Building an Aerial–Ground Robotics System for Precision Farming

©SHUTTERSTOCK.COM/ZAPP2PHOTO

An Adaptable Solution

By Alberto Pretto, Stéphanie Aravecchia, Wolfram Burgard, Nived Chebrolu, Christian Dornhege, Tillmann Falck, Freya Fleckenstein, Alessandra Fontenla, Marco Imperoli, Raghav Khanna, Frank Liebisch, Philipp Lottes, Andres Milioto, Daniele Nardi, Sandro Nardi, Johannes Pfeifer, Marija Popović, Ciro Potena, Cédric Pradalier, Elisa Rothacker-Feder, Inkyu Sa, Alexander Schaefer, Roland Siegwart, Cyrill Stachniss, Achim Walter, Wera Winterhalter, Xiaolong Wu, and Juan Nieto

Digital Object Identifier 10.1109/MRA.2020.3012492
Date of current version: 25 August 2020

The application of autonomous robots in agriculture is gaining increasing popularity thanks to the high impact it may have on food security, sustainability, resource-use efficiency, reduction of chemical treatments, and optimization of human effort and yield. With this vision, the Flourish research project aimed to develop an adaptable robotic solution for precision farming that combines the aerial survey capabilities of small autonomous unmanned aerial vehicles (UAVs) with targeted intervention performed by multipurpose unmanned ground vehicles (UGVs). This article presents an overview of the scientific and technological advances and outcomes obtained in the project. We introduce multispectral-perception algorithms and aerial and ground-based systems developed to monitor crop density, weed pressure, and crop nitrogen (N)-nutrition status and to accurately classify and locate weeds. We then introduce the navigation and mapping systems tailored to our robots in the agricultural environment as well as the modules for collaborative mapping. We finally present the ground-intervention hardware, software solutions, and interfaces we implemented and tested in different field conditions and with different crops. We describe a real use case in which a UAV collaborates with a UGV to monitor the field and perform selective spraying without human intervention.

Background

Collaborative aerial- and ground-based robotic systems offer significant benefits to many practical applications, as they can merge the advantages of

multiple heterogeneous platforms. This is especially useful for precision agriculture scenarios, where the areas of interest are usually vast. For example, a UAV allows for rapid inspections of large areas, e.g., mapping weed distributions or crop nutrition-status indicators. This information can then be shared with a UGV, which can perform targeted actions, e.g., selective weed treatment or fertilizer applications on required areas, with relatively high operating times and payload capacities. One of the main objectives of the Flourish project (see Figure 1) [1] is to exploit this combined workflow in an autonomous robotic system for precision agriculture to achieve high yields while minimizing on-field chemical applications via targeted intervention.

This article presents an overview of the scientific and technical outcomes obtained within the Flourish project, providing insights and practical details on the lessons learned in several areas ranging from robot navigation, mapping, and coordination up to robot vision, multispectral data analysis, and phenotyping.

Robotics in Agriculture: An Overview of Recent Projects

Robotic applications in agriculture have significant potential to improve field monitoring and intervention procedures. However, these technologies are still in a development phase, with many possible uses yet to be explored.

A project similar to Flourish is Robot Fleets for Highly Effective Agriculture and Forestry Management (RHEA) [2], which aims at diminishing the use of agricultural chemical inputs by 75%, improving crop quality, enhancing human health and safety, and reducing production costs by means of sustainable crop management using a fleet of small heterogeneous robots (ground and aerial) equipped with advanced sensors, enhanced end effectors, and improved decision-control algorithms. Likewise, the Precision Farming of Hazelnut Orchards (PANTHEON) project [3] looks to develop a supervisory control and data acquisition system for precision farming in hazelnut orchards using a team of aerial–ground robots.

Other recent projects dealing with the development of autonomous ground platforms are the Ground Robot for Vineyard Monitoring and Protection (GRAPE) [4] and Sweet Pepper Harvesting Robot (SWEEPER) [5] projects. The former aims at creating agricultural service companies and equipment providers to develop vineyard robots that can increase the cost-effectiveness of their products as compared to traditional practices. In particular, the project addresses the market of instruments used for biological control by developing the tools required to execute (semi) autonomous vineyard monitoring and farming tasks with UGVs, thereby reducing environmental impact with respect to traditional chemical control. The SWEEPER's main

