Supplementary Material: Validation on a Second Plant

I. Experimental Setup

To check that the pipeline is not over-fitted to a single specimen, we apply it without modification to a second potted plant. The recording is made with the same Intel RealSense D435i camera on a tripod, and the branch is excited by the same constant-velocity blower used in the main paper. The captured RGB video has resolution of 640x480 at 30 FPS. The branch is annotated as a chain of 25 nodes with 24 edges in the first frame and tracked through 590 frames using TAPIR. Edges are grouped into the same three anatomical zones as in the main paper: *proximal* (edges 1-7), *mid* (edges 8-15), and *distal* (edges 16-24). Because each bending angle requires a grandparent-parent-child triplet, edge 1 has no measurable bending, and per-edge compliance is reported for the remaining 23 edges. All hyperparameters are inherited unchanged from the main paper ($\lambda_m = 1.0$, $\lambda_v = 0.35$, $\lambda_a = 0.2$, $\lambda_e = 0.45$ for the smoother; shrinkage prior $\lambda = 12$, block size $B = 25$, 250 bootstrap repetitions for compliance estimation).

II. Projected Compliance and Uncertainty Overlays

Fig. 1 projects the estimated κ_e and the corresponding bootstrap uncertainty U_e onto the first frame of the input video, in the same format used in the main paper. The compliance overlay (a) localises the most compliant and stiffest segments anatomically, while the uncertainty overlay (b) makes immediately visible which edges of the recovered map should be trusted and which should be treated as inconclusive. The numerical content of these overlays is the same as the per-edge plots in Fig. 4.

III. Impact of Smoothing in Noise Reduction

The factor-graph smoother again removes a substantial fraction of the geometric noise from the raw tracks while preserving the underlying bending signal (Fig. 2). On plant 2, it reduces the mean bending RMS from $4.42°$ to $2.11°$ and the mean edge-length CV from 0.077 to 0.044. Both reductions are larger than the corresponding numbers reported for plant 1 in the main paper. This is not a sign that the method behaves differently on the two specimens; it is a direct consequence of the raw input quality. Plant 2's raw tracks contain along-axis spikes of up to $\sim 10°$, whereas plant 1's raw tracks were already comparatively clean. TAPIR drifts more aggressively along the branch axis when the surrounding texture is uniform, which is the case for the longer, thinner stems of plant 2. The smoother, therefore, has more noise to remove, and its relative impact is larger. Importantly, the mean smoothed node motion (7.86 px) is essentially identical to the raw mean (8.02 px), confirming that the smoother attenuates noise rather than the underlying motion field.

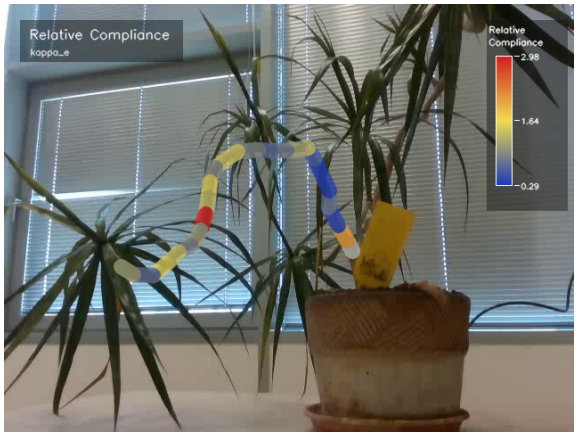

(a)

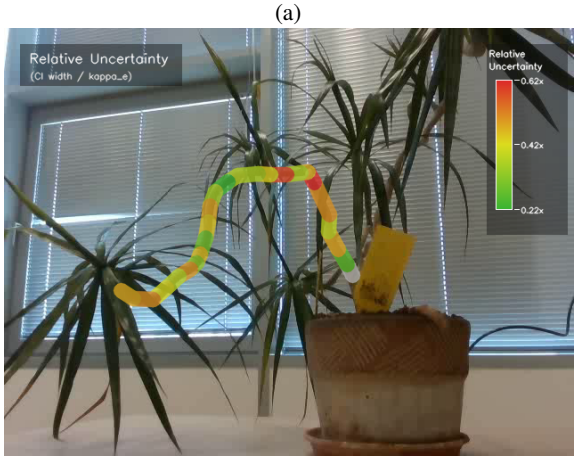

(b)

Fig. 1: plant 2 anatomical overlays. (a) Relative projected compliance κ_e (blue: stiff, red: compliant). (b) Relative bootstrap CI width U_e (green: well-constrained, red: unreliable).

