

DynamicRoots: a Software Platform for the Reconstruction and Analysis of Growing Plant Roots

Olga Symonova¹, Christopher N. Topp², Herbert Edelsbrunner¹,

1 IST Austria, Klosterneuburg, Austria

2 Donald Danforth Plant Science Center, St. Louis, Missouri, USA

Abstract

We present a software platform for reconstructing and analyzing the growth of a plant root system from a time-series of 3D voxelized shapes. It aligns the shapes with each other, constructs a geometric graph representation together with the function that records the time of growth, and organizes the branches into a hierarchy that reflects the order of creation. The software includes the automatic computation of structural and dynamic traits for each root in the system enabling the quantification of growth on fine-scale. These are important advances in plant phenotyping with applications to the study of genetic and environmental influences on growth.

Introduction

Root system architecture plays a key role in plant fitness and crop productivity. Climate change, a growing global population, and unsustainable use of fertilizers are responsible for the pressing need to understand how root systems grow and interact with their environment, with the goal to develop robust and efficient crops [1–3]. Studying plant roots systems has been hindered by their complex three-dimensional branching topology, environmental growth plasticity, as well as the opacity of the soil. In recent years, however, various imaging technologies have been applied to monitor the formation and development of roots non-disruptively and throughout a time period: X-ray [4–8], MRI [9], PET-MRI [10], laser [11], and optical imaging [12,13]. These technologies reconstruct the 3D shape of a root system and enable the assessment of traits that were impossible to quantify by hand measurements before [5,12–15]. Most of the existing tools compute traits for static shapes, and if more than one shape representing the same root system is available, the global traits are computed without separating the contributions of individual branches [16]. Software that analyzes a time-series of 2D root images but requires significant user input to track growth has been introduced in [17,18]. Several high-throughput platforms were developed for the analysis of video clips of growing dicot roots [19–21]. Advances in the fields of computer graphics and computer vision make it possible to quickly align large shapes that are at least partially similar [22,23]. These methods can be used to improve the analysis of time-series but have not yet been exploited for plant phenotyping.

In this paper, we present the `DynamicRoots` software designed to automatically process a time-series of 3D reconstructions of a growing root system. It registers the 3D shapes in a common coordinate system and creates a data structure that records the

growth. Using this data structure, we estimate dynamic traits of individual branches, which we aggregate to global traits. Our contributions to the state-of-the-art are:

- a geometric graph representation of a root system that reflects its hierarchical structure;
- an algorithm that guarantees temporal coherence of the hierarchical decomposition and computes the time function recording the growth process;
- a suite of structural and dynamic root traits including elongation rates and branching frequency.

The computed traits facilitate the analysis of root systems in unprecedented resolution of time and space. We illustrate the software by analyzing several time-series of reconstructed rice and maize roots.

Outline. The Methods section gives a detailed description of the algorithm that constructs the growth record for a time-series of 3D root shapes, and lists the structural and dynamic traits computed from this record. The Results section explains how we validate the software and presents experimental results for rice and maize root systems. We conclude the paper with a discussion of current limitations and future research directions.

Methods

The goal of this work is to acquire the shape of a growing root system, to compactly represent it in a unified data structure, and to compute traits that reflect the structural and dynamic properties of the root system. We divide the process into five steps:

1. Take 2D images of root systems and construct a time-series of 3D shapes.
2. Align the shapes to define depth as the geodesic distance from consistently detected seed areas.
3. Use the depth function to decompose each shape into a hierarchy of branches.
4. Construct the time function and reorganize the hierarchy by repairing inconsistencies between depth and time of growth.
5. Use the time function on the branch hierarchy to compute structural and dynamic traits of the root system.

Correspondingly, we have five subsections explaining the steps in appropriate detail. Step 1 relies on prior work presented elsewhere, while the other four steps are novel contributions reported for the first time in this paper. Steps 2 to 4 construct the growth record, and Step 5 makes use of this unifying data structure.

Background and Assumptions

We review the prior work to the extent necessary to provide the context for our work. To overcome the obstacle of opaque soil, we grow rice and maize plants in gel with nutrients inside transparent cylindrical containers [13], monitoring the development of the root system for several weeks. To construct the 3D shape of a root system, we place the container on a rotating table connected to a camera and a computer, and we take 40 images at 9° increments. Currently, 20 of these images are used to reconstruct the shape of the root system represented as a collection of voxels in 3D space [24]. Letting t be the time of the acquisition, we write \mathbb{V}_t for the reconstructed shape. Depending on the purpose, we will think of \mathbb{V}_t as a subset of \mathbb{R}^3 , a collection of voxels, or a graph

whose nodes are the voxels and whose arcs connect neighboring voxels. Drawn as straight line segments ending at voxel centers, the arcs have length 1, $\sqrt{2}$, or $\sqrt{3}$ times the length of a voxel edge.

