

Grow with the Flow: 4D Reconstruction of Growing Plants with Gaussian Flow Fields

Weihan Luo^{1*} Lily Goli^{1,2} Sherwin Bahmani^{1,2}
Felix Taubner^{1,2} Andrea Tagliasacchi^{1,3} David B. Lindell^{1,2}

¹University of Toronto ²Vector Institute ³Simon Fraser University
weihanluo.ca/growflow/

Abstract. Modeling the time-varying 3D appearance of plants during growth poses unique challenges: unlike most dynamic scenes, plants continuously generate new geometry as they expand, branch, and differentiate. Existing dynamic scene representations are ill-suited to this setting: deformation fields provide insufficient constraints to yield physically plausible scene dynamics, and 4D Gaussian splatting represents the same physical structures with different Gaussian primitives at different times, breaking temporal consistency. We introduce GROWFLOW, a dynamic representation that couples 3D Gaussian primitives with a neural ordinary differential equation to model plant growth as a continuous flow field over geometric parameters (position, scale, and orientation). Our representation enables consistent appearance rendering and models nonlinear, continuous-time growth dynamics with full temporal correspondences for every primitive. To initialize a sufficient set of Gaussian primitives, we first reconstruct the mature plant and then learn a reverse-growth process, effectively simulating the plant’s developmental history in reverse. GROWFLOW achieves superior image quality and geometric coherence compared to prior methods on a new, multi-view timelapse dataset of plant growth, and provides the first temporally coherent representation for appearance modeling of growing 3D structures.

1 Introduction

Accurately modeling plant growth has wide-reaching implications for plant phenotyping, agriculture, and biological research, where understanding the temporal development of plant structures is essential for analyzing morphology, function, and environmental response [8, 19, 33, 36, 37]. Unlike most dynamic scenes, plant growth is inherently non-rigid and involves continuous structural change: new leaves and branches emerge gradually, altering both geometry and topology over time [7, 16, 25, 41, 44]. We address the problem of reconstructing time-varying 3D representations of plant growth from multi-view time-lapse imagery, with a particular focus on capturing temporally coherent geometry throughout development.

Contemporary dynamic scene representations fall broadly into two families, neither of which is well-suited to our problem. Deformation-based methods [46, 48] map a canonical representation to scene structure at each timestep

* weihan262144@outlook.com

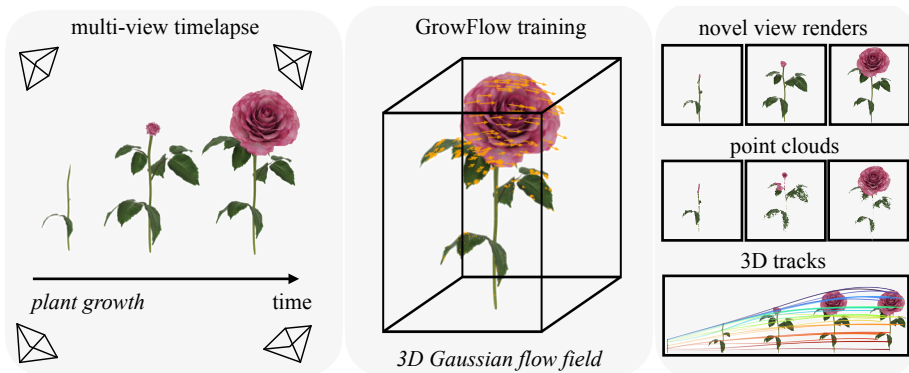


Fig. 1: GROWFLOW. We propose GROWFLOW, a method for reconstructing high-fidelity geometry of plant growth. Given multi-view timelapse images of a plant, our method accurately reconstructs the dynamic structure using a set of 3D Gaussian primitives and a flow field defined over their parameters. Our continuous flow field further enables temporal interpolation of both geometry and appearance between frames. We can also track structures during a plant’s growth by visualizing the positions of the 3D Gaussian primitives, as shown above for the synthetic *rose* plant. Please see the Supp. Webpage for video results.

via a learned deformation field, but impose little constraint on the smoothness or physical plausibility of the field—nothing prevents learning geometrically implausible mappings that merely minimize the photometric loss. Methods based on 4D Gaussians with temporal masking [11, 26, 47] are even less constrained: geometry is discarded and introduced across time with no notion of correspondence. Most closely related to our setting, GrowSplat [1] applies 3D Gaussian Splatting (3DGS) [21] to plant growth, but produces independent per-timestep reconstructions that similarly lack temporal correspondences. In growth modeling, tracking the development of individual leaves and branches over time is as important as rendering quality—and previous work does not meet this requirement.

We propose a new perspective: plant growth can be modeled as a continuous dynamical system, where each scene element follows a smooth trajectory through space and time, governed by an underlying vector field. We parameterize this vector field as a neural ordinary differential equation (ODE) [5], whose integration naturally enforces smooth, continuous evolution, as the Gaussian trajectories are constrained to follow a consistent vector field—providing an inductive bias that unconstrained deformation fields lack.

Building on this insight, we present GROWFLOW, a novel dynamic representation that couples 3D Gaussian primitives with a neural ODE to learn this growth vector field, yielding a temporally coherent and biologically plausible evolution of plant geometry, as shown in Fig. 1. A key challenge in this setting is how to continuously introduce new geometry as the plant grows: directly adding new Gaussians is non-differentiable and hard to optimize. We sidestep this by

reconstructing the mature plant and learning growth in reverse—modeling the plant’s developmental history backwards through time. Concretely, we learn a continuous ODE flow field over the position, scale, and orientation of 3D Gaussian primitives, while keeping color and opacity fixed, then reverse this process to recover a realistic growth trajectory. Because all Gaussians persist throughout the ODE trajectory, each primitive maintains a consistent identity across time, enabling the kind of geometric coherence that existing methods cannot provide. While this restricts GROWFLOW to monotonic growth, where plant structure only accumulates over time, this assumption holds broadly in plant phenotyping and agricultural settings, where GROWFLOW achieves state-of-the-art performance in both novel-view and novel-time synthesis. In summary, we make the following contributions:

- We introduce GROWFLOW, a dynamic scene representation that couples 3D Gaussians with neural ODEs to model the continuous, non-rigid evolution of plant growth from multi-view time-lapse images.
- We propose a reverse-growth formulation that sidesteps non-differentiable topology changes and enables end-to-end training of a continuously evolving scene representation.
- To the best of our knowledge, we present the first multi-view timelapse dataset of real growing plants, comprising three plant species (*blooming flower*, *corn*, and *paperwhite*) recorded using a calibrated single-camera turntable system.

Opportunities for future research. Modeling the dynamic topological changes associated with plant growth is a challenging problem, but research in this direction has strong potential for scientific impact. We therefore overview the limitations of our current formulation alongside the many opportunities it opens for future research. First, GROWFLOW is optimized for monotonic growth scenarios. While we show that the approach performs well on real captured data, extending it to processes involving structural loss, such as leaf senescence or petal drop, is a natural and promising direction. Second, while this work focuses on temporally coherent geometry reconstruction and novel-view synthesis, coupling our representation with explicit trait extraction modules could unlock direct recovery of morphogenetic quantities—such as stem length, leaf area, and branching angles—opening exciting new avenues for automated plant phenotyping and monitoring. To facilitate future work, we will publicly release all code and data.

