

GrowFields: Compositional 4D Neural Fields for Topology-Changing Plant Growth

Joaquin Gajardo¹, Michele Volpi², Marko Mihajlovic¹, Siyu Tang¹,
Lukas Roth^{1,3}, and Sergey Prokudin¹

¹ ETH Zürich, Switzerland

² Swiss Data Science Center, Switzerland

³ Martin-Luther-Universität Halle-Wittenberg, Germany

<https://joaquin-gajardo.github.io/growfields>

Abstract. Quantifying plant growth dynamics from sparse longitudinal 3D observations is fundamental for agriculture and plant sciences. Yet, plants pose unique challenges: they undergo intricate non-rigid deformations, exhibit changing topology as new organs emerge, and often lack explicit temporal correspondences between consecutive data acquisitions due to newly formed tissue. Methods designed for general scenes struggle to model topology changes and asynchronous organ growth characteristic of plants. To address these challenges, we introduce GrowFields, a compositional dynamic neural field representation for organ-aware 4D plant growth modelling from point cloud time series. Our approach decomposes a plant into its constituent organs and aligns each organ into its own canonical coordinate frame, isolating intrinsic growth patterns from global plant motion. We then learn a shared continuous neural deformation field that models temporal dynamics across all organs, conditioned on learnable per-organ latent codes capturing organ identity and growth characteristics. The resulting modular yet unified representation naturally accommodates the asynchronous development of plant organs while remaining grounded in the practical setting of organ-level plant tracking. We evaluate GrowFields on growth sequences from four plant species, assessing geometric fitting and organ tracking accuracy using manually annotated leaf-tip trajectories. Results demonstrate consistent improvements in spatial precision, temporal coherence, and morphological fidelity over a range of existing representations.

1 Introduction

Understanding and quantifying plant growth dynamics is central to agriculture and plant sciences. Plants are complex, dynamic, multi-scale organisms whose response to environmental conditions can be partly characterised through the temporal dynamics of observable organ-level traits, such as the timing of organ emergence and leaf expansion rates [22]. While 3D sensing technologies and plant modelling methods have advanced significantly, enabling detailed reconstructions of plant geometry [9, 45, 74], the temporal dimension of plant growth




Fig. 1: Overview of GrowFields. Given a time series of segmented plant point clouds, the plant is decomposed into organs and each organ is aligned to a canonical coordinate frame (A). For an initial point $\mathbf{x}_0$ belonging to organ i , its trajectory over time is predicted by integrating a shared neural deformation field f_θ (B), conditioned on a learnable organ-specific latent code $\mathbf{z}_i$: $\mathbf{x}(t, \mathbf{z}_i, \mathbf{x}_0) = \mathbf{x}_0 + \int_0^t f_\theta(\mathbf{x}(\tau), \tau, \mathbf{z}_i) d\tau$. The predicted organ motions are recomposed in the global coordinate frame to produce a temporally coherent 4D reconstruction of the full plant (C). The SIREN-based field [60] produces smooth trajectories for all points, modelling continuous organ growth dynamics (D).

remains underexplored. Modelling plants in 4D from sparse temporal observations could enable continuous measurement of organ-specific growth rates, early stress detection, and quantitative analysis of plant development, providing a powerful new tool for plant scientists and agronomists [18].

Faithfully capturing these dynamics at the organ level remains challenging. Plants undergo complex non-rigid deformations, exhibit changing topology as new organs appear and old ones decay, and data acquisition is often sparse and noisy. Furthermore, newly formed tissue has no direct temporal correspondence, breaking assumptions typically made in non-rigid registration and dynamic scene reconstruction.

General-purpose 4D reconstruction techniques, based on temporal Signed Distance Functions (SDFs) [55, 67], or flow- [43] and point-based dynamic models [52], usually assume globally consistent topology or smoothly varying deformation fields. Recent work such as GrowFlow [37] models plant growth through a continuous deformation field learned in reverse time to mitigate the appearance of new geometry. While effective for modelling smooth plant deformation, such global formulations do not explicitly model organ identity or asynchronous

organ development, making it difficult to represent organ emergence, large non-rigid growth, and long-term organ-level dynamics. Other plant-specific methods leverage domain-specific cues such as organ segmentation [38] or skeleton extraction [6, 26, 46, 76], but often rely on multi-stage pipelines and manually engineered matching strategies that may not generalise across species or sensing conditions. Moreover, most existing approaches address reconstruction, tracking, and dynamic modelling in isolation, and therefore struggle to jointly capture continuous plant deformation, organ identity, and topology changes under realistic temporally sparse acquisition settings.

Motivated by the observation that plant organs grow asynchronously and follow distinct deformation patterns, we propose *GrowFields*, a part-based dynamic neural field framework for continuous 4D plant reconstruction from longitudinal point clouds (Figure 1). Instead of modelling plants as a single deforming structure, our approach leverages their natural modular organisation: the plant is decomposed into organs whose geometry and motion can be analysed independently while remaining part of a coherent whole. By representing organ dynamics in canonical reference frames and learning a shared neural model of growth conditioned on learnable latent codes representing organ identity, the framework captures common geometric and temporal patterns while allowing individual organs to evolve according to their own trajectories. Operating directly on segmented point clouds aligns with common phenotyping workflows [31, 75], while the compositional neural field formulation provides a flexible framework for learning nonlinear plant dynamics at the organ level.

To rigorously assess the effectiveness of GrowFields, we conduct a comprehensive evaluation across multiple plant species exhibiting diverse morphologies, growth patterns, and environmental conditions. Our experimental setup includes newly curated leaf-tip annotations on relevant public datasets, enabling quantitative assessment of temporal coherence, organ-level tracking, and geometric stability. This application-grounded evaluation design, combined with physiologically relevant metrics, provides a thorough empirical validation of the proposed framework and establishes a challenging benchmark for future research. We further probe the method’s practicality through experiments on robustness to noise and missing data, temporal extrapolation, and automatically segmented and tracked organ labels.

To summarise, our main contributions are:

- We introduce **GrowFields**, a compositional neural field framework for modelling topology-changing plant growth from sparse longitudinal point clouds as a modular dynamical system. It combines organ canonicalisation with a shared continuous velocity field conditioned on per-organ latent codes, effectively capturing large non-rigid growth, continuous deformation, and asynchronous organ development.
- We present an **open-source benchmark and evaluation protocol** for 4D plant growth modelling based on annotated leaf-tip trajectories, enabling quantitative evaluation of temporal coherence and organ-level tracking accuracy.

2 Related work

Understanding plant development from 3D observations requires representations able to capture geometric structure, temporal evolution, and organ-level organisation under real-world data acquisition constraints. Thus, we review prior work on plant reconstruction, dynamic plant models, organ segmentation and tracking, and dynamic and compositional neural representations.

Static plant models. Plant modelling has long been studied in computer graphics and plant sciences [12, 45, 66]. Recent approaches reconstruct plant geometry directly from images or point clouds [9, 34, 36, 73, 74]. Notably, 3D Gaussian Splatting (3DGS) [25] has recently become popular for high-fidelity plant reconstruction [3, 31, 44, 57, 58, 62] and yield-component organs [13, 24, 75], but does not explicitly encode plant structure. Structured alternatives have started to emerge, including neural leaf models [73], plant mesh reconstruction methods [36], and parametric plant representations such as Demeter [9] and Gaussian-Plant [74]. While these methods achieve high-quality 3D geometry, they remain restricted to static scenes. We instead model plant development over time, enabling the characterisation of dynamic plant behaviour and the computation of continuous growth variables.