In order to analyze the growth of a root system, the same plant is reconstructed at different moments in time. We thus get a time-series of shapes \mathbb{V}_t , for $t = 0, 1, \dots, T$, which provides a much richer source of information about the root system than a single such shape. However, harvesting this information is more difficult since we need to compare, quantify similarities and differences, and resolve occasional contradictions. To do this automatically within our software platform, we make use of the following three assumptions:

- A₁: A root system grows in only two ways: by generating a new branch at a fork, and by elongating an existing branch at the tip.
- A₂: A root system is connected and contractible. In other words, the branches do not form loops, and they do not surround voids in space.
- A₃: The time-series is dense enough to observe the correct root hierarchy: when a side branch first appears at a fork, it is shorter than the parent branch.

While we have biological reasons to believe that A₁ and A₂ are true, both assumptions are sometimes violated by the reconstructed shapes. Our software can tolerate mild violations of all of the three assumptions but will fail when the reconstructed shapes contain gross contradictions to our expectations.

Depth, Seed, and Alignment

In this and the next two subsections, we explain how to construct the *growth record* of a root system, which provides a decomposition into a hierarchy of branches together with the time of growth. As a first step, we equip each reconstructed shape with a function $\varphi_t: \mathbb{V}_t \rightarrow \mathbb{R}$ that maps each voxel to its geodesic distance from a *seed area*, a concept we will explain shortly. We call this distance the *depth* of the voxel, and we call φ_t the *depth function* of \mathbb{V}_t . After fixing a set of voxels $S_t \subseteq \mathbb{V}_t$ as the seed area, we compute φ_t using Dijkstra's shortest path algorithm [25]. Running this algorithm on the graph of voxels is straightforward, but the result crucially depends on the choice of S_t , which is inspired by the biological idea that the root grows from a particular first location. It will be more important that the location of this area is consistent among all shapes in the time-series than to correctly identify the biological beginnings of the root system. We therefore compute S_T , and propagate its location to the other shapes after computing an initial alignment of all shapes in the time-series. We will explain the alignment first and return to computing the seed area thereafter.

Note that Assumption A₁ implies $\mathbb{V}_s \subseteq \mathbb{V}_t$ for all $0 \leq s \leq t \leq T$. To align the two shapes, we can therefore use an algorithm that requires that the first shape be contained in the second. Our algorithm of choice is the *4-point congruent sets (4PCS) algorithm* described in [22]. To apply it, we think of \mathbb{V}_s as a collection of points in \mathbb{R}^3 , namely the centers of its voxels. The algorithm selects four approximately coplanar points from \mathbb{V}_s and searches for a congruent configuration of four points in \mathbb{V}_t . Let M be the matrix of the rigid motion that moves the 4-point configuration in \mathbb{V}_t to that in \mathbb{V}_s . While the 4PCS algorithm is fast, the alignment it produces can often be improved, and we use the *iterative closest point (ICP) algorithm* originally described in [26] for this purpose. It computes a second matrix M' , and we set $M_{st} = M' \cdot M$ so that $M_{st}(\mathbb{V}_t)$ is aligned with \mathbb{V}_s . To align all $T + 1$ shapes in the time-series, it suffices to compute T rigid motions, and we choose to compute M_{0t} for $t = 1, 2, \dots, T$. To align \mathbb{V}_t with \mathbb{V}_s , we then use $M_{st} = M_{0s}^{-1} \cdot M_{0t}$. Suppose now that we have computed the seed

area in the last shape, $S_T \subseteq \mathbb{V}_T$. Using the rigid motions, we get $S_t \subseteq \mathbb{V}_t$ as the set of voxels closest to $M_{tT}(S_T)$, for $0 \leq t \leq T$. 116

To detect the seed area in \mathbb{V}_T , we perform *principal component analysis* (PCA) on a small neighborhood of a radius R around every voxel. Covariance matrices with three similar eigenvalues indicate a similar spread of points along the principal axes, which in 3D space corresponds to a spherical neighborhood. We find the biggest connected component of voxels with roughly spherical neighborhoods and form the seed area by adding voxels with offset at most R from this component. For the results reported in this paper, we used $R = 3$ and we suggest that this parameter is set to a value from the interval $[r, 2r]$, in which r is the average radius of the branches. If the seed was not detected – for example when the set of voxels with roughly spherical neighborhoods is empty or larger than an allowed fraction of the total volume – we choose a voxel in the center of the topmost horizontal slice of the root system as the seed area. 117
118
119
120
121
122
123
124
125
126
127
128