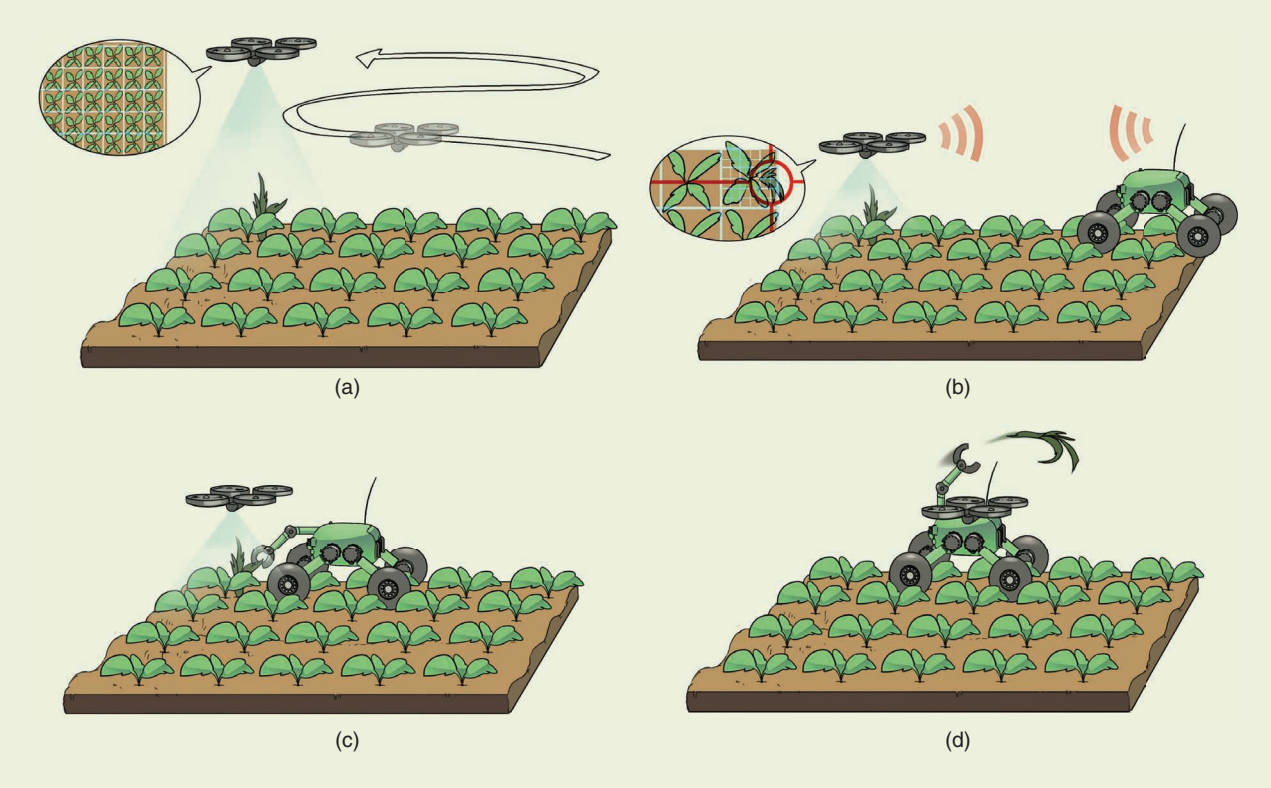

Figure 1. A conceptual overview of the Flourish project. (a) A UAV continuously surveys a field during the growing season. (b) The UAV collects data about crop density and weed pressure and (c) coordinates and shares information with a UGV, which is used for (d) targeted intervention and data analysis. The gathered and merged information is then delivered to farm operators for high-level decision making.

objective is to put first-generation greenhouse harvesting robots on the market.

UAVs are increasingly used in many agricultural robotics applications, e.g., for tree 3D reconstruction and canopy estimation [6], fruit counting [7], yield estimation [8], and automated monitoring using lightweight devices [9].

On the industry side, several start-ups have been initiated, and many more are expected to be funded. The major services provided are UGVs for weed removal [10]–[12] and in-season data analytics or early pest and disease detection from aerial or satellite imagery.

Experimental Platforms

The Flourish project exploited existing state-of-the-art farming and aerial robots, extending them in various ways to improve both autonomous navigation and environment modeling capabilities and enable them to perform robust plant classification and/or selective weed-removal operations.

Multirotor Used in the Flourish Project

The main UAV platform in the project is a fully sensorized DJI Matrice 100 [see Figure 2(a) and (b)]. The platform includes an Intel NUC i7 computer for onboard processing, a NVIDIA TX2 GPU for real-time weed detection, a GPS module, and a visual-inertial (VI) system for egomotion estimation. We employed a VI sensor developed at the Autonomous Systems Lab [13] and also tested and integrated a commercially available sensor, the Intel ZR300, for wider usage.

Ground Vehicle

The BoniRob Farming Robot

Bosch Deepfield Robotics' BoniRob [see Figure 2(c)] is a flexible research platform for agricultural robotics. Its four wheels can be independently rotated around the vertical axis—resulting in omnidirectional driving capabilities—and are mounted at the end of lever arms, letting the robot adjust its track width from 1 to 2 m. To execute complex tasks, the BoniRob carries a multitude of sensors: GPS, real-time kinematic (RTK) GPS, a push broom lidar, two omnidirectional lidars, red-green-blue (RGB) cameras, a VI system, hyperspectral cameras, wheel odometers, and so forth. These sensors are directly connected to a set of onboard PCs that run the Robot Operating System (ROS) and communicate through an internal network. The BoniRob's batteries are complemented by a backup generator that facilitates long-term field application.

Weed-Intervention Module

Supporting the target use case of selective weed intervention, the robot is equipped with an extension module, the weed-intervention module (Figure 3), which consists of a perception system for weed classification, multimodal actuation systems, and their supporting aggregates.

The main design objectives of this unit are high weed throughput, precise treatment, and flexibility. The weeds are treated mechanically with two ranks of stampers or chemically with one rank of sprayers. The weeds are detected and tracked in real time using three cameras with nonoverlapping fields of view (FoV).

The weed unit's perception system consists of three ground-facing global shutter cameras and three narrow-beam sonars. To protect this perception system from natural light sources, the weed-control unit is covered, and artificial lights are installed to control the illumination. A first RGB + near-infrared response camera is used for weed detection and tracking, while the other two RGB cameras are used for tracking. The sonars help recover the absolute scale of the camera images. Further details about the weed-intervention module can be found in the "Selective Weed Removal" section.

Data Analysis and Interpretation in a Farming Scenario

Precision farming applications aim to improve farm productivity while reducing the use of fertilizers, herbicides, and pesticides. To meet these challenges, in-field measurements of plant-vitality indicators and weed density are required. We addressed both of these requirements from a "robotic" point of view by proposing a set of methods to 1) accurately detect plants and classify them as crops and weeds (see the "Crop and Weed Detection" and "Automatic Synthetic Data Set Generation" sections) and 2) automatically analyze the N status of crops from multispectral aerial images (see the "Multispectral N-Status Detection and Phenotyping" section).

Crop and Weed Detection

A prerequisite for selective and plant-specific treatments with farming robots is an effective plant-classification system that provides the robot with information on where and when to trigger its actuators to perform the desired action in real time.

In the Flourish project, we focus on vision-based approaches for plant classification and use machine learning techniques to effectively cope with the large variety of different crops and weeds as well as with changing environmental conditions. Figure 4(a)–(c) illustrates the results obtained by our plant-classification systems for both the UGV and UAV platforms. The further distinction between weeds and grass-like weeds allows our system to perform different treatments in a targeted manner depending on the type of weed. For example, local mechanical treatments are most effective when applied to the stem location of the plants. In contrast, grass-like weeds can be effectively treated by spraying herbicides on their leaf surfaces.

We developed several data-driven plant-classification systems, ranging from approaches based on handcrafted features and random forests [14], [15] to deep learning

approaches based on lightweight, fully convolutional networks (FCNs) [16]. The latter showed superior performance and better generalization capability.

