

# Automated Pollen Detection with an Affordable Technology

Nam Cao<sup>1</sup>, Matthias Meyer<sup>2</sup>, Lothar Thiele<sup>2</sup>, Olga Saukh<sup>1,3</sup>

<sup>1</sup>Institute of Technical Informatics, Graz University of Technology, Austria

<sup>2</sup>Computer Engineering and Networks Lab, ETH Zurich, Switzerland

<sup>3</sup>Complexity Science Hub Vienna, Austria

{cao.nam, saukh}@tugraz.at, {matthias.meyer, thiele}@tik.ee.ethz.ch

## Abstract

Airborne pollen cause seasonal allergies and the number of people affected increases yearly due to global warming and urbanization. Governmental pollen sensing stations are sampling traps which require manual pollen identification and counting by trained personnel in the lab. In the past years, a number of researchers and startups started working towards automated pollen measurements by exploring a wide range of techniques. Many solutions reported in the literature are expensive or work for a limited number of pollen species. In this paper, we present the design of a prototype of an *automated and affordable* pollen detection device built from off-the-shelf components. The design consists of three subsystems operating in the field and communicating the data to the backend server: (1) a particle trap with automatic filtering, (2) a particle concentration subsystem, and (3) a digital transmitted light microscope with layer-wise focus. The prototype shows effective particle gathering, filtering and concentration in a tiny-sized area. As a result, we reduce particle loss and improve image quality taken by the optical system when searching and focusing on pollen grains. The test results show that our device achieves high efficiency with up to 150 l/min air flow rates, evaluates over 90 % of captured pollen grains, and achieves 1 h measurement delay on average (2 h at maximum). The prototype collects raw time-stamped microscopic images of pollen with 5-60 depth layers per sample depending on the number of objects contained in one sample. All images are transmitted to the backend server where we run a pollen detection algorithm to extract individual pollen grains from every image. We achieve 0.90 average precision and F1-score of 0.88 when detecting pollen in the images of individual layers taken in the field. Our prototype successfully operated in the wild for 115 days between April and August 2019, and shows high stability under a wide range of varying weather conditions, little maintenance need and low device-to-device variation.

## 1 Introduction

Airborne pollen is one of the most common triggers of seasonal allergies. It affects over 24 % of Europeans [52] and over 20 % Americans [55], with children quota up to 40 % [52]. Global warming impacts pollination and causes higher pollen concentrations and longer pollen seasons. Growing carbon dioxide (CO<sub>2</sub>) concentrations encourage plants to produce more pollen [8]. High concentrations of pollen in the air can spread for several kilometers, even indoors if structures are not sealed [55], and trigger allergic reactions. Allergies occur when the body's immune system overreacts to an otherwise harmless substance. This can cause mild annoyances like hives or itchy eyes, or life-threatening issues such as anaphylaxis, where blood pressure plummets, airways start swelling and shut. Some types of pollen can travel deep into the lungs and cause irritation, even for people who do not have allergies [55]. Moreover, rising CO<sub>2</sub> concentrations increase the allergenic peptides on pollen, which are known to trigger the body's immune system. More peptides on a pollen grain increase the severity of the allergy [55]. Thus, with global warming allergic diseases are increasing with unprecedented complexity and severity.

There are three big peaks in pollen production throughout the year. In Europe and in the US, *trees* like oak, ash, and birch see pollen surges in spring. Pollen from *grasses* peaks over summer, and *ragweed* pollen spikes in fall [55]. The effect of CO<sub>2</sub> on pollen can be observed at small scales as well: grasses and ragweed plants increase their pollen production in response to local exposure to CO<sub>2</sub>, e.g., from the exhaust of cars along a highway [8, 55]. Urban areas are particularly affected due to binding of urban pollutants with allergenic proteins [22]. Moreover, pollen concentrations and the pollen types in the air decrease gradually with increasing height [58]. The outdoor pollen concentration at low level is six to seven times higher than at high altitude [15]. This motivates the need for making timely, accurate and spatially-resolved information about pollen distribution available to aerobiologists, health experts, and the general public. It can be used to setup early warning systems and to give personal advice to allergy sufferers.

Conventional pollen measurement devices are particle traps that require manual pollen identification and counting on the daily or weekly basis by trained personnel. For

example, the European pollen measurement network comprises over 400 pollen traps [2] operated by governmental institutions in different countries. In the past few years several automatic pollen measurement systems were announced [42, 18, 52, 7, 56]. These systems explore a wide range of technologies including elastic scattering, microscopy and holography. A few of these devices are extremely expensive [42, 7, 52] (up to 100'000 \$) while other use affordable methods based on holographic imaging [56]. The latter work with low air flow rates (from 13 l/min [56] to 40 l/min [52]). The technology described in this paper is on the opposite side of the design spectrum. We make a set of ultimately different design choices. We use high air flow rates (up to 150 l/min) for a better overall system efficiency, a unique particle gathering and concentration solutions, and a low-cost transmitted light microscope to generate layered images of pollen grains at different focus depths.

**Challenges.** Automatic pollen sensing requires efficient solutions with respect to both hardware and software. *On the hardware side*, it is essential to choose a technology which can measure relevant morphological characteristics of pollen grains [6], such as their size, shape, pores and furrows, surface patterning, and surface sculpture [26] with sufficient level of detail. There are over 50'000 species of trees in the world, but luckily only about 100 cause pollen allergies [24]. Nevertheless, many pollen types have similar features, *e.g.*, size, which exhibits notable natural variability [11, 18]. This causes significant challenge for experts to make correct pollen identification. Even manual identification requires specialized training and is supported by illustrated handbooks [27] and online databases [26, 43]. The literature covers a broad range of pollen sensing systems [42, 18, 52] exploring different technologies to successfully capture specific morphological features of pollen grains crucial for pollen identification. An additional challenge is the overall system stability and a steady performance under varying weather conditions. *On the software side*, deep neural networks are becoming more and more popular in the context of automatic pollen detection and identification [31, 51]. However, many recent works focus on performing these tasks on high-quality and clean microscopic images found in online databases, *e.g.*, [28, 33, 38]. These results can not be directly applied to images taken with a low-cost microscope in the field. An additional challenge of training these machine learning algorithms is that the available expert knowledge is noisy, limited and scarce. This is especially problematic in the face of high variability of pollen grains and the fact that recent works claim their deep models significantly outperform experts on the pollen identification task [53].

**Contributions and road-map.** Our aim is to design an *automated and affordable* pollen sensing system based on an efficient sampling and detection of pollen grains in microscopic images. The developed system puts value on inexpensive technologies and subsystems, unique design choices, and the use of machine learning algorithms. In this paper, we describe our design in detail, perform rigorous subsystem and end-to-end system tests, and present the first field

performance results from the past pollen season from April to August 2019. This paper makes the following contributions:

- We provide details on the prototype pollen detection system in Sec. 3. The prototype features several unique design choices different from the state-of-the-art automatic pollen measurement stations. In particular, we rely on an *omnidirectional cyclone* to improve system's efficiency and achieve air flow rates of up to 150 l/min. *Particle concentration* solution and *optimized control* allow increasing the number of pollen grains falling into a tiny-sized area visible through the microscope, yet minimizing the processing time of every sample.
- Sec. 4 presents a *pollen detection algorithm* designed to process noisy field images. It achieves the average precision of 0.90 and the F1-score of 0.88 when tested on over 5'000 images taken in the wild.
- We present the results of operating the prototype in the field in Sec. 5. The system shows robust and stable performance under a range of varying weather conditions, low device-to-device variation and little maintenance need.

We share gathered microscopic images of pollen gains through a public platform<sup>1</sup> to involve interested domain experts to help us with pollen identification and further automation of the system. Sec. 2 surveys the state-of-the-art in pollen sensing and Sec. 6 concludes this paper with a discussion of the trade-offs we make and outlines our future research directions.

## 2 Related Work

In Europe pollen is monitored by a network of manually operated 400+ pollen traps [2]. The most widely used pollen trap in the world is the volumetric *Hirst-type*. It exists in different versions such as VPPS1000, VPPS2000, VPPS2010 [5]. Non-volumetric air sampling devices such as Durham [40], Cour [10, 41] and Rotorod [23] samplers are other popular pollen traps used for short-term exposure measurements. *Ionic spore trap* makes use of electrostatics to attract particles. Since pollen is rarely present in the ambient air, the airflow generated by a trap impacts the overall system's efficiency. The airflow rate of a *Hirst-type trap* is fairly low with about 10-20 l/min (liters per minute), whereas an *Ionic spore trap* [4] achieves up to 660 l/min. Every 1-7 days gathered samples are processed in the laboratory by trained aerobiologists. Manual sample inspection limits the number of covered locations and the inherent resolution and timeliness of pollen exposure models and forecasts.

