


Graph-Enhanced Optimal Transport: Leveraging Structural Features for Plant Matching

Aakaash Radhoe



Graph-Enhanced Optimal Transport: Leveraging Structural Features for Plant Matching

by

Aakaash Radhoe

to obtain the degree of Master of Science in Robotics
at the Delft University of Technology,
to be defended publicly on Thursday December 19, 2024 at 9:30 AM.

Student number:	4450108	
Project duration:	February, 2024 – December, 2024	
Faculty:	Mechanical Engineering (ME)	
Department:	Cognitive Robotics (CoR)	
Thesis committee:	Prof. Dr. ing. J. Kober,	TU Delft, supervisor
	Dr. Ing. G. Franzese,	TU Delft, supervisor
	Dr. Ing. C. Della Santina,	TU Delft

Graph-Enhanced Optimal Transport: Leveraging Structural Features for Plant Matching

Aakaash Radhoe
Mechanical Engineering

Delft University of Technology
Delft, Netherlands

A

Abstract—This paper presents a novel Graph Optimal Transport (Graph OT) framework for analyzing and aligning plant structures across different growth stages and transformations. Our method extends existing graph matching techniques by incorporating domain-specific botanical features and employing a multi-scale matching strategy that captures both local and global structural characteristics. The framework combines multiple feature representations, including node descriptors, spectral embeddings, Node2Vec embeddings, and relative positions, to construct an augmented cost matrix for optimal transport-based matching. We evaluated our approach on a dataset of 50 distinct plant structures under various transformations, including rotation, deformation, and partial matching scenarios.

The results indicate that our Graph OT framework significantly outperforms traditional optimal transport (OT) methods, achieving node-matching accuracy scores of 0.75 for rotated, 0.74 for deformed structures, 0.67 for cut structures, and 0.71 for structures with skipped nodes. Our approach demonstrates particular robustness in handling complex transformations. This method provides a powerful tool for botany applications such as crop management, growth modeling, and automated pruning systems.

Index Terms—Optical Transport, Graph Theory, Spectral Embedding, Plant Matching, Structural Alignment

I. INTRODUCTION

The study of plant structures and their growth patterns is a crucial area of research in botany, agriculture, and environmental science [1]. Plant registration, the process of aligning and matching structural representations between different specimens, enables transformative applications in plant science and agricultural technologies.

Traditional approaches to plant structure analysis have been constrained by manual measurement techniques and simplified representations that inadequately capture the intricate morphological complexities of plant architectures. In recent years, the advent of high-resolution 3D scanning technologies has enabled researchers to capture detailed representations of plant structures, creating an urgent need for more sophisticated computational methods to analyze and compare complex plant datasets.

The potential of plant registration extends beyond mere structural analysis. By establishing precise structural correspondences, robot policies can be transferred between different plant specimens, enabling comparative analysis of growth patterns and architectural characteristics. This approach paves the




Fig. 1: The registration of plant structures involves aligning graphs, where the red graph represents the source and the blue graph represents the target. Meanwhile, green lines indicate the matches identified between these graphs.

way for developing generalized learning algorithms capable of operating across multiple plant types, bridging the gap between individual plant studies and broader botanical understanding.

This paper presents a novel approach to perform registration between two plants using Graph Optimal Transport (Graph OT), a framework that combines graph theory, optimal transport, and machine learning techniques. Our method aims to address several key challenges in plant structure analysis:

- 1) Examination and alignment of plant structures from various angles;
- 2) Comparison and alignment of plant structures throughout various deformations, as a plant is not a rigid object;
- 3) Assessment and alignment of plant structures in trimmed versions of the plants.

Using recent advancements in graph comparison and optimal transport theory, our approach offers a robust and flexible framework for plant registration. The proposed method extends existing graph matching techniques by incorporating domain-specific features relevant to plant morphology and by employing a multi-scale matching strategy that captures both local and global structural characteristics. Through this work, our aim is to contribute to the field of computational botany by providing a powerful tool for analyzing plant structures, which can aid in various applications such as crop harvesting,

growth modeling, and plant pruning. The main contributions of this paper are:

- 1) Development and use of a novel 2D plant graph dataset to test plant structure analysis techniques.
- 2) An extended Graph Optimal Transport framework tailored for plant registration, incorporating domain-specific features and multi-scale matching.
- 3) A comprehensive evaluation demonstrating effectiveness across various transformation scenarios.