To effectively generalize to new conditions (different fields, weather, and so on), we exploited geometric patterns resulting from the fact that several crops were sown in rows. Within a field of row crops, the plants share a similar lattice distance along the row, whereas weeds appear randomly. In contrast to visual cues, this geometric signal is much less affected by changes in visual

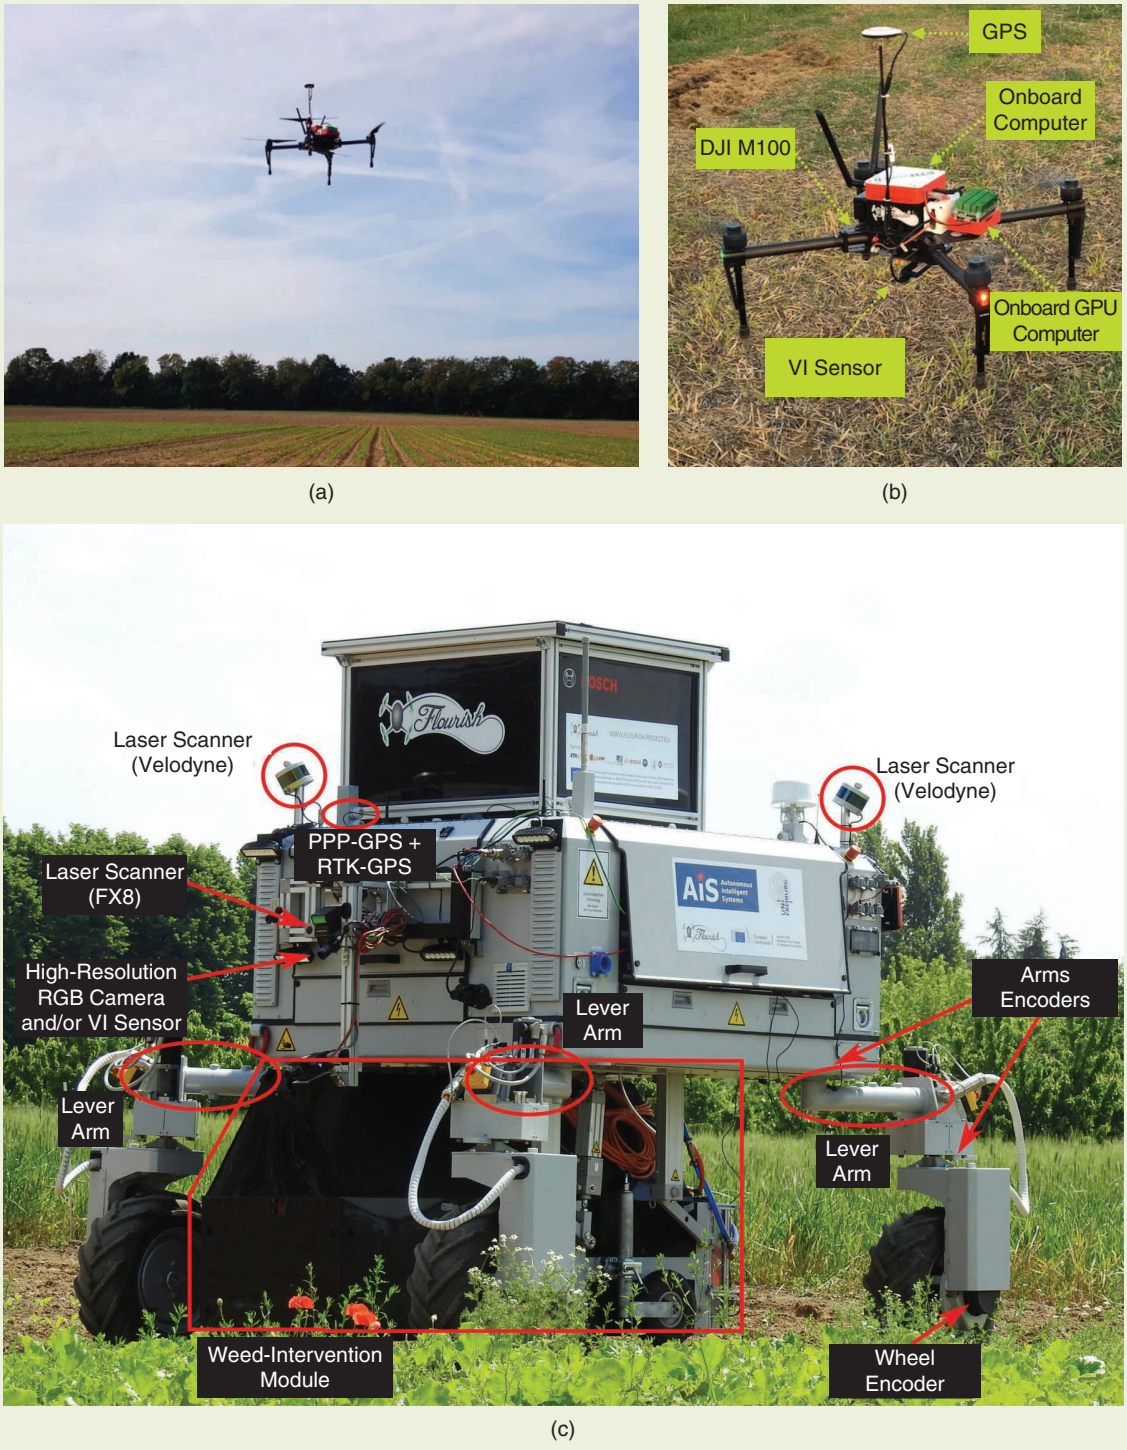

Figure 2. The two main robots used in the experiments and demonstrations: (a) a DJI Matrice 100 UAV multirotor performing an autonomous flight over a sugar beet field, (b) the UAV with the installed sensors, and (c) the Bosch BoniRob farming UGV. PPP: Precise Point Positioning.

appearance. We propose a semisupervised online approach [17] that exploits additional arrangement information about the crops to adapt the visual classifier. We also successfully tested approaches that operated on image sequences obtained along crop rows, enabling the classifier to learn features that describe the plant arrangement [see Figure 5(a) and (b)] [16]. The image sequence reveals that crops grow along the row and have similar spacing, whereas the weeds appear randomly in the field strip. We show that incorporating this geometric information boosts the classification performance and generalization capabilities of the plant classifiers.