Dynamic plant models. Early work analysed temporal plant point clouds to detect growth events or track organs across time [33, 47]. Subsequent approaches focused on skeleton- or graph-based registration [6, 41, 46, 76], typically extracting plant topology from point clouds via Laplacian contraction [5] or L_1 -medial contraction [23]. In parallel, visual SLAM [14] or 3DGS [1] have also been used to visualise temporally consistent appearance changes. Some work has also sought to interpolate registered 3D scans by simple spherical linear interpolation (SLERP) [46, 76] or optimal transport [21]. More recently, GrowFlow [37] models plant growth in reverse time by combining 3DGS with a HexPlane [4] deformation field guided by Neural Ordinary Differential Equations (Neural ODEs) [8], achieving improved interpolation results over general-purpose 4D Gaussian Splatting (4DGS) [69]. However, evaluation is limited to synthetic plants and real plants without large organ-level changes. Despite these advances, most existing approaches do not support organ-level modelling and tracking, which is essential for phenotyping and plant analysis [18]. In contrast, we learn compositional dynamic neural fields from segmented longitudinal data, enabling structure-aware modelling of asynchronous multi-organ plant development.

Organ segmentation and tracking. Organ-level decomposition is central to plant phenotyping. Methods include direct 3D instance segmentation of point clouds [9, 16, 29, 61] using annotated datasets [40, 56, 78], and image-based segmentation projected into 3D [39, 72]. Recent approaches leverage vision foundation models [59, 71] or specialised segmentation networks [7], often combined with neural rendering frameworks such as PlantSegNeRF [72]. Segmentation has also enabled organ tracking and emergence detection, using bipartite matching [17, 30], task-specific detection frameworks [28], or visual place recognition

and feature matching [19, 35, 54]. We build upon these foundations to learn plant growth dynamics from sparse temporal data, motivating the use of dynamic neural representations capable of capturing large, non-linear developmental changes at the organ-level.

Dynamic and compositional neural fields. Implicit neural representations such as SIREN [60, 70] model high-frequency signals, including geometry, as continuous coordinate-based functions. Dynamic extensions incorporate deformation fields or velocity-driven ODEs to capture temporal shape evolution [32, 53, 55, 64, 67]. However, the fine and slender structures common in plants pose challenges for state-of-the-art SDF-based formulations [36]. Point-based neural fields [52, 69] provide higher fidelity for modelling such geometry and have recently been applied to plant dynamics [37]. Additionally, several works have also explored part-aware neural representations for humans and articulated objects. NASA [11] pioneered this direction using separate neural implicit functions for individual body parts, but such designs scale poorly with increasing numbers of parts and may introduce inconsistencies at part boundaries. Subsequent approaches, including COAP [42], DANBO [63], and Neural Body [49, 50], instead learn shared compositional neural fields conditioned on part-specific embeddings. These methods demonstrate the effectiveness of compositional neural fields for modelling articulated motion under *fixed* topology. In contrast, plant development involves continuous growth and topology changes as new organs emerge. We extend the compositional neural field paradigm to this setting, conditioning a single shared continuous velocity field on per-organ latent codes to model asynchronous *organ growth* in sparse longitudinal point clouds, capturing the non-rigid, topology-changing nature of plant development.

3 Method

3.1 Intuition

Plant growth is inherently modular: each organ (stem or leaf) evolves according to its own geometry and structural constraints. At the same time, organs are not independent; they grow in coordination as parts of a single organism. Our dynamic plant model is designed to reflect both properties. We operate on sequences of organ-segmented plant point clouds and align each organ independently into a canonical local coordinate frame. This isolates intrinsic per-organ growth patterns from global plant motion and provides stable trajectories for modelling. We then represent the temporal evolution of all organs with a single shared continuous neural field conditioned on learnable per-organ latent codes. The shared network captures geometric and temporal priors common across organs, while the latent codes encode organ identity and growth characteristics. Together, canonical alignment and latent-conditioned shared dynamics yield a unified yet modular representation of complex, asynchronous plant development.

3.2 Data preparation and alignment

Our input is a time-ordered sequence of 3D point clouds $\{\mathbf{X}^{(t)}\}_{t=0}^T$, each representing a plant at time t . Every point $\mathbf{x} \in \mathbf{X}^{(t)}$ carries a semantic instance label identifying its organ (stem or leaf). These labels provide temporally consistent organ-level trajectories [30].

We represent the plant at time t as a set of organ-specific point clouds $\mathbf{X}_i^{(t)}$:

$$\mathbf{X}^{(t)} = \bigcup_{i=1}^{N^{(t)}} \mathbf{X}_i^{(t)}, \quad \mathbf{X}_i^{(t)} \in \mathbb{R}^{N_i^{(t)} \times 3}, \quad (1)$$

where $N^{(t)}$ is the number of organs observed at time t and $N_i^{(t)}$ is the number of points belonging to organ i at that time.

Leaf-level alignment. Leaves exhibit large pose variation due to growth and self-motion. To obtain a stable reference frame for intrinsic shape modelling, we apply a PCA-based alignment inspired by physiological leaf coordinates [27]. Each leaf is oriented such that its principal axis aligns with the global Z -axis, while its base is translated to the origin. This canonicalisation removes global pose variation while preserving intrinsic geometry. For each timepoint, we store the rigid transformations $\{\mathbf{R}_i^{(t)}, \mathbf{t}_i^{(t)}\}$ for later reprojection.

3.3 Organ-conditioned dynamic neural field

Formulation. After canonical alignment, organ growth is modelled as a shared continuous neural deformation field conditioned on organ identity. We define a sinusoidal network [60] (SIREN) f_θ that predicts a time-dependent velocity field:

$$f_\theta : (\mathbf{x}, t, \mathbf{z}_i) \mapsto \mathbf{v}(\mathbf{x}, t, \mathbf{z}_i) \in \mathbb{R}^3, \quad (2)$$

where $\mathbf{x}$ denotes a 3D point in canonical coordinates, $t \in [0, 1]$ denotes normalised time per-organ, and $\mathbf{z}_i \in \mathbb{R}^d$ is a learnable latent code assigned to organ i . The latent code is concatenated with spatial and temporal inputs to the network, yielding $[\mathbf{x}; t; \mathbf{z}_i] \in \mathbb{R}^{4+d}$. The network parameters θ are shared across all organs.

Latent code design. We follow the auto-decoder paradigm [48]: no encoder is used, and the latent codes $\{\mathbf{z}_i\}_{i=1}^N$ are optimised jointly with the network parameters. Each organ instance receives its own code, which are initialised from a zero-mean Gaussian $\mathcal{N}(\mathbf{0}, \sigma_z^2 \mathbf{I})$ and regularised globally toward the origin to prevent unbounded growth:

$$\mathcal{L}_{\text{code}} = \frac{1}{N} \sum_{i=1}^N \|\mathbf{z}_i\|_2^2. \quad (3)$$

Discrete dynamics. Starting from the initial organ point cloud $\mathbf{X}_i^{(0)}$, we propagate its geometry by accumulating predicted velocities through Euler integration:

$$\mathbf{x}^{(t)} = \mathbf{x}^{(0)} + \sum_{\tau=0, \Delta t, \dots, t-\Delta t} f_{\theta}(\mathbf{x}^{(\tau)}, \tau, \mathbf{z}_i) \Delta t, \quad (4)$$

which approximates the continuous-time formulation:

$$\frac{d\mathbf{x}}{dt} = f_{\theta}(\mathbf{x}, t, \mathbf{z}_i) \quad \mathbf{x}^{(t)} = \mathbf{x}^{(0)} + \int_0^t f_{\theta}(\mathbf{x}^{(\tau)}, \tau, \mathbf{z}_i) d\tau. \quad (5)$$

Training. The shared field f_{θ} and latent codes $\{\mathbf{z}_i\}$ are optimised jointly. At each iteration, we sample a mini-batch of organs $\mathcal{B} \subset \{1, \dots, N\}$. For each organ $i \in \mathcal{B}$, we select source and target timepoints (s, k) with $k > s$, randomly sample points $\mathbf{x}^{(s)} \in \mathbf{X}_i^{(s)}$, and propagate them forward through the learned velocity field.