The paper is organized as follows: Section II provides mathematical background on Optimal Transport. Section III summarizes the main related works in the field. Section IV describes the proposed Graph Optimal Transport (Graph OT) framework in detail. Section V introduces our novel 2D plant graph dataset created for evaluating graph matching algorithms. Section VI presents experimental results and comparisons with existing methods, and a detailed analysis of the Graph OT framework’s registration capabilities across different plant structure transformations. Finally, Section VIII concludes the paper with a discussion of the implications of our work and potential directions for future research.

II. BACKGROUND

Optimal Transport (OT) is a mathematical framework that has gained significant attention in various fields, including machine learning, computer vision, and data analysis [2]. It provides a principled way to compare and align probability distributions, making it particularly useful for comparing structured data such as graphs [3].

The concept of Optimal Transport dates back to the work of Gaspard Monge in the 18th century [4], who formulated it as a problem of efficiently moving piles of soil. In the 1940s, Leonid Kantorovich generalized this problem, leading to what is now known as the Monge-Kantorovich problem [5].

Let μ and ν be two probability measures defined in the spaces \mathcal{X} and \mathcal{Y} , respectively. The Optimal Transport problem aims to find a transport plan π that minimizes the cost of moving the mass from μ to ν . This can be expressed as follows

$$\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y) \quad (1)$$

where $c(x, y)$ is a cost function that defines the cost of moving the mass from x to y , and $\Pi(\mu, \nu)$ is the set of all joint probability measures in $\mathcal{X} \times \mathcal{Y}$ with marginals μ and ν .

A particularly important concept in OT is the Wasserstein distance, also known as the Earth Mover’s Distance (EMD) [6]. For discrete probability distributions p and q , the Wasserstein distance of order p is defined as:

$$W_p(p, q) = \left(\inf_{\pi \in \Pi(p, q)} \sum_{i, j} \pi_{ij} d(x_i, y_j)^p \right)^{1/p} \quad (2)$$

where $d(x_i, y_j)$ is a distance metric between points x_i and y_j .

Optimal transport has been successfully applied to graph comparison tasks [3]. In these applications, graphs are typically represented as probability distributions over their nodes

or edges, and the OT framework is used to find an optimal alignment between these distributions.

The Gromov-Wasserstein distance [7] is particularly relevant for graph comparison, as it allows for comparing distributions that lie in different metric spaces. This is crucial for graphs, as they often do not share the same set of nodes.

While the OT problem has a clear mathematical formulation, solving it efficiently for large-scale problems can be challenging. Various algorithms have been proposed to address this, including: 1) The Hungarian algorithm for the linear assignment problem [8], 2) the Sinkhorn algorithm for entropic regularized OT [9] and 3) approximation methods for large-scale OT [10].

These computational advances have made OT increasingly practical for real-world applications, including the comparison and alignment of complex structures like graphs and plants.

In this work, we explore how these concepts from Optimal Transport can be applied and extended to the specific problem of graph registration, incorporating additional graph-theoretic features to enhance the alignment process.

As a matter of fact, graph theory provides a powerful framework for representing and analyzing complex structures, including plant architectures. A graph $G = (V, E)$ consists of a set of vertices V and a set of edges E connecting these vertices. In the context of plant structures, the vertices can represent branch junctions or endpoints, while the edges represent the connections between these points.

Graph-based representations of plants offer several advantages:

- 1) They capture the topological relationships between different parts of the plant.
- 2) They can incorporate additional attributes such as branch lengths, angles, and other morphological features.
- 3) They allow for the application of well-established graph algorithms and analysis techniques.

The next section highlights the main application developments of graph-based optimal transport related to plant matching.

III. RELATED WORK

Previous work has explored the use of graphs in plant structure analysis, including the development of multiscale tree graphs [11] and the application of graph-based methods for plant registration [12].

The application of Optimal Transport (OT) to graph comparison tasks has led to the development of Graph Optimal Transport (Graph OT) methods. These approaches aim to find correspondences between nodes in different graphs by formulating the problem as an optimal transport task.

Recent work in this area includes Vayer et al. [3], who proposed a framework for comparing graphs using the Gromov-Wasserstein distance, allowing for the comparison of distributions in different metric spaces. Additionally, Margetic et al. [13] introduced a GOT framework that leverages both local and global graph structures for more accurate matching. Moreover, Margetic et al. [14] improved their GOT framework

by presenting fGOT, an approach that utilizes filters to refine the algorithm’s performance and manage incomplete graphs.

While GOT methods have not been widely applied to plant structure analysis, efforts have been made to develop registration techniques for plant point clouds. Recent research by the University of Bonn has made significant strides in plant registration techniques. Chebrolu et al. [12] developed a novel approach to register spatio-temporal plant point clouds, focusing on phenotyping applications. Their method employs deep learning to learn optimal features for matching plant structures across different time points. This work addresses similar challenges in plant structural comparison but differs methodologically from GOT approaches.

