Curriculum Vitae

Alexis Joly

2020

Sommaire / ***SUMMARY***

Formulaire 1 — Curriculum Vitae 3
Form 1 — Curriculum Vitae 3

Formulaire 2 — Description synthétique de l'activité antérieure 10
Form 2 — Summary of your past activity 10

Formulaire 3 — Contributions majeures 13
Form 3 — Major contributions 13

Formulaire 4 — Liste complète des contributions 19
Form 4 — Complete list of contributions 19

Formulaire 1 — CURRICULUM VITAE
Form 1 — CURRICULUM VITAE

1) Parcours Professionnel / *Professional history*

2020 – 2020	**Permanent researcher** (DR2 Inria employee) Inria, Sophia-Antipolis research center, ZENITH team, LIRMM, Montpellier, France
2010 – 2019	**Permanent researcher** (CR1 Inria employee) Inria, Sophia-Antipolis research center, ZENITH team, LIRMM, Montpellier, France
2006 – 2010	**Permanent researcher** (CR2 Inria employee) Inria, Rocquencourt research center, IMEDIA team, Paris, France
2005 – 2005	**Post-doc position** (Inria employee) associated team between (i) National Institute of Informatics (NII), Tokyo (4 months) and (ii) Inria Rocquencourt, France (8 months)
2003 – 2005	**PhD fellowship** (INA PhD contract) collaboration between (i) National Institute of Audiovisual (INA), France and (ii) La Rochelle Université

2) Prix et distinctions / *Prizes and awards*

- La Recherche Prize, 2016 (13-th edition), delivered by the French magazine La Recherche, jury chaired by Pascale COSSART, perpetual secretary at the French Académie des sciences, awarded paper: Alexis Joly et al. “A look inside the Pl@ntNet experience”. In: *Multimedia Systems* 22.6 (2016), pp. 751–766. DOI: 10.1007/s00530-015-0462-9. URL: https://hal.inria.fr/hal-01182775
- Futur en Seine Awards, 2014, Collaborative Research Award, delivered by the jury of Futur en Seine en Seine festival, awarded innovation: Pl@ntNet platform

3) Encadrement d’activités de recherche / *Supervision of research activities*

- *Supervision of PhDs*

2020 – 2023 **Camille Garcin**, *Set-valued classification in the case of long-tail distributions*, supervision rate of 33%, University of Montpellier. Camille’s PhD just started and is funded in the context of the 3IA program program (chair of J. Salmon). He is currently working on a survey of top-K classification losses in the context of long-tail distributions (plant classification data).

2019 – 2022 **Quentin Leroy**, *Active learning of unseen classes - application to French TV archives*, supervision rate of 50%, University of Montpellier. Quentin’s PhD is funded in the context of an industrial contract with INA. He is currently working on the training of semi-supervised representations from labeled and unlabeled datasets (using self-supervision approaches).

2019 – 2022 **Benjamin Deneu**, *Interpretability of distribution models of plant species communities learned through deep learning - application to crop weeds in the context of agro-ecology*, supervision rate of 50%, University of Montpellier. Benjamin’s PhD is funded in the context of the #DigitAg national initiative dedicated to digital agriculture. He is currently studying the impact of the spatial structure of the environment on plant species distribution through convolutional neural networks. An article was recently submitted to Methods in Ecology and Evolution.

2018 – 2021 **Alena Shilova**, *Scheduling Strategies for High Performance Deep Learning*, supervision rate of 30%, Université de Bordeaux. Alena’s PhD is funded in the context of a national Inria project related to the convergence of HPC, Big Data and machine learning. She’s currently working on optimal checkpoint strategies (based on dynamic programming) for training convolutional neural networks in limited memory. She recently submitted an article to Journal of Machine Learning Research [my12].

2017 – 2020 **Titouan Lorieul,** *Uncertainty issues in fine-grained image classification*, supervision rate of 100% , University of Montpellier. Titouan started with a 6 month internship during his last year of Telecom ParisTech and then continued on to a PhD degree. He is working on the problem of disentangling data ambiguity from model uncertainty for classification with reject option and set-valued classification. A paper was recently submitted to Journal of Machine Learning Research, two journal articles [my23, my65] and two conference papers were published beforehand [my80, my22]. He is currently writing his PhD thesis.

2016 – 2019 **Christophe Botella**, *Species Distribution Modelling from Crowdsourced Data*, supervision rate of 33%, University of Montpellier. In the aim to study the impact of observation bias in Pl@ntNet data for ecological study, I proposed to François Munoz (ecologist) and Pascal Monestiez (statistician) to co-supervise a PhD in the context of an INRA/Inria grant. A journal article [my31], a book chapter [my32] and several conference/workshop papers were published [my35, my14, my30].

2016 – 2019 **Jose Mario Carranza Rojas**, *Digitized Herbarium data analysis*, PhD student at Instituto Tecnológico de Costa Rica, supervision rate of 25%. Erick Mata (former director of Encyclopedia of Life) proposed me to co-supervise the PhD thesis of Jose who spent a 4 months internship in our team. An article was published in BMC Evolutionary Biology journal [my49], another one as a book chapter [my34] and a last one has been recently accepted for publication in APPS journal.

2013 – 2016 **Valentin Leveau**, *Large-scale visual entities recognition*, supervision rate of 50%, University of Montpellier, **defended** November 2016. Valentin is one of the three PhD students I supervised in the context of a long-term collaboration with INA (through an industrial PhD contract). Three publications were accepted at ACM ICMR conference [my74, my88, my106].

2011 – 2013 **Pierre Letessier**, *Frequent Visual Objects Mining in Large Multimedia Collections*, supervision rate of 50%, Telecom ParisTech, **defended** March 2013. Pierre was the first PhD student I supervised in the context of my collaboration with INA (CIFRE). Six papers were published in ACM Multimedia and ACM ICMR conferences [my124, my136, my87, my117, my114] and Multimedia Tools and Applications journal [my99].

2011 – 2013 **Riadh Mohamed Trad**, *Visual-based event mining in Social Media*, supervision rate of 90%, Telecom ParisTech, **defended** May 2013. I supervised Riadh in the context of a EU project on event-based multimedia information retrieval (Glocal). To efficiently solve such issues, he attacked the problem of distributed KNN-graph approximation via hashing. Two publications were accepted at ACM ICMR conference [my139, my128].

2010 – 2012 **Amel Hamzaoui**, *Shared-Neighbours methods for visual content structuring and mining*, supervision rate of 90%, Université Paris-Sud, **defended** May 2012. I supervised Amel in the context of a EU project on cross-modal information retrieval. Her main contribution was a multi-source shared-nearest neighbor clustering algorithm which was applied to various multimodal information retrieval challenges. Two publications were accepted in Multimedia Tools and Applications journal [my134, my99] as well as several demos and workshop papers [my121, 140, my117].

2009 – 2011 **Ahmed Rebai**,*Interactive object retrieval using interpretable visual models*, supervision rate of 90% , university of Paris-sud, **defended** April 2011. Ahmed was the first PhD student I supervised (in the context of the same EU project than Amel). His work was among the first ones considering the learning of sparse representations as a way to improve interpretability. One publication was accepted in Pattern Recognition journal [my126] and two at ACM ICMR conference [my138, my156].

- *Supervision of Masters and internships (not already listed as PhDs)*

2019 – **Mathilde Negri**, *Study of the spatial transferability of species distribution models*, AgroParisTech, 6 months MsC internship, 1 publication.

2019 – **Lucas Bernigaud Samatan**, *Study of the beta-diversity between agricultural and semi-natural environments based on the predictions of a deep convolutional model*, AgroParisTech, 3 months Graduate internship.

2017 – **Valentin Lhoste**, *Data analysis through Citizen science: survey and perspectives*, University of Montpellier, 5 months MsC internship, 1 publication.

2015 – **Ghofrane Ben Ahmed**, Development of a gamified crowdsourcing mobile application, *National Institute of Applied Sciences and Technologies (Tu)*, 5 months graduate internship.

2012 – **Saloua Litayem**, Hash-based Super Vector Machines, Telecom ParisTech, 4 months graduate + 6 months MsC internship, 2 publications.

- *Supervision of Post-docs & engineers*

I have been supervising **7 postdoctoral fellows**: Lee Sue Han (2018-2019), Maximilien Servajean (2015-2016), Julien Champ (2012-2014), Hervé Goëau (2010-2014), Vera Bakic (2011-2013), Sébastien Poullot (2011-2012), Laurent Joyeux (2009-2012). I will recruit two new post-doc in 2020 in the context of EU and French projects.

I have been supervising **about 10 engineers** within EU and French R&D projects.

- *Ability to motivate young researchers*

Six of the PhDs & post-docs I supervised found a job as permanent researcher (Maximilien Servajean at Université Paul Valéry, Hervé Goëau at CIRAD, Pierre Letessier at INA, Riadh Trad and Sebastien Poullot in the media lab of Dassault Systems company, Valentin Leveau at Seenovate).

4) Responsabilités collectives / *Responsibilities*

- *Participation in committees, tasks of collective interest*

  - **Scientific Advisory Board** of the EU FP7 **REVEAL** project (2014-2017), 11 EU partners, 7M euros

  - **Scientific Advisory Board** of the EU FP7 **SocialSensor** project (2013-2015), 12 EU partners, 10M euros

  - **Scientific Advisory Board** of the French National Research Agency Challenge AI for Biodiversity (IA-Biodiv), 2019

  - **Reviewer/Expert** for National and Regional projects: French National Research Agency, National HPC grand equipment (GENCI), French International Cooperation program for south America (**STIC AmSud**)

  - **Writing of white papers and reports for the EU commission:** Cross-disciplinary Challenges and Recommendations regarding the Future of Multimedia Search Engines (2010, available on EU bookshop), Latest Trends in Multimedia Search Computing (2012, white paper), Recommendations about benchmarking campaigns as a tool to foster multimedia search technology transfer at the European level (2013, [my122])

  - Reviewer for impactful journals (selection): Transactions on Pattern Analysis and Machine Intelligence (A$*$, SJR =5.3), International Journal of Computer Vision (A, SJR =3.6), Information Sciences (SJR =2.0), Transactions on Image Processing (A$*$, SJR =1.8), Plant Methods (SJR =1.5), Computers and Electronics in Agriculture (SJR =1.0), Ecological Informatics (SJR =1.0), Machine Learning (A$*$, SJR =1.0)

  - Program committee of major conferences (selection): AAAI, ACM MM, ACM ICMR, CVPR, EDBT, VLDB, TODS, ICIP, ICDE, CLEF, ICASSP

  - Participation to research coordination actions: GdR PARCS (Participatory Research & Citizen Sciences, CNRS working group), **GdR MADICS** (Big Scientific Data, CNRS working group), **GdR ISIS** (Information, Signal, Image and Vision, CNRS working group), **SABIOD** (Scaled Acoustic Biodiversity, CNRS coordination action)

  - PhD thesis examiner: Zongyuan Ge (Queensland university, 2017), Roman Fedorov (Politecnico di Milano, 2017), Lee Sue Han (University of Malaya, 2018)

  - Jury for associate professors: Two at University of Toulon (18-th of May 2017 and 3rd of May 2018), two at CIRAD (Nov. 2020 and 15-th of Feb. 2016)

- *Participation in colloquia organization*

- ACM Multimedia 2018 (~ *800 participants*), organizing committee, area chair
- CLEF 2014-2018 (~ *250 participants yearly*), organizing committee, session chair
- ACM ICMR 2017 (~ *400 participants*), organizing committee, tutorial program chair
- ACM Multimedia 2013 (~ *800 participants*), organizing committee, video program chair
- CBMI 2012 (~ *200 participants*)**,** organizing committee, panel chair
- Workshops: LifeCLEF workshop 2014-2020 *(~ 20-40 participants yearly)*, Floris'Tic 2016-2018 (~ *80 participants yearly*), Environmental Multimedia Retrieval (in the scope of ACM ICMR 2014-2015, ~ *30 participants yearly*), special session on Image Processing and Pattern Recognition for Ecological Applications (ICIP 2013,~ *20 participants*), ISIS workshop on multimedia data linking 2013 (~ *20 participants*), International Workshop on Search Computing 2012 (~ *80 participants*), ImageCLEF 2011-2013 (~ *50 participants yearly*), CHORUS+ think tank on Multimedia Search technology transfer (WWW 2012 conference 2012, ~ *20 participants*), special session on Information Retrieval from Scientific Multimedia Data (CLEF 2011, ~ *20 participants*)

- *Invitations*

- invited keynote speaker: keynote at EGC 2021, keynote at Digital Infrastructures 2018, keynote at Journée scientifique de la Faculté des Sciences et Technologies de Lyon 2014
- invited panelist: HPC-AI-BigData Convergence days (Conv' 2019), ACM Multimedia 2013
- invited talks in conferences: British Ecological Society conference 2019 (Belfast, thematic session on ecology in the digital age), HPC-AI-BigData Convergence days (Conv' 2019, Paris), Forum Teratec 2018 (at Institut Polytechnique de Paris), TDWG 2016 conference (Costa-Rica) IBC 2013 (Amsterdam), Future Internet Assembly 2010 (Ghent), NEM Summit 2008 (Saint-Malo)
- invited talks in workshops/seminars (selection of): University of Florida Phenology workshop 2018, The Open University seminar (2017), ENPC seminar in 2011, CWI seminar in 2009, TRECVID 2008 workshop
- invited demos (selection of): SIDO 2019, Milan universal expo 2015, WWW 2012 conference
- invited lecturer: ERMITES summer school 2015-2016

5) Management (le cas échéant) / *Management (if applicable)*

Indiquez les équipes de recherche et les projets de recherche que vous avez encadrés. Séparez les deux catégories et donnez des éléments permettant d'apprécier l'importance de l'encadrement : taille des équipes, volumes des projets,...

Indicate the research teams and research projects that you have managed. Separate the two categories, and provide elements of appreciation on the importance of the management: size of teams or projects, volume of the projects...

- *Management of research platforms*

2014 – 2020 Foundator & **Scientific coordinator** of LifeCLEF international research platform dedicated to the *computer-assisted identification of living organisms*: ~ 10 collaborators for the organization, ~ 100 registrants each year, ~ 100 publications, ~ 1000 citations

2011 – 2020 Scientific coordinator of the Pl@ntNet platform dedicated to the *crowdsourcing of botanical data through automated plant identification*. Currently: 3 permanent researchers, 3 engineers, 2 PhD students, 1 postdoc, tens of national and international partners, millions of end-users.

- *Management of international projects*

2019 – 2022 Steering board of Cos4cloud, H2020 research project (6M euros) aimed at integrating citizen science in the European Open Science Cloud (EOSC) through the co-design of innovative services to solve challenges faced by citizen observatories, while bringing Citizen Science (CS) projects as a service for the scientific community and the society and providing new data sources. I am leading the work package related to the development of the new services.

2012 – 2014 Steering board of Glocal, EU FP7 research project (7.7M euros) on *event-based information retrieval,* 12 partners (including Yahoo and AFP). I was leading a work package related to visual information retrieval.