Chamfer supervision. We supervise the predicted trajectory using bidirectional Chamfer distance computed at every predicted timestep:

$$\mathcal{L}_{\text{CD}}^{(i)} = \frac{1}{k-s} \sum_{t=s+1}^k \text{CD}(\hat{\mathbf{X}}_i^{(t)}, \mathbf{X}_i^{(t)}). \quad (6)$$

The Chamfer distance between two point sets $\hat{\mathbf{X}}$ and $\mathbf{X}$ is defined as

$$\text{CD}(\hat{\mathbf{X}}, \mathbf{X}) = \frac{1}{|\hat{\mathbf{X}}|} \sum_{\mathbf{x} \in \hat{\mathbf{X}}} \min_{\mathbf{y} \in \mathbf{X}} \|\mathbf{x} - \mathbf{y}\|_2^2 + \frac{1}{|\mathbf{X}|} \sum_{\mathbf{y} \in \mathbf{X}} \min_{\mathbf{x} \in \hat{\mathbf{X}}} \|\mathbf{x} - \mathbf{y}\|_2^2. \quad (7)$$

Total objective. The Chamfer loss is computed per organ and averaged over the sampled mini-batch, yielding the following final training objective:

$$\mathcal{L}_{\text{total}} = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \mathcal{L}_{\text{CD}}^{(i)} + \lambda_{\text{code}} \mathcal{L}_{\text{code}}. \quad (8)$$

3.4 Full-plant dynamics

Composition. A complete 4D reconstruction of the plant requires merging the continuous organ trajectories into a single, temporally coherent structure. For each organ i , we select its earliest observed frame and use the corresponding point cloud $\mathbf{X}_i^{(0)}$ as initialisation. Note that our model does not *predict* organ emergence; instead, each newly observed organ is incorporated into the shared dynamical representation upon its first observation. This sidesteps the ill-posed problem of synthesizing unseen geometry while still yielding a coherent, topology-changing full-plant reconstruction. Given the shared neural field

f_θ and the organ’s latent code $\mathbf{z}_i$, we propagate the geometry forward through the field, integrating its velocity over time:

$$\hat{\mathbf{X}}_i^{(t+\Delta t)} = \hat{\mathbf{X}}_i^{(t)} + f_\theta(\hat{\mathbf{X}}_i^{(t)}, t, \mathbf{z}_i) \Delta t, \quad (9)$$

producing a continuous temporal trajectory within the canonical coordinate frame. Each predicted organ state $\hat{\mathbf{X}}_i^{(t)}$ is then transformed back to the global plant space via its rigid alignment parameters:

$$\hat{\mathbf{X}}_{i,\text{global}}^{(t)} = \hat{\mathbf{X}}_i^{(t)} \mathbf{R}_i^{(t)\top} + \mathbf{t}_i^{(t)}. \quad (10)$$

The rigid poses $\{\mathbf{R}_i^{(t)}, \mathbf{t}_i^{(t)}\}$ are recovered from the canonical alignment of each observed scan (subsection 3.2); for frames with no observed scan, such as dense interpolation, they are obtained via SLERP on rotations and linear interpolation on translations.

Plant synthesis. The full plant model at time t is obtained by combining all reconstructed organ point sets:

$$\hat{\mathbf{X}}^{(t)} = \bigcup_{i=1}^{N^{(t)}} \hat{\mathbf{X}}_{i,\text{global}}^{(t)}. \quad (11)$$

This synthesis yields a complete 4D plant trajectory that integrates asynchronous, organ-specific growth patterns into a coherent dynamic representation.

4 Experiments

4.1 Experimental setup

Dataset and evaluation protocol. We evaluate on the TrackPlant3D dataset [30], comprising daily LiDAR scans of four species (maize, sorghum, tobacco, tomato) with per-organ instance labels, collected from two public sources [10, 56]. We use eight test and four validation sequences (one per species, for hyperparameter tuning), spanning 5–17 timepoints each; to enable tracking evaluation, we additionally annotated leaf tips on all scans where they could be clearly identified. We report Chamfer distance (CD, mm^2) for geometric fit and End-Point Error (EPE, mm) on the annotated leaf tips for tracking accuracy, averaged over timesteps per sequence and computed on test points only. Unless otherwise stated, all reported metrics are averaged over five independent runs, and predicted geometry is reprojected using the available ground-truth per-frame poses so that the metrics isolate shape and growth from pose. Further dataset, implementation, and baseline details are in the supplementary material.

Following prior neural-field works [77], we randomly split each organ’s point cloud into 50% train and 50% test points independently at each timestep, ensuring that the 234 leaf-tip points are always in the test set. This split is fixed and

Table 1: Full plant reconstruction results. Mean per-sequence temporal metrics across eight growth sequences: CD (mm^2) for geometric accuracy and EPE (mm) on annotated leaf tips for tracking accuracy. Full-plant baselines take no segmentation input, whereas the part-aware COAP and our models use the same ground-truth organ segmentation. Our full model achieves the best overall performance; DSR is a 4D-implicit method with no deformation field, so EPE is undefined. Colours indicate column-wise ranking: 1st, 2nd, and 3rd. All metrics are averaged over five independent runs.

Method	MaC2		MaC3		SoC2		SoH2		TbC1		TbS3		TmIH3		TmIS1		Mean	
	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓
<i>Full-plant baselines</i>																		
NDF [65]	14,496.5	73.6	1,227.9	95.8	1,313.4	54.9	303.0	54.3	290.2	14.0	135.6	17.9	14.1	10.4	52.7	22.0	2,229.2	42.9
NVFi [32]	2,178.9	72.2	1,495.3	97.4	873.0	56.3	161.8	56.0	103.0	15.1	86.0	18.2	32.6	9.4	137.3	21.9	633.5	43.3
DSR [64]	481.3	—	194.9	—	186.4	—	61.1	—	29.9	—	25.0	—	13.1	—	84.0	—	134.5	—
CanFields [67]	2,025.7	77.6	1,774.8	98.9	611.4	56.4	267.6	59.5	102.1	16.6	67.3	21.1	23.0	10.4	53.5	14.4	615.7	44.4
DPF [52]	8.17	23.18	5.00	5.30	1.21	5.78	0.92	57.22	0.36	7.58	0.46	16.21	0.17	1.93	3.47	3.97	2.47	15.15
<i>Part-aware methods</i>																		
COAP [42]	4.84	7.57	2.48	3.87	1.31	7.76	1.02	5.61	1.03	3.22	0.65	2.55	0.38	1.40	1.82	1.88	1.69	4.23
Ours (per-organ)	2.41	3.68	1.52	2.07	0.87	1.47	1.01	3.44	0.70	3.28	0.48	2.67	0.13	0.74	0.33	1.59	0.93	2.37
Ours (full)	2.51	2.38	1.43	2.01	0.69	1.34	0.74	0.97	0.59	1.66	0.45	1.30	0.12	0.80	0.32	1.34	0.86	1.47

shared across all methods. At inference we predict *full trajectories*: from the first observed frame (t_0) we integrate the deformation field multi-step and evaluate at every timestep against the ground truth, rather than only single-step-ahead predictions.

Baselines. We compare against general 4D shape-interpolation methods that operate on the *full plant* without segmentation input: CanFields [67], Neural Velocity Fields (NVFi) [32], Neural Diffeomorphic Flow (NDF) [65], a 4D SDF model (DSR) [64], and chained Dynamic Point Fields (DPF) [52], where each frame transition is modelled by a separate MLP (also evaluated per organ). As *part-aware* baselines that use the same ground-truth organ segmentation as our method, we extend the compositional COAP [42] encoder-decoder into a dynamic deformation field trained like GrowFields.

