

Supplementary Figure S1. Basics of leaf architecture.
Silhouette of an Arabidopsis thaliana (A), sessile oak (B) and rose (C) leaf. The Arabidopsis and sessile oak leaves are simple leaves formed by a petiole supporting a single leaf blade which margin is dissected into small teeth in the case of Arabidopsis and larger lobes in the case of sessile oak. Rose has a compound leaf formed by several leaflets united by the rachis. The margin of the leaflets are dissected into numerous tiny serrations.


## Supplementary Figure S2. Morphometrics of leaves L11, L12 and L13.

Morphological parameters were extracted for young leaves L11, L12 and 13 based on biologically relevant landmarks determined by the MorphoLeaf application. Measures were performed on the whole leaf (A-C) or on tooth 1 (D-G, teeth numbered according to their position on each side of the leaf blade starting from the tip to the base, which also correspond to their order of initiation). Blade width (A), blade area (B) and teeth number (C) plotted against blade length. Tooth width (D), tooth height (E), relative proximo-distal position of the distal sinus of teeth $1(\mathbf{F})$ and tooth height to width ratio (G) plotted against blade length. $n=114$ for L11, $n=46$ for L12 and $n=47$ for L13


Supplementary Figure S3. Hierarchical organisation of leaf margin outgrowths.
The hierarchical organisation of the leaf margin in northern red oak (Quercus rubra, A) and Arabidopsis thaliana ( $\mathbf{B}$ ). The first panels (i) show the leaves, the second panels (ii) show the results of the segmentation and landmark identification (junctions between the petiole and the blade (grey diamonds), leaf tip (pink diamond) and sinuses (green dots)). The third panels (iii) show the results of the naive identification of the teeth (in blue) as structures contained between two successive sinuses. The fourth panels (iv) show the correct identification of the teeth, with secondary structures (light blue) supported by first order teeth (blue). Note that the whole tooth corresponds to light blue and blue regions.


Acer negundo


Ailanthus altissima


Alnus orientalis



Fagus sylvatica


Hamamelis japonica


Koelreuteria paniculata


Mahonia aquifolium


Populus tremula


Sorbus intermedia


Supplementary Figure S4. Analysis of leaf shape of different species using MorphoLeaf.
For each species, the input leaf image (binary image, left panel) and the result of the analysis performed with MorphoLeaf (right panel) are shown. Note that no manual corrections have been done, and that the results shown are those directly obtained by the MorphoLeaf.
Landmark description: blue diamonds: petiole-blade junctions; blue cubes: leaf tips, green dots: primary tooth/lobe sinuses; red triangles: primary tooth/lobe tips; orange dots secondary tooth/lobe sinuses; orange triangles secondary tooth/lobe tips.
The leaf pictures are from the Middle European Woods data base from the Department of Image Processing at the Institute of Information Theory and Automation of the ASCR, Czech Republic.

Novotný, P. and Suk, T. (2013) Leaf recognition of woody species in Central Europe. Biosys. Eng. 115, 444-452.


## Supplementary Figure S5. Current limitations of MorphoLeaf.

For each species, a leaf (left panels, $\mathbf{i}$ ), its contour with the biologically relevant landmarks (blue diamonds: petiole-blade junctions; blue cubes: leaf tips, green dots: primary tooth/lobe sinuses; red triangles: primary tooth/lobe tips; orange dots secondary tooth/lobe sinuses; orange triangles secondary tooth/lobe tips, central panels, ii) and mean contours (grey: generated without reparametrisation; blue: generated after leaf tip, tooth sinuses and tips guided reparametrisation, right panels, iii) are shown. Scale bars=1cm
The hierarchy of palmately lobed leaves with several levels of dissection (such as sycamore maple, Acer pseudoplatanus Ai) cannot be established with MorphoLeaf. Nevertheless, appropriate setting of the parameters allows the detection of the sinuses on the main lobes ( $\mathbf{A}$-ii) and thus proper quantitative analyses and mean leaf shape reconstruction (A-iii, $n=4$ ).
In MorphoLeaf two neighbouring structures such as teeth or lobes share the same sinus between them. Because of this definition the MorphoLeaf cannot be used to properly analyse pinnately compound leaves (such as rose, Rosa sp, Bi), in which two neighbouring leaflets are separated by a stretch of the rachis. Running MorphoLeaf leads to the identification of extra structures corresponding to rachis segments ( $\mathbf{B}-\mathrm{ii}$ ) and hence producing false quantitative analyses. Nevertheless, the mean shape that is reconstructed by MorphoLeaf is accurate (B-iii, $n=3$ ).
MorphoLeaf allows the identification of complex patterns of dissections (such as those of European white birch, Betula pendula, C-i and C-ii). However, while using the primary sinuses and peaks to guide the reconstruction increases the quality of the mean shape in the distal part of the leaf, it degrades the quality in the proximal part (compare grey and blue contours in C-iii, $n=3$ ). Because of the heterogeneity in the sizes of primary teeth (alternatively large and small teeth in the proximal region of the leaf) and the variability in the total number of teeth ( 10 or 11 per half-leaf here), dissimilar structures are put into correspondence to reconstruct the mean shape. To avoid this, the user can manually homogenise the number of teeth per half leaf on each sample (by adding or removing teeth in the distal part where they are small), which solves the homology issue and leads to an accurate reconstruction of the mean shape in both the distal and proximal part of the leaf (Civ).