The efforts to automate pollen sample processing date back to 1990s. Many proposed solutions are based on image recognition [54, 21, 12, 28, 33, 38], whereas other use biomolecular analysis, *i.e.*, DNA [32] or chemical identification [49, 13]. The performance of all these methods is high, but they all require manual treatment and careful preparation of samples. Thus, these semi-automatic methods experience significant processing delays and require humans to handle parts of the process.

<sup>1</sup><https://zenodo.org/record/3572653>

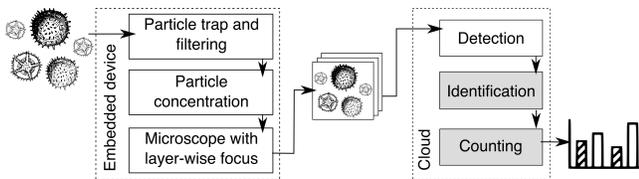


Figure 1. System architecture composed of an embedded devices with three subsystems (particle trap, concentrator and digital microscope), and cloud service for data storage and processing (including pollen detection).

In recent years, progress has been made in testing, evaluation and development of instruments and methods for a full automation of pollen sensing. Amongst the proposed systems one can distinguish between devices using microscopy and automatic image recognition such as *Hund BAA500* (2015) [42] and [7], systems based on air flow cytometry such as *Yamatronics KH-3000* (2007) [30], detectors which extensively leverage holography, e.g., *Swisens Poleno Basic* (2018) [52] and [56], and optical detectors that make use of the fluorescence of bioparticles like *Plair PA-300* (2016) [18] and *Swisens Poleno Standard* (2018) [52]. The systems based on high-quality microscopes and fluorescence are expensive and so far have not received wide acceptance. Promising newly emerged automatic pollen sensing systems [52, 56, 57] extensively leverage holography. In contrast, the technology described in this paper (1) works with significantly higher air flow rates, (2) generates up to 4 times better resolved images with an off-the-shelf transmitted light microscope, and (3) provides 5-60 focus depth layers through every sample to help achieve high accuracy of pollen identification.

Automatic pollen identification in microscopic images has a long history. Identification quality heavily depends on the,—possibly manual,—sample preprocessing and on the type and quality of the technology used to record microscopic images. Early methods worked on images taken with expensive scanning electron microscopes and performed statistical classification of texture parameters achieving 94.3 % accuracy on six classes [34]. Later works explored pollen identification in light microscope images using neural networks with careful manual feature selection and showed 87 %–97 % accuracy on three to five classes [60, 47]. All these methods work with carefully prepared images centered at a pollen grain [19]. Pollen detection within an image was investigated in [9] and reported a 90 % detection rate on three classes. The latest detection results using holographic images report up to 98 % detection precision on five classes [56] and 87 % classification accuracy on 15 classes [53].

In recent years object and pattern detection based on deep neural networks has become state-of-the-art [35, 25, 37]. The description and comparison of the latest architectures can be found in [17]. The authors show that the performance of the recently suggested detectors vary in terms of both their accuracy, training and inference speed. The YOLOv3 detector [45], adapted in this work, achieves a good trade-off between accuracy and execution.

### 3 Automated Pollen Sensing Device

The overall architecture of our pollen sensing system is presented in Fig. 1 and comprises an embedded device and

a set of data processing algorithms running in the cloud. Pollen sensing device consists of three subsystems depicted in Fig. 2: (1) a particle trap with automatic filtering (in Fig. 2b), (2) a particle concentrator (in Fig. 2c), and (3) a digital microscope with layer-wise focus (in Fig. 2a ⑤). Layered images of pollen grains are transmitted to the cloud infrastructure for permanent storage and further processing. This section describes all subsystems and evaluates their performance in isolation. We use *dry bee flower pollen* and *flour flecks* in extensive subsystem tests for convenience reasons, even though the properties of fresh pollen contained in the ambient air are different: (i) both dry bee flower pollen and the flour appear to be more sticky, and (ii) the internal structures of dry pollen are often damaged, which makes their detection and identification more challenging than in the case of fresh pollen. A very preliminary version of the system with notable differences is described in [14].

#### 3.1 Cyclone-based Trap and Particle Filter

Common pollen grain size falls in the range 10-100  $\mu\text{m}$  with the smallest grain shape being 5  $\mu\text{m}$   $\times$  2.4  $\mu\text{m}$  and the largest over 200  $\mu\text{m}$  in diameter [6]. The mass of particles typically varies between 25  $\mu\text{g}$  and 75  $\mu\text{g}$ . We chose a cyclone-based particle gathering and filtering design (Fig. 2a ⑨ and Fig. 2b) since the method can filter particles with a diameter of down to 0.5  $\mu\text{m}$  [59]. Conventional cyclone-based particle traps are directional: they have one inlet and are equipped with a vane to rotate and suck in the surrounding air from the wind direction [1]. This makes the pollen trap sample from one direction at a time. In contrast, our design adapts an omnidirectional cyclone by implementing six equally spaced inlets to cover all 360 degrees. To further increase the system efficiency, we follow the ideas in [59]: The inlets are split in 3 pairs with each pair comprising a main and an auxiliary inlet placed on the opposite sides of the cyclone body (Fig. 2a ⑨). An auxiliary inlet is located 5 mm higher than the main one to stimulate the air flow. A brushless fan for pumping the air is placed on top of the cyclone. Fig. 2b shows the 3D design of the cyclone which we print using the PLA plastic. The system is capable of generating air flows up to 150 l/min. This makes our design more efficient than the majority of existing pollen sampling systems. The efficiency of the system is determined by the amount of air it processes [16]. Higher air flow rates through a measurement device result in larger volumes of air being processed and, thus, in a more accurate estimation of pollen concentration in the environment.

**Cyclone efficiency.** We test the filtering system with flour flecks of similar size as pollen grains. A small portion of flour (10 g) is manually fed into one of the cyclone’s inlets. We gather the filtered air at the cyclone outlet and evaluate the total weight of flecks contained in the air, which the cyclone system fails to capture. The experiments are carried out at four different air flow rates of 100, 125, 140 and 150 l/min. The test reveals that with the highest air flow rate of 150 l/min only <1 % of tiny flecks was contained in the output air.

We also evaluated the percentage of the flour flacks that get stuck on the walls of the inner cyclone body, which

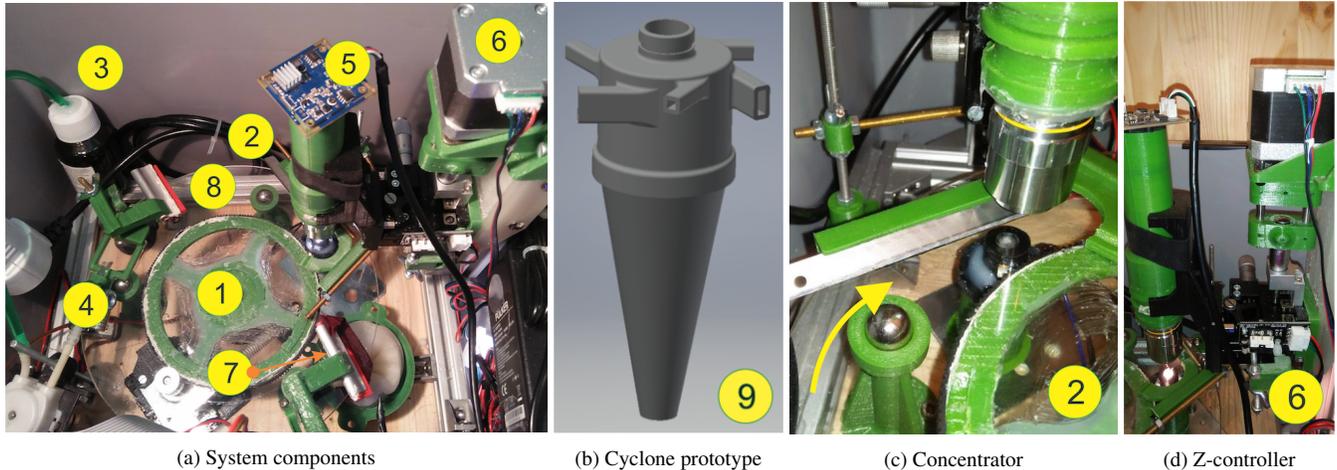


Figure 2. Subsystem components: ① rotating glass plate, ② particle concentrator, ③ glycerine jar, ④ peristaltic pump, ⑤ digital microscope, ⑥ Z-controller, ⑦ waste remover, ⑧ particle drop down area, ⑨ cyclone prototype.