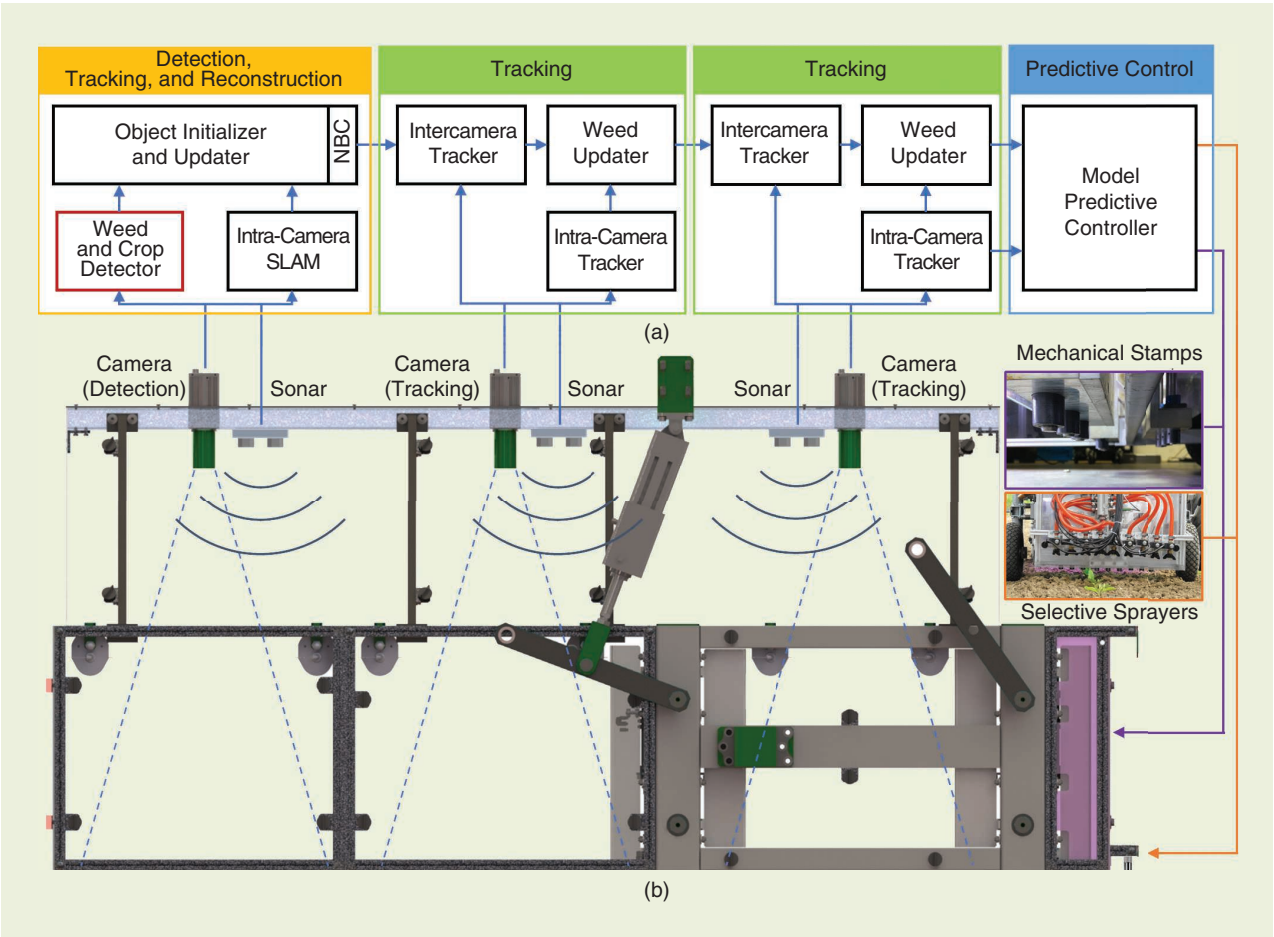

Figure 3. (a) A schematic rendering of the weed-intervention module. (b) An overview of our proposed weed-control system, composed of weed detection, tracking, and predictive control modules. The weeds are tracked across the cameras and finally fed into a predictive control module to estimate the time and position of treatment at which they will be approaching the weeding tools. SLAM: simultaneous localization and mapping; NBC: naive Bayes classifier.

Figure 4. The example results obtained by our plant-classification systems. (a) A UGV-based semantic segmentation into crop, weed, and grass-like weed. (b) Stem detection providing the accurate location of crops and weeds. (c) A UAV-based semantic segmentation.

The same underlying lightweight FCN structure was deployed on our UAV systems [18]. Here, we used the FCN in a classical single-image fashion, as the larger footprint of the camera implicitly covered enough information about the plant arrangement. Through our crop and weed classification systems, we enabled UGVs to perform plant-specific, high-precision in-field treatments and transformed the UAVs into an efficient system for crop-monitoring applications.

Automatic Synthetic Data Set Generation

High-performing, data-driven plant-classification approaches usually require large annotated data sets acquired from across different plant growth stages and weather conditions. Annotating such data sets at a pixel level is an extremely time-consuming task.

We address this problem by proposing an automatic, model-based data set generation procedure [19] that produces large synthetic training data sets by rendering a large number of photo-realistic views of an artificial agricultural scenario using a modern 3D graphic engine. To do so, we randomize a few key parameters (e.g., the size, deformation, and distribution of leaves; the structure of plants; types of soil; and so on) using a bottom-up procedure (see Figure 6):

- We model the leaves of the target plants using kinematic chains upon which we apply a few real-world RGB textures.
- We model the plants through radially distributed leaf layers, where the number of leaves per layer depends on plant species and their growth stages.
- A virtually infinite number of realistic agricultural scenes can be then rendered by adding random soil backgrounds and sampling random illumination conditions and plant distributions; the ground-truth segmentation masks are automatically generated by the graphic engine;.
- The generated synthetic data sets can be used to effectively train modern deep learning-based image-segmentation architectures.

