



PlantNet Participation at LifeCLEF2014 Plant Identification Task

Hervé Goëau, Alexis Joly, Itheri Yahiaoui, Vera Bakić, Anne Verroust-Blondet, Pierre Bonnet, Daniel Barthélémy, Nozha Boujemaa, Jean-François Molino

► To cite this version:

Hervé Goëau, Alexis Joly, Itheri Yahiaoui, Vera Bakić, Anne Verroust-Blondet, et al.. PlantNet Participation at LifeCLEF2014 Plant Identification Task. CLEF: Conference and Labs of the Evaluation Forum, Sep 2014, Sheffield, United Kingdom. pp.724-737. halsde-01064569

HAL Id: halsde-01064569

<https://hal.science/halsde-01064569>

Submitted on 16 Sep 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pl@ntNet's participation at LifeCLEF 2014

Plant Identification Task

Hervé Goëau¹, Alexis Joly^{1,2}, Itheri Yahiaoui¹, Vera Bakić¹, Anne Verroust-Blondet¹, Pierre Bonnet³, Daniel Barthelemy⁴, Nozha Boujemaa⁵, and Jean-François Molino⁶

¹ Inria, France, `name.surname@inria.fr`

² LIRMM, Montpellier, France

³ CIRAD, UMR AMAP, France, `pierre.bonnet@cirad.fr`, <http://amap.cirad.fr>

⁴ CIRAD, BIOS Direction and INRA, UMR AMAP, F-34398, France, `daniel.barthelemy@cirad.fr`, <http://amap.cirad.fr/fr/index.php>

⁵ INRIA, Direction of Saclay Center, `nozha.boujemaa@inria.fr`, <http://www-saclay.inria.fr>

⁶ IRD, UMR AMAP, France, `jean-francois.molino@ird.fr`, <http://amap.cirad.fr>

Abstract. This paper describes the participation of Inria within the Pl@ntNet project⁷ at the LifeCLEF2014 plant identification task. The aim of the task was to produce a list of relevant species for each plant observation in a test dataset according to a training dataset. Each plant observation contains several annotated pictures with organ/view tags: Flower, Leaf, Fruit, Stem, Branch, Entire, Scan (exclusively of leaf). Our system treated independently each category of organ/view and then a late hierarchical fusion is used in order to combine the results on visual content analysis from the most local level analysis in pictures to the highest level related to a plant observation. For the photographs of flowers, leaves, fruits, stems, branches and entire views of plants, a large scale matching approach of local features extracted using different spatial constraints is used. For scans, the method combines the large scale matching approach with shape descriptors and geometric parameters on shape boundary. Then, several fusion methods are experimented through the four submitted runs in order to combine hierarchically the local responses to the final response at the plant observation level. The four submitted runs obtained good results and got the 4th to the 7th place over 27 submitted runs by 10 participating teams.

Keywords: Pl@ntNet, Inria, LifeCLEF, plant, leaves, flowers, fruits, stem, bark, scan, branch, multi-organ, image, collection, identification, classification, evaluation, benchmark

1 Introduction

The plant identification task of LifeCLEF2014 [?] [?] was organized as a plant species retrieval task over 500 species with visual content being the main available information. The aim of the task was to produce a list of relevant species

⁷ <http://www.plantnet-project.org/>

for each observation in a test dataset containing 8163 plant observations based on 13146 images, where each observation can contain several pictures of detailed views on various organs: Flower, Fruit, Leaf, Stem, Entire, Branch and Scan (scans or scan-like pictures of leaf exclusively). The training dataset contains numerous plant observations related to 47815 images tagged with these organ/view information.

We present in this paper the methods used in order to produce the four submitted runs by Inria within the PlantNet project. It is based more or less on the same system used during our previous participations to the ImageCLEF2011-2013 plant tasks [?] but with some improvements mainly related to new types of combination of image list results provided by the different image queries from a same plant observation.

2 Description of the system

Each tested observation can potentially contain several pictures on several organs/views. Our system treated independently each organ/view and then a late hierarchical fusion framework is used in order to combine the visual content analysis from the most local level in pictures to the highest level related to one same plant observation.