Branch Hierarchy 129

As a second step in the construction of the growth record, we decompose each shape in the time-series into branches, which we organize hierarchically. This decomposition is inspired by the path decompositions of trees used in the design of efficient data structures; see e.g. [27, Chapter 5]. Each branch starts at its tip at the bottom and extends upward, ending right before its fork, which lies somewhere between the two ends of the parent branch; we therefore represent each branch by the triplet of these entities: $\mathbb{B}^i = (\text{Tip}^i, \text{Fork}^i, \text{Parent}^i)$. An exception is the topmost branch, which does not have a parent. Among all possible branch decompositions, we are interested in the one defined by the depth function. 130
131
132
133
134
135
136
137
138

To construct the decomposition, we traverse the voxels in the order of decreasing depth, adding an arc right after its second endpoint. With each visited voxel, we create a new component, and with each visited arc, we either form an additional connection within the component, or we merge two components into one. After visiting all voxels and arcs, they are part of a single component representing the entire root system, and we recover the branches and the relations between them by analyzing the history of events. To make this concrete, we equip each voxel, v , with a pointer to another voxel, $\rho(v)$. In the beginning, each voxel points to itself. After completion of the algorithm, each voxel points to the tip of its branch, except for the tip itself, which points to the tip of the parent branch. A *tip* is therefore identified by having other voxels point to it, a *branch* is represented by its tip and consists in addition of all voxels pointing to this tip, and the *parent branch* is given by the pointer stored at the tip. 139
140
141
142
143
144
145
146
147
148
149
150

The only interesting event during the course of the algorithm is when it encounters an arc that connects two components. Let its endpoints be the voxels u and v , and write $p = \rho(u)$ and $q = \rho(v)$ for the corresponding tips, noting that $p \neq q$ by assumption. Assume without loss of generality that q is deeper than p : $\varphi(p) < \varphi(q)$. If the difference in depth between u and p is smaller than some fixed threshold, then we declare the entire branch to be a part of the branch of q by updating the pointer: $\rho(p) = q$. Otherwise, we consolidate the branch of p by declaring v as its fork and the branch of q as its parent. 151
152
153
154
155
156
157
158

As mentioned earlier, each branch is represented by its tip, p , and consists of all voxels u with $\rho(u) = p$, together with p . To avoid possible confusions arising from the two different uses of the ρ -pointer, we save the pointer to the parent branch by setting $\text{Parent}(p) = \rho(p)$, and thereafter set $\rho(p) = p$ to indicate that it belongs to the branch it represents. In addition, we store a pointer $\text{Fork}(p)$ to the fork in the parent branch. The thus linked branches form what we call a *branch hierarchy*. At this moment, the hierarchy is determined by depth: a branch has a tip that is necessarily less deep than the tip of the parent branch. Ultimately, we would like the hierarchy be determined by 159
160
161
162
163
164
165
166

age. In the next section, we will explain how the hierarchy can be modified to achieve this goal.

Time Function

We are now ready to compute the time function of the last shape in the time-series, $\tau: \mathbb{V} \rightarrow \mathbb{R}$ with $\mathbb{V} = \mathbb{V}_T$. Compressing the information of the time-series into a single shape, it maps every voxel of \mathbb{V} to the moment in time the voxel was created. The time function is piecewise constant, with possibly large steps in particular but not exclusively at forks. We note that the time function is the central piece of the growth record that enables the dynamic analysis of the root system.

To prepare the computation, we initialize the time function to $\tau(v) = T$ for every voxel $v \in \mathbb{V}$. Traversing the other shapes in the time-series from back to front, we use the Euclidean distance between voxels to match the branches of \mathbb{V}_t with those of \mathbb{V} . It is therefore important that the shapes in the time-series are aligned and registered into a common coordinate system, which we henceforth assume. Specifically, for each tip in \mathbb{V}_t , we find the closest voxel $v \in \mathbb{V}$, and we match the branch of the tip with the branch \mathbb{J} in \mathbb{V} that contains v . Finally, we set $\tau(u) = t$ for all voxels u of \mathbb{J} whose depth satisfies $\varphi_T(u) \leq \varphi_T(v)$.

Figure 1. A switch at the fork v .

After running the algorithm, the set of voxels $u \in \mathbb{V}$ with $\tau(u) \leq t$ would ideally be the set of voxels in \mathbb{V}_t , for every $0 \leq t \leq T$. Of course, this will never be the case in practice, but there is a reason other than the inevitable noise in the system that may lead to discrepancies between the computed and the ideal time function, namely possible inconsistencies between the decompositions into branches between \mathbb{V}_t and \mathbb{V} . Instead of detecting inconsistencies between the branch decompositions, we repair the time function. Suppose $v \in \mathbb{V}$ is a fork at which the classification into main branch and side branch is incorrect. By Assumption A₃, we get a conflict between the depth and the time. Specifically, we observe neighbors u of v for which

$$\varphi(u) < \varphi(v) \quad \text{and} \quad \tau(u) > \tau(v). \tag{1}$$

If we detect such a fork, we repair the inconsistency by an operation we call a *switch*. Letting p and q be the tips of the two branches at v , with $v = \text{Fork}(p)$ but $v \neq w = \text{Fork}(q)$, we note that $q = \text{Parent}(p)$. To perform the switch, we reassign the part of the branch of q between v and w to the branch of p ; see Fig. 1. In addition, the time at any voxel u within this part is set to the minimum of the current value and the value at v .