2010 – 2012 Steering board of CHORUS+, EU FP7 research Coordination Action (1.3M euros) on *multimedia information retrieval and next generation search engines* coordinating seven EU FP7 research projects and tens of major research actors in Europe. I was leading the work package related to the evaluation of multimedia search methods and technologies.

2009 – 2012 Steering board of VITALAS, EU FP7 research project (8.1M euros) on *cross-media information retrieval,* 11 partners (including Fraunhofer, EADS, CWI, Robotiker). I was leading the work package on scalable content-based information retrieval.

- *Management of French projects*

2020 – 2023 Management of a research axis in the context of **CAMELOT** (CooperAtive MachinE Learning and OpTimization), a 4-years research and educational project funded by French National Research Agency in the context of the 3IA program program (chair of J. Salmon).

2019 – 2021 Management of two projects (60K euros each) funded by the French Ministry of Culture toward two major actors of the French cultural domain: the French National Library (BNF) and the French National institute of audio-visual (INA). Thematic: active learning in TV archives (INA) and digitized collections (BNF).

2018 – 2021 Steering board of WeedElec, 4-years trans-disciplinary project on *automated weeding via agricultural robots and deep learning,* funded by several French Agencies (French National Research Agency, AFB) and government departments (agriculture, high-school teaching, ecology), 5 partners (IRSTEA, INRA, CIRAD, Agrial, Inria). I am leading the computer vision activities of the project.

2018 – 2019 Scientific coordinator *of* **Pl@ntHealth**, 2-years project on *plant disease recognition,* funded by 3 *laboratories of excellence* of different domains: Agropolis (agronomy), NUMEV (computer science) and CEMEB (biodiversity) and involving 4 partners (LIRMM, CEFE, AMAP and Tela Botanica).

2014 – 2018 Steering board of Floris'Tic, 4-years project on *computer & citizen sciences in botany* funded by the French Future Investments Program (PIA), 8 partners (Inria, INRA, CIRAD, IRD, Tela Botanica, CNRS, Univ. of Montpellier, Agropolis foundation). I was leading one of the three components of the project.

2010 – 2013 Steering board of **OTMedia**, 3-years research project on *trans-media data mining* funded by the French National Research Agency (ANR), 7 partners (INA, Inria, AFP, LIA, La Sorbonne, Syllabs, ParisTech LATTS). I was leading the work-package on scalability issues.

2008 – 2010 Steering board of **R2I**, 3-years research project on *Interactive Information Retrieval* funded by the French National Research Agency (ANR), 4 partners (univ. of Caen, Inria IMEDIA, Inria LEAR, Exalead). I was leading the work package on interactive search aspects.

6) Collaborations / *Collaborations*

- Since 2018, I am collaborating with a group of american biologists and computer scientists involved in the iDigBio initiative, in particular **Pamela Soltis** (University of Florida), **Charles Davis** (Harvard university), **Patrick W. Sweeney** (Yale university) and **Elizabeth R. Ellwood** (La Brea Tar Pits and Museum). We co-published several papers on deep learning for phenology and they invited me to several workshops in the US.
- Since 2015, I am collaborating with **Dennis ShaSha** from New York University. We jointly work on different topics in particular on *crowdsourcing models & algorithms* that involve thousands of specialized categories (a publication was accepted in *Transactions on Multimedia*). Recently, we did start another piece of work related to set-valued classfication. A publication was recently submitted to the prestigious journal TPAMI.
- Since 2016, I am collaborating with **Erick Mata** (professor at Instituto Tecnológico de Costa Rica, former director of Encyclopedia of Life). Two of his PhD students and one MsC student did spend a 4 months internship under my co-supervision. I invited him as a keynote speaker of LifeCLEF 2016 workshop. He invited me at the TDWG 2016 international conference and he introduced me to the responsibles of several major international initiatives in biodiversity informatics (GBIF, Encyclopedia of Life, ebird, iDigBio, BHL, etc.). Besides, we are currently working on the creation of a new working group at TDWG dedicated to *machine learning for biodiversity informatics* (jointly

with **Francisco Pando** from Real Jardín Botánico-CSIC, Spain, **Luiz Gadelha** from National Laboratory for Scientific Computing, Brazil). Luiz Gadelha included me in a brazilian program for visiting researchers (visit to be organized in 2019).

- Since 2010, I have a long-term collaboration with **Henning Müller** from HES-SO, Switzerland. We first met within a EU coordination action (CHORUS+) and we continued working together since that time (mainly on the evaluation of multimedia information retrieval in particular in the context of trans-disciplinary research). In 2012, he invited me to organize a track within ImageCLEF. In 2014, I proposed him to co-lead a new lab within CLEF (LifeCLEF). We co-published tens of papers.
- In 2018, I have been collaborating with **Vikram A. Saletore** (Intel, US, AI Products Group) and **Valeriu Codreanu** (Surfsara, ND) on a large-scale experiment aimed at training a world's flora model on a supercomputer of 1656 nodes Intel Skylake8168-24cores. Several talks about this work were given and a submission to *Science* magazine is in preparation.
- In 2014, I started collaborating with **Stefanos Vrochidis** (Scientific Manager of the EU FP7 project MULTISENSOR). We co-edited a special issue related to environmental multimedia data and we co-organized several workshops and special sessions on the topic. Few weeks ago, we finalized the co-edition of a Springer book on *Multimedia Technologies for Environmental & Biodiversity Informatics* (involving about 50 contributors from all over the world).
- In 2008, after I finished my PhD, I have been invited by **Paul Over** from the US NIST (National Institute of Standards and Technology) to co-organize a challenge related to video copy detection in the context of the TRECVID conference series.

7) Enseignement / *Teaching*

- 2020: AgroParisTech (section IODAA) - Deep learning - 10 hours
- 2020: InnObs technical school (for professionals) - Innovations in the observation of seasonal biological events and associated data management - 6 hours.
- 2018-2020: Université Paul Valéry (MsC Miashs) - Machine learning - 10 hours per year
- 2018-2020: Polytech school (University of Montpellier) - Image representation - 6 hours per year
- 2015-2017: University of Montpellier- MsC IMAGINA- Content-based Image Retrieval - 3 hours per year
- 2013-2014: University of Montpellier- MsC Big data management - Content-based Information Retrieval - 3 hours per year
- 2008: University Paris 13 (IUT de Saint-Denis) - Content-based Image Retrieval applied to TV monitoring - 2 hours
- 2006: Centre International de Rencontres Mathematiques- Similarity Search Algorithms and Structures - 5 hours

8) Diffusion de l'information scientifique / *Dissemination of scientific knowledge*

- **LABbota 2018-2020:** I participate to the regional project LABbota supported by the Tela Botanica association in partnership with Science Animation, CNRS, Inria and CIRAD. This regional project aims to (i) federating CSTI actors in the Occitania region around participatory sciences, (ii) fostering the mobilization of citizens and researchers, (iii) spreading plant sciences culture. It is based on 4 national platforms combining participatory science, games and collaborative tools related to plants: Smart'Flore, Observatoire Des Saisons, Pl@ntNet and ThePlantGame. In this context, I participate in the setup and animation of events such a 2-day training course on biodiversity IT tools dedicated to animators and relay actors (at « La Gazette de Montpellier »), a workshop and public debate on science and citizens organized at « La Halle Tropisme », a treasure hunt in the Montpellier plant garden organized in the context of a summer school (with Polytech), an experiment mixing Pl@ntNet and the ESRI Inc. collector application led by a teacher from Lyon.
- **Floris'Tic 2014-2018:** I participate to an educational project aimed at boosting environmental and digital education through innovative tools, especially with young audiences. The project develops and disseminates several tools in this regard including Tela Botanica 's MOOC platform, Smart'Flore, ThePlantGame and Pl@ntNet. The three main objectives of the project are: (i) to form relay actors to the use of these tools (teachers, educators, animators, facilitators, etc.) (ii) to organize educational events in schools, parks, expos, etc. and (iii) to set up and animate citizen science programs. For instance, Floris'Tic co-organizes a program for reducing early school leaving with the French National Education (using Smart'Flore and Pl@ntNet). My role in the project is two folds: (i) I participate to the project management and to the strategic decisions (jointly with Christel Vignau and Pierre Bonnet), (ii) I participate to the setup and animation of some of the events (*e.g.* two working sessions in elementary schools, two days of training organized by the association *les écologistes de l'Euzière*, a participatory workshop in *salon de l'écologie 2016*, yearly animations in *fête de la science*, an itinerant animation *Expomobile Homonumericus 2016-2018*, an animation in the garden of the French Pavilion of *Expo Milano 2015*).

- **Community management**: I spend several hours a week animating Pl@ntNet user community. This includes: (i) animating the community of developers using Pl@ntNet API (about 200 users), (ii) animating Pl@ntNet's social networks (twitter account @PlantNetProject, facebook account @PlantNetProject) and (iii) managing the mailbox contact@plantnet-project.org.
- **Consulting & Interviews**: (i) participation to the scientific advisory board of the AI Family Challenge organized by the NGO Iridescent that *supports girls, children, and their families to identify problems in their communities and find technology based solutions*, (ii) interview on data science in the context of agronomy by Agroof (2019) office.
- **Demos and animation in expos (before or out of scope of above-mentioned projects)**: *European research expo* 2008-2009 (2-days demo on Inria stand), *Salon de l'agriculture 2013* (3-days animation on CIRAD stand), *Berges en seine festival 2013* (setup and training of facilitators), *Futur en Seine festival* 2014 (2-days animation), *50 ans de l'Inria 2017* (2-days demo)
- **Participation to the realization of popularization movies:** " *Pl@ntNet, the application that helps people identify plants*" (Inria production, 900 000 views on Youtube), "*Visual-based Transmedia Event Detection*" (INA production, 500 views on Youtube)
- **Popularization articles (as author):** "Botanists of the twenty-first century" , *Proceedings of the UNESCO, 2014*, " Pl@ntNet, un réseau et des outils pour une recherche participative" , *interstice*, 2016

9) Eléments divers / *Other relevant information*

- Excerpt from the Inria 2019 evaluation of ZENITH team - evaluating experts: Ihab Ilyas, Tiziana Catarci, Marc Boullé: *The evaluation committee has been impressed by the focus and the achievements of this team and recommended setting it as a great example of achieving perfect balance between theory and practice. The contributions are related to both the "progress of knowledge" and the "advancement of technology". It is noteworthy that the focus on technology is remarkable, and that it might be explained by a good ratio between researchers and engineers. The team demonstrates a strong system building culture as a way to test and validate the developed technology. This culture makes it extremely easy to find adoption and collaboration channels with industry and the user community.*

Formulaire 2 — DESCRIPTION SYNTHÉTIQUE DE L'ACTIVITÉ ANTÉRIEURE

Form 2 — SUMMARY OF YOUR PAST ACTIVITY

In a nutshell: My central research field is *multimedia information retrieval* with connex research interests in *machine learning* and *computer vision*. My approach to this field is guided by my knowledge acquired on *high-dimensional data*, *efficient representation learning*, and *probabilistic queries*. Since 2011, I am mostly applying my research to *biodiversity informatics*, a transdisciplinary field in which I am increasingly involved.

Summary of activity: Rather than restricting search to the use of metadata, *multimedia information retrieval* attempts to *index*, *search* and *analyse* digital objects such as images or videos by means of efficient *representations* of their actual content. This is a hard problem because of two fundamental issues: the *dimensionality curse* (*i.e.* the complexity of analyzing and organizing data in *high-dimensional spaces* and the *semantic gap* (*i.e.* the difference in meaning between constructs formed within different representation systems).

During my PhD thesis and the next few years, I addressed these two fundamental problems in the specific context of *content-based copy detection* which aims at retrieving near-duplicates of a candidate image or video within a very large database. In that case, the *semantic gap* can be restricted to the problem of finding *invariants* to some predefined transformations such as translation, cropping, re-scaling, rotation, etc. Therefore, I introduced several *invariant detectors* and *descriptors* [my150, my162, my167, my160] as well as new *matching algorithms* allowing to compute image similarity scores based on epipolar geometry [my148, my106, my136, my142]. Regarding the *dimensionality curse*, my main contributions were about *approximate nearest neighbors search* methods, *i.e.* methods allowing to trade quality for time to efficiently solve nearest neighbor queries in the high dimensional representation space. In particular, I have been working on the principle of *probabilistic queries* within space partitions such as the ones induced by *locality sensitive hashing* or *space-filling* curves [my145, my163, my166]. Overall, my contributions to *content-based copy detection* have had a pretty good impact. I obtained among the best results in several benchmarks or comparative studies [my147, my152, my149, 151]. My contributions have been the basis of many following works in particular two other PhD thesis in the continuity (Julien Law-to, Sebastien Poullot). A patent of which I am the author was deposited by INA [my144] and a product called INA Signature was commercialized and is still used by several major companies in the TV, film and web industry as well as by the French government.

After my work on *content-based copy detection*, I've been progressively moving to more general content-based information retrieval problems for which the semantic gap can not be tackled solely by invariants. In particular, I've been working on *learning-based hash functions* [my135, my128] aimed at preserving the semantic information of the original feature space rather than its original metric structure. Besides, I've been working on the acceleration of the prediction of learned classifiers typically for interactive search applications such as relevance feedback [my125, my143]. Noticeably, we did show in [my125] that any linear classifier of the form $h(\mathbf{x}) = sgn(\mathbf{w}.\mathbf{x} + b)$ can be approximated through hashing by a range query in the Hamming cube (with proved convergence). Last but not least, I spent several years on *shared nearest neighbors* based representations [my126, my134, my88], in particular in the context of fine-grained image classification [my74, my86, my121, my138]. My last contributions to this domain in the 2013-2015 period include: (i) the kernelization of spatially consistent nearest neighbors matches [my74], (ii) the learning of spatially localized visual dictionaries [my74, my99] or (iii), the extension of such representation to the temporal domain for bird sounds recognition [my86]. Nowadays, such representations are outperformed in most cases by representations learned through deep learning. However, there is still a number of tasks for which they are more effective, in particular the ones involving small and rigid visual objects such as logos [my74, my87] or whale's biomarkers [my73]. For instance, some of the algorithms I developed are still in use in some innovative products of INA such as DigInPix [my87].

With the advent of deep learning in the 2010's, I've been progressively moving to that kind of representation learning method within my research activities. In particular, I conducted several pioneering studies on fine-grained classification based on deep convolutional neural networks (CNN) in the biodiversity domain. In [my68] and [my81], we reported for the first time the superiority of deep learning models for plant identification and bird identification. In [my52] and [my67], we did show that CNNs are amazingly robust to the presence of noise and to a dramatic scaling of the number of species (up to 10,000 of species). In a recent study [my29], we came to the conclusion that the performance of recent deep learning models is now very close to the expertise of the best French botanists. In [my65], we did show that it is even possible to identify very hard plant groups such as crop varieties. In [my23, my49, my24], we reported high performance gains regarding the difficult task of annotating herbarium specimens in natural-history collections (as relayed by a news article of Nature journal web site). Two last examples of work with strong potential impact in agriculture are: (i) the *instance segmentation* of crops and weeds by precision agricultural robots [**champ2020weedelc**] and (ii), the identification of plant

disease through the analysis of visual symptoms [my9].