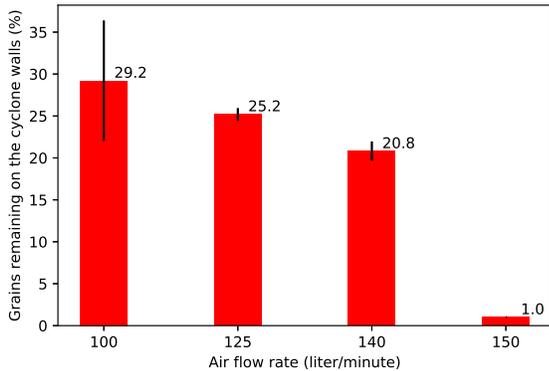


Figure 3. Evaluation of the cyclone efficiency: with air flow rate  $\geq 150$  l/min almost no particles get stuck to the cyclone walls.

we remove with a brush after the experiment when opening the cyclone. Fig. 3 summarizes the results of five experiments conducted with air flow settings in the range from 100 to 150 l/min under the room temperature conditions (22 °C, 40 % relative humidity). It suggests that there is almost no flour ( $< 1$  %) attached to the walls when the cyclone speed approaches 150 l/min. Note that the stickiness of the flour is weather-dependent, but is generally higher than of the fresh pollen. We can thus expect a better field performance than obtained in this test.

**Optimal air flow rate.** We choose a 10 W brushless DC motor to ventilate the air. This fan creates the maximum throughput of 150 l/min with the 6 inlets. At this air flow rate we keep the system at a relatively low power consumption. On the one hand, the fan alone consumes only about 3 W for the following reason: As the environment inside the sealed box is close to vacuum, the fan can not keep up with the nominal power even when the power supply comes close to the nominal voltage of 12 V. On the other hand, further increase of the throughput would require larger air inlets and thus increase the chance that unexpected objects beside pollen,—such as dust and small insects,—get inside the system. As

a consequence, larger objects may obstruct the pollen and impact the accuracy of the imaging subsystem.

### 3.2 Particle Concentration

After the particles are captured by the cyclone, they drop down driven by gravity and the air force inside the cyclone. A glass plate surface of 200 mm in diameter is placed at the bottom of the box to catch the filtered particles as can be seen in Fig. 2a ①. We create a thin layer of glycerine to keep the particles on the surface of the glass plate. Moreover, glycerine helps to considerably improve the quality of the microscopic images: Pollen grains are being drawn into the liquid, and the light refraction helps to illuminate all their sides. We rely on glycerine in our design due to its attractive chemical properties of being transparent, non-evaporating, dense and exhibiting sufficient thermal stability for outdoor operation as will be evaluated later in the paper.

Filtered particles are spread over a large area comparable in size to the bottom opening of the cyclone which makes 15 mm diameter area in our design (Fig. 2a ⑧). Processing the whole area with a digital microscope would require the imaging system to make sequential sweeps over the area to obtain pictures of all captured pollen grains. We solve the problem by designing a concentrator which moves the particles to a tiny-sized area to mitigate the need of shifting the microscope. This reduces the delay of processing a glycerine sample. The concentrator leverages a paper cutting blade attached by magnets to a plastic 3D printed holding frame as depicted in Fig. 2a ② and Fig. 2c. Two cylindrical magnets of 4mm in diameter are integrated into the ends of the holding frame to keep the blade in the frame. The holding frame makes a 35-40° angle with the glass surface. The magnetic force makes the lower edge of the blade tightly touch the glass surface while the upper edge is kept inside the frame. When the glass plate rotates the particles move in the clockwise direction and form a *narrow high-density track* along the blade. An example of the track is depicted in Fig. 4. The microscope in Fig. 2a ⑤ is aligned with the blade edge so that a small part of the blade is visible in the image (see

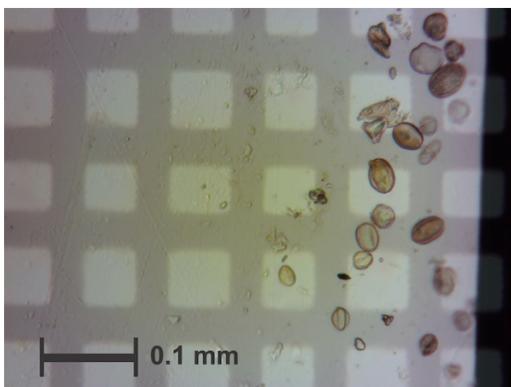


Figure 4. A microscopic view of a sample high-density track built by the concentration subsystem. Grid cell dimensions are 0.1 mm  $\times$  0.1 mm.

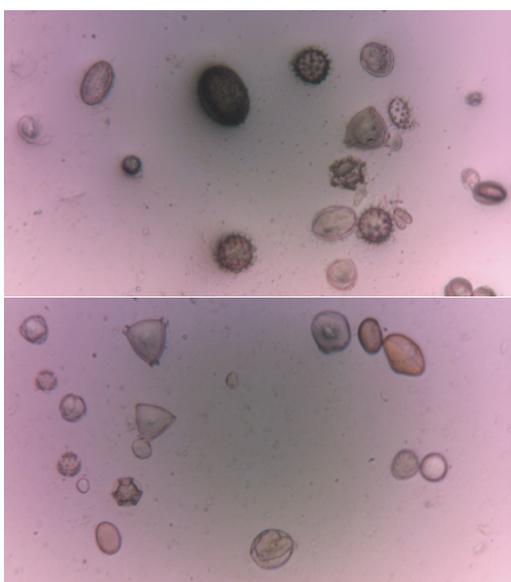


Figure 5. Sample microscopic images with insufficient, 1500 steps/droplet (top) and sufficient, 1000 steps/droplet (bottom) amount of glycerine.

Fig. 4). This design decision achieves the following goals: (1) the initial autofocus calibration is easy and fast due to the presence of an object to focus on instead of having a solely transparent background; (2) since the edge of the blade always appears in the screen, adjusting the focus at runtime to compensate for the glass plate irregularities and possible inclination is also fast; (3) a track misalignment can quickly be detected while the system is running in the wild.

**Stable glycerine supply.** A peristaltic pump is used to apply and spread the glycerine on the glass surface at periodic time intervals through a 4 mm tube shown in Fig. 2a ④. The glycerine is stretched by a piece of rubber to create a thin sticky layer to capture filtered particles. It is well known that the glycerine facial tension depends on the temperature and decreases with high temperatures. On the one hand, at high temperatures ( $\geq 40^\circ\text{C}$ ) the same amount of glycerine covers a larger area, yet the glycerine layer gets thinner and causes degradation of the image quality. In Fig. 5, we show an example of a microscopic image with sufficient (bottom) and insufficient (top) amounts of glycerine tested with dry flower

pollen. Note that the appearance of back and intransparent pollen grains that cast shadows are signs of insufficient glycerine supply. On the other hand, glycerine surplus makes the glycerine layer more fluid and results in a larger width of the concentration track. These two factors hamper the work of the imaging system and increase the time of processing a sample. As a consequence, the end-to-end measurement delay is increased. To keep the performance of the glycerine layer stable,—in terms of both the narrowest possible track and a sufficient amount of glycerine for a high image quality,—we carry out experiments described in the following section to find the optimal amount of glycerine.