4.2 Results

Plant-level results. We present full plant reconstruction results in Table 1 and qualitative comparisons in Figure 2. Methods that model scene dynamics with global deformation or motion fields, such as NDF [65] and NVFi [32], struggle with the rapidly changing topology of plant growth. When new organs emerge, these models must explain previously unseen geometry through a continuous deformation of existing structure, which often leads to unstable or degenerate solutions.

SDF-based approaches such as DSR [64] and canonical-shape interpolation methods such as CanFields [67] improve geometric fitting in some cases but still produce large tracking errors. These methods favour coherent global surface evolution and therefore struggle to model thin structures such as leaves and stems whose motion and growth are largely organ-specific.




Fig. 2: Qualitative full-plant reconstruction. Colours denote organ identities from the ground-truth segmentation, except for CanFields [67], where they denote the parts discovered by its own internal motion-segmentation. Compared to per-organ MLPs, our shared latent-conditioned field resolves part boundaries correctly, illustrating the effectiveness of the latent-code formulation.

Point-based models provide a stronger baseline. Dynamic Point Fields (DPF) [52] represent geometry with explicit point primitives while learning an implicit deformation model, which improves reconstruction accuracy. However, without an explicit mechanism for modelling organ emergence and independent growth, performance still degrades when plant structure changes substantially across timesteps.

Compositional models are more robust in this setting. Our point-based adaptation of the part-aware COAP framework [42] retains the part decomposition but replaces the occupancy representation with per-part point deformation fields. This decomposition better reflects plant morphology and allows individual organs to evolve independently.

Our method further improves on these approaches by combining per-organ canonicalisation with a shared latent-conditioned deformation field. This formulation models organ-specific motion while capturing global correlations between plant parts. As shown in Table 1, the full model achieves the best overall performance across sequences and metrics, improving both reconstruction accuracy (CD) and leaf-tip tracking accuracy (EPE).

Table 2: Per-organ reconstruction results. We report mean temporal metrics for individual organ sequences grouped by their parent plant sequence. The *Mean* column reports the average across the eight plant sequences (each plant weighted equally). Compared to full-plant evaluation (Table 1), most baselines show only modest improvements despite operating on isolated organs. Our full model achieves the best overall performance, with the lowest mean CD (0.87) and EPE (1.62), indicating improved geometric reconstruction and leaf-tip tracking across organ sequences. Note that DSR is a 4D-implicit method with no deformation field, so point tracking (EPE) is undefined. All metrics are averaged over five independent runs.

Method	MaC2		MaC3		SoC2		SoH2		TbC1		TbS3		Tm1H3		Tm1S1		<i>Mean</i>	
	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓
NDF [65]	484.6	65.9	317.0	104.4	45.9	69.7	17.6	29.8	29.5	14.0	15.2	6.8	1.5	2.7	3.6	6.8	114.4	37.5
NVFi [32]	540.1	68.7	190.1	100.4	37.5	75.0	13.7	28.6	13.9	15.3	7.2	6.4	1.8	2.5	5.3	7.7	101.2	38.1
DSR [64]	41.9	—	21.6	—	26.9	—	10.1	—	15.3	—	6.5	—	1.3	—	2.1	—	15.7	—
CanFields [67]	245.0	52.1	355.7	77.1	174.8	55.3	294.6	31.1	18.9	7.1	6.6	4.7	2.0	2.7	6.1	5.1	138.0	29.4
DPF [52]	2.32	2.05	1.67	2.60	1.30	1.43	0.87	1.43	0.69	3.18	0.56	1.57	0.30	2.01	0.68	2.23	1.05	2.06
COAP [42]	4.17	8.20	2.55	5.68	2.22	16.63	0.93	8.82	0.99	5.01	0.64	2.02	0.42	1.75	1.60	1.93	1.69	6.26
Ours (per-organ MLPs)	2.10	3.47	1.52	2.89	1.13	1.59	0.96	2.34	0.56	4.13	0.50	1.88	0.24	1.42	0.51	1.84	0.94	2.44
Ours (full)	2.26	2.59	1.48	2.46	0.93	1.51	0.64	0.85	0.49	2.04	0.45	1.20	0.20	0.98	0.46	1.33	0.87	1.62

Organ-level results. Since several baselines struggle in the full-plant setting (Table 1), we additionally test the models by providing per-organ sequences independently as input, in order to probe their behaviour in a reduced, simpler, topology-invariant setting. Table 2 reports reconstruction metrics computed per organ sequence. To ensure fair comparison with baselines, we consider only organ sequences with at least three frames. This reduces the evaluation set from 49 to 45 sequences and from a total of 192 to 189 leaf-tip error measurements.

Despite operating on isolated organs, most baselines improve only modestly compared to the full-plant setting. For example, DPF achieves a mean CD of 1.05 and EPE of 2.06, while COAP obtains 1.69 CD and 6.26 EPE. Our per-organ MLP variant improves these results to 0.94 CD and 2.44 EPE. The full model achieves the best overall performance with a mean CD of 0.87 and EPE of 1.62. This indicates that modelling interactions between organs remains beneficial even when evaluation is performed at the organ level.

Interpolation. To evaluate temporal generalisation, we train the models using only every second timestep and evaluate reconstruction on both observed and interpolated frames. Results are summarised in Table 3, where we group metrics separately for *observed* (training) frames and the held-out *interpolated* frames; per-organ interpolation results are reported in the supplementary material.

Because most sequences are short, typically providing only four to five training timesteps, interpolation is challenging. Under these conditions, DPF (per-organ) remains a strong baseline and achieves the best overall interpolation performance (mean CD 10.24, EPE 4.00). This behaviour reflects the strong local fitting ability of DPF, whereas our model prioritises globally consistent growth dynamics. Our model remains competitive and performs particularly well on longer sequences such as Tm1H3, which contains 16 timesteps (8 used for training). These results indicate that the proposed representation captures smooth temporal dynamics while maintaining accurate geometric reconstruction.

Table 3: Full plant interpolation results. Models are trained using every second timestep and evaluated on both observed (training) and interpolated frames. DPF (per-organ) achieves the strongest interpolation performance overall, while our full model remains competitive among continuous neural-field approaches and performs particularly well on longer sequences. All metrics are averaged over five independent runs.

Method	MaC2		MaC3		SoC2		SoH2		TbC1		TbS3		Tm1H3		Tm1S1		Mean	
	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓
<i>Training frames</i>																		
COAP [42]	3.20	11.87	0.82	3.08	3.04	9.60	1.11	2.20	0.91	2.39	0.56	2.34	0.15	2.46	0.59	1.40	1.30	4.42
DPF (per-organ)	2.62	1.94	0.77	1.67	0.51	0.85	0.63	2.87	0.83	1.98	0.57	1.58	0.13	1.26	0.41	1.67	0.81	1.73
Ours (per-organ MLPs)	2.59	2.21	0.97	5.49	0.59	1.28	0.72	4.73	0.77	3.44	0.47	2.05	0.12	1.14	0.37	1.93	0.82	2.78
Ours (full)	2.24	2.14	0.70	1.72	0.42	0.71	0.44	0.56	0.57	5.39	0.39	1.70	0.10	0.58	0.35	1.24	0.65	1.75
<i>Interpolated frames</i>																		
COAP [42]	19.30	14.63	9.57	12.97	25.42	10.20	31.07	9.97	4.16	3.90	1.88	2.82	0.83	2.55	1.62	1.63	11.73	7.33
DPF (per-organ)	17.75	8.06	8.93	4.72	18.38	5.26	29.38	5.74	3.23	3.13	1.88	1.89	0.89	1.63	1.47	1.60	10.24	4.00
Ours (per-organ MLPs)	18.00	8.04	9.47	13.32	19.17	5.55	29.42	9.73	2.84	4.59	1.76	2.57	0.91	1.56	1.53	2.28	10.39	5.96
Ours (full)	18.01	7.76	9.18	11.82	19.53	5.34	28.60	9.69	2.93	6.49	1.75	2.48	0.89	1.12	1.42	1.88	10.29	5.82

tion. Qualitative interpolation and point-tracking comparisons are provided in the supplementary material.