I've recently become interested in *uncertainty issues* related to representation learning and fine-grained image classification. *Uncertainty* is actually crucial when considering classification problems that involve thousands of domain specific labels. A picture of a plant, for instance, contains only a partial information that is usually not sufficient to determine its scientific name with certainty. In [my58], I have been working on modelling such uncertainty in the context of a crowdsourcing system involving experts as well as non expert annotators. Therefore, we introduced a new Bayesian inference framework allowing to learn the annotators' confusion and to optimally assign them new items to be validated. In the continuity of this work, I have been working on a non-parametric Bayesian model allowing to combine annotators' suggestions even when the number of possible labels is undetermined and might change over time [my41]. In mirror to this research, I also started working on the uncertainty of machine-learning based classifiers, in particular deep CNNs. I conducted some experimental studies on the topic [my67, my29, my23] and I currently work with one of my PhD students (Titouan Lorieul) on new methods aimed at disentangling data ambiguity from model uncertainty in deep NNs. In [my22][0], in particular, we were interested in classification with reject option which is one way to address the problem of estimating the uncertainty of a classifier (and which is of very practical usage for any further scientific analysis of automatically classified data). Recent approaches addressing this problem use criteria based on either a confidence measure or a measure of dispersion. However, neither approach explicitly combines the two main sources of uncertainty: task ambiguity, which is intrinsic to the task, and model uncertainty, which arises from data sampling and learning stochasticity. In [my22], we explore how these two quantities can be merged to establish more effective rejection criteria. In particular, we propose a series of methods that combine measures of disagreement and estimates of ambiguity using an ensemble of models. Nevertheless, further analysis shows empirical evidence that the ideal rejector may be a more complex function than such criteria, and may even be counter-intuitive in some cases, implying that it is difficult to find a theoretical rejector criterion that would generalize across tasks and that this remains an open problem.

Since 2017, I am working on a new topic illustrating my progressive thematic mobility towards *biodiversity* and *ecological* informatics: species distribution modelling. Species distribution models (SDM) are widely used for ecological research and conservation purposes. Given a set of spatial occurrences of a set of species, the aim is to infer their spatial distribution and ecological niche over a given territory. This is usually achieved through environmental niche modeling approaches, *i.e.* by predicting the distribution in the geographic space on the basis of a mathematical representation of their known distribution in environmental space (= realized ecological niche). The environment is in most cases represented by climate data (such as temperature, and precipitation) and other variables such as soil type or land cover. My first contribution to this field is to develop a big data approach to the problem [my31, my71, my32]. In [my31], for instance, I've been working on the first niche modelling study purely based on automatically identified species occurrences. It did reveal some new areas of invasion of alien species on the French territory. In [my32, my19, my25], I've been rather working on deep learning approaches to species distribution modelling (*i.e.* to learn efficient representation in the multimodal environmental space itself). The problem of classical ecological approaches is actually that they do not well capture the deviation between the theoretical fundamental niche of the species and the realized abundance in the environmental feature space that can have a very irregular shape due to environment perturbation history or species propagation constraints. Deep neural networks on the other side, are able to learn complex non-linear transformations as well as the spatial autocorrelation of species distribution [my16]. Currently, I am working on the problem of the observation bias in such species distribution models [my13] and on the interpretability of the learned models (article under revision).

As a continuation of my research activities, I am also very involved in software development and dissemination. In particular, for the past ten years, several of my works have been integrated into the Pl@ntNet platform that I coordinate with my botanist colleague Pierre Bonnet from CIRAD. The most visible part of this collaborative platform is a million users mobile application which allows you to identify plants with your smartphone and share your observations in a citizen science approach dedicated to monitoring plant biodiversity. The server-side identification engine (Snoop) is based on my work in high-dimensional similarity search (*e.g.* [my145, my135]) and in representation learning for plant identification (*e.g.* [my29, my67]). Other softwares of the platform also integrates some of my research results, *e.g.* ThePlantGame (70K users) that is based on my work in crowdsourcing [my58], or Smart'Flore (75K users) that integrates some of my results in species distribution modelling [my32]. More generally, my activities in the Pl@ntNet platform are varied in nature. As technical manager of the platform, I supervise the engineers who develop the applications and decide on the main orientations in terms of architecture, infrastructure, new functionalities and technological choices. At present, the volume of data stored is about 110 million plant observations, representing more than 170 million images and about 20 TB of data. With the help of my colleague Pierre Bonnet and the services of the platform's organizations (Inria, INRA, IRD, CIRAD), I also devote a significant amount of time to the data management plan (licenses, privacy, RGPD, security) as well as to dissemination activities such as the organization of field events, the animation of the Pl@ntNet user community or communication actions. Finally, Pl@ntNet is a great lever to weave collaborations in many fields and all over the world. A growing part of my activity over the last two years is to respond to these numerous requests and to ensure the follow-up of

[0] english version under review

those that we follow up on. This includes scientific players of course, but also industry, associations, citizen observatories, museums, managers of natural areas, parks & gardens, and last but not least, international platforms such as GBIF, iDigBio, IUCN or CABI.

Formulaire 3 — CONTRIBUTIONS MAJEURES

Form 3 — MAJOR CONTRIBUTIONS

Major contribution 1 : A posteriori multi-probe locality sensitive hashing

1. Description de la contribution / *Description of the contribution*

This work introduces a new algorithm that overcomes the main drawback of the classical Locality Sensitive Hashing (LSH), *i.e.* its memory consumption that is overlinear in the dataset size. By probing multiple buckets within each hash table in an optimized way, the proposed algorithm allows reducing drastically the number of required hash tables (with a proved lower complexity).

2. Contribution personnelle de la candidate ou du candidate / *Personal contribution of the applicant*

I am at the origin of the theoretical contributions of this work, as well as the main actor of its implementation, experimentation and publication.

3. Originalité et difficulté / *Originality and difficulty*

Contrary to previous multi-probe methods (including some of my previous work) that were based on likelihood criteria that a given bucket contains query results, AMP-LSH introduces a more reliable a posteriori model taking into account some prior about the queries and the searched objects. This prior knowledge allows a better quality control of the search and a more accurate selection of the most probable buckets. In this way, it is closer to a metric learning approach and it offers a wider variety of usage than nearest neighbors search in Euclidean space. The most difficult part and originality of this work was to understand that the classical multi-probe approach relies on a so called *prosecutor's fallacy* in its reasoning, and then to generalize it to any prior distribution. The second main difficulty was to find an algorithm allowing to select the most probable buckets in an efficient way and without computing the probability of all buckets of the partition. This was achieved through the use of look-up tables for the marginal probabilities and through an extension of the initial sequence probing algorithm of Lv et al. [155] for the classical multi-probe LSH.

4. Validation et impact / *Validation and impact*

We did show that our a posteriori scheme outperforms other LSH-based search algorithms both in space and time efficiency while offering a better quality control. It has been the source of inspiration of various following works on high-dimensional data hashing such as *density sensitive hashing*, *topology preserving hashing* or *complementary projection hashing*. The method was integrated in INA Signature, a video copy detection technology commercialized by INA company and still used by several major companies in the TV, film and web industry (TF1, Canal+, M6, Dailymotion, TDF, EuropaCorp) and by the French government (in the context of Hadopi law). It was also integrated within Snoop, a C++ image retrieval framework developed jointly with Olivier Buisson (researcher/engineer at INA) and Jean-Christophe Lombardo (engineer at Inria). Snoop is part of the back office of the Pl@ntNet application as well as other applications such as DigInPix (diffused by INA) or a whale identification prototype currently developed in collaboration with Cetamada NGO.

5. Diffusion / *Dissemination*

The initial publication [my145] was cited by more than 183 scientific papers according to Google scholar. Some of my works making use of AMP-LSH also received a pretty good dissemination such as [my142] about logo retrieval (*162 citations*) or [my139] about visual-based event retrieval (*24 citations*).

Major contribution 2 : Random Maximum Margin Hashing

1. Description de la contribution / *Description of the contribution*

This work introduces a new hash function family composed of learned randomized projections. Similarly to Locality Sensitive Hashing (LSH), it relies on random projections as a way to embed high-dimensional feature vectors into compact hash codes that approximately preserve the topology of the original space. The main difference, however, is that the random projections are learned through machine learning techniques in order to better fit the data distribution and provide a better selectivity. More precisely, RMMH family is defined by the set of hyperplanes maximizing the margin between randomly selected sets of training samples:

$$h(\mathbf{x}) = \mathbf{w}_m.\mathbf{x} + b_m$$
$$(\mathbf{w}_m, b_m) = \arg\max_{\mathbf{w},b} \frac{1}{\|\mathbf{w}\|} \quad s.t. \ (\mathbf{w}.\mathbf{x}_j + b)y_j \geq 1 \forall j$$

where the x_j are selected at random from the data and the labels y are selected at random from a Bernouilli distribution with probability $p = 0.5$.

2. Contribution personnelle de la candidate / du candidat / *Personal contribution of the candidate*

I am at the origin of the theoretical contributions of this work, as well as the main actor of its implementation, experimentation and publication.

3. Originalité et difficulté / *Originality and difficulty*

Contrary to other high-dimensional data hashing methods attempting to control the collision probability of close points, our method focus on maximizing the entropy of the learned representation space. By training purely random splits of the data, regardless the closeness of the training samples, it is indeed possible to generate consistently more independent hash functions. On the other side, the use of support vectors machines allows to maintain good generalization performance on any data type, in particular in kernelized spaces. More generally, this was the first work introducing the principle of *randomly trained classifiers* (*i.e.* classifiers trained using random labelling of the training data). Whereas it can appear as counter intuitive at a first glance, we did show that this is a powerful statistical learning paradigm allowing to map any dataset onto an approximately uniform representation space. Recently, this principle of training classifiers using random labelling has been explored further [79, 45]. Zhang et al. [79] in particular, show that convolutional networks can be easily fit on randomly labeled training data and discuss the contradiction of this phenomenon with the classical statistical learning theory of generalization.

4. Validation et impact / *Validation and impact*

The efficiency and genericity of RMMH for *approximate nearest neighbors search* has been shown in many works. In particular, it was re-implemented by several authors of impactful publications such as *spherical hashing*, *hash bit selection*, *reciprocal hash tables* or *collaborative hashing*. Because of the high entropy of the hash codes learned by RMMH, we also did show in following works that it is particularly efficient for large-scale data analysis problems [my128, my124] in particular in distributed contexts. Recently, I have been contacted by Peter Christen (Australian National University), a specialist of record linkage issues and who did show that RMMH consistently outperforms state-of-the-art privacy-preserving records linkage approaches based on bloom filters. Besides, RMMH has been integrated within Snoop software and is currently running within the back end of the Pl@ntNet application.

5. Diffusion / *Dissemination*

The initial publication [my135] was cited by 170 scientific papers according to Google scholar. Some of my works making use of RMMH also received a pretty good dissemination such as [my103] about plant images retrieval (*90 citations*) or [my116] about visual objects mining (*27 citations*).

Major contribution 3 : Representation learning and interactive search for plant identification

1. Description de la contribution / *Description of the contribution*

This work is a set of algorithmic and methodological contributions related to plant identification.

2. Contribution personnelle de la candidate / du candidat / *Personal contribution of the candidate*

I am the principal investigator of this research line. I am at the origin of most algorithmic and methodological contributions as well as one of the main actor of their implementation, experimentation and publication.

3. Originalité et difficulté / *Originality and difficulty*

When I started working on plant identification in 2010, the state-of-the-art in this domain was mostly based on *leaf shape boundary features* and *shape matching* methods. The studied datasets (such as the *Smithsonian leaf dataset*) were usually created in a well controlled environment and were far away from what should be expected to set up a real-world system. In the context of the research activities I supervised and conducted within the Pl@ntNet project, we did contribute to deeply transform this domain in a few years. First of all, we did show that other organs of the plant than leaves can be much more discriminant and that it is highly beneficial to combine several views of a specimen to identify it (*e.g.* in [my119, my81, my29]). During decenies, computer scientists focused solely on the leaf because it was easier to digitize and easier to process with state-of-the-art techniques (segmentation, shape matching, etc.). In parallel, we attacked the hard problem of the bias and of the lack of diversity in the training data. Most studies were actually completely biased by some specific acquisition protocols, such as taking many pictures of the same few specimens with the same acquisition device. To overcome this issue, we did set up a large-scale crowdsourcing platform that is nowadays still producing visual training data available for the scientific community [my132, my70, my71]. As another contribution, we did show the importance of interactive search mechanisms in the identification process (*e.g.* in [my103, my44]). Visualizing the images of the training set that are the most similar to the query is actually crucial for refining and/or validating the results of the automatic classification (a single species can actually have a very diverse morphology). Last but not least, we did progressively scale up the number of species by several orders of magnitude compared to previous work (*e.g.* 10K species in [my52, my29] or 1K species in [my46] based on herbarium data). In 2018, we succeeded in training a convolutional neural network at the scale of the world's flora (300K species), thanks to a long-term experiment making use of several supercomputers provided by the French national HPC equipment (genci).

4. Validation et impact / *Validation and impact*

Most of my contributions to plant identification have been integrated within the Pl@ntNet platform. When we released the first mobile application in 2013, it was a real challenge because the brute force identification performance was still limited. However, this was a necessity to prime the pump and collect training data. Since that time, the number of users doubled every year. In 2019, 9.5M of people used the application at least once and 1.6M of them use it more than 20 times. About 12% of the users declare using Pl@ntNet for their professional activity (in various domains such as agronomy, education, tourism or plant's trade). In particular, a large number of teachers of all grade levels use it in the context of their courses. Besides, Pl@ntNet data is shared through the world's largest collection of biodiversity occurrences (GBIF), supported by governments. Significant scientific findings in ecology start to be produced such as [my31] where we did show that Pl@ntNet data stream is highly relevant for the monitoring of invasive species.

5. Diffusion / *Dissemination*

The methodology and data resulting from Pl@ntNet are at the basis of hundreds of publications related to plant identification. The majority of work published after 2016 on plant identification cite Pl@ntNet as a reference and experiment their approach on Pl@ntNet data. More broadly, this has given birth to the LifeCLEF evaluation forum that is interested in the identification of living organisms in general. The bird sound challenge (BirdCLEF), for instance, has become one of the high points of bio-acoustics and has pioneered the use of deep learning in this field. Thanks to the high dissemination of Pl@ntNet and LifeCLEF, we initiated strong collaborations with researchers in ecology and biology, in particular for the analysis of natural collections (Harvard, Yale, university of Florida, Kew botanical garden, university of Lausane, CSIC, etc.). We also interact with some big industrial players who who want to showcase their technology on Pl@ntNet data (Intel, Google) or simply be associated as sponsors (Microsoft). Finally the diffusion of Pl@ntNet in the civil society and in the media is also very strong (web, TV, radio, press, social networks, etc.). We collaborate ourselves with a significant number of associations, teachers and natural areas managers but the number of people or organizations setting up a project on their own based on Pl@ntNet is even much bigger (we can infer easily such usages from the collected data).