- **local level** (see subsection ??): we extract various types of local features around interest points on training and test images. Then, considering one single test image with one organ/view tag, and only one type of feature, we search the knn local features in the training dataset with the same organs/views tag, and combine them in order to obtain a ranked training image list.
- **image level**: (late fusion of features, see subsection ??) then we combine the ranked training image list produced for all types of features.
- **multi-image level**: (see subsection ??) the results from multiple test images from a same plant observation within a same organ/view are combined.
- **multi-image level**: (see subsection ??) the results from several organs/views from a same plant observation are combined.
- **observation level**: (see subsection ??) the combined results are transformed in a ranked list of species.

2.1 Local level: large scale matching approach

Interest points detection (*Run 1, 2, 3 & 4*)

We explored the fact that in most cases, one organ or one group of organs, is generally centred in the pictures of the dataset. Harris corners were used at four distinct resolutions with multiple orientations ([?]). In addition, as in [?], to minimize the effect of the cluttered background, a rhomboid-shaped mask was applied to the input image and we used a Gaussian-like distribution in a 7x7 grid in order to take more points at the center of the pictures (the figure ??(a))




Fig. 1. Examples of interest points detection with rhomboid-based mask and with grid-based weighting

illustrates the detected points). About 100 points per image were output, and according to the multiple-orientations, between 150 and 200 local features are extracted in patches around these points (see subsection ??), depending of the visual content of the picture analysed.

Local features extraction (*Run 1, 2, 3 & 4*)

Several local features are extracted around each interest point from an oriented and scaled patch: rotation invariant Local Binary Pattern (ri-LBP) [?]; SURF [?] [?] (sums of 2D Haar wavelet responses); a 20-dim. Fourier histogram [?]; an 8-dim. Edge Orientation Histogram (EOH); a 27-bin weighted RGB histogram (wght-RGB) [?], a 27-bin weighted LUV histogram (wght-LUV) and a 30-bin HSV histogram. After a series of preliminary evaluations tests with the training data, we concluded that not all type of local features should be used for all views (see subsection ??). In the specific case of *Scan* we used a concatenation of the SURF, EOH, Fourier and a 16-dim. histogram based on the Hough transform [?] (we will call these concatenated descriptors "SEFH").

Visual index of local features (*Run 1, 2, 3 & 4*)

For each type of view/organ, and for each type of features, local features are computed, compressed and indexed using RMMH method [?,?] using a 256-bit hash code, which led to a total of 43 unique hash tables (6 organs/views x 7 types of features + 1 separate table for scans), but around half of them are finally considered for computing the submitted runs (see subsection ??).

Local similarity search (*Run 1, 2, 3 & 4*)

Considering a test image I_q , its associated organ/view, and considering one visual index of one type of features, the local similarity search produces a response list R composed of similar training images sorted by a score S_i related to a matching function on each training image i . Each local query feature q of the image I_q is compressed with RMMH and its approximate k -nearest neighbors are searched by probing multiple neighboring buckets in the consulted hash table (according to the a posteriori multi-probe algorithm described in [?]). Two types of votes and parameters of k were used through in the submitted runs in order

to produce scores and sort the training image response list: *Vote On Best* and *Ranking Dominate*.

Vote On Best (VOB) (Run 1, 2 & 4)

The score S_i here is directly the number of matches between a training image I_i and the query image I_q . More precisely it is the number of instances of image I_i retrieved through the $k = 30$ nearest neighbors lists of each local feature of the query image Q . Finally the list is limited to the 300 best training images.

Ranking Dominate(RD) (Run 3)

