Using Pure Pollen Species When Training a CNN to Segment Pollen Mixtures

Nana Yang[1,*], Victor Joos[1], A-L Jacquemart[2], Christel Buyens[2], C.De Vleeschouwer[1]
[1]ICTEAM Institute, UCLouvain, Belgium
[2]ELI Institute, UCLouvain, Belgium
{nana.yang, christophe.devleeschouwer}@uclouvain.be

Abstract

Recognizing the types of pollen grains and estimating their proportion in pollen mixture samples collected in a specific geographical area is important for agricultural, medical, and ecosystem research. Our paper adopts a convolutional neural network for the automatic segmentation of pollen species in microscopy images, and proposes an original strategy to train such network at reasonable manual annotation cost. Our approach is founded on a large dataset composed of pure pollen images. It first (semi-)manually segments foreground, i.e. pollen grains, and background in a fraction of those images, and use the resulting annotated dataset to train a universal pollen segmentation CNN. In the second step, this model is used to automatically segment a large number of additional pure pollen images, so as to supervise the training of a pollen species segmentation model. Despite the fact that it has been trained from pure images only, the model is shown to provide accurate segmentation of species in pollen mixtures. Our experiments also demonstrate that dedicating a model to the segmentation of a subset of the available pure pollen species makes it possible to train a bin pollen class, corresponding to pollen species that are not in the subset of species recognized by the model. This strategy is useful to cope with unexpected species in a mixture.

1. Introduction

Sexual reproduction in flowering plants involves the production, transfer and deposition of pollen grains to fertilise the ovules and thus enable the formation of seeds. The morphology of pollen grains depends on the pollination mode and the identity of the plant species. For example, the pollen of wind-pollinated trees, such as pines, has two air sacs that allow it to be dispersed over long distances. The identification of pollen types allows many applications such as the reconstruction of the history and evolution of ecosystems since the ice ages, including the influences of human populations through pollen trapped in peat accumulations (palaeoecology), alerts of allergic peaks by monitoring pollen counts in the air (allergology), pollen trapped in honey for certification of single-flower or controlled origin honey (melisso-palynology) or pollen transported and collected by insect pollinators as food resources (pollination ecology and pollinator health).

The traditional methods of recognizing pollen types use light microscopy on fresh pollen, acetolysed pollen, or scanning electron microscopy [10]. Recognition is based on the size, shape, thickness and ornamentation of the exine, the number and the shape of germination pits or pores. Such identification is time-consuming and requires considerable recognition expertise. New DNA barcoding techniques are being developed, allowing precise identification from DNA profile banks [6]. They are expensive and time-consuming. The development of techniques based on automatic recognition offers a remarkable opportunity and is highly valued by palynology researchers [24].

Features such as shape [19], texture [19], contour [24] have been used to segment pollen in early machine learning research. With the emergence of deep learning [12], convolutional neural networks [1] are now widely used due to its ability to learn end-to-end, to directly convert image pixels into predictions of interest. However, deep learning requires a large amount of labeled data [2], and the annotation of mixture pollen images is extremely difficult. Only a high-level expert can distinguish different pollen in an image. This is due to the relatively high similarity of pollen and to the fact that pollen particles are 3D objects, which might appear quite differently when observed from different angles. Hence, labeling mixed pollen images is a difficult, tedious, and time-consuming task.

In this paper, as shown in Fig. 1, a method is proposed to learn how to segment individual pollen species in pollen mixtures, based on a large dataset composed of pure pollen images. To save manual annotation time, our method first trains a universal i.e. species-agnostic, pollen segmentation CNN based on the (semi-)manual segmentation of pollen

*The author is funded by China Scholarship Council.

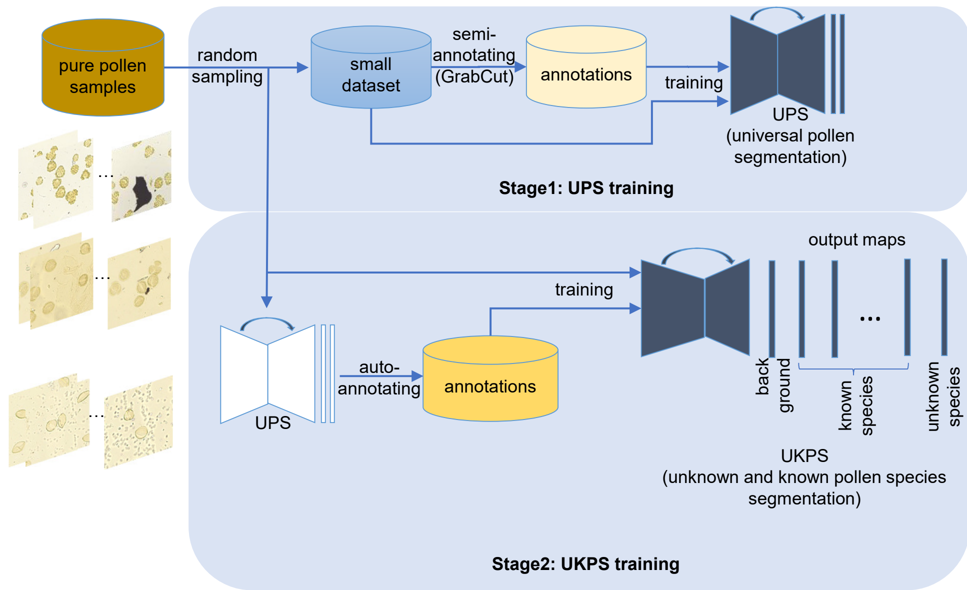

Figure 1. Overview of our method. In the first stage, we (semi-)manually annotate a small subset using the GrabCut algorithm and then train UPS model capable of distinguishing pollen grains and background. In the second stage, the UPS model is used to automatically annotate a large number of additional pure pollen images to train UKPS model, which can differentiate between known pollen species and unknown species.