Major contribution 4 : Crowdsourcing Thousands of Specialized Labels

1. Description de la contribution / *Description of the contribution*

This contribution is a set of theoretical contributions and data-driven algorithms to allow the crowdsourcing of thousands of specialized labels thanks to the pro-active training of the annotators. The whole framework allows to (i) train annotators on how to disambiguate among automatically generated candidate labels, (ii) evaluate the quality of annotators' label suggestions and (iii) weight predictions. The framework relies on deep learning, variational Bayesian inference and task assignment to adapt to the skills of each annotator both in the questions asked and the weights given to their answers. The underlying judgements are Bayesian, based on adaptive priors. To achieve live experiments, the whole framework has been implemented in a serious game available on the web (ThePlantGame).

2. Contribution personnelle de la candidate / du candidat / *Personal contribution of the candidate*

I am at the origin of the global architecture and of half of the theoretical and algorithmic contributions of this work. I participated to their implementation and experimentation jointly with Maximilien Servajean.

3. Originalité et difficulté / *Originality and difficulty*

In classical crowdsourcing frameworks, the labels correspond to well known or easy-to-learn concepts so that it is straightforward to train the annotators by giving a few examples with known answers. Neither is true when there are thousands of complex domain-specific labels. The originality of this work is to focus on such annotations that usually require hard expert knowledge (such as plant species names, architectural styles, medical diagnostic tags, etc.). We consider that common knowledge is not sufficient to perform the task but any people can be taught to recognize a small subset of domain-specific concepts. In such a context, it is best to take advantage of the various capabilities of each annotator through teaching (annotators can enhance their knowledge), assignment (annotators can be focused on tasks they have the knowledge to complete) and inference (different annotator propositions can be aggregated to enhance labeling quality). One of the hardest difficulty was to scale the Bayesian inference to the millions of parameters of the confusion matrices of the annotators. The second difficulty was to take into account the uncertainty due to the partial knowledge of the annotators' confusion when they are trained on quizz of limited size.

4. Validation et impact / *Validation and impact*

The live experiments conducted via ThePlantGame did show that pro-actively training the annotators and optimizing the tasks assignment allow a considerable performance gain compared to classical crowdsourcing frameworks. It is sufficient to have a very small fraction of expert annotators in the pool to reach a high accuracy overall. Because of its success, theplantgame has been kept running after the experimentation phase (after a slight refactoring). Nowadays, it accounts for 70K registered players who annotated more than 350K plant observations coming from different sources (Tela Botanica, iSpot & Pl@ntNet).

5. Diffusion / *Dissemination*

The initial publication was accepted for publication in *IEEE transactions on multimedia* [my58] and a variant of the inference algorithm was accepted in the high impact *information sciences* journal [my41]. ThePlantGame has been integrated within several courses of Tela Botanica's MOOC platform dedicated to various domains related to plants. Moreover, it is used by a growing number of teachers of all grade levels.

Major contribution 5 : A deep learning approach to species distribution modelling

1. Description de la contribution / *Description of the contribution*

This work is the first attempt in applying deep feedforward neural networks and convolutional neural networks in particular to species distribution modeling (SDM). It introduces and evaluates several architectures based on a probabilistic modeling suited for regression on count data (Poisson regression).

2. Contribution personnelle de la candidate / du candidat / *Personal contribution of the candidate*

I am at the origin of half of the theoretical, methodological and algorithmic contributions of this work. The implementation and experimentation were mainly achieved by Christophe Botella (under my supervision). Pierre Bonnet and François Munoz contributed to the ecological interpretation and publication of the results.

3. Originalité et difficulté / *Originality and difficulty*

SDM is usually achieved through *environmental niche modeling* approaches, *i.e.* by predicting the distribution in the geographic space on the basis of a representation in the environmental space. This environmental space is in most cases represented by climate data and other variables such as soil type, land cover, distance to water, etc. Then, the objective is to learn a function that takes the environmental feature vector of a given location as input and outputs an estimate of the abundance of the species. Classical SDM approaches postulate that the relationship between output and environmental variables is relatively simple, typically based on a linear combination of the input variables and of their pairwise interactions. However, it sets a strong prior constraint without a clear theoretical founding as the explanatory factors of a species presence can be related to complex environmental patterns. To overcome this limitation, we propose using deep feedforward neural networks (DNN) because they favor high order interaction effects between the input variables, without constraining too much their functional form. Beyond, we proposed to train convolutional neural networks (CNN) on multi-channel spatial rasters in order to capture the spatial autocorrelation of the environmental variables (still with a Poisson regression loss for the optimization).

4. Validation et impact / *Validation and impact*

Our experiments on 50 plant species of the French flora confirmed that DNN and CNN models clearly outperform classical approaches used in ecology (*e.g.* Maxent). They are able to learn non-linear transformations of input environmental features that are relevant for every species without having to think a priori about variables correlation or selection. CNN can capture extra information contained in spatial patterns of environmental variables and surpass DNN. We also did show that the models trained on higher number of species in output (from 50 to 1000) stabilize predictions across species and improve them globally. This is probably the most important outcome of our study. It opens new opportunities for the development of ecological studies based on the use of CNN and DNN (*e.g.* for the study of communities).

5. Diffusion / *Dissemination*

This work was first published in a Springer book about *multimedia tools for environmental & biodiversity informatics* [my32] and a deeper analysis of the proposed method has been recently subbmited to Methods in Ecology and Evolution. It has inspired the creation of a new challenge in the context of LifeCLEF organized jointly with Elijah Cole (from Caltech) with the support of Microsoft AI for Earth's program.

Formulaire 4 — LISTE COMPLÈTE DES CONTRIBUTIONS[1]

Form 4 — COMPLETE LIST OF CONTRIBUTIONS[1]

1. Publications caractéristiques/*representative publications*

Alexis Joly and Olivier Buisson. “Random Maximum Margin Hashing”. In: *CVPR’11 - IEEE Computer Vision and Pattern Recognition*. Colorado springs, United States: IEEE, June 2011, pp. 873–880. DOI: 10.1109/CVPR.2011.5995709. URL: https://hal.inria.fr/hal-00642178

Maximilien Servajean et al. “Crowdsourcing Thousands of Specialized Labels: A Bayesian Active Training Approach”. In: *IEEE Transactions on Multimedia* 19.6 (June 2017), pp. 1376–1391. DOI: 10.1109/TMM.2017.2653763. URL: https://hal.archives-ouvertes.fr/hal-01629149

Alexis Joly et al. *Multimedia Tools and Applications for Environmental & Biodiversity Informatics*. OUV. Springer International Publishing, 2018. DOI: 10.1007/978-3-319-76445-0. URL: https://hal-lirmm.ccsd.cnrs.fr/lirmm-01959343

Christophe Botella et al. “Bias in presence-only niche models related to sampling effort and species niches: Lessons for background point selection”. In: *PLoS ONE* 15.5 (May 2020), e0232078. DOI: 10.1371/journal.pone.0232078. URL: https://hal.archives-ouvertes.fr/hal-02639237

Pierre Bonnet et al. “How citizen scientists contribute to monitor protected areas thanks to automatic plant identification tools”. In: *Ecological Solutions and Evidence* 1.2 (2020). DOI: 10.1002/2688-8319.12023. URL: https://hal.inrae.fr/hal-02937618

2. Développements technologiques : logiciel ou autre réalisation / *Technology development : software or other realization*

The following list of technology developments is sorted by decreasing importance of their impact. For each development, we used the auto-evaluation criteria for software characterization.

- Pl@ntNet: Image-based Plant identification (**software characterization:** A-5, SO-4, SM-3-up4, EM-3-up4, SDL-5; **own contribution**: DA-4, CD-2, MS-2, TPM-4) -
- ThePlantGame: A gamified crowdsourcing application (**software characterization:** A-4, SO-4, SM-2-up3, EM-3, SDL-4; **own contribution**: DA-4, CD-1, MS-2, TPM-4) - it relies on deep learning, variational Bayesian inference and tasks assignment as a way to proactively train the annotators to solve complex tasks (*i.e.* plant identification). Released in April 2016, it accounts for more than 70K players and it is extensively used in the context of the MOOC’s developed by Tela Botanica.
- Snoop: Content-based Image Retrieval Framework (**software characterization:** A-4, SO-4, SM-4, EM-3, SDL-3; **own contribution**: DA-3, CD-3, MS-1, TPM-4) - it integrates most of my works on hashing, approximate k-nn search & visual retrieval as well as other methods of the state-of-the-art. One of its main advantage is its genericity in terms of query types, metrics and data types. In particular, it integrates a deep learning module allowing to learn supervised or unsupervised visual features instead of relying on hand-crafted visual features. It is integrated in Pl@ntNet and DigInPix (see below) and there are currently discussions to integrate it into at BNF, Musée d’Orsay and INHA in the context of a collaboration between Inria and French Ministry of Culture.
- INA Signature: content-based video copy detection framework (**software characterization:** A-4, SO-4-down3, SM-4, EM-3, SDL-3; **own contribution**: DA-4, CD-2, MS-1, TPM-1) - a video copy detection technology integrating my PhD work and commercialized by INA company (one of the world leader in digital archiving). The technology is still used by several major companies in the TV, film and web industry (TF1, Canal+, M6, Dailymotion, TDF, EuropaCorp) and by the French government (in the context of Hadopi law).
- Smart’Flore: an Android mobile application for the discovery of the surrounding vegetal biodiversity (**software characterization:** A-4, SO-2-up4, SM-3, EM-3, SDL-4; **own contribution**: DA-3, CD-1, MS-1, TPM-4) - Smart’Flore app includes three main features: (i) the geo-based exploration of the world’s largest repository of biodiversity occurrences (GBIF), (ii) the exploration of virtual botanical trails (created offline through a dedicated web application hosted by Tela Botanica NGO) and (iii) the access to a variety of information about the plants. Smart’Flore is the first mobile app in the world making use of the GBIF web services which makes it a remarkable and possibly highly visible realization. The first public version of the app was released in March 2016. Nowadays, it has been downloaded by more than 75K users and the daily number of sessions is about 250.

- DigInPix: Visual Named-Entities Identification (**software characterization:** A-3, SO-3, SM-3, EM-3, SDL-3; **own contribution**: DA-3, CD-1, MS-1, TPM-1) - This is a public web application developed on top of Snoop by one of my former PhD student (Pierre Letessier) at INA. It allows recognizing thousands of named-entities (companies, brands, organizations) based on their visual identity (logos, banners, templates, etc.).
- TRECVID imedia tool: synthetic video copies generator developed for TRECVID evaluation campaigns (**software characterization:** A-3, SO-2, SM-2, EM-1, SDL-1; **own contribution**: DA-4, CD-2, MS-1, TPM-4) - it was used by tens of research teams in the world to evaluate content-based video detection technologies in particular within the TRECVID campaign organized by the US National Institute of Standards and Technology between 2008 and 2011.

3. Impact socio-économique et transfert / *Socio-economic impact and transfer*

3.1 What kind of action are you involved in?

I am involved in the transfer and valorization of PI@ntNet, a large-scale participatory platform relying on image-base plant identification as a way to involve non-expert users in the observation of plant biodiversity.

3.2 Under what conditions, in what framework, is the action done? Who are your interlocutors ?

PI@ntNet is sustained as one of the actions of InriaSOFT program. InriaSOFT aims to promote the software developed within the research teams of Inria and its partners, by building consortia around some major software packages. These are aimed at organizing communities of developers and users of these software to promote industrialization, sustainability of technical support and future developments. For this, the consortia fund the engineering costs necessary to support the software and animate the community through an annual membership fee. Figure 1 displays the structure of the consortium for PI@ntNet (with three different levels of membership). The five founder organisms already signed the consortium agreement and discussions with several other potential members are well advanced.

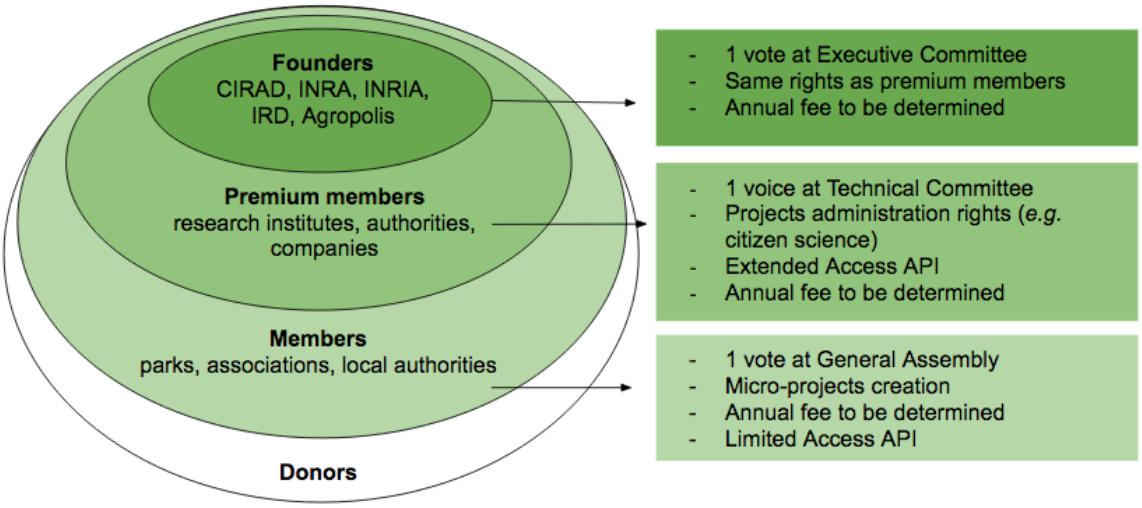

Figure 4.1 – Structure of PI@ntNet consortium

The economic model of the consortium is based on 4 main concepts of value creation:

For research organisations: a service for collecting and managing participatory science data on a large scale and shared in order to reduce the cost. Access to this service is mainly through membership fees to the consortium (10K or 20K euros per year depending on the level of membership).

For environmental education and territorial management stakeholders: attractive tools that are easily adaptable and simple to implement. Revenues in this context come from the contextualization of PI@ntNet tools to specific projects implemented by stakeholders (associations, managers of natural or anthropized areas, educational organizations, local authorities, specialized services). This is done either through direct contract with them or through the participation to call for projects or call for tenders (about 50K euros in 2019).

For companies developing applications: access to specialised and high-performance services at a progressive cost. The use of the PI@ntNet platform functionalities as a service is invoiced in pay-as-you-go mode. The use of the PI@ntNet API is free of charge up to 50 identification requests per day and for specific scientific and educational

purposes. Beyond that, we ask for a contribution of 3 euros per 1000 identification requests. There are currently nearly two hundred users of the API, including about ten in paid mode (having generated a few thousand euros in 2019).