2 Related Work

Dynamic novel view synthesis. Recent work in dynamic 3D scene modeling has largely shifted from Neural Radiance Fields (NeRFs) [31, 35] to 4D extensions of 3D Gaussian Splatting (3DGS) [21], which offer superior rendering quality and computational efficiency. The most common strategy is to learn a deformation field that maps a single set of canonical Gaussians to their state at each observed timestep [12, 18, 28, 46, 48]. This process is often accelerated using compact and efficient neural representations such as HexPlanes [3, 15]. However,

deformation-based representations learn independent per-timestep deformations from a canonical space; as a result, they do not explicitly introduce new structure or capture the local spatio-temporal dependencies and monotonic growth inherent in plant growth.

Another line of work optimizes 4D spatio-temporal Gaussians to represent the scene’s evolution [11, 26, 47]. A related approach models the continuous trajectory of each Gaussian’s parameters over time, often using simple functions such as polynomials [27, 45]. Finally, some methods adopt a sequential strategy, propagating Gaussian parameters from one frame to the next to enforce temporal consistency [30]. However, these methods often rely on auxiliary inputs (e.g., optical flow or depth) or use masks to remove "inactive" Gaussians, which breaks explicit 3D correspondences between timesteps; sequential methods additionally assume persistent structures, and cannot account for new geometry emerging over time. In contrast, our approach models plant growth as a continuous, temporally coherent 3D Gaussian flow, enabling both the introduction of new structures and accurate prediction of unseen timesteps.

Continuous-time dynamics models. Continuous-time dynamical systems can be mathematically represented as ordinary differential equations (ODEs), where the rate of change of the system state is described as a function of the current state and time. Neural ODEs [6] parametrize the underlying flow field using a neural network and recover the trajectory of the system by integration. Several extensions focus on improving optimization stability [13, 14], computational efficiency [20, 22, 32], or adapting them to irregularly sampled data [17, 38].

Our work is most closely related to methods that model continuous-time dynamics of 3D scenes using neural ODEs. For example, Du et al. [10] learn a velocity field by integrating an ODE over point tracks, but they require dense point correspondences as input. More recently, Wang et al. [43] combined latent ODEs with 3D Gaussians for temporal forecasting; however, their primary goal is motion extrapolation beyond observed trajectories, whereas we introduce a new dynamic 3D Gaussian representation and a multi-stage optimization procedure specifically designed to capture plant growth.

While several prior techniques [1, 4, 9, 29, 34, 50] tackle plant growth reconstruction, these methods rely on point cloud registration rather than modeling continuous-time dynamics with 3D Gaussians, limiting their ability to interpolate between observations and to guarantee smooth trajectories, as our neural ODE representation does.

3 Method

Given a set of posed images I_p^t of a growing plant observed over multiple timesteps $t \in 0, \dots, T$ and multiple views p , our goal is to reconstruct the plant’s growth in 3D such that the reconstruction faithfully follows its natural trajectory. In particular, we seek a representation that evolves smoothly over time while ensuring that the visible volume of the plant is monotonically non-decreasing, consistent with natural growth. To this end, we adopt 3D Gaussian splats [21] as our un-

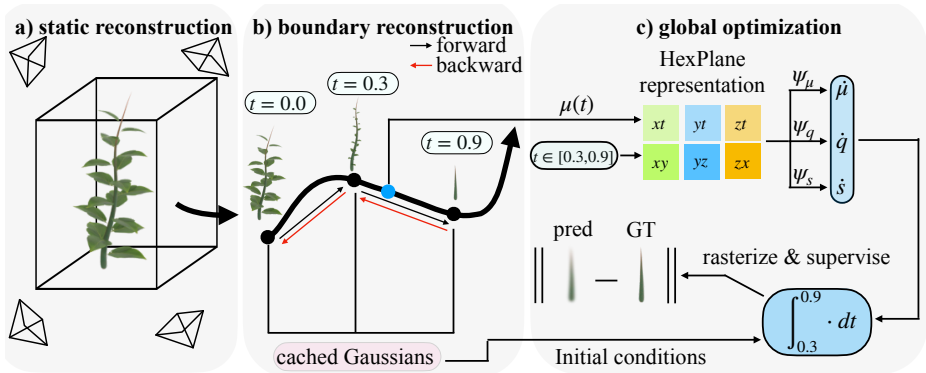


Fig. 2: Method overview. (a) Our method first optimizes a set of 3D Gaussians on the fully-grown plant. (b) Using the optimized 3D Gaussians from the fully-grown plant, we progressively train the dynamics model to learn the state of the plant at each timestep. After each reconstructed timestep, we cache the Gaussians for that timestep and use them as initial conditions to optimize for the next timestep. (c) During the global optimization step, we randomly sample a timestep t_k and integrate to t_{k+1} , leveraging the cached Gaussians from the boundary reconstruction step as initial conditions. We then optimize the dynamics model to enforce consistency between rendered and captured measurements.

derlying 3D representation and optimize a flow field that continuously evolves the Gaussian particles over time to model plant growth. Achieving such smooth temporal evolution is non-trivial: while existing approaches to dynamic 3D reconstruction allow arbitrary deformations either from a canonical template [46] or between discrete timesteps [30], these formulations are not well-suited to modeling growth. Instead, plant growth should evolve continuously from one timestep to the next, following a smooth and monotonic trajectory rather than resetting from a canonical state or diverging unpredictably across timesteps.

To address this challenge, we first introduce a differentiable approach to modeling growth with 3D Gaussian particles in Section 3.1. We then develop a time-integrated neural field that produces a smooth trajectory of growth across all timesteps in Section 3.2. Finally, we present a training strategy that ensures stable optimization in Section 3.3.

3.1 3D Gaussian Flow Fields

We represent the underlying 3D structure using 3D Gaussian Splatting (3DGS) [21], a high-quality representation that enables real-time rendering. Specifically, the 3D scene is modeled with a set of N Gaussians \mathbf{G}_i , each parameterized by a center $\mu_i \in \mathbb{R}^3$, rotation quaternion $q_i \in \mathbb{R}^4$, scale $s_i \in \mathbb{R}^3$, opacity $o_i \in \mathbb{R}$, and color coefficients $c_i \in \mathbb{R}^r$, represented via rank- r spherical harmonics. These Gaussians are projected into a given view using a linearized projection model [51] and then alpha-blended in depth order to render the target image.