IV. Motion Amplification Along the Branch

The relative motion RMS at each node, plotted as a function of root distance in Fig. 3(a), accumulates near-monotonically from the proximal zone to the tip. The trend is distinctly cleaner than the corresponding plot for plant 1. The main reason for observation is the skeleton of plant 2. Given a longer stem, the lever-arm effect predicted by the cantilever model has a larger range over which to express itself.

The per-joint bending in Fig. 3(b), in contrast, is again *not* a monotonic function of root distance. The largest bend occurs at edge 17 in the distal zone, while several proximal and mid bends sit close to or below $1.5°$. This reproduces the qualitative observation from the main paper that branch deformation is heterogeneous along the skeleton, and that which segments bend the most is determined by local material and geometric factors, not solely by how far they are from the clamp.

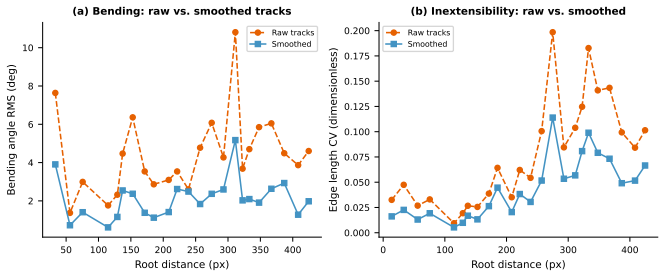

Fig. 2: Effect of factor-graph smoothing on plant 2. Left: bending angle RMS as a function of root distance. Right: edge-length coefficient of variation. Smoothed tracks (blue) substantially reduce both metrics relative to the raw moving-averaged tracks (orange), without distorting the underlying motion signal.

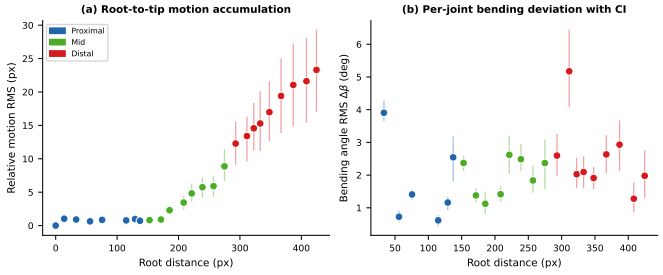

Fig. 3: plant 2 motion analysis. (a) Node relative motion RMS as a function of root distance, showing near-monotonic root-to-tip amplification. (b) Per-joint bending angle RMS $\Delta\beta$ with 95% bootstrap confidence intervals.

V. Relative Projected Compliance Map

Fig. 4 shows the per-edge projected relative compliance and its bootstrap uncertainty for all 23 non-root edges. The most compliant edge of the entire branch is edge 17 in the distal zone, with $\kappa_e = 2.98$ and a bending RMS of 5.17°. This edge is well-supported (573 of 590 inlier frames) and has a tight confidence interval, so its high compliance is an authentic structural feature rather than a small-sample artifact.

The stiffest section of the branch lies in the proximal zone: edges 5, 3, and 6 all sit clearly below the neutral line, with κ_e ranging from 0.29 to 0.51. Anatomically, these correspond to the part of the stem closest to the soil-pot interface, where the trunk is thickest and boundary condition is closest to a clamped beam. This is the prior expectation for a potted plant, and it is what the recovered map shows.

It is worth noting that this is a genuine difference from plant 1, where the stiff sections were located in the mid and distal zones rather than at the base. Plant 1 had a shorter and relatively uniform basal stem with two locally reinforced sections further along the branch, mainly due to the ∪ shape geometry of the branch. Whereas plant 2 has longer and more taper branch , with most of its rigidity concentrated near the pot mostly because of it's ∩ shape geometry. The pipeline does not impose a fixed expectation about where stiff segments should be; it reports what the bending data say, and the two plants happen to differ.

