

Flower Mapping in Grasslands with Drones and Deep Learning

Johannes Gallmann¹ (Orcid ID: 0000-0003-4782-2768)

Beatrice Schüpbach² (Orcid ID: 0000-0003-4090-8155)

Katja Jacot²

Matthias Albrecht² (Orcid ID: 0000-0001-5518-3455)

Jonas Winizki²

Helge Aasen³ (Orcid ID: 0000-0003-4343-0476)

Correspondence: Helge Aasen, Universitätstrasse 2, 8092 Zürich, Switzerland,

helge.aasen@usys.ethz.ch

Headline: Automated Flower Abundance Mapping

¹Department of Computer Science, ETH Zürich, Zürich, Schweiz

²Agricultural Landscape and Biodiversity Group, Agroscope, Zürich, Schweiz

³Department of Agricultural Science, ETH Zürich, Zürich, Schweiz

Abstract

- Manual assessment of flower abundance of different flowering plant species in grasslands is a time consuming process.
- We present an automated approach to determine the flower abundance in grasslands from drone images using a deep learning (Faster R-CNN) object detection approach, which is trained and evaluated on data of five flights and two sites. Our deep learning network is able to identify and classify individual flowers.
- The novel method allows generating spatially explicit maps of flower abundance that meets or exceeds the accuracy of the manually counted extrapolation method and is less labor intensive. The results are very good for some types of flowers with precision and recall being close to or higher than 90 %. Other flowers are detected poorly due to reasons such as lack of enough training data, appearance changes due to phenology or flowers being too small to be reliably distinguishable on the aerial images.
- The method is able to give precise estimates of the abundance of many flowering plant species. The collection of more training data will allow better predictions in the future for the flowers that are not well predicted yet. The developed pipeline can be applied to any sort of aerial object detection problems.

1 Keywords

Aerial Images, Drones, Faster R-CNN, Flower Abundance Mapping, Machine Learning, Object Detection, Remote Sensing, Unmanned Aerial Vehicles

26 **2 Introduction**

27 The service done by pollinators in farmlands is estimated to value more than
28 150 Billion Euros a year worldwide (Gallai, Salles, Settele, & Vaissière 2009).
29 Their declining numbers (Hallmann et al. 2017) motivate many ecologists to
30 study their interplay with the environment. This includes the assessment of
31 flower abundance and distribution, which is an extremely time consuming task.

32 In the last 10 years, rapid development in sensor technology and robotics
33 have enhanced the capabilities of unmanned aerial vehicles (UAVs) (Anderson
34 & Gaston 2013; Pajares 2015; Sanchez-Azofeifa et al. 2017; Aasen, Honkavaara,
35 Lucieer, & Zarco-Tejada 2018). Today it is both technologically possible and
36 affordable to take ultra-high spatial resolution images of large areas (several
37 deca-ha with ground resolution of 1 cm / pixel). When flying lower and slower,
38 even resolutions of down to millimetres can be reached. Consequently, UAVs
39 have also been used in many ecological settings. These include invasive species
40 mapping (Martin et al. 2018; Müllerová et al. 2017; Hill et al. 2017; Kattenborn,
41 Eichel, & Fassnacht 2019; de Sá et al. 2018), wild live assessment (Andrew &
42 Shephard 2017; Hollings et al. 2018; Christiansen et al. 2019; Eikelboom et al.
43 2019; Rey, Volpi, Joost, & Tuia 2017) and plant biodiversity estimation (Getzin,
44 Wiegand, & Schöning 2012).

45 Recently, deep learning based classification methods have appeared that are
46 able to utilize the details of ultra high resolution image data. We use the Faster
47 R-CNN object detection pipeline (Ren, He, Girshick, & Sun 2015). It utilizes
48 deep convolutional neural networks (CNNs) to detect and classify objects in
49 RGB images. A deep CNN is a network with many layers. It takes the pixels
50 of an image as input and as output predicts the likelihood for each class label
51 it has been trained on. Internally it applies thousands of learned filters to all
52 regions of the image and in the end combines them to find the likelihood of

53 each class label. Recently, such approaches have also been introduced to detect
54 and count animals (Eikelboom et al. 2019; Rey, Volpi, Joost, & Tuia 2017) and
55 plants (Eikelboom et al. 2019; Kattenborn, Eichel, & Fassnacht 2019; Osco et al.
56 2020) in an ecological context.

57 Remote flower mapping in a grassland containing many species is a challeng-
58 ing task, since the structures are fine and flowers might be occluded by other
59 plants. Current approaches of automated flower mapping work with image res-
60 olutions in the range of centimeters or even meters per pixel (Landmann et al.
61 2015; Chen, Jin, & Brown 2019; Abdel-Rahman et al. 2015) and are therefore
62 not suited to detect individual flowers and differentiate between flower species of
63 similar color. Other approaches are handcrafted for a single species (Campbell
64 & Fearn 2018; Horton, Cano, Bulanon, & Fallahi 2017) and are not applicable
65 to a wide range of use cases.

66 In this article we present a deep learning based method to collect information
67 about flower abundance and distribution in grasslands from drone images. To
68 evaluate its performance we address several questions:

- 69 (i) Is it possible to identify flowers in overhead images with flowers spanning
70 only a few pixels at a ground resolution of 1.5 mm per pixel using deep
71 CNNs?
- 72 (ii) How do UAV based automated counts compare to manual in field mea-
73 surements by an expert?
- 74 (iii) How does automatically generated flower abundance maps of a whole field
75 compare to the educated guess by an expert?

76 **3 Materials and Methods**

77 **3.1 Overview**

78 The proposed method can be divided into the three main phases of data collec-
 79 tion (section 3.2), model training (section 3.3) and application to unseen images
 80 (section 3.4) as depicted in fig. 1.

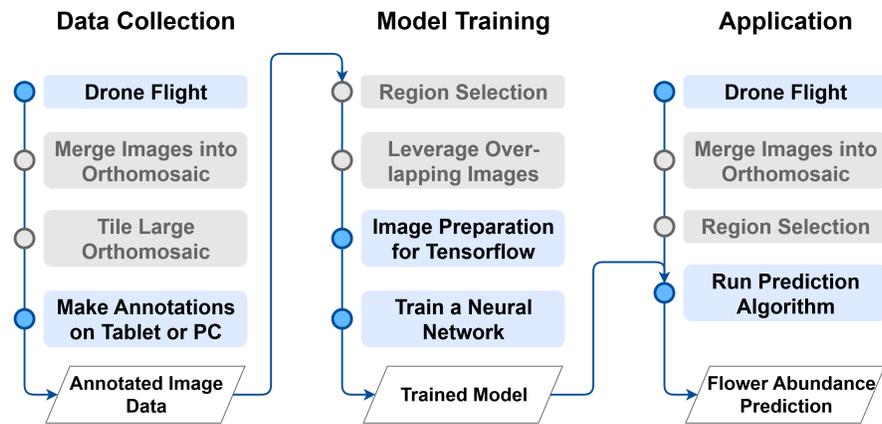


Figure 1: Overview of the proposed method. Grey colored steps might not be necessary for some use cases.

<i>Ranunculus</i> sp (n = 474)	<i>Lotus corniculatus</i> (3271)	<i>Galium mollugo</i> (659)	<i>Crepis biennis</i> (159)	<i>Centaurea jacea</i> (805)
- <i>Ranunculus bulbosus</i> (442)	- <i>Lotus corniculatus</i> (2926)	- <i>Galium mollugo</i> (202)	- <i>Crepis biennis</i> (89)	- <i>Centaurea jacea</i> (786)
- <i>Ranunculus friesianus</i> (8)	- <i>Lathyrus pratensis</i> (345)	- <i>Achillea millefolium</i> (338)	- <i>Leontodon hispidus</i> (10)	- <i>Lychnis flos-cuculi</i> (19)
- <i>Ranunculus acris</i> (24)		- <i>Daucus carota</i> (65)	- <i>Tragopogon pratensis</i> (8)	
		- <i>Carum carvi</i> (54)	- <i>Picris hieracioides</i> (52)	

Table 1: Plant species that are combined into one group.

81 **3.2 Data Collection**

82 **3.2.1 Dataset**

83 The dataset on which the method is evaluated consists of 10000 annotated flow-
84 ers. The aerial images are captured at two sites and five days from a flight height
85 of 19 meters with a ground sampling distance of approximately 1.5 mm/pixel.
86 For the collection of the flower dataset a drone model called TransformerUAV
87 (Copting GmbH 2017) as well as a DJI Matrice 600 PRO (SZ DJI Technology
88 Co., Ltd. 2018) were used. Both drones are programmed to fly along a prede-
89 fined route such that the area is fully covered and the images have an overlap
90 of 60 % to 90 %. Attached to the drone is a Sony ILCE-7RM2 (Sony Corpo-
91 ration 2015) camera that takes 42.2 Megapixel photos in combination with a
92 Zeiss Batis 1.8/85 telephoto lens (Carl Zeiss AG 2017). The weather was sunny
93 on all flight days. One of the two sites has been managed extensively during
94 the last 15 years such that the plant diversity in this meadow is very high. A
95 total of 40 different flowering plant species have been found between May 23rd
96 and July 3rd. Approximately half of these 40 flower species are omitted in the
97 analysis because too few samples are present in the dataset to reasonably train
98 a neural network. We excluded all flowers with less than 50 samples in total
99 from the experiments. As summarized in table 1, some flowers are combined
100 into groups because they have few annotated samples or they look similar to
101 other flowers. Because the individual flowers within an inflorescence can rarely
102 be identified in the drone images, all inflorescences are annotated as one flower
103 instance. Subsequently, when referring to the term flower, inflorescences are
104 included as well.

105 **3.2.2 Traditional Data Acquisition**

106 Traditionally, information about flower abundance is acquired by counting or
107 estimating the flowers by hand within survey plots which are distributed inside
108 the area of interest. We used 15 survey plots that are one by one meters wide.
109 Once the flowers present within the survey plots are counted or estimated, these
110 numbers are extrapolated to the size of the whole area of interest. If the positions
111 of the survey plots are well chosen, this method should produce a good estimate
112 of the abundance of flowers. We carried out the traditional approach of manual
113 counting in parallel to each iteration of the drone based data acquisition method
114 and used it as a baseline.