## A Leaf Parameters






B Tooth 1 Parameters





C Tooth 2 Parameters





D Tooth 3 Parameters





Supplementary Figure S6. Morphometrics of mature leaves L01, L03, L05, L07, LO9 and L11.
(A) Quantitative parameters associated with the entire leaf. All these parameters (except the teeth number) were determined on the reconstructed mean shapes. The number of teeth was calculated from the leaves that were used to generate the mean shapes, error bars are standard errors. (B,C,D) Quantitative parameters associated with tooth 1,2 and 3 determined from the reconstructed mean shapes. $n=10$ for L01, L03, $n=11$ for L07, L09, L11 and $n=12$ for L05.



Supplementary Figure S7. Developmental trajectories of leaf LO1, L03, L05, L07, L09 and L11 For each leaf rank (L01, L03, L05, L07, L09 and L11), 10 mean shapes were reconstructed using the normalization method based on bins are shown. For each leaf, the first five mean shapes and last five mean shapes are at the same scale. The blade length (in $\mu \mathrm{m}$ ) of the mean shapes is indicated inside of the contours. The tooth sinuses that become shallower during later stages of development and are hardly detectable at the mature stage are circled in red.


B




Supplementary Figure S8. PCA Analysis of L01, L03, L07 and L11
(A) PCA analysis of the whole sample of registered and reparametrised leaf contours (between 100 and $2000 \mu \mathrm{~m}$ ). The mean shape is represented by a solid line while fine and gross doted lines represent leaf shapes obtained by varying the PC by +1 SD and -1 SD, respectively. This analysis identified three principal components that explained more than $80 \%$ of the total shape variability. The first axis (PC1, 55,5\% of total variability) may be interpreted as corresponding mostly to differences in teeth positions along the proximo-distal axis of the leaf. The second axis (PC2, 14\% of total variability) may be linked to the size of the teeth. The third axis (PC3, 11\% of total variability) may correspond to a global shape feature.
(B) Variations of PC1, PC2 and PC3 during the course of leaf development (error bars are standard errors).


Supplementary Figure S9. Relative position of teeth 1,2 and 3 sinus along the proximo-distal leaf axis during entire $\mathbf{L 1 1}$ development.
The relative position of the distal sinus of teeth 1,2 , and 3 along the proximo-distal axis of the leaf is shown during the entire development of L11. Below, three size-fitted average shapes at four different developmental stages are shown and the trajectories of the distal sinuses of teeth 1,2 and 3 are indicated using the same colour code as in the graph. The region of the leaf located between the distal sinus of tooth 3 (blue line) and the base of the leaf is the region which shows the strongest relative increase in size during the development of leaves $>2 \mathrm{~mm}$ in length.

|  | Leaf Shape |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | lanceolate | elliptical | ovate | obovate | cordate | palmate | lobed | palmately compound | pinnately compound |
|  | smooth |  |  |  |  |  |  |  | Choisya ternate |  |
|  | denticulate |  | Sambucus nigra |  |  |  | Acer sp. |  |  | Rosa sp |
|  | crenate |  |  | Alnus orientalis |  |  |  |  |  |  |
|  | sinuate |  | Hamamelis japonica | Populus tremula |  | Cercidiphyllum japonicum |  |  |  |  |
|  | serrate | Forsythia intermedia | Arabidopsis thaliana Fagus sylvatica Mahonia aquifolium | Ilex aquifolium |  |  | Acer pseudoplatanus |  |  |  |
|  | spiny |  | Castanea sativa |  |  |  |  |  |  |  |
|  | lobate | Acer negundo Ailanthus altissima |  |  |  |  | Acer campestre | Quercus petraea |  |  |
|  | lobate serrate |  |  | Koelreuteria paniculata |  |  |  | Acer tataricum |  |  |
|  | doubly dentate |  |  |  | Alnus glutinosa | Corylus colurna |  |  |  |  |
|  | doubly serrate |  | Betula pendula | Sorbus intermedia |  | Ulmus laevis |  | Quercus rubra |  |  |

Table S1. Summary of the species analysed. For each species the leaf shape and the margin shape are indicated

## Movies 1-12

Movies are available at http://morpholeaf.versailles.inra.fr/video/videoArabidopsis.html.

