

Statistical shape representations for temporal registration of plant components in 3D

Karoline Heiwolt¹, Cengiz Öztireli², Grzegorz Cielniak¹

Abstract—Plants are dynamic organisms and understanding temporal variations in vegetation is an essential problem for robots in the wild. However, associating repeated 3D scans of plants across time is challenging. A key step in this process is re-identifying and tracking the same individual plant components over time. Previously, this has been achieved by comparing their global spatial or topological location. In this work, we demonstrate how using shape features improves temporal organ matching. We present a landmark-free shape compression algorithm, which allows for the extraction of 3D shape features of leaves, characterises leaf shape and curvature efficiently in few parameters, and makes the association of individual leaves in feature space possible. The approach combines 3D contour extraction and further compression using Principal Component Analysis (PCA) to produce a shape space encoding, which is entirely learned from data and retains information about edge contours and 3D curvature. Our evaluation on temporal scan sequences of tomato plants shows, that incorporating shape features improves temporal leaf-matching. A combination of shape, location, and rotation information proves most informative for recognition of leaves over time and yields a true positive rate of 75%, a 15% improvement on state-of-the-art methods. This is essential for robotic crop monitoring, which enables whole-of-lifecycle phenotyping.

I. INTRODUCTION

Monitoring plants across their whole life cycle is an important task in modern agriculture and agronomy. The ability to observe dynamic changes in plant morphology will allow farmers, breeders, and researchers to measure plant performance, understand genome-environment interactions, and identify indicators for plant stresses [1]. Robotic automation plays an increasingly important role in enabling high-throughput phenotyping directly in the field, by recording richer, more accurate information at an increased observation frequency [2]. In robotics, the use of three-dimensional (3D) sensors for constructing (e.g. [3]) and maintaining (e.g. [4]) models of relatively rigid and structured environments over time are very active areas of research. However, plants are non-rigid, growing organisms with high architectural complexity. Rapid changes in appearance, morphology, and topology, such as the emergence of new organs, are common.

This work was supported by the Engineering and Physical Sciences Research Council [EP/S023917/1]

For the purpose of open access, the author(s) has applied a Creative Commons Attribution (CC BY) license to any Accepted Manuscript version arising.

¹KH and GC are with the Lincoln Centre for Autonomous Systems, University of Lincoln, Lincoln LN6 7TS, United Kingdom kheiwort@lincoln.ac.uk, gcielniak@lincoln.ac.uk

²CÖ is with the Department of Computer Science and Technology, University of Cambridge, Cambridge, United Kingdom aco41@cam.ac.uk

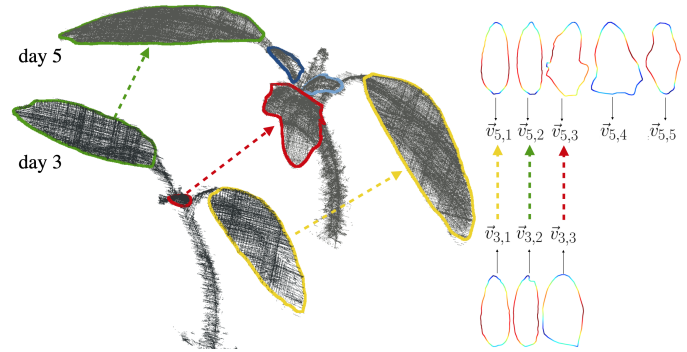


Fig. 1. Left: 3D point clouds of the same tomato plant recorded two days apart, overlaid with 3D edge contours extracted using our proposed method. Right: 3D leaf outlines of both scans are compressed into feature vectors $v_{t,i}$ and matched in feature space.

Additionally, external factors including lighting, temperature, and wind conditions cause both short-term and permanent deformations [5]. A non-rigid spatio-temporal registration, which can locate correspondences between scans, while allowing for growth and appearance changes is needed.

To cope with non-uniform growth patterns, segmenting plants into individual organs and tracking single components over time has proved helpful to simplify the registration problem [6], [7], [8]. One essential step is the re-identification and tracking of the same organs in consecutive 3D scans. Existing approaches base the pairwise matching of unique leaf instances on similarities in their global location in 3D space [7] or their relative topological location within the graph-like branching structure of the plant [9], [10]. Location proves informative across short time intervals and in controlled lab conditions with minimal external interference, such as wind. In the field, however, leaf location changes quickly and non-uniformly. Correspondence between leaves can not be guaranteed by a similar location. Especially if the time span between observations is increased, global location and topology are not reliable to robustly re-identify specific organ instances. Furthermore, topology is difficult to extract if the plant's branches can not fully be observed. In practice, due to heavy occlusions, the canopy can be observed more reliably than the hidden connectivity of leaves through a network of branches (e.g. [11]).