To model plant growth, we adapt this representation so that it evolves over time, allowing new structures to emerge gradually and coherently rather than being introduced abruptly. Growth can manifest in two ways: (i) increasing the scale of existing particles, thereby expanding the volume, or (ii) introducing new particles. While scale growth suffices at early stages, it cannot account for the formation of new matter and quickly degrades visual quality without particle addition. Conversely, densification in 3DGS is a discrete, non-differentiable process, making optimization challenging. To address this, we reverse the problem: instead of modeling forward growth, we model backward shrinkage from the final state (time $t=T$) to the initial state ($t=0$). This assumes that all matter required for the plant is already represented at T , eliminating the need for discrete particle addition. The task then reduces to making Gaussians disappear or “shrink” smoothly, either by scaling them down to zero or by becoming occluded within existing matter. This disappearance process is differentiable, making it well-suited for gradient-based optimization. Consequently, the problem reduces to modeling the temporal deformation of Gaussian parameters that govern geometry while keeping appearance fixed. Concretely, we allow the center, rotation, and scale of each Gaussian to evolve over time, while assuming that color and opacity remain constant under fixed lighting conditions. This assumption is practical for our controlled capture setup, though the framework can naturally be extended to model time-varying appearance by including color in the flow field integration. Each Gaussian is thus represented as

$$\mathbf{G}_i^{(t)} = (\mu_i^{(t)}, q_i^{(t)}, s_i^{(t)}, o_i, c_i), \quad (1)$$

where $\mu_i^{(t)}$, $q_i^{(t)}$, and $s_i^{(t)}$ are time-varying geometric parameters, and o_i and c_i are time-invariant appearance parameters.

3.2 Time-Integrated Velocity Field

Our goal is to obtain a smooth trajectory of growth by continuously deforming the geometry of Gaussians as they shrink backward in time. To this end, we model the velocities of Gaussian geometric parameters: translational velocity $\dot{\mu}_i(t)$, rotational velocity $\dot{q}_i(t)$, and volumetric velocity $\dot{s}_i(t)$. We define a time-dependent velocity field F_ϕ governing the dynamics of each Gaussian:

$$\dot{\theta}_i(t) = F_\phi(\mu_i(t), t), \quad \theta_i(t) = \theta_i(T) + \int_T^t F_\phi(\mu_i(\tau), \tau) d\tau, \quad (2)$$

where $\theta_i(t)$ denotes the geometric parameters of Gaussian i at time t . We require F_ϕ to be at least C^0 -continuous in both space and time. This guarantees that integrating the velocity field produces C^1 -continuous trajectories, yielding smooth temporal evolution of centers, rotations, and scales. This design avoids sudden or unpredictable changes between timesteps, ensuring that the reconstructed plant evolves along smooth and differentiable trajectories. We model the velocity field F_ϕ using a spatio-temporal HexPlane encoder followed by multi-layer perceptron (MLP) decoders, similar to [3, 46], as shown in Fig. 2. The HexPlane encoder

interpolates features from a continuous spatio-temporal grid, which are then decoded by MLP heads into the geometric velocities. Formally, given Gaussian centers $\mu_i(t)$ and time t , we extract a latent feature \mathbf{z}_i via:

$$\mathbf{z}_i = \psi(\text{HexInterp}(\mu_i(t), t)), \quad (3)$$

where HexInterp denotes interpolation from a multi-level HexPlane grid. Features are bilinearly interpolated from the six spatio-temporal planes (x, y) , (y, z) , (x, z) , (x, t) , (y, t) , (z, t) , combined via a product across planes, and concatenated across L resolution levels before being fed to the MLP ψ . The latent feature \mathbf{z}_i is then decoded into per-parameter velocities:

$$\dot{\mu}_i = \psi_\mu(\mathbf{z}_i), \quad \dot{q}_i = \psi_q(\mathbf{z}_i), \quad \dot{s}_i = \psi_s(\mathbf{z}_i), \quad (4)$$

where ψ_μ , ψ_q , and ψ_s are independent MLP decoders. To recover Gaussian parameters at any future time t_1 from an initial state, we integrate velocity:

$$\theta_i(t_1) = \theta_i(t_0) + \int_{t_0}^{t_1} F_\phi(\mu_i(t), t) dt, \quad (5)$$

which can be solved using standard ODE solvers such as Runge–Kutta [2, 24, 39].

3.3 Training Dynamics

Static reconstruction. We first optimize a static 3DGS model on the fully-grown plant at timestep T , following standard procedure as in [21], optimizing a mixture of L1 and SSIM losses. After optimization, we obtain a set of Gaussians $\mathbf{G}^{t_0} = \{\mu^{t_0}, q^{t_0}, s^{t_0}, c, o\}$.

Boundary reconstruction. In principle, integrating from $t_0 = T$ backward to all timesteps could produce the entire trajectory. However, directly optimizing such long-range ODE integration leads to unstable training, with vanishing gradients and accumulated numerical error. To address this, we adopt a piecewise integration strategy: instead of integrating across the full sequence, we train progressively from T to earlier steps t_1, t_2, \dots , caching intermediate states as boundary conditions. At each stage, the Gaussian state from the previous boundary condition \mathbf{G}^{t_k} serves as the initial condition, and we integrate the velocity field through a single timestep to obtain $\mathbf{G}^{t_{k+1}}$:

$$\mathbf{G}^{t_{k+1}} = \mathbf{G}^{t_k} + \int_{t_k}^{t_{k+1}} F_\phi(\mu(t), t) dt. \quad (6)$$

This reduces the depth of recursive integration, stabilizes optimization, and ensures that each segment remains well-conditioned. Importantly, although integration is performed in a piecewise manner, the velocity field F_θ is shared across all segments, which guarantees continuity of the underlying dynamics. At each timestep, we supervise the predicted boundary state with an L1 loss against the ground-truth images of that timestep, and progressively expand the cache of boundary states as training proceeds.

Global optimization. After recovering and storing all boundary states in the cache, we perform a global optimization of the trajectory. At each iteration, we randomly sample a timestep t_k and integrate the velocity field between t_k and t_{k+1} using the cached boundary \mathbf{G}^{t_k} as the initial condition:

$$\tilde{\mathbf{G}}^{t_{k+1}} = \mathbf{G}^{t_k} + \int_{t_k}^{t_{k+1}} F_\phi(\mu(t), t) dt. \quad (7)$$

The predicted Gaussians $\tilde{\mathbf{G}}^{t_{k+1}}$ are then rasterized and supervised against the ground truth images at timestep t_{k+1} using an L1 penalty between the rendered and ground-truth pixel values.

4 Multi-View Plant Growth Dataset

Simulated dataset. We construct a simulated multiview timelapse dataset in Blender by porting seven distinct plant-growth scenes—*clemtis*, *tulip*, *plant1*, *plant2*, *plant3*, *plant4*, and *plant5*—originally created by artists on Blender Market. For each scene, we render 70 timesteps of growth from 34 camera viewpoints uniformly distributed along a full 360° orbit around the plant, at a resolution of 400×400 . This synthetic setup provides full control over geometry, materials, and lighting, enabling quantitative evaluation of reconstruction accuracy. For the spatial split, we use 31 views for reconstruction and 3 held-out views for novel-view evaluation at each timestep. For evaluation, we train on every 6th timestep (12 training timesteps, 372 training images per scene) and evaluate across 69 of 70 timesteps, of which 58 are unseen during training.