For end users of Pl@ntNet applications: a collaborative encyclopedia and a pocket botanist. The revenue generated by end users is mainly through the collection of donations (60K euros in 2019).
The increase in these different sources of revenue will gradually cover the human and IT infrastructure costs necessary for the smooth running and evolution of the platform over the long term. Eventually, in 2023/2024, it is envisaged that 3 people will be employed on permanent contracts by Inria as part of the InriaSOFT scheme, for an annual budget of around 300K euros. If these objectives are achieved, the question of the creation of an autonomous foundation could be addressed.

3.3 Who are the targeted end users?

Pl@ntNet accounts for million of end-users (*e.g.* 9M in 2019). A majority of users exploit Pl@ntNet for their recreation or personal interest (88%). Among the remaining 12% who use it for professional purposes, the most frequently represented activity is landscape management (34.6%). It includes landscape workers, managers and architects, as well as foresters. The second category is more concerned with the production and/or transfer of knowledge (23.5%), *i.e.* teachers (in botany, biology, horticulture but also at primary school or middle school), students (in horticultural production for example), trainers (landscape management, aromatherapy, herbal medicine, etc.), facilitators (botanists, nature guides) and scientists (biologists mainly). The category of ground workers represents 16% of professional respondents. This category includes farmers, nurserymen, horticulturists and gardeners.

3.4 What is your involvement?

I am the chief scientist of the whole platform since 2011. Besides scientific and management aspects, my main role in Pl@ntNet can be qualified as *chief technology officer*: I am in charge of the technical innovation and the deployment of technologies adapted to the development and the efficiency of the platform. As in a start-up however, I am involved in a wider variety of activities including front-ends design, infrastructure management, fundraising, press & public relations, juridical issues, etc. Concerning the transfer and valorization of Pl@ntNet in the context of the InriaSOFT action, I am the principal interlocutor.

3.5 From your viewpoint, have you reached your objective? What is the real impact?

The initial objectives of Pl@ntNet have clearly been surpassed. There is already an observable societal impact in that Pl@ntNet contributes to a better knowledge of biodiversity in several domains including education, landscape management and agriculture. More indirect impacts could be observed in the long-term, in particular on the development of agro-ecology and precision agriculture for which plant identification is a central issue.
From a research perspective, Pl@ntNet already has a strong impact in biodiversity informatics and we observe a growing impact in other fields such as ecology, evolutionary biology or agronomy. In particular, several new research projects involving Pl@ntNet in that domains were accepted recently or are under construction (one regional project on plant epidemiology, two French projects on precision agriculture, two international projects on phenology and ecology).

4. Toutes les publications/*All publications*

4.1 Revues internationales/*International journals*

→ The following publications are sorted by descending year of publication. Top quality journals and articles with a high impact are highlighted in bold.

[my1] Pierre Bonnet, Alexis Joly, Jean-Michel Faton, Susan Brown, David Kimiti, Benjamin Deneu, Maximilien Servajean, Antoine Affouard, Jean-Christophe Lombardo, Laura Mary, et al. "How citizen scientists contribute to monitor protected areas thanks to automatic plant identification tools". In: *Ecological Solutions and Evidence* 1.2 (2020), e12023.

[my2] Pierre Bonnet, Alexis Joly, Jean-Michel Faton, Susan Brown, David Kimiti, Benjamin Deneu, Maximilien Servajean, Antoine Affouard, Jean-Christophe Lombardo, Laura Mary, Christel Vignau, and François Munoz. "How citizen scientists contribute to monitor protected areas thanks to automatic plant identification tools". In: *Ecological Solutions and Evidence* 1.2 (2020). DOI: 10.1002/2688-8319.12023. URL: https://hal.inrae.fr/hal-02937618.

[my3] **Christophe Botella, Alexis Joly, Pascal Monestiez, Pierre Bonnet, and François Munoz. "Bias in presence-only niche models related to sampling effort and species niches: Lessons for background point selection". In: *PLoS ONE* 15.5 (May 2020), e0232078. DOI: 10.1371/journal.pone.0232078. URL: https://hal.archives-ouvertes.fr/hal-02639237.**

[my4] Julien Champ, Adán Mora-Fallas, Hervé Goëau, Erick Mata-Montero, Pierre Bonnet, and Alexis Joly. "Instance segmentation for the fine detection of crop and weed plants by precision agricultural robots". In: *Applications in Plant Sciences* 8.7 (2020). DOI: 10.1002/aps3.11373. URL: https://hal.inrae.fr/hal-02910844.

[my5] **Charles C. Davis, Julien Champ, Daniel Park, Ian Breckheimer, Goia Lyra, Junxi Xie, Alexis Joly, Dharmesh Tarapore, Aaron M Ellison, and Pierre Bonnet. "A New Method for Counting Reproductive Structures in Digitized Herbarium Specimens Using Mask R-CNN". In: *Frontiers in Plant Science* 11 (July 2020). DOI: 10.3389/fpls.2020.01129. URL: https://hal.inrae.fr/hal-02909794.**

[my6] Hervé Goëau, Adán Mora-Fallas, Julien Champ, Natalie L. Rossington Love, Susan Mazer, Erick Mata-Montero, Alexis Joly, and Pierre Bonnet. "A new fine-grained method for automated visual analysis of herbarium specimens: A case study for phenological data extraction". In: *Applications in Plant Sciences* 8.6 (June 2020), #e11368. DOI: 10.1002/aps3.11368. URL: https://hal.inrae.fr/hal-02894994.

[my9] **Sue Han Lee, Hervé Goëau, Pierre Bonnet, and Alexis Joly. "New perspectives on plant disease characterization based on deep learning". In: *Computers and Electronics in Agriculture* 170 (Mar. 2020), p. 105220. DOI: 10.1016/j.compag.2020.105220. URL: https://hal.umontpellier.fr/hal-02470280.**

[my10] Katelin D. Pearson, Gil Nelson, Myla Aronson, Pierre Bonnet, Laura Brenskelle, Charles C. Davis, Ellen Denny, Elizabeth R. Ellwood, Hervé Goëau, J Mason Heberling, Alexis Joly, Titouan Lorieul, Susan Mazer, Emily Meineke, Brian Stucky, Patrick W. Sweeney, Alexander White, and Pamela S. Soltis. "Machine Learning Using Digitized Herbarium Specimens to Advance Phenological Research". In: *Bioscience* 70.7 (2020), pp. 610–620. DOI: 10.1093/biosci/biaa044. URL: https://hal.umontpellier.fr/hal-02573627.

[my23] **Titouan Lorieul, Katelin D. Pearson, Elizabeth R. Ellwood, Hervé Goëau, Jean-Francois Molino, Patrick W. Sweeney, Jennifer M. Yost, Joel Sachs, Gil Nelson, Pamela S. Soltis, Pierre Bonnet, Alexis Joly, and Erick Mata-Montero. "Toward a large-scale and deep phenological stage annotation of herbarium specimens: Case studies from temperate, tropical, and equatorial floras". In: *Applications in Plant Sciences* 7.3 (Mar. 2019), e01233. DOI: 10.1002/aps3.1233. URL: https://hal.umontpellier.fr/hal-02137748.**

[my26] Kwangjin Park, Alexis Joly, and Patrick Valduriez. "DZI: An Air Index for Spatial Queries in One-dimensional Channels". In: *Data and Knowledge Engineering* 124 (Nov. 2019), p. 101748. DOI: 10.1016/j.datak.2019.101748. URL: https://hal-lirmm.ccsd.cnrs.fr/lirmm-02386429.

[my28] Sylvie Blangy, Valentin Lhoste, Céline Arnal, Jennifer Carré, Anaïs Chapot, Isabelle Chuine, Gaëlle Darmon, Alexis Joly, Pascal Monestiez, and Pierre Bonnet. "Au-delà de la collecte des données dans les projets de sciences citoyennes : ouvrir le champ de l'analyse et de l'interprétation des données aux citoyens." In: *Technologie et innovation* 18- 3.Innovations citoyennes (Sept. 2018). DOI: 10.21494/ISTE.OP.2018.0289. URL: https://hal.archives-ouvertes.fr/hal-01824900.

[my31] Christophe Botella, Alexis Joly, Pierre Bonnet, Pascal P. Monestiez, and François Munoz. "Species distribution modeling based on the automated identification of citizen observations". In: *Applications in Plant Sciences*. Green Digitization: Online Botanical Collections Data Answering Real-World Questions 6.2 (Mar. 2018), pp. 1–11. DOI: 10.1002/aps3.1029. URL: https://hal.umontpellier.fr/hal-01739481.

[my41] **Maximilien Servajean, Romain Chailan, and Alexis Joly. "Non-parametric Bayesian annotator combination". In: *Information Sciences* 436-437 (Apr. 2018), pp. 131–145. DOI: 10.1016/j.ins.2018.01.020. URL: https://hal-lirmm.ccsd.cnrs.fr/lirmm-01703020.**

[my49] **Jose Carranza-Rojas, Hervé Goëau, Pierre Bonnet, Erick Mata-Montero, and Alexis Joly. "Going deeper in the automated identification of Herbarium specimens". In: *BMC Evolutionary Biology* 17.1 (Dec. 2017), p. 181. DOI: 10.1186/s12862-017-1014-z. URL: https://hal.inria.fr/hal-01580070.**

[my58] **Maximilien Servajean, Alexis Joly, Dennis Shasha, Julien Champ, and Esther Pacitti. "Crowdsourcing Thousands of Specialized Labels: A Bayesian Active Training Approach". In: *IEEE Transactions on Multimedia* 19.6 (June 2017), pp. 1376–1391. DOI: 10.1109/TMM.2017.2653763. URL: https://hal.archives-ouvertes.fr/hal-01629149.**

[my63] Pierre Bonnet, Alexis Joly, Hervé Goëau, Julien Champ, Christel Vignau, Jean-François Molino, Daniel Barthélémy, and Nozha Boujemaa. "Plant identification: Man vs. Machine". In: *Multimedia Tools and Applications*. LifeCLEF 2014 plant identification challenge 75.3 (2016), pp. 1647–1665. DOI: 10.1007/s11042-015-2607-4. URL: https://hal.inria.fr/hal-01182778.

[my65] **Julien Champ, Titouan Lorieul, Pierre Bonnet, Najate Maghnaoui, Christophe Sereno, Thierry Dessup, Jean-Michel Boursiquot, Laurent AUDEGUIN, Thierry Lacombe, and Alexis Joly. "Categorizing plant images at the variety level: Did you say fine-grained?" In: *Pattern Recognition Letters* 81 (Oct. 2016). UMR AGAP - équipe DAAV - Diversité, adaptation et amélioration de la vigne, pp. 71–79. DOI: 10.1016/j.patrec.2016.05.022. URL: https://hal-lirmm.ccsd.cnrs.fr/lirmm-01348914.**

[my70] Alexis Joly, Pierre Bonnet, Hervé Goëau, Julien Barbe, Souheil Selmi, Julien Champ, Samuel Dufour-Kowalski, Antoine Affouard, Jennifer Carré, Jean-François Molino, Nozha Boujemaa, and Daniel Barthélémy. "A look inside the Pl@ntNet experience". In: *Multimedia Systems* 22.6 (2016), pp. 751–766. DOI: 10.1007/s00530-015-0462-9. URL: https://hal.inria.fr/hal-01182775.

[my78] Stefanos Vrochidis, Kostas D. Karatzas, Ari Karppinen, and Alexis Joly. "Guest Editorial: Environmental Multimedia Retrieval". In: *Multimedia Tools and Applications* 75.3 (Feb. 2016), pp. 1557–1562. DOI: 10.1007/s11042-016-3256-y. URL: https://hal.archives-ouvertes.fr/hal-01373782.

[my99] Amel Hamzaoui, Pierre Letessier, Alexis Joly, Olivier Buisson, and Nozha Boujemaa. "Object-based visual query suggestion". In: *Multimedia Tools and Applications*. Multimedia Tools and Applications 68.2 (Jan. 2014), pp. 429–454. DOI: 10.1007/s11042-012-1340-5. URL: https://hal.inria.fr/hal-00806635.

[my103] **Alexis Joly, Hervé Goëau, Pierre Bonnet, Vera Bakić, Julien Barbe, Souheil Selmi, Itheri Yahiaoui, Jennifer Carré, Elise Mouysset, Jean-François Molino, Nozha Boujemaa, and Daniel Barthélémy. "Interactive plant identification based on social image data". In: *Ecological Informatics* 23 (2014), pp. 22–34. DOI: 10.1016/j.ecoinf.2013.07.006. URL: https://hal.archives-ouvertes.fr/hal-00908872.**

[my126] **Ahmed Rebai, Alexis Joly, and Nozha Boujemaa. "BLasso for object categorization and retrieval: Towards interpretable visual models". In: *Pattern Recognition* 45.6 (June 2012), pp. 2377–2389. DOI: 10.1016/j.patcog.2011.11.022. URL: https://hal.inria.fr/hal-00739706.**

[my134] Amel Hamzaoui, Alexis Joly, and Nozha Boujemaa. "Multi-source shared nearest neighbours for multi-modal image clustering". In: *Multimedia Tools and Applications* 51.2 (2011), pp. 479–503.

[my146] Alexis Joly, Olivier Buisson, and Carl Frélicot. "Content-based Copy Retrieval using Distortion-based Probabilistic Similarity Search". In: *IEEE Transactions on Multimedia* (June 2008). URL: https://hal.inria.fr/hal-02420864.

[my148] **A. Joly, O. Buisson, and C. Frelicot. "Content-Based Copy Retrieval Using Distortion-Based Probabilistic Similarity Search". In: *IEEE Transactions on Multimedia* 9.2 (Feb. 2007), pp. 293–306. ISSN: 1941-0077. DOI: 10.1109/TMM.2006.886278.**

[my149] Alexis Joly. "ImagEval task 1: Retrieving transformed images with new local descriptors based on dissociated dipoles". In: (2007).

4.2 Conférence internationales avec comité de lecture/*Reviewed international conferences*

→ The following publications are sorted by descending year of publication. Rank A conferences are highlighted in bold.

[my7] Alexis Joly, Hervé Goëau, Christophe Botella, Rafael Ruiz De Castaneda, Hervé Glotin, Elijah Cole, Julien Champ, Benjamin Deneu, Maximilien Servajean, Titouan Lorieul, Willem-Pier Vellinga, Fabien-Robert Stöter, Andrew Durso, Pierre Bonnet, and Henning Müller. "LifeCLEF 2020 Teaser: Biodiversity Identification and Prediction Challenges". In: *42nd European Conference on IR Research on Advances in Information Retrieval (ECIR)*. Vol. Lecture Notes in Computer Science. Advances in Information Retrieval. Proceedings, Part II 12036. Lisbon, Portugal, Apr. 2020, pp. 542–549. DOI: 10.1007/978-3-030-45442-5_70. URL: https://hal.inrae.fr/hal-02873670.

[my13] Christophe Botella, Alexis Joly, Pierre Bonnet, François Munoz, and Pascal Monestiez. "A new multi-species method to correct for sampling bias in presence only niche models, and its application to Pl@ ntNet citizen science data in France". In: *2019 ESA Annual Meeting (August 11–16)*. ESA. 2019.