In this paper, we propose to use 3D shape features to track unique plant components over time. We present a novel feature extraction pipeline to compute concise shape descriptors from 3D point clouds of plants and explore the use of these shape representations for temporal matching

(Fig. 1). We define the pairwise organ association task as a minimum bipartite graph matching problem, introduce a suitable feature space cost metric, and achieve robust association of leaves over time at 75% accuracy without any manual intervention. In summary, our contributions include:

- A novel leaf shape compression pipeline, combining 3D contour extraction and subsequent PCA encoding.
- A cost metric designed to compare and associate plant components in shape feature space.
- Ablation and comparison studies demonstrating the applicability of this shape space representation to temporal leaf matching.

II. RELATED WORK

A. Spatio-temporal registration of plant parts

On top of a growing body of work focused on temporal registration on a whole-plot scale (e.g. [12]), recently, two approaches were introduced to address the spatio-temporal registration of individual plant point clouds: In [9], the authors perform point-wise binary semantic segmentation into leaves and stems, extract minimal skeleton representations of consecutive plant scans, and find pairwise correspondences between nodes of temporally separated plant skeletons by describing the association as a Hidden Markov Model, with the hidden states representing the correspondences between nodes. The emission and the transition cost metrics exploit the skeleton’s topology, i.e. the position of each node in the global graph-like organisation of the plant and their connectivity, as well as their location in Euclidean 3D space, and semantic class. Magistri and colleagues [7] pair organs of the same semantic class by minimising the overall Euclidean distance between organ point clusters in two consecutive scans. Optimal assignment is computed using the Hungarian method [13] on a cost matrix consisting of Euclidean distances between organ centroids. In both cases, local leaf shape is not used explicitly.

B. Analysing leaf shape

In several areas of plant sciences, edge contours are used as an informative proxy for leaf shape [14]. On 2D images, leaf contours are found by manually marking predefined landmarks around the leaf perimeter (e.g. [15]), or in an automated fashion, using tracing algorithms on binary images of single leaves [14]. A standard approach to measuring variation and comparing shape in 2D is learning a statistical shape model for an object class from data. Most commonly, PCA is applied to compute the orthogonal vectors of greatest variance in a training data set and project the data onto those basis vectors. Individual examples of the class can then be described as a linear combination of shape feature components. In [16], the authors measure the distances between nine landmarks along the perimeter of leaves and perform PCA on those distances to detect shape differences in leaves grown in different environmental conditions. In [14], the authors perform PCA directly on outline coordinates to distinguish Arabidopsis leaf-shape mutants. [17] introduces a method to perform PCA on

Fourier Harmonics extracted from outline gradients, producing scale- and rotation-invariant shape descriptors to discover and distinguish interpretable shape categories in data sets of unlabelled 2D images of leaves, without relying on domain expertise. In each case, a low-dimensional vector space is found to concisely capture variability in shape within the given training set.

However, 2D image projections inevitably fail to represent the characteristic curvature of leaves and can not capture the full growth process in 3D space. Related work on learning shape spaces from 3D data can be found in the fields of computer graphics, e.g. for recognition of faces [18], human skeleton tracking [19], and detecting shape anomalies in medical imaging [20]. For PCA analysis all 3D shapes have to be aligned and described with the same corresponding topology, i.e. the same number of nodes with identical matrix connections. Objects are often registered by fitting a template mesh to each example and stacking the vertex coordinates into a single input vector [19]. With our approach, we aim to extract a set of coordinates along the leaf edge, which captures both edge contour and 3D curvature. This representation is equivalent to fitting a template mesh, consisting of a single chain of vertices and edges of fixed length, to each leaf and thus establishing correspondence of coordinates.

III. METHOD

In the following, we first give a formal description of the organ assignment problem. Next, we present our proposed method, including 3D leaf outline extraction, shape space encoding via PCA, and computation of the feature space cost metric used to perform pairwise leaf matching. Finally, we define the baseline methods used for comparison.

A. Problem formulation

The pairwise association of leaves between two subsequent scans can be described as a Minimum Bipartite Graph Matching problem. A bipartite graph of leaves $G = (B, A, E)$ has two disjoint sets of vertices B (before) and A (after), such that each edge $(b_i, a_j) \in E$ connects a vertex $b_i \in B$ with one $a_j \in A$. Each edge in graph G has an associated cost $c_{ij} \in \mathbb{R}^+$. A matching M is a subset of edges E , such that no two edges in M share a common vertex. The minimum cost bipartite graph matching is a matching for which the sum of costs across all edges is the minimum

$$c(M) = \sum_{(i,j) \in M} c_{ij}. \quad (1)$$

This optimisation can be solved in polynomial time using the Hungarian method [13]. An example is illustrated in the left panel of Fig. 4.