Captured dataset. Our captured dataset consists of three plant scenes—*blooming flower*, *corn*, and *paperwhite*—captured with a Raspberry Pi HQ camera [42] (Fig. 3). The three species were chosen to represent a diverse range of growth patterns and temporal scales, with each sequence focused on the most dynamic phase of development: the *blooming flower* undergoes rapid petal expansion, *corn* exhibits strong vertical elongation and leaf splitting, and *paperwhite* displays complex branching with multiple structures emerging simultaneously. Plants are placed on a motorized turntable; at each timestep, we capture 50 images at fixed elevation with 7.2° angular spacing, yielding full 360° coverage. We use 43 views for reconstruction and 7 held-out views for novel-view evaluation at each timestep. Images are captured at a resolution of 1200×1200 . Capture frequency is adapted to each species’ growth rate: for the *blooming flower*, we capture every 15 minutes for 86 timesteps (4,300 total im-




Fig. 3: Multi-view timelapse capture setup. A Raspberry Pi-controlled turntable and camera autonomously capture multi-view images of the plant over multiple weeks.

ages); for *corn*, every hour for 71 timesteps (3,550 total images); and for *paperwhite*, every hour for 50 timesteps (2,500 total images). For evaluation, we train on a sparse subset of timesteps and evaluate across the full sequence. For *blooming flower*, *corn*, and *paperwhite*, we train on every 17th, 10th, and 7th timestep respectively (6, 8, and 8 training timesteps; 258, 344, and 344 training images), evaluating on all 86, 71, and 50 timesteps, of which 80, 63, and 42 are unseen.

To get poses for training, we run COLMAP [40] on all images of the first timestep and propagate them to the other timesteps as the viewpoints are the same throughout.

5 Experiments

Implementation details. For static reconstructions of fully grown plants, we use 3DGS with default training settings and the Adam [23] optimizer, training each model for 30K iterations. During the boundary reconstruction phase, we optimize each boundary timestep for 300 iterations using the adjoint method [6], with relative and absolute tolerances of 10^{-4} and 10^{-5} , respectively, for the neural ODE solver. The dynamic reconstruction phase uses the same solver configuration and is trained for 30K iterations.

Baselines. We compare our method against state-of-the-art methods in dynamic reconstruction: Dynamic 3DGS [30], 4D-GS [46], and 4DGS [47]. For all results, we use the corresponding open source implementations of these methods. For timestep interpolation, our method, 4D-GS, and 4DGS inherently support querying intermediate timesteps. For Dynamic 3DGS, which does not natively support continuous time, we perform interpolation between learned timesteps by fitting a third-degree polynomial to the Gaussian centers and colors. Rotations are interpolated using spherical linear interpolation (slerp), while scales and opacities are kept fixed, consistent with the original implementation.

Metrics. We employ two complementary measures to evaluate reconstruction methods. Since our goal is to recover geometrically faithful growth rather than only achieving photometric accuracy, we introduce a geometric accuracy metric based on Chamfer Distance (CD). We track foreground Gaussians by matching each to its nearest vertex on the ground-truth plant mesh at the first timestep. Per-timestep Chamfer Distance is then computed between these foreground Gaussians and their corresponding mesh vertices, averaged across time. For 4DGS, we apply their temporal masking before computing distances. In addition, we evaluate the photometric quality of test views using standard image-based metrics: PSNR, LPIPS, and SSIM.

5.1 Simulated Results

Qualitative comparisons. Fig. 4 presents qualitative and quantitative comparisons against baseline methods for plant-growth reconstruction. Our method

Method	Training times				Interpolation times				Combined			
	PSNR (dB)↑	SSIM↑	LPIPS↓	CD↓	PSNR (dB)↑	SSIM↑	LPIPS↓	CD↓	PSNR (dB)↑	SSIM↑	LPIPS↓	CD↓
4D-GS	33.04	0.946	0.094	0.73	32.77	0.944	0.094	0.78	32.81	0.944	0.094	0.77
4DGS	30.19	0.939	0.107	12.00	29.11	0.905	0.145	11.95	29.29	0.910	0.138	11.96
Dynamic 3DGS	32.48	0.912	0.154	13.18	32.03	0.908	0.158	13.64	32.11	0.909	0.157	13.56
Ours	35.43	0.957	0.065	0.10	34.93	0.955	0.066	0.11	35.02	0.956	0.066	0.11

Fig. 4: Results on synthetic data. We compare results on both seen and interpolated times averaged over synthetic scenes. GrowFlow achieves stable geometry, unlike prior methods that show visually correct renderings for training frames but struggle on interpolation frames. Yellow marks interpolated frames, and ↓ next to a metric indicates that a lower value is better. Please see the Supp. Webpage for video results.

yields geometrically coherent trajectories: Gaussian centers closely follow the plant’s true surface over time and produce high-quality novel-view renderings. In contrast, baseline approaches exhibit pronounced geometric drift, with Gaussian centers gradually detaching from the plant surface or floating in space as time progresses. Dynamic 3DGS [30] and 4D-GS [46] frequently displace Gaus-

Method	Training times			Interpolation times			Combined		
	PSNR (dB) ↑	SSIM ↑	LPIPS ↓	PSNR (dB) ↑	SSIM ↑	LPIPS ↓	PSNR (dB) ↑	SSIM ↑	LPIPS ↓
4D-GS	30.51	0.989	0.036	24.70	0.977	0.045	25.36	0.978	0.044
4DGS	31.25	0.991	0.031	26.86	0.979	0.046	27.37	0.981	0.044
Dynamic 3DGS	27.49	0.981	0.049	23.86	0.960	0.075	24.27	0.963	0.072
Ours	28.80	0.987	0.031	27.28	0.984	0.033	27.47	0.984	0.033

Fig. 5: Results on captured data. We compare results on both seen (“training”) and interpolated times averaged over all captured scenes. GrowFlow achieves stable, coherent geometry, unlike prior methods that struggle with renderings and reconstructed geometry on the interpolated frames. Yellow marks interpolated frames, and ↓ next to a metric indicates that a lower value is better. Please see the Supp. Webpage for video results.

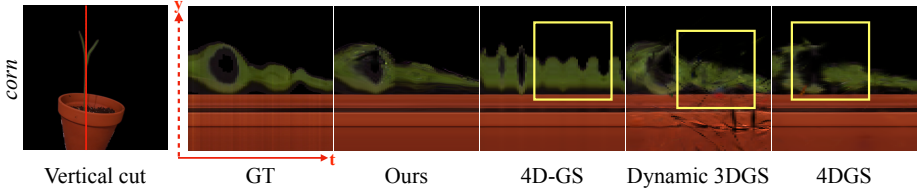


Fig. 6: Temporal slice visualization. We analyze the accuracy of reconstructed motion by tracking a vertical cut from the predicted images of the *corn* scene through time. Our method shows more faithful alignment with GT, while baselines exhibit noisy temporal dynamics (yellow boxes).

sians corresponding to shrunken or disappearing structures into the far field or behind background elements, rather than shrinking them downward as the plant regresses. As illustrated in Fig. 4, these Gaussians often remain at roughly their original height but are simply pushed behind the scene, making them invisible in the renderings. Furthermore, 4DGS [47] leverages different Gaussians to model different frames separately, limiting its ability to track the same set of Gaussians throughout time.