This new vote introduces a more advanced image ranking method based on the number of *dominating local features*. Inspired by the concept of top-k dominating queries in the database domain [?], we consider that a given local feature \mathbf{x} is *dominating* another local feature \mathbf{y} if it is closer to a query feature \mathbf{q} (according to the distance used as matching function). Two images I_1 and I_2 can then be compared with reference to a query image I_q by counting the number of query local features \mathbf{q}_i for which a local feature of I_1 is dominating all features of I_2 (or conversely). More precisely, images are ranked using a quick sort algorithm based on the following comparison function:

$$f(I_1, I_2) = \text{sgn}\left(\sum_{i=1}^n \text{sgn}(r_i^1 - r_i^2)\right) \quad (1)$$

where r_i^1 is the rank of the best match of I_1 in the k -nns of \mathbf{q}_i and r_i^2 is the rank of the best match of I_2 in the k -nns of \mathbf{q}_i . One advantage of this ranking function is that it is much less sensitive to the value of k . When k grows, the ranking of the top-K images becomes more and more stable. Ideally, the image ranking could even be computed from the complete rankings of all local features in the dataset. But for efficiency reasons, this is clearly not manageable, and we choose a value of $k = 600$.

2.2 Shape boundary description on scans

(Run 1, 2, 3 & 4)

In the case of leaf scans and scan-like pictures, additional boundary descriptors are used since it has been showed to provide a consistent improvement of the global performances in the previous ImageCLEF2011 dataset [?]. Involving a first step of automatic leaf boundary detection, these descriptors are the Directional Fragment Histogram (DFH) describing the leaf margin, and several geometric parameters embedded in one single vector (Shapes6) of 6 dim. measuring 6 standard geometric parameters on the shape: *rectangularity*, *convexity*, *solidity*, *circularity*, *Sphericity* and *Ellipse variance*.

2.3 Image level: late fusion of features

At this step, we have now for each image query I_q several image training response lists, one for each type of (local or global) features. The next step is now to

combine these responses in order to obtain a single response. Thus, we use a late fusion approach with different methods experimented through the four submitted runs.

CIN and WP (*Run 1*)

For a query image I_q , according to its organ/view, the basic algorithm applied is: (i) starting from the retrieved lists of similar training images for each type of feature, (ii) transform each list of images into a probability distribution by species with an adaptive rule *CIN* (see below), (iii) merge probabilities for all types of features with a Weighted Probability approach (WP).

CIN It an adaptive KNN rule established in this particular context of plant observations [?]. For converting an image response list to a species probability distribution, we use this adaptive rule focusing on plant observations (rather than images). Indeed, the more a species will be represented in a response through various plants observed by distinct users at different dates and locations, the more the associated images will be informative for predicting a species. In contrast, numerous redundant near duplicate images from the same plant observation will be less informative for predicting a species. Instead of using a basic KNN rule focusing on images, we search for top- K'' *classes* (species) represented with at least K' different plant observations. The values of K' and K'' are determined empirically based on the given training database, and are constant for a database: K' is a percentage of the average size of the training class, while K'' is a percentage of the number of training classes. The response R is scanned from the most to the least similar image, the counter of the number of classes with at least K' images is incremented accordingly, and when we find K'' such classes, we stop the scanning of the response. The adaptive criterion is important in order to avoid the noise in the final response: (a) had we searched for a fixed (and not dependent on the training data) number of classes with at least K' plant observations, we would often output classes that are not relevant to fill-in the pre-defined requirement, or (b) had we output a fixed number of most similar classes, we would give more weight to the classes with small number of plant observations and would not reward the fact that some classes are well represented in terms of plant observations. Finally, our K'' per class resulted in the K per image of 85, ranging from 15 to 205, for the most to the least different-plant-observation-containing response; organ-wise, the values of K ranged from 66 for Stem to 101 for Fruit.

Moreover, to eliminate the redundant images, we consider only two most similar images per one plant observation: the score per plant observation S^m is the average of the image scores S_i . The score for a class is the sum of the scores of its plant observations S^m : this step actually favours the well-represented classes, and penalizes the classes with small number of plant observations. Finally, the classes with only one image coming from one plant observation are removed from the list as outliers.