B. Our approach

Our approach takes a time series of 3D plant point clouds with instance annotations for all leaves as input. To extract a concise summary of leaf shape, we first extract an outline representation of each leaf, consisting of a point set of $n = 500$ coordinates sampled uniformly along its 3D edge

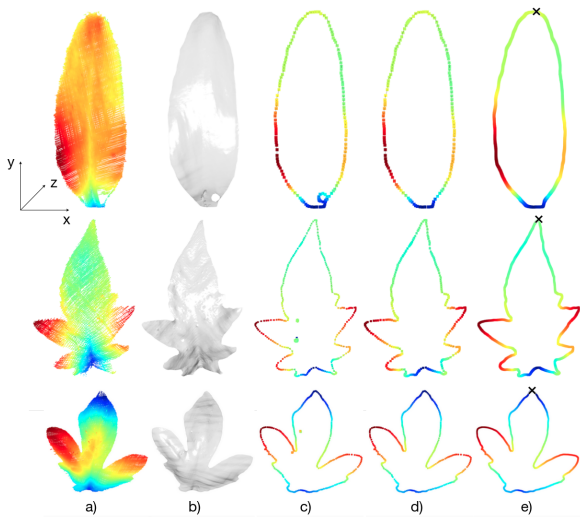


Fig. 2. The processing stages of boundary extraction, shown on three different leaf shapes. The columns show a) the aligned input point cloud, b) the triangulated mesh object, c) all mesh boundary loops, d) a single cleaned boundary loop, e) and a point set of $n = 500$ outline coordinates, with the apex marked as x . Colours are added to illustrate depth.

contour. From a training set of outlines, we learn a statistical shape space, which allows the representation of individual leaves through a linear combination of concise shape descriptors. Pairwise assignment of leaves is performed using a cost metric based on distances between examples in shape feature space.

1) *Outline extraction*: The workflow comprises three steps: Determination of the main leaf axes, surface triangulation of the point cloud to obtain a concave triangle mesh, and extraction of a point set along the 3D mesh boundary. Starting with the raw point clouds, all individual leaf point clusters are first aligned into a shared local coordinate frame, following the Leaf Axis Determination procedure in [21]. The local right-hand coordinate frame for each leaf is defined by the leaf's principal axis, pointing in the direction of the largest variance in 3D space and oriented outward from the plant emergence point, and the axis perpendicular to the leaf surface (Fig. 2a). Next, a triangle mesh object is computed from each point cloud via Delaunay Triangulation in 2.5D on the x-y-leaf plane (Fig. 2b). A crucial parameter for mesh generation is the maximum permitted length l_{max} for triangle edges, which represents a trade-off between maintaining fine detail in the contour along the edge of leaves and creating holes within the boundaries of leaves wherever the point sampling density is low. In favour of maintaining boundary detail, we set the maximum edge length low at $l_{max} = 1\text{mm}$, and instead apply smoothing and hole closing operations to avoid holes in the mesh. Before triangulation, we apply Laplacian smoothing exclusively in the z-dimension, i.e. perpendicular to the leaf surface, with a radius of 1.2 and smoothing factor of 0.2, for eight iterations. Following Delaunay Triangulation, a second step of Laplacian smoothing is applied in all three dimensions for 20 iterations. Finally, a hole filling operation is employed to

close any holes with up to 10 edges.

In the resulting mesh, a boundary loop is any chain of edges connecting those vertices, which are not enclosed within another complete loop of edges connecting its neighbouring vertices without connecting back to the vertex of interest (Fig. 2c). To filter out irrelevant closed loops around any remaining holes within the mesh, we only retain the single connected boundary loop with the largest spatial extent (Fig. 2d). Additionally, any leaf outlines for which at least one vertex is connected with two or more edges are excluded. This connectivity implies sub-loops within the contour and prevents uniform sampling along the edge in the next step.