Root Traits

Using the growth record explained above, we compute a suite of traits that describe a root system captured in a time-series. Denoting the growth record by $\mathbb{V} = \mathbb{V}_T$, we make essential use of the decomposition into branches, which we denote as \mathbb{B}^i . Recall that $\mathbb{B}_t^i \subseteq \mathbb{B}^i$ is the sublevel set of the time function on the branch. We call the voxel $u \in \mathbb{B}_t^i$ with maximum depth the *tip* of the branch at time t . We compute the following traits for every branch \mathbb{B}^i and at every moment of time $0 \leq t \leq T$:

- volume:** the number of voxels in \mathbb{B}_t^i ;
- depth of tip:** the geodesic depth of the tip of \mathbb{B}_t^i ;
- location of tip:** the three Cartesian coordinates of the tip;

length: the difference between the geodesic depth values of the tip and the fork; 200
switch event: a logical variable that indicates whether or not \mathbb{B}_t^i is longer than its 201
parent branch; 202
tortuosity: the ratio between the length of a branch and the Euclidean distance 203
between tip and fork; 204
average radius: the square root of the volume divided by π times the length; 205
angle to gravity: the angle between the major PCA axis of \mathbb{B}_t^i and the vertical 206
direction; 207
angle to parent branch: the angle between the major PCA axes of \mathbb{B}_t^i and of its 208
parent branch; 209
number of children: the number of branches that have \mathbb{B}_t^i as parent branch. 210

There is an ambiguity about which part of a branch should be taken into account when 211
computing the root emergence angle or direction of growth. In order to avoid this 212
uncertainty, we use the major PCA axis of the branch which shows the main direction 213
of the spread of all the voxels forming this branch. 214

Only two of the measurements relate the root system to the surrounding, namely the 215
location of the tip and the angle to gravity. Both are useful to study the reaction of the 216
root system to nutrients and other local environmental conditions. The measurements 217
we output for each branch can of course be accumulated to obtain overall traits of the 218
root system, such as the **total volume**, the **total length**, the **total average radius**, 219
and so on. Similarly, we can take differences to compute dynamic traits, as for example 220
the **amount of growth** from time t to time $t + 1$. Beyond the traits listed above, we 221
compute the following characteristics of the seed area: 222

seed volume: the number of voxels in the seed area of the root system; 223
seed orientation: the angle between the major PCA axis of the seed area and the 224
vertical direction. 225

Comparing the seed volume with the total volume, we get the **volume fraction** of the 226
seed area. The concrete representation of the seed area is useful in turning the 227
geometric hierarchy of the branches into a biologically more meaningful one. The former 228
is defined by assigning the index 0 to the topmost branch and increasing the index by 1 229
whenever we pass from a branch to a child branch. A biologically more meaningful 230
classification is into 231

- *crown roots*, which emerge above the seed area; 232
- *seminal roots* with forks located within the seed area; 233
- the *primary root* corresponding to the topmost branch; 234
- *lateral roots*, which are the children of other branches. 235

Unfortunately, crown roots often emerge from the stem above the reconstructed space, 236
which results in disconnected branches. The reconstruction software used in this 237
study [24] enforces connectivity and tries to find the shortest path between the biggest 238
component of the root system and a disconnected branch. This shortest path does not 239
always join the branch in a biologically correct manner, making the classification into 240
biological classes more challenging. At the moment, we explore other ways to 241
differentiate between crown, seminal, and lateral roots. In particular, elongation rates 242
might provide a good clue for detecting crown roots as their initiation is developmentally 243
delayed but typically proceeds faster than of the other root types, see Fig. 2. 244

Figure 2. Structural and dynamical analysis of a root system. *Top row:* images of a maize plant imaged between the first and ninth days after planting. *Bottom row:* images of a rice plant imaged between the seventh and eleventh days after planting. *Left:* decomposition into a seed area and branches distinguished by color. *Center:* the color-coded time function. *Right:* the color-coded branch hierarchy.

Results

Applying our methods, we obtain insights into the growth behavior of rice and maize plants. We note, however, that this is primarily a Methods paper, so we have limited the experiments to a few examples selected to illustrate useful features of the `DynamicRoots` software. The source code and the pre-compiled executable file for 64-bit Windows can be downloaded from <http://dynamicroots.sourceforge.net/>. The software has a command-line interface. In addition, we provide a python script, which organizes the input files and passes them as arguments to scripts that perform format conversion, data alignment, and time-series processing. The `DynamicRoots` software was developed in C++ using the Approximate Nearest Neighbor (<http://www.cs.umd.edu/~mount/ANN/>) and ALGLIB (<http://www.alglib.net/>) libraries. `DynamicRoots` is an open source software distributed under the terms of the BSD license.