Although their work addresses similar challenges in plant structural comparison, there are key methodological differences. Chebrolu et al. [12] transform point cloud data into skeletal graph representations as an intermediate step, whereas our approach operates directly on graph structures. This distinction is crucial: we focus on graph-level registration techniques, treating the graph extraction as a preprocessing step outside the scope of our primary registration method. Also, the matching itself is done with non-rigid ICP, which is not the case for us.

Research focusing on spatial-temporal matching of plant structures provides valuable information on the challenges of the problem of plant alignment [15], [16]. These works are complementary to our graph-based optimal transport approach. By targeting different aspects of the plant registration problem, these studies collectively advance our understanding of computational methods for analyzing plant morphology.

The related work section highlights the potential of employing Graph Optimal Transport techniques for plant registration. By integrating the advantages of graph-based representations, optimal transport theory, and domain-specific plant characteristics, our research seeks to improve current advances in plant registration. In the subsequent section, we will elaborate on our novel Graph OT-based approach to plant registration and alignment.

IV. GRAPH OPTIMAL TRANSPORT FOR PLANT STRUCTURE ALIGNMENT

This section presents a novel algorithm for aligning and comparing plant structures using an extended Graph Optimal Transport (Graph OT) framework. Our approach builds upon recent advancements in graph comparison using Optimal Transport, while incorporating domain-specific features relevant to plant structures.

Figure 2 presents the complete pipeline of our proposed method. The algorithm consists of several key stages, each designed to capture different aspects of the plant structure comparison problem.

The pipeline begins with two input graphs (G_1 and G_2) representing the plant structures to be compared. These graphs are processed in parallel during the feature extraction phase, which calculates four different feature types: Node Descriptors, Spectral Embeddings, Node2Vec Embeddings, and Rel-

ative Positions. Each feature type will be discussed in greater detail later in this section.

These features are then combined to construct an augmented cost matrix, which serves as input to the hierarchical matching process. The matching process employs Optimal Transport principles to find correspondences between nodes while maintaining both local and global structural consistency.

The final stage involves a RANSAC-based refinement step that helps to account for geometric transformations between the plant structures. The pipeline outputs the matched nodes, which can then be used for accuracy calculation and subsequent analysis tasks.

This integrated approach allows us to leverage both structural and spatial information effectively, making the matching process robust to various types of variations in plant structures, including rotations, deformations, and partial matches.

A. Feature Extraction

To capture the structural and positional information of the plant, we compute several node-level features and embeddings, which are comprised of the following steps.

1) *Relative Positions*: To mitigate the sensitivity to global graph transformations, we compute the positions of the nodes relative to the centroid of the graph. Given a graph’s original node positions of a graph $p_v, v \in V$, we first normalize these positions. We then compute the centroid of the graph c and define the relative position of each node as $p'_v = p_v - c$. This transformation ensures that the graph structure is invariant to global translations, the topological relationships are preserved and the node positions are comparable across different graph instances.

2) *Node Descriptors*: In our novel approach to enhance the graph matching process, we introduce a comprehensive set of node-level features that capture both local and global structural information of the plant. For each node $v \in V$, we compute the following unique combination of graph-theoretic features:

$$f(v) = [d(v), c(v), e(v), b(v), ec(v), cl(v), dr(v)] \quad (3)$$

For each node $v \in V$, we compute a set of local and global topological features. The degree $d(v)$ represents the number of edges connected to the node, capturing the branching complexity at each point. The clustering coefficient $c(v)$ measures the interconnectedness of a node’s neighboring nodes, indicating the compactness of local branch structures. Eigenvector centrality $e(v)$ assigns importance to a node based on the significance of its neighbors, highlighting key structural junction points. Betweenness centrality $b(v)$ quantifies how often a node lies on the shortest paths between other nodes, identifying critical connective branches. Eccentricity $ec(v)$ measures the maximum distance from a node to all other nodes, characterizing the node’s structural extremity. Closeness centrality $cl(v)$ captures a node’s proximity to all other nodes, indicating its centrality within the plant structure. The distance from the root $dr(v)$, calculated as the shortest path distance from the node to the plant’s root, provides information on the hierarchical positioning of the node.