115 **3.2.3 Drone Based Data Acquisition**

116 Before the drone flight, ground control points (GCPs) are placed inside the test
117 region. For planning a proper placement of the GCPs, refer to Roth, Hund,
118 & Aasen (2018). GCPs are small signs with a unique pattern facing upwards
119 so that they can be recognized on the drone images. The exact GPS positions
120 of all these GCPs are collected with a differential GPS with a precision of a
121 few centimeters. Later they are used in the Agisoft software (Agisoft 2019)
122 as described below. Having the GCPs in place, the drone can be flown along
123 a predefined route across the field with a camera attached that takes a large
124 amount of highly overlapping aerial images of the field. Using a drone with
125 RTK GNSS potentially allows to omit the need for GCPs.

126 After the flight, the relative positions of the large amount of overlapping
127 aerial images are reconstructed and merged together into a large orthomosaic.
128 An orthomosaic is a detailed, accurate photo representation of an area, created
129 out of many photos that are stitched together and geometrically corrected. We
130 use the structure from motion (SfM) approach (Ullman 1979; Harwin & Lucieer

131 2012) implemented in the software Agisoft Metashape Version 1.5.3 (Agisoft
132 2019). Agisoft takes all aerial images as inputs. It aligns all photos and generates
133 a point cloud model. We used a sparse point cloud. From the sparse point
134 cloud either a digital surface model or a mesh of the topography of the research
135 area can be created. Then, based on the topography and the reconstructed
136 relative positions and orientations of the images, an orthomosaic is generated.
137 We enabled the option *blending disabled* to use the original information of the
138 images in the orthomosaic.

139 Agisoft automatically detects the unique pattern on the GCPs to map the
140 GPS coordinates to each of them. The advantage of providing the positions of
141 the GCPs in the field is that Agisoft creates an orthomosaic that is orthorec-
142 tified and georeferenced. Georeferencing of the orthomosaic is later needed to
143 display the user’s position in the Android FieldAnnotator application as well
144 as to be able to copy annotations to the single orthorectified images that are
145 georeferenced (cf. 3.2.4 and 3.2.5 for further reading).

146 **3.2.4 Annotating**

147 Having an orthomosaic of the region of interest, flowers have to be annotated.
148 The annotated flowers are needed as training data for the machine learning
149 model. Since the survey plots needed for the traditional data acquisition are
150 visible on the drone images, we annotated all flowers within these plots. This
151 allows us to verify whether the number of flowers visible on the drone images
152 are comparable to the number of flowers that are manually counted by hand.

153 For annotating, we use the LabelMe program (Wada 2016) and an Android
154 tablet application called FieldAnnotator which we specifically developed for this
155 purpose. The advantage of being able to make the annotations on a tablet is that
156 they can be made directly in the field. This might be necessary because some
157 flowers can be very hard to distinguish in the image alone. If one can compare

158 the image to the actual flowers on site, the quality of the training data can be
159 improved and it is made sure that the number of false annotations is minimized.
160 Android tablets are not capable of handling large orthomosaics (around 50000
161 times 50000 pixels for a 30 times 30 meters area). Therefore a script tiles the
162 orthomosaic into small chunks of 256 times 256 pixels in various zoom levels
163 before these tiles are then imported into the FieldAnnotator application. The
164 resulting annotations are stored in a json file.

165 **3.2.5 Leveraging Overlapping Images**

166 Since the camera attached to the drone captures a large amount of highly over-
167 lapping images, the idea is to use the overlapping images as additional training
168 data. Since the flowers are pictured from a slightly different angle on each im-
169 age and the background changes from image to image, this provides valuable
170 additional training data. Grasslands have a very complex structure and it is
171 hard to reconstruct the exact geometry of the images. Therefore, the copied
172 annotations are slightly shifted within the overlapping images. To correct for
173 the shift, a script lets the user view and adjust all annotations in the LabelMe
174 application. These slight adjustments of the annotations take significantly less
175 time than collecting new data.

176 **3.3 Model Training**

177 **3.3.1 Selecting Regions of Interest in Annotated Images**

178 In case images are only partly annotated, we developed a script that allows the
179 user to cut out certain regions (polygon shaped) from the images. Only the
180 image pixels within these selected regions are kept while the rest of the image
181 pixels are overridden with black. This ensures that the Tensorflow model (Abadi
182 et al. 2015) does not learn to classify non-annotated flowers as the background

183 class.

184 **3.3.2 Image Preparation for Tensorflow**

185 The training data consisting of image files alongside with json files containing the
186 annotations has to be converted into a format that is supported by Tensorflow.

187 To do so, our pipeline automatically carries out the subsequent steps. First,
188 the images are split up into tiles. The default tile size is set to 450 times 450
189 pixels. These image tiles are then upscaled by a factor of two to 900 times 900
190 pixel tiles as suggested by Hu & Ramanan (2017) and justified in supplementary
191 material A.1.1. The tiles are overlapping such that flowers positioned on the
192 edge of two tiles are not lost as training data but are always present as a whole
193 in at least one tile. Additionally, all annotations (including point and polygon
194 annotations) are converted to bounding boxes. Finally, the images are split up
195 into train, test and validation set.

196 **3.3.3 Neural Network Training**

197 The core of the pipeline consists of a CNN. We use the Faster R-CNN archi-
198 tecture. This architecture outputs the bounding box coordinates of the objects
199 it recognizes on an input image. The Faster R-CNN architecture requires more
200 compute power than other architectures but it has been shown that it performs
201 well on aerial and other high resolution images (Carlet & Abayowa 2017; Huang
202 et al. 2017). Since the default configuration of the Faster R-CNN architecture
203 is not optimized to detect very small objects (Huang et al. 2017; Zhang et al.
204 2017) of only a few pixels in diameter, such as flowers in aerial images, we
205 adjusted some parameters (cf. A.1.1 in the supplementary materials for exper-
206 iment results on different parameter combinations). Additionally, we use data
207 augmentation techniques to increase the diversity of our dataset. The following
208 data augmentation options are used: random horizontal and vertical flips, ran-

209 dom brightness adjustments, random contrast adjustments, random saturation
210 adjustments and random box jittering.

211 During training, the validation set is used to decide when to change the
212 learning rate and when to stop training. Every 2500 steps the training is paused
213 and the prediction algorithm followed by the evaluation algorithm is run on
214 the validation set. The learning rate is adjusted if for the last 15000 steps
215 no further improvements were made. After adjusting the learning rate two
216 times from 3×10^{-4} to 3×10^{-5} and from 3×10^{-5} to 3×10^{-6} , the training is
217 stopped if for 15000 steps again no improvement on the performance has been
218 made. The number of 15000 steps is chosen empirically. Reducing the learning
219 rate twice by a factor of 10 is directly adapted from the Faster R-CNN default
220 configuration. The evaluation metric can be chosen as either the F1 score or
221 the mean average precision (mAP). Section 3.4 further explains the prediction
222 and evaluation process.

223 The number of training examples can vary heavily from class to class. There-
224 fore each class is assigned a weight. The weight is inversely proportional to the
225 number of training examples and influences the loss function during training.
226 This ensures that the network does not just optimize to detect the most com-
227 mon classes. Each mistake on a less common class has a much higher penalty
228 to the loss function as a consequence. Once a network is fully trained it can be
229 exported as an inference graph. This exported inference graph is then used by
230 the prediction and evaluation scripts described in section 3.4.

231 **3.4 Application to Unseen Images**

232 **3.4.1 Predictions**

233 The trained network can be used to make predictions on images of arbitrary
234 size (e.g. orthomosaics) provided they have a similar ground sampling distance

235 to the training images. The pipeline handles the tiling of large images as well
236 as the reassembling of the prediction results from the single tiles. Optionally
237 a region of interest can be selected within an image. As a consequence, only
238 the flower abundance within this region of interest is assessed by the prediction
239 algorithm.

240 The prediction algorithm draws the bounding boxes of all detected flowers
241 onto the image and saves the statistics about the flower abundance to a json
242 file. To improve the prediction accuracy, the tiles have an overlap of 100 pixels
243 by default. This ensures that as long as a flower is not larger than 100 pixels in
244 diameter, it is fully visible on at least one tile. Error prone predictions close to
245 or on the edge of a tile can therefore be ignored because they are fully covered
246 on the adjacent tile. Nevertheless, having this overlap introduces the problem of
247 duplicate predictions. This is mitigated by applying non maximum suppression
248 with an intersection-over-union threshold of 0.3 similar to Ozge Unel, Ozkalayci,
249 & Cigla (2019). Meaning that for all predictions that have an overlap of more
250 than 30 %, only the one with the highest confidence score is kept.

251 **3.4.2 Evaluations**

252 To evaluate the performance of a model, the predictions on the test set are
253 compared to the ground truth of the test set. The main metrics of interest are
254 precision and recall. To compute precision and recall values, the true positive
255 (TP), false positive (FP) and false negative (FN) predictions have to be known.
256 In order to obtain these values the predictions are sorted by their confidence.
257 Then it is looped through all the predictions and each of them is compared to
258 all ground truth bounding boxes of the same label. To compare two bounding
259 boxes, the intersection-over-union (*iou*) formula is used:

$$iou = \frac{\textit{intersection area}}{\textit{area of union}}$$

260 If the greatest *iou* value is greater than some threshold (default of 0.3),
261 the corresponding ground truth box is marked as used and the prediction is
262 marked as true positive. If the largest *iou* value is lower than the threshold,
263 the prediction is marked as false positive. After this process is done for each
264 prediction, all ground truth entries that are not marked as used are counted as
265 false negatives. Having the TP, FP and FN numbers, the precision and recall
266 values can easily be calculated using the following formulas:

$$\textit{precision} = \frac{TP}{TP + FP}$$

267

$$\textit{recall} = \frac{TP}{TP + FN}$$

268 A good way to rate the performance of a model is to compute the F1 score.
269 The F1 score is calculated as follows:

$$F1 = 2 \cdot \frac{\textit{precision} \cdot \textit{recall}}{\textit{precision} + \textit{recall}}$$

270 The better the precision and recall values are, the better the F1 score gets.
271 It rates precision and recall equally and reaches its maximum of 1 at perfect pre-
272 cision and recall. As an alternative to the F1 score, the mean average precision
273 (mAP) as defined in the PASCAL VOC Challenge Development Kit (Evering-
274 ham & Winn 2011) can be used to rate a model's performance.