These behaviors highlight a key limitation of approaches that do not explicitly model continuous growth: they prioritize reproducing photorealistic appearance in training views at the expense of temporally coherent geometry. Our representation optimizes a smooth flow field over Gaussian parameters, allowing superior novel view synthesis capabilities, but most importantly, reconstructing physically plausible growth.

Quantitative comparisons. Quantitatively, our approach outperforms all baselines by a substantial margin in both image-quality metrics and Chamfer Distance. This demonstrates that GROWFLOW achieves superior geometric fidelity and photometric consistency not only at supervised training timesteps but also at the 58 interpolated timesteps unseen during training.

5.2 Captured Results

Qualitative comparisons. Fig. 5 presents qualitative and quantitative comparisons against baseline methods on the *blooming flower* and *paperwhite* scenes. While baselines render novel views at training timesteps well, their quality degrades when rendering novel views at interpolated timesteps. 4D-GS [46] fails most notably during interpolation: rather than producing smooth shrinkage, the reconstructed plant oscillates between growing and shrinking. Dynamic 3DGS [30] assumes fixed Gaussian sizes over time and thus cannot model the shrinking plant; it instead turns affected Gaussians black to match the background, minimizing photometric loss at the cost of physical plausibility. In contrast, our method produces temporally smooth and physically plausible interpolations throughout.

Quantitative comparisons. We omit the Chamfer Distance calculation as we do not have ground-truth mesh for the captured data. Overall, our method achieves higher quality novel view renderings compared to baseline methods. Despite achieving slightly lower PSNR and SSIM on the training timesteps, our LPIPS is comparable to baselines. Because our neural ODE optimizes for a continuous flow field of Gaussian parameters rather than overfitting to individual training timesteps, it trades slightly lower performance on training timesteps for superior interpolation quality on real-world plants. Nonetheless, it produces more plausible growth geometry versus baselines.

Temporal slice visualization. To further evaluate motion accuracy, Fig. 6 visualizes a tracked horizontal slice of the plant across timesteps in a novel rendered viewpoint for the *corn* scene. Our method closely matches the ground-truth motion, whereas baselines exhibit significant structural distortions and temporal misalignment.

5.3 Ablation Study

HexPlane. Neural ODE frameworks are often parameterized using MLPs. However, as shown in the insets of Fig. 7, substituting our spatio-temporal HexPlane encoder with an MLP leads to noticeably degraded reconstruction quality, e.g., the flower bud exhibits more artifacts and temporal instability. HexPlane provides a higher-quality inductive bias for capturing spatial and temporal variations, enabling smoother and more consistent Gaussian trajectories. The quantitative results in Tab. 1 confirm this, i.e., the HexPlane achieves superior image fidelity and improves geometric accuracy compared to the MLP alternative.

Table 1: Ablation on the *clematis* scene.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	CD \downarrow
Ours	33.05	0.947	0.071	0.02
w/o HexPlane	32.18	0.944	0.076	0.03
w/o boundary	28.52	0.914	0.097	36.47

Boundary reconstruction. The boundary reconstruction stage is essential for stable optimization of the neural ODE. Without it, the model must rely on long-range integration from the final timestep to all earlier states, which leads to accumulated numerical errors, vanishing gradients, and poor convergence. Although the model can eventually produce reasonable photometric reconstructions, it struggles to maintain geometric consistency, resulting in drifting Gaussians and degraded temporal coherence. As shown in Fig. 7 and Tab. 1, removing the boundary reconstruction step substantially harms both image quality and geometric fidelity, highlighting its importance in accurately modeling continuous plant growth.

6 Conclusion

In this work, we propose GROWFLOW, the first continuous dynamic 3D representation for plant growth, combining 3D Gaussians with neural ODEs to model the non-rigid evolution of plant growth from multi-view time-lapse images. By

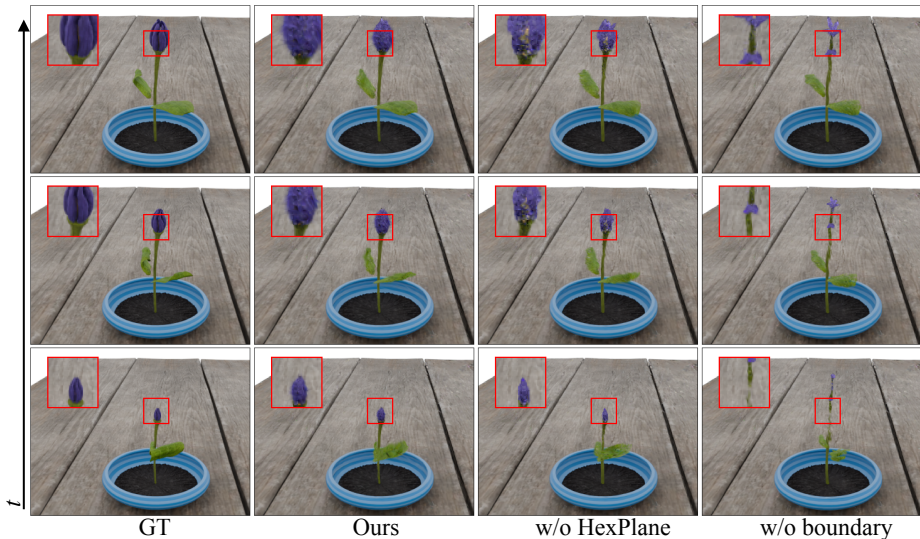


Fig. 7: Qualitative ablations. Replacing our HexPlane representation with an MLP with Fourier encodings reduces capacity and degrades rendering quality. Skipping the boundary reconstruction stage causes the reconstructed geometry to break down.

learning a continuous 3D Gaussian flow field, GROWFLOW captures the underlying growth vector field, enabling temporally coherent reconstruction of plant geometry. To address the challenge of continuously emerging structures, we introduce a reverse-growth formulation, training the model to shrink 3D Gaussians over time and later reversing this flow to recover realistic growth trajectories. We validate our method on both synthetic scenes and a real-world captured dataset of three plant species — *blooming flower*, *corn*, and *paperwhite* — recorded with a calibrated single-camera turntable system, demonstrating superior geometric accuracy and photometric quality compared to existing baselines.

GROWFLOW is designed under the assumption of monotonic growth, which is directly relevant to many plant phenotyping and agricultural applications, and is practical for species exhibiting predominantly additive growth. We view this as a natural starting point for this problem, and encourage future work to relax this assumption to handle non-monotonic phenomena such as leaf senescence and pruning. Other promising directions include incorporating biologically motivated priors and extending the framework to other dynamic objects whose geometry emerges over time, e.g., growing crystals, developing embryos, or erupting geological formations.

Acknowledgements. DBL acknowledges support of NSERC under the RG-PIN program. DBL also acknowledges support from the Canada Foundation for Innovation and the Ontario Research Fund.