## Supplementary Methods

## MorphoLeaf software methods

## Blade contour segmentation and landmark detection

## Leaf contour segmentation

On each image of a data set, the contour of the leaf is automatically extracted using the watershed method (Vincent and Soille, 1991). Briefly, two markers are located on the image, one within the leaf area and one within the background. These markers are obtained after an automated thresholding of the image intensities, which allows to roughly separate the regions corresponding to the leaf and the background. Then, the watershed algorithm allows to efficiently detect the limit of the leaf (Figure SM1, contour in yellow). In order to remove nonrelevant details along the contour, the MorphoLeaf application proposes a tool to simplify the contour. This simplification is performed by retaining only the first elliptical Fourier descriptors that encode the contour in the frequency domain (Kuhl and Giardina, 1982). In the MorphoLeaf application, the number of descriptors (i.e. the degree of simplification) is controlled by the "Fidelity" parameter, which corresponds to the proportion of the power spectrum that must be retained during the simplification .

## Leaf blade identification

Once the leaf contour is extracted, the limit between the petiole and the blade is manually determined by the user. This is not automated due to the difficulty to find a common criterion to various species and developmental stages. In practice, the user sets two landmarks on the contour on both sides of the petiole-blade limit. The blade is then defined as the half contour on either side of the Petiole landmarks that defines the largest area (the other one, which corresponds to the petiole of the leaf, is not further analyzed). Next, the leaf apex is automatically determined as the position on the blade contour that is at the maximal Euclidean distance from the midpoint of the segment defined by the petiole landmarks (see Figure SM1).


Figure SM1. Segmentation of the leaf blade and determination of sinus landmarks. The leaf contour is automatically segmented (contour in yellow). The blade contour is delimited by the manually positioned limits of the petiole (purple circles). The leaf apex (blue circle) is determined as the contour point at the greatest distance from the midpoint of the petiole limit (dashed segment). In each significant concave region (contour portions in green), a single tooth sinus is identified as the position with maximal negative curvature (green circles). Scale bar: $100 \mu \mathrm{~m}$.

## Teeth sinuses

The identification of the teeth sinuses is based on the curvature of the contour. The local curvature is estimated at each position over the contour. The contour can then be decomposed into continuous convex or concave regions, the latter containing each a single sinus. To reduce the influence of insignificant contour oscillations, we only consider concave regions where the absolute curvature remains above a threshold value. In practice, the local curvature is computed from the outer angle formed by the position of interest and two neighboring contour points located within a given distance of the point of interest: the sharper the angle, the higher the curvature. In the MorphoLeaf application, the corresponding distance and threshold curvature are controlled by the "Half Neighborhood" and "Maximal Negative Curvature" parameters. Within each concave interval, the point with the maximal curvature (sharpest angle) is selected as a sinus (see Figure SM1).

## Teeth tips

A single tooth tip is identified in each interval delimited by two consecutive sinuses (except between the two sinuses separated by the leaf apex in the case of pinnate leaves like in Arabidopsis). We developed two strategies to identify the tooth tip depending on the shape of the teeth:

- Method 1. For pointy teeth, the tip is identified as the point of maximal curvature, which is assessed using the same curvature measure used to identify the sinus positions but for inner angles. The Half Neighborhood parameter, which is expressed here as a fraction of the length of the leaf perimeter contained between two consecutive sinuses is the parameter controlling the calculation of the local curvature. Using such a parameter proportional to the size of the tooth instead of an absolute parameter enables a more accurate sinus tip identification when the size of the teeth is heterogeneous.
- Method 2. For rounded teeth, as for example those appearing at the margin of Arabidopsis primordia (Figure SM2A), the method described above was not accurate enough. We thus developed an alternative method based on local symmetry, starting from the observation that a rounded tooth can be seen (at least in the vicinity of its tip) to emerge from a shape that is symmetrical with respect to the tip. The recursive method we designed consists in identifying the optimal symmetry axis, which crosses the tooth contour at the tip location. The first candidate as a tip between two sinuses is the intersection between the tooth contour and the perpendicular bisector of the basis of the tooth (Figure SM2B). Besides, a reference sinus for the tooth of interest is identified as the sinus delimiting the largest half-tooth defined by the


Figure SM2. Determination of the tooth tip based on local symmetry. A: a leaf contour (yellow curve) and teeth sinuses (green circles). B: zoom corresponding to the rectangle in A. The tooth basis (thick red segment) and the perpendicular bisector (thin red segment). C: search of the tooth tip which maximizes the local symmetry criterion (blue circle). White arrowhead: tip determined by the maximal curvature criterion. Scale bar: $100 \mu \mathrm{~m}$.
bisector (sinus on the right in Figure SM2B). Starting from the opposite sinus, successive points over the contour are considered, defining new bases (thick segments in Figure SM2C). At each step, the local symmetry is estimated by computing the ratio of the areas of the two regions delimited by the bisector and the basis. The contour position intersected by the axis with the maximal symmetry (e.g., a ratio close to 1 ) is chosen as the tooth tip (blue circle in Figure SM2C, to be compared with the solution given by the maximal curvature criterion, indicated by the arrow).