grains in a small subset of images randomly selected in the dataset. This model is then used to automatically segment a large number of additional pure pollen images, so as to supervise the training of our pollen species segmentation CNN. Despite it is trained on pure images, our model is shown to accurately segment species in pollen mixtures. To handle scenarios for which all possible pollen species are not known/available a priori, we propose an alternative multi-model training strategy, in which each model is trained to recognize only a subset of the species available in the training set, while assigning an 'unknown' label to other species. This strategy makes it possible to segment and identify as 'unknown' a pollen grain that has not been seen during training. To the best of our knowledge, our work is the first one to formulate and address mixture pollen analysis as a segmentation problem. This is in contrast with most previous works, which have to isolate individual pollen grains, before recognizing them based on a classifier.

The rest of the paper is organized as follows. Section 2 surveys the related works. Section 3 describes our proposed method. Section 4 demonstrates the relevance of our approach on a dataset including more than one hundred pollen species.

2. Related Work

Current pollen classification methods are mainly based on manually extracted features and automatically learned features [3]. They generally consider images that correspond to a bounding box, including a single pollen grains. In [14] hand-crafted texture features extracted, to be processed based on Fisher Discriminant Analysis (FDA) [26], to finally classify the pollen with K-Nearest Neighbors (KNN) [27]. The authors in [2, 24] implemented pollen classification on single pollen grain dataset by investigating different contour and texture extraction methods. Since the manual selection of appropriate feature requires the combined expertise of biologists and computer vision scientists, it remains a tedious task, especially because it needs to be repeated each time a novel pollen species has to be handled. Therefore, deep learning has been considered to define features automatically, based on examples.

In [3, 7, 23], researchers proposed to use convolutional neural networks to classify pollen images, but in these methods, the image background in the dataset is uniform, and there is only one pollen grain in the images. The above method fails if there are multiple overlapping pollen grains or multiple pollen types in the image, or when there are pollen types in the image that are unknown to the model. To address this limitation, we propose to formulate the pollen recognition problem as a pixel-wise [4] segmentation problem. In this formulation, the CNN predicts a label in each image pixel. This label determines whether the pixel covers the background or a pollen grain and, in the latter, specifies to which pollen species the pixel corresponds. However, training such a CNN using conventional supervision methods, whilst possible and effective [11, 13], requires the

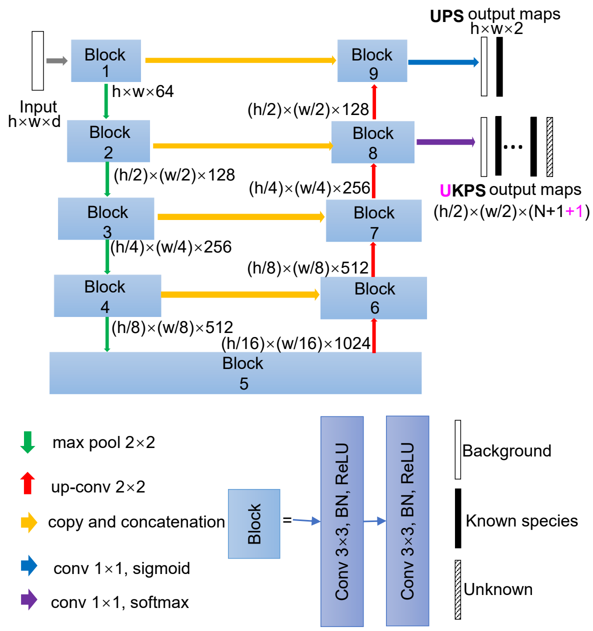

Figure 2. Architectures of our models. We have three kinds of models. They use the same architecture for encoder. UPS has 9 blocks, and KPS(without unknown species map compared with UKPS) and UKPS have 8 blocks. In the figure, h, w, d represent the height, width and channel of the input image respectively. N is the number of known pollen species by the model. Block1 - Block9 are convolution+batchnormalization+ReLU.

annotation of a large amount of images containing the various pollen species that have to be recognized by the model. Such annotation of pollen mixture images requires biological expertise, which is time-consuming and error-prone.

In order to solve the tedious task of labeling large dataset, we first (semi-)manually segment foreground pixels, i.e. the pollen grains, on a small set of images, and use the resulting annotations to train a Universal Pollen Segmentation (UPS) model. This model is then used to automatically segment a large number of additional pure pollen images to supervise a multi-class segmentation model that is designed to differentiate pollen species and, optionally, to identify pollen grains that correspond to species that have not been considered during training.