Finally, the boundary loop is sorted into clockwise order around the leaf in its local frame, starting at the maximum point along the main axis, which usually corresponds to the leaf apex, assuming proper leaf axis alignment. We then sample $n = 500$ points around the outline, at uniform distances along the loop, starting at the apex (Fig. 2e). The coordinates are normalised by the total length of the outline to render the representation of shape scale invariant. Furthermore, the coordinate set can optionally be transformed back into global frame to re-introduce information about its global transform, specifically its location, relative rotation, and scale. The global frame origin is placed at the plant emergence point, which is considered a temporally consistent reference point to align all scans of the same plant over time. In summary, the outline representation retains information about the edge contour shape and 3D curvature of the leaf. If given in the leaf's local coordinate frame with respect to the leaf apex at its origin, it is invariant to the scale, global location, and global orientation. Alternatively, we optionally translate, rotate, and scale outlines back to the original scan frame, retaining its original location, orientation, and scale relative to the plant emergence point. The representation achieves correspondence between points on the outlines of differently shaped leaves, which allows for the application of PCA.

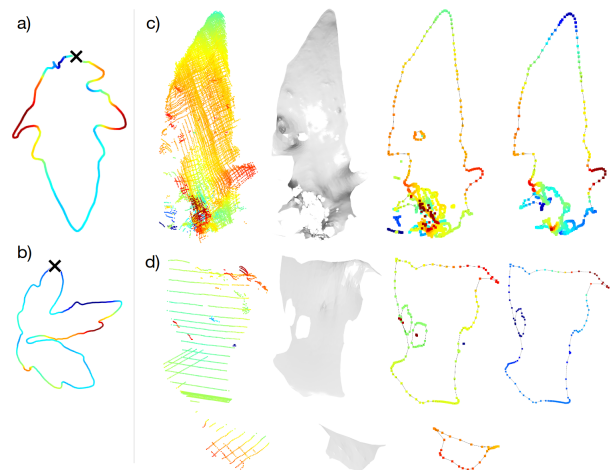


Fig. 3. Error cases in the outline extraction procedure: In a) and b), axis alignment failed because of violated assumptions (13% of leaves). In c) the sampling density is uneven across the leaf, and the leaf captured in d) is very small and thus the sampling density is low across the entire leaf.

2) *Limitations of outline extraction:* The assumption that leaves grow in an outward direction from the plant emergence point is violated for an average of 13% of all leaves across any scans in the Pheno4D data set. These leaf point clouds are then oriented upside-down (Fig. 3a) or less frequently sideways (Fig. 3b) for the remainder of the outline extraction and as a consequence the apex point is misidentified. For our evaluation, we chose not to exclude these cases, but we found that shape space learning is improved when the correspondence between single leaves can be established with correctly aligned axes more consistently (see IV-D). Furthermore, due to inadequacies of the sensing setup, the outline of leaves is sometimes ambiguous. Where the leaves are very small and/or the sensing density is low, it is difficult to extract complete mesh objects (see examples c) and d) in Fig. 3). Future iterations of this work will explore alternative meshing algorithms so that more data can be considered for training. In our study, clean outlines could be extracted from 68% of the original annotated data set. The exclusion of those leaves emulates difficulties that are realistic to be expected in the field as well, considering that new leaves emerge and old leaves frequently die or are trimmed between time steps. A robust temporal matching algorithm is expected to cope with missing leaf instances.

3) *PCA shape encoding:* PCA computes a linear transformation which can be used to represent an input outline $o \in \mathbb{R}^{3n}$ with n coordinates as a linear combination of d shape parameters, which form a feature vector $v \in \mathbb{R}^d$. The features are analogous to scalar weights v_k in a d -dimensional subspace, spanned by the orthogonal vectors V , such that

$$o(v) = \bar{o} + \sum_{k=1}^d v_k V_k, \quad (2)$$

where \bar{o} is the mean of all outlines in the training set. The shape parameters are fit to training data and inversely, a shape vector can be projected back into 3D space. We apply z -score-normalisation to standardise feature scaling of the input outline vectors before fitting a PCA model to the training set.

4) *Cost metric:* The same feature scaling and PCA encoding are applied to the test set in all following experiments. To compute our cost metric, which is inversely proportional to the similarity between leaves in this multivariate shape feature space, we use the Mahalanobis distance metric. The Mahalanobis distance is a unitless and scale-invariant metric, measuring each feature's z -score along the corresponding axis, taking into account single features' distributions and potential correlations across the training data set. For a pair of PCA-encoded feature vectors $v_{t,i}$ and $v_{t+1,j}$ the cost is computed as

$$c_{i,j} = \sqrt{(v_{t,i} - v_{t+1,j})K^{-1}(v_{t,i} - v_{t+1,j})^T}, \quad (3)$$

where K is the covariance matrix of features computed across the training data set.