Weighted Probability combination (WP) At this step we have several species probability distributions, one for each type of features and we use a weighted fusion in order to obtain a final probability distribution. Let us define F as a set of local features and $P(C_f^k)$ as probability of class C^k for feature $f \in F$. In order to reflect the discriminating power of each local feature, we define the final probability:

$$P(C^k) = \sum_{f \in F} w(f) * P(C_f^k) \quad (2)$$

where

$$w(f) = \max_{\forall k} P(C_f^k) / \sum_{f \in F} \max_{\forall k} P(C_f^k) \quad (3)$$

BordaMNZ & IrpMNZ (*BordaMNZ: Run 2, 3 & 4, IrpMNZ: Run 4*)

In [?] the authors proposed similarity ranking lists fusion algorithms in order to merge some multi-feature similarity lists into a final overall similarity ranking list. They showed good performances on their experimental results with basic algorithms working directly with the ranks, i.e. without complicated and empirical transformation of similarity lists into probability distributions. We experimented here with two algorithms in order to merge the image training response lists given by each type of features into one single list: the Borda Count Algorithm taken from social theory in voting, and the *Inverse Rank Position Algorithm* (Irp). Basically a score is computed for each training image contained in the response lists to combine, and then used for sorting a final merged list of training images.

Considering an image query Q , a training image I_i , and several image training response lists R , one for each feature f from a set of available features F , the score used in the Borda Count algorithm is the sum of the rank positions of the training image in the lists to merge:

$$Borda(Q, I_i) = \sum_{f=1}^F rankposition(I_i, f) \quad (4)$$

The score used in the Irp algorithm is the inverse of the sum of inverses of the rank positions of the training image in the different lists to merge:

$$Irp(Q, I_i) = \frac{1}{\sum_{f=1}^F \frac{1}{rankposition(I_i, f)}} \quad (5)$$

Like in the previous work mentioned in [?], we used the MNZ extension of these two algorithms (*BordaMNZ* and *IrpMNZ*) in order to heavily weights common training images in the lists to combine. Indeed, the hypothesis is that different lists contain similar sets of relevant training images while they contain different sets of non-relevant training images. Basically, the weights are based on the number of Non-Zero entries of each training images in the lists to combine, i.e. the number of times a training image appears in the ranked lists.

$$BordaMNZ(Q, I_i) = Borda(Q, I_i) \times \alpha \quad (6)$$

$$IrpMNZ(Q, I_i) = Irp(Q, I_i) \times \alpha \quad (7)$$

where α is the number of feature response lists in which a training image was found.

2.4 Intermediate levels: multiple-images and organs/views

At this step, we have now a *knowledge* at the image query level combining the different type of features for each image query from a same observation plant query O_q . This knowledge can be a probability distribution if the Weighted Probability approach was used, or a ranked training image list if the BordaMNZ or IrpMNZ methods were used. The image queries from a same plant O_q can be related to the same or distinct organs and views showing different angles of the plant and we try here to take advantage from the complementarity of these views in order to compute a final *knowledge* at the observation level. Next combinations of knowledge are performed in this order:

- image level to organ/view level (if the query plant O_q has several images from one same organ/view V_q),
- then, organ/view level to observation level (if the query plant O_q as several organs V).

Only the Weighted Probability approach (*Run 1*) and the BordaMNZ approach (*Run 2, 3 & 4*) are used in these upper levels of combination following the same formulas in previous subsections ??.

For instance, at organ level with the Weighted Probability approach, given several probability distributions from a set of different organs/views V , we apply the same the weighted fusion used as in subsection ?? where local features F are replaced by V .

2.5 Observation level and decision: the final knowledge to a species list

(*Run 1, 2, 3 & 4*)

At this final step, we have for one query observation plant O_q , a distribution of probabilities if the Weighted Probability approach was used in the lower levels of fusion (*Run 1*), or a final list of training image response if the BordaMNZ was used instead.

In the case of the WP approach, the final list of species is directly given by the probability distribution following a decreasing order of probabilities. In the case of the BordaMNZ approach, we use a last step for converting the image list response into a probability distribution with the CIN rule (see previous section ??).