3. Methods

This section introduces the scheme proposed to quantify the amount of individual pollen species in a pollen mixture microscopy image. It relies on pixel-wise labelling of microscopy images, using convolutional neural networks (CNNs). As shown in Fig. 1, the method considers the training of a universal, species agnostic, pollen segmentation CNN model (named UPS in the following), to automatically annotate pure pollen species images. Those images and their annotations are then used to supervise the training of a model that is able to discriminate between species. Optionally, an additional output label is considered to indicate that a pixel corresponds to an *unknown* species of pollen, i.e. a pollen species that has not been considered during training. Depending whether this *unknown* species option is activated or not, we refer to those species segmentation models as UKPS (for Unknown and Known Pollen Species) or KPS, respectively.

3.1. Pollen Grain Segmentation

As depicted in Fig. 2, our universal pollen segmentation CNN consists in a 9 layer U-Net [20], followed by two output feature maps connected to a sigmoid to predict a binary value in each image pixel. Training this CNN requires the definition of pollen grain segmentation masks in a number of image samples. This foreground/background segmentation is done (semi-)manually using the GrabCut algorithm [21]. GrabCut takes as minimal input a bounding-box that is manually-defined around the foreground object of interest, and splits the bounding-box inner pixels in two connected regions that most differ in terms of their distribution. The annotation procedure is both effective and efficient since the foreground pollen grains are highly contrasted compared to the background, making GrabCut appropriate for this task. In practice, only a few hours have been required to label 6000 grains in 1000 images.

3.2. Pollen Species Segmentation

The CNN adopted to recognize and segment pollen species in pollen mixtures is depicted in Fig. 2. It adopts a similar U-Net backbone than the one adopted for pollen grain segmentation, but omits the last decoder layer to limit the resolution and thus the memory footprint associated to the output maps of the network. In practice, one output map is devoted to each of the C possible categories that the model is expected to recognize. A softmax operator is applied to the C activations outputted by the network in each pixel, so as to predict a vector with positive components that sum to one. Each of these vectors is close to one-hot vector that is supposed to encode the category of its corresponding pixel.

Two types of pollen species segmentation CNNs have been envisioned in our study, depending on how they define the output categories they aim at recognizing.

The first one, named *known pollen species* (KPS), is certainly the most natural. It assigns one specific output to each of the N pollen species represented in the training set, and considers an additional bin class category to identify the pixels that do not correspond to one of the known species (combining background pixels and unknown pollen grains). However, the experiments presented in Section 4 (particularly Fig. 5) reveal that such a model fails in labelling a pollen grain that is not part of the N known species properly, i.e. as bin class category.

The second strategy, named *unknown and known pollen species* (UKPS), aims at addressing this limitation. It proposes to partition the N pollen species into K disjoint subsets, and to learn a model for each subset, thereby leading to K models. Specifically, similar to KPS, the model associated to the ith subset assigns one output map to each of the species in the subset but, in contrast to KPS, assigns non-pollen background pixels and unknown pollen species to two distinct outputs. This is made possible since the pollen grains of $(N - N_i)$ species that are not in the ith subset can be used to supervise the training of the unknown pollen category.

4. Experiments and Discussion

In this section, we present experiments to demonstrate that our data annotation and training strategies are effective in learning convolutional neural networks that are able to segment known and unknown pollen species in pollen mixtures.

4.1. Implementation Details

Dataset. A total of 149 pure pollen files with inconsistent backgrounds as shown in Fig. 1, and 10 mixture pollen files were scanned and uploaded to Cytomine [15]. Each file corresponds to a large-scale image. Some of them have resolutions over 30000 by 30000. Each pure pollen file includes a single species of pollen. Those pure files include species from the Asteraceae [9], Brassicaceae [22], Fabaceae [17], Lamiaceae [18], Rosaceae [5] families (the number of pollen species in each family is 10, 21, 26, 23, and 12, respectively) and 67 species outside this set of families. Large-scale images implies higher computational and memory requirements during model training. Therefore, we crop a large-scale image into multiple 800 by 800 images. We split our 149 pure pollen species into known (131 species, used for training) and unknown (18 species, kept for testing, and never seen at training) species. A small set of ten 800x800 images was randomly selected among the known species to train the UPS network. To train and evaluate the KPS model, about 30 images have been randomly selected for training, 10 images for validation and 10 images for testing. When considering UKPS, we have considered K=5 subsets of known species, corresponding to the Asteraceae, Brassicaceae, Fabaceae, Lamiaceae, and Rosaceae. Each subset was used to learn the known species their respective model, and the species remaining out of the 131 species available for training were used to supervise the unknown pollen class, thereby resulting in 5 models, denoted by the name of their associated family.

Semi-manual annotation. We selected a small dataset containing 10 800x800 images in each of the 131 pure species pollen images. We used Grabcut to segment the pollen grains in those images in an effective manner. It took about 8 hours to annotate the whole set of images.

Data augmentation. The researchers pointed out in [25] that the more data the model can access, the more effective the model is. To further increase the amount of data that the model can access, we use random rotate90, flip, transpose, and brightness contrast as data augmentation methods during training.

Loss. The loss function adopted in this paper is defined in Eq. (3). It combines the cross-entropy and the Dice Loss.

$$L_{CE} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} y_i^c \cdot \log p_i^c \tag{1}$$

$$L_{Dice} = 1 - \frac{2\sum_{i=1}^{N}\sum_{c=1}^{C} y_i^c \cdot p_i^c}{\sum_{i=1}^{N}\sum_{c=1}^{C} y_i^{c2} + \sum_{i=1}^{N}\sum_{c=1}^{C} p_i^{c2}} \tag{2}$$

$$Loss = 0.9 \cdot L_{CE} + L_{Dice} \tag{3}$$

In Eqs. (1) and (2), y_i^c is a binary indicator (0 or 1) indicating whether pixel i lies in class c or not, and p_i^c is the softmax or sigmoid output of the network for pixel i and class c.