Extrapolation. We additionally test longer-horizon extrapolation by training on all timesteps *except* the last three and evaluating exclusively on those held-out final frames (SoH2 is excluded, as it has only five timesteps). Predictions are multi-step trajectories rolled out from t_0 beyond the training horizon. As shown in Table 4, GrowFields outperforms COAP across all sequences; the large absolute errors on some sequences are partly explained by new organs emerging within the held-out frames, a limitation we discuss in subsection 5.1. A qualitative example is shown in Figure 3.

Table 4: Extrapolation results. Models are trained on all but the last three timesteps and evaluated only on those held-out final frames; predictions are multi-step rollouts from t_0 . We report CD ↓ / EPE ↓. GrowFields outperforms COAP on every sequence. All metrics are averaged over five independent runs.

Method	MaC2		MaC3		SoC2		TbC1		TbS3		Tm1H3		Tm1S1		Mean	
	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓
<i>Training frames</i>																
COAP [42]	1.34	12.26	0.57	5.01	0.49	1.02	0.18	1.13	0.26	1.24	0.22	1.15	0.54	1.54	0.51	3.34
Ours	1.06	1.81	0.40	0.92	0.39	1.08	0.15	0.80	0.18	1.67	0.14	0.79	0.30	1.33	0.37	1.20
<i>Extrapolated frames</i>																
COAP [42]	768.08	47.17	1632.07	103.33	620.43	42.77	79.27	9.23	15.82	5.04	8.59	5.92	33.40	6.25	451.09	31.39
Ours	356.13	20.96	321.08	69.77	337.39	27.04	15.38	4.06	8.09	4.14	2.80	2.52	6.56	4.46	149.63	18.99

4.3 Analysis

Ablation study. We analyse the contribution of the main design components in Table 5. Removing the organ’s latent code leads to the largest drop in reconstruction accuracy, increasing the mean error from 0.86/1.47 to 1.05/3.44




Fig. 3: Qualitative extrapolation. GrowFields prediction on `maize_control_plant2`, trained on all but the last three timesteps and rolled out from t_0 across the full sequence (left to right); colours denote organ identities, and the dashed box marks the three held-out frames evaluated for extrapolation. Although the model predicts plausible continued organ elongation into the held-out frames; without generative priors, geometric accuracy rapidly degrades (see subsection 5.1).

(CD/EPE). This result highlights the importance of conditioning the deformation field on a latent representation of organ state. Removing canonicalisation produces the largest drop in tracking accuracy and also substantially degrades reconstruction (mean error rises from 0.86/1.47 to 0.96/4.23 (CD/EPE)), confirming that aligning organs to a shared canonical space stabilises deformation learning across timesteps. Replacing our latent conditioning with Feature-wise Linear Modulation (FiLM) conditioning [51], which conditions the network by modulating intermediate feature activations through learned feature-wise affine transformations, similarly reduces performance (0.96/2.47). This suggests that directly injecting the latent code provides a stronger inductive bias for modelling plant motion than conditioning through feature modulation alone. Finally, removing the multi-step Chamfer loss leads to moderate but consistent degradation (0.91/1.89), indicating that it improves temporal stability and geometric consistency. Overall, the ablation results confirm that latent conditioning and canonicalisation are the most important components of the proposed formulation.

Robustness. Our method assumes organ-segmented inputs, a reasonable requirement given the rapidly maturing literature on automated 3D plant organ segmentation, which is already of sufficient quality for organ tracking [19, 30, 35, 56]. To verify that GrowFields does not rely on ground-truth labels, we re-run it on *automatically* segmented point clouds (*w/ model segmentation* in Table 5), obtained with a combination of a retrained PSegNet [29] model on all available sequences, and the TrackPlant3D organ tracking method [30]. As illustrated qualitatively in Figure 4, the method remains effective under realistic prepro-




Fig. 4: Qualitative result under automated segmentation. GrowFields prediction on `maize_control_plant2` using *automatically* segmented organ labels as input; colours denote the input organ identities. Despite noisier labels than the ground-truth segmentation, the method still recovers coherent organ-level growth.

Table 5: Ablations results and evaluation on full plants. Removing latent conditioning and canonicalisation leads to the most noticeable drop in reconstruction and tracking accuracy, highlighting the importance of these design choices. *Note that *w/model segmentation* is not directly comparable to results with GT segmentation given different averaging or number of leaf-tip evaluations (only provided as reference values). All metrics are averaged over five independent runs.

Method	MaC2		MaC3		SoC2		SoH2		TbC1		TbS3		Tm1H3		Tm1S1		Mean	
	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓	CD ↓	EPE ↓
Ours (full)	2.51	2.38	1.43	2.01	0.69	1.34	0.74	0.97	0.59	1.66	0.45	1.30	0.12	0.80	0.32	1.34	0.86	1.47
<i>w/o</i> canonicalisation	2.76	3.14	1.58	2.28	0.93	3.41	0.79	13.24	0.60	5.79	0.47	1.40	0.21	3.18	0.34	1.42	0.96	4.23
<i>w/o</i> latent code	2.77	8.04	1.76	3.70	0.83	1.43	1.00	3.32	0.83	3.72	0.61	3.31	0.15	1.28	0.47	2.72	1.05	3.44
<i>w/o</i> multi-step CD loss	2.87	2.79	1.45	2.21	0.68	1.37	0.75	1.04	0.58	3.29	0.46	1.65	0.14	1.34	0.38	1.48	0.91	1.89
<i>w/o</i> $\mathcal{L}_{code}$	2.52	2.46	1.54	2.14	0.71	1.38	0.74	0.99	0.61	3.42	0.45	1.06	0.12	0.85	0.32	1.62	0.88	1.74
<i>w/ FiLM</i> conditioning [51]	2.62	5.74	1.67	1.93	0.83	1.39	0.92	3.32	0.67	3.04	0.47	2.04	0.13	0.92	0.33	1.35	0.96	2.47
<i>w/ model segmentation</i> *	1.77	2.17	1.21	1.94	0.68	1.32	0.48	0.80	2.01	1.59	0.36	1.55	0.12	0.61	0.27	1.36	0.86	1.42

cessing, yielding a mean CD 0.86 / EPE 1.42 versus 0.86 / 1.47 with ground-truth labels. We further probe robustness to imperfect geometry by adding Gaussian noise and by removing local point neighbourhoods (*holes*) around test regions. GrowFields degrades gracefully in both cases; notably, the holes experiment confirms it learns genuine growth dynamics rather than exploiting nearest-neighbour train–test proximity, with full details in the supplementary material.

Statistical significance. Given the limited number of full-plant sequences, we assess significance at the organ level, considering the results of Table 2 (45 organ sequences in total). A Wilcoxon signed-rank test on per-organ metrics (Shapiro–Wilk $p < 0.001$, indicating non-normal distributions) shows that GrowFields significantly outperforms COAP on both CD and EPE (both $p < 0.001$), and DPF on CD ($p < 0.001$) and EPE ($p = 0.004$).

5 Conclusion

We introduced GrowFields, a compositional neural field framework for modelling 4D plant growth from sparse longitudinal point clouds. By canonicalising plant organs and learning a shared latent-conditioned deformation field, the model captures asynchronous organ development and topology changes within a unified representation. Experiments across multiple plant species show that this formulation improves both geometric reconstruction and organ-level tracking compared to existing dynamic reconstruction methods. Our evaluation protocol based on annotated leaf-tip trajectories further provides a practical benchmark for studying plant growth dynamics. While we focus on plants, the formulation is general and could apply to other dynamic point clouds composed of independently evolving parts, such as 4D cell or vessel growth in the medical domain [68]. We hope this work contributes a useful representation for learning structure-aware models of plant development and enables future research on data-driven plant phenotyping and analysis.