2.6 Summary of the different submitted runs

Preliminary evaluation and feature selection Considering the available number of visual feature for each images and the distinct methods of combinations, we supposed here that each type of organ/views has a its own relevant subset of visual features exploited no necessarily with the same combination method. We made some intensive preliminary evaluations on the training dataset in order to find these subsets and methods for each type of organ/views. We used a Leave One (Observation) Out procedure: focusing on one organ/view, for image from the training dataset, we use it as an internal query and excluded from the returned ranked image response list all the images belonging to the same plant observation as the query image, in order to remove a bias related to near duplicate images form a same plant. We concluded that in most of the cases, it was more or less the same subsets of features which was pointed out by the 3 fusion methods BordaMNZ, IrpmNZ and WP. Since the BordaMNZ demonstrated the best performances for mostly every organ/view, we choose to select its subset of visual features. As might be expected, color is dominant for the flower but may lead to confusion for leaves, while texture plays an important role for stems and fruits. The features finally selected are presented in table ??:

Table 1. Selected features for each type of organ/view according to the preliminary evaluation on the training dataset with a Leave One (Observation) Out procedure.

	SURF	Fourier	EOH	riLbp	wRGB	wLUV	HSV	SEFH	DFH	Shapes6
Branch			✓	✓	✓	✓				
Entire	✓		✓	✓	✓	✓				
Flower	✓	✓	✓	✓	✓	✓	✓			
Fruit	✓			✓	✓	✓				
Leaf	✓		✓	✓		✓	✓			
Scan								✓	✓	✓
Stem	✓	✓		✓	✓	✓				

Recap chart of the different submitted runs The Table 2 summary of the different methods used at each level of fusion for each submitted run.

PlantNet Run 1 is actually more or less the same method used for the submitted run *Inria PlantNet Run 1* during the previous ImageCLEF 2013 Plant Task [?]. We used this configuration as a baseline in order to see the score this year despite 250 more new species, and to see if we can obtain better scores with the 3 new approaches experimented in the 3 other submitted runs.

PlantNet Run2 is the BordaMNZ version of the *PlantNet Run 1*, where for each step of fusion we changed the Weighted Probability WP combination method by the BordaMNZ method. In this way, only list responses are merged at each step, and finally we use the CIN approach in order to obtain a list of probability of species, and thus, a list of ranked of species according to these probability values.

PlantNet Run 3 is like the *PlantNet Run2* except that we changed the local responses combination with the Ranking Dominate vote (see previous subsection ??).

PlantNet Run 4 is supposed to be the best configuration, where for each level of combination we choose the best method according to the preliminary evaluations on the training dataset. The BordaMNZ is most of the time the retained method, except for scans and stem pictures at the feature combination level for which we had noticed better performances with the IprMNZ combination. At the time of writing the paper, we noticed a mistake with the Fruit organ where we selected a Weighted Probability approach and which probably gave bad combinations in upper level since image response lists were expected instead of probability distributions.

Table 2. Summary of the different methods used at each level of fusion for each submitted run.

Run name	Local	Features	Images	Organs/views	Observation (decision)
PlantNet Run 1	VOB	CIN+WP	WP	WP	
PlantNet Run 2	VOB	BordaMNZ	BordaMNZ	BordaMNZ	CIN
PlantNet Run 3	RD	BordaMNZ	BordaMNZ	BordaMNZ	CIN
PlantNet Run 4	VOB	Fruit: WP Scan: IprMNZ Stem: IprMNZ Others: BordaMNZ	BordaMNZ	BordaMNZ	CIN

Submitted "image" runs For each submitted run, it was recommended but not mandatory to submit a second set of run files detailing species prediction at image level like in the previous ImageCLEF 2013 Plant Task [?]. The aim was to see if the participants can take benefit from the combination of multiple images from a same plant. In order to produce these complementary "image" runs, we simplify the hierarchical fusion of treatment, stopping at the feature combination level for one same image query (see Table 3). Unfortunately, we encountered a problem during the generation of the file (Image) PlantNet Run 3 and we were not able to submit it.

3 Results

Table ?? resumes the official scores for the first 7 submitted runs out of 27 total. Figure ?? gives an overview of all results obtained by the participants. Our four runs give generally good results and are placed in the 7 top runs, and we can say that our team take the lead of the pack of other participating teams, but far behind the impressive results obtained by the IBM AU team. As expected, we slightly improved performances with the new methods introduced this year, the previous method experimented in the *PlantNet Run 1* being the less effective one. Best results are obtained by the full BordaMNZ approach in *PlantNet Run*

Table 3. Summary of the different methods used at each level of fusion for each submitted "image" run.