Metric. Mean intersection over union is a standard metric to assess segmentation models. It calculates the intersection over union ratio of ground-truth and predicted regions for each class, and then averages it over the classes. Formally, it is defined as:

$$\text{mIoU} = \frac{1}{C+1}\sum_{i=0}^{C}\frac{|G_i \cap P_i|}{|G_i \cup P_i|} \tag{4}$$

$$= \frac{1}{C+1}\sum_{i=0}^{C}\frac{TP_i}{FN_i + FP_i + TP_i} \tag{5}$$

with G_i denoting the ground truth, P_i the model prediction associated to class i, and C the number of classes. TP_i, FN_i and FP_i represent true positive, false negative and false positive detection of pixel belonging to class i.

Training. The model parameters are initialized using Kaiming normal [8]. We use SGD with 0.9 momentum and 0.0001 weight decay. UPS training uses a batch size of 2, and the initial learning rate is 0.01. The learning rate is

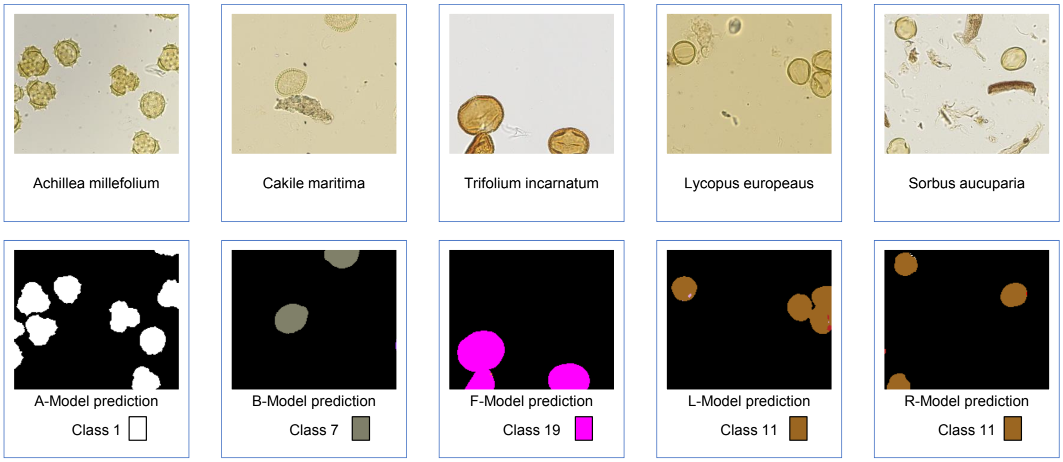

Figure 3. Predictions on pure pollen images of known species by our proposed UKPS models. The legend under each prediction defines the UKPS model that has been used (see text for details), and shows the color-label that corresponds to the species present in each image. We can see that while the models are mostly accurate, they have trouble in some cases to correctly segment the borders of pollen grains, given them a different species label (as can be seen by the wrong colors appearing in the last two images)