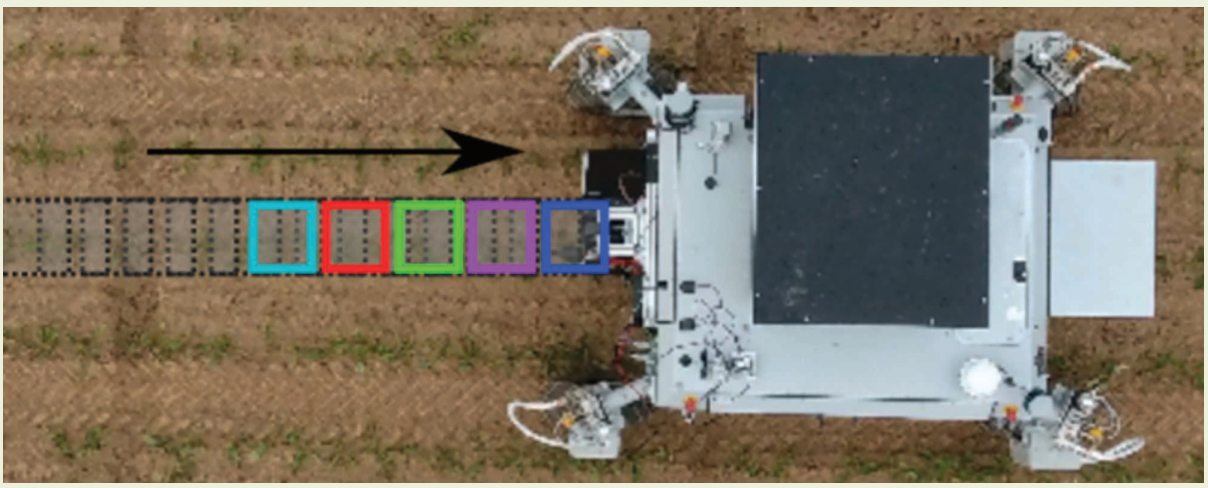

(a)

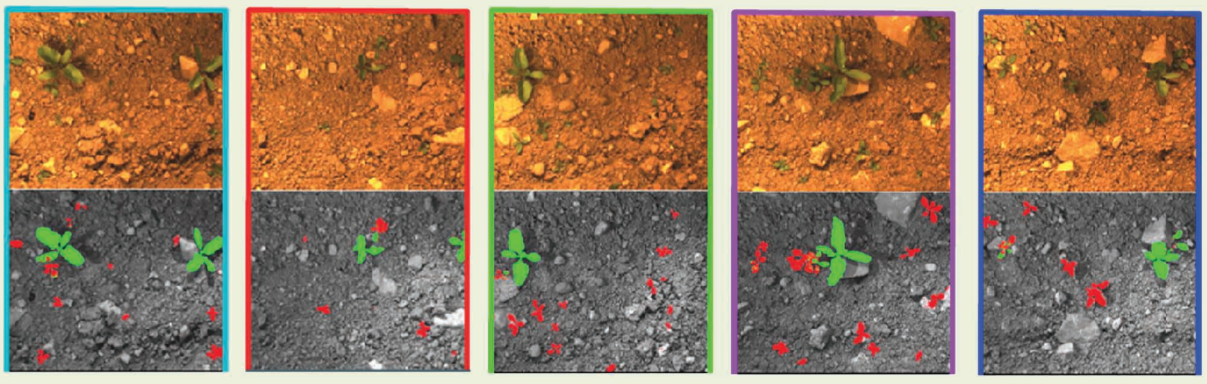

(b)

Figure 5. (a) The BoniRob UGV acquiring images while driving along the crop row. Our approach [16] exploits an image sequence by selecting those images from the history that do not overlap in object space. (b) The exemplary prediction of crop plants and weeds for the entire image sequence. Note that the model was trained on data acquired in a different field.

Multispectral N-Status Detection and Phenotyping

The nutrition status of a crop is linked to yield formation and the environmental footprint of agronomy. Well-fertilized crops produce optimal yield and quality and are more stress resilient. Fertilizer deficiency hampers yield, whereas a surplus supply increases the risk of nutrient loss to the environment and increases susceptibility to pests and diseases. N plays a prominent role in the management of most crops because of crops' generally high demand for N and its very high mobility in soil. Although sugar beet N demand is relatively low, its yield and quality are strongly dependent on N management: too low an N application limits the tuber yield, while a high N application reduces the extractable sugar content in the tuber [20].

Therefore, it is important to apply N fertilizer at the right time, rate, and place. These decisions can be supported by optical remote sensing tools that make use of visible or nonvisible parts of the spectral reflection of crops, as employed for N-status detection in the Flourish project