[my14] Christophe Botella, Maximilien Servajean, Pierre Bonnet, and Alexis Joly. “Overview of GeoLifeCLEF 2019: plant species prediction using environment and animal occurrences”. In: *Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 2380. Lugano, Switzerland, Sept. 2019. URL: https://hal.archives-ouvertes.fr/hal-02190170.

[my16] Benjamin Deneu, Maximilien Servajean, Christophe Botella, and Alexis Joly. “Evaluation of Deep Species Distribution Models using Environment and Co-occurrences”. In: *CLEF 2019 - Conference and Labs of the Evaluation Forum*. Lugano, Switzerland, Sept. 2019. URL: https://hal.inria.fr/hal-02290310.

[my17] Hervé Goëau, Pierre Bonnet, and Alexis Joly. “Overview of LifeCLEF Plant Identification Task 2019: diving into Data Deficient Tropical Countries”. In: *CLEF 2019 - Conference and Labs of the Evaluation Forum*. Vol. 2380. Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum. Linda Cappellato and Nicola Ferro and David E. Losada and Henning Müller. Lugano, Switzerland: CEUR, Sept. 2019, pp. 1–13. URL: https://hal.umontpellier.fr/hal-02283184.

[my18] **Alexis Joly, Hervé Goëau, Christophe Botella, Stefan Kahl, Marion Poupard, Maximilien Servajean, Hervé Glotin, Pierre Bonnet, Willem-Pier Vellinga, Robert Planqué, Jan Schlüter, Fabien-Robert Stöter, and Henning Müller. “LifeCLEF 2019: Biodiversity Identification and Prediction Challenges”. In: *ECIR 2019 - 41st European Conference on IR Research*. Ed. by Leif Azzopardi, Benno Stein, Norbert Fuhr, Philipp Mayr, and Claudia Hauff. Vol. LNCS. 11438. Cologne, Germany, Apr. 2019, pp. 275–282. DOI: 10.1007/978-3-030-15719-7_37. URL: https://hal.umontpellier.fr/hal-02273257.**

[my19] Alexis Joly, Hervé Goëau, Christophe Botella, Stefan Kahl, Maximilien Servajean, Hervé Glotin, Pierre Bonnet, Robert Planqué, Fabian-Robert Stöter, Willem-Pier Vellinga, and Henning Müller. “Overview of LifeCLEF 2019: Identification of Amazonian Plants, South & North American Birds, and Niche Prediction”. In: *CLEF 2019 - Conference and Labs of the Evaluation Forum*. Ed. by Fabio Crestani, Martin Brascher, Jacques Savoy, Andreas Rauber, Henning Müller, David E. Losada, Gundula Heinatz Bürki, Gundula Heinatz Bürki, Linda Cappellato, and Nicola Ferro. Vol. LNCS. Experimental IR Meets Multilinguality, Multimodality, and Interaction 11696. Lugano, Switzerland, Sept. 2019, pp. 387–401. DOI: 10.1007/978-3-030-28577-7_29. URL: https://hal.umontpellier.fr/hal-02281455.

[my21] Stefan Kahl, Fabian-Robert Stöter, Hervé Goëau, Hervé Glotin, Robert Planque, Willem-Pier Vellinga, and Alexis Joly. “Overview of BirdCLEF 2019: Large-Scale Bird Recognition in Soundscapes”. In: *Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 2380. Cappellato, L. and Ferro, N. and Losada, D. E. and Müller, H. Lugano, Switzerland: CEUR, Sept. 2019, pp. 1–9. URL: https://hal.umontpellier.fr/hal-02345644.

[my24] Adán Mora-Fallas, Hervé HG Goëau, Susan Mazer, Natalie Love, Erick Mata-Montero, Pierre Bonnet, and Alexis AJ Joly. “Accelerating the Automated Detection, Counting and Measurements of Reproductive Organs in Herbarium Collections in the Era of Deep Learning”. In: vol. 3. Pensoft Publishers, 2019, e37341.

[my25] Mathilde Negri, Maximilien Servajean, Benjamin Deneu, and Alexis Joly. “Location-Based Plant Species Prediction Using A CNN Model Trained On Several Kingdoms - Best Method Of GeoLifeCLEF 2019 Challenge”. In: *Working Notes of CLEF 2019 - Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 2380. Lugano, Switzerland, Sept. 2019. URL: https://hal.archives-ouvertes.fr/hal-02392637.

[my30] Christophe Botella, Pierre Bonnet, François Munoz, Pascal P. Monestiez, and Alexis Joly. “Overview of GeoLifeCLEF 2018: location-based species recommendation”. In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshops Proceedings. 2125. Avignon, France: CEUR-WS, Sept. 2018. URL: https://hal.archives-ouvertes.fr/hal-01913238.

[my35] Benjamin Deneu, Maximilien Servajean, Christophe Botella, and Alexis Joly. “Location-based species recommendation using co-occurrences and environment-GeoLifeCLEF 2018 challenge”. In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 2125. Avignon, France, Sept. 2018. URL: https://hal.archives-ouvertes.fr/hal-01913241.

[my36] Hervé Goëau, Pierre Bonnet, and Alexis Joly. “Overview of ExpertLifeCLEF 2018: how far automated identification systems are from the best experts?” In: *CLEF - Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 2125. Avignon, France, Sept. 2018. URL: https://hal.archives-ouvertes.fr/hal-01913244.

[my37] Hervé Goëau, Stefan Kahl, Hervé Glotin, Robert Planqué, Willem-Pier Vellinga, and Alexis Joly. “Overview of BirdCLEF 2018: monospecies vs. soundscape bird identification”. In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshops Proceedings. 2125. Avignon, France, Sept. 2018. URL: https://hal.archives-ouvertes.fr/hal-02189229.

[my39] Alexis Joly, Hervé Goëau, Christophe Botella, Hervé Glotin, Pierre Bonnet, Willem-Pier Vellinga, Robert Planqué, and Henning Müller. "Overview of LifeCLEF 2018: A Large-Scale Evaluation of Species Identification and Recommendation Algorithms in the Era of AI". In: *CLEF: Cross-Language Evaluation Forum*. Vol. LNCS. Experimental IR Meets Multilinguality, Multimodality, and Interaction 11018. Avignon, France, Sept. 2018, pp. 247–266. DOI: 10.1007/978-3-319-98932-7_24. URL: https://hal.archives-ouvertes.fr/hal-01913231.

[my44] **Antoine Affouard, Hervé Goëau, Pierre Bonnet, Jean-Christophe Lombardo, and Alexis Joly. "Pl@ntNet app in the era of deep learning". In: *ICLR: International Conference on Learning Representations*. Toulon, France, Apr. 2017. URL: https://hal.archives-ouvertes.fr/hal-01629195.**

[my46] Pierre Bonnet, Alexis Joly, Hervé Goëau, Jean-Christophe Lombardo, Antoine Affouard, Sen Wang, Rémi Knaff, Jean-François Molino, and Daniel Barthélémy. "Going deeper in the automated identification of Herbarium specimens". In: *Botany - Botanical Crossroads*. Forth Worth, TX, United States, June 2017. URL: https://hal.archives-ouvertes.fr/hal-01837387.

[my50] Jose Carranza-Rojas, Alexis Joly, Pierre Bonnet, Hervé Goëau, and Erick Mata-Montero. "Automated Herbarium Specimen Identification using Deep Learning". In: *TDWG: Biodiversity Information Standards*. Vol. 1. Data Integration in a Big Data Universe: Associating Occurrences with Genes, Phenotypes, and Environments. Ottawa, Canada, Oct. 2017, e20302. DOI: 10.3897/tdwgproceedings.1.20302. URL: https://hal.archives-ouvertes.fr/hal-01629142.

[my52] Hervé Goëau, Pierre Bonnet, and Alexis Joly. "Plant identification based on noisy web data: the amazing performance of deep learning (LifeCLEF 2017)". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 1866. Dublin, Ireland, Sept. 2017. URL: https://hal.archives-ouvertes.fr/hal-01629183.

[my53] Hervé Goëau, Hervé Glotin, Willem-Pier Vellinga, Robert Planqué, and Alexis Joly. "LifeCLEF Bird Identification Task 2017". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 1866. Dublin, Ireland, Sept. 2017. URL: https://hal.archives-ouvertes.fr/hal-01629175.

[my55] **Alexis Joly, Pierre Bonnet, Antoine Affouard, Jean-Christophe Lombardo, and Hervé Goëau. "Pl@ntNet -My Business". In: *MM: Multimedia*. Mountain View, United States, Oct. 2017, pp. 1–11. URL: https://hal.inria.fr/hal-01638263.**

[my56] Alexis Joly, Hervé Goëau, Hervé Glotin, Concetto Spampinato, Pierre Bonnet, Willem-Pier Vellinga, Jean-Christophe Lombardo, Robert Planque, Simone Palazzo, and Henning Müller. "LifeCLEF 2017 Lab Overview: Multimedia Species Identification Challenges". In: *CLEF: Cross-Language Evaluation Forum*. Ed. by Gareth J.F. Jones, Séamus Lawless, Julio Gonzalo, Liadh Kelly, Lorraine Goeuriot, Thomas Mandl, Linda Cappellato, and Nicola Ferro. Vol. LNCS. Experimental IR Meets Multilinguality, Multimodality, and Interaction 10456. Dublin, Ireland: Springer, Sept. 2017, pp. 255–274. DOI: 10.1007/978-3-319-65813-1_24. URL: https://hal.archives-ouvertes.fr/hal-01629191.

[my62] Pierre Bonnet, Alexis Joly, Julien Champ, Hervé Goëau, Samuel Dufour-Kowalski, Antoine Affouard, Jean-Christophe Lombardo, Sen Wang, Jennifer Carré, Jean-François Molino, Nozha Boujemaa, and Daniel Barthélémy. "An image-based Plant identification platform for thousands of species". In: *ATBC: Association for Tropical Biology and Conservation*. Tropical Ecology and Society Reconciliating Conservation and Sustainable use of Biodiversity. Montpellier, France, June 2016. URL: https://hal.archives-ouvertes.fr/hal-01837355.

[my64] Julien Champ, Hervé Goëau, and Alexis Joly. "Floristic participation at LifeCLEF 2016 Plant Identification Task". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 1609. Évora, Portugal, Sept. 2016, pp. 450–458. URL: https://hal.archives-ouvertes.fr/hal-01373776.

[my67] Hervé Goëau, Pierre Bonnet, and Alexis Joly. "Plant Identification in an Open-world (LifeCLEF 2016)". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 1609. Évora, Portugal, Sept. 2016, pp. 428–439. URL: https://hal.archives-ouvertes.fr/hal-01373780.

[my68] Hervé Goëau, Hervé Glotin, Willem-Pier Vellinga, Robert Planqué, and Alexis Joly. "LifeCLEF Bird Identification Task 2016: The arrival of Deep learning". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 1609. Évora, Portugal, Sept. 2016, pp. 440–449. URL: https://hal.archives-ouvertes.fr/hal-01373779.

[my71] **Alexis Joly, Hervé Goëau, Julien Champ, Samuel Dufour-Kowalski, Henning Müller, and Pierre Bonnet. "Crowdsourcing Biodiversity Monitoring: How Sharing your Photo Stream can Sustain our Planet". In: *MM: Conference on Multimedia*. Amsterdam, Netherlands: ACM, Oct. 2016, pp. 958–967. DOI: 10.1145/2964284.2976762. URL: https://hal.inria.fr/hal-01373762.**

[my72] Alexis Joly, Hervé Goëau, Hervé Glotin, Concetto Spampinato, Pierre Bonnet, Willem-Pier Vellinga, Julien Champ, Robert Planqué, Simone Palazzo, and Henning Müller. "LifeCLEF 2016: Multimedia Life Species Identification Challenges". In: *CLEF: Cross-Language Evaluation Forum*. Ed. by Norbert Fuhr, Paulo Quaresma, Teresa Gonçalves, Birger Larsen, Krisztian Balog, Craig Macdonald, Linda Cappellato, and Nicola Ferro. Vol. LNCS. Experimental IR Meets Multilinguality, Multimodality, and Interaction 9822. Évora, Portugal: Springer, Sept. 2016, pp. 286–310. DOI: 10.1007/978-3-319-44564-9_26. URL: https://hal.archives-ouvertes.fr/hal-01373781.

[my73] Alexis Joly, Jean-Christophe Lombardo, Julien Champ, and Anjara Saloma. "Unsupervised Individual Whales Identification: Spot the Difference in the Ocean". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 1609. Évora, Portugal, Sept. 2016, pp. 469–480. URL: https://hal.archives-ouvertes.fr/hal-01373777.

[my74] **Valentin Leveau, Alexis Joly, Olivier Buisson, and Patrick Valduriez. "Spatially Localized Visual Dictionary Learning". In: *ICMR: International Conference on Multimedia Retrieval*. New York, United States: ACM, June 2016, pp. 367–370. DOI: 10.1145/2911996.2912070. URL: https://hal.archives-ouvertes.fr/hal-01373778.**

[my76] **Maximilien Servajean, Alexis Joly, Dennis Shasha, Julien Champ, and Esther Pacitti. "ThePlantGame: Actively Training Human Annotators for Domain-specific Crowdsourcing". In: *MM: Conference on Multimedia*. Amsterdam, Netherlands: ACM, Oct. 2016, pp. 720–721. DOI: 10.1145/2964284.2973820. URL: https://hal.inria.fr/hal-01373769.**

[my80] Julien Champ, Titouan Lorieul, Maximilien Servajean, and Alexis Joly. "A comparative study of fine-grained classification methods in the context of the LifeCLEF plant identification challenge 2015". In: *CLEF: Conference and Labs of the Evaluation Forum*. Ed. by CEUR-WS. Vol. CEUR Workshop Proceedings. 1391. Toulouse, France, Sept. 2015. URL: https://hal.inria.fr/hal-01182788.

[my81] Hervé Goëau, Pierre Bonnet, and Alexis Joly. "LifeCLEF Plant Identification Task 2015". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. Working Notes of CLEF 2015 - Conference and Labs of the Evaluation forum 1391. Toulouse, France, Sept. 2015. URL: https://hal.inria.fr/hal-01182795.

[my82] Hervé Goëau, Hervé Glotin, Willem-Pier Vellinga, Robert Planqué, Andreas Rauber, and Alexis Joly. "LifeCLEF Bird Identification Task 2015". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 1391. Toulouse, France, Sept. 2015. URL: https://hal.inria.fr/hal-01182796.

[my85] Alexis Joly, Hervé Goëau, Hervé Glotin, Concetto Spampinato, Pierre Bonnet, Willem-Pier Vellinga, Robert Planqué, Andreas Rauber, Simone Palazzo, Bob Fisher, and Henning Müller. "LifeCLEF 2015: Multimedia Life Species Identification Challenges". In: *CLEF: Cross-Language Evaluation Forum*. Vol. LNCS. Working Notes of CLEF 2015 - Conference and Labs of the Evaluation forum 9283. Université de Toulouse. FRA. Toulouse, France, Sept. 2015, pp. 462–483. DOI: 10.1007/978-3-319-24027-5_46. URL: https://hal.inria.fr/hal-01182782.