**Concentration track size.** The size of the concentration track directly affects the efficiency of the overall measurement system. We strive for the narrowest possible track to ensure all captured pollen fall inside the visible microscopic image area. In the following experiment we explore the dependency of the track size on the ambient temperature. We operate the system for an hour to ensure a smooth and stable glycerine layer. We feed dry pollen into the cyclone inlet. As expected, the concentrator creates a narrow track. We use the following method to measure the size of the track: The 10x lens is used to observe the microscopic area of 0.6 mm  $\times$  0.4 mm containing the track. The glass plate is rotated at a slow speed of 3 sec/step, with the step length of 0.25 mm measured along the track line. The image and the corresponding temperature are saved every 20 rotations. For measuring the track size, we use a 0.1 mm  $\times$  0.1 mm measurement grid printed on a transparent film as a ruler. The track size images are blended with the reference grid image and used to estimate the track size at various temperatures. Fig. 4 shows a sample track including the measurement grid. In the image, the width of the track is around 0.3 mm with a few outliers. The track size depends on the amount of glycerine and the ambient temperature. In Fig. 6 we show an empirically established dependency. By experimenting with different amounts of glycerine we found that the concentration track becomes stable,—in terms of its size and size variation as well as a low number of outliers ( $<1\%$ ),—with the glycerine supply frequency in the range from 1000 to 1500 steps/droplet. Fig. 6 shows that with the temperature increase, the concentration track size gets somewhat larger. This is because the facial tension of glycerine decreases, which leads to the same amount of glycerine covering a larger area and thus the pollen grains are spread over a larger track. We fix the glycerine supply at 1000 steps/droplet for the whole range of working temperatures. When we reduce the amount of glycerine, the glycerine layer becomes too thin and leads to the degradation of quality of the microscopic images. With 1'000 steps/droplet we estimate the amount of glycerine required for the whole pollen season from February to the end of September is 400 ml. The glycerine is stored in a jar shown in Fig. 2a ③.

### 3.3 Digital Microscope

Our imaging system relies on a digital 8.0 mega pixel camera module equipped with a IMX179 Sony sensor [3] with the maximum resolution of 3'264  $\times$  2'448 pixels. The camera is fixed on the one end of the 115 mm long pipe

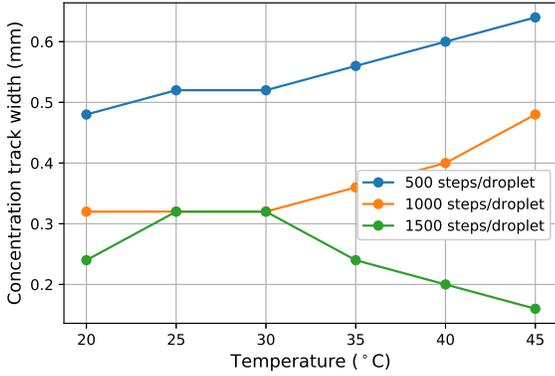


Figure 6. Particle concentration track width depends on temperature

shown in Fig. 2d. The other end of the pipe contains an objective with a Bresser planachromatic 10x lens. The parallel light source of a warm light LED (230-240 Lumen, color temperature 3'000-3'500 K) is placed under the glass plate. The light beam goes through a sample towards the microscope. A stepper motor slowly rotates the glass plate with the glycerine-soaked pollen grains on its surface. To minimize the impact of the inertial force which causes vibration waves on the glycerine surface and makes the grains move after each rotation step, the plate is accelerated and decelerated at a slow rate of 100 steps/sec<sup>2</sup> and 50 steps/sec<sup>2</sup> respectively. After each rotation step a glycerine sample in the current view of a tiny-sized window of the digital microscope is processed by executing a sequence of focusing steps for taking pollen images at different focus depths. The images are both stored locally and are being transmitted to the cloud for permanent storage and further processing. After a sample has been processed, it is removed from the surface with a rubber shown in Fig. 2a ⑦. The waste is gathered in a tray under the glass plate. Processed pollen particles stay bounded to glycerine which prevents them from taking off. Below we present the sample processing algorithm executed by the imaging subsystem. Delay optimization caused by the algorithm results in a variable sample processing time. We thus implement a measurement timestamp reconstruction routine to compensate for the processing delays and determine an accurate time of every measurement.

A glycerine sample may contain pollen grains at different depths. We use a self-designed slider to change the focus when processing a glycerine sample. The slider comprises a stepper and a manual LGX40-C micro stage. The two components are attached to each other by a 3D frame printed from the PLA plastic (Fig. 2a ⑥ and Fig. 2d). The combination helps to increase the shifting resolution up to 0.15  $\mu\text{m}/\text{step}$  in the 1/16 step mode of the stepper controller A4988. Processing a sample means shooting a sequence of microscopic images,—referred to as *layers*,—with different focus setting while sequentially changing the focus depth. Since many depth layers do not contain any pollen grain in focus, we use the following algorithm to minimize the number of steps we make through a sample in order to reduce the time required to process it. The algorithm is listed in Alg. 1. It estimates the *sharpness* of each depth layer over

### Algorithm 1 Shooting layered microscopic images

**Input:** Sharpness of an empty image  $\sigma_{\min}$ ,  
operation interval of the z-controller  $[0, p_{\max}]$   
**Output:** A set of images of the depth layers  
Set z-controller to its initial position  $p = 0$   
 $\sigma_L^{-1} = \sigma_{\max} = \sigma_{\min}$   
**while**  $p < p_{\max}$  **do**  
  Compute Laplacian variance  $\sigma_L$  of the current layer  
  **if**  $p > 0$  and  $|\sigma_L^{-1} - \sigma_L| \geq \epsilon$  **then**  
    **if**  $\sigma_L > \sigma_{\max}$  **then**  
       $\sigma_{\max} = \sigma_L$   
    **end if**  
    Record image of the current depth layer  
    Advance with a small step  $p += 5 \mu\text{m}$   
  **else**  
    Advance with a big step  $p += 10 \mu\text{m}$   
     $\sigma_{\max} = \sigma_L$   
  **end if**  
   $\sigma_L^{-1} = \sigma_L$   
**end while**

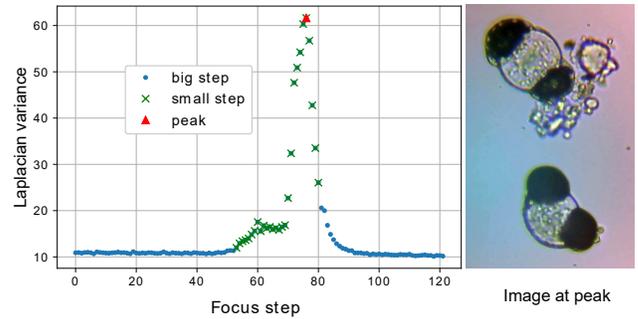


Figure 7. Sample execution of the layer-wise focus algorithm (1).

a sequence of stepper steps  $p \in [0, p_{\max}]$  by computing the *Laplacian variance*  $\sigma_L$  of each layer. If the sharpness difference  $|\sigma_L^{-1} - \sigma_L|$  between two consecutive layers is low, a larger stepper step of 10  $\mu\text{m}/\text{step}$  can safely be taken without missing pollen grains. A step-by-step algorithm execution is exemplified in Fig. 7. Once the sharpness difference is higher than a threshold  $\epsilon$ , the stepper makes small steps of 5  $\mu\text{m}/\text{step}$  to take depth-layered images through a potentially found pollen grain. As the sharpness of consecutive layers decreases for the subsequent stepper positions, the algorithm switches back to making bigger steps. The process stops once the step count reaches the upper bound  $p_{\max}$ .