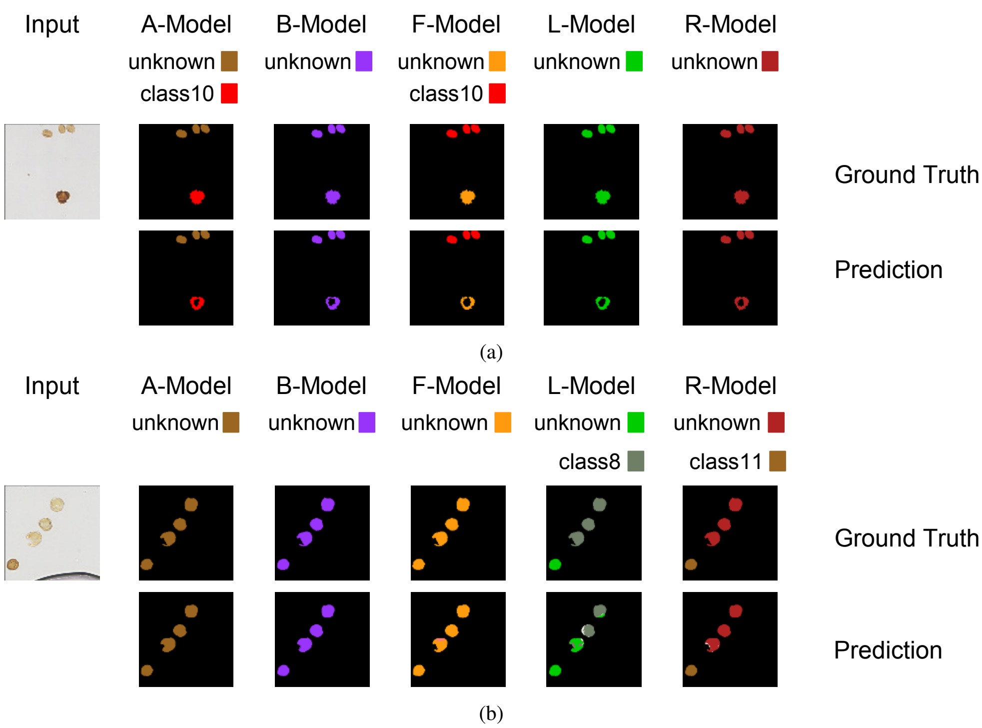

Figure 4. Visualization of the predictions for images of a mixture of pollen grains using multiple family-specific models based on UKPS. (a) There are two types of pollen: Tanacetum vulgare (class 10, in the Asteraceae family) and Lotus corniculatus (class 10, in the Fabaceae family). They are correctly classified as known by their specific models (A and F), while they are classified as unknown by the other models. (b) As in (a) the two types of pollen, Lamium galeobdolon (class 8, Lamiaceae) and Sorbus aucuparia (class 11, Rosaceae) are segmented as known species only by their respective models (the L-model and R-model).

reduced by a factor of 0.97 at each epoch, using an exponential learning rate scheduler. To train KPS and UKPS, to include more variety in each batch, we resize the images from 800×800 to 400×400 before network input. The batch size can therefore be increased to 16. The initial learning rate is set to 0.01. The exponential learning rate decay factor is set to 0.99. We use early stopping to stop training if the metric does not improve on the validation set

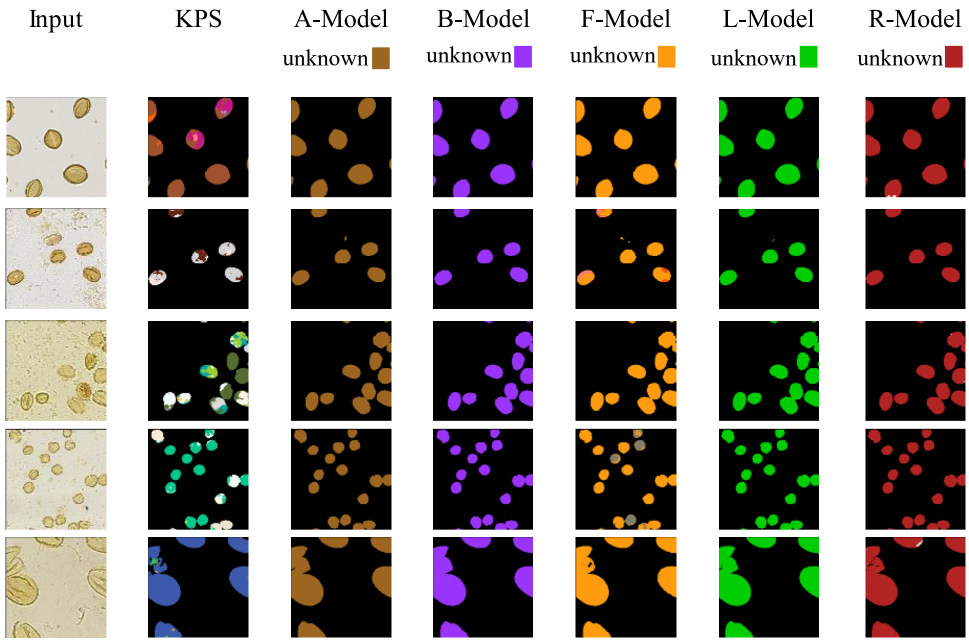

Figure 5. Predictions for unknown species using different models.

for 30 epochs.

Hardware. All models are trained with Pytorch [16] on Nvidia GeForce GTX 1080 Ti 11Gb GPUs.

4.2. Visual Assessment

A few representative examples of pollen species segmentation results are presented in Fig. 3. For convenience, here we name the family-specific UKPS models based on the first letter of the corresponding family name. Specifically, Asteraceae model, Brassicaceae model, Fabaceae model, Lamiaceae model, and Rosaceae model are referred to as A-Model, B-Model, F-Model, L-Model, and R-Model respectively). We observe that our proposed UKPS can accurately recognize known pollen species in pure images, even though they have trouble in some cases to correctly segment the borders of pollen grains(as can be seen in the last two images in Fig. 3).

More interestingly, Fig. 4 considers the segmentation of pollen mixtures. There, we observe that the proposed UKPS can successfully segment known pollen species seen by the model in the mixed pollen grain images, while only ever having been trained on pure pollen images. Fig. 4a shows the ground truth and corresponding prediction for the image containing two pollen species, Tanacetum vulgare (known by A-Model) and Lotus corniculatus species (known by F-Model). In Fig. 4b, the input image includes Lamium galeobdolon (known by L-Model) and Sorbus aucuparia grains (known by R-Model). Although the central region of Tanacetum vulgare grain in Fig. 4a is misclassified as background and some pixels of Lamium galeobdolon in Fig. 4b are segmented as unknown type, it is reasonable because our mixture pollen images and the pure pollen images used for training have been collected and observed with different set-up, which results in significant visual differences between the scan samples provided by biologists. The discrepancy between training and testing images is illustrated by the images presented in Fig. 6.

A side contribution of this paper consists in proposing a solution to recognize unknown pollen species as pollen grains, without assigning them to one of the (known) species considered during training. As a baseline, we consider a combination of UPS and KPS. The approach is depicted in Fig. 7. The UPS model predicts whether a pixel is part of a pollen grain or not. Instead, KPS differentiates the known pollen species, and is supposed to assign pixels that do not correspond to a known species, either because they correspond to background or to an unknown pollen species, to the additional bin class of KPS. Hence, as shown in Fig. 7, combining the predictions from UPS and KPS should make it possible to identify unknown pollen species, which are the pixels labelled as pollen by UPS and bin class by KPS.