## Identification of the tooth hierarchy

A primary tooth is defined as growing on the main leaf margin, while a secondary tooth is formed on a primary tooth. The objective is to retrieve this hierarchical structure by identifying secondary sinuses, i.e. sinuses that delimit secondary teeth only (as opposed to primary sinuses, at the basis of primary teeth). A proper hierarchy identification is crucial because the positioning of the sinuses is not sufficient to properly characterize the leaf serrations when secondary teeth occur on the contour (see Figure SM3). Note that a single sinus can delimit primary and secondary teeth at the same time.

To determine the tooth hierarchy, we designed two different approaches. Both are based on the observation that the shape defined by the primary sinuses is well aligned with the general leaf contour, i.e. a contour in which all teeth have been erased. On the contrary, secondary


Figure SM3. The sinuses identification is not sufficient to properly quantify leaf blade serrations. A: contour of an Arabidopsis leaf, with the serrations of the upper half-leaf highlighted (green curve). Sinus positions are showed (green circles). B: zoom on the rectangle in A. Due to the presence of a secondary tooth, the position of the teeth is incorrect if we consider the sinuses sequentially. In gray: the numbers indicate four successive teeth that are incorrectly identified; their basal limits are indicated with continuous gray segments. In blue: the four teeth are correctly identified; their basal limits are indicated with blue dashed segments. Teeth 1,2 , and 3 are primary teeth while tooth 4 is a secondary tooth that is part of the primary tooth 1 . Scale bar: $100 \mu \mathrm{~m}$.
sinuses are located further from this general contour. The first method (Method 1, described below) is efficient on a large variety of species, but fails on Arabidopsis because its leaves show a broad variety of tooth size along the blade during their development. We therefore developed a specific method to determine the hierarchy of developing teeth in Arabidopsis young leaves, based on a recursive process (Method 2, described below).

The two approaches described below apply to half-leaves (portion of the contour from the petiole to the leaf apex). The sinuses are ordered relatively to their successive positions on the contour (from the leaf apex to the petiole).

Method 1. The principle is to iteratively remove secondary sinuses until only primary sinuses remain on the contour. The method relies on the property that, contrary to secondary sinuses, primary sinuses are well positioned along the global leaf contour (i.e., a contour without teeth). To quantify this, we use a curvature criterion: we compare the alignment of each sinus with its two neighbors and the local curvature at approximately the same position of the global leaf contour. These two measures should be similar at the position of primary sinuses, and significantly different for secondary sinuses. The shape of the global contour is approximated by drastically smoothing the initial contour (see Figure SM4AB). The smoothing is done using the strategy based on the elliptical Fourier descriptors introduced above, and the retained proportion of the power spectrum (degree of smoothing) is fixed as $50 \%$. Initially, all sinuses are primary sinuses. The two bounding sinuses (first and last ones on the half-blade) are always primary. Sinus positions are projected onto the global smooth contour (see Figure SM4B, blue crosses). For each (non-bounding) sinus, we compute the angle formed with the two neighboring sinuses and the angle formed at the position of the projected sinus with two points in its vicinity on the global smooth contour. At each step, the candidate secondary sinus is the one for which the difference between these two angles is the highest (the pairs of compared angles for each candidate sinus are indicated by the same number in Figures SM4B and SM4C). If this difference is large enough (greater than a fixed threshold, the "Stringency Factor" in the MorphoLeaf application), the sinus is identified as secondary and removed from the set of sinuses (see the successive steps in Figure SM4C). Then, because a sinus is removed, the neighborhood relationship is changed for its two neighboring sinuses, thus the corresponding angles are recalculated accordingly. Otherwise, if the


Figure SM4. Hierarchy computation: the iterative method. A: northern red oak (Quercus rubra) leaf image with the leaf contour (in red) and the teeth sinuses (light blue squares). B: Same segmentations as in A, with the projections of the sinuses (blue crosses) onto the general leaf shape (smooth contour, in orange). The numbered angles are related to the different steps of the method (see below). C: successive steps of the algorithm, applied to the sinuses of the upper half-leaf (red portion of the contour). Secondary sinuses are iteratively removed (gray squares). The green line connects the remaining sinuses. At each step, the candidate sinus (marked by an arrowhead) corresponds to the sinus that is the more distant from the general contour. Practically, the alignment of consecutive sinuses (angles in C) is compared to the curvature at the corresponding positions in the smooth contour (angles in B). At each step, the candidate is characterized by the greatest difference between the two angles (the pairs of angles corresponding to the successive candidates are indicated by a same number in Figs. B and C). The procedure stops when the angle difference is not significant anymore (step 4). Scale bar: $500 \mu \mathrm{~m}$.
difference is small, the procedure stops and the remaining sinuses are labeled as primary (see the final result in the last panel in Figure SM4C). In the MorphoLeaf application, the sole parameter of the method is the Stringency Factor.