[my86] Alexis Joly, Valentin Leveau, Julien Champ, and Olivier Buisson. "Shared nearest neighbors match kernel for bird songs identification -LifeCLEF 2015 challenge". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 1391. Toulouse, France, Sept. 2015. URL: https://hal.inria.fr/hal-01182784.

[my87] **Pierre Letessier, Nicolas Hervé, Alexis Joly, Hakim Nabi, Mathieu Derval, and Olivier Buisson. "DigInPix: Visual Named-Entities Identification in Images and Videos". In: *ICMR: International Conference on Multimedia Retrieval*. Shanghai, China: ACM, June 2015, pp. 661–664. DOI: 10.1145/2671188.2749369. URL: https://hal.inria.fr/hal-01182780.**

[my88] **Valentin Leveau, Alexis Joly, Olivier Buisson, and Patrick Valduriez. "Kernelizing Spatially Consistent Visual Matches for Fine-Grained Classification". In: *ICMR: International Conference on Multimedia Retrieval*. Shangai, China: ACM, June 2015, pp. 155–162. DOI: 10.1145/2671188.2749328. URL: https://hal.inria.fr/hal-01145988.**

[my92] Daniel Barthélémy, Nozha Boujemaa, Jean-François Molino, Alexis Joly, Hervé Goëau, Vera Bakić, Souheil Selmi, Julien Champ, Jennifer Carre, Mathias Chouet, Aurélien Perronnet, Christelle Vignau, Samuel Dufour-Kowalski, Antoine Affouard, Julien Barbe, and Pierre Bonnet. "Pl@ntNet, une plate-forme innovante d'agrégation et partage d'observations botaniques". In: *International Conference 'Botanists of the Twenty-first Century'*. Ed. by Noëline R. Rakotoarisoa, Stephen Blackmore, and Bernard Riera. UNESCO. Paris, France, Sept. 2014, pp. 191–197. URL: https://hal.archives-ouvertes.fr/hal-01837361.

[my93] **Julien Champ, Alexis Joly, and Pierre Bonnet. "Fine-grained Visual Faceted Search". In: *MM: Conference on Multimedia*. Orlando, FL, United States: ACM, Nov. 2014, pp. 721–722. DOI: 10.1145/2647868.2654875. URL: https://hal.inria.fr/hal-01088822.**

[my95] **Hervé Goëau, Pierre Bonnet, Alexis Joly, Antoine Affouard, Vera Bakić, Julien Barbe, Samuel Dufour-Kowalski, Souheil Selmi, Yahiaoui Itheri, Christel Vignau, Daniel Barthélémy, and Nozha Boujemaa. "Pl@ntNet Mobile 2014: Android port and new features". In: *ICMR: International Conference on Multimedia Retrieval*. Glasgow, United Kingdom: ACM - Association for Computing Machinery, Apr. 2014, pp. 527–530. DOI: 10.1145/2578726.2582618. URL: https://hal.inria.fr/hal-01088819.**

[my96] Hervé Goëau, Hervé Glotin, Willem-Pier Vellinga, Robert Planqué, Andreas Rauber, and Alexis Joly. "LifeCLEF Bird Identification Task 2014". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. 1180. Sheffield, United Kingdom, Sept. 2014, pp. 585–597. URL: https://hal.inria.fr/hal-01088829.

[my97] Hervé Goëau, Alexis Joly, Pierre Bonnet, Souheil Selmi, Jean-François Molino, Daniel Barthélémy, and Nozha Boujemaa. "LifeCLEF Plant Identification Task 2014". In: *CLEF: Conference and Labs of the Evaluation Forum*. Ed. by L. Cappellato, N. Ferro, M. Halvey, and W. Kraaij. Vol. CEUR Workshop Proceedings. Working Notes for CLEF 2014 Conference 1180. University of Scheffield. GBR. Sheffield, United Kingdom, Sept. 2014, pp. 598–615. URL: https://hal.archives-ouvertes.fr/halsde-01064599.

[my98] Hervé Goëau, Alexis Joly, Itheri Yahiaoui, Vera Bakić, Anne Verroust-Blondet, Pierre Bonnet, Daniel Barthélémy, Nozha Boujemaa, and Jean-François Molino. "PlantNet Participation at LifeCLEF2014 Plant Identification Task". In: *CLEF: Conference and Labs of the Evaluation Forum*. Ed. by L. Cappellato, N. Ferro, M. Halvey, and W. Kraaij. Vol. CEUR Workshop Proceedings. Working Notes for CLEF 2014 Conference Proceedings 1180. Sheffield, United Kingdom, Sept. 2014, pp. 724–737. URL: https://hal.archives-ouvertes.fr/halsde-01064569.

[my102] Alexis Joly, Julien Champ, and Olivier Buisson. "Instance-based Bird Species Identification with Undiscriminant Features Pruning". In: *CLEF: Conference and Labs of the Evaluation Forum*. Vol. CEUR Workshop Proceedings. Information Access Evaluation. Multilinguality, Multimodality, and Interaction 1180. Sheffield, United Kingdom: Springer International Publishing, Sept. 2014, pp. 625–633. URL: https://hal.inria.fr/hal-01088798.

[my104] Alexis Joly, Hervé Goëau, Hervé Glotin, Concetto Spampinato, Pierre Bonnet, Willem-Pier Vellinga, Robert Planque, Andreas Rauber, Bob Fisher, and Henning Müller. "LifeCLEF 2014: Multimedia Life Species Identification Challenges". In: *CLEF: Cross-Language Evaluation Forum*. Vol. LNCS. Information Access Evaluation. Multilinguality, Multimodality, and Interaction 8685. Sheffield, United Kingdom: Springer International Publishing, Sept. 2014, pp. 229–249. DOI: 10.1007/978-3-319-11382-1_20. URL: https://hal.inria.fr/hal-01075770.

[my105] Alexis Joly, Henning Müller, Hervé Goëau, Hervé Glotin, Concetto Spampinato, Andreas Rauber, Pierre Bonnet, Willem-Pier Vellinga, and Bob Fisher. "LifeCLEF: Multimedia Life Species Identification". In: *EMR: Environmental Multimedia Retrieval*. Vol. CEUR Workshop Proceedings. Proceedings of the 1st International Workshop on Environnmental Multimedia Retrieval co-located with ACM International Conference on Multimedia Retrieval (ICMR 2014) 1222. EMR co-located with ACM International Conference on Multimedia Retrieval (ICMR). Glasgow, United Kingdom, Apr. 2014, pp. 7–13. URL: https://hal.inria.fr/hal-01088812.

[my106] **Valentin Leveau, Alexis Joly, Olivier Buisson, Pierre Letessier, and Patrick Valduriez. "Recognizing Thousands of Legal Entities through Instance-based Visual Classification". In: *MM: Conference on Multimedia*. Orlando, FL, United States: ACM, Nov. 2014, pp. 1029–1032. DOI: 10.1145/2647868.2655038. URL: https://hal.inria.fr/hal-01077508.**

[my110] Vera Bakić, Sofiène Mouine, Saloua Ouertani-Litayem, Anne Verroust-Blondet, Itheri Yahiaoui, Hervé Goëau, and Alexis Joly. "Inria's participation at ImageCLEF 2013 Plant Identification Task". In: *CLEF (Online Working Notes/Labs/Workshop) 2013*. Valencia, Spain, Sept. 2013. URL: https://hal.inria.fr/hal-00874330.

[my111] Caputo Barbara, Muller Henning, Thomee Bart, Villegas Mauricio, Roberto Paredes, David Zellhofer, Hervé Goëau, Alexis Joly, Pierre Bonnet, Jesus Martinez Gomez, Ismael Garcia Varea, and Miguel Cazorla. "Imageclef 2013: the vision, the data and the open challenges". In: *CLEF 2013 - 4th Conference and Labs of the Evaluation Forum : Information Access Evaluation meets Multilinguality, Multimodality, and Visualization*. Vol. 8138. CLEF 2013: Information Access Evaluation. Multilinguality, Multimodality, and Visualization. Valencia, Spain: Springer, Sept. 2013, pp. 250–268. DOI: 10.1007/978-3-642-40802-1_26. URL: https://hal.inria.fr/hal-00908881.

[my112] **Hervé Goëau, Pierre Bonnet, Alexis Joly, Vera Bakić, Julien Barbe, Souheil Selmi, Jennifer Carré, Daniel Barthélémy, Nozha Boujemaa, Jean-François Molino, Grégoire Duché, and Aurélien Perronet. "Pl@ntNet Mobile App". In: *MM: Multimedia*. Barcelone, Spain: ACM, Oct. 2013, pp. 423–424. DOI: 10.1145/2502081.2502251. URL: https://hal.inria.fr/hal-00908910.**

[my113] Hervé Goëau, Pierre Bonnet, Alexis Joly, Vera Bakić, Daniel Barthélémy, Nozha Boujemaa, and Jean-François Molino. "The ImageCLEF 2013 Plant Identification Task". In: *CLEF*. Valencia, Spain, Sept. 2013. URL: https://hal.inria.fr/hal-00960929.

[my114] **Nicolas Hervé, Marie-Luce Viaud, Jérôme Thievre, Agnès Saulnier, Pierre Letessier, Julien Champ, Olivier Buisson, and Alexis Joly. "OTmedia: The French Transmedia News Observatory". In: *ACM Multimedia*. Barcelone, Spain, Oct. 2013, pp. 441–442. DOI: 10.1145/2502081.2502260. URL: https://hal.inria.fr/hal-00908925.**

[my115] Alexis Joly, Hervé Goëau, Pierre Bonnet, Vera Bakić, Jean-François Molino, Daniel Barthélémy, and Nozha Boujemaa. "The Imageclef Plant Identification Task 2013". In: *MAED: Multimedia Analysis for Ecological Data*. Barcelone, Spain, Oct. 2013, pp. 23–28. DOI: 10.1145/2509896.2509902. URL: https://hal.inria.fr/hal-00908934.

[my117] **Pierre Letessier, Nicolas Hervé, Champ Julien, Alexis Joly, Olivier Buisson, and Amel Hamzaoui. "Small objects query suggestion in a large web-image collection". In: *ACM Multimedia*. Barcelone, Spain: ACM, Oct. 2013. DOI: 10.1145/2502081.2502248. URL: https://hal.inria.fr/hal-00908891.**

[my118] Vera Bakić, Itheri Yahiaoui, Sofiène Mouine, Saloua Litayem Ouertani, Wajih Ouertani, Anne Verroust-Blondet, Hervé Goëau, and Alexis Joly. "Inria IMEDIA2's Participation at ImageCLEF 2012 Plant Identification Task". In: *CLEF (Online Working Notes/Labs/Workshop) 2012*. Rome, Italy, Sept. 2012. URL: https://hal.inria.fr/hal-00744901.

[my119] Hervé Goëau, Pierre Bonnet, Julien Barbe, Vera Bakić, Alexis Joly, and Jean-François Molino. "Multi-Organ Plant Identification". In: *ACM International Workshop on Multimedia Analysis for Ecological Data*. Nara, Japan, Oct. 2012, pp. 41–44. URL: https://hal.inria.fr/hal-00739724.

[my120] Hervé Goëau, Pierre Bonnet, Alexis Joly, Itheri Yahiaoui, Daniel Barthélémy, Nozha Boujemaa, and Jean-François Molino. "The ImageCLEF 2012 Plant Identification Task". In: *CLEF'2012: Conference and Labs of the Evaluation Forum*. Rome, Italy, Sept. 2012. URL: https://hal.inria.fr/hal-00960918.

[my121] Amel Hamzaoui, Alexis Joly, and Boujemaa Nozha. "Plant Leaves Morphological Categorization with Shared Nearest Neighbours Clustering". In: *ACM International Workshop on Multimedia Analysis for Ecological Data*. Nara, Japan, Oct. 2012, pp. 37–40. URL: https://hal.inria.fr/hal-00739715.

[my123] **Alexis Joly, Julien Champ, Pierre Letessier, Nicolas Hervé, Olivier Buisson, and Marie-Luce Viaud. "Visual-Based Transmedia Events Detection". In: *MM '12: 20th International conference on Multimedia*. MM '12 Proceedings of the 20th ACM international conference on Multimedia. ACM. Nara, Japan: ACM, Nov. 2012, pp. 1351–1352. DOI: 10.1145/2393347.2396480. URL: https://hal.inria.fr/hal-00755696.**

[my124] **Pierre Letessier, Alexis Joly, and Olivier Buisson. "Scalable Mining of Small Visual Objects". In: *MM'12: 20th ACM international conference on Multimedia*. Nara, Japan, Oct. 2012, pp. 599–608. URL: https://hal.inria.fr/hal-00739735.**

[my125] **Saloua Litayem Ouertani, Alexis Joly, and Nozha Boujemaa. "Hash-Based Support Vector Machines Approximation for Large Scale Prediction". In: *BMVC: British Machine Vision Conference*. Surrey, United Kingdom: BMVA Press, Sept. 2012, pp. 86.1–86.11. URL: https://hal.inria.fr/hal-00733912.**

[my127] Riadh Mohamed Trad, Alexis Joly, and Nozha Boujemaa. "Distributed approximate KNN Graph construction for high dimensional Data". In: *BDA: Bases de Données Avancées*. Clermont-Ferrand, France, Oct. 2012. URL: https://hal.inria.fr/hal-00756624.

[my128] **Riadh Mohamed Trad, Alexis Joly, and Boujemaa Nozha. "Distributed KNN-graph approximation via hashing". In: *ICMR - International Conference on Multimedia Retrieval - 2012*. Proceedings of the 2nd ACM International Conference on Multimedia Retrieval. Hong-Kong, Hong Kong SAR China: ACM, June 2012, 43:1–43:8. DOI: 10.1145/2324796.2324847. URL: https://hal.inria.fr/hal-00739713.**

[my129] Daniel Barthélémy, Nozha Boujemaa, Daniel Mathieu, Jean-François Molino, Alexis Joly, Hervé Goëau, Pierre Bonnet, Elise Mouysset, Philippe Birnbaum, Violette Roche, Wajih Ouertani, and Pierre Grard. "The Pl@ntNet project: Plant Computational Identification and Collaborative Information System". In: *Biosystematics 2011 - Digital identification methods in the digital age - morphological and molecular, computer-aided and fully automated, research and citizen science*. Berlin, Germany, Feb. 2011. URL: https://hal.inrae.fr/hal-02808663.

[my130] Daniel Barthélémy, Nozha Boujemaa, Daniel Mathieu, Jean-François Molino, Alexis Joly, Elise Mouysset, Philippe Birnbaum, Hervé Goëau, Pierre Bonnet, and Violette Roche. "The Pl@ntnet project: plant computational identification and collaborative information system". In: *IBC 2011 - XVIII International Botanical congress*. Melbourne, Australia, July 2011. URL: https://hal.inrae.fr/hal-02810776.