Datasets and Validation

We use the software to analyze four different time-series of 3D shapes, three obtained by imaging growing rice and maize plants, and one artificially created to validate the software. The latter consists of three models forming a caricature of a growing root system; see Fig. 3. Created from a collection of discretized algebraic curves, the models

Figure 3. A series of three 3D models mimicking a growing root system designed to facilitate the validation of our software. The branches are labeled as in Table 1. The color-coded model at the far *right* illustrates the time function.

are obtained by sweeping balls of varying radii along the curves to get branches with prescribed thickness. Observe that the time-series satisfies Assumptions A_1 to A_3 stated in Section . After designing the models, we 3D-printed them from resin, imaged the three shapes, reconstructed the 3D shapes from the 2D images, and finally processed the result with our software. Table 1 compares the hand measurements of *branch length*, *depth of tip*, *volume*, *average radius*, and *tortuosity* with the values computed by our software. The agreement in branch length and depth of tip is remarkably high, with a somewhat larger relative error only for very short branches. The tortuosity of a branch depends on the estimation of its length and agrees fairly well with our hand measurements, with only 6% average relative error. The volume and therefore the average radius are slightly underestimated by the software because the reconstruction of the 3D shapes tends to lose some of the thickness due to small misalignments of the 2D images. Nonetheless, the total volume estimated by our software differs on average only by 5.2% from the volume computed as a product of the mass of the printed model and the resin density. In order to validate the values of the angle to parent branch, we represent the each branch as a line segment passing through the tip and the fork and compute the angle between them. We note that the software computes the angle between the branches using PCA and it is impossible to measure this value by hand. Nevertheless, we observe a similar behavior in the values measured by hand and computed with the software: the angles decrease as branches elongate. Unfortunately,

there is no a common way to measure branch angles, and `DynamicRoots` offers a new approach to estimate this trait.

283
284

Table 1. Comparison of traits measured by hand and computed with `DynamicRoots`. Length, vertical depth, and radius are given in millimeters (mm), volume in cubic millimeters (mm³), angle in degrees, and tortuosity is dimensionless. The last column gives the average error of the `DynamicRoots` results compared to hand measurements.

Trait	Hand measurements			DynamicRoots results			Error	
	<i>t</i> = 0	<i>t</i> = 1	<i>t</i> = 2	<i>t</i> = 0	<i>t</i> = 1	<i>t</i> = 2		
branch r0	Length	63.800	71.100	95.220	64.068	72.275	98.383	0.018
	Depth	63.800	71.100	95.220	62.083	69.683	95.348	0.016
	Volume	450.976	502.576	673.071	444.766	491.641	621.281	0.037
	Radius	1.500	1.500	1.500	1.488	1.473	1.418	0.027
	Tortuosity	1.000	1.000	1.000	1.030	1.040	1.030	0.033
	Angle to parent	0.000	0.000	0.000	0.000	0.000	0.000	0.000
branch r1	Length	35.770	76.110	102.680	36.198	84.020	112.870	0.072
	Depth	36.860	76.620	99.610	34.443	75.388	100.940	0.032
	Volume	112.375	239.107	322.579	96.219	203.297	248.516	0.174
	Radius	1.000	1.000	1.000	0.920	0.878	0.838	0.122
	Tortuosity	1.140	1.050	1.058	1.120	1.120	1.120	0.048
	Angle to parent	41.957	27.139	25.140	27.600	20.360	18.700	0.283
branch r2	Length	42.980	66.770	99.060	46.173	72.208	98.910	0.052
	Depth	36.530	55.890	78.230	32.820	53.085	74.885	0.065
	Volume	135.026	209.764	311.206	131.938	187.922	229.125	0.130
	Radius	1.000	1.000	1.000	0.953	0.910	0.858	0.093
	Tortuosity	1.075	1.136	1.237	1.120	1.160	1.180	0.036
	Angle to parent	57.645	49.580	40.981	55.200	43.570	36.840	0.088
branch r3	Length		7.000	37.550		4.335	39.833	0.221
	Depth		57.960	83.350		55.803	81.983	0.027
	Volume		17.224	92.395		8.516	54.703	0.457
	Radius		0.885	0.885		0.790	0.660	0.181
	Tortuosity		1.127	1.097		1.000	1.110	0.062
	Angle to parent		90.000	40.030		45.360	15.980	0.548
branch r4	Length			9.270			7.940	0.143
	Depth			66.480			65.203	0.019
	Volume			13.665			6.875	0.497
	Radius			0.685			0.525	0.234
	Tortuosity			1.150			1.070	0.070
	Angle to parent			38.432			15.49	0.597
branch r5	Length		4.300	27.830		3.538	32.415	0.171
	Depth		51.190	71.190		47.538	69.125	0.050
	Volume		6.339	41.025		2.484	19.609	0.565
	Radius		0.685	0.685		0.473	0.440	0.334
	Tortuosity		0.883	0.984		1.000	1.110	0.130
	Angle to parent		62.507	60.122		24.510	8.940	0.730
branch r6	Length	27.490	35.010	53.730	25.753	36.245	56.103	0.048
	Depth	30.260	36.970	54.370	24.633	33.623	50.628	0.115
	Volume	86.362	109.987	168.798	76.719	103.422	155.781	0.083
	Radius	1.000	1.000	1.000	0.975	0.953	0.940	0.044
	Tortuosity	1.180	1.113	1.130	1.070	1.100	1.110	0.041
	Angle to parent	49.957	46.420	35.918	43.190	36.200	27.890	0.193
total	Length	170.040	260.290	425.340	172.190	272.620	446.453	0.037
	Volume	784.738	1084.997	1622.737	749.641	997.281	1335.891	0.101
	Mass × Density	694.000	967.000	1397.00	749.641	997.281	1335.891	0.052