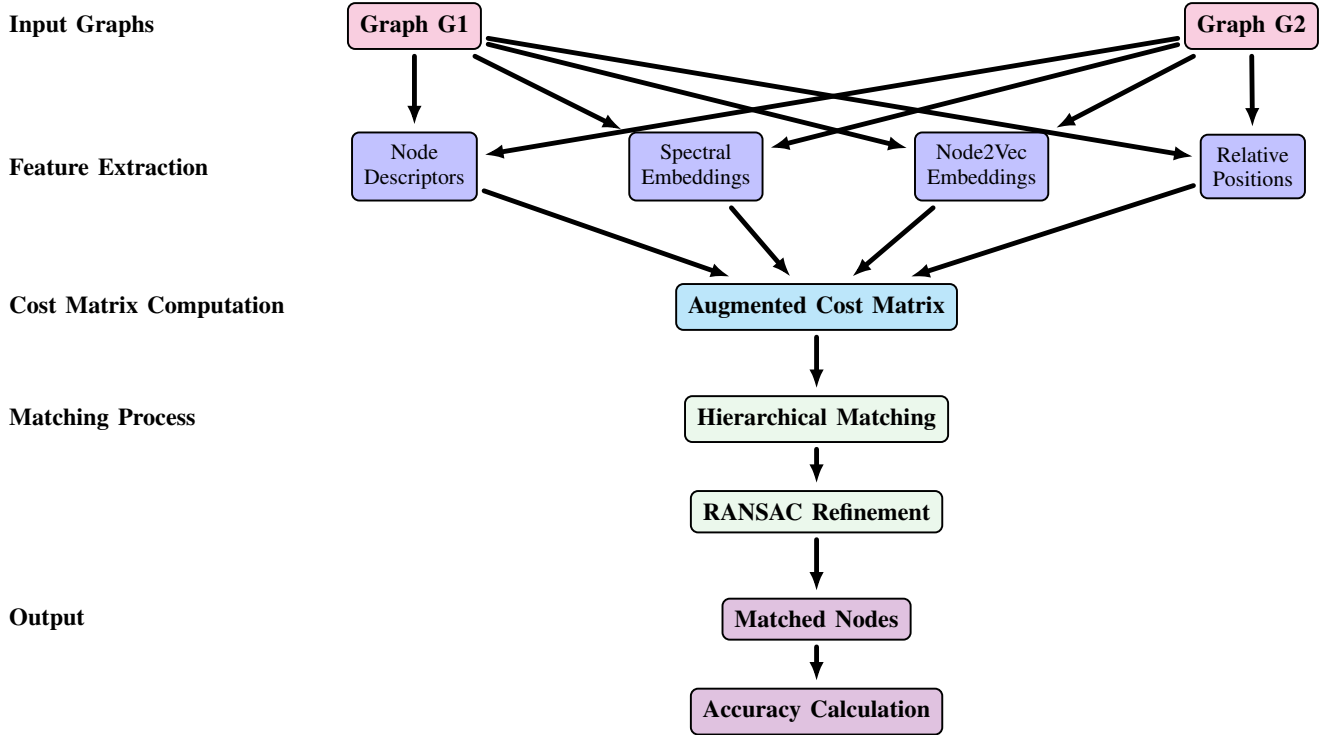


Fig. 2: Pipeline for graph matching and analysis, with aligned labels and connecting lines.

3) *Spectral Embeddings*: We compute the spectral embedding of the graph using the eigenvectors of the normalized Laplacian matrix L , defined as

$$L = I - D^{-1/2}AD^{-1/2} \quad (4)$$

where A is the adjacency matrix and D is the degree matrix. We use the eigenvectors as a low-dimensional representation of the graph structure.

4) *Node2Vec Embeddings*: To capture higher-order neighborhood information, we employ a simplified version of random walk-based embedding, inspired by Node2Vec [17]. Our implementation differs from the original Node2Vec in several key aspects:

Unbiased Random Walks: Unlike Node2Vec, which uses biased random walks controlled by parameters p and q , our implementation uses simple unbiased random walks. For each node v , we perform a random walk of fixed length:

$$walk(v) = (v_0 = v, v_1, \dots, v_l) \quad (5)$$

where l is the walk length, and each v_{i+1} is chosen uniformly at random from the neighbors of v_i . **Fixed Walk Length**: Our walks have a fixed length equal to the number of nodes in the graph, whereas Node2Vec typically uses shorter, parameterized walk lengths. **Embedding Generation**: We use Word2Vec [18] to generate embeddings from these walks, similar to Node2Vec. For each node v , we compute an embedding $\phi(v) \in \mathbb{R}^d$, where d is the embedding dimension.