C. Cost metrics based on transform

In [7] the optimal pairwise matching is found using a cost based on the mean location of each leaf. Additionally, global rotation and scale information can be easily extracted from the leaf point cloud transform in the full plant's global reference frame. We represent any combination of location, rotation, and scale information as a one-dimensional feature vector. Location is given as the leaf centroid $m \in \mathbb{R}^3$, rotation as the three vectors describing the axes of the local leaf coordinate frame, given with respect to the global frame $x', y', z' \in \mathbb{R}^3$, and finally scale $s \in \mathbb{R}$, a scalar equivalent to the total length of the outline loop in the original coordinate frame. Any combination of these transform descriptors are stacked in a one-dimensional vector $v = [m, x', y', z', s]$ with length between one and 13 elements. The cost c is computed as the Mahalanobis distance, as defined in Eq. 3.

IV. EXPERIMENTAL EVALUATION

In the following, we evaluate the quality of feature encoding and performance of our leaf-matching approach based on shape characteristics and various combinations of information about their global location, rotation, and size.

A. Data set

For testing, we use time-series 3D point cloud data of real tomato plants (*Solanum Lycopersicum*) taken from the Pheno4D set [22]. It contains point clouds of 7 plants, captured once per day, across 20 days. 77 of the total 140 point clouds are annotated with consistent organ instance labels. From those annotated scans, our outline extraction procedure yields 649 boundary instances, including leaves from all 7 unique plants, and spread across 11 time steps (on average two days apart). For the following experiments, we use leaf outlines from two plants for training (corresponding to 25% of available data), and the remaining five for testing. This split ensures a sufficiently diverse training set of leaves from all growth stages.

B. Performance metrics

Typically, the Hungarian algorithm is applied to assign a set of n jobs to a set of m machines. If $n \neq m$, the algorithm finds a matching consisting of $\min\{m, n\}$ edges, so that every vertex of the smaller set is covered. In the case of plants, new components emerge and old components die or are trimmed frequently in between scans. These appearing and disappearing graph vertices are still included in the minimum cost matching (in Fig. 4 leaves 3 and 5 are missing at time t , leaf 4 is missing at time $t+1$). As a result, the algorithm naturally predicts additional expected false positive (FP) pairings, representing the best-fitting correspondences for any single leaves present in only one of the two scans. Thus, FP do not directly reflect the quality of the matching performance. However, we expect a suitable cost metric to yield a significantly higher associated cost for a false edge (F), compared to true edges (T), which include true positives (TP) and false negatives (FN, i.e. where a true pairing was not selected, see Fig. 4b for an example). If the degree of

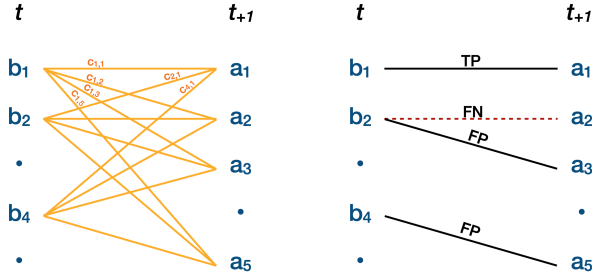


Fig. 4. **Left:** Example graph G with all edges E shown in yellow. Their associated costs $c_{i,j}$ are indicated in an exemplary way for the vertices b_1 and a_1 . The vertices 3, 4, and 5 are only present in one of the two time steps t or t_{+1} . **Right:** A predicted minimum cost matching (black lines) including examples of a true positive (TP) and false negative (FN), as well as an expected and an unexpected false positive (FP) pairing.

separation between cost distributions for true corresponding edges compared to all possible false correspondence edges in E is high, then a suitable cost metric could be used analogously to a confidence metric for suggested pairings, allowing for a threshold to be placed on the associated cost for predicted pairings in the future. To reflect this requirement, we evaluate the sensitivity, or true positive rate, as

$$\text{sensitivity} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (4)$$

and the Distribution-Free Overlapping Index [23] between observed cost distributions for true and false correspondence edges in E across the test set as

$$\eta = \frac{\sum_{i=1}^b \min\{f_i^T, f_i^F\}}{\min\{T, F\}}, \quad (5)$$

where f^T and f^F are the observed discrete frequency densities of the cost associated with true and false edges respectively, discretised into the same $b = 50$ bins. The measure is normalised to values between 0 and 1 by dividing by the number of edges in the smaller set. A value of 0 indicates perfect separation, 1 indicates perfect overlap of the cost distributions.

C. Analysis of shape space encoding

Fig. 5 shows the t-distributed stochastic neighbour embedding (*t-SNE*) for four individual leaves across time. This visualisation models the multivariate feature space distances between leaf instances in 2D and shows, how similar shapes cluster together. In the given example, differentiation between the four unique shown leaves in shape feature space is possible and thus the learned shape representation appears to capture useful information for associating leaves by shape similarities.