Method 2. The idea behind this recursive method is to detect sinuses that are not well aligned with their two neighboring sinuses. The two extreme sinuses on either side of the contour at the current recursion are called bounding sinuses (initial bounding sinuses are pointed by arrows in Figure SM5A), and the portion of contour delimited by bounding sinuses is called a


Figure SM5. Hierarchy computation: the recursive algorithm of method 2. A: a leaf contour (in gray) and the portion of the contour processed to build the hierarchy (in C), with the corresponding sinuses (green circles). Arrows: initial bounding sinuses. B: steps of the recursive procedure (the gray arrows illustrate the flow of the steps). The contour is recursively split into two sub-contours at the position of the sinus that maximizes the angle with bounding sinuses (red dash lines, maximal angle indicated by an arc). According to the angle value, a change in the hierarchy level is detected (star, step 1-2) or not (steps 1 and 1-1). The last row shows the final teeth, numbered in the order they were identified (tooth 3 is secondary). In step 1-2-2, the contour considered by the algorithm is indicated by the continuous green line, but the corresponding tooth (4) includes the green dashed line. C: corresponding hierarchy tree, with the same numbering as the one on the last row in B. Scale bar: $100 \mu \mathrm{~m}$.
lobe. The objective is to recursively build the hierarchy tree: the root corresponds to the base contour of the leaf (level 0) and the other nodes correspond each to a tooth, whose rank is determined by its level in the tree (a node of level 1 corresponds to a primary tooth, etc.). If the current lobe contains at least an inner sinus (i.e., distinct from the bounding sinuses; see, e.g., step 1 in Figure SM5B), the deepest sinus is determined as the sinus that forms the largest inner angle with the bounding sinuses (the deepest sinuses are indicated by arcs in Figure SM5B). This angle is used as a measure of alignment. The selected inner sinus splits the current lobe into two consecutive sub-lobes (in Figure SM5B, the split of the initial lobe yields the two sub-lobes highlighted in green in steps 1-1 and 1-2), and it becomes a bounding sinus in each of the two sub-lobes. If this sinus is sufficiently deep (if the angle is larger than a threshold, see below), no higher hierarchy is detected (steps 1 and 1-1 in Figure SM5B) and the same procedure is applied recursively on each of the two sub-lobes, with no change in the hierarchy (they are sister lobes). Alternatively, if the sinus is insufficiently deep (if the angle is smaller than the threshold, see below), the level of the sublobe with the smallest area is increased by 1, while the level of the other sub-lobe is unchanged (step 1-2 in Figure SM5B). In parallel, a new node is created at the current level in the tree that corresponds to the lobe with the largest area; this node is the mother of the sublobe, whose level is thus increased by 1. Next, the procedure is recursively repeated on each of the sub-lobes. The procedure stops when there is no more inner sinus, so that the current
lobe is a tooth (and a terminal node in the tree). The threshold angle determining whether the sinus is deep enough is chosen by the user in the MorphoLeaf application (Limit angle parameter).

## Validation of automatic landmark detection

The correct identification of the teeth over the blade is crucial for the proper quantification of the blade shape. The teeth are defined by both the sinus positioning and their hierarchical organization. Therefore, we analyzed the performances of the MorphoLeaf application for the sinus detection and for the hierarchy identification by comparing automatic and manual results.

## Validation of the sinus detection

To test the sinus detection procedure, two datasets were processed. Fifty images of Arabidopsis developing leaves and 10 of Northern red oak mature leaves, with blade contours and petiole limits previously segmented, were presented separately to 4 biologist experts. Using the Free-D software interface, each expert manually pointed the sinus positions over the leaf contours in the two datasets. In parallel, the automatic sinus segmentation of MorphoLeaf was applied to the data, with different values for the two parameters used in the module to determine sinuses positions: the Half Neighborhood and the angle defining the Maximal Negative Curvature.