[my131] Hervé Goëau, Pierre Bonnet, Alexis Joly, Nozha Boujemaa, Daniel Barthélémy, Jean-François Molino, Philippe Birnbaum, Elise Mouysset, and Marie Picard. "The ImageCLEF 2011 plant images classification task". In: *ImageCLEF 2011*. Amsterdam, Netherlands, Sept. 2011. URL: https://hal.inria.fr/hal-00642197.

[my132] **Hervé Goëau, Alexis Joly, Souheil Selmi, Pierre Bonnet, Elise Mouysset, and Laurent Joyeux. "Visual-based plant species identification from crowdsourced data". In: *MM'11 - ACM Multimedia 2011*. Scottsdale, United States: ACM, Nov. 2011, pp. 0-0. DOI: 10.1145/2072298.2072472. URL: https://hal.inria.fr/hal-00642236.**

[my133] Hervé Goëau, Alexis Joly, Itheri Yahiaoui, Pierre Bonnet, and Elise Mouysset. “Participation of INRIA & Pl@ntNet to ImageCLEF 2011 plant images classification task”. In: *ImageCLEF 2011*. Amsterdam, Netherlands, Sept. 2011, pp. 0-0. URL: https://hal.inria.fr/hal-00642239.

[my135] **Alexis Joly and Olivier Buisson. “Random Maximum Margin Hashing”. In: *CVPR’11 - IEEE Computer Vision and Pattern Recognition*. Colorado springs, United States: IEEE, June 2011, pp. 873–880. DOI: 10.1109/CVPR.2011.5995709. URL: https://hal.inria.fr/hal-00642178.**

[my136] **Pierre Letessier, Olivier Buisson, and Alexis Joly. “Consistent Visual Words Mining with Adaptive Sampling”. In: *ICMR’11 - International Conference on Multimedia Retrieval*. Trento, Italy: ACM, Apr. 2011, pp. 0-0. DOI: 10.1145/1991996.1992045. URL: https://hal.inria.fr/hal-00642202.**

[my137] Jean-François Molino, Daniel Barthélémy, Nozha Boujemaa, Daniel Mathieu, Alexis Joly, and Elise Mouysset. “Access to botanical data through multimedia and collaborative retrieval”. In: *Conference on Multilingual and Multimodal Information Access Evaluation (CLEF)*. Amsterdam, Netherlands, Sept. 2011. URL: https://hal.inrae.fr/hal-02804437.

[my138] **Ahmed Rebai, Alexis Joly, and Nozha Boujemaa. “Interpretable Visual Models for Human Perception-Based Object Retrieval”. In: *ICMR’11 - First ACM International Conference on Multimedia Retrieval*. Trento, Italy: ACM, Apr. 2011, 21:1–21:8. DOI: 10.1145/1991996.1992017. URL: https://hal.inria.fr/hal-00642232.**

[my139] **Riadh Mohamed Trad, Alexis Joly, and Nozha Boujemaa. “Large scale visual-based event matching”. In: *ICMR’11 - International Conference on Multimedia Retrieval*. Trento, Italy: ACM, Apr. 2011, 53:1–53:7. DOI: 10.1145/1991996.1992049. URL: https://hal.inria.fr/hal-00642210.**

[my142] **Alexis Joly and Olivier Buisson. “Logo retrieval with a contrario visual query expansion”. In: *Proceedings of the 17th ACM international conference on Multimedia*. ACM. 2009, pp. 581–584.**

[my143] **Saloua Litayem Ouertani, Alexis Joly, and Nozha Boujemaa. “Interactive Objects Retrieval with Efficient Boosting”. In: *MM’09 - Proceedings of the 17th ACM international conference on Multimedia*. Beijing, China: ACM, Oct. 2009, pp. 545–548. DOI: 10.1145/1631272.1631352. URL: https://hal.inria.fr/hal-00724876.**

[my145] **Alexis Joly and Olivier Buisson. “A posteriori multi-probe locality sensitive hashing”. In: *Proceedings of the 16th ACM international conference on Multimedia*. ACM. 2008, pp. 209–218.**

[my150] **Alexis Joly. “New local descriptors based on dissociated dipoles”. In: *Proceedings of the 6th ACM international conference on Image and video retrieval*. ACM. 2007, pp. 573–580.**

[my152] **Julien Law-To, Li Chen, Alexis Joly, Ivan Laptev, Olivier Buisson, Valerie Gouet-Brunet, Nozha Boujemaa, and Fred Stentiford. “Video copy detection: a comparative study”. In: *Proceedings of the 6th ACM international conference on Image and video retrieval*. ACM. 2007, pp. 371–378.**

[my153] Julien Law-To, LI Chen, Alexis Joly, Ivan Laptev, Olivier Buisson, Valérie Gouet-Brunet, Nozha Boujemaa, and Fred Stentiford. “Video Copy Detection: a Comparative Study”. In: *CIVR 2007*. Amsterdam, France, July 2007. URL: https://hal.inria.fr/hal-02420846.

[my154] Julien Law-To, Li Chen, Alexis Joly, Ivan Laptev, Olivier Buisson, Valérie Gouet-Brunet, Nozha Boujemaa, and Fred Stentiford. “Video copy detection”. In: *ACM International Conference on Image and Video Retrieval (CIVR’07)*. Amsterdam, Netherlands: ACM Press, July 2007, pp. 371–378. DOI: 10.1145/1282280.1282336. URL: https://hal.archives-ouvertes.fr/hal-02338328.

[my156] **Ahmed Rebai, Alexis Joly, and Nozha Boujemaa. “Interpretability based interest points detection”. In: *Proceedings of the 6th ACM international conference on Image and video retrieval*. ACM. 2007, pp. 33–40.**

[my162] Alexis Joly and Olivier Buisson. “Discriminant local features selection using efficient density estimation in a large database”. In: *Proceedings of the 7th ACM SIGMM international workshop on Multimedia information retrieval*. ACM. 2005, pp. 201–208.

[my163] Alexis Joly, Olivier Buisson, and Carl Frelicot. “Statistical similarity search applied to content-based video copy detection”. In: *21st International Conference on Data Engineering Workshops (ICDEW’05)*. IEEE. 2005, pp. 1285–1285.

[my164] Alexis Joly, Carl Frélicot, and Olivier Buisson. “Content-based video copy detection in large databases: A local fingerprints statistical similarity search approach”. In: *IEEE International Conference on Image Processing 2005*. Vol. 1. IEEE. 2005, pp. I–505.

[my166] Alexis Joly, Carl Frélicot, and Olivier Buisson. “Feature statistical retrieval applied to content based copy identification”. In: *2004 International Conference on Image Processing, 2004. ICIP’04*. Vol. 1. IEEE. 2004, pp. 681–684.

[my167] **Alexis Joly, Carl Frélicot, and Olivier Buisson. "Robust content-based video copy identification in a large reference database". In: *International Conference on Image and Video Retrieval*. Springer. 2003, pp. 414–424.**

4.3 Livres et chapitres de livre/*Books and book chapters*

[my8] Alexis Joly, Hervé Goëau, Stefan Kahl, Benjamin Deneu, Willem-Pier Vellinga, Maximilien Servajean, Elijah Cole, Lukáš Picek, Rafael Ruiz de Castañeda, Isabelle Bolon, Andrew Durso, Titouan Lorieul, Christophe Botella, Hervé Glotin, Julien Champ, Ivan Eggel, Pierre Bonnet, and Henning Müller. "Overview of LifeCLEF 2020: A System-Oriented Evaluation of Automated Species Identification and Species Distribution Prediction". In: *Experimental IR Meets Multilinguality, Multimodality, and Interaction. 11th International Conference of the CLEF Association, CLEF 2020, Thessaloniki, Greece, September 22–25, 2020, Proceedings*. Vol. 12260. Lecture Notes in Computer Science. Springer, Sept. 2020, pp. 342–363. DOI: 10.1007/978-3-030-58219-7_23. URL: https://hal.inrae.fr/hal-02945382.

[my20] Alexis Joly, Hervé Goëau, Hervé Glotin, Concetto Spampinato, Pierre Bonnet, Willem-Pier Vellinga, Jean-Christophe Lombardo, Robert Planqué, Simone Palazzo, and Henning Müller. "Biodiversity Information Retrieval Through Large Scale Content-Based Identification: A Long-Term Evaluation". In: *Information Retrieval Evaluation in a Changing World: Lessons Learned from 20 Years of CLEF*. Ed. by Nicola Ferro and Carol Peters. Vol. 41. The Information Retrieval Series. AcknowledgementsThe organization of the PlantCLEF task is supported by the French project Floris'Tic (Tela Botanica, INRIA, CIRAD, INRA, IRD) funded in the context of the national investment program PIA. The organization of the BirdCLEF task is supported by the Xeno-Canto foundation for nature sounds as well as the French CNRS project SABIOD.ORG and EADM MADICS, and Floris'Tic. The annotations of some soundscapes were prepared with the late wonderful Lucio Pando at Explorama Lodges, with the support of Pam Bucur, Marie Trone and H. Glotin. The organization of the SeaCLEF task is supported by the Ceta-mada NGO and the French project Floris'Tic. 2019, pp. 389–413. DOI: 10.1007/978-3-030-22948-1_16. URL: https://hal.umontpellier.fr/hal-02273280.

[my29] Pierre Bonnet, Hervé Goëau, Siang Thye Hang, Mario Lasseck, Milaň Sulc, Valéry V. Malécot, Philippe Jauzein, Jean-Claude Melet, Christian You, and Alexis Joly. "Plant Identification: Experts vs. Machines in the Era of Deep Learning". In: *Multimedia Tools and Applications for Environmental & Biodiversity Informatics*. Vol. Chapter 8. Multimedia Systems and Applications Series. Editions Springer, June 2018, pp. 131–149. DOI: 10.1007/978-3-319-76445-0_8. URL: https://hal.archives-ouvertes.fr/hal-01913277.

[my32] Christophe Botella, Alexis Joly, Pierre Bonnet, Pascal Monestiez, and François Munoz. "A deep learning approach to Species Distribution Modelling". In: *Multimedia Tools and Applications for Environmental & Biodiversity Informatics*. Ed. by Alexis Joly, Stefanos Vrochidis, Kostas Karatzas, Ari Karppinen, and Pierre Bonnet. Multimedia Systems and Applications Series. Springer, 2018, pp. 169–199. DOI: 10.1007/978-3-319-76445-0_10. URL: https://hal.archives-ouvertes.fr/hal-01834227.

[my34] Jose Carranza-Rojas, Alexis Joly, Hervé Goëau, Erick Mata-Montero, and Pierre Bonnet. "Automated identification of herbarium specimens at different taxonomic levels". In: *Multimedia Tools and Applications for Environmental & Biodiversity Informatics*. Vol. Multimedia Systems and Applications. Multimedia Systems and Applications Series. Editions Springer, June 2018, pp. 151–167. DOI: 10.1007/978-3-319-76445-0_9. URL: https://hal.archives-ouvertes.fr/hal-01913272.

[my38] Alexis Joly, Pierre Bonnet, Stefanos Vrochidis, Kostas Karatzas, and Ari Karppinen. "Multimedia Tools and Applications for Environmental & Biodiversity Informatics Introduction". In: *Multimedia Tools and Applications for Environmental & Biodiversity Informatics*. Multimedia Systems and Applications Series. Editions Springer, 2018, p. 7. DOI: 10.1007/978-3-319-76445-0_1. URL: https://hal.inrae.fr/hal-02791676.

[my40] Alexis Joly, Stefanos Vrochidis, Kostas Karatzas, Ari Karppinen, and Pierre Bonnet. *Multimedia Tools and Applications for Environmental & Biodiversity Informatics*. OUV. Springer International Publishing, 2018. DOI: 10.1007/978-3-319-76445-0. URL: https://hal-lirmm.ccsd.cnrs.fr/lirmm-01959343.

4.4 Workshop internationaux avec comité de lecture/*Revised international workshop proceedings*

[my12] Olivier Beaumont, Lionel Eyraud-Dubois, Julien Herrmann, Alexis Joly, and Alena Shilova. *Optimal checkpointing for heterogeneous chains: how to train deep neural networks with limited memory*. Research Report RR-9302. Inria Bordeaux Sud-Ouest, Nov. 2019. URL: https://hal.inria.fr/hal-02352969.

[my57] Valentin Leveau and Alexis Joly. *Adversarial autoencoders for novelty detection*. Research Report. Inria - Sophia Antipolis, Feb. 2017. URL: https://hal.inria.fr/hal-01636617.

[my116] Pierre Letessier, Olivier Buisson, and Alexis Joly. *Scalable Mining of Small Visual Objects (with new experiments)*. Research Report. Lirmm, Dec. 2013. URL: https://hal.inria.fr/hal-00912560.

[my122] Alexis Joly. *Recommendations about benchmarking campaigns as a tool to foster multimedia search technology transfer at the European level*. Technical Report. Inria, July 2012. URL: https://hal.inria.fr/hal-02428503.

[my141] Amel Hamzaoui, Alexis Joly, and Nozha Boujemaa. *Multi-source shared nearest neighbours for multi-modal image clustering*. Research Report RR-7351. INRIA, June 2010. URL: https://hal.inria.fr/inria-00496170.

[my147] Alexis Joly, Julien Law-To, and Nozha Boujemaa. *INRIA-IMEDIA TRECVID 2008: Video Copy Detection*. Tech. rep. 2008.

[my160] Alexis Joly, Ahmed Rebai, and Boujemaa Nozha. *Constant Tangential Angle Elected Interest Points*. Research Report. 2006, p. 22. URL: https://hal.inria.fr/inria-00085764.

4.5 Brevets/*Patents*

[my144] Alexis Joly. *Search of similar features representing objects in a large reference database*. US Patent 7,400,784. 2008.

4.6 Revues nationales/*National journals*

[my101] Alexis Joly, Pierre Bonnet, and Hervé Goëau. “Pl@ntNet, un réseau et des outils pour une recherche participative”. In: *Interstices* (Mar. 2014). URL: https://hal.inria.fr/hal-01350266.

4.7 Conférence nationales avec comité de lecture/*Reviewed national conferences*

[my22] Titouan Lorieul and Alexis Joly. “Vers un désenchevêtrement de l’ambiguïté de la tâche et de l’incertitude du modèle pour la classification avec option de rejet à l’aide de réseaux neuronaux”. In: *CAp 2019 - Conférence sur l’Apprentissage automatique*. Toulouse, France, July 2019. URL: https://hal.archives-ouvertes.fr/hal-02421210.

[my108] Maximilien Servajean, Esther Pacitti, Miguel Liroz-Gistau, Alexis Joly, and Julien Champ. “PlantRT : a Distributed Recommendation Tool for Citizen Science”. In: *BDA: Gestion de Données - Principes, Technologies et Applications*. Autrans, France, Oct. 2014, pp. 48–50. URL: https://hal.archives-ouvertes.fr/hal-01163819.