References

1. Adebola, S., Xie, S., Kim, C.M., Kerr, J., van Marrewijk, B.M., van Vlaardingen, M., van Daalen, T., van Loo, E., Rincon, J.L.S., Solowjow, E., et al.: GrowSplat: Constructing temporal digital twins of plants with Gaussian splats. In: Proc. CASE (2025)
2. Butcher, J.C.: A history of Runge-Kutta methods. *Applied numerical mathematics* **20**(3), 247–260 (1996)
3. Cao, A., Johnson, J.: Hexplane: A fast representation for dynamic scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 130–141 (2023)
4. Chebrolu, N., Låbe, 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)
5. Chen, R.T.: torchdiffeq, 2018. URL <https://github.com/rtqichen/torchdiffeq> **14** (2018)
6. Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. *Adv. Neural Inform. Process. Syst.* **31** (2018)
7. Coen, E., Cosgrove, D.J.: The mechanics of plant morphogenesis. *Science* **379**(6631), eade8055 (2023)
8. Dhondt, S., Wuyts, N., Inzé, D.: Cell to whole-plant phenotyping: the best is yet to come. *Trends in plant science* **18**(8), 428–439 (2013)
9. 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 (2017)
10. Du, Y., Zhang, Y., Yu, H.X., Tenenbaum, J.B., Wu, J.: Neural radiance flow for 4d view synthesis and video processing. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 14304–14314. IEEE Computer Society (2021)
11. Duan, Y., Wei, F., Dai, Q., He, Y., Chen, W., Chen, B.: 4d-rotor gaussian splatting: towards efficient novel view synthesis for dynamic scenes. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–11 (2024)
12. Duisterhof, B.P., Mandi, Z., Yao, Y., Liu, J.W., Seidenschwarz, J., Shou, M.Z., Ramanan, D., Song, S., Birchfield, S., Wen, B., et al.: Deformgs: Scene flow in highly deformable scenes for deformable object manipulation. arXiv preprint arXiv:2312.00583 (2023)
13. Dupont, E., Doucet, A., Teh, Y.W.: Augmented neural odes. *Advances in neural information processing systems* **32** (2019)
14. Finlay, C., Jacobsen, J.H., Nurbekyan, L., Oberman, A.: How to train your neural ODE: the world of Jacobian and kinetic regularization. In: International conference on machine learning. pp. 3154–3164. PMLR (2020)
15. Fridovich-Keil, S., Meanti, G., Warburg, F.R., Recht, B., Kanazawa, A.: K-planes: Explicit radiance fields in space, time, and appearance. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12479–12488 (2023)

16. Geng, C., Zhang, Y., Wu, S., Wu, J.: Birth and death of a rose. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 26102–26113 (2025)
17. Goyal, P., Benner, P.: Neural odes with irregular and noisy data. arXiv preprint arXiv:2205.09479 (2022)
18. Huang, Y.H., Sun, Y.T., Yang, Z., Lyu, X., Cao, Y.P., Qi, X.: Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4220–4230 (2024)
19. Ijiri, T., Yoshizawa, S., Yokota, H., Igarashi, T.: Flower modeling via x-ray computed tomography. *ACM Transactions on Graphics (TOG)* **33**(4), 1–10 (2014)
20. Kelly, J., Bettencourt, J., Johnson, M.J., Duvenaud, D.K.: Learning differential equations that are easy to solve. *Advances in Neural Information Processing Systems* **33**, 4370–4380 (2020)
21. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 139–1 (2023)
22. Kidger, P., Chen, R.T., Lyons, T.J.: "hey, that's not an ode": Faster ode adjoints via seminorms. In: ICML. pp. 5443–5452 (2021)
23. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
24. Kutta, M.W.: Beitrag zur näherungsweise integration totaler differentialgleichungen. *Zeitschrift für Mathematik und Physik* **46**, 435–453 (1901)
25. 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 (TOG)* **32**(6), 1–10 (2013)
26. Li, Z., Chen, Z., Li, Z., Xu, Y.: Spacetime gaussian feature splatting for real-time dynamic view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8508–8520 (2024)
27. Lin, Y., Dai, Z., Zhu, S., Yao, Y.: Gaussian-flow: 4d reconstruction with dynamic 3d gaussian particle. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21136–21145 (2024)
28. Liu, I., Su, H., Wang, X.: Dynamic gaussians mesh: Consistent mesh reconstruction from dynamic scenes. *ICLR* **5**, 6 (2025)
29. Lobefaro, L., Malladi, M.V., Guadagnino, T., Stachniss, C.: Spatio-temporal consistent mapping of growing plants for agricultural robots in the wild. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 6375–6382. IEEE (2024)
30. Luiten, J., Kopanas, G., Leibe, B., Ramanan, D.: Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In: 2024 International Conference on 3D Vision (3DV). pp. 800–809. IEEE (2024)
31. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
32. Norcliffe, A., Deisenroth, M.P.: Faster training of neural odes using gaussian-legendre quadrature. arXiv preprint arXiv:2308.10644 (2023)
33. Owens, A., Cieslak, M., Hart, J., Classen-Bockhoff, R., Prusinkiewicz, P.: Modeling dense inflorescences. *ACM Transactions on Graphics (TOG)* **35**(4), 1–14 (2016)
34. Pan, H., Hétyroy-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)

35. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5865–5874 (2021)
36. Pound, M.P., Atkinson, J.A., Townsend, A.J., Wilson, M.H., Griffiths, M., Jackson, A.S., Bulat, A., Tzimiropoulos, G., Wells, D.M., Murchie, E.H., et al.: Deep machine learning provides state-of-the-art performance in image-based plant phenotyping. *Gigascience* **6**(10), gix083 (2017)
37. Rincón, M.G., Mendez, D., Colorado, J.D.: Four-dimensional plant phenotyping model integrating low-density lidar data and multispectral images. *Remote Sensing* **14**(2), 356 (2022)
38. Rubanova, Y., Chen, R.T., Duvenaud, D.K.: Latent ordinary differential equations for irregularly-sampled time series. *Advances in neural information processing systems* **32** (2019)
39. Runge, C.: Über die numerische auflösung von differentialgleichungen. *Mathematische Annalen* **46**, 167–178 (1895). <https://doi.org/10.1007/BF01446807>
40. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4104–4113 (2016)
41. Sinnott, E.W.: Plant morphogenesis. (1960)
42. Upton, E., Halfacree, G.: Raspberry Pi user guide. John Wiley & Sons (2016)
43. Wang, D., Rim, P., Tian, T., Wong, A., Sundaramoorthi, G.: Ode-gs: Latent odes for dynamic scene extrapolation with 3d gaussian splatting. arXiv preprint arXiv:2506.05480 (2025)
44. Wang, H., Zhang, B., Klein, J., Michels, D.L., Yan, D.M., Wonka, P.: Autoregressive generation of static and growing trees. In: Proceedings of the SIGGRAPH Asia 2025 Conference Papers. pp. 1–12 (2025)
45. Wang, Q., Ye, V., Gao, H., Austin, J., Li, Z., Kanazawa, A.: Shape of motion: 4d reconstruction from a single video. arXiv preprint arXiv:2407.13764 (2024)
46. 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. pp. 20310–20320 (2024)
47. Yang, Z., Pan, Z., Zhu, X., Zhang, L., Jiang, Y.G., Torr, P.H.: 4d gaussian splatting: Modeling dynamic scenes with native 4d primitives. arXiv preprint arXiv:2412.20720 (2024)
48. Yang, Z., Gao, X., Zhou, W., Jiao, S., Zhang, Y., Jin, X.: Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20331–20341 (2024)
49. Ye, V., Li, R., Kerr, J., Turkulainen, M., Yi, B., Pan, Z., Seiskari, O., Ye, J., Hu, J., Tancik, M., et al.: gsplat: An open-source library for gaussian splatting. *Journal of Machine Learning Research* **26**(34), 1–17 (2025)
50. Zheng, Q., Fan, X., Gong, M., Sharf, A., Deussen, O., Huang, H.: 4d reconstruction of blooming flowers. In: *Computer Graphics Forum*. vol. 36, pp. 405–417. Wiley Online Library (2017)
51. Zwicker, M., Pfister, H., Van Baar, J., Gross, M.: Ewa volume splatting. In: Proceedings Visualization, 2001. VIS’01. pp. 29–538. IEEE (2001)