275 **3.4.3 Visualizations**

276 The pipeline offers various options for visualizing the results. Apart from draw-
277 ing the predictions as colored bounding boxes onto the images, erroneous pre-
278 dictions can be highlighted. Additionally, heatmaps that visualize the density
279 distribution of the flowers can be generated from the prediction output. The
280 size of the kernel for the flower density mapping is customizable. Optionally the
281 heatmap can be drawn directly onto the image. The heatmaps can be generated
282 for an individual class or for all classes. If the input images are georeferenced,
283 there is the option to generate one heatmap from a collection of images. If the
284 images are overlapping, the heatmap indicates the average number of flowers
285 found at a particular position. Furthermore, the user can provide the geo co-
286 ordinates of the upper left and lower right corner of the desired output region.
287 The script will then output a heatmap of exactly that region. This allows for
288 time series generations. Example results of such time series generations can be
289 viewed in section 4.3.

290 **4 Results**

291 **4.1 Manual Counting vs Drone Image Based Tablet An-** 292 **notations**

293 Since the exact same areas are annotated on the tablet as they are manually
294 counted by hand, we are able to directly compare the numbers of flowers an-
295 notated on the drone images on the tablet to the numbers of flowers manually
296 counted by hand. Table 2 lists a representative subset of all flowers found within
297 the test fields.

298 Some flowers are hardly visible on the drone images and therefore signif-
299 icantly less instances are counted in the tablet annotations compared to the

Flower	Manually Counted	Tablet Annotations
<i>Leucanthemum vulgare</i>	724	960
<i>Onobrychis viciifolia</i>	483	105
<i>Lotus corniculatus</i>	1943	748
<i>Salvia pratensis</i>	142	127
<i>Ranunculus sp</i>	431	474
<i>Knautia arvensis</i>	371	471
<i>Trifolium pratense</i>	129	72
<i>Medicago lupulina</i>	117	5
<i>Centaurea jacea</i>	25	28

Table 2: Comparison of selected manually counted total numbers to tablet annotations.

300 manually counted data. *Onobrychis viciifolia*, *Medicago lupulina* and to some
301 extent *Trifolium pratense* fall under this category. The flowers of *Medicago*
302 *lupulina* are too small to be reliably identifiable on the drone images. *Trifolium*
303 *pratense* and *Onobrychis viciifolia* would be large enough but often they are
304 hardly distinguishable from the background. Refer to table 2 for visualizations
305 of 25 flowers found within the test fields. For other flowers (*Leucanthemum vul-*
306 *gare*, *Ranunculus sp*, *Knautia arvensis* and *Centaurea jacea*) there are more
307 flowers annotated on the tablet than manually counted by hand (cf. section
308 5.1).

309 4.2 Prediction on a Meadow

310 The idea of the experiments in this section is to simulate an as realistic as
311 possible situation. The data of one of the five flights is entirely used as test
312 data. 90 % of the data of all other flights is used as training data and 10 % as
313 validation data.

314 4.2.1 Performance inside Survey Plots

315 We compared the tablet annotations with the deep learning predictions within
316 the survey plots. The prediction performance for each flower can be obtained
317 from table 4. A prediction is considered if its confidence score is greater than

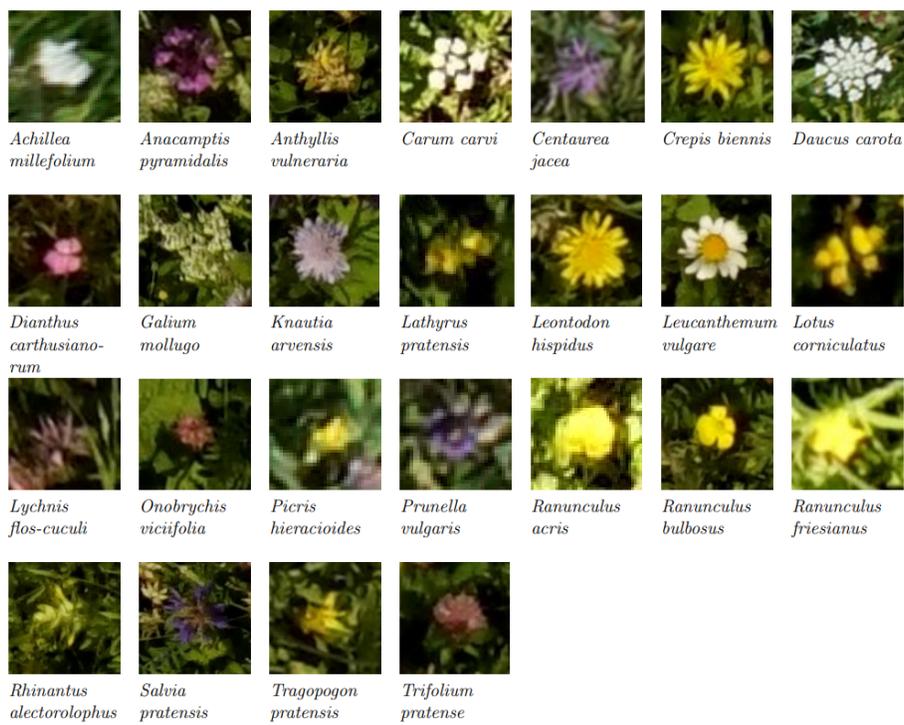


Figure 2: Excerpts from aerial images of the most common flowers.

0.2. The overall precision and recall are 87 % and 84.2 % respectively. The vast majority of the flowers present in the test data of June 14th are *Knautia arvensis*, *Leucanthemum vulgare* and *Lotus corniculatus*. These three flowers perform well and therefore the good overall score is mainly determined by these three flowers. All the other flowers perform worse than the overall performance indicates.

Table 3 shows the confusion matrix of this experiment. It is striking that there are only a few confusions between different flowers (brown). The much more common cases are that flowers are predicted where there are none (red) and flowers are not predicted where they should be (orange). The green entries denote the correctly predicted flowers.

Table 4 shows that the flowers with little training data tend to not perform well. The question is whether this is due to the lack of enough training data or because assigning an inversely proportional weight to each class during training

	<i>A. vulneraria</i>	<i>C. jacea</i>	<i>C. biennis</i>	<i>D. carthusianorum</i>	<i>G. mollugo</i>	<i>K. arvensis</i>	<i>L. vulgare</i>	<i>L. corniculatus</i>	<i>O. viciifolia</i>	<i>P. vulgaris</i>	<i>Ranunculus sp</i>	<i>R. alectorolophus</i>	<i>S. pratensis</i>	<i>T. pratense</i>	False Negatives
<i>A. vulneraria</i>	1	-	-	-	-	-	-	3	-	-	-	-	-	-	2
<i>C. jacea</i>	-	27	-	-	-	17	-	-	-	1	-	-	3	3	2
<i>C. biennis</i>	-	-	14	-	-	-	-	5	-	-	-	-	-	-	2
<i>D. carthusianorum</i>	-	3	-	8	-	1	-	-	10	-	-	-	-	6	6
<i>G. mollugo</i>	-	-	-	-	8	-	-	-	-	-	-	-	-	-	8
<i>K. arvensis</i>	-	-	-	-	-	412	1	-	-	-	-	-	2	-	23
<i>L. vulgare</i>	-	-	1	-	1	4	906	-	-	-	-	-	1	-	109
<i>L. corniculatus</i>	-	-	6	-	-	-	1	877	-	-	-	-	-	-	142
<i>O. viciifolia</i>	-	1	-	-	-	1	-	-	45	-	-	-	-	11	37
<i>P. vulgaris</i>	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
<i>Ranunculus sp</i>	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
<i>R. alectorolophus</i>	-	-	-	-	-	-	-	1	-	-	-	8	-	-	17
<i>S. pratensis</i>	-	-	-	-	-	-	-	-	-	-	-	-	12	-	3
<i>T. pratense</i>	-	-	-	-	-	1	-	-	-	-	-	-	-	4	2
False Positives	4	6	17	-	32	24	31	117	3	-	1	5	6	18	-

Table 3: The table shows the confusion matrix. The columns represent what the model predicted and the rows represent what the model should have predicted (the ground truth). The green, red, orange and brown numbers denote TP, FP, FN and confusions between two flowers respectively.

332 is not sufficient to regularize the loss function. Therefore we trained a separate
 333 network in which the three best performing flowers (*Leucanthemum vulgare*, *Lo-*
 334 *tus corniculatus* and *Knautia arvensis*) are ignored and treated as background.
 335 With the mAP rising from 25.2 % to 31.5 % (f1 score improves from 47 % to
 336 51.3 %) a certain improvement can be seen but the performance is still signif-
 337 icantly below what is satisfactory. Therefore the possibility of leveraging two
 338 separately trained networks is not further evaluated.