B. Cost Matrix Definition

This simplified approach retains the ability to capture some higher-order neighborhood information while being computationally less expensive than the full Node2Vec algorithm. The resulting embeddings contribute to the cost matrix used in the optimal transport problem:

$$C_{ij} = \dots + w_n \|\phi_i - \phi_j\|_2 + \dots \quad (6)$$

where w_n is the weight for the embedding component in the cost matrix. Although this approach loses some of the fine-grained control over the exploration-exploitation trade-off that Node2Vec offers, it provides a computationally efficient way to incorporate higher-order structural information into our graph matching process.

In short, the proposed cost matrix C for Graph Optimal Transport is constructed by combining the different feature representations:

$$C_{ij} = w_s \|s_i - s_j\|^2 + w_f \|f_i - f_j\|^2 + w_n \|\phi_i - \phi_j\|^2 + w_p \|p_{rel,i} - p_{rel,j}\|^2 \quad (7)$$

where s_i, s_j are spectral embeddings ; f_i, f_j are node descriptor features ; ϕ_i, ϕ_j are Node2Vec embeddings ; $p_{rel,i}, p_{rel,j}$ are relative positions ; w_s, w_f, w_n, w_p are weights for each component.

C. Hierarchical Matching

The hierarchical matching approach employs a sophisticated node neighborhood analysis to refine graph matching results. The core innovation lies in using a "hop-based descriptor" strategy that captures the local topological context of each node beyond simple structural features. This approach involves neighborhood exploration, where the algorithm computes descriptors by counting the number of unique neighbors within progressively larger hop distances (1-hop, 2-hop, etc.).

The process begins by solving the Optimal Transport problem using the Earth Mover's Distance (EMD):

$$\min_{\pi \in \Pi(p,q)} \sum_{i,j} \pi_{ij} C_{ij} \quad (8)$$

where p and q are uniform distributions over the nodes of the two graphs. Following this, we extract the initial matches greedily from the transport plan π .

After initial matching using optimal transport, the algorithm applies a similarity criterion. It checks the hop-based descriptors of matched node pairs, retaining nodes only if their neighborhood structures are sufficiently similar, using a similarity threshold. This refinement process effectively filters out potentially incorrect initial matches by ensuring that matched nodes have comparable local graph structures, not just similar individual node properties. Let $H(n, k)$ represent the number of unique neighbors of node n within k hops. The hop similarity is defined as:

$$\text{Hop Similarity} = 1 - \frac{\|H(n_1, \text{max_hops}) - H(n_2, \text{max_hops})\|}{\text{max_hops}} \quad (9)$$

By incorporating this hierarchical matching strategy, the algorithm enhances matching precision by considering not just individual node characteristics, but their broader topological context within the graph. The method is particularly valuable in scenarios with complex graph structures where simple one-to-one node matching might fail to capture intricate structural relationships. This approach allows for both local and global consistency in matching, addressing the multi-scale nature of plant structures and providing a robust method for comparing and aligning plant structures represented as graphs.

D. RANSAC Refinement

In order to incorporate possible geometric changes between the plant structures, we utilize a RANSAC-based refinement step. Initially, we used preliminary matches to approximate the transformation between the two graphs. The transformation is then applied to the unmatched nodes. In addition, additional matches are identified on the basis of proximity within the transformed space.

V. DATASET

To evaluate graph matching algorithms, we propose a novel 2D plant graph dataset. The dataset consists of 50 two-dimensional plant graphs, where each graph $G = (V, E)$ represents a plant structure. Vertices V represent branch junctions or endpoints, while edges E represent the physical connections

between these points. Each vertex contains spatial coordinates (x, y) . To maintain compactness and enhance computational efficiency, the graphs are restricted to at most 25 vertices.

To evaluate algorithm performance under different scenarios commonly encountered in plant analysis, we apply four types of transformations, which are shown in Figure 3:

- 1) **Rotation:** Each base graph is rotated at angles $\theta \in \{0^\circ, 10^\circ, 20^\circ, \dots, 350^\circ\}$ around its centroid.
- 2) **Deformation:** Simulates natural plant movement through local perturbations and global warping.
- 3) **Cutting:** Simulates partial plant structures by removing sections of the graph while preserving the largest connected component.
- 4) **Node Skipping:** Simulates missing or occluded nodes by removing nodes and reconnecting edges to the nearest remaining nodes.