Supplementary Material

A Video Results

We include an extensive set of results in the Supp. Webpage. There, we show novel view and geometry comparisons against baseline methods on synthetic and captured data. We further show the produced flow field from our trained model.

B Implementation Details

In this section, we provide a detailed description of the network architecture. We implement our dynamic Gaussian representation using the open-sourced Gaussian Splatting implementation `gsplat` [49] and the neural ODE codebase `torchdiffeq` [5]. Our HexPlane architecture follows closely [3, 46], where the spatial resolutions are set to 64 and the temporal resolution is set to 25, which are upsampled by 2. The learning rate of the HexPlane is set to 1.6×10^{-3} , and the learning rate of the MLP decoder is set to 1.6×10^{-4} , both of which are exponentially decayed by a factor of 0.1 until the end of training, for 30K iterations. Unlike [46], we omit the total variation loss, as it does not bring additional improvement. We use a batch size of 30 viewpoints for both our boundary reconstruction stage and dynamic optimization stage, but keep the temporal batch size to 1. The MLP decoders consist of a two-layer MLP with 64 units and a ReLU activation function.

After static reconstruction, we fix the background Gaussians and optimize only the foreground Gaussians within a manually defined bounding box. This constrains the neural ODE to modeling foreground flow, greatly easing optimization.

C Dataset

Extra details of all the simulated and captured datasets can be found in Table S1 and Table S2.

Hardware. Our setup consists of a Raspberry Pi 5 (16GB) with an Active Cooler, powered by a 27W USB-C supply, and an HQ Camera CS with a 6mm wide-angle lens connected via a 300mm cable and stabilized on a tripod. The Pi sends commands to a programmable motorized turntable (ComXim) to rotate the plant and triggers the camera to capture images at each position. To prevent plants wobbling during captures, we set the velocity of the turntable to be the lowest and wait a few seconds after rotations before doing a capture. A pseudo-code of the capture process is illustrated 1.

Algorithm 1 Real Data Collection for GROWFLOW

```

1: Hardware components: Raspberry Pi, HQ Camera CS, motorized turntable.
2: Set turntable velocity to be lowest.
3: for  $t = 1$  to  $n_{\text{timesteps}}$  do
4:   for  $p = 1$  to  $n_{\text{views}}$  do
5:     Send rotation command to turntable: rotate by  $\frac{360}{n_{\text{views}}}$  degrees
6:     Wait for turntable to stabilize
7:     Trigger camera to capture image  $I_p^t$ 
8:   end for
9:   Wait until next timestep
10: end for
11: Output: Multi-view image set  $\{I_p^t\}$  for all timesteps  $t$  and views  $p$ 

```

Table S1: Descriptions of simulated scenes. All scenes sit in a blue vase on top of a wooden table.

	Scene description
Clematis	A purple clematis flower with six-pointed petals and yellow-white stamens at the center, growing on a thin green stem with small leaves.
Tulip	A pink tulip with partially open petals, showing a lighter pink/white interior, on a green stem with two long tulip-shaped leaves.
Plant1	A small, green, young seedling with thin stems and small jagged leaves.
Plant2	A young plant with a tall, dark central stem and several branches bearing smooth, rounded green leaves at varying heights.
Plant3	A tall, slender plant with a pink-red stem and narrow, elongated leaves in shades of pink and purple, arranged sparsely along the stem.
Plant4	A young, slender seedling with a single upright stem and several pairs of bright green, oval leaves arranged oppositely along the stalk.
Plant5	A small plant with distinctively split and fenestrated dark green leaves, resembling a juvenile <i>Monstera deliciosa</i> , with multiple leaves branching from a central stem.

D Additional Results

D.1 Synthetic Results

Tables S3, S4, S5, S6 provide a breakdown of the quantitative results in simulation across all scenes. Overall, our method achieves state-of-the-art performance across all scenes compared to baselines. Please refer to the Supp. Webpage for additional video results and comparisons to baselines.

D.2 Captured Results

Tables S7, S8, S9 provide a breakdown of the quantitative results across all captured scenes. Furthermore, Figure S1 compares the reconstructed *corn* scene across all baselines. Consistent with the results in the main text, our method

Table S2: Descriptions of captured scenes. The growth time refers to the total duration from planting to the end of the capture period.

	Scene description	Growth time
Paperwhite	A paperwhite narcissus with several brown bulbs sitting in a shallow white pot filled with gravel. Multiple green shoots are emerging from the bulbs, with one tall, slender stem extending prominently upward, topped with a closed bud.	one month
Corn	A corn seedling with two narrow, upright green leaves forming a V-shape, planted in a terracotta pot filled with dark soil.	three weeks
Blooming flower	A vibrant pink-red flower in full bloom with wide, open petals, growing from a bulbous beige ceramic vase-shaped pot sitting on an orange saucer.	one week

Table S3: PSNR (dB) results across different synthetic scenes for combined (training + interpolation) frames.

Method	Clematis	Plant1	Plant2	Plant3	Plant4	Plant5	Tulip	Average
4D-GS	31.10	34.11	33.11	32.98	34.30	32.16	31.94	32.81
4DGS	27.62	29.78	29.24	29.73	30.06	29.50	29.13	29.29
Dynamic 3DGS	30.56	33.64	31.46	32.64	33.80	31.58	31.07	32.11
Proposed	33.05	38.12	32.73	35.50	37.54	33.30	34.90	35.02

reconstructs more accurate novel view renders and plant geometry over the training and interpolated timesteps. Please refer to the Supp. Webpage for additional video results and comparisons to baselines.