339 When looking at the predictions, there are various sources of errors appar-
 340 ent. Some examples can be seen in fig. 3. For *Leucanthemum vulgare*, a typical
 341 error occurs where two instances are very close to each other as in image a).
 342 In that case often only one of the two flowers is detected. The missing annota-
 343 tion is not caused by the non maximum suppression algorithm as a closer look
 344 discloses. Another typical source of errors are flowers that are on the verge of
 345 fading. In the case of image b) two flowers are detected that are not anno-
 346 tated in the ground truth because the botanical expert considered the flowers
 347 to be faded already. Even when manually counting the flowers by hand it is
 348 sometimes difficult to decide if a flower should be counted or not because of

Flower	Train Instances	Test Instances	Precision	Recall	mAP	F1 Score
<i>A. vulneraria</i>	196	6	20.0 %	16.7 %	0.056	0.182
<i>C. jacea</i>	742	53	73.0 %	50.9 %	0.382	0.6
<i>C. biennis</i>	124	21	36.8 %	66.7 %	0.325	0.475
<i>D. carthusianorum</i>	20	34	100.0 %	23.5 %	0.235	0.381
<i>G. mollugo</i>	546	16	19.5 %	50.0 %	0.1	0.281
<i>K. arvensis</i>	429	438	89.6 %	94.1 %	0.879	0.918
<i>L. vulgare</i>	928	1022	96.5 %	88.6 %	0.861	0.924
<i>L. corniculatus</i>	2153	1026	87.4 %	85.5 %	0.772	0.864
<i>O. viciifolia</i>	92	95	77.6 %	47.4 %	0.407	0.588
<i>R. alectorolophus</i>	23	26	61.5 %	30.8 %	0.218	0.41
<i>S. pratensis</i>	133	15	50.0 %	80.0 %	0.436	0.615
<i>T. pratense</i>	109	7	9.5 %	57.1 %	0.104	0.163
Overall	5495	2759	87.0 %	84.2 %	0.398	0.855

Table 4: Performance of the prediction algorithm on all flower species present in the field on June 14th. The numbers in the *Train Instances* and *Test Instances* columns refer to the ground truth annotations. The overall scores of the performance metrics are weighted means.

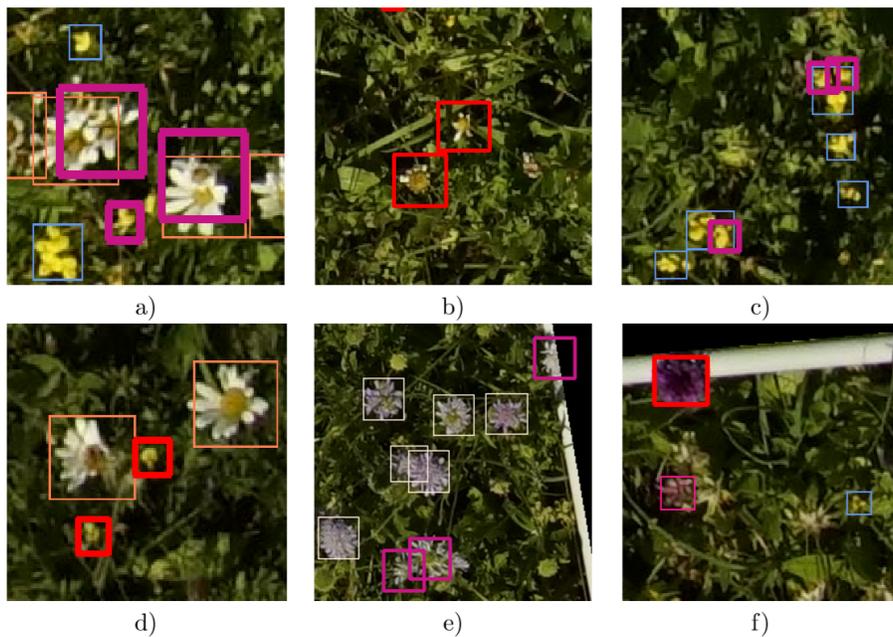


Figure 3: Selection of typical mispredictions. All thin bounding boxes are correct predictions. The bold red bounding boxes denote false positive and the bold violet bounding boxes denote false negative predictions. There are various explanations for the mispredictions: Overlapping flowers (a), partially withered flowers (b and e), collections of flowers (c), missing ground truth annotations (d) and flowers that are missing in the training data (f).

349 the seamless transition from blooming to faded. Two main problems exist for
 350 *Lotus corniculatus*. Firstly, the blooms of *Lotus corniculatus* are often arranged
 351 as small inflorescences as visible in image a) in the bottom left or in image c).
 352 In some cases the network predicts the blooms of an inflorescence as individual
 353 instances while in the ground truth the whole inflorescence is annotated as one
 354 instance. The opposite case is common as well. The second problem of *Lotus*
 355 *corniculatus* are false positive predictions caused by missing ground truth an-
 356 notations (as in image d)). These problems are further discussed in section 5.1.
 357 The main error source of *Knautia arvensis* is again blooms that look different
 358 because they are wilting as for example in image e). In image f) the model er-
 359 roneously predicts a *Knautia arvensis* where there is a *Anacamptis pyramidalis*.
 360 *Anacamptis pyramidalis* is not included in the training because too few training
 361 instances exist.

362 4.2.2 Performance outside Survey Plots

363 We compared the predictions of the deep learning model on the full test field to
 364 the extrapolation of the manually counted flowers. The numbers of manually
 365 counted flowers are extrapolated to the size of the whole field which is 730
 366 square meters. Table 5 lists all flowers that are detected reasonably well inside

Flower	Drone Based Prediction	Extrapolation of Manual Counting	Relative Difference
<i>Centaurea jacea</i>	456	505	10.7 %
<i>Knautia arvensis</i>	8059	8308	3.1 %
<i>Leucanthemum vulgare</i>	7044	10778	53.0 %
<i>Lotus corniculatus</i>	50365*	51139	1.5 %
<i>Onobrychis viciifolia</i>	595	3761	532 %
<i>Salvia pratensis</i>	209	673	222 %

Table 5: Predictions on the whole field of 730 square meters. The 50365 predicted *Lotus corniculatus* are calculated as the multiplicative of the actual predictions of the network (19389) and a ratio of 2.6. The numbers in table 2 suggest that there are 2.6 blooms per prediction on average.

367 the survey plots by the deep learning model. For each flower species the number
368 of deep learning detections in the whole field is listed as well as the number of
369 flowers predicted by the extrapolation of the manual counting.

370 For *Centaurea jacea*, *Knautia arvensis* and *Lotus corniculatus* the number
371 of drone based predictions is very similar to the extrapolation of the manually
372 counted number of flowers. The results are within 11 %, 3 % and 2 % respec-
373 tively. According to heatmaps generated from the drone based predictions (cf.
374 section 4.3), these are also the flowers that are relatively evenly distributed.
375 The extrapolation of the manually counted number of *Leucanthemum vulgare*
376 is 53 % higher than the number of drone based predictions. The question is,
377 which prediction is more accurate. Assuming that the performance of the pre-
378 diction algorithm is similar on the whole field as it is inside the annotated survey
379 plots, the extrapolation of the manual counting must be inaccurate. Even when
380 adding 8 % to the number of drone based predictions to compensate for the rel-
381 atively low recall value of *Leucanthemum vulgare*, the results still have a 47 %
382 gap. The extrapolation is based on the manually counted number of flowers
383 which is lower than the number of tablet annotations within the survey plots as
384 pointed out in section 4.1. If the tablet based numbers were taken, the result
385 of the extrapolation would be an additional 51 % higher making them a total
386 of 131 % higher than the drone based predictions.

387 The main reason for the bad results of *Onobrychis viciifolia* is that it is
388 very hard to distinguish on the drone images. The most probable reason for
389 the unsatisfactory results of *Salvia pratensis* is that the amount of training
390 data is too low to accomplish good results. A likely additional reason is an
391 unrepresentative choice of survey plot locations for these flowers.

392 **4.3 Density Distribution Maps**

393 The heatmaps in fig. 4 depict the abundance of some selected individual flowers
394 in one of our test fields on June 14th. The three heatmaps for *Leucanthemum*
395 *vulgare*, *Lotus corniculatus* and *Knautia arvensis* are generated from the ortho-
396 mosaic.

397 Table 5 contains a time series of an excerpt of our main test site. It illustrates
398 the difference of the abundance evolution of *Leucanthemum vulgare* and *Lotus*
399 *corniculatus*. It is conspicuous that *Lotus corniculatus* is much more evenly
400 distributed than *Leucanthemum vulgare*. While *Leucanthemum vulgare* has a
401 peak population on June 6th, on July 3rd the population is almost completely

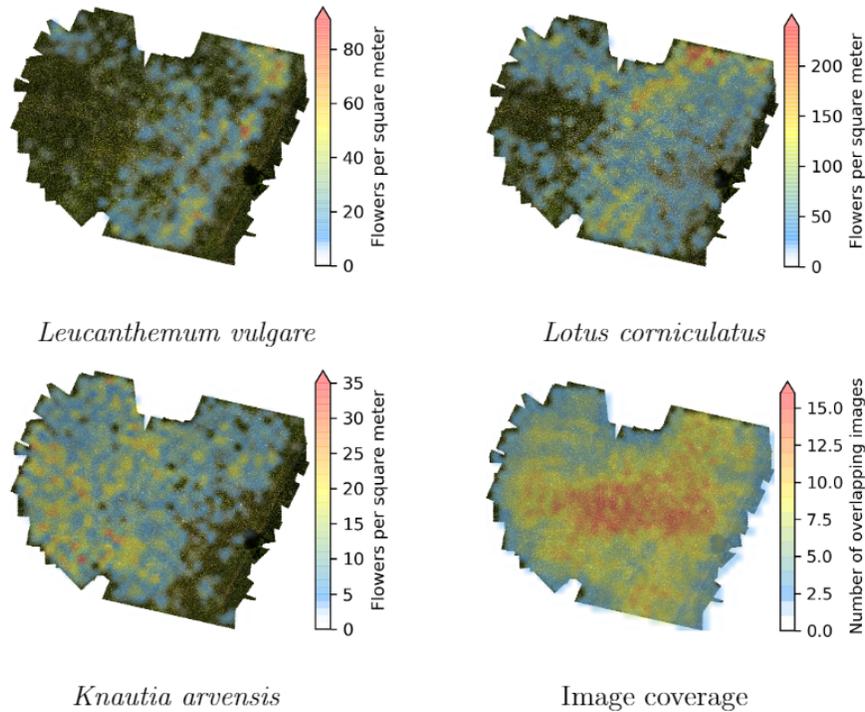


Figure 4: Heatmaps of our main test site for various flower species. The last image depicts the image coverage of the field. In the images, survey plots as well as GCPs are visible.