Fig. 5 shows the predictions of five unknown pollen species. The KPS of baseline can classify N known pollen species, and is supposed to assign a bin class label to pixels that do not correspond to a known pollen species, either because they correspond to background or are unknown.

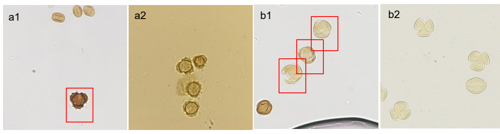

Figure 6. The discrepancy between training and testing images. In the figure, a1 and b1 are testing images, a2 is Tanacetum vulgare(boxed in a1) image of the training, and b2 is Lamium galeobdolon(boxed in b1).

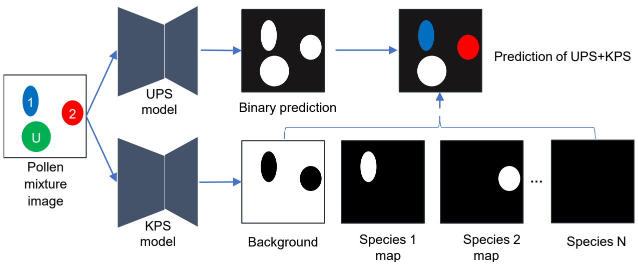

Figure 7. Baseline.The UPS model predicts whether a pixel is part of a pollen grain or not. KPS differentiates the known pollen species, and is supposed to assign pixels that do not correspond to a known species as background. Green pollen grain is unknown in the figure, so it should be background in the prediction of KPS.

We however observe in Fig. 5 that KPS segments unknown species as known types rather than the background. In contrast, each UKPS model appears to successfully segment unseen species into the appropriate unknown class.

4.3. Quantitative Assessment

To provide a quantitative evaluation of our method, we have tested the baseline and five UKPS models. The IoU of each class is shown in Figs. 8 to 12 generated based on test images that only contain pollen species that are available at training (meaning from the 131 species considered when training KPS and UKPS). We observe that, in general, most of the species that each UKPS model can see in the corresponding family has a slightly higher IoU than the baseline, and what is important is that the UKPS model can identify unknown species well. And the baseline can theoretically identify unknown pollen types, but experiments show that it fails. Tab. 1 presents the IoU of background and unknown category on the whole 18 absolutely unknown species. The results show that our proposed UKPS can segment unknown pollen types well, compared with the baseline.

5. Conclusion

With the wide application of deep learning in computer vision, automatic pollen segmentation becomes realistic. However, annotating pollen grain images is a very time-consuming task, especially when dealing with pollen mixtures. Moreover, the dataset collected can not include all

	IoU (%)	
Model	Background	Unknown
UPS+KPS(see Fig. 7)	88.06	5.71
Asteraceae model	99.26	94.97
Brassicaceae model	99.13	69.20
Fabaceae model	99.21	88.82
Lamiaceae model	99.24	90.98
Rosaceae model	99.18	91.19

Table 1. The IoU of background and unknown type on unknown pollen species images.

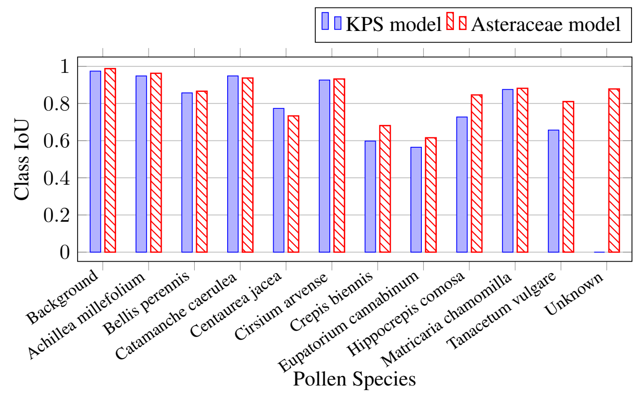

Figure 8. Class IoU of the Asteraceae family using KPS and Asteraceae model

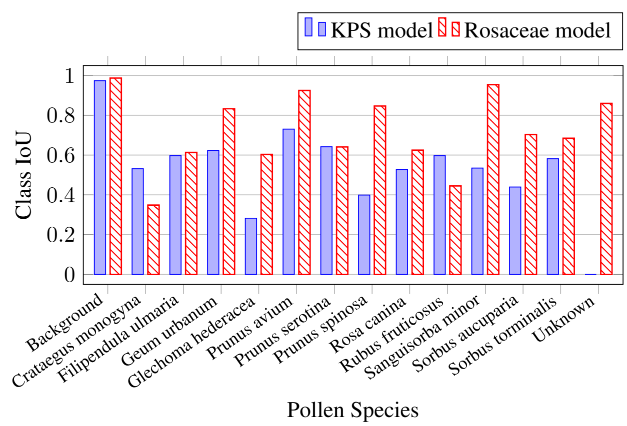

Figure 9. Class IoU of the Rosaceae family using KPS and Rosaceae model

pollen species, and there is some bias between the training set and that in the real application. In this research, we propose to train a universal pollen segmentation model using a small set of data that can be annotated at reasonable cost. This model is then used to automatically annotate a large dataset of images that have the particularity to include a single pollen species. Hence, the pollen segmentation mask