The first of the three natural datasets consists of 478 time-series of 3 shapes each, reconstructing root systems of growing rice plants on Days 12, 14, 16 after planting. Using static traits, this data has been used in [16] to study the genetic control of shape and development of the root architecture. We intend to extend this study with the dynamic traits and the information about the branch hierarchy available from our

285
286
287
288
289

software. The second natural dataset consists of 4 time-series of 25 shapes each, reconstructing root systems of growing rice plants imaged every two hours from 8am to 4pm for five days starting from Day 7 after planting. The third natural dataset consists of 3 time-series of 23 shapes each, reconstructing root systems of growing maize plants images from Day 1 to Day 9 after planting.

The relatively high temporal resolution in the last two datasets leads to detailed characteristics of the root development not available in the first dataset.

Traits

The traits described in Section are output in a file for each time-series; see the supplementary material, which includes the files for the maize plant shown in the top row of Fig. 2. For complete disclosure, our software also outputs files containing the color-coded geometry of the root system, representing the time function and the branch hierarchy, among other information; see Fig. 2.

We use the remainder of this subsection to demonstrate how `DynamicRoots` can be used to analyze growth dynamics. As a first example, we look at the dynamic addition of volume to the shapes for the two displayed time-series. Each vertical bar in Fig. 4 represents the root system at a moment in time, with the colored segments representing branches in the decomposition. The displayed behavior reveals a remarkable dynamics in the growth process: the volume of the maize root almost doubles in the course of just a few hours, with the bulk of the additional volume contributed by newly formed branches (see Fig. 4 top, between 72 and 96 hours).

Figure 4. Volume dynamics of the root systems in Fig. 2; maize at the top and rice at the bottom. The colored segments within each vertical bar show the volume contributed by individual branches.

As a second example, we look at the branching dynamics, and in particular at the difference between the number of branches formed during day- and night-time. Fig. 5 shows the number of branches as graphs over time for four rice plants. Observing the root systems during five consecutive days, we note that there are more branches formed during day-time than during night-time. The rice plants grew in the 12-hour cycle with the temperature of 28°C and 25°C at day and night. In [28] the authors studied diurnal elongation rates in *Arabidopsis thaliana* and observed higher elongation rates at night time than during the day. In [29] the authors demonstrated that growth conditions have a great impact on the diurnal growth patterns in sorghum and rice plants, they did not find differences in day-night elongation rates when the temperature was constant. In our data, we did not see differences in elongation rates at day and night time for plants in Fig. 5. While more experiments are needed to understand the branching dynamics of root systems, it is clear that `DynamicRoots` can be of valuable assistance in collecting and analyzing the data.

Figure 5. The branching dynamics of rice plants during five consecutive days. The shading distinguishes day- from night-time, highlighting that more branches are formed during day-time.

Performance

We distinguish between performance characteristics that pertain to the speed and the correctness of the software. In addition to the traits discussed above, the software outputs performance statistics, including the **running time**, the **average distance**

between aligned 3D shapes, and the **number of switch events**. For example, the average running time for one time-series from our first dataset, including file format conversion, graphical file output, shape alignment and processing, is 51 seconds on an Intel Core 3.4GHz desktop computer.