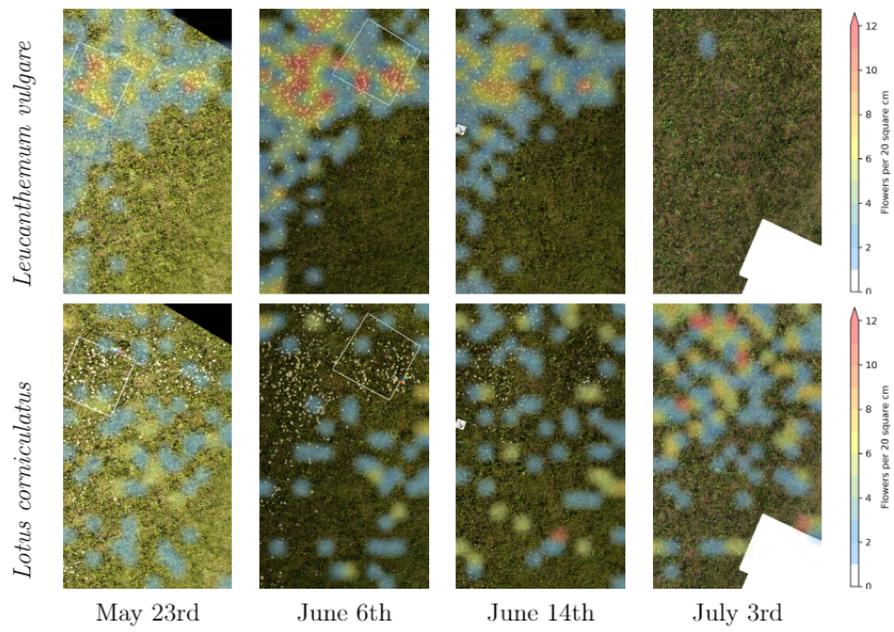


Figure 5: Timeseries of the distribution of *Leucanthemum vulgare* and *Lotus corniculatus* in our main test field.

402 faded. The peak population of *Lotus corniculatus* is much less pronounced.

403 5 Discussion

404 5.1 Different Approaches for Flower Abundance Mapping

405 We evaluated different approaches to map flowers in grasslands. We used manual
406 counting inside survey plots as a baseline and compared it to tablet annotations
407 on images of the survey plots and automated deep learning based mapping inside
408 the survey plots from drone images. Then we compared the extrapolations of
409 the manual counting to the predictions of the deep learning model on the whole
410 test field. Section 4.1 shows that some flowers have more tablet annotations in
411 the images than are manually counted by hand inside the survey plots. This can
412 be explained by the fact that manually counting flowers by hand requires a high
413 level of concentration. Mistakes happen very easily if a lot of flowers are present
414 within a small area. Annotating on an image has the advantage that flowers are
415 marked and therefore the risk of counting twice or forgetting to count a flower is
416 minimized. When combining these falsely counted numbers with non optimally
417 chosen survey plot locations, the extrapolations of the manually counted flowers
418 have the potential to be very inaccurate.

419 With a reliable flower detection model, the results can be much more accu-
420 rate than with the extrapolation from the manual counting. Moreover, the drone
421 based approach has other advantages. The potential to have spatially explicit
422 maps of flowers goes beyond what can be done with the traditional approach of
423 extrapolating the manually counted numbers of flowers within the survey plots.
424 Once a trained network is available, manually labelling the species to train the
425 network is no more necessary. It is sufficient to fly the drone over the meadow
426 and let the deep learning algorithm predict the species. The prediction time

427 of the trained deep learning network for one square meter is approximately 7.4
428 seconds using a GTX 1080 GPU (Nvidia Corporation 2016). On the contrary,
429 manually counting the flowers by hand within a survey plot can take between
430 one and ten minutes, depending on the flower density. The predictions of the
431 network have to be controlled by a good botany expert.

432 Whether it is possible to achieve reliable predictions for a certain flower
433 on drone images depends on several factors. First, enough training data of
434 the flower in question needs to be available. The results suggest that with a
435 few hundred instances good performance can be achieved. Second, also the
436 morphology of the flower has an impact. Flowers such as *Galium mollugo* are
437 difficult for an object detection network to predict reliably. The cause seems
438 to be that this flower can sometimes be very small and in other cases multiple
439 instances of the same flower species cover a large area of partly overlapping
440 inflorescences in which it is difficult to separate the single instances. In such
441 cases it would be interesting to see how an image segmentation network such as
442 U-Net (Ronneberger, Fischer, & Brox 2015), which predicts regions (pixels) that
443 belong to a certain class, would perform. Third, the size of a flower should span
444 a certain minimum amount of pixels. The good results of *Lotus corniculatus*
445 suggest that a diameter of around 5 to 10 pixels is sufficient. These results are
446 likely to be positively influenced by the distinct color and the strong contrast
447 to the background of *Lotus corniculatus*. Other flowers of similar size such as
448 *Onobrychis viciifolia* or *Trifolium pratense* perform significantly worse. These
449 flowers are much harder to distinguish from the background. It is evident that
450 distinguishability (mainly driven by contrast) is the fourth main factor which
451 determines the prediction performance of a network for a particular flower.

452 When taking a closer look at the results, a substantial portion of mispredic-
453 tions that negatively influences the model performance scores such as mAP and

454 F1 score is not fatal. This includes for example false positives that are in fact
455 missing annotations in the ground truth such as the examples in table 3. False
456 positive predictions of flowers that are on the verge of fading fall under this
457 category as well. The mispredictions caused by the confusion between single
458 flowers and inflorescences of *Lotus corniculatus* as described in section 4.2.1 are
459 not severe either. If such mispredictions were ignored, the performance scores
460 would be better.

461 These mispredictions exemplify the challenges that exist for the training data
462 collection. Even when being able to directly compare the image on the tablet
463 to the flowers on site, it is sometimes not clear how to annotate a flower. *Lotus*
464 *corniculatus* is a good example. They are often arranged as inflorescences. It is
465 not uncommon however that there are single flowers that do not belong to the
466 same inflorescence. Since it is often not possible to distinguish the single flow-
467 ers within an inflorescence, the whole inflorescence is annotated as one flower
468 instance. Unfortunately there are border cases in which a single flower very
469 close to another inflorescence is annotated as a separate instance in the ground
470 truth but the prediction algorithm includes that flower in the inflorescence and
471 predicts only one bounding box. This results in false negative predictions for
472 the single flowers very close to the collection as the examples in image c) in
473 table 3 show. The opposite case that multiple single flowers are predicted sepa-
474 rately while they are annotated as an inflorescence with a single bounding box is
475 common as well. The second main problem for *Lotus corniculatus* is that some
476 instances are hardly visible on the images because they are very small. Some-
477 times they are partly hidden by other vegetation and occasionally weak motion
478 blur is present which makes it even harder to distinguish between flower and
479 background. This problem also manifests itself in false positive and false neg-
480 ative predictions. False positive predictions are mainly caused by background

481 areas that look similar to a blurred flower and real flowers which are not present
482 in the ground truth annotations (as in image d)). The false negative predictions
483 are often flowers that are small and hardly distinguishable.

484 As demonstrated on the example of *Lotus corniculatus* in table 4, an average
485 number of flowers per annotation can be calculated from the training data and
486 manually counted data. This value can then be multiplied with the total amount
487 of predictions to get the number of flowers.

488 **5.2 Influences of the Network Configuration and Image** 489 **Resolution**

490 It is advised to scale up all images with objects that are smaller than 40 pixels
491 in diameter by a factor of two in order to improve the performance of a network
492 (Hu & Ramanan 2017). This is the case for the vast majority of flowers dealt
493 with in this study. The Faster R-CNN architecture is not designed to detect
494 very small objects such as flowers of just a few pixels in decimeter (Huang et al.
495 2017; Zhang et al. 2017). Therefore scaling up the images is an appropriate
496 counter measure which helped to improve our results.

497 Data Augmentation options are a convenient way of artificially increasing
498 the amount of training data. One should be careful with applying too many
499 augmentation options. Since the flowers do not span a large number of pixels,
500 they are predicted based on minuscule details. Changing these details too much
501 might be counterproductive. Flips and random box jittering can be applied
502 without hesitation. They do not alter the important details but only the orien-
503 tation or the position of the bounding box. Brightness, contrast and saturation
504 adjustment should be applied moderately. In our experiments the maximal
505 change is a delta of 25 %.

506 **5.3 Practical Considerations**

507 Our main test grassland site was around 30 times 30 meters large. In order to
508 have enough overlapping images to generate an orthomosaic of this area, a drone
509 has to fly over the meadow for about 20 minutes. This means that it is difficult
510 to scale this approach to larger areas. A way of overcoming this problem is to
511 take sample pictures with less or no overlap or at random locations of a larger
512 meadow and therefore omitting the generation of an orthomosaic. Knowing
513 the flight height and the lens angle of the camera, one can calculate the covered
514 area of the image. Running the prediction algorithm on these sample images and
515 extrapolating the numbers of predictions to the size of the whole meadow can
516 still achieve very good results. The advantage over the manual counting flower
517 abundance determination approach is that a much larger sample size can be
518 collected. The effort to collect the vegetation data is smaller and more precise.
519 This enables to spend more time for controlling, extrapolating and analysing
520 the data, which finally earns a better result. What remains to be evaluated is
521 whether the prediction algorithm generates similar results close to the edges of
522 an image compared to the center. The viewing angle changes across an image
523 which changes also the appearance of the imaged objects (Aasen & Bolten 2018;
524 Roth, Aasen, Walter, & Liebisch 2018; Aasen 2016). Consequently there could
525 be a degradation in prediction performance. The orthomosaics are created only
526 from the center regions of the single images.

527 Various metrics are used to describe a model's performance. Precision, re-
528 call, F1 score and mAP all describe a certain aspect of a model's performance.
529 It depends on the application case, which metric is most important. Precision
530 and recall can easily be controlled with the minimum confidence parameter.
531 The higher the minimum confidence parameter of the prediction script is set,
532 the higher the precision gets. Lowering the minimum confidence score increases

533 the recall. For the abundance determination use case as in this study a bal-
534 anced precision/recall ratio is advantageous, because false negatives and false
535 positives are likely to cancel each other out and therefore a good estimate of the
536 abundance can be given. The F1 score is mainly determined by precision and
537 recall. The higher these two values are, the higher is the F1 score. A balanced
538 ratio of precision and recall rewards the score even more. Consequently, the F1
539 score is a good indicator of a model’s performance.