**Sample processing speed and delay.** Let a particle exit the cyclone at time  $t$ . The sample containing the particle gets processed once it appears in a tiny-sized window inspected by the microscope. For this reason, we have (1) to minimize the measurement delay by reducing the distance the sample has to travel to get processed, and (2) to reconstruct the correct measurement time  $t$ . In our design, the distance between the cyclone dropdown area and the microscope is 64 mm, which results in an average delay of about 80 min. Additionally, we propose a correction routine presented in Alg. 2 to compensate for the variable delay caused by the op-

---

**Algorithm 2** Reconstructing measurement timestamps

---

**Input:** An array  $\{C_{180}, \dots, C_{i-1}, C_i, C_{i+1}, \dots, C_1\}$  of processing times of the past 180 samples

**Output:** Corrected timestamp  $t$  of the last processed sample

Get current time:  $t_{now} = \text{now}()$

Average measurement delay:  $C_{mean} = \sum_{i=1..160} C_i$

Max overestimation:  $C_o = \sum_{i=140..160} C_i$

Max underestimation:  $C_u = \sum_{i=160..180} C_i$

$t = (t_{now} - C_{mean}) - \{-C_o, +C_u\}$

---

Pollen type	Family	Concentration in 1 m <sup>3</sup> of air
Vitis	Vitaceae	42.97
Oleaceae	Olive	615.76
Poaceae	Grasses	767.33
Urticaceae	Nettle	227.67
Pinaceae	Pine	44.58
Castanea	Chestnut	998.74
Quercus	Oak	133.15

Table 1. The average concentration of pollen grains in the ambient air in the Braga region, Portugal during the corresponding pollen seasons in 1999 (source: [46]).

timizations introduced in Alg. 1. We consider every sample takes  $C_i$  time units to be processed. The processing time depends on the number of steps that need to be taken through the sample by Alg. 1 and the number of depth-layered images to be stored and transmitted from the camera to the embedded device. Capturing and transmitting an image is an expensive operation in Alg. 1 and takes 30% of the total time of the focusing algorithm. The rest of the time is spent on moving the stepper, focusing, resizing the image and computing the Laplacian variance over it. The size of the microscopic window is 0.6 mm × 0.4 mm and the rotation step of the glass table is 0.4 mm/step. This takes 64 mm / 0.4 mm/step = 160 rotation steps on average for a sample to be processed. Under- and overestimation of the processing time is limited by the diameter of the cyclone dropdown area, which makes 15 mm in our design. We store an array  $\{C_{180}, \dots, C_{i-1}, C_i, C_{i+1}, \dots, C_1\}$  of processing times of the last 180 samples. The average delay is then the sum  $C_{mean} = \sum_{i=1..160} C_i$  over the last 160 samples and the error is limited by  $C_{mean} + \{-C_o, +C_u\}$ .

Alg. 1 takes images of the layers that correspond to the sharpness peaks including several adjacent layers before and after the peak. The more prominent the peak is the more adjacent images are taken. The location of the sharpness peak is highlighted in the picture with a red triangle. The total number of recorded images depends on the number of peaks in the image and on their prominence. The empirical average time of processing a sample in our system is around 30 seconds. If a sample contains many objects, the focusing algorithm takes longer to execute and yields a higher number of depth-layered images. However, the concentration of pollen grains in the ambient air is low. Table 1 presents several examples of the average concentration of pollen grains in the ambient air in the Braga region, Portugal [46] during the respective pollen seasons. In our system around 50 %

of images are empty. Our empirical mean sample processing delay  $C_{mean}$  reported by Alg. 2 is 1 h and the maximum delay  $C_{mean} + C_u$  is 2 h. These results are in line with the performance of the other automatic pollen sensing devices [42, 29].

**Quality of layered images.** Evaluating the quality of the obtained microscopic images is a difficult task. On the one hand, the estimation of the image quality by experts severely relies on their expertise, experience, and qualification. This results in generally noisy and subjective judgements. Given a large variety and local variability of pollen, no expert is trained in identifying all possible types. On the other hand, recently emerged pollen measurement systems claim being able to achieve a significantly better pollen identification accuracy than manual evaluation by the experts [53]. In this section, we evaluate the quality of images by analyzing the presence of the pollen distinguishing features in our microscopic images.

The morphological description of pollen grains presented in [6] lists eight visible features of pollen that help to identify and classify them in the microscopic images. These features include: pollen units, polarity, symmetry, shape, size, apertures, sub-divisions of the pollen surface and exine ornamentation. Scientific literature suggest that transmitted light microscopes are well suited to capture deep and surface textures of pollen but provide limited information on shape [50]. Our system yields up to 60 layer-wise images of each sample with 5 μm resolution including the top view and the equatorial view of every pollen grain. The important visual features are contained in the top and equatorial views of each pollen grain, yet with a varying level of sharpness and detail. Fig. 8 presents the top and the equatorial views of sample pollen grains taken by our system. The resolution of images is adequate to be processed by a machine learning algorithm and to be evaluated by human eyes. According to the literature [6], the size of pollen grains from the smallest (5 μm × 2.4 μm) to the largest (over 200 μm × 200 μm) corresponds to 28 × 14 pixels and over 1'143 × 1'143 pixels in our images respectively. The size of the most airborne pollen in our region ranges from 10 μm to 100 μm in diameter which corresponds to the diameter range from 58 to 580 pixels in our images respectively. Our image resolution is 0.17 μm/pixel. The samples shown in Fig. 8 were gathered during our field tests from April to August 2019. In each pair of images, the left image shows the top view and the right image shows the equatorial view of a sample grain.

### 3.4 Embedded Control and Cloud Storage

Raspberry Pi 3.0 running Raspbian Embedded Linux controls the measurement loop and connects to the Internet over a mobile 4G network to communicate the data to the cloud. The system provides continuous operation and delivers pollen measurements with high temporal resolution of up to 1 sample every 30 s on average.

Our prototype consumes 6.0 W of power in total with the major consumption being spent on powering the fan (3.0 W), the Raspberry Pi (2.0 W), and the motor (1.0 W). Although in most cases the system is permanently powered, low power consumption allows running the system in the wild for a few

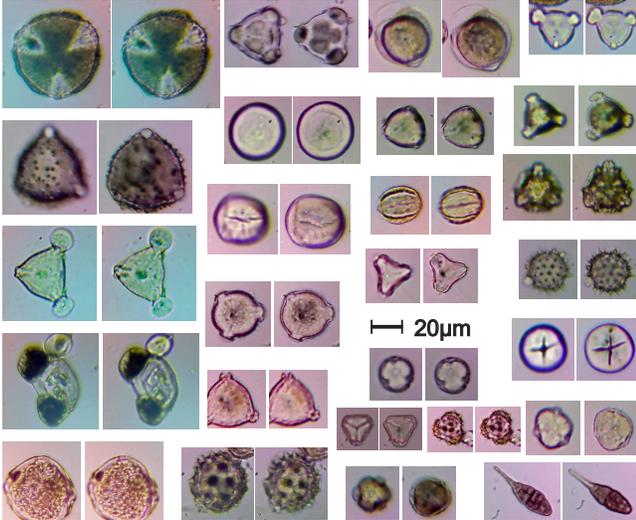


Figure 8. Pollen and spore top view and equatorial views. The samples were gathered in the field between April and August 2019.

Name	Training	Validation	Testing
Date and split	May 26 (80%)	May 26 (20%)	May 23 (100%)
#samples	456	115	506
#total depth layers	4234	1167	3769
#layers with pollen	2302	688	1982
#total pollen	4569	1376	3114

Table 2. Statistics on the annotated data set (ground truth).

days on a moderate-sized battery. The system requires maintenance and cleaning up the waste at the end of the measurement season.

The total material cost of the system is 700 \$, with 150 \$ being spent on the 10x lens and 70 \$ on the transmitted light microscope. The dimensions of the measurement prototype are 30 cm×40 cm×22 cm and the total weight is 8 kg. This makes the system fairly portable. Further miniaturization is limited by the size, geometry and efficiency of the cyclone subsystem.

We store gathered microscopic images in the cloud for further analysis. The average transmission speed is 1 GB/day. The data, including the timestamped depth-layered images along with the temperature and humidity sensor readings are communicated to the cloud once per minute. We rely on the Google Drive service and the dedicated APIs to push the data to the server. The experienced delay is insignificant despite relying on the free version of the service. It takes about 30 seconds to transmit and store 100 images on Google Drive, which is adequate for our application. After pushing collected data to the cloud, the local images on the Raspberry Pi are erased. The next section describes pollen detection in the gathered field images.

## 4 Pollen Detection

We augmented our pollen measurement system with a server-based pollen detection software based on an existing, accurate and efficient object detector YOLOv3 [45]. Our

pollen detection model determines the bounding box around each pollen grain given an image of a single depth layer. To train a pollen detection model, we manually annotate a subset of the acquired data. Since manual image annotation comes with a significant time and labor overhead and requires expert knowledge, the depth layers from only two days are annotated (see Table 2). The days are chosen from the second half of May, which is the end of the tree pollen season and the beginning of the grass pollen season. Therefore, both pollen types may occur in gathered samples. We chose two sunny days that are in close succession to minimize the impact of seasonal changes, ensure similar pollen types, and similar weather conditions for training and testing. Significantly fewer pollen is contained in the ambient air on a rainy day. Our labelled data set contains 1’077 processed samples comprising 9’170 depth layers and containing 9’059 images of pollen grains to be detected. Note that 46% of layers do not contain any pollen. The data set covers over 15 different types of pollen. 80% of one day are used for training and the other 20% are used to select the best performing model during training. All samples from the other day are held out for testing. The allocation of training and testing sets on two separate days is chosen to provide a meaningful measure of detection accuracy since it enables the capture of the daytime-dependant pollen distribution. An iterative annotation process with a human-in-the-loop similar to [39] has been used in which false positives and false negatives are presented to the annotator for label verification – an important step to improve annotation quality. However, a perfect alignment of ground truth bounding box and pollen is not guaranteed for all images and thus the bounding boxes should be considered imprecise.