Fig. 7 shows an example of how the cost metric is then applied in practice to relate the organs of two consecutive 3D scans. The figure shows all extracted leaf outlines (left) and the confusion matrix of pairwise feature space distance or cost $c_{i,j}$ associated with each possible edge connecting vertices of B and A . Like in Fig. 4, black lines are used to show the matches predicted by our method. All true

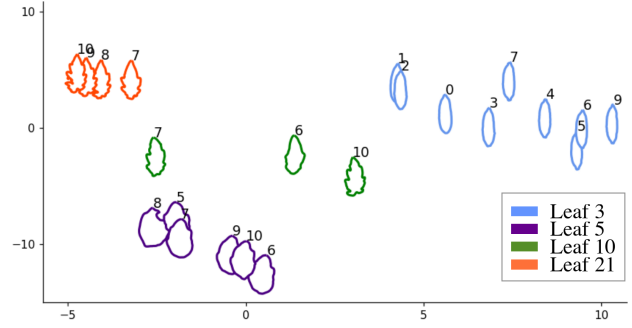


Fig. 5. *t-SNE* visualisation of the learned shape representations for all instances of four example leaves. All outlines of the same colour are temporally separated instances of the same unique leaf. Number annotations indicate at which time step each outline was captured.

pairs are detected (sensitivity = 1) and one additional expected false positive is predicted, representing the best-fitting pair of leaves present in only one of the two scans, as discussed in section IV-B. Note that leaf 7, which is misaligned and processed upside-down in both scans, is still assigned correctly. Across the full data set, we observed that this is generally the case, as long as both instances of the leaf are misaligned similarly.

Matching performance is further dependent on the amount of compression applied via PCA. Fig. 6 illustrates the sensitivity of our outline matching procedure using shape only as a function of the number of principal components considered. We found that performance peaked in the range between 19 and 28 features. All following experiments were performed using the first 23 principal components for feature space compression.

D. Ablation study

Table I shows the sensitivity and overlapping index observed for the test set of 489 leaves taken from 5 plants across 11 time steps. The combination of shape, location, and orientation information proves most informative, achieving an overall true positive rate of 75.4%. While the difference in means between observed discrete frequency densities of cost associated with true and false edges respectively proved very highly significant for all ablations in independent samples *t*-tests (all $p < 0.001$), the proportion of overlap varies broadly. The biggest separation between observed cost distributions is achieved by our model using shape, location, and orientation information. We further observed that adding scale information appears to hinder the successful learning of informative

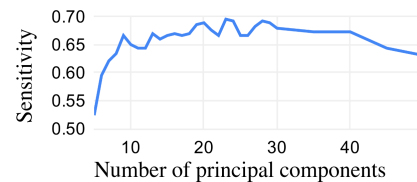


Fig. 6. Sensitivity of outline matching procedure using shape, location, and rotation, as a function of the number of principal components considered.

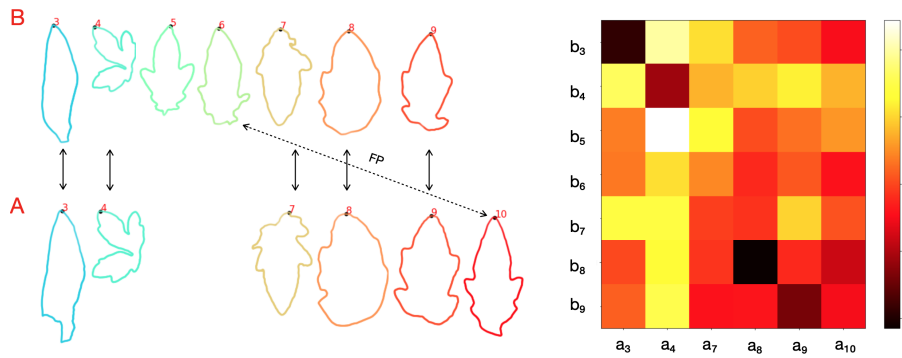


Fig. 7. Example visualisation of the matching step. **Left:** All extracted leaf outlines observed at two consecutive time steps (B - before and A - after) as vertices of a bipartite graph. The arrows indicate edges, which are part of the minimum cost matching predicted by our method, including one natural false positive (FP) representing the best fit between new or disappearing leaves. **Right:** The confusion matrix of costs associated with each possible edge connecting elements of B and A .

