tree-shape space for visualization using the one-to-one inversion discussed in Sec. 4.2. Note that this generation process is random since the real values τ_i are randomly picked. However, one can control the generation process, *e.g.*, by restricting these values to be within a specific range to ensure the plausibility of the generated samples.

6 Results and discussions

We demonstrate the proposed framework and evaluate its performance on the Pheno4D dataset [16], which contains 4D models of seven tomato and seven maize plants. All the models are not registered onto each other, neither spatially nor temporally; see the Supplementary Material for further details and results with ablation studies.

	Before reg.			After reg. (ours)			After reg. [3]			After reg. [15]		
	Mean	Median	Std.	Mean	Median	Std.	Mean	Median	Std.	Mean	Median	Std.
Tomato	325	347	52	136	136	27	416	430	85	261	245	56
Maize	35	30	8	10	9	2.5	554	496	216	19	18	5

Table 1: Statistics of the error between pairs of source and target 4D plants before spatial registration and after registration using our method, [3], and [15].

6.1 Spatiotemporal Registration

(1) Spatial registration. Fig. 2 shows an example of the spatial registration of four complex 3D tree-shaped structures of tomato plants. For clarity, we only show the branch-wise correspondences up to the second layer. Figs. 2, 3, and 4 in the Supplementary Material show, respectively, an example of the spatial registration within and across two 4D tomato plants as well as a zoom-in onto individual trees. This result shows that our framework is efficient in finding correct branch-wise correspondences between complex tree-shaped structures. We evaluate the quality of the proposed spatial registration method and compare it with state-of-the-art techniques such as [3, 15] on eight pairs of 4D tomato plants and nine pairs of 4D maize plants. We use three quantitative measures:

- *Geodesic length between 4D tree-shapes.* Tab. 1 reports the mean,

median, and standard deviation of the geodesic distances before and after spatial registration of our method. We can see that the geodesic distance becomes much smaller after spatial registration with our method, compared to [3] and [15] as branches get correctly aligned across 3D trees. The residual error is due to differences in structure and growth rates between 4D plants.

- *Cycle consistency error.* To demonstrate further the quality of our registration, we have measured the cycle consistency error of the spatial registration. Given a source and a target 3D trees, the registration maps a point \mathbf{x} on the source to a point \mathbf{y} on the target. We then map \mathbf{y} , using the registration procedure, back onto the source tree to lead to a point \mathbf{x}' . The registration procedure is accurate if \mathbf{x}' and \mathbf{x} are very close to each other. Thus, we define the registration error as the percentage of points \mathbf{x} whose

ϵ	Our method			Pan <i>et al.</i> [15]		
	Mean	Median	Std.	Mean	Median	Std.
0.1	0.94%	0.08%	2.01%	10.33%	7.09%	10.6%
0.05	2.00%	0.15%	3.62%	15.89%	17.88%	10.79%
0.02	2.87%	0.65%	4.17%	20.39%	22.22%	9.83%
0.01	5.45%	2.76%	6.15%	21.95%	22.22%	8.64%

Table 2: Cycle consistency errors (\downarrow).

(a) Between a source and a target 4D tree.

	Before reg.			After reg.		
	Mean	Median	Std.	Mean	Median	Std.
Tomato	382	384	69	279	282	47
Maize	59	60	12	34	37	11

(b) Between a 4D tree and its randomly warped version.

	Before reg.			After reg.		
	Mean	Median	Std.	Mean	Median	Std.
Tomato	356	375	138	26	27	5
Maize	71	70	8	15	16	3

Table 3: Statistics of the error between pairs of 4D source and target plants before and after their temporal registration.

cycle consistency distance $\|\mathbf{x} - \mathbf{x}'\|$ is higher than a threshold ϵ . Tab. 2 summarizes the mean, median, and standard deviation of this error for our method and [15] (the lower the error, the better) for different values of ϵ . Our method significantly outperforms the state-of-the-art.

- *Computation time.* On average, [3] requires 30s, [15] requires 20.1s, while our method requires only 2.3s to register two 3D trees. This shows that the proposed method is significantly faster than [3] and [15].