To summarize, the complete dataset includes:

- 50 base graphs
- 1750 rotated variants (50 graphs \times 35 rotation angles)
- 50 deformed variants
- 50 cut variants
- 50 skipped node variants

VI. RESULTS

The evaluation encompasses 50 distinct plant structures, each represented as a graph with varying complexity and structural characteristics. Our primary metric is node matching accuracy, which quantifies the proportion of correctly identified correspondences between nodes in pairs of plant graphs. This metric directly reflects the algorithm's ability to identify and match corresponding structural elements across different plant representations. Our experimental evaluation demonstrates the effectiveness of the proposed Graph Optimal Transport (Graph OT) framework for plant structure alignment. By conducting thorough evaluations on a varied dataset of plant configurations, we measure the performance of our approach relative to a traditional Optimal Transport (OT), ICP and PRBoNN [12].

A. 2-D Plant Registration with Graph OT

We evaluated the performance of our Graph Optimal Transport (Graph OT) framework in matching plant structures under rotation transformations. As shown in Figure 4, our method demonstrates robust performance when matching original plant structures to their rotated counterparts. The accuracy distribution across all plants reveals strong matching capabilities, with the majority of plants achieving accuracy scores above 0.7. The Graph OT framework consistently outperforms the baseline OT method, which struggles to establish correct correspondences under rotation (average accuracy ≈ 0.1).

The distribution plot indicates that while performance varies across different plant structures, the Graph OT framework maintains reliable matching capabilities. This robust performance under rotation transformations is particularly noteworthy as it establishes a strong foundation for handling




Fig. 3: All the different transformation types applied on the point clouds.




Fig. 4: Accuracy Distribution across all plants: Original \Rightarrow Rotated

more complex structural variations in plant matching scenarios. Beyond rotation, we also evaluated our framework on three additional transformation types: deformation, cutting, and node skipping. The Graph OT framework maintains strong performance across all transformation types, achieving average accuracies of 0.74 for rotated structures, 0.67 for cut structures, and 0.71 for skipped node transformations.

The experimental results reveal a substantial performance advantage of our Graph OT framework over the baseline OT method. As illustrated in Figure 4, our approach consistently achieves significantly higher accuracy scores across the entire dataset. The Graph OT method demonstrates strong performance, with average accuracy values of 0.75 across different plant structures. Representing a marked improvement over the baseline OT method.

B. Comparisons with State-of-the-Art Methods

Table I presents a comprehensive comparison between our Graph OT framework and two state-of-the-art methods: ICP

and PRBoNN. The results demonstrate that Graph OT consistently outperforms both alternatives in most transformation types, with particularly notable improvements in handling deformed and cut structures. While ICP shows competitive performance in rotation scenarios, our method maintains superior performance in more complex transformations where traditional approaches struggle.

The performance distribution exhibits several notable characteristics:

- The Graph OT algorithm frequently achieves near-perfect matching scores (accuracy ≈ 1.0) for numerous plant structures, particularly evident in standard configurations.
- Even in challenging cases, the Graph OT framework maintains acceptable performance levels with accuracy scores consistently above 0.4.
- The integration of multiple feature types enables robust performance across varying plant structures.

The stark contrast in performance can be attributed to

TABLE I: Comparison of matching accuracy ($mean \pm std$) across different algorithms and transformation types.

Method	Original ↓ Rotated	Rotated ↓ Deformed	Rotated ↓ Cut	Rotated ↓ SkippedNode
OT	0.06 ± 0.05	0.06 ± 0.05	0.08 ± 0.07	0.07 ± 0.05
Graph OT (Ours)	0.75 ± 0.20	0.74 ± 0.21	0.67 ± 0.24	0.71 ± 0.18
ICP	0.67 ± 0.14	0.61 ± 0.12	—	—
PRBoNN [12]	0.33 ± 0.42	0.14 ± 0.24	0.21 ± 0.30	0.27 ± 0.39

several key methodological innovations. Our integration of multiple feature types, including node descriptors, spectral embeddings, and relative positions, enables the framework to capture both local and global structural properties effectively. This comprehensive feature representation, combined with our hierarchical matching approach, proves to be particularly advantageous when handling complex plant structures.

C. Ablation study of the Graph OT

To fully comprehend the role of each element in our Graph OT framework, we executed a comprehensive ablation study. We evaluated the performance of different method variants by removing or altering specific features to understand each component’s individual and collective significance. The analysis yielded detailed insights into each feature’s contribution to the alignment of plant structures. The results of the ablation study are summarized in Table II.

Excluding node descriptor features, which encapsulate both local and global graph-theoretical properties, resulted in a marked decrease in accuracy, particularly for intricate plant structures. This highlights the crucial role of detailed node-level data in distinguishing structurally similar branches.