540 We have used independent train, test and validation sets for our evaluations.
541 In the future, the results should be validated in more ways, e.g. by using cross-
542 validation or by testing the models on more unseen test sites as well as including
543 data with different environmental conditions.

544 The method developed in this study opens a wide range of use cases beyond
545 the substitution of manual flower counting. Weed control could be realized in a
546 precision agriculture setting. Detecting invasive neophyte plants in difficult-to-
547 access areas could replace manual checks. The multitemporal abundance maps
548 have the potential to map flowering dynamics quantitatively and spatially assess
549 co-occurrence of different flowers and assess the influence of climate conditions
550 of different years on the abundance. By detecting certain indicator species,
551 conclusions may be drawn about the soil properties. The presence of *Leucan-*
552 *themum vulgare* for example is an indicator for nutrient-poor meadows. In the
553 context of quality assessment of meadows in connection with direct payments by
554 the state, drone usage is imaginable. Apart from flowering plant detection, the
555 method can be applied to other areas such as monitoring of wildlife aggregations
556 as described by Lyons et al. (2019).

557 For some use cases it might be beneficial to have real time detections. The
558 method developed in this study is not designed for that. By using the default
559 configuration of the Faster R-CNN architecture without upscaling the images,

560 the prediction algorithm can be sped up by a factor of four at least. The
561 drawback is that the accuracy is lower. Nevertheless, for some use cases this
562 might be acceptable. Using a more light weight object detection network design
563 such as the SSD architecture (Liu et al. 2016) can deliver further speed ups.
564 However, the accuracy is expected to be lower than with Faster R-CNN.

565 More training data would have been beneficial to better train the underrep-
566 resented flowers and catch flowers during their entire phenology. Unfortunately,
567 this was not possible due to the failure of the initially used drone. However,
568 with the now designed framework new training data can be created and pooled
569 with the current training data to expand the training dataset and allow better
570 predictions in the future. The suite of tools developed in this study is easy to
571 install and can be applied to any sort of object detection problem on aerial im-
572 ages. The time consuming task of training data collection by annotating aerial
573 images can be carried out on the FieldAnnotator application for Android or
574 with the widely used LabelMe application for desktop operating systems. The
575 script that copies annotations onto overlapping images can be a powerful way
576 of increasing the amount of training data without major efforts.

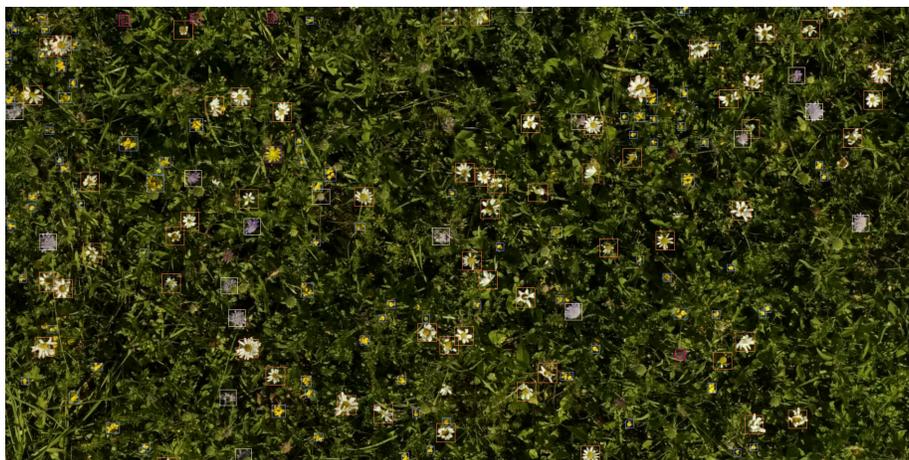


Figure 6: Typical prediction example.

577 **6 Acknowledgements**

578 We thank Bettina Keller and Pascal Kipf who helped assessing the data in
579 the field. Thanks also to Alexander Indermaur, who assisted the drone flights.
580 Special thanks also to Norbert Kirchgessner who organized the computer setup,
581 Lukas Roth who assembled the drone of ETH Zürich and Achim Walter who
582 hosted Johannes Gallmann in his group during the work.

583 **7 Author's contributions**

584 Helge Aasen, Johannes Gallmann, Matthias Albrecht and Beatrice Schüpbach
585 conceived the ideas and designed the methodology; Johannes Gallmann, to-
586 gether with Katja Jacot from Agroscope, collected the data; Jonas Winizki
587 was responsible for the drone flights and Agisoft image processing; Johannes
588 Gallmann analyzed the data, programmed the software and led the writing of
589 the manuscript. All authors contributed critically to the drafts and gave final
590 approval for publication.

591 **8 Data Availability**

592 All code as well as documentation of the code and tutorial videos how to ap-
593 ply the code to new use cases are available on Github ([https://github.com/](https://github.com/tschutli/Phenotator-Toolbox)
594 [tschutli/Phenotator-Toolbox](https://github.com/tschutli/Phenotator-Toolbox)). All our data including images and trained
595 networks can be found at [https://datadryad.org/stash/share/0501neE_](https://datadryad.org/stash/share/0501neE_h0QvhFOBR73XEU1HD0ZChlwYdxKbA0yH1I)
596 [h0QvhFOBR73XEU1HD0ZChlwYdxKbA0yH1I](https://datadryad.org/stash/share/0501neE_h0QvhFOBR73XEU1HD0ZChlwYdxKbA0yH1I).

597 **References**

- 598 Aasen, H. (2016). Influence of the viewing geometry on hyperspectral data
599 retrieved from uav snapshot cameras. *ISPRS Annals of Photogramme-*
600 *try, Remote Sensing & Spatial Information Sciences*, 3(7). doi:10.5194/
601 isprs-annals-III-7-257-2016.
- 602 Aasen, H. & Bolten, A. (2018). Multi-temporal high-resolution imaging spec-
603 troscopy with hyperspectral 2d imagers—from theory to application. *Remote*
604 *sensing of environment*, 205, 374–389. doi:10.1016/j.rse.2017.10.043.
- 605 Aasen, H., Honkavaara, E., Lucieer, A., & Zarco-Tejada, P.J. (2018). Quantita-
606 tive remote sensing at ultra-high resolution with uav spectroscopy: A review
607 of sensor technology, measurement procedures, and data correction workflows.
608 *Remote Sensing*, 10(7). doi:10.3390/rs10071091.
- 609 Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z. et al. (2015). Ten-
610 sorFlow: Large-scale machine learning on heterogeneous systems. Software
611 available from tensorflow.org.
- 612 Abdel-Rahman, E.M., Makori, D.M., Landmann, T., Piironen, R., Gasim, S.
613 et al. (2015). The utility of aisa eagle hyperspectral data and random forest
614 classifier for flower mapping. *Remote Sensing*, 7(10), 13298–13318. doi:10.
615 3390/rs71013298.
- 616 Agisoft, L. (2019). Agisoft metashape user manual. *Professional Edition, Ver-*
617 *sion 1.5*, 1, 71.
- 618 Anderson, K. & Gaston, K.J. (2013). Lightweight unmanned aerial vehicles
619 will revolutionize spatial ecology. *Frontiers in Ecology and the Environment*,
620 11(3), 138–146. doi:10.1890/120150.

- 621 Andrew, M.E. & Shephard, J.M. (2017). Semi-automated detection of eagle
622 nests: an application of very high-resolution image data and advanced image
623 analyses to wildlife surveys. *Remote Sensing in Ecology and Conservation*,
624 3(2), 66–80. doi:10.1002/rse2.38.
- 625 Campbell, T. & Fearn, P. (2018). Simple remote sensing detection of corymbia
626 calophylla flowers using common 3 –band imaging sensors. *Remote Sensing*
627 *Applications: Society and Environment*, 11, 51 – 63. doi:10.1016/j.rsase.2018.
628 04.009.
- 629 Carl Zeiss AG (2017). Zeiss Batis 1.8/85, Technische Daten/Technical
630 Specifications. [https://www.zeiss.com/content/dam/camera-lenses/
631 files/service/download-center/datasheets/batis-lenses/
632 datasheet-zeiss-batis-1885.pdf](https://www.zeiss.com/content/dam/camera-lenses/files/service/download-center/datasheets/batis-lenses/datasheet-zeiss-batis-1885.pdf). Accessed: 2020-01-24.
- 633 Carlet, J. & Abayowa, B. (2017). Fast vehicle detection in aerial imagery. *CoRR*,
634 abs/1709.08666.
- 635 Chen, B., Jin, Y., & Brown, P. (2019). An enhanced bloom index for quan-
636 tifying floral phenology using multi-scale remote sensing observations. *IS-*
637 *PRS Journal of Photogrammetry and Remote Sensing*, 156, 108 – 120. doi:
638 10.1016/j.isprsjprs.2019.08.006.
- 639 Christiansen, F., Sironi, M., Moore, M.J., Di Martino, M., Ricciardi, M. et al.
640 (2019). Estimating body mass of free-living whales using aerial photogramme-
641 try and 3d volumetrics. *Methods in Ecology and Evolution*, 10(12), 2034–2044.
642 doi:10.1111/2041-210X.13298.
- 643 Copting GmbH (2017). TransformerUAV. [https://www.copting.de/index.
644 php/produktuebersicht/uav-copter-drohnen/transformer-uav](https://www.copting.de/index.php/produktuebersicht/uav-copter-drohnen/transformer-uav). Ac-
645 cessed: 2020-01-24.

646 de Sá, N.C., Castro, P., Carvalho, S., Marchante, E., López-Núñez, F.A. et al.
647 (2018). Mapping the flowering of an invasive plant using unmanned aerial
648 vehicles: Is there potential for biocontrol monitoring? *Frontiers in Plant*
649 *Science*, 9, 293. doi:10.3389/fpls.2018.00293.