(2) Temporal registration. To demonstrate the proposed temporal registration, we first randomly resample the 4D sequences to simulate unsynchronised 4D plants and then take pairs of these re-sampled 4D sequences and realign them, temporally, using the proposed temporal registration tools. Fig. 3 shows an example of a temporal registration between complex 4D tree-shaped structures that grow at different rates. We can see that, after the temporal registration, the growth rate becomes closer to the groundtruth (see also Fig. 6 of the Supplementary Material). The Supplementary Material also provides more visual results. Similar to the spatial registration, we use the geodesic distance between 4D trees, before and after temporal registration, as a measure of the quality of registration (the smaller the better). Tab. 3 reports the mean, median, and standard deviation of those errors for three different cases. For all of the cases, our proposed temporal registration significantly reduces the geodesic distances between 4D plants: *In Tab. 3-(a)*, we randomly choose six pairs of 4D tomato plants and four pairs of maize plants. Then, we compute errors between each pair before and after temporal registration. We observe that the proposed temporal registration makes the target closer to the source 4D plant. Note that there is a residual error which is due to the differences in the branching structures. *In Tab. 3-(b)*, we register randomly warped 4D plants to their ground truth. Since both 4D plants have the same branching structures and are only temporarily warped, temporal registration significantly reduces the error, which demonstrates that the proposed temporal registration can align two 4D shapes with minimal error.

Overall, the temporal registration takes 6 to 37ms; see Table 1 in the Supplementary Material, which also provides an ablation study that analyzes the spatio-temporal registration of 4D tree-shaped objects.

6.2 Geodesics between 4D Trees

Fig. 8 in the Supplementary Material shows a geodesic path, before spatiotemporal registration, between a source 4D plant (top row) and a target 4D plant (bottom row). The in-between rows correspond to 4D plants sampled at equidistances along the geodesic path between the source and target. In this example, the target has a different growth rate than the source. It also has missing samples. Thus, we interpolate it and sample it at the same time intervals as the source but we do not perform spatiotemporal registration. Consequently, the target follows a different growth pattern than the source. As we can see in this figure, since the two 4D plants are not spatio-temporally registered, the 4D geodesic

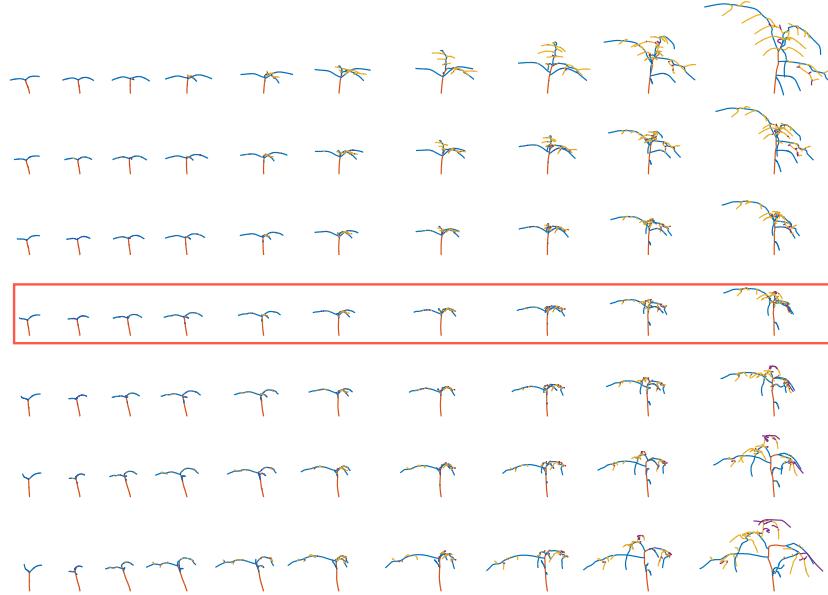


Fig. 4: The 4D geodesic between the two 4D tree shapes of Fig. 8 in the Supplementary Material after spatiotemporal registration. The highlighted row is the mean 4D tree. This geodesic computation requires on average 0.025s for 4D tomato plants.



Fig. 5: The mean 4D plant shape of the seven registered 4D tomato plants in Fig. 10-b in the Supplementary Material. The computation time is in the order of 0.0006s.

does not look realistic as the intermediate 4D plants along the geodesic exhibit significant shrinkage. Fig. 4 shows the geodesic between the same 4D plants but after spatiotemporal registration using the proposed framework. We can clearly see that the target 4D plant (last row) is temporally well aligned with the source 4D plant (first row). Also, we can also see that the branches of the intermediate plants along the geodesic path do not suffer from the shrinkage observed in Fig. 8 in the Supplementary Material prior to the registration.