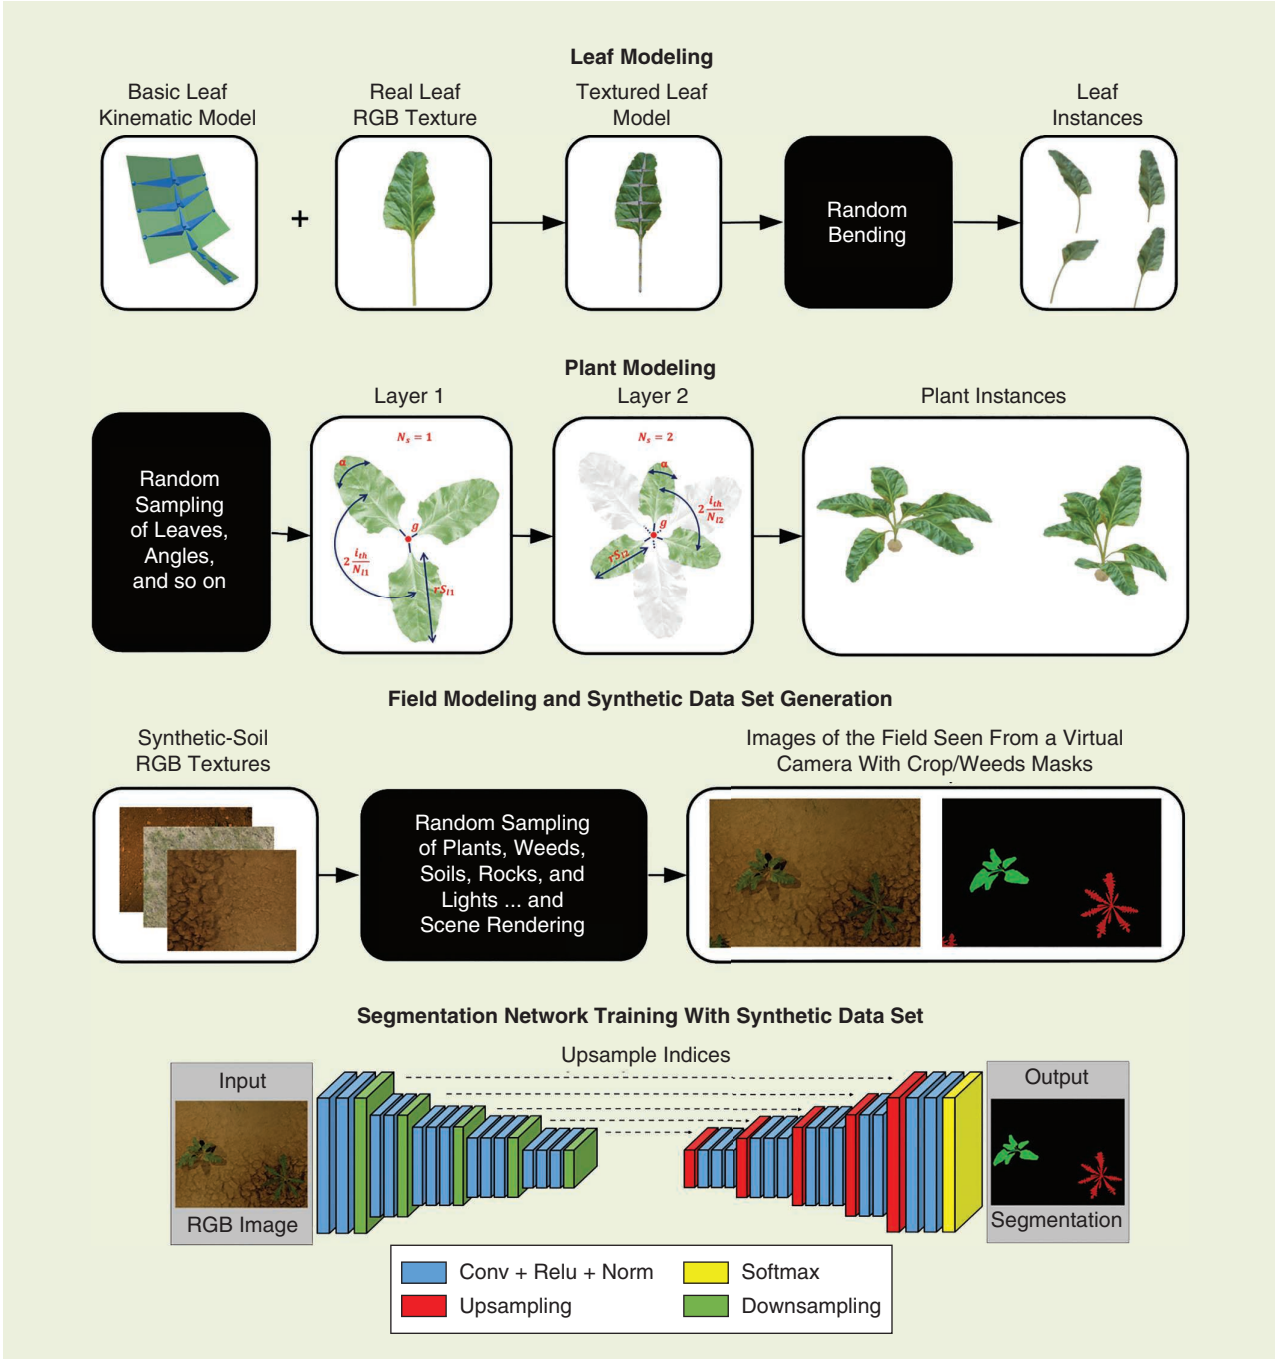

Figure 6. An overview of the automatic, model-based data set generation procedure. Conv: convolution; Relu: rectified linear unit; Norm: batch normalization.

[21]. It serves as an example among others for the image-based assessment of plant traits, which play an increasingly prominent role in the development of sustainable agronomic practices in precision farming [22], [23].

To validate the spectral and imaging methodology for sugar beets in the Flourish project, randomized field trials were established in commercial sugar beet fields, and different N-input treatments were applied from 2015 to 2017. Aerial-image spectroscopy was realized with a multispectral Gamaya VNIR 40 camera mounted on a UAV. For the ground-truth plant N status, tuber yield and sugar content were measured. As for other crops, our results show red-edge-based spectral indices, such as the simple ratio and the normalized-difference red-edge ratio [24]; this indicates the N status in sugar beets successfully obtained from the UAV-based sensor, resulting in useful N-fertilizer application maps.

Positioning and Environment Modeling

The ability to localize and build a model of the surrounding environment is an essential requirement to support the reliable navigation of an autonomous robot. Such tasks are even more challenging in a farming scenario where 1) the environment is mainly composed of repetitive patterns with no distinctive landmarks and 2) multispectral information should be included in the modeling process to support decision making for farm management. Moreover, in a multirobot setup as in the Flourish project, the UAV and the UGV should be able to cooperatively build a shared model of the environment. The following presents the main contributions we proposed for localizing and modeling cultivated fields using a UAV (the "UAV Localization and Mapping" section) and a UGV (the "UGV Global Positioning and Mapping" section) as well as in terms of fusing this information among robots (the "Cooperative UAV–UGV