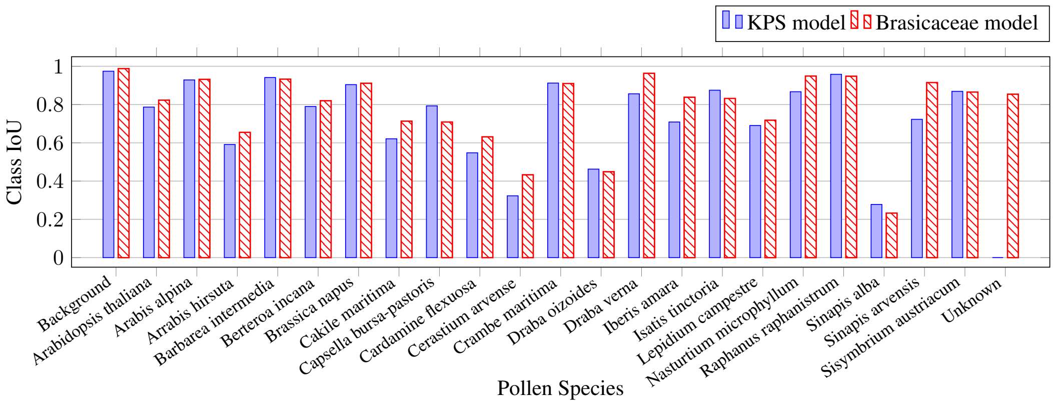

Figure 10. Class IoU of the Brasicaceae family using KPS and Brasicaceae model

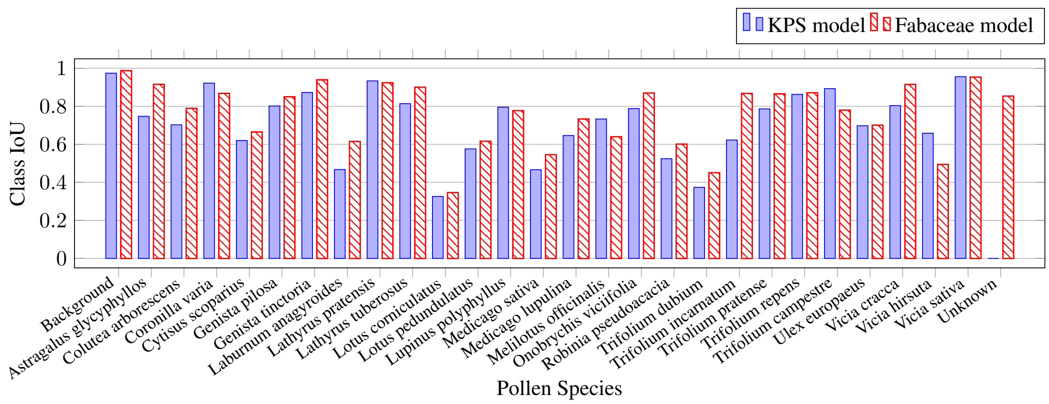

Figure 11. Class IoU of the Fabaceae family using KPS and Fabaceae model

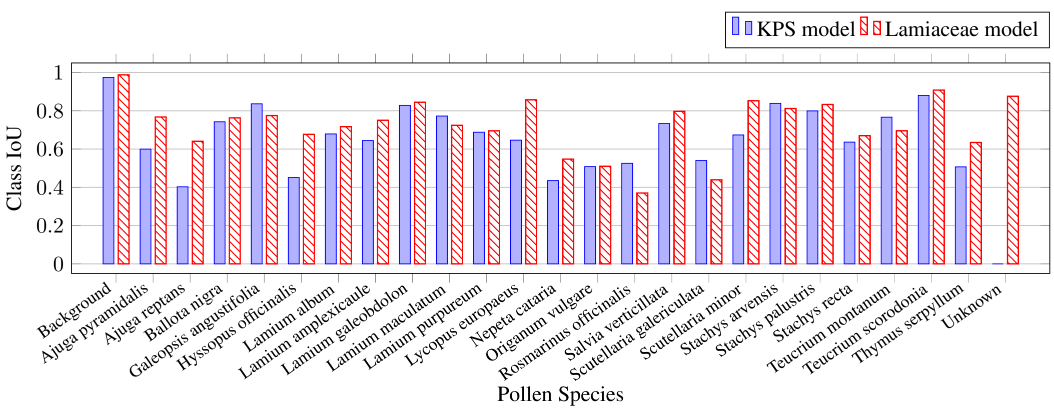

Figure 12. Class IoU of the Lamiaceae family using KPS and Lamiaceae model

actually provides a pollen species segmentation, which can be used to train a pollen species segmentation model. The UKPS can successfully assign pollen species unseen by the model into unknown class rather than known classes. The resulting model is shown to provide accurate segmentation of species, including in pollen mixtures. Our experiments also demonstrate that dedicating a model to the segmentation of a subset of the pure pollen species available at training makes it possible to train a bin pollen class, dedicated to all pollen classes (including the one unavailable at training) that are not part of the model-specific subset of species.