For each image of the datasets, the first step consisted in automatically matching all segmented points (either defined by the experts or by MorphoLeaf) that correspond to the same biological landmark (see Figure SM6). Practically, points were sorted into different classes (showed with different colors in Figure SM6), so that each class corresponded to a single identified biological feature. Thus, a class contained at most one point from a given source (an expert or a given set of parameters). Each class that does not contain at least one point identified by an expert contains only false positives (false detection of a sinus; classes colored in white in Figure SM6). Conversely, each class that does not contain exactly the total number of sources contains at least one false negative (undetected sinus; for example, in Figure SM6, three pairs of parameters yielded false negatives in the class colored in light


Figure SM6. Determination of homologous landmarks classes from different sinus segmentations. The positions of sinuses segmented over a blade contour (curve in yellow), either manually (here, by three biologist experts) or automatically (here, with 12 different pairs of parameters, same as in Figure SM7), are represented by colored dots. Points corresponding to a same biological landmark are put into the same sinus class and displayed with a specific color. Classes of false positives are all represented by white dots (arrowheads). A: standard 2D view of the contour and segmentations. B: tilted 3D view of $A$; for the sake of illustration, a different altitude (in the z-dimension) was assigned to each segmentation source, which is positive for each expert (points above the contour) and negative for the automatic method (points below the contour). The altitudes assigned to the different sets of parameters (from top to bottom) follows the same order than in Figure SM7. This representation was also used to visually inspect the accuracy of the sinus assignments to the different classes. Scale bar: $100 \mu \mathrm{~m}$.
blue). Once each detected point assigned to a sinus class, we considered three criteria to compare manual and automatic detections: accuracy, number of false positives and number of false negatives.

For all pairs of parameters, we measured the distances between the automatic detections and the sinuses identified by three out of the four experts (this set of three referent experts was thus considered as ground truth). Practically, within each landmark class (except false positive classes), the geodesic distance (distance along the contour) from each segmented point to the barycenter of the points from the set of referent experts was computed. For each set of parameters, these measures were averaged across all classes and all leaves in the data set and compared to the same measures computed for the remaining expert. Applied successively to all possible sets of three referent experts, this allowed us to compare the performances of each expert individually with the ones of the automatic method. The results are shown in Figures SM7A and SM8A. For the Arabidopsis dataset (Figure SM7A), the mean distances do not exceed $3 \mu m$. This is remarkably low when compared to the leaf contour length in the data set (which ranges from $553 \mu \mathrm{~m}$ to $3930 \mu \mathrm{~m}$, with an average length of


Figure SM7. Comparison between manual and automatic sinus detections in Arabidopsis developing leaves. A set of Arabidopsis leaf images was manually and independently analyzed by four biologist experts to extract the position of sinuses. In each of the four columns, the results obtained by one expert (black square) and by automatic detections with different parameters (colored dots) are compared to the detections performed by the set of the three reference experts. We considered three criteria: precision (distances in $\mu \mathrm{m}$ ) (A), number of false positive (B) and of false negative (C) detections. See the text for details.
$1738 \mu \mathrm{~m})$. For all experts, the accuracy was close to the one reached by automatic detections, whatever the parameter values. This confirmed that the automatic method can be as accurate as expert detections. For the oak dataset (Figure SM8A), the contour lengths were longer (ranging from 68 cm to 108 cm for 15 cm -long blades in average, with an average contour length of 90 cm ), but the mean distances always remained below 1 mm , with three sets of


Expert compared to the automatic method with different parameters
Figure SM8. Comparison between manual and automatic sinus detections in Northern red oak mature leaves. A set of oak leaf images was manually and independently analyzed by four biologist experts to extract the position of sinuses. In each of the four columns, the results obtained by one expert (black square) and by automatic detections with different parameters (colored dots) are compared to the detections performed by the set of the three reference experts. We considered three criteria: precision (distances in $\mu \mathrm{m}$ ) (A), number of false positive (B) and of false negative (C) detections. See the text for details.
parameters that showed the same level of performance than the experts.

For Arabidopsis, a total number of $383,405,407$ and 390 sinuses were detected over the 50
blade contours by Experts 1, 2, 3, and 4, respectively (in average, a number of 396 sinuses). The analysis of the number of false positive and of false negative detections (Figures S7BC) both showed a great variability according to the parameter values. Nevertheless, for each criterion, a set of parameters that provides satisfactory results that are comparable to those of the expert can be found. Concerning false positive detections, the 7 sets of tested parameters that provided the best results yielded less than 20 false positive points in total (i.e., $5 \%$ or less of the mean number of sinuses detected by the experts). These detections mainly occurred in the apical region of the blade, due to insignificant contour oscillations or to the two hollows on both sides of the leaf apex. The level of over-detection can be reduced by increasing the smoothing of the contours, but at the risk of eroding the real teeth. The number of false negative detections was slightly higher and was comprised between 20 and 80 for a large majority of parameters (between $5 \%$ and $20 \%$ of the mean number of sinuses detected by the experts). Experts 1, 2, 3, and 4 yielded 45, 23, 21 and 38 false detections, respectively. Hence, for both the experts and the automatic method, there was more false negatives than false positives. This is probably due to the specificity of the Arabidopsis dataset, which contains images of developing leaves in which the limits of emerging teeth are not pronounced and the sinus sometimes hardly visible.