For implementing the detection system we use an existing PyTorch [44] YOLOv3 implementation [36], pre-trained on the ImageNet [48] and adapted to work with our data set. Image pre-processing includes resampling the image to a squared form and a random horizontal flipping. Similar to the original YOLOv3 implementation, multi-scale training is used and the image size is randomly varied within the range [320, 512] with a step size of 32. For further implementation details and the values of all hyperparameters an interested reader is referred to our publicly available source code<sup>2</sup>.

**Data post-processing.** In order to evaluate the performance of our pollen detection model we run a sequence of post-processing steps on the predicted bounding boxes and on the manually labelled data which we treat as the ground truth: (1) bounding box merging, and (2) removal of bounding boxes with an area smaller than a threshold. The first post-processing step is necessary, since the detection model may predict multiple bounding boxes per pollen grain. We perform non-maximum suppression to transform multiple (possibly imprecise) prediction bounding boxes into one single bounding box. An *intersection over union* (IoU) factor of 0.5 is used as a threshold to determine if two bounding boxes have to be merged or not. In the second post-processing step we remove the potential pollen grains that are too small (smaller than 0.5% of the image area). The reason behind

<sup>2</sup><https://github.com/osaukh/pollenpub>

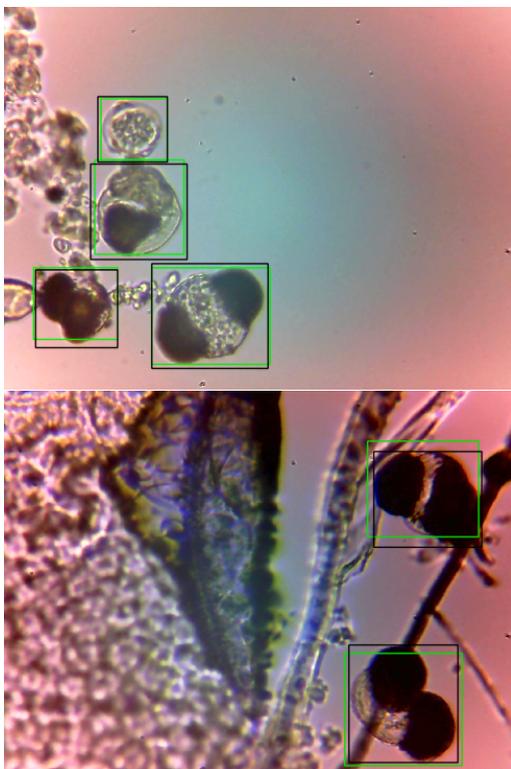


Figure 9. Examples of automatically predicted bounding boxes (green rectangles) and manually annotated ground truth (black rectangles). A huge fraction of the bottom image is covered by unwanted, additional material. Note the varying precision of the ground truth bounding boxes.

this filtering step is the observation that it becomes hard to identify pollen with the necessary certainty if an object in the picture is small, damaged or partially visible. The thresholds were established empirically based on our training and validation data sets.

**Detection performance results.** Following the pre-processing step, the remaining predicted bounding boxes are used to validate and test the performance of the pollen detection model. Our model yields 0.9 average precision and F1-score of 0.88 on the test data set. If the IoU of a predicted bounding box and the ground truth bounding box is above the threshold of 0.5 we consider the detection to be correct. This threshold has been chosen according to the literature [20] and the fact that we are only interested in the correct detection and not in an accurate mapping of the predicted bounding box to the imprecise ground truth bounding box. Based on these criteria we calculate the average precision to be 90% on the test data set. As discussed in Sec. 2, our detection precision is in line with the state-of-the-art.

## 5 Overall Performance and Field Tests

This section evaluates the overall system performance and summarizes our measurement experiences in the field. We estimate and empirically measure the overall system efficiency in terms of the percentage of captured pollen being processed, and the total measurement delay. Since we aim for a large-scale deployment in the future, we also evaluate the device-to-device variation to ensure measurement con-



Figure 10. Hibiscus flower and its pollen grains under the microscope.

sistency across multiple devices and robustness of our design. Finally, we describe our field experiments from spring and summer 2019 and make selected data available through a public platform to get expert community support with pollen labelling and identification.

### 5.1 End-to-End System Tests

We use *freshly gathered* hibiscus pollen (see Fig. 10) to perform the end-to-end system tests with two identical prototypes of our system. We chose hibiscus, since the grains are sufficiently large (up to 200  $\mu\text{m}$  in diameter) and can thus be easily split, counted and feed into the cyclone inlets. Moreover, hibiscus was still to find in Graz, Austria in September. Although the grains are visible with a bare eye it is hard to determine the exact number of grains in a pile without the help of a microscope.

To test the end-to-end system performance, we feed one hibiscus grain into one of the cyclone inlets every 5 min and wait for the images of the grains to appear on Google Drive. The pictures are then processed by the pollen detection model. In total, we use 18 hibiscus grains in the test. We compare the timestamps and the number of pollen at the system's input and output and summarize the obtained results below.

**Measurement efficiency.** As demonstrated in Sec. 3.3, the maximum obtained particle concentration track size is  $<0.5$  mm (see Fig. 6), which perfectly matches our microscopic window size of  $0.6$  mm  $\times$   $0.4$  mm when using the 10x lens. The percentage of outliers in the concentration subsystem tests is around 1%. Given the particle loss of  $<2\%$  in the cyclone subsystem (including the particle filtering loss of  $<1\%$  and the measurement that up to 1% of particles get stuck inside the cyclone), we expect 97% of pollen grains be correctly captured and counted. Our end-to-end test reveals that 16 out of 18 hibiscus grains were correctly processed and counted by the system. This yields an empirical end-to-end system efficiency of 89%. We also extended the experiments with the dry bee flower pollen to ensure the small-sized grains can be accurately processed by the system as well. Since the commercial bee flower pollen are stored in agglomerates, we can only break these into smaller parts of a few dozen grains each. We feed these into the cyclone and wait until the grains become visible in the microscopic view. We managed to count 20 out of 22 high-density pollen agglomerates and, therefore, infer the 90% efficiency of the system. Both tests with large and small pollen grains show consistent results that confirm the system is 90% efficient.



Figure 11. Deployment on the rooftop of the Information Technology campus at Graz University of Technology, Austria.

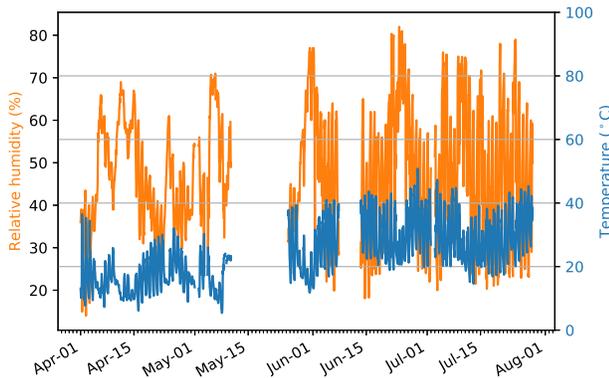


Figure 12. Temperature and humidity variations during the deployed time period. Measurement gaps correspond to system upgrades.

**Measurement delay.** Our system features an average measurement delay of 1 h, which is mostly determined by the length of the path between the cyclone and the imaging subsystem. The measured end-to-end delay in the test with manually fed hibiscus grains appeared to be 75 min on average, with a standard deviation of 6 min.

**Device-to-device variation.** We compare the overall system efficiency and the measured delay when running a sequence of 2 end-to-end tests with two identical pollen sensing systems. All tests were run sequentially using hibiscus pollen grains. We achieve 90 % and 86 % average efficiency with a standard deviation of 6 % and 6 % respectively. The obtained delays are 75 min and 78 min on average and the standard deviations are 6 min and 5 min respectively.