Run name	Local	Features	Decision
(Image) PlantNet Run 1	VOB	CIN+WP	
(Image) PlantNet Run 2	VOB	BordaMNZ	CIN
(Image) PlantNet Run 3	RD	BordaMNZ	CIN
(Image) PlantNet Run 4	VOB	Fruit: WP Scan: IprMNZ Stem: IprMNZ Others: BordaMNZ	CIN

2. It is disappointing that the new Ranking Vote experimented in *PlantNet Run 3* didn't show some improvement and is strangely very close to the second run *PlantNet Run 2*. Finally, our supposed best run *PlantNet Run 4*, with the best combination at each level of fusion, didn't show the best performances, but it is still better of our baseline *Run PlantNet 1*. Table ?? resumes the scores obtained

Run name	Score
IBM AU Run 4	0,471
IBM AU Run 3	0,459
IBM AU Run 2	0,454
PlantNet Run 2	0,289
PlantNet Run 3	0,289
PlantNet Run 4	0,282
PlantNet Run 1	0,278

Table 4. Official scores of the LifeCLEF 2014 Plant Task obtained by the the first 7 submitted runs out of 27 total.

for the first 6 submitted runs out of 14 total. Figure ?? gives an overview of all results obtained by the participants who submitted these optional runs. Our 3 submitted "image" runs by our team are in the top 6. For each method, like for most of other methods used by the other participants, it highlights the fact that we obtained some slightly improvement by combining the pictures and the organs/views from a same plant observation query.

Table ?? resumes the detailed scores by organs/views obtained for the first 10 submitted runs out of 14 total and Figure ?? gives an overview of all results obtained by the participants who submitted these optional runs. In most cases, our 3 submitted "image" runs by our team are in the top 6, expect for the *Entire* and *Flower* where the 4th run of IBM AU obtained best results, and for the *Stem* where the FINKI team used some better approaches.




Fig. 2. Overview of the official scores of the LifeCLEF 2014 Plant Task. The four submitted runs by our team are in the top 7.

Run	Score on images	Official score (on observations)
IBM AU Run 4	0,456	0,471
IBM AU Run 3	0,446	0,459
IBM AU Run 2	0,438	0,454
PlantNet Run 2	0,28	0,289
PlantNet Run 4	0,276	0,282
PlantNet Run 1	0,271	0,278

Table 5. "Image" scores of the LifeCLEF 2014 Plant Task obtained by the the first 6 submitted "image" runs out of 14 total.

4 Conclusions

Inria within the PlantNet project submitted four runs, and 3 complementary "image" runs, that used the same hierarchical fusion framework, and where each run was related to distinct configurations of methods. All the submitted runs obtained good results and are ranked from the 4th to the 7th place compared to the other submitted run, but clearly behind the best results obtained by the IBM AU team. The first run was more or less the same approach used in the previous ImageCLEF 2013 plant identification task, and was considered as a baseline in our work. We try to improve this approach with simple but effective fusion approaches of ranked lists of results based on the Borda Count and the Inverse




Fig. 3. Overview of the image scores of the LifeCLEF 2014 Plant Task compared with the official scores related to observation.

Run	Branch	Entire	Flower	Fruit	Leaf	Leaf Scan	Stem
IBM AU Run 4	0,292	0,333	0,585	0,339	0,318	0,64	0,269
IBM AU Run 3	0,298	0,34	0,57	0,326	0,304	0,614	0,267
IBM AU Run 2	0,294	0,335	0,555	0,317	0,3	0,612	0,267
PlantNet Run 4	0,112	0,167	0,366	0,197	0,165	0,541	0,152
PlantNet Run 2	0,112	0,181	0,376	0,22	0,164	0,453	0,156
PlantNet Run 1	0,112	0,168	0,366	0,197	0,165	0,449	0,133
IBM AU Run 1	0,103	0,193	0,389	0,161	0,103	0,278	0,138
FINKI Run 1	0,088	0,117	0,255	0,177	0,16	0,4	0,157
FINKI Run 2	0,108	0,099	0,187	0,16	0,14	0,399	0,18
FINKI Run 3	0,088	0,117	0,255	0,177	0,162	0,36	0,159

Table 6. Detailed image scores by organ/views obtained by the the first 10 submitted runs out of 14 total.