For the oak dataset, a total number of $347,367,351$ and 376 sinuses were detected over the 10 blade contours by Experts 1, 2, 3 and 4, respectively (in average, a number of 360 sinuses). If the numbers of false positive points were similar to the ones obtained for Arabidopsis, the level of false negative detections was significantly lower (Figures SM8BC), probably because the ambiguity in detecting teeth is lower in mature leaves than in developing ones. Four parameter sets generated 22 or less false positives points ( $6 \%$ or less of the mean total number of sinuses detected by the experts), while the majority of parameters (5 out of 6) yielded 36 or less false negative detections ( $10 \%$ or less of the mean total number of sinuses detected by the experts). As in the case of Arabidopsis, several sets of parameters provided a number of false positive or negative detections comparable to the results from the experts.

These results emphasize the compromise to be found between the amounts of false positives and of false negatives. In practice, we recommend the use of parameter values that minimize
the number of false negative detections at a price of a reasonable number of false positive points. In our opinion, it is indeed easier to visually detect and then remove false positive points with the tools provided by the Free-D software. It is generally possible to identify eligible parameter values by manually testing several parameters and by visually inspecting the corresponding results. In the case of our Arabidopsis dataset, an angle of 60 degrees and a neighborhood size of 30 pixels seems to be a good compromise that yields about 40 false detections and about 30 false positive points, which represent $10 \%$ and $7 \%$ of the true total number of sinuses, respectively. Concerning the oak dataset, an angle of 145 degrees and a neighborhood size of 25 or 50 pixels provided good overall results.

## Validation of the hierarchy procedure

We quantified the performances of the two procedures available for the automatic tooth hierarchy identification. The method 1 was evaluated on a set of 10 northern red oak leaf images. The method 2 was evaluated on a set of 40 leaf images from Arabidopsis mir164a-4 mutants, which present an increased level of leaf serrations (Nikovics et al., 2006). The two data sets were previously analyzed to extract the blade contours and the positions of petiole, apex and tooth tips and sinuses. Next, both data sets were presented separately to four experts. Each expert manually identified, among all the sinuses, the pairs of sinuses delimiting secondary teeth. For the two species, the hierarchy identification was identical for all experts, so that their results were considered as ground truth in the following analyses. Based on the method previously presented to match homologous segmented points (Figure SM6), we compared the manual identifications of secondary teeth to the ones obtained using the automatic methods. For this, we quantified the number of true and false positive ( $T P$ and $F P$ ) and negative ( $T N$ and $F N$ ) detections of secondary teeth, and computed the sensitivity (true positive rate, $[T P /(T P+F N)]$ ) and specificity (true negative rate, [TN/ $(T N+F P)])$. Note that here, the TN detections correspond to primary teeth detections.

For northern red oak leaves, we applied the iterative method with four different values for the Stringency Factor. The results are presented in Table 1. All tested stringency values yielded satisfactory results with sensitivity and specificity both above 95\%, with an optimal parameter value of 0.25 .

| Northern red oak leaves (10 leaves) |  |  |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Stringency Factor | Total number of teeth | Total number of primary teeth | Total number of secondary teeth | Number of segmented secondary teeth | True positive (TP) | False positive (FP) | True negative (TN) | False negative (FN) | Sensitivity | Specificity |
| 0.15 |  |  |  | 207 | 201 | 6 | 143 | 0 | 100\% | 96\% |
| 0.20 |  |  |  | 203 | 201 | 2 | 147 | 0 | 100\% | 99\% |
| 0.25 |  |  | 201 | 201 | 201 | 0 | 149 | 0 | 100\% | 100\% |
| 0.30 |  |  |  | 194 | 194 | 0 | 149 | 7 | 97\% | 100\% |

Table 1. Performance quantification of the iterative hierarchy method for secondary teeth detection. The secondary teeth automatically segmented in northern red oak leaves by the iterative method, with four different parameters, were compared to the true segmentations defined by experts.

For Arabidopsis leaves, we applied the method with four different values for the Limit Angle parameter. The results are presented in Table 2. The evaluation of the recursive method showed that an angle around 45 degrees provides very satisfactory results, with a low number of false negative and positive detections, and both a sensitivity and a specificity above $90 \%$.