## 5.2 Field Performance

From April to August 2019 our prototype was running on the rooftop of the Information Technology campus at Graz University of Technology, Austria. The deployed prototype system is shown in Fig. 11. This time period covers a wide range of temperatures (10 – 50 °C) and humidities (10 – 90 % relative humidity) measured inside the box and plotted in Fig. 12. Measurement gaps in the plot correspond to system upgrades. The main improvements were related to the imaging subsystem, and the addition of the peristaltic pump for fine-grained control of the glycerine supply. This was a necessary step due to high temperatures in the beginning

of summer that led to a considerable drop of image quality, as was explained in detail in Sec. 3.2. There were no unexpected system failures during the field operation.

Our gathered data set contains 390'000 raw images with 3'300 of new images being generated per day on average. Each processed sample consists of 10 depth layers on average. On a sample day in May, we could count over ~300 captured pollen grains of at least 15 different types recorded as ~3'000 layers (see Table 2). Naturally, the number of captured pollen is highly weather-dependent. Our data set also contains rainy days with few captured grains. Fig. 9 shows examples of pollen grains from the observed period. Although the system also captures dust and dirt in addition to pollen, the appearance of larger waste, as shown the bottom image in Fig. 9 is rare. The majority of pollen pictures are either empty or indeed contain pollen grains brought by the wind.

## 5.3 Open Design, Detection Model and Data

We aim for the development of an open system by sharing our hardware design and the detection model with the community. We also open source the data set we used to train and test the detection algorithm<sup>3</sup>. Although pollen identification is left for future work, with this paper we hope to get attention of the domain scientists to help us identify and label gathered pollen images. For this reason, we publish a subset of gathered data<sup>4</sup> and look forward to collaboration with experts on this challenging task. The lack of open databases of microscopic images is explicitly pointed out in the recent technology review [29].

## 6 Conclusion and Future Work

Automatic pollen identification is an active field of research due to (1) the increasing severity and complexity of pollen allergies over the past years triggered by the global warming, and (2) a significant sparsity of manually operated pollen traps. This paper presents a prototype of an automated, continuous and affordable pollen sensing system comprising a particle trap, a concentrator, a low-cost transmitted light microscope and a pollen detection algorithm. We describe the function of all subsystems in detail and run exhaustive subsystem tests to show their efficiency despite the low cost of the used components and material. The system was successfully operated in the field during five months covering mainly the tree and grass pollen seasons. Although automatic pollen identification and the prototype performance comparison to conventional pollen traps is left for future work, we evaluate the device-to-device variation to show high consistency of pollen counts across two identical prototypes. Our system design features unique particle filtering and concentration solutions not seen in existing automatic pollen identification systems. These features allow for an automatic in-the-field operation of a low-cost transmitted light microscope and successful detection of pollen grains in noisy microscopic images. As a result of these design choices, the key performance characteristics of the built prototype can be summarized as follows: (1) The

<sup>3</sup><https://github.com/osaukh/pollenpub>

<sup>4</sup><https://zenodo.org/record/3572653>

system is operated at an air flow rate of 150 l/min, which makes it highly efficient compared to the state-of-the-art. (2) We ensure that >90 % of gathered pollen grains appear in the taken microscopic images. (3) The system is fairly lightweight, low-power and low-cost with the total size of 30 cm×40 cm×22 cm, weight of 8 kg, power consumption of 6 W, and the total material cost of 700 \$. (4) The worst-case measurement delay is capped by 2 h, whereas the mean measurement delay due to algorithmic optimization stays around 1 h. We also provide a measurement timestamp reconstruction procedure to provide an accurate measurement time. (5) The pollen detection model optimized for noisy layered images taken with a low-cost transmitted light microscope achieves 0.90 average precision and F1-score of 0.88.

**Future work.** Our future research will focus on an automatic pollen identification. Towards this end, we hope to leverage the achieved maturity of the prototype and an open data set gathered in the field to assign correct pollen labels with help of the expert community of aerobiologists. To minimize the experts' time and effort, we plan to augment a labelling software with unsupervised clustering and interactive learning to improve efficiency of the labelling process. Additionally, we plan to scale up our deployment to ten stations deployed in Graz before the next pollen season begins.

## Acknowledgments

This research is partly funded by the Vietnam International Education Cooperation Department, Ministry of Education and Training, Project 911.

## 7 References

- [1] Burkard air sampler Hirst type. <http://bit.ly/2kKqoCv>.
- [2] European aeroallergen network. <https://www.ean-net.org>.
- [3] IMX179 Sony sensor. <https://bit.ly/2ND0wWM>.
- [4] Ionic spore trap. <http://bit.ly/2kiHU00>.
- [5] Lanzoni VPPS. <http://lanzoni.it/vpps.html>.
- [6] Morphological characteristics of pollen grains. <https://bit.ly/2Dy3eWq>.
- [7] PollenSense. <https://pollensense.com>.
- [8] J. Albertine, W. Manning, M. DaCosta, M. Muilenberg, K. Stinson, and C. Rogers. Projected carbon dioxide to increase grass pollen and allergen exposure despite higher ozone levels. *PLoS ONE*, 9:e11712, 2014.
- [9] G. Allen, R. Hodgson, S. Marsland, and J. Flenley. Machine vision for automated optical recognition and classification of pollen grains or other singulated microscopic objects. In *15th International Conference on Mechatronics and Machine Vision in Practice, M2VIP'08*, pages 221 – 226, 2009.
- [10] J. Belmonte, M. Canela, and R. A. Guàrdia. Comparison between categorical pollen data obtained by Hirst and Cour sampling methods. *Aerobiologia*, 16(2):177–185, 2000.
- [11] H.-J. Beug. *Leitfaden der Pollenbestimmung für Mitteleuropa und angrenzende Gebiete*. Verlag Friedrich Pfeil, Mnich, 2004.
- [12] A. Boucher, P. J. Hidalgo, M. Thonnat, J. Belmonte, C. Galan, P. Bonton, and R. Tomczak. Development of a semi-automatic system for pollen recognition. *Aerobiologia*, 18(3-4):195–201, 2002.
- [13] J. T. Buters, M. Thibaudon, M. Smith, R. Kennedy, A. Rantio-Lehtimäki, R. Albertini, G. Reese, B. Weber, C. Galan, R. Brandao, C. M. Antunes, S. Jäger, U. Berger, S. Celenk, Ł. Grewling, B. Jackowiak, I. Sauliene, I. Weichenmeier, G. Pusch, H. Sarioglu, M. Ueffing, H. Behrendt, M. Prank, M. Sofiev, and L. Cecchi. Release of Bet v 1 from birch pollen from 5 European countries. Results from the HIALINE study. *Atmospheric Environment*, 55:496–505, 2012.
- [14] N. Cao, O. Saukh, and L. Thiele. An automated real-time and affordable airborne pollen sensing system: Poster abstract. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks, IPSN*, pages 321–322, 2019.
- [15] P. Capone, L. Boccacci, S. D. Renzi, R. Ferrante, M. Brighetti, M. Serra, R. Sisto, A. Pelliccioni, A. Travaglini, and M. D'Ovidio. 878?airborne pollen sampling at two different heights: variation of concentrations in indoor and outdoor environments and implication for occupational health. *Occupational and Environmental Medicine*, 75:A52, 2018.
- [16] J. Chen, Z. A. Jiang, and J. Chen. Effect of inlet air volumetric flow rate on the performance of a two-stage cyclone separator. *ACS Omega*, 3(10):13219–13226, 2018.
- [17] K. Chen, J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, Z. Zhang, D. Cheng, C. Zhu, T. Cheng, Q. Zhao, B. Li, X. Lu, R. Zhu, Y. Wu, J. Dai, J. Wang, J. Shi, W. Ouyang, C. C. Loy, and D. Lin. MMDetection: Open MMLab Detection Toolbox and Benchmark. *arXiv:1906.07155 [cs, eess]*, June 2019.
- [18] B. Crouzy, M. Stella, T. Konzelmann, B. Calpini, and B. Clot. All-optical automatic pollen identification: Towards an operational system. *Atmospheric Environment*, 140:202–212, 2016.
- [19] M. Del Pozo Banos, J. Ticay-Rivas, J. Cabrera-Falcon, J. Arroyo, C. Travieso, L. Sanchez-Chavez, S. Perez, J. Alonso, and M. Ramirez-Bogantes. Image processing for pollen classification. *Biodiversity Enrichment in a Diverse World*, 2012.
- [20] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010.
- [21] I. France, A. W. G. Duller, G. A. T. Duller, and H. F. Lamb. A new approach to automated pollen analysis. *Quaternary Science Reviews*, 19(6):537–546, 2000.
- [22] U. Frank and D. Ernst. Effects of NO<sub>2</sub> and ozone on pollen allergenicity. *Frontiers in Plant Science*, 7(2):2–5, 2016.
- [23] D. A. Frenz. The effect of windspeed on pollen and spore counts collected with the Rotorod Sampler and Burkard spore trap. *Annals of Allergy, Asthma and Immunology*, 85(5):392–394, 2000.
- [24] The Allergy Sufferer's Guide to Trees, Plants and Flowers. <https://bit.ly/2lPoJvz>, 2015.
- [25] R. Girshick. Fast R-CNN. In *The IEEE International Conference on Computer Vision (ICCV)*, Dec. 2015.
- [26] Tiny Grains, Big Data: The Global Pollen Project. <https://bit.ly/2kOS8Wq>, 2017.
- [27] M. Hesse, H. Halbritter, M. Weber, R. Buchner, A. Frosch-Radivo, S. Ulrich, and R. Zetter. *Pollen Terminology. An illustrated handbook*. Springer-Verlag Wien, 2009.
- [28] K. Holt, G. Allen, R. Hodgson, S. Marsland, and J. Flenley. Progress towards an automated trainable pollen location and classifier system for use in the palynology laboratory. *Review of Palaeobotany and Palynology*, 167(3-4):175–183, 2011.
- [29] J. Huffman, A. Perring, N. Savage, B. Clot, B. Crouzy, F. Tummon, O. Shoshanim, B. Damit, J. Schneider, V. Sivaprakasam, M. Zawadowicz, I. Crawford, M. Gallagher, D. Topping, D. Doughty, S. Hill, and Y. Pan. Real-time sensing of bioaerosols: Review and current perspectives. *Aerosol Science and Technology*, 0(0):1–31, 2019.
- [30] S. Kawashima, B. Clot, T. Fujita, Y. Takahashi, and K. Nakamura. An algorithm and a device for counting airborne pollen automatically using laser optics. *Atmospheric Environment*, 41(36):7987–7993, 2007.
- [31] N. Khanzhina, E. Putin, A. Filchenkov, and E. Zamyatina. Pollen grain recognition using convolutional neural network. *European Symposium on Artificial Neural Networks*, (April):25–27, 2018.
- [32] K. Kraaijeveld, L. A. de Weger, M. Ventayol García, H. Buermans, J. Frank, P. S. Hiemstra, and J. T. den Dunnen. Efficient and sensitive identification and quantification of airborne pollen using next-generation DNA sequencing. *Molecular Ecology Resources*, 15(1):8–16, 2015.
- [33] R. Lagerstrom, K. Holt, Y. Arzhaeva, L. Bischof, S. Haberle, F. Hopf, and D. Lovell. Pollen image classification using the classifyfynder system: Algorithm comparison and a case study on New Zealand honey. *Advances in Experimental Medicine and Biology*, 823:207–226, 2015.
- [34] M. Langford, G. Taylor, and J. Flenley. Computerized identification of pollen grains by texture analysis. *Review of Palaeobotany and Palynology*, 64:197–203, 1990.