References

[1] Saad Albawi, Tareq Abed Mohammed, and Saad ALZAWI. Understanding of a Convolutional Neural Network. Aug. 2017. 1

[2] Sebastiano Battiato, Alessandro Ortis, Francesca Trenta, Lorenzo Ascari, Mara Politi, and Consolata Siniscalco. Detection and Classification of Pollen Grain Microscope Images. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 4220–4227, Seattle, WA, USA, June 2020. IEEE. 1, 2

[3] André R. de Geus, Celia A.Z. Barcelos, Marcos A. Batista, and Sérgio F. da Silva. Large-scale Pollen Recognition with Deep Learning. In *2019 27th European Signal Processing Conference (EUSIPCO)*, pages 1–5, Sept. 2019. 2

[4] M. Dyrmann, A. K. Mortensen, H. Midtiby, and R. Jørgensen. Pixel-wise classification of weeds and crops in images by using a Fully Convolutional neural network. *undefined*, 2016. 2

[5] P. Garcia-Oliveira, M. Fraga-Corral, A. G. Pereira, C. Lourenço-Lopes, C. Jimenez-Lopez, M. A. Prieto, and J. Simal-Gandara. Scientific basis for the industrialization of traditionally used plants of the Rosaceae family. *Food Chemistry*, 330:127197, Nov. 2020. 4

[6] Morgan R Gostel and W John Kress. The expanding role of dna barcodes: Indispensable tools for ecology, evolution, and conservation. *Diversity*, 14(3):213, 2022. 1

[7] James A. Grant-Jacob, Matthew Praeger, Robert W. Eason, and Ben Mills. Semantic segmentation of pollen grain images generated from scattering patterns via deep learning. *J. Phys. Commun.*, 5(5):055017, May 2021. 2

[8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1026–1034, 2015. 4

[9] P. P. J. Herman and N. Swelankomo. Asteraceae. *Bothalia*, 41(1):176–178, Dec. 2011. 4

[10] Kainat Kanwal, Muhammad Zafar, Amir Muhammad Khan, Tariq Mahmood, Qamar Abbas, Fethi Ahmet Ozdemir, Mushtaq Ahmad, Shazia Sultana, Anam Fatima, Aqsa Aziz, et al. Implication of scanning electron microscopy and light microscopy for oil content determination and seed morphology of verbenaceae. *Microscopy research and technique*, 85(2):789–798, 2022. 1

[11] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9404–9413, 2019. 2

[12] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, May 2015. 1

[13] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 2

[14] J. Víctor Marcos, Rodrigo Nava, Gabriel Cristóbal, Rafael Redondo, Boris Escalante-Ramírez, Gloria Bueno, Óscar Déniz, Amelia González-Porto, Cristina Pardo, François Chung, and Tomás Rodríguez. Automated pollen identification using microscopic imaging and texture analysis. *Micron*, 68:36–46, Jan. 2015. 2

[15] Raphaël Marée, Loïc Rollus, Benjamin Stévens, Renaud Hoyoux, Gilles Louppe, Rémy Vandaele, Jean-Michel Begon, Philipp Kainz, Pierre Geurts, and Louis Wehenkel. Collaborative analysis of multi-gigapixel imaging data using Cytomine. *Bioinformatics*, 32(9):1395–1401, May 2016. 4

[16] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 6

[17] Samuel Paul Raj, Pravin Raj Solomon, and Baskar Thangaraj. Fabaceae. In Samuel Paul Raj, Pravin Raj Solomon, and Baskar Thangaraj, editors, *Biodiesel from Flowering Plants*, pages 291–363. Springer, Singapore, 2022. 4

[18] R. R. Raja. Medicinally potential plants of Labiatae (Lamiaceae) family: An overview. *Research Journal of Medicinal Plant*, 6(3):203–213, 2012. 4

[19] M. Rodríguez-Damián, E. Cernadas, Arno Formella, Manuel Fernandez-Delgado, and Pilar Sa-Otero. Automatic detection and classification of grains of pollen based on shape and texture. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 36(4):531–542, Aug. 2006. 1

[20] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. 3

[21] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. ”GrabCut”: Interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.*, 23(3):309–314, Aug. 2004. 3

[22] Renate Schmidt and Ian Bancroft. *Genetics and Genomics of the Brassicaceae*. Jan. 2011. 4

[23] Víctor Sevillano and José L. Aznarte. Improving classification of pollen grain images of the POLEN23E dataset through three different applications of deep learning convolutional neural networks. *PLOS ONE*, 13(9):e0201807, Sept. 2018. 2

[24] C. Travieso, J. Briceño, J. R. Ticay-Rivas, and J. B. Alonso. Pollen classification based on contour features. *2011 15th IEEE International Conference on Intelligent Engineering Systems*, 2011. 1, 2

[25] Jason Wang, Serra Mall, and Luis Perez. The Effectiveness of Data Augmentation in Image Classification using Deep Learning. page 8. 4

[26] Haishan Ye, Yujun Li, Cheng Chen, and Zhihua Zhang. Fast Fisher Discriminant Analysis with Randomized Algorithms. *Pattern Recognition*, 72, June 2017. 2

[27] Shichao Zhang, Xuelong Li, Ming Zong, Xiaofeng Zhu, and Debo Cheng. Learning k for kNN Classification. *ACM Transactions on Intelligent Systems and Technology*, 8:1–19, Jan. 2017. 2