5.1 Limitations and outlook

Our approach has three main limitations. First, because we warp the initial organ point cloud forward without adding points, regions undergoing rapid growth become spatially sparse relative to the ground-truth scans (*e.g.* the leaf at day 7 in Figure 2); a promising direction is adaptive densification analogous to 3DGS [25] and 4DGS [69]. Second, very short organ sequences (≤ 5 training timesteps) leave the velocity field underconstrained, as reflected in the extrapolation experiment and the SoH2 sequence. Third, as noted in section 3, GrowFields does not predict organ emergence; newly observed organs are only incorporated upon first observation. Incorporating priors from L-systems or generative models to anticipate morphogenesis is an interesting avenue for future work [20]. Beyond these limitations, given the increasing availability of 4D plant point-cloud datasets, a further promising direction is to harmonise and combine multiple datasets to learn generative geometric models of 4D plant growth that generalise across species and environmental conditions [2, 15].

Acknowledgements

This work was supported by funding from the *Smart Sustainable Farming Research Program* from the World Food System Center from ETH Zürich.

References

1. Adebola, S., Xie, S., Kim, C.M., Kerr, J., van Marrewijk, B.M., van Vlaardingen, M., van Daalen, T., van Loo, E., Susa Rincon, J.L., Solowjow, E., van Zedde, R., Goldberg, K.: Growsplat: Constructing temporal digital twins of plants with

- gaussian splats. In: 2025 IEEE 21st International Conference on Automation Science and Engineering (CASE). pp. 1766–1773 (2025). <https://doi.org/10.1109/CASE58245.2025.11163998>
2. Ahmed, F., Li, D., Zhao, B., Wang, Z., Huang, J., Li, T., Huang, J., Hou, J., Jobaer, S., Yan, H.: Pepper-4D: Spatiotemporal 3D Pepper Crop Dataset for Phenotyping. *Plants* **15**(4), 599 (Feb 2026). <https://doi.org/10.3390/plants15040599>
 3. An, Q., Li, G., Deng, X., Cui, Y., Guo, Q., Liu, Y., Zhou, L., Wei, L., Zhao, B.: AgriGaussian: A low-cost 3D reconstruction method for high-fidelity plant architecture analysis. *Artificial Intelligence in Agriculture* p. S2589721726000036 (Jan 2026). <https://doi.org/10.1016/j.aiia.2026.01.006>
 4. Cao, A., Johnson, J.: HexPlane: A Fast Representation for Dynamic Scenes. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 130–141. IEEE, Vancouver, BC, Canada (Jun 2023). <https://doi.org/10.1109/CVPR52729.2023.00021>
 5. Cao, J., Tagliasacchi, A., Olson, M., Zhang, H., Su, Z.: Point cloud skeletons via laplacian based contraction. In: 2010 Shape Modeling International Conference. pp. 187–197 (2010). <https://doi.org/10.1109/SMI.2010.25>
 6. Chebrolu, N., Magistri, F., Läbe, T., Stachniss, C.: Registration of spatio-temporal point clouds of plants for phenotyping. *PLOS ONE* **16**(2), e0247243 (Feb 2021). <https://doi.org/10.1371/journal.pone.0247243>
 7. Chen, F., Tsafaris, S.A., Giuffrida, M.V.: GMT: Guided Mask Transformer for Leaf Instance Segmentation. In: 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 1217–1226. IEEE, Tucson, AZ, USA (Feb 2025). <https://doi.org/10.1109/WACV61041.2025.00126>
 8. Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural Ordinary Differential Equations. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 31. Curran Associates, Inc. (2018)
 9. Cheng, T., Zhai, A.J., Chen, E.Z., Zhou, R., Deng, Y., Li, Z., Zhao, K., Shiu, J., Zhao, Q., Xu, Y., Wang, X., Shen, Y., Wang, S., Ainsworth, L., Guan, K., Wang, S.: Demeter: A parametric model of crop plant morphology from the real world. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 28740–28751 (October 2025)
 10. Conn, A., Pedmale, U.V., Chory, J., Navlakha, S.: High-resolution laser scanning reveals plant architectures that reflect universal network design principles. *Cell Systems* **5**(1), 53–62.e3 (2017). <https://doi.org/https://doi.org/10.1016/j.cels.2017.06.017>
 11. Deng, B., Lewis, J.P., Jeruzalski, T., Pons-Moll, G., Hinton, G., Norouzi, M., Tagliasacchi, A.: NASA Neural Articulated Shape Approximation. In: *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII*. pp. 612–628. Springer-Verlag, Berlin, Heidelberg (2020). https://doi.org/10.1007/978-3-030-58571-6_36, event-place: Glasgow, United Kingdom
 12. Deussen, O., Hanrahan, P., Lintermann, B., Mëch, R., Pharr, M., Prusinkiewicz, P.: Realistic modeling and rendering of plant ecosystems. In: *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*. pp. 275–286. SIGGRAPH '98, Association for Computing Machinery, New York, NY, USA (1998). <https://doi.org/10.1145/280814.280898>
 13. Ding, Y., Li, Y., Li, X., Guo, Z., Li, X., Li, Z.: GRPE: High-fidelity 3D Gaussian reconstruction for plant entities. *Computers & Graphics* p. 104277 (Jun 2025). <https://doi.org/10.1016/j.cag.2025.104277>

14. Dong, J., Burnham, J.G., Boots, B., Rains, G., Dellaert, F.: 4D crop monitoring: Spatio-temporal reconstruction for agriculture. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). pp. 3878–3885. IEEE, Singapore, Singapore (May 2017). <https://doi.org/10.1109/ICRA.2017.7989447>
15. Drees, L., Weber, I., Rufwurm, M., Roscher, R.: Time Dependent Image Generation of Plants from Incomplete Sequences with CNN-Transformer. In: Andres, B., Bernard, F., Cremers, D., Frintrop, S., Goldlücke, B., Ihrke, I. (eds.) Pattern Recognition, vol. 13485, pp. 495–510. Springer International Publishing, Cham (2022). https://doi.org/10.1007/978-3-031-16788-1_30, series Title: Lecture Notes in Computer Science
16. Du, R., Ma, Z., Xie, P., He, Y., Cen, H.: PST: Plant segmentation transformer for 3D point clouds of rapeseed plants at the podding stage. *ISPRS Journal of Photogrammetry and Remote Sensing* **195**, 380–392 (Jan 2023). <https://doi.org/10.1016/j.isprsjprs.2022.11.022>
17. Freeman, H., Kantor, G.: Transformer-Based Spatio-Temporal Association of Apple Fruitlets. arXiv:2503.03200 [cs] (Mar 2025), arXiv: 2503.03200
18. Furbank, R.T., Tester, M.: Phenomics – technologies to relieve the phenotyping bottleneck. *Trends in Plant Science* **16**(12), 635–644 (Dec 2011). <https://doi.org/10.1016/j.tplants.2011.09.005>
19. Fusaro, D., Magistri, F., Behley, J., Pretto, A., Stachniss, C.: Horticultural temporal fruit monitoring via 3D instance segmentation and re-identification using colored point clouds. *Computers and Electronics in Agriculture* **247**, 111723 (Jun 2026). <https://doi.org/10.1016/j.compag.2026.111723>
20. Ghrer, S., Godin, C., Wuhler, S.: Learning to infer parameterized representations of plants from 3d scans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 42236–42245 (June 2026)
21. Golla, T., Kneiphof, T., Kuhlmann, H., Weinmann, M., Klein, R.: Temporal Upsampling of Point Cloud Sequences by Optimal Transport for Plant Growth Visualization. *Computer Graphics Forum* **39**(6), 167–179 (Sep 2020). <https://doi.org/10.1111/cgf.14009>
22. Granier, C., Tardieu, F.: Multi-scale phenotyping of leaf expansion in response to environmental changes: the whole is more than the sum of parts. *Plant, Cell & Environment* **32**(9), 1175–1184 (2009). <https://doi.org/https://doi.org/10.1111/j.1365-3040.2009.01955.x>, <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1365-3040.2009.01955.x>
23. Huang, H., Wu, S., Cohen-Or, D., Gong, M., Zhang, H., Li, G., Chen, B.: L_1 -medial skeleton of point cloud. *ACM Transactions on Graphics* **32**(4), 1–8 (Jul 2013). <https://doi.org/10.1145/2461912.2461913>
24. Jiang, L., Sun, J., Chee, P.W., Li, C., Fu, L.: Cotton3DGaussians: Multiview 3D Gaussian Splatting for boll mapping and plant architecture analysis. *Computers and Electronics in Agriculture* **234**, 110293 (Jul 2025). <https://doi.org/10.1016/j.compag.2025.110293>
25. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (July 2023)
26. Khanam, T., Laga, H., Bennamoun, M., Wang, G., Sohel, F., Boussaid, F., Wang, G., Srivastava, A.: A Riemannian Approach for Spatiotemporal Analysis and Generation of 4D Tree-Shaped Structures. In: Computer Vision - ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LXVII. pp. 326–341. Springer-Verlag, Berlin, Heidelberg (2024). https://doi.org/10.1007/978-3-031-72855-6_19, event-place: Milan, Italy