Spectral embeddings, which capture fundamental topological aspects, showed a significant impact when omitted. The absence of a low-dimensional spectral representation, which captures global structural similarities, led to a noticeable performance deterioration.

The omission of Node2Vec embeddings, which capture higher-order neighborhood information, marginally improved the matching accuracy, as can be seen in Table II. This suggests that such embeddings are not essential for simple graphs; nevertheless, they contribute to improved matching accuracy in more complex graphs.

The features of the spatial relationship, denoted by the relative positions of the nodes, were also identified as crucial. Removing these geometric layout features greatly compromised the framework’s performance, emphasizing the importance of spatial relationship inclusion in distinguishing structurally similar but geometrically distinct branches. The hierarchical matching approach, which integrates global alignment with detailed local refinement, proved to be superior, clearly exceeding simpler one-level matching methods. This multi-scale strategy effectively tackles the complex, multi-dimensional challenges in plant structure comparisons.

The findings from the ablation study highlight the synergistic benefits of the Graph OT framework components. By integrating node descriptors, spectral embeddings, Node2Vec

representations, spatial data, and hierarchical matching, our method effectively captures both local and global structural attributes. This comprehensive approach enables robust alignment of plant structures across a variety of transformational scenarios, from rotational changes to partial structural distortions. These results not only support our proposed methodology, but also offer valuable insights for future advancements. Researchers interested in graph-based plant structure analysis can take advantage of these component-wise insights to create more advanced alignment and comparison methods.

VII. DISCUSSION

While our method demonstrates strong overall performance, certain patterns emerge in the error cases that warrant discussion. The higher variance observed in lower performing plants suggests that specific structural configurations remain a challenge for our approach. This is particularly evident in the left portion of the accuracy distribution, where the method shows increased variability in matching results.

The presence of occasional outliers in the performance distribution indicates specific instances where the matching process encounters difficulties. These cases often correspond to particularly complex plant structures or situations where the structural similarity between different parts of the plant leads to ambiguous matching scenarios.

The stability of our method is demonstrated through the consistent performance across different plant structures and conditions. Trend analysis reveals that the method maintains relatively consistent performance levels with fewer significant drops in accuracy. This stability is particularly important for applications that require reliable alignment of plant structure.

In conclusion, our experimental results demonstrate that the proposed Graph OT framework provides a robust and effective solution for plant structure alignment. Significant performance improvements over traditional approaches, combined with consistent accuracy across various plant structures, establish our method as a viable tool for practical applications in plant matching.

VIII. CONCLUSION

Our Graph OT framework introduces a novel approach to plant structure matching, addressing critical challenges in computational botany and plant structure analysis. By developing a sophisticated method that integrates multiple feature representations and a hierarchical matching strategy, we have demonstrated a significant advancement in the accurate alignment and comparison of plant structures across various transformational scenarios.

The core strength of our approach lies in its comprehensive feature integration. By combining node descriptors, spectral embeddings, Node2Vec representations, and relative positioning, we have created a robust framework capable of capturing both local and global structural characteristics of plant architectures. The ablation study provides critical insights into the contribution of each feature type, highlighting the synergistic nature of our multidimensional approach.

Enhanced Descriptors	Node2Vec	Hierarchical Matching	Position Cost	Mean Accuracy	Std Deviation
True	False	True	True	0.8127	0.2020
True	True	True	True	0.7932	0.1720
True	True	True	False	0.7897	0.1715
True	False	False	True	0.7863	0.2117
True	False	True	False	0.7856	0.2004
True	True	False	True	0.7793	0.1756
True	True	False	False	0.7686	0.1794
False	True	True	True	0.7611	0.1920
True	False	False	False	0.7560	0.2095
False	True	True	False	0.7450	0.1970
False	True	False	True	0.7366	0.1936
False	True	False	False	0.7215	0.2010
False	False	True	True	0.7090	0.1993
False	False	True	False	0.6797	0.1925
False	False	False	True	0.6639	0.2036
False	False	False	False	0.6298	0.1934

TABLE II: Ablation Study Configuration Performance




Fig. 5: Graph of a symmetric plant, with the target rotated by 180 deg

Our experimental results underscore the effectiveness of the framework in different transformation types. With node matching accuracies of 0.75 for rotated structures, 0.74 for deformed structures, 0.67 for cut structures, and 0.71 for structures with skipped nodes, the method consistently outperforms traditional optimal transport and other state-of-the-art techniques. This performance demonstrates the framework’s capability to manage complex structural variations, establishing it as a useful tool for practical applications in plant analysis and plant matching.