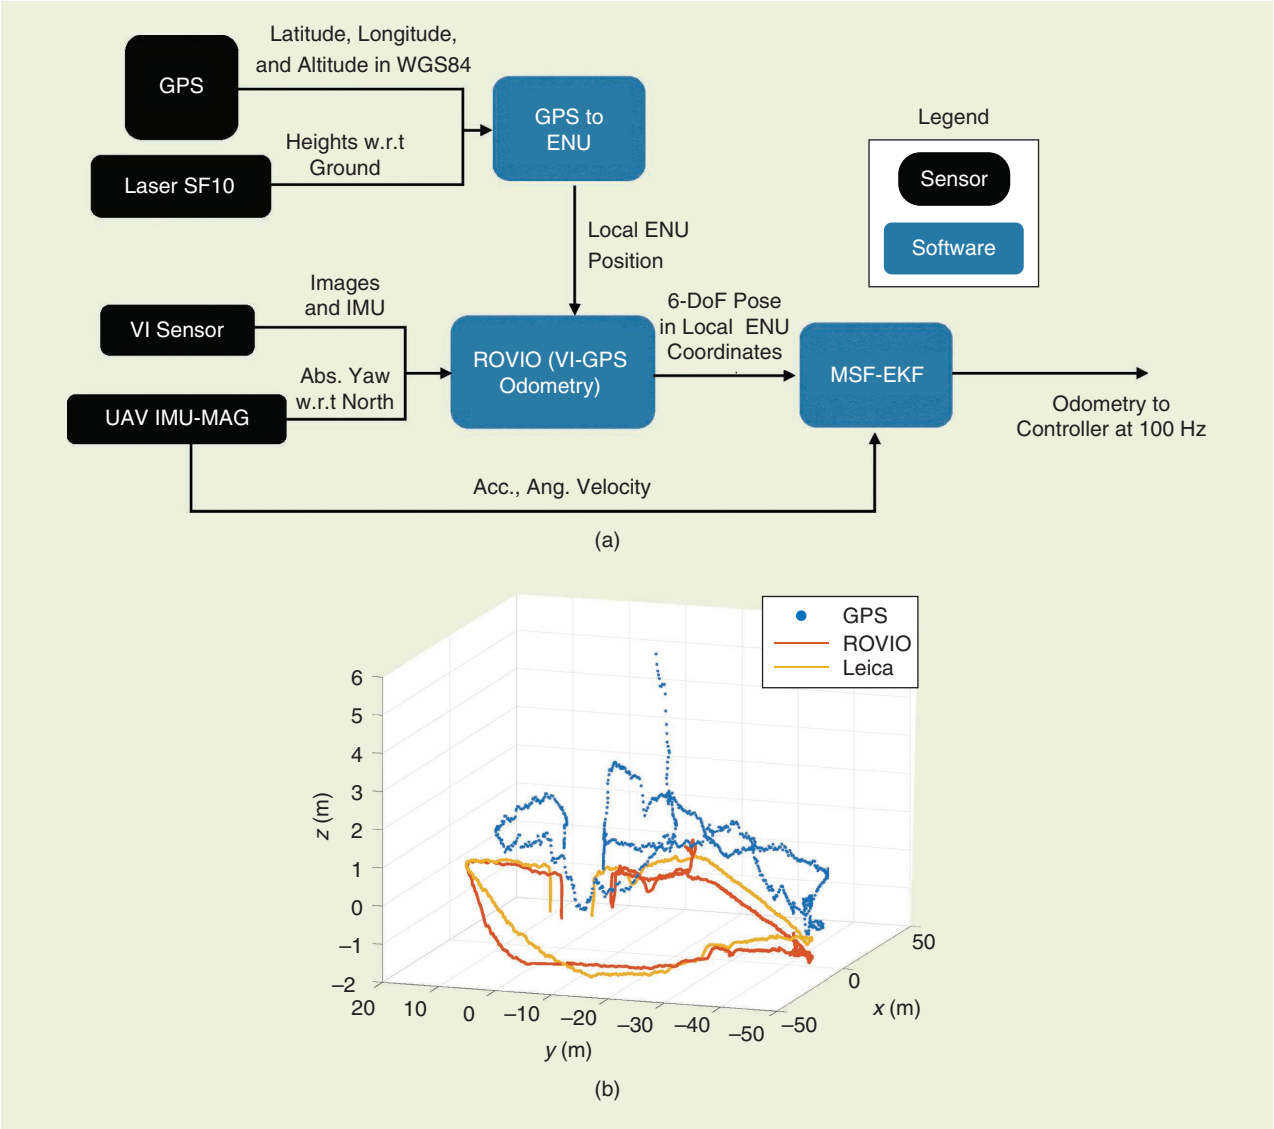

Figure 7. (a) A block diagram of our UAV state-estimation framework depicting the sensor suite and major software components. (b) A comparison of our VI-GPS fusion-based state estimation with raw GPS and ground truth. w.r.t.: with respect to; WGS: World Geodetic System; MAG: magnetometer; ENU: east, north, up coordinates; Abs.: absolute; Acc.: acceleration; Ang.: angular.

Environment Modeling" section) and across time (the "Long-Term Temporal Map Registration" section).

UAV Localization and Mapping

The aim of the UAV perception system is to collect high-resolution spatiotemporal, multispectral maps of the field. These data are critical as they allow for mission planning before the UGV actually enters the field, thereby optimizing the time/location of ground-intervention procedures without the risk of crop damage and soil compaction. The perception pipeline requires two main competencies: 1) motion estimation and precise localization within the field and 2) multiresolution, multispectral aerial mapping based on the indicators needed to assess plant health.

The key challenges for on-field vision-based localization are the homogeneous appearance of crops and the accuracy of the GPS which, when used alone, is not sufficient to construct maps for defining paths for UGV intervention. To address this, we developed an onboard state-estimation system that combines data from a synchronized VI sensor, a GPS sensor, the UAV inertial measurement unit (IMU), and, optionally, a laser altimeter to estimate the six degree-of-freedom (6-DoF) pose. Figure 7(a) provides an overview of the system. The robust visual inertial odometry (ROVIO) [25] framework is used to produce a 6-DoF pose output based on raw images and the IMU data from the VI sensor. The multisensor fusion framework [26] then combines the ROVIO output and the UAV IMU data to obtain state estimates that are passed to our model predictive controller for trajectory tracking. To improve its accuracy and robustness, we integrated our system with MAPLAB [27], a framework that has map-maintenance and processing capabilities. On-field results using the AscTec NEO and DJI Matrice 100 UAV platforms demonstrate a high state-estimation accuracy compared to the ground truth from a Leica Geosystems Total Station [see Figure 7(b)].