We use the remainder of this subsection to illustrate how we make use of the performance measurements. A delicate step in our algorithm is the alignment of the 3D shapes in a time-series. We underline the fact that the alignment performs quite accurately due to the property of the 4PCS algorithm which can robustly align shapes even with small overlap, noise, and outliers [22]. To shed light on the quality of the alignment, we compute the average distance between the aligned shapes and provide the result as an indicator of the quality of the achieved alignment. More precisely, we sample 10% of the voxels in the first shape, \mathbb{V}_0 , and we compute the Euclidean distance to the closest voxels in \mathbb{V}_1 to \mathbb{V}_T . For a high quality alignment, the average distance will be small, namely less than the edge length of a voxel. In the first dataset, there are several time-series for which the accuracy is compromised by technical difficulties with the imaging platform. For these series, the average distance we observe is significantly larger than for series that do not suffer from these imaging artefacts: e.g. 5.03 versus 0.84 times the length of an edge. In our analysis of the first dataset, we have used a threshold of 10.00 edge lengths for the average distance of the alignment beyond which a shape was excluded from the computations.

Another measurement that varies significantly with the quality of the alignment is the number of switch events. On average, we observe 9.04 switches between shapes in the first dataset. However, in the time-series affected by the mentioned image artefacts, this number increases to 58.96. The number of switches by itself can be an interesting biological trait as it measures how often a branch out-grows its parent branch. However, if the number is significantly higher than the average, then this may be reason for concern since it might be caused by technical difficulties in the acquisition pipeline rather than biological reality.

Discussion

The **DynamicRoots** software provides novel capabilities for the analysis of growing root systems of plants. The analyzed shapes come from plants grown in a gel substrate, but can be derived from any growth and imaging system. A logical next step is the genome-wide analysis of dynamic traits for root systems of agricultural plants. Given the availability of the decomposition into branches and the detailed per-branch information, it is time to study the local and global response of plants to localized nutrients in the environment.

While we hope that our software finds many users and contributes to our understanding of root systems, it is important to point out that there are pitfalls in using it. Most root systems we considered have 3D reconstructions with touching branches that lead to violations of our Assumption A_2 . As a consequence, there are loops in the reconstructed shape which leads to faulty decompositions into branches. Fortunately, such mistakes are local and do not affect branches away from the loops. Ideally, the software would have the built-in capability to remove loops by separating touching branches, but this is a challenging problem that awaits a satisfactory solution.

Nonetheless, the framework described here represents a major advance in root phenotyping methodology. For the first time, the growth dynamics of each branch in a complex three-dimensional root system can be automatically tracked and quantified over time. Thus, **DynamicRoots** will enable high-throughput analysis of local growth response to a wide-range of biotic and abiotic plant-environment interactions, as well as provide a powerful tool to answer fundamental questions such as: "How is root growth

regulated over the circadian clock?”, and “How do local growth decisions of each root tip contribute to global root architecture?”.

Acknowledgments

The first author thanks Niloy Mitra for sharing the 4CPS alignment code and Szymon Rusinkiewicz for sharing his implementation of the ICP algorithm. CNT thanks Ni Jiang for providing the 3D reconstructions of the validation models shown in Fig. 3.

Supporting Information

Root traits computed by DynamicRoots for the maize root shown in Fig. 2. The names and values of the traits are listed in columns, rows correspond to individual branches.

References

1. S. DE DORLODOT, B. FORSTER, L. PAGES, A. PRICE, R. TUBEROSA AND X. DRAYE. Root system architecture: opportunities and constraints for genetic improvement of crops. *Trends in Plant Science* **12** (2007), 474–481.
2. J. LYNCH. Root architecture and plant productivity. *Plant Physiol.* **109** (1995), 7–13.
3. M. TESTER AND P. LANGRIDGE. Breeding technologies to increase crop production in a changing world. *Science* **327** (2010), 818–822.
4. P. J. GREGORY, D. J. HUTCHISON, D. B. READ, P. M. JENNESON, W. B. GILBOY AND E. J. MORTON. Non-invasive imaging of roots with high resolution X-ray micro-tomography. *Plant and Soil* **255** (2003) 351–359.
5. C. E. HARGREAVES, P. J. GREGORY AND A. G. BENGOUGH. Measuring root traits in barley (*Hordeum vulgare* ssp *vulgare* and ssp *spontaneum*) seedlings using gel chambers, soil sacs, and X-ray microtomography. *Plant and Soil* **316** (2009), 285–297.
6. D. A. HEERAMAN, J. W. HOPMANS AND V. CLAUSNITZER. Three dimensional imaging of plant roots in situ with x-ray computed tomography. *Plant and Soil* **189** (1997), 167–179.
7. J. S. PERRET, M. E. AL-BELUSHI AND M. DEADMAN. Non-destructive visualization and quantification of roots using computed tomography. *Soil Biology & Biochemistry* **39** (2007),391–399.
8. S. R. TRACY, J. A. ROBERTS, C. R. BLACK, A. MCNEILL, R. DAVIDSON AND S. J. MOONEY. The X-factor: visualizing undisturbed root architecture in soils using X-ray computed tomography. *Journal of Experimental Botany* **61** (2010), 311–313.
9. A. H. VAN, T. SCHEENEN AND F. J. VERGELDT. MRI of intact plants. *Photosynthesis Research* **102** (2009), 213–222.
10. S. JAHNKE, M. I. MENZEL, C. VAN DUSSCHOTEN, G. W. ROEB, J. BÖHLER, S. MINWUYELET, ET AL. Combined MRI-PET dissects dynamic changes in plant structures and functions. *Plant Journal* **59** (2009), 634–644.
11. S. FANG, X. YAN AND H. LIAO. 3D reconstruction and dynamic modeling of root architecture in situ and its application to crop phosphorus research. *Plant Journal* **60** (2009), 1096–1108.
12. R. T. CLARK, R. B. MACCURDY, J. K. JUNG, J. E. SHAFF, S. R. MCCOUCH, D. J. ANESHANSLEY, ET AL. Three-dimensional root phenotyping with a novel imaging and software platform. *Plant Physiology* **156** (2011), 455–465.
13. A. IYER-PASCUZZI, O. SYMONOVA, Y. MILEYKO, Y. HAO, H. BELCHER, J. HARER, ET AL. Imaging and analysis platforms for automatic phenotyping and classification of plant root systems. *Plant Physiol.* **152** (2010), 1148–1157.