27. Kirchgeßner, N.: Extraktion physiologischer Koordinatensysteme von Pflanzenwurzeln und -blättern aus Bildsequenzen. Dissertation, Universität Heidelberg (2003). <https://doi.org/10.11588/heidok.00003482>
28. Li, D., Ahmed, F., Wang, Z.: 3D-NOD: 3D new organ detection in plant growth by a spatiotemporal point cloud deep segmentation framework. *Plant Phenomics* **7**(1), 100002 (Mar 2025). <https://doi.org/10.1016/j.plaphe.2025.100002>
29. Li, D., Li, J., Xiang, S., Pan, A.: PSegNet: Simultaneous Semantic and Instance Segmentation for Point Clouds of Plants. *Plant Phenomics* **2022**, 9787643 (2022). <https://doi.org/10.34133/2022/9787643>
30. Li, D., Liu, L., Xu, S., Jin, S.: TrackPlant3D: 3D organ growth tracking framework for organ-level dynamic phenotyping. *Computers and Electronics in Agriculture* **226**, 109435 (Nov 2024). <https://doi.org/10.1016/j.compag.2024.109435>
31. Li, J., Qi, X., Nabaei, S.H., Liu, M., Chen, D., Sun, Q., Zhang, X., Yin, X., Li, Z.: A survey on 3D reconstruction techniques in plant phenotyping: From classical methods to Neural Radiance Fields (NeRF), 3D Gaussian Splatting (3DGS), and beyond. *Plant Phenomics* **7**(4), 100137 (Dec 2025). <https://doi.org/10.1016/j.plaphe.2025.100137>
32. Li, J., Song, Z., Yang, B.: NVFi: Neural velocity fields for 3d physics learning from dynamic videos. In: *Thirty-seventh Conference on Neural Information Processing Systems* (2023)
33. Li, Y., Fan, X., Mitra, N.J., Chamovitz, D., Cohen-Or, D., Chen, B.: Analyzing growing plants from 4D point cloud data. *ACM Transactions on Graphics* **32**(6), 1–10 (Nov 2013). <https://doi.org/10.1145/2508363.2508368>
34. Liu, Z., Cheng, Z., Yokoya, N.: Neural Hierarchical Decomposition for Single Image Plant Modeling. In: *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 733–742 (2025). <https://doi.org/10.1109/CVPR52734.2025.00077>
35. Lobefaro, L., Sodano, M., Fusaro, D., Magistri, F., Malladi, M.V.R., Guadagnino, T., Pretto, A., Stachniss, C.: Spatio-Temporal Consistent Semantic Mapping for Robotics Fruit Growth Monitoring. *IEEE Robotics and Automation Letters* **10**(9), 9470–9477 (Sep 2025). <https://doi.org/10.1109/LRA.2025.3594985>
36. Luo, C., Umetani, N.: LeafFit: Plant Assets Creation from 3D Gaussian Splatting. *arXiv:2602.11577 [cs:GR]* (Feb 2026), note: accepted to Eurographics 2026.
37. Luo, W., Goli, L., Bahmani, S., Taubner, F., Tagliasacchi, A., Lindell, D.B.: Grow with the Flow: 4D Reconstruction of Growing Plants with Gaussian Flow Fields. *arXiv:2602.08958 [cs:CV]* (Feb 2026), *arXiv: 2602.08958*
38. 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, Las Vegas, NV, USA (Oct 2020). <https://doi.org/10.1109/IROS45743.2020.9340918>
39. van Marrewijk, B.M., van Daalen, T., Xin, B., van Henten, E.J., Polder, G., Kootstra, G.: 3D plant segmentation: Comparing a 2D-to-3D segmentation method with state-of-the-art 3D segmentation algorithms. *Biosystems Engineering* **254**, 104147 (Jun 2025). <https://doi.org/10.1016/j.biosystemseng.2025.104147>
40. Mertoğlu, K., Şalk, Y., Sarıkaya, S.K., Turgut, K., Evrenosoğlu, Y., Çevikalp, H., Gerek, Ö.N., Dutağacı, H., Rousseau, D.: Planest-3d: A new annotated data set of 3d color point clouds of plants. In: *2023 31st Signal Processing and Communications Applications Conference (SIU)*. pp. 1–4 (2023). <https://doi.org/10.1109/SIU59756.2023.10223838>

41. Meyer, L., Gilson, A., Scholz, O., Stamminger, M.: CherryPicker: Semantic Skeletonization and Topological Reconstruction of Cherry Trees. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 6244–6253. IEEE, Vancouver, BC, Canada (Jun 2023). <https://doi.org/10.1109/CVPRW59228.2023.00664>
42. Mihajlovic, M., Saito, S., Bansal, A., Zollhöfer, M., Tang, S.: COAP: Compositional Articulated Occupancy of People. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13201–13210 (Jun 2022)
43. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Occupancy flow: 4d reconstruction by learning particle dynamics. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5379–5389 (2019)
44. Ojo, T., La, T., Morton, A., Stavness, I.: Splanting: 3D plant capture with gaussian splatting. In: SIGGRAPH Asia 2024 Technical Communications. pp. 1–4. ACM, Tokyo Japan (Dec 2024). <https://doi.org/10.1145/3681758.3698009>
45. Okura, F.: 3D modeling and reconstruction of plants and trees: A cross-cutting review across computer graphics, vision, and plant phenotyping. *Breeding Science* **72**(1), 31–47 (2022). <https://doi.org/10.1270/jsbbs.21074>
46. Pan, H., Hetroy-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, London, United Kingdom (Dec 2021). <https://doi.org/10.1109/3DV53792.2021.00041>
47. Paproki, A., Sirault, X., Berry, S., Furbank, R., Fripp, J.: A novel mesh processing based technique for 3D plant analysis. *BMC Plant Biology* **12**(1), 63 (2012). <https://doi.org/10.1186/1471-2229-12-63>
48. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165–174 (2019)
49. Peng, S., Geng, C., Zhang, Y., Xu, Y., Wang, Q., Shuai, Q., Zhou, X., Bao, H.: Implicit Neural Representations With Structured Latent Codes for Human Body Modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(8), 9895–9907 (2023). <https://doi.org/10.1109/TPAMI.2023.3245815>
50. Peng, S., Zhang, Y., Xu, Y., Wang, Q., Shuai, Q., Bao, H., Zhou, X.: Neural Body: Implicit Neural Representations With Structured Latent Codes for Novel View Synthesis of Dynamic Humans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9054–9063 (Jun 2021)
51. Perez, E., Strub, F., de Vries, H., Dumoulin, V., Courville, A.: FiLM: Visual reasoning with a general conditioning layer. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018). <https://doi.org/10.1609/aaai.v32i1.11671>
52. Prokudin, S., Ma, Q., Raafat, M., Valentin, J., Tang, S.: Dynamic point fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 7964–7976 (October 2023)
53. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10318–10327 (2021)
54. Riccardi, A., Kelly, S., Marks, E., Magistri, F., Guadagnino, T., Behley, J., Bennewitz, M., Stachniss, C.: Fruit Tracking Over Time Using High-Precision Point Clouds. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 9630–9636 (2023). <https://doi.org/10.1109/ICRA48891.2023.10161350>