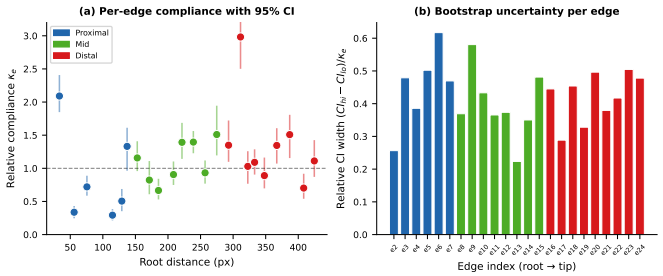

Fig. 4: plant 2 relative compliance estimation. (a) κ_e per edge as a function of root distance with 95% moving-block bootstrap confidence intervals; the dashed line marks neutral compliance ($\kappa = 1$). (b) Relative bootstrap CI width U_e per edge.

One isolated proximal edge deserves a separate note. Edge 2 sits very close to the root but shows above-average compliance ($\kappa_e = 2.09$) with a tight confidence interval. Where it shows the limitations of factor graph based smoothing. We can see in the Figure 2 (a) that the factor graph reduces the noise in edge 2 but it's still very high, because the raw TAPIR tracks are relatively bad compare to other regions due to lack of proper features and visually similar background branches. We can remove these artifacts by making vision based tracking pipeline more robust. We can either user recent and more advanced methods than TAPIR or use artificial illumination sources, and capturing videos with higher resolution and better framerate.

VI. Uncertainty Analysis

The median relative bootstrap CI width across the 23 edges of plant 2 is 0.43, noticeably higher than the 0.29 reported for plant 1. Two factors directly account for this gap. First, plant 2 was recorded over 590 frames, roughly half the 1104 frames available for plant 1, so each edge has approximately half the temporal evidence available for bootstrap resampling. Second, several mid-zone edges of plant 2 lose their tracks for extended periods because of leaf occlusion; edges 9 and 10, for example, retain only 428 inlier frames each. Both effects directly inflate the bootstrap variability, which is exactly what the moving-block bootstrap is meant to do.

Of the 23 edges, four are flagged as unreliable by the $U_e > 0.5$ threshold. Two of these (edges 5 and 6 in the proximal zone) have intrinsically small bending RMS, so their signal-to-noise ratio is poor regardless of how many frames are available. The other two (edge 9 in the mid zone and edge 23 in the distal zone) suffer either from reduced track support due to occlusion or from a low absolute bending signal. The shrinkage prior pulls these weakly-supported edges toward $\kappa_e = 1$, and the bootstrap CIs widen accordingly. Together, shrinkage and bootstrap produce a self-consistent uncertainty. Where the data is informative, the estimate is sharp (the lowest U_e in the recording is 0.22, on edge 13); where the data is poor, the estimate is conservatively pulled toward neutral and flagged.

VII. Shared Excitation Proxy

Fig. 5 shows the estimated shared excitation z_t for plant 2. The waveform exhibits the same broad-gust structure as the

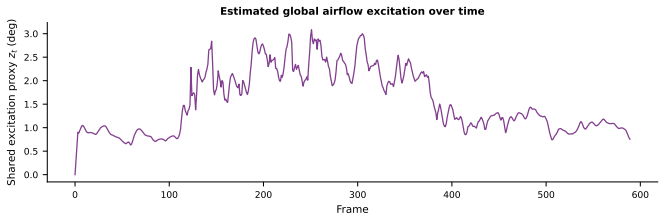

Fig. 5: Estimated shared excitation proxy z_t for plant 2 over the 590-frame sequence, capturing the periodic airflow from the blower.

corresponding signal for plant 1, with an RMS of $1.70°$. Two minor differences are worth noting. First, the RMS is slightly higher than in plant 1 ($1.18°$), which is consistent with the larger raw bending response of plant 2 noted earlier. Second, plant 1's recording opened with a roughly eight-second silent interval before the first gust, whereas plant 2's recording has an interval of around 4 seconds, which is consistent with the video.