650 Eikelboom, J.A.J., Wind, J., van de Ven, E., Kenana, L.M., Schroder, B. et al.
651 (2019). Improving the precision and accuracy of animal population estimates
652 with aerial image object detection. *Methods in Ecology and Evolution*, 10(11),
653 1875–1887. doi:10.1111/2041-210X.13277.

654 Everingham, M. & Winn, J. (2011). The pascal visual object classes challenge
655 2012 (voc2012) development kit. *Pattern Analysis, Statistical Modelling and*
656 *Computational Learning, Tech. Rep.*

657 Gallai, N., Salles, J.M., Settele, J., & Vaissière, B.E. (2009). Economic valuation
658 of the vulnerability of world agriculture confronted with pollinator decline.
659 *Ecological Economics*, 68(3), 810 – 821. doi:10.1016/j.ecolecon.2008.06.014.

660 Getzin, S., Wiegand, K., & Schöning, I. (2012). Assessing biodiversity in forests
661 using very high-resolution images and unmanned aerial vehicles. *Methods in*
662 *Ecology and Evolution*, 3(2), 397–404. doi:10.1111/j.2041-210X.2011.00158.x.

663 Hallmann, C., Sorg, M., Jongejans, E., Siepel, H., Hofland, N. et al. (2017).
664 More than 75 percent decline over 27 years in total flying insect biomass in
665 protected areas. *PLoS ONE*, 12, 1–21. doi:10.1371/journal.pone.0185809.

666 Harwin, S. & Lucieer, A. (2012). Assessing the accuracy of georeferenced point
667 clouds produced via multi-view stereopsis from unmanned aerial vehicle (uav)
668 imagery. *Remote Sensing*, 4(6), 1573–1599. doi:10.3390/rs4061573.

669 Hill, D.J., Tarasoff, C., Whitworth, G.E., Baron, J., Bradshaw, J.L. et al. (2017).
670 Utility of unmanned aerial vehicles for mapping invasive plant species: a case

- 671 study on yellow flag iris (*iris pseudacorus* l.). *International Journal of Remote*
672 *Sensing*, 38(8-10), 2083–2105. doi:10.1080/01431161.2016.1264030.
- 673 Hollings, T., Burgman, M., van Andel, M., Gilbert, M., Robinson, T. et al.
674 (2018). How do you find the green sheep? a critical review of the use of
675 remotely sensed imagery to detect and count animals. *Methods in Ecology*
676 *and Evolution*, 9(4), 881–892. doi:10.1111/2041-210X.12973.
- 677 Horton, R., Cano, E., Bulanon, D., & Fallahi, E. (2017). Peach flower monitoring
678 using aerial multispectral imaging. *Journal of Imaging*, 3(1). doi:10.3390/
679 jimaging3010002.
- 680 Hu, P. & Ramanan, D. (2017). Finding tiny faces. *2017 IEEE Conference on*
681 *Computer Vision and Pattern Recognition (CVPR)*, 1522–1530. doi:10.1109/
682 CVPR.2017.166.
- 683 Huang, J., Rathod, V., Sun, C., Zhu, M., Korattikara, A. et al. (2017).
684 Speed/accuracy trade-offs for modern convolutional object detectors. *Pro-*
685 *ceedings of the IEEE conference on computer vision and pattern recognition*,
686 7310–7311. doi:10.1109/CVPR.2017.351.
- 687 Kattenborn, T., Eichel, J., & Fassnacht, F.E. (2019). Convolutional neural
688 networks enable efficient, accurate and fine-grained segmentation of plant
689 species and communities from high-resolution uav imagery. *Scientific Reports*,
690 9(1), 2045–2322. doi:10.1038/s41598-019-53797-9.
- 691 Landmann, T., Piironen, R., Makori, D.M., Abdel-Rahman, E.M., Makau, S.
692 et al. (2015). Application of hyperspectral remote sensing for flower mapping
693 in african savannas. *Remote Sensing of Environment*, 166, 50 – 60. doi:
694 <https://doi.org/10.1016/j.rse.2015.06.006>.
- 695 Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S. et al. (2016). Ssd:

696 Single shot multibox detector. *European conference on computer vision*, 21–
697 37. Springer. doi:10.1007/978-3-319-46448-0_2.

698 Lyons, M.B., Brandis, K.J., Murray, N.J., Wilshire, J.H., McCann, J.A. et al.
699 (2019). Monitoring large and complex wildlife aggregations with drones.
700 *Methods in Ecology and Evolution*, 10(7), 1024–1035. doi:10.1111/2041-210X.
701 13194.

702 Martin, F.M., Müllerová, J., Borgniet, L., Dommange, F., Breton, V. et al.
703 (2018). Using single- and multi-date uav and satellite imagery to accurately
704 monitor invasive knotweed species. *Remote Sensing*, 10(10). doi:10.3390/
705 rs10101662.

706 Müllerová, J., Brůna, J., Bartaloš, T., Dvořák, P., Vítková, M. et al. (2017).
707 Timing is important: Unmanned aircraft vs. satellite imagery in plant inva-
708 sion monitoring. *Frontiers in Plant Science*, 8, 887. doi:10.3389/fpls.2017.
709 00887.

710 Nvidia Corporation (2016). NVIDIA GeForce GTX 1080 User Guide. [https://](https://www.nvidia.com/content/force-gtx/GTX_1080_User_Guide.pdf)
711 www.nvidia.com/content/force-gtx/GTX_1080_User_Guide.pdf. Ac-
712 cessed: 2020-02-03.

713 Osco, L.P., dos Santos de Arruda, M., Junior, J.M., da Silva, N.B., Ramos,
714 A.P.M. et al. (2020). A convolutional neural network approach for counting
715 and geolocating citrus-trees in uav multispectral imagery. *ISPRS Journal of*
716 *Photogrammetry and Remote Sensing*, 160, 97 – 106. doi:10.1016/j.isprsjprs.
717 2019.12.010.

718 Ozge Unel, F., Ozkalayci, B.O., & Cigla, C. (2019). The power of tiling for small
719 object detection. *Proceedings of the IEEE Conference on Computer Vision*
720 *and Pattern Recognition Workshops*.

- 721 Pajares, G. (2015). Overview and current status of remote sensing applications
722 based on unmanned aerial vehicles (uavs). *Photogrammetric Engineering &*
723 *Remote Sensing*, 81(4), 281–330. doi:10.14358/PERS.81.4.281.
- 724 Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster r-cnn: Towards real-
725 time object detection with region proposal networks. *Advances in neural*
726 *information processing systems*, 91–99. doi:10.1109/TPAMI.2016.2577031.
- 727 Rey, N., Volpi, M., Joost, S., & Tuia, D. (2017). Detecting animals in african
728 savanna with uavs and the crowds. *Remote Sensing of Environment*, 200, 341
729 – 351. doi:10.1016/j.rse.2017.08.026.
- 730 Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks
731 for biomedical image segmentation. N. Navab, J. Hornegger, W.M. Wells,
732 & A.F. Frangi, editors, *Medical Image Computing and Computer-Assisted*
733 *Intervention – MICCAI 2015*, 234–241. Springer International Publishing,
734 Cham. doi:10.1038/s41592-018-0261-2.
- 735 Roth, L., Aasen, H., Walter, A., & Liebisch, F. (2018). Extracting leaf area
736 index using viewing geometry effects—a new perspective on high-resolution
737 unmanned aerial system photography. *ISPRS journal of photogrammetry and*
738 *remote sensing*, 141, 161–175. doi:10.1016/j.isprsjprs.2018.04.012.
- 739 Roth, L., Hund, A., & Aasen, H. (2018). Phenofly planning tool: flight planning
740 for high-resolution optical remote sensing with unmanned areal systems. *Plant*
741 *methods*, 14(1), 116. doi:10.1186/s13007-018-0376-6.
- 742 Sanchez-Azofeifa, A., Antonio Guzmán, J., Campos, C.A., Castro, S., Garcia-
743 Millan, V. et al. (2017). Twenty-first century remote sensing technologies
744 are revolutionizing the study of tropical forests. *Biotropica*, 49(5), 604–619.
745 doi:10.1111/btp.12454.

- 746 Sony Corporation (2015). Sony ILCE-7RM2 User Manual. <https://www.sony.com/electronics/support/res/manuals/W001/W0014549M.pdf>. Accessed: 2020-01-24.
- 747
748
- 749 SZ DJI Technology Co., Ltd. (2018). DJI Matrice 600 Pro User Manual. https://dl.djicdn.com/downloads/m600%20pro/20180417/Matrice_600_Pro_User_Manual_v1.0_EN.pdf. Accessed: 2020-01-24.
- 750
751
- 752 Ullman, S. (1979). The interpretation of structure from motion. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 203(1153), 405–426. doi:10.1098/rspb.1979.0006.
- 753
754
- 755 Wada, K. (2016). LabelMe: Image Polygonal Annotation with Python. <https://github.com/wkentaro/labelme>. Accessed: 2020-01-24.
- 756
- 757 Zhang, S., Zhu, X., Lei, Z., Shi, H., Wang, X. et al. (2017). S³fd: Single shot scale-invariant face detector. *2017 IEEE International Conference on Computer Vision (ICCV)*, 192–201. doi:10.1109/ICCV.2017.30.
- 758
759

	Original Image	Upscaled Image						Various Scales	Image
			Yes	Yes	Yes	Yes	Yes	Yes	Yes
Train Overlap			Yes	Yes		Yes	Yes	Yes	Yes
Data Augmentation						Yes	Yes		
Anchor Size 128						Yes			
Stride 16							Yes		
Multiple Scales									Yes
Precision	81.7 %	82.1 %	82.3 %	84.7 %	85.2 %	85.2 %	86.1 %	85.7 %	73.4 %
Recall	79.5 %	83.9 %	81.8 %	84.3 %	83.8 %	83.6 %	81.7 %	76.6 %	70.9 %
F1 Score	0.806	0.830	0.820	0.845	0.845	0.844	0.838	0.809	0.721
mAP	0.575	0.668	0.634	0.686	0.680	0.671	0.619	0.588	0.485

Table 6: Performance of various parameter configurations.