6.3 Summary statistics and 4D tree-shae synthesis and generation

We took seven unregistered 4D tomato plants from the Pheno4D dataset and computed their mean and modes of variations using the proposed framework. Fig. 10 in the Supplementary Material shows the seven 4D tomato plants before (Fig. 10-a) and after (Fig. 10-b) co-registering them onto each other using the proposed spatiotemporal registration. Fig. 5 shows the mean 4D plant of these

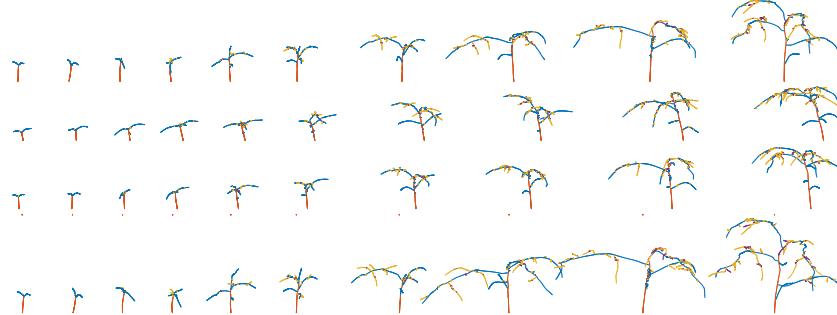


Fig. 6: Randomly generated 4D tomato plants (one 4D plant per row). The model requires on average 0.0001s to generate a 4D tomato plant.

seven 4D plants computed using the proposed framework. We can see that the mean captures the main characteristics of these seven 4D plants. Fig. 12 in the Supplementary Material shows the first and second principal directions of variation that represent the spatiotemporal shape variability in the seven 4D tomato plants. The Supplementary Material provides more results.

Fig. 6 shows four randomly generated 4D tomato plants. In this experiment, the generative model is learned from seven 4D tomato plants. To obtain plausible random 4D plants, we restricted the randomization within -3 to $+3$ times the standard deviation along each principal direction of variation. The Supplementary Material provides additional examples of 4D tree-shaped structures randomly generated from the learned generative model.

7 Conclusion

We have proposed a novel theoretical framework and a set of new computational tools for analyzing tree-shaped 4D structures, *i.e.*, 3D objects that have a tree structure and grow and deform over time. These tools include mechanisms for the spatiotemporal registration of and geodesics computation between tree-shaped 4D structures. These are the key building blocks for learning statistical models and generative models from collections of tree-shaped 4D structures such as growing plants. Our key contribution is to model the spatiotemporal variability in 4D tree shapes as trajectories in the SRVF space, reducing the problem to that of analyzing curves in an Euclidean space. As demonstrated in the experiments, the proposed computational tools can handle complex tree-shaped 4D objects that undergo complex non-rigid and topological deformations. Although we evaluated the framework on 4D plants, it is general and can be used for other types of tree-shaped structures such as neuronal structures in the brain, blood vessels, and airway trees. The current framework is limited to the skeleton structure of tree-shaped 4D objects. In the future, we plan to extend the framework to handle the full 3D geometry of the branches in tree-shaped 3D and 4D objects. **Potential negative impact.** To the best of our knowledge, this work may not have any negative impact on society.

Acknowledgements

This work is supported by the Australian Research Council (ARC) Discovery Projects no. DP220102197, DP210101682, and DP210102674. Tahmina Khanam is funded by the Murdoch University International Postgraduate Scholarship. Guan Wang is supported by the Scientific Research Foundation for Yangtze Delta Region Institute of University of Electronic Science and Technology of China (Huzhou) under Grant U032200114. Anuj Srivastava is supported by NSF DMS 1953087 and NSF DMS 241374. The authors would like to thank Adam Duncan for sharing with us the code of [4]. The source code of this project is publicly available, for research purposes, at https://github.com/Tahmina979/4Dtreeshape_project.