- [35] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal Loss for Dense Object Detection. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2999–3007, Oct. 2017.
- [36] E. Linder-Norén. Eriklindernoren/PyTorch-YOLOv3, Sept. 2019.
- [37] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. SSD: Single Shot MultiBox Detector. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *Computer Vision – ECCV 2016*, pages 21–37. Springer International Publishing, 2016.
- [38] J. V. Marcos, R. Nava, G. Cristóbal, R. Redondo, B. Escalante-Ramírez, G. Bueno, Ó. Déniz, A. González-Porto, C. Pardo, F. Chung, and T. Rodríguez. Automated pollen identification using microscopic imaging and texture analysis. *Micron*, 68:36–46, 2015.
- [39] M. Meyer, S. Weber, J. Beutel, and L. Thiele. Systematic Identification of External Influences in Multi-Year Microseismic Recordings Using Convolutional Neural Networks. *Earth Surface Dynamics*, 7(1):171–190, 2019.
- [40] Y. Okamoto, S. Horiguchi, H. Yamamoto, S. Yonekura, and T. Hanazawa. Present situation of cedar pollinosis in Japan and its immune responses. *Allergology International*, 58(2):155–162, 2009.
- [41] F. Orlandi, J. Oteros, F. Aguilera, A. Ben Dhiab, M. Msallem, and M. Fornaciari. Design of a downscaling method to estimate continuous data from discrete pollen monitoring in Tunisia. *Environmental Sciences: Processes and Impacts*, 16(7):1716–1725, 2014.
- [42] J. Oteros, G. Pusch, I. Weichenmeier, U. Heimann, R. Möller, S. Röseler, C. Traidl-Hoffmann, C. Schmidt-Weber, and J. Buters. Automatic and online pollen monitoring. *Int. Archives of Allergy and Immunology*, 167(3):158–166, 2015.
- [43] Paldat - palynological database.
- [44] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. Oct. 2017.
- [45] J. Redmon and A. Farhadi. YOLOv3: An Incremental Improvement. 2018.
- [46] H. Ribeiro, M. Cunha, and I. Abreu. Airborne pollen concentration in the region of Braga, Portugal, and its relationship with meteorological parameters. *Aerobiologia*, 19(1):21–27, 2003.
- [47] M. Rodríguez-Damin, E. Cernadas, A. Formella, M. Fernández-Delgado, and P. 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, 2006.
- [48] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [49] F. Schulte, J. Lingott, U. Panne, and J. Kneipp. Chemical characterization and classification of pollen. *Analytical Chemistry*, 80(24):9551–9556, 2008.
- [50] M. Sicaguru, L. Mander, G. Fried, and S. Punyasena. Capturing the surface texture and shape of pollen: A comparison of microscopy techniques. *PLoS one*, 7:e39129, 2012.
- [51] T. Sladevic. The application of convolutional neural network for pollen bearing bee classification. *IEEE Workshop on Advances in Information, Electronic and Electrical Engineering*, 2018.
- [52] Swisens Poleno. <https://bit.ly/2lWkbUo>, 2018.
- [53] Actual results of the first series of Swisens Poleno Bioaerosol Monitor. <https://bit.ly/2nAjnFp>, 2019.
- [54] E. L. Vezey and J. J. Skvarla. Computerized feature analysis of exine sculpture patterns. *Review of Palaeobotany and Palynology*, 64(1-4):187–196, 1990.
- [55] It’s not your imagination. Allergy season gets worse every year. <https://bit.ly/2Kn1YLd>, 2019.
- [56] Y. Wu, A. Calis, Y. Luo, C. Chen, M. Lutton, Y. Rivenson, X. Lin, H. Koydemir, Y. Zhang, H. Wang, Z. Göröcs, and A. Ozcan. Label-free bioaerosol sensing using mobile microscopy and deep learning. *ACS Photonics*, 5(11):4617–4627, 2018.
- [57] Y. Wu, Y. Luo, G. Chaudhari, Y. Rivenson, A. Calis, K. de Haan, and A. Ozcan. Bright-field holography: cross-modality deep learning enables snapshot 3d imaging with bright-field contrast using a single hologram. *Light: Science and Applications*, 8(1), 2019.
- [58] X. Xiao, A. Fu, X. Xie, M. Kang, D. Hu, P. Yang, and Z. Liu. An investigation of airborne allergenic pollen at different heights. *International archives of allergy and immunology*, 160:143–151, 09 2012.
- [59] H. Yoshida, Y. Inada, K. Fukui, and T. Yamamoto. Improvement of gas-cyclone performance by use of local fluid flow control method. *Powder Technology*, 193(1):6–14, 2009.
- [60] Y. Zhang, D. W. Fountain, R. M. Hodgson, J. R. Flenley, and S. Gunetileke. Towards automation of palynology 3: Pollen pattern recognition using Gabor transforms and digital moments. *Journal of Quaternary Science*, 19(8):763–768, 2004.