14. T. GALKOVSKIY, Y. MILEYKO, A. BUCKSCH, B. MOORE, O. SYMONOVA, C. A. PRICE, ET AL. GiA Roots: software for the high throughput analysis of plant root system architecture. *BMC Plant Biology* **12** (2012), 116–127.
15. M. P. POUND, A. P. FRENCH, J. ATKINSON, D. M. WELLS, M. J. BENNETT AND T. PRIDMORE. RootNav: Navigating images of complex root architectures. *Plant Physiol.* **164** (2013), 1802–14.
16. C. N. TOPP, A. S. IYER-PASCUZZI, J. T. ANDERSON, C.-R. LEE, P. R. ZUREK, O. SYMONOVA, ET AL. 3D phenotyping and quantitative trait locus mapping identify core regions of the rice genome controlling root architecture. *PNAS* **110** (18) E1695–E1704, 2013.
17. J. LE BOT, V. SERRA, J. FABRE, X. DRAYE, S. ADAMOWICZ AND L. PAGES. DART: a software to analyse root system architecture and development from captured images. *Plant and Soil* **326** (2010), 261–273.
18. G. LOBET, L. PAGES AND X. DRAYE. A novel image-analysis toolbox enabling quantitative analysis of root system architecture. *Plant Physiology* **157** (2011), 29–39.
19. A. FRENCH, S. UBEDA-TOMAS, T. J. HOLMAN, M. J. BENNETT AND T. PRIDMORE. High-throughput quantification of root growth using a novel image-analysis tool. *Plant Physiology* **150** (2009), 1784–1795.
20. N. D. MILLER, T. L. DURHAM BROOKS, A. H. ASSADI AND E. P. SPALDING. Detection of a gravitropism phenotype in glutamate receptor-like 3.3 mutants of *Arabidopsis thaliana* using machine vision and computation. *Genetics* **186** (2010), 585–593.
21. N. YAZDANBAKHSI AND J. FISAHN. High throughput phenotyping of root growth dynamics, lateral root formation, root architecture and root hair development enabled by PlaRoM. *Functional Plant Biology* **36** (2009), 938–946.
22. D. AIGER, N. J. MITRA AND D. COHEN-OR. 4-points congruent sets for robust pairwise surface registration. *ACM Trans. Graph.* **27** (2008), 85:1–85:10.
23. S. RUSINKIEWICZ AND M. LEVOY. Efficient variants of the ICP algorithm. *In Proc. of 3DIM* (2001), 145–152.
24. Y. ZHENG, S. GU, H. EDELSBRUNNER, C. TOMASI AND P. BENFEY. Detailed reconstruction of 3D plant root shape. *In “Proc. 13th Internat. Conf. Comput. Vision, 2011”*, 2026–2033.
25. E. W. DIJKSTRA. A note on two problems in connexion with graphs. *Numerische Mathematik* **1** (1959), 269–271.
26. P. J. BESL AND N. D. MCKAY. A method for registration of 3-D shapes. *IEEE Trans. Patt. Anal. Mach. Intel.* **14** (1992), 239–256.
27. R. E. TARJAN. *Data Structures and Network Algorithms*. SIAM, Philadelphia, Pennsylvania, 1983.
28. N. YAZDANBAKHSI, J. FISAHN. Stable diurnal growth rhythms modulate root elongation of *Arabidopsis thaliana*. *Plant Root* **5** (2011), 17–23.
29. M. IJIMA, Y. ORIBE, Y. HORIBE AND Y. KONO. Time lapse analysis of root elongation rates of rice and sorghum during the day and night. *Annals of Botany* **81** (1998), 603–607.