760 A Supplementary Materials

761 A.1 Supplementary Results

762 A.1.1 Faster R-CNN Parameter Tuning

763 Table 6 shows the effects of certain parameter choices for the Faster R-CNN
764 architecture. For each column in table 6, a network is trained and evaluated on
765 the same data. All results are weighted average values across all flower species
766 meaning that all TP, FP and FN values are summed up across all flower species
767 and then the precision, recall and F1 score are calculated using these summed
768 up values. All predictions having a confidence score below 0.2 are ignored for
769 the evaluation. Having a limited amount of data available, the decision is made
770 to use 70 % of the images of each flight for training, 20 % for testing and 10 %
771 for validating. This splitting strategy ensures that the majority of the data is
772 used for training but still large enough portions are available for validation and
773 testing.

774 As the first column of table 6 shows, using the original image size mainly
775 hurts the recall metric. The low recall value indicates that many flowers are not
776 detected by the trained network. Upscaling the images mitigates this problem
777 as can be seen from the other columns of table 6. In all configurations marked
778 with *Upscaled Image*, the image data is tiled into 450 times 450 pixel tiles. Each
779 such tile is then scaled up to a 900 times 900 pixel tile. The drawback is that

780 this slows down the training and prediction process.

781 The effect of adding a train overlap is not clear. While the recall, F1 score
782 and mAP slightly drop compared to the configuration with only upscaled images
783 (column 3 vs. column 2 in table 6), in combination with data augmentation
784 options, the performance of these three metrics is better with a train overlap
785 (column 5 vs column 4). The opposite happens with the precision metric.

786 Using data augmentation techniques and a train overlap within the training
787 pipeline results in the best performance in terms of recall, F1 score and mAP.
788 The precision metric is only insignificantly lower than in other configurations.
789 In all other experiments the configuration with data augmentation and train
790 overlap is used.

791 The fourth last column contains the results with the base anchor size set to
792 128 pixels instead of default value of 256. Recall, F1 score and mAP are slightly
793 worse than with a base anchor size of 256 pixels. The third last column shows
794 the performance with the `first_stage_features_stride`, `height_stride` and
795 `width_stride` parameters set to 16. The `first_stage_features_stride` de-
796 fines the output stride of the extracted region proposal network feature map.
797 A bigger `first_stage_features_stride` value has the consequence that the
798 region proposal network of the Faster R-CNN architecture outputs a feature
799 map with a lower resolution. The `height_stride` and the `width_stride` vari-
800 ables control the distance in pixels of two consecutive anchors. An anchor is a
801 location within the image from which various sizes of possible bounding boxes
802 to be evaluated are spanned. Having a large distance between two such anchors
803 might cause the network to miss flowers that are placed in between two such
804 anchors. These changes are suggested in general by Ren, He, Girshick, & Sun
805 (2015) and specifically for small objects by Zhang et al. (2017). The results
806 show that the performance is worse than with these values set to the minimum

807 of 8. The recall value is notably lower than in the configuration with the stride
808 values set to 8. This is an indication that some flowers are missed by the pre-
809 diction network due to the lower coverage of anchors and the lower resolution
810 of the feature map.

811 Training networks with images of multiple scales did not result in promising
812 prediction performance. Neither did applying a Lab or HSV color space trans-
813 formation to the images improve the detection results (experiment results not
814 included in table 6).

815 A.1.2 Predictions on Simulated Resolutions

816 The higher the drone can fly the more area can be covered with a single drone
817 flight. Table 7 demonstrates the effect of decreasing ground resolution on an
818 example excerpt of an aerial image containing a *Leucanthemum vulgare* flower
819 and a *Lotus corniculatus* inflorescence. Figure 7 and 8 illustrate the effect of
820 decreasing ground resolution on the F1 score and the mAP respectively. Both
821 figures show that down to a ground resolution of 5 mm per pixel, there is just
822 a marginal decrease in prediction performance. Further decreasing the ground
823 resolution to 10 mm and 20 mm per pixel however has noticeable negative
824 effects on the model’s performance. As expected, the performance of small
825 flowers such as *Lotus corniculatus* decreases disproportionately because at a
826 certain ground resolution they simply get indistinguishable. The average size
827 of a *Lotus corniculatus* flower is around 16 mm. The performance of larger

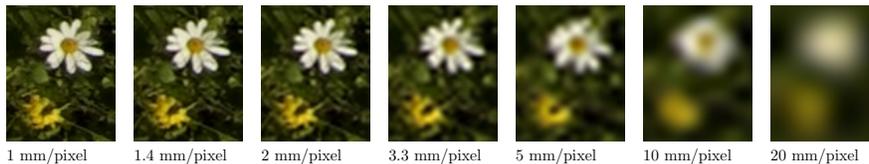


Table 7: Resolution degradation on an excerpt of an aerial image.

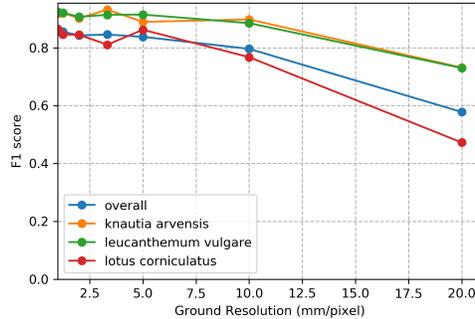


Figure 7: Evolution of F1 score over various simulated ground resolutions.

828 flowers such as *Leucanthemum vulgare* (40 mm) and *Knautia arvensis* (34 mm)
 829 degrades notably slower. The graphs for the precision and recall metrics are
 830 omitted since the trends are equivalent to the trends of the F1 score and the
 831 mAP metric.

832 Each training, test and validation image is first scaled down to the desired
 833 ground resolution and then scaled up again. After upscaling, all datasets have
 834 the same ground sampling distance (pixel size) as the original images again.
 835 This ensures that the flower’s sizes (in image pixels) are large enough to be
 836 detectable by the faster R-CNN network architecture and prevents performance
 837 losses caused by this problem as described by Hu & Ramanan (2017). For each
 838 ground resolution a network is trained and evaluated with the processed training
 839 images.

840 A.2 FieldAnnotator Android Application

841 For the annotations, an Android tablet application called FieldAnnotator has
 842 been developed. It can be downloaded from the Google Playstore. The advan-
 843 tage of being able to make the annotations on a tablet is that they can be made
 844 directly in the field. This is necessary because some flowers can be very hard

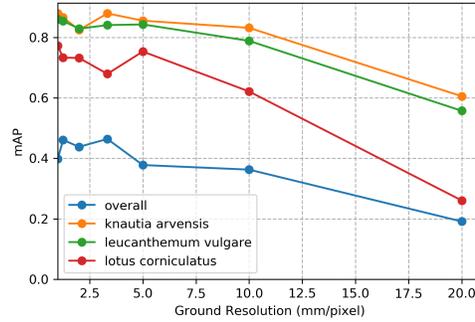


Figure 8: Evolution of mAP over various simulated ground resolutions.

845 to distinguish in the image alone. If one can compare the image to the actual
 846 flowers on site, the quality of the training data can be improved and it is made
 847 sure that the number of false annotations is minimized.

848 A screenshot of the main window of the annotation application can be seen in
 849 figure 9. The output folder created by the image preprocessing tool mentioned
 850 in the previous subsection can be copied onto the Android tablet and imported
 851 into the annotation application (1). The orthophoto is then displayed to the
 852 user. If the geo information is included in the metadata file, the user's GPS
 853 location is indicated on top of the image (2). This helps the user navigate
 854 through the field. The displayed image can be zoomed up to a level where the
 855 individual pixels are visible. If the user clicks on any location in the image,
 856 the annotation settings on the right appear. The user can select the type of
 857 flower from the list (3). Next to each flower in the list, the number of already
 858 recorded occurrences are indicated in brackets. If necessary, the position of an
 859 annotation can be fine tuned by using the four buttons on the bottom (4). The
 860 user can dismiss the annotation in processing by clicking on 'cancel/delete' (5)
 861 or save it by clicking on 'save' (6). Optionally, instead of a point annotation, a
 862 polygon can be drawn around a region. To do so, the switch at the bottom (7)

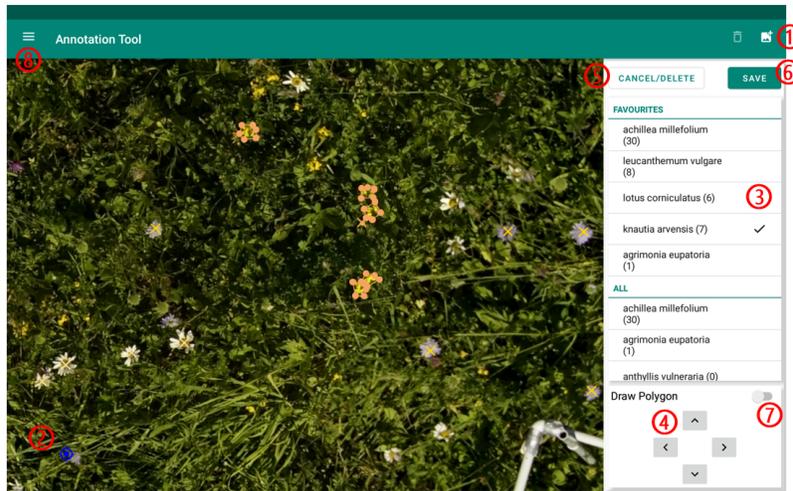


Figure 9: Screenshot of the main window of the FieldAnnotator application for Android.

863 has to be activated. Already saved annotations can also be edited or deleted
 864 again by simply clicking on them inside the image. By clicking on the menu
 865 button on the upper left, the settings screen can be opened. There the user can
 866 edit the list of flowers either manually or by importing a predefined list from a
 867 csv file. Also an export of the flower list to a csv file is possible. Furthermore,
 868 some zoom settings such as the maximal zoom level or the zoom level at which
 869 the annotations should be displayed to the user can be set. The application
 870 continuously saves the annotations to a json file in the project folder. The
 871 application is programmed in Kotlin.