TABLE I
ABLATION STUDY

| shape | Ablations | | | sensitivity | η |
|-------|-----------|----------|-------|---------------|---------------|
| | location | rotation | scale | | |
| ✓ | | | | 0.6958 | 0.658 |
| ✓ | ✓ | | | 0.7087 | 0.6084 |
| ✓ | | ✓ | | 0.6958 | 0.67 |
| ✓ | | | ✓ | 0.6343 | 0.9935 |
| ✓ | ✓ | ✓ | | 0.7540 | 0.5275 |
| ✓ | ✓ | | ✓ | 0.6958 | 0.8673 |
| ✓ | | ✓ | ✓ | 0.5955 | 0.9935 |
| ✓ | ✓ | ✓ | ✓ | 0.7152 | 0.8382 |

TABLE II
COMPARISON STUDY

| shape | Method | | | sensitivity | η | |
|-------|----------|----------|-------|---------------|---------------|--------|
| | location | rotation | scale | | | |
| ✓ | ✓ | ✓ | | 0.7540 | 0.5275 | (ours) |
| | ✓ | | | 0.6051 | 0.9288 | [7] |
| | | ✓ | | 0.4822 | 0.9676 | |
| | | | ✓ | 0.2654 | 1 | |
| | ✓ | ✓ | | 0.6117 | 0.8382 | |
| | ✓ | | ✓ | 0.6634 | 0.9029 | |
| | | ✓ | ✓ | 0.5307 | 0.8997 | |
| | ✓ | ✓ | ✓ | 0.6246 | 0.8188 | |

shape representations. By normalising outline length and thus removing uniform scale differences, the shape model learned through PCA only considers shape differences and not size differences. Reintroducing scale information, however, might interfere with the effort to establish correspondence between outline points and impair the model’s ability to sufficiently capture shape. Additional to the results shown in table I, we found that removing any misaligned leaves such as depicted in 3a) and b) improved our best model by a further 10 percentage point margin to 85.5% sensitivity and an overlapping index of 0.4298. However, this was achieved by visual examination and manual exclusion of individual leaves and an automatic filter does not currently exist in our implementation.

E. Comparison study

This experiment is designed to support the claim that shape features are a useful similarity measure for spatio-temporal tracking of unique plant components between temporally separated point clouds. We compare our feature space cost metric to alternative cost metrics based on transform information in Table II. In our tests, for tomato plants scanned on average two days apart, our best method constitutes a 15% sensitivity increase from alternative methods found in literature and 9% improvement from the best alternative exploiting global leaf transform information. Note also that the shape-only ablation of our method still outperforms any method without shape information. In this experiment, again all methods produced cost distribution for true and false

graph edges, which differed significantly in their mean, but the overlap between the cost distributions could be reduced notably from 0.9288 when considering location only to 0.5275 with our approach.

V. CONCLUSION

We propose a novel workflow to track individual leaves in sequences of 3D scans of plants by computing correspondences in shape feature space. We confirm our hypothesis that shape information is more informative for the task than location or additional leaf transform information. Our experiments suggest, that our approach outperforms the state of the art and a combination of shape, location, and rotation proved most informative in an ablation study. Future work will include improvements to the boundary extraction pipeline, including exploring alternative meshing algorithms. Furthermore, PCA is a common but limited statistical tool to learn shape spaces, this could be extended by employing feature learning such as through autoencoder networks. On top of that, alternative boundary representations might be employed to ensure full rotational invariance. Eventually, we aim to use the data association achieved with this method to establish time-series prediction models for plant growth. Ultimately, robust temporal registration of plant components in 3D is an important underpinning technology needed to enable sustained crop monitoring with a large potential positive impact on the integration of agricultural robots into a sustainable global food chain, as well as applications in vegetation mapping for robots in natural environments.