E GROWFLOW Training Algorithm

We begin with a detailed outline of the training algorithm of our pipeline in Algorithm 2. The first phase is the static reconstruction stage, where we optimize a set of 3D Gaussians on posed images of the fully grown plant. By the end, we

Table S4: SSIM results across different synthetic scenes for combined (training + interpolation) frames.

Method	Clematis	Plant1	Plant2	Plant3	Plant4	Plant5	Tulip	Average
4D-GS	0.933	0.952	0.948	0.946	0.951	0.942	0.939	0.944
4DGS	0.887	0.922	0.911	0.914	0.921	0.910	0.908	0.910
Dynamic 3DGS	0.900	0.922	0.903	0.913	0.920	0.905	0.901	0.909
Proposed	0.947	0.968	0.943	0.963	0.966	0.941	0.962	0.956

Table S5: LPIPS results across different synthetic scenes for combined (training + interpolation) frames.

Method	Clematis	Plant1	Plant2	Plant3	Plant4	Plant5	Tulip	Average
4D-GS	0.102	0.087	0.095	0.095	0.089	0.097	0.095	0.094
4DGS	0.158	0.129	0.139	0.136	0.130	0.140	0.135	0.138
Dynamic 3DGS	0.162	0.148	0.165	0.156	0.152	0.161	0.155	0.157
Proposed	0.071	0.051	0.082	0.061	0.055	0.089	0.055	0.066

Table S6: CD results across different synthetic scenes for combined (training + interpolation) frames.

Method	Clematis	Plant1	Plant2	Plant3	Plant4	Plant5	Tulip	Average
4D-GS	0.21	0.20	2.03	0.22	0.17	2.42	0.12	0.77
4DGS	42.63	3.98	2.82	14.25	2.78	10.56	6.72	11.96
Dynamic 3DGS	79.26	0.79	2.32	1.98	0.22	0.40	9.98	13.56
Proposed	0.02	0.08	0.10	0.28	0.11	0.12	0.02	0.11

Table S7: PSNR (dB) results across different captured scenes for combined (training + interpolation) frames.


Method	Blooming flower	Corn	Paperwhite	Average
4D-GS	23.11	29.22	23.74	25.36
4DGS	26.87	30.12	25.11	27.37
Dynamic 3DGS	25.26	23.46	24.09	24.27
Proposed	26.77	30.68	24.97	27.47

Table S8: SSIM results across different captured scenes for combined (training + interpolation) frames.

Method	Blooming flower	Corn	Paperwhite	Average
4D-GS	0.983	0.982	0.970	0.978
4DGS	0.984	0.984	0.976	0.981
Dynamic 3DGS	0.976	0.946	0.966	0.963
Proposed	0.990	0.986	0.977	0.984

Table S9: LPIPS results across different captured scenes for combined (training + interpolation) frames.

Method	Blooming flower	Corn	Paperwhite	Average
4D-GS	0.032	0.052	0.048	0.044
4DGS	0.034	0.053	0.046	0.044
Dynamic 3DGS	0.044	0.111	0.062	0.072
Proposed	0.020	0.043	0.036	0.033

**Fig. S1:** We show our method’s novel view renders against baselines on trained and interpolated timesteps. Our method more faithfully reconstructs the *corn* at interpolated timesteps compared to baselines (images indicated with a yellow border are novel view renders of interpolated times).

have optimized a set of Gaussians at timestep t_0 , which we denote as \mathbf{G}^{t_0} . For the subsequent training phases, we freeze color c and opacity o . Next, for the boundary reconstruction, we integrate backwards in time, one timestep at a time and cache the optimized Gaussians for each timestep. After this phase, we have a set of cached Gaussians for each timestep. Finally, during the global optimization step, we randomly sample a timestep, and leverage the cached Gaussian at that timestep to optimize the neural ODE. The result is a trained neural ODE F_ϕ able to interpolate over unseen timepoints.

F Additional Visualizations

Adaptability to difficult scenes. Our method can also reconstruct a variety of difficult plants such as color-varying plants, multiple plant growth, and complex

Algorithm 2 Training Loop for GROWFLOW

- 1: **Input:** Gaussians \mathbf{G} , posed images I_p^t , neural ODE F_ϕ , number of timesteps N .
 - 2: **Parameters:** $n_{\text{static}} = 30000$, $n_{\text{boundary}} = 300$, $n_{\text{global}} = 30000$.
 - 3:
 - 4: **Step 1: Static Reconstruction**
 - 5: **for** $epoch = 1$ to n_{static} **do**
 - 6: Pick last timestep ground truth image $I_{\text{last}} = I_p^T$
 - 7: $I_{\text{pred}} \leftarrow \text{Rasterize}(\mathbf{G})$
 - 8: Compute $L \leftarrow \text{loss}(I_{\text{pred}}, I_{\text{last}})$
 - 9: Update \mathbf{G}
 - 10: **end for**
 - 11: **Output:** $\mathbf{G}^{t_0} = (\mu^{t_0}, q^{t_0}, s^{t_0}, c, o)$
 - 12:
 - 13: **Step 2: Boundary Reconstruction**
 - 14: **for** $k \in \{0, \dots, N-1\}$ **do** ▷ Backwards in time
 - 15: **for** $epoch = 1$ to n_{boundary} **do**
 - 16: Pick ground truth image $I^{t_{k+1}}$
 - 17: $\mathbf{G}^{t_{k+1}} = \mathbf{G}^{t_k} + \int_{t_k}^{t_{k+1}} F_\phi(\mu(t), t) dt$
 - 18: $I_{\text{pred}} \leftarrow \text{Rasterize}(\mathbf{G}^{t_{k+1}})$
 - 19: Compute $L \leftarrow \text{loss}(I_{\text{pred}}, I^{t_{k+1}})$
 - 20: Update F_ϕ
 - 21: **end for**
 - 22: Cache $\mathbf{G}^{t_{k+1}}$
 - 23: **end for**
 - 24: **Output:** a set of cached Gaussians for each timestep $\{\mathbf{G}^{t_k}\}_k$
 - 25:
 - 26: **Step 3: Global Optimization**
 - 27: Re-initialize new F_ϕ
 - 28: **for** $epoch = 1$ to n_{global} **do**
 - 29: Randomly sample timestep t_k
 - 30: Pick ground truth image $I^{t_{k+1}}$
 - 31: $\tilde{\mathbf{G}}^{t_{k+1}} = \mathbf{G}^{t_k} + \int_{t_k}^{t_{k+1}} F_\phi(\mu(t), t) dt$
 - 32: $I_{\text{pred}} \leftarrow \text{Rasterize}(\tilde{\mathbf{G}}^{t_{k+1}})$
 - 33: Compute $L \leftarrow \text{loss}(I_{\text{pred}}, I^{t_{k+1}})$
 - 34: Update F_ϕ
 - 35: **end for**
 - 36: **Output:** Optimized F_ϕ
-




Fig. S2: Difficult scenes. Our method also works on color-varying plants, multiple plant growth, and complex branching.

branching (see Fig. S2). To model color-varying plants, we add an additional MLP, $\hat{c} = \psi_c(\mathbf{z})$, integrated alongside other parameters.