55. Sang, L., Canfes, Z., Cao, D., Marin, R., Bernard, F., Cremers, D.: 4deform: Neural surface deformation for robust shape interpolation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6542–6551 (June 2025)
56. Schunck, D., Magistri, F., Rosu, R.A., Cornelißen, A., Chebrolu, N., Paulus, S., Léon, J., Behnke, S., Stachniss, C., Kuhlmann, H., Klingbeil, L.: Pheno4D: A spatio-temporal dataset of maize and tomato plant point clouds for phenotyping and advanced plant analysis. PLOS ONE **16**(8), e0256340 (Aug 2021). <https://doi.org/10.1371/journal.pone.0256340>
57. Shen, P., Jing, X., Deng, W., Jia, H., Wu, T.: PlantGaussian: Exploring 3D Gaussian splatting for cross-time, cross-scene, and realistic 3D plant visualization and beyond. The Crop Journal **13**(2), 607–618 (Apr 2025). <https://doi.org/10.1016/j.cj.2025.01.011>
58. Shen, Y., Zhou, H., Yang, X., Lu, X., Guo, Z., Jiang, L., He, Y., Cen, H.: Biomass phenotyping of oilseed rape through UAV multi-view oblique imaging with 3DGS and SAM model. Computers and Electronics in Agriculture **235**, 110320 (Aug 2025). <https://doi.org/10.1016/j.compag.2025.110320>
59. Singh, R., Bidese, R., Dhakal, K., Sornapudi, S.: Few-Shot Adaptation of Grounding DINO for Agricultural Domain. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 5371–5381 (Jun 2025)
60. Sitzmann, V., Martel, J.N.P., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. NIPS '20, Curran Associates Inc., Red Hook, NY, USA (2020)
61. Song, H., Wen, W., Wu, S., Guo, X.: Comprehensive review on 3D point cloud segmentation in plants. Artificial Intelligence in Agriculture p. S2589721725000066 (Jan 2025). <https://doi.org/10.1016/j.aiia.2025.01.006>
62. Stuart, L.A.G., Wells, D.M., Atkinson, J.A., Castle-Green, S., Walker, J., Pound, M.P.: High-fidelity wheat plant reconstruction using 3D Gaussian splatting and neural radiance fields. GigaScience **14**, giaf022 (Jan 2025). <https://doi.org/10.1093/gigascience/giaf022>
63. Su, S.Y., Bagautdinov, T., Rhodin, H.: DANBO: Disentangled Articulated Neural Body Representations via Graph Neural Networks. In: Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II. pp. 107–124. Springer-Verlag, Berlin, Heidelberg (2022). https://doi.org/10.1007/978-3-031-20086-1_7, event-place: Tel Aviv, Israel
64. Sun, D., Huang, H., Li, Y., Gong, X., Ye, Q.: Dsr: Dynamical surface representation as implicit neural networks for protein. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) Advances in Neural Information Processing Systems. vol. 36, pp. 13873–13886. Curran Associates, Inc. (2023)
65. Sun, S., Han, K., Kong, D., Tang, H., Yan, X., Xie, X.: Topology-preserving shape reconstruction and registration via neural diffeomorphic flow. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 20845–20855 (June 2022)
66. Walter, A., Studer, B., Kölliker, R.: Advanced phenotyping offers opportunities for improved breeding of forage and turf species. Annals of Botany **110**(6), 1271–1279 (02 2012). <https://doi.org/10.1093/aob/mcs026>
67. Wang, M., Li, C., Vaxman, A.: CanFields: Consolidating diffeomorphic flows for non-rigid 4d interpolation from arbitrary-length sequences. In: Proceedings of the

- IEEE/CVF International Conference on Computer Vision (ICCV). pp. 28587–28598 (October 2025)
68. Wiesner, D., Suk, J., Dummer, S., Nečasová, T., Ulman, V., Svoboda, D., Wolterink, J.M.: Generative modeling of living cells with $SO(3)$ -equivariant implicit neural representations. *Medical Image Analysis* **91**, 102991 (Jan 2024). <https://doi.org/10.1016/j.media.2023.102991>
 69. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 20310–20320 (June 2024)
 70. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. *Computer Graphics Forum* **41**(2), 641–676 (2022). <https://doi.org/https://doi.org/10.1111/cgf.14505>
 71. Xing, J., Miyakawa, R., Yang, Y., Liu, X., Shinoda, R., Santo, H., Toda, Y., Okura, F.: Zero-shot Hierarchical Plant Segmentation via Foundation Segmentation Models and Text-to-image Attention. *arXiv:2509.09116 [cs:CV]* (Dec 2025), *arXiv: 2509.09116*
 72. Yang, X., Du, R., Huang, H., Xie, J., Xie, P., Fang, L., Guo, Z., Jiang, N., Jiang, Y., Cen, H.: PlantSegNeRF: A few-shot, cross-species method for plant 3D instance point cloud reconstruction via joint-channel NeRF with multi-view image instance matching. *Artificial Intelligence in Agriculture* **16**(1), 546–564 (Mar 2026). <https://doi.org/10.1016/j.aiia.2025.11.009>
 73. Yang, Y., Mao, D., Santo, H., Matsushita, Y., Okura, F.: Neuraleaf: Neural parametric leaf models with shape and deformation disentanglement. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 28167–28176 (October 2025)
 74. Yang, Y., Shinoda, R., Santo, H., Okura, F.: GaussianPlant: Structure-aligned Gaussian Splatting for 3D Reconstruction of Plants. *arXiv:2512.14087 [cs:CV]* (Dec 2025), *arXiv: 2512.14087*
 75. Zhang, D., Gajardo, J., Medic, T., Katircioglu, I., Boss, M., Kirchgessner, N., Walter, A., Roth, L.: Wheat3DGS: In-field 3D Reconstruction, Instance Segmentation and Phenotyping of Wheat Heads with Gaussian Splatting. In: *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR) Workshops*. pp. 5369–5379 (June 2025), comment: Daiwei Zhang and Joaquin Gajardo contributed equally to this work.
 76. 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 (Nov 2023). <https://doi.org/10.1016/j.isprsjprs.2023.10.009>
 77. Zhang, Y., Prokudin, S., Mihajlovic, M., Ma, Q., Tang, S.: Degrees of freedom matter: Inferring dynamics from point trajectories. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 25332–25342 (June 2024)
 78. Zhu, J., Zhai, R., Ren, H., Xie, K., Du, A., He, X., Cui, C., Wang, Y., Ye, J., Wang, J., Jiang, X., Wang, Y., Huang, C., Yang, W.: Crops3D: a diverse 3D crop dataset for realistic perception and segmentation toward agricultural applications. *Scientific Data* **11**(1), 1438 (Dec 2024). <https://doi.org/10.1038/s41597-024-04290-0>