However, the method is not without limitations. Performance can be compromised with extreme rotations ($> 90^\circ$) and highly symmetric structures, as can be seen in Figure 5. Furthermore, computational complexity increases significantly with graph size, potentially limiting applications to very large plant structures (> 1000 nodes).

This study suggests several promising directions for future research. First, extend this framework to the matching of three-dimensional plant structures. Second, scale it up to accommodate larger plant structures. Third, incorporate domain-specific botanical constraints to enhance matching accuracy. An additional crucial step is the extraction of graphs from point cloud data generated by sensors.

The promising results of our Graph OT framework open

new avenues for computational plant analysis. By providing a powerful tool for structural alignment, our approach can significantly contribute to the field of computational botany and its practical applications in agriculture and botanical research.

REFERENCES

- [1] C. Godin, “Representing and encoding plant architecture: A review,” *Annals of Forest Science*, vol. 57, no. 5, pp. 413–438, 2000. DOI: 10.1051/forest:2000132.
- [2] G. Peyré and M. Cuturi, “Computational optimal transport,” *Foundations and Trends in Machine Learning*, vol. 11, no. 5-6, pp. 355–607, 2019.
- [3] T. Vayer, L. Chapel, R. Flamary, R. Tavenard, and N. Courty, “Optimal transport for structured data with application on graphs,” *arXiv preprint arXiv:1805.09114*, 2019.
- [4] G. Monge, “Mémoire sur la théorie des déblais et des remblais,” *Histoire de l’Académie Royale des Sciences de Paris*, 1781.
- [5] L. Kantorovitch, “On the translocation of masses,” *Management Science*, vol. 5, no. 1, pp. 1–4, 1958, ISSN: 00251909, 15265501. [Online]. Available: <http://www.jstor.org/stable/2626967> (visited on 10/11/2024).
- [6] Y. Rubner, C. Tomasi, and L. J. Guibas, “The earth mover’s distance as a metric for image retrieval,” *International journal of computer vision*, vol. 40, no. 2, pp. 99–121, 2000.
- [7] F. Mémoli, “Gromov–wasserstein distances and the metric approach to object matching,” *Foundations of computational mathematics*, vol. 11, no. 4, pp. 417–487, 2011.
- [8] H. W. Kuhn, “The hungarian method for the assignment problem,” *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [9] M. Cuturi, “Sinkhorn distances: Lightspeed computation of optimal transport,” in *Advances in neural information processing systems*, 2013, pp. 2292–2300.

- [10] J. Altschuler, J. Weed, and P. Rigollet, “Near-linear time approximation algorithms for optimal transport via sinkhorn iteration,” *arXiv preprint arXiv:1705.09634*, 2017.
- [11] C. Godin and Y. Caraglio, “A multiscale model of plant topological structures,” *Journal of Theoretical Biology*, vol. 191, no. 1, pp. 1–46, 1998. DOI: 10.1006/jtbi.1997.0561.
- [12] N. Chebrolu, F. Magistri, T. Labe, and C. Stachniss, “Registration of spatio-temporal point clouds of plants for phenotyping,” *PloS one*, vol. 16, no. 2, e0247243, 2021.
- [13] H. P. Maretic, M. El Gheche, G. Chierchia, and P. Frossard, “Got: An optimal transport framework for graph comparison,” *arXiv preprint arXiv:1906.00200*, 2019.
- [14] H. P. Maretic, M. E. Gheche, G. Chierchia, and P. Frossard, “Fgot: Graph distances based on filters and optimal transport,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 7, pp. 7710–7718, 2022. DOI: 10.1609/aaai.v36i7.20738.
- [15] L. Lobefaro, M. V. R. Malladi, O. Vysotska, T. Guadagnino, and C. Stachniss, “Estimating 4d data associations towards spatial-temporal mapping of growing plants for agricultural robots,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 4212–4218. DOI: 10.1109/IROS55552.2023.10342449.
- [16] L. Lobefaro, M. Malladi, T. Guadagnino, and C. Stachniss, “Spatio-temporal consistent mapping of growing plants for agricultural robots in the wild,” in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2024.
- [17] A. Grover and J. Leskovec, *Node2vec: Scalable feature learning for networks*, 2016. arXiv: 1607.00653 [cs.SI]. [Online]. Available: <https://arxiv.org/abs/1607.00653>.
- [18] T. Mikolov, K. Chen, G. Corrado, and J. Dean, *Efficient estimation of word representations in vector space*, 2013. arXiv: 1301.3781 [cs.CL]. [Online]. Available: <https://arxiv.org/abs/1301.3781>.