Rank Position Algorithms. These two approaches performed slightly better, and we obtained the best results for the full BordaMNZ approach.

Acknowledgments. Part of this work was funded by the Agropolis foundation through the project Pl@ntNet (<http://www.plantnet-project.org/>)




Fig. 4. Overview of the detailed image scores by organ/view.

References

1. Aslam, J.A., Montague, M.: Models for metasearch. In: Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 276–284. SIGIR '01, ACM, New York, NY, USA (2001), <http://doi.acm.org/10.1145/383952.384007>
2. Bakic, V., Mouine, S., Ouertani-Litayem, S., Verroust-Blondet, A., Yahiaoui, I., Goëau, H., Joly, A.: Inria's participation at imageclef 2013 plant identification task. In: CLEF (Online Working Notes/Labs/Workshop) 2013. Valencia, Espagne (2013)
3. Bakić, V., Yahiaoui, I., Mouine, S., Ouertani-Litayem, S., Ouertani, W., Verroust-Blondet, A., Goëau, H., Joly, A.: Inria imedia2's participation at ImageCLEF 2012 plant identification task. In: CLEF (Notebook Papers/Labs/Workshop) (2012)
4. Bay, H., Ess, A., Tuytelaars, T., Van Gool, L.: Surf: Speeded up robust features. In: Computer Vision and Image Understanding (CVIU) (2008)
5. Evans, C.: Notes on the opensurf library. Tech. rep., University of Bristol (2009)
6. Ferecatu, M.: Image retrieval with active relevance feedback using both visual and keyword-based descriptors. Ph.D. thesis, Univ. Versailles St-Quentin (2005)
7. Goëau, H., Bonnet, P., Barbe, J., Bakić, V., Joly, A., Molino, J.F., Barthélémy, D., Boujemaa, N.: Multi-organ plant identification. In: Workshop on Multimedia Analysis for Ecological Data Proceedings. MAED '12 (2012)
8. Goëau, H., Joly, A., Bonnet, P., Selmi, S., Molino, J.F., Barthélémy, D., Boujemaa, N.: Lifeclef plant identification task 2014. In: CLEF working notes 2014 (2014)
9. Goëau, H., Joly, A., Yahiaoui, I., Bonnet, P., Mouysset, E.: Participation of INRIA& Pl@ntNet to ImageCLEF 2011 plant images classification task. In: CLEF (Notebook Papers/Labs/Workshop) (2011)

10. Joly, A., Buisson, O.: Random maximum margin hashing. In: CVPR 2011 (2011)
11. Joly, A., Buisson, O.: A posteriori multi-probe locality sensitive hashing. In: Proceedings of the 16th ACM international conference on Multimedia. MM '08 (2008)
12. Joly, A., Müller, H., Goëau, H., Glotin, H., Spampinato, C., Rauber, A., Bonnet, P., Vellinga, W.P., Fisher, B.: Lifeclef 2014: multimedia life species identification challenges. In: Proceedings of CLEF 2014 (2014)
13. Jovic, M., Hatakeyama, Y., Dong, F., Hirota, K.: Image retrieval based on similarity score fusion from feature similarity ranking lists. In: Wang, L., Jiao, L., Shi, G., Li, X., Liu, J. (eds.) Fuzzy Systems and Knowledge Discovery, Lecture Notes in Computer Science, vol. 4223, pp. 461–470. Springer Berlin Heidelberg (2006)
14. Ojala, T., Pietikäinen, M., Mäenpää, T.: Gray scale and rotation invariant texture classification with local binary patterns. In: Computer Vision - ECCV 2000, Lecture Notes in Computer Science, vol. 1842. Springer Berlin Heidelberg (2000)
15. Yahiaoui, I., Mzoughi, O., Boujemaa, N.: Leaf shape descriptor for tree species identification. In: ICME (2012)
16. Yiu, M.L., Mamoulis, N.: Efficient processing of top-k dominating queries on multi-dimensional data. In: Proceedings of the 33rd international conference on Very large data bases. pp. 483–494. VLDB Endowment (2007)