REFERENCES

- [1] S. Das Choudhury, A. Samal, and T. Awada, "Leveraging image analysis for high-throughput plant phenotyping," *Frontiers in plant science*, vol. 10, p. 508, 2019.
- [2] K. M. F. James, D. J. Sargent, A. Whitehouse, and G. Cielniak, "High-throughput phenotyping for breeding targets—current status and future directions of strawberry trait automation," *PLANTS, PEOPLE, PLANET*, vol. 4, no. 5, pp. 432–443, 2022. [Online]. Available: <https://nph.onlinelibrary.wiley.com/doi/abs/10.1002/ppp3.10275>
- [3] M. Zollhöfer, P. Stotko, A. Görlitz, C. Theobalt, M. Nießner, R. Klein, and A. Kolb, "State of the art on 3d reconstruction with rgb-d cameras," in *Computer graphics forum*, vol. 37, no. 2. Wiley Online Library, 2018, pp. 625–652.
- [4] T. Krajník, J. P. Fentanes, J. M. Santos, and T. Duckett, "Fremen: Frequency map enhancement for long-term mobile robot autonomy in changing environments," *IEEE Transactions on Robotics*, vol. 33, no. 4, pp. 964–977, 2017.
- [5] F. Tardieu, L. Cabrera-Bosquet, T. Pridmore, and M. Bennett, "Plant phenomics, from sensors to knowledge," *Current Biology*, vol. 27, no. 15, pp. R770–R783, 2017.
- [6] F. Okura, "3d modeling and reconstruction of plants and trees: A cross-cutting review across computer graphics, vision, and plant phenotyping," *Breeding Science*, vol. 72, no. 1, pp. 31–47, 2022.
- [7] F. Magistri, N. Chebrolu, and C. Stachniss, "Segmentation-based 4d registration of plants point clouds for phenotyping," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 2433–2439.
- [8] A. Zaganidis, L. Sun, T. Duckett, and G. Cielniak, "Integrating deep semantic segmentation into 3-d point cloud registration," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 2942–2949, 2018.
- [9] N. Chebrolu, F. Magistri, T. Läbe, and C. Stachniss, "Registration of spatio-temporal point clouds of plants for phenotyping," *PLOS ONE*, vol. 16, no. 2, p. 1–25, Feb 2021.
- [10] H. Pan, F. Hétyroy-Wheeler, J. Charlaix, and D. Colliaux, "Multi-scale space-time registration of growing plants," in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 310–319.
- [11] E. Marks, F. Magistri, and C. Stachniss, "Precise 3d reconstruction of plants from uav imagery combining bundle adjustment and template matching," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2259–2265.
- [12] J. Dong, J. G. Burnham, B. Boots, G. Rains, and F. Dellaert, "4d crop monitoring: Spatio-temporal reconstruction for agriculture," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 3878–3885.
- [13] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [14] A. Backhaus, A. Kuwabara, M. Bauch, N. Monk, G. Sanguinetti, and A. Fleming, "Leafprocessor: a new leaf phenotyping tool using contour bending energy and shape cluster analysis," *New phytologist*, vol. 187, no. 1, pp. 251–261, 2010.
- [15] E. Biot, M. Cortizo, J. Burguet, A. Kiss, M. Oughou, A. Maugarny-Calès, B. Gonçalves, B. Adroher, P. Andrey, A. Boudaoud, *et al.*, "Multiscale quantification of morphodynamics: Morpholeaf software for 2d shape analysis," *Development*, vol. 143, no. 18, pp. 3417–3428, 2016.
- [16] J. P. Young, T. A. Dickinson, and N. G. Dengler, "A morphometric analysis of heterophyllous leaf development in *ranunculus flabellaris*," *International Journal of Plant Sciences*, vol. 156, no. 5, pp. 590–602, 1995.
- [17] J. Victorino and F. Gómez, "Contour analysis for interpretable leaf shape category discovery," *Plant methods*, vol. 15, no. 1, pp. 1–12, 2019.
- [18] V. Blanz and T. Vetter, "A morphable model for the synthesis of 3d faces," in *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, 1999, pp. 187–194.
- [19] B. Allen, B. Curless, and Z. Popović, "The space of human body shapes: reconstruction and parameterization from range scans," *ACM transactions on graphics (TOG)*, vol. 22, no. 3, pp. 587–594, 2003.
- [20] M. Rychlik, W. Stankiewicz, and M. Morzyński, "Applications of 3d pca method for extraction of mean shape and geometrical features of biological objects set," *Mathematical Modelling and Analysis*, vol. 13, no. 3, pp. 413–420, 2008.
- [21] R. Ando, Y. Ozasa, and W. Guo, "Robust surface reconstruction of plant leaves from 3d point clouds," *Plant Phenomics*, vol. 2021, 2021.
- [22] D. Schunck, F. Magistri, R. A. Rosu, A. Cornelißen, N. Chebrolu, S. Paulus, J. Léon, S. Behnke, C. Stachniss, H. Kuhlmann, *et al.*, "Pheno4D: A spatio-temporal dataset of maize and tomato plant point clouds for phenotyping and advanced plant analysis," *PLOS ONE*, vol. 16, no. 8, pp. 1–18, 2021.
- [23] M. Pastore and A. Calcagni, "Measuring distribution similarities between samples: A distribution-free overlapping index," *Frontiers in psychology*, vol. 10, p. 1089, 2019.