References

1. Besl, P., McKay, N.: A method for registration of 3-d shapes, *ieee t. pattern anal.*, **14**, 239–256 (1992)
2. Chebrolu, N., Läbe, T., Stachniss, C.: Spatio-temporal non-rigid registration of 3d point clouds of plants. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 3112–3118. IEEE (2020)
3. Chebrolu, N., Magistri, F., Läbe, T., Stachniss, C.: Registration of spatio-temporal point clouds of plants for phenotyping. *PloS one* **16**(2), e0247243 (2021)
4. Duncan, A., Klassen, E., Srivastava, A.: Statistical shape analysis of simplified neuronal trees. *The Annals of Applied Statistics* **12**(3), 1385–1421 (2018)
5. Feragen, A., Hauberg, S., Nielsen, M., Lauze, F.: Means in spaces of tree-like shapes. In: 2011 International Conference on Computer Vision. pp. 736–746. IEEE (2011)
6. Feragen, A., Lauze, F., Lo, P., de Bruijne, M., Nielsen, M.: Geometries on spaces of treelike shapes. In: Computer Vision–ACCV 2010: 10th Asian Conference on Computer Vision, Queenstown, New Zealand, November 8–12, 2010, Revised Selected Papers, Part II 10. pp. 160–173. Springer (2011)
7. Feragen, A., Lo, P., de Bruijne, M., Nielsen, M., Lauze, F.: Toward a theory of statistical tree-shape analysis. *IEEE transactions on pattern analysis and machine intelligence* **35**(8), 2008–2021 (2012)
8. Feragen, A., Owen, M., Petersen, J., Wille, M.M., Thomsen, L.H., Dirksen, A., de Bruijne, M.: Tree-space statistics and approximations for large-scale analysis of anatomical trees. In: Information Processing in Medical Imaging: 23rd International Conference, IPMI 2013, Asilomar, CA, USA, June 28–July 3, 2013. Proceedings 23. pp. 74–85. Springer (2013)
9. Jiang, B., Ren, X., Dou, M., Xue, X., Fu, Y., Zhang, Y.: Lord: Local 4d implicit representation for high-fidelity dynamic human modeling. In: European Conference on Computer Vision. pp. 307–326. Springer (2022)
10. Jiang, B., Zhang, Y., Wei, X., Xue, X., Fu, Y.: H4d: human 4d modeling by learning neural compositional representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19355–19365 (2022)
11. Laga, H.: A survey on nonrigid 3d shape analysis. *Academic Press Library in Signal Processing*, Volume 6 pp. 261–304 (2018)

12. Laga, H., Padilla, M., Jermyn, I.H., Kurtek, S., Bennamoun, M., Srivastava, A.: 4d atlas: Statistical analysis of the spatiotemporal variability in longitudinal 3d shape data. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(2), 1335–1352 (2022)
13. Lobefaro, L., Malladi, M.V., Vysotska, O., Guadagnino, T., Stachniss, C.: Estimating 4d data associations towards spatial-temporal mapping of growing plants for agricultural robots. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 4212–4218. IEEE (2023)
14. Magistri, F., Chebrolu, N., Stachniss, C.: Segmentation-based 4d registration of plants point clouds for phenotyping. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 2433–2439. IEEE (2020)
15. Pan, H., Hétroy-Wheeler, F., Charlaix, J., Colliaux, D.: Multi-scale space-time registration of growing plants. In: 2021 International Conference on 3D Vision (3DV). pp. 310–319. IEEE (2021)
16. Schunck, D., Magistri, F., Rosu, R.A., Cornelissen, A., Chebrolu, N., Paulus, S., Léon, J., Behnke, S., Stachniss, C., Kuhlmann, H., et al.: Pheno4d: A spatio-temporal dataset of maize and tomato plant point clouds for phenotyping and advanced plant analysis. *Plos one* **16**(8), e0256340 (2021)
17. Srivastava, A., Klassen, E., Joshi, S.H., Jermyn, I.H.: Shape analysis of elastic curves in euclidean spaces. *IEEE transactions on pattern analysis and machine intelligence* **33**(7), 1415–1428 (2010)
18. Tsumura, R., Morishima, Y., Koseki, Y., Yoshinaka, K.: Body surface registration considering individual differences with non-rigid iterative closest point. *International Journal of Computer Assisted Radiology and Surgery* pp. 1–10 (2023)
19. Wang, D., Puttonen, E., Casella, E.: Plantmove: A tool for quantifying motion fields of plant movements from point cloud time series. *International Journal of Applied Earth Observation and Geoinformation* **110**, 102781 (2022)
20. Wang, G., Laga, H., Jia, J., Xie, N., Tabia, H.: Statistical modeling of the 3d geometry and topology of botanical trees. In: *Computer graphics forum*. vol. 37, pp. 185–198. Wiley Online Library (2018)
21. Wang, G., Laga, H., Srivastava, A.: Elastic shape analysis of tree-like 3d objects using extended srvf representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)
22. Wang, G., Laga, H., Xie, N., Jia, J., Tabia, H.: The shape space of 3d botanical tree models. *ACM Trans. Graph.* **37**(1) (jan 2018)
23. Zhang, J., Yao, Y., Deng, B.: Fast and robust iterative closest point. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(7), 3450–3466 (2021)
24. Zhang, T., Elnashef, B., Filin, S.: Spatio-temporal registration of plants non-rigid 3-d structure. *ISPRS Journal of Photogrammetry and Remote Sensing* **205**, 263–283 (2023)