## Arabidopsis thaliana leaves (40 leaves)

| Limit Angle parameter (in degrees) | Total number of teeth | Total number of primary teeth | Total number of secondary teeth | Number of segmented secondary teeth | True positive (TP) | False positive (FP) | True negative (TN) | False negative (FN) | Sensitivity | Specificity |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 35 | 345 | 317 | 28 | 12 | 12 | 0 | 317 | 16 | 43\% | 100\% |
| 40 |  |  |  | 24 | 22 | 2 | 315 | 6 | 79\% | 99\% |
| 45 |  |  |  | 35 | 26 | 9 | 308 | 2 | 93\% | 97\% |
| 50 |  |  |  | 53 | 28 | 25 | 292 | 0 | 100\% | 92\% |

Table 2. Performance quantification of the recursive hierarchy method for secondary teeth detection. Secondary teeth automatically identified in Arabidopsis thaliana leaves were compared to the true segmentation defined by experts (third and fourth columns). Numbers of true and false positive (TP and FP) and negative (TN and FN) detections were evaluated, and the method sensitivity and specificity were computed for each parameter value (last two columns).

In conclusion, for both red oak and Arabidopsis leaves, the two methods for the automatic detection of secondary teeth showed very good performances.

## Estimation of growth trajectories

The blade shape evolution during growth can be estimated by averaging contours of different lengths during development. First, contours are sorted according to the blade length, so that
they are all distributed over a "length axis". In parallel, all blade contours are reparametrized so that the landmarks which are biologically homologous across the data are put into correspondence. After this reparametrisation procedure, the contour portions comprised between the homologous landmarks in different leaves are composed of the same number of points. This ensures that the $p$-th point in a contour is homologous to the $p$-th points in all other contours. Thus, relevant mean shapes can be computed. Before the averaging, contours are aligned using a group-wise registration procedure (Maschino et al., 2006). Then, we proposed two approaches to estimate a growth trajectory based on the computation of mean blade contours. The first method is based on assigning leaves to different size classes and the second one on a moving (or sliding) average approach. For both methods, it should be stressed that a sufficiently dense sampling (with no gap) of the contours according to the blade length is critical for a proper estimation of the growth trajectory.

Bin-based method. The strategy consists in sorting the contours into different length classes ("bins") on the length axis, and then in averaging all contours within each class. In the MorphoLeaf application, the user can specify the number of desired average contours (i.e., the number of classes). Then, the limits of the bins are computed so that all contain the same number of contours (total number of contours divided by the number of classes). Alternatively, the limits of the intervals can be manually specified, by loading in the application a file that contains the chosen limit values.

Moving average method. Instead of fixing intervals, this second approach (also called Sliding average) consists in using an adaptive kernel strategy (Parzen, 1962) for the contour averaging. At a given blade length $L$, all contours are used to compute the corresponding average. A weight is affected to each contour according to a Gaussian kernel function centered on $L$. Thus, only the contours whose lengths are close to $L$ significantly contribute to the averaging. The bandwidth parameter which controls the width of the kernel is computed locally and is equal to the distance from $L$ to the length of the $k$-th nearest contour on the length axis. This allows the bandwidth to adapt to the local density of contours on the length axis: the bandwidth is smaller when the number of contours with a length close to $L$ is high, and conversely. Besides, $k$ is a smoothing parameter: small $k$ values result in narrower bandwidths while high $k$ values lead to larger bandwidths (Figure SM9). The user can control


Figure SM9. Effects of varying the neighbor rank in the sliding averaging on the mean shape. Average Arabidopsis thaliana leaves of $504 \mu \mathrm{~m}$-long were reconstructed using the sliding average method from a set of 207 leaves of lengths ranging from $109 \mu \mathrm{~m}$ to $2383 \mu \mathrm{~m}$. For each reconstructed average shape, the neighbor rank used and the resulting bandwidths are indicated.
the bandwidth by defining the $k$ value (called Neighbor Rank in the MorphoLeaf application). It is possible to specify either the desired number of average contours (which are then equally distributed over the range of all lengths) or a set of specific lengths, by loading a file with the chosen values in the application.

A file with the local bandwidth at each mean shape is generated and can be used by the user to check the contribution of individual leaves to the mean shape.

## Bibliography

Kuhl, F. P. and Giardina, C. R. (1982). Elliptic Fourier features of a closed contour. Computer Graphics and Image Processing, 18, 236-258.

Maschino E., Maurin Y. and Andrey P. (2006). Joint registration and averaging of multiple 3D anatomical surface models. Computer Vision and Image Understanding, 101, 16-30.

Nikovics, K., Blein, T., Peaucelle, A., Ishida, T., Morin, H., Aida, M. and Laufs, P. (2006) The balance between the MIR164A and CUC2 genes controls leaf margin serration in Arabidopsis

The Plant Cell 18(11), 2929-45.

Parzen E. (1962). On estimation of a probability density function and mode. The Annals of Mathematical Statistics. 33(3), 1065-1076.

Vincent L. and Soille P. (1991). Watersheds in digital spaces: an efficient algorithm based on immersion simulations. In IEEE Transactions on Pattern Analysis and Machine Intelligence, 13(6), 583-598.