High-resolution field map models are a key prerequisite for enabling robots in precision agriculture. To this end, we developed a UAV environmental modeling framework using the pose estimate from our localization system as well as color and multispectral camera information over multiple flights to create spatiotemporal spectral field models. Figure 8 depicts our pipeline. Taking raw RGB and multispectral images and UAV poses as inputs, we radiometrically corrected the spectral data to create a spatial field model in the form of a dense point cloud. For each point, the spectral reflectances in the multispectral wavelength bands were estimated and stored. The field evolution over time can be viewed through layered orthomosaics generated from this data through a custom browser-based visualization module, as shown in Figure 9(a) and (b). We used higher-quality RGB camera images, which recovered high-resolution field geometry, and also used the relative position and orientation between the RGB and multispectral camera to estimate its spectral

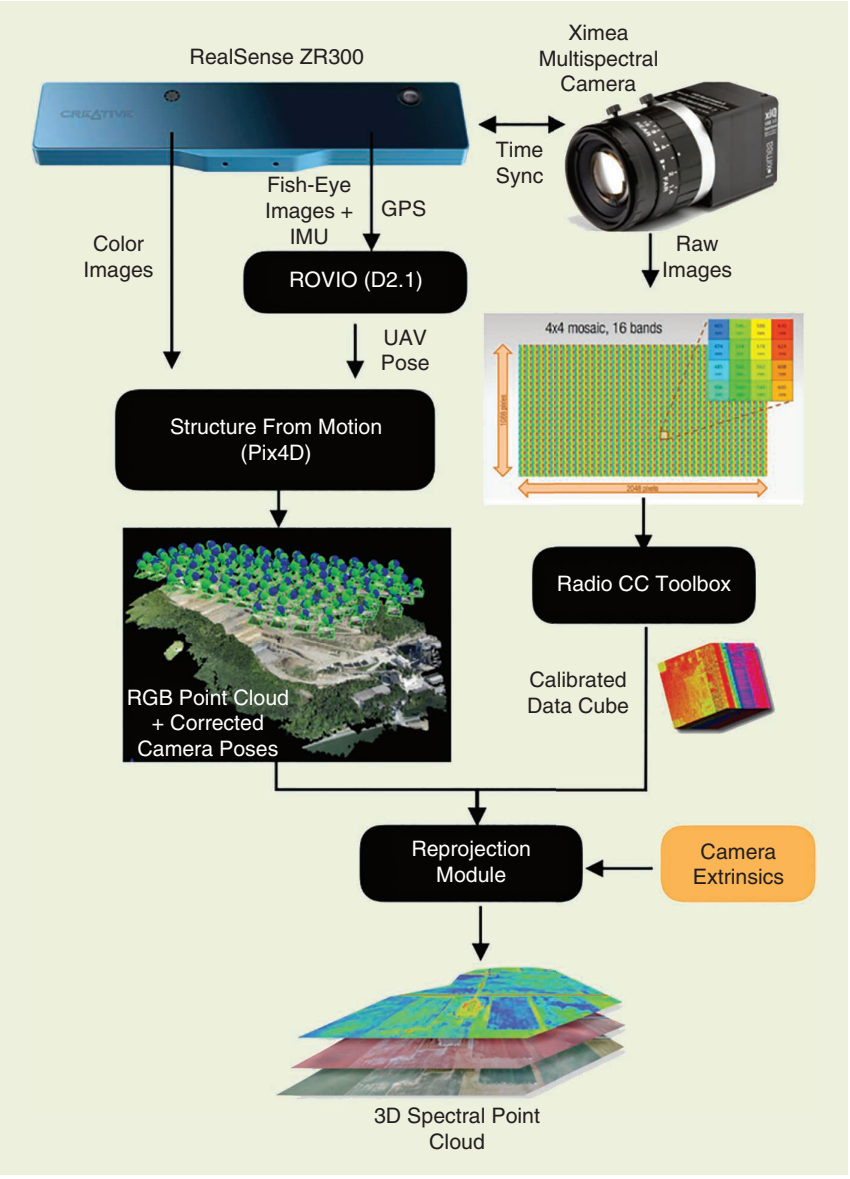

Figure 8. A block diagram of our environment modeling framework depicting the sensor suite and major software components. CC: camera calibration. (Images courtesy of Intel, Pix4D SA, and Ximea.)

reflectance. Importantly, our strategy eliminated the need for a separate reconstruction step for each band and the subsequent alignment step.

UGV Global Positioning and Mapping

Currently, most positioning systems used in commercial farming UGV rely on high-end RTK-GPSs that are, however, not robust enough against base station signal loss or multipath interference and cannot provide the full 6D position (translation and rotation) of the vehicle. We tackle this problem by proposing a UGV positioning system [28] that effectively fuses several heterogeneous cues extracted from consumer-grade sensors and exploits the specific characteristics of the agricultural context with a few additional constraints. We formulate the global localization problem as a 6D pose graph optimization problem. The constraints between consecutive nodes [Figure 10(a)] are represented by motion estimations (the wheel and visual odometries and so on). Noisy, but drift-free GPS and IMU readings are directly integrated as prior nodes. Driven by the fact that both GPS and visual odometry (VO) provide poor estimates along the z-axis (i.e., parallel to the gravity vector), we introduce two additional altitude constraints:

- an altitude prior, provided by a digital elevation model
- a smoothness constraint for the altitude of adjacent nodes.

The integration of such constraints improves the accuracy of the altitude estimation and also benefits the estimation of the remaining state components. The optimization problem is cyclically solved online by using a sliding-window strategy, as presented in Figure 10(b).

Cooperative UAV–UGV Environment Modeling

Building a shared map of the environment is an essential but challenging task: the UAV can provide a coarse reconstruction of large areas, which should be updated with more detailed information collected by the UGV. We introduced Aerial–Ground Collaborative 3D Mapping for Precision Farming (AgriColMap) [29], an effective map-registration pipeline that registers heterogeneous maps built by the UGVs and UAVs.

AgriColMap leverages a multimodal field representation and formulates the data-association problem as a large displacement dense optical flow (LDOF) estimation. The complete pipeline is schematized in Figure 11(a). We assume that both the UAV and UGV can generate colored, georeferenced point clouds of a farm environment, $\mathcal{M}_A$ and $\mathcal{M}_G$, e.g., using photogrammetry-based 3D reconstruction. Our goal is to estimate an affine transformation $F: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that allows their accurate alignment by compensating for the geo-tags misalignments and the reconstruction and scale errors. We start looking for a set of point correspondences $m_{A,G} = \{(p, q) : p \in \mathcal{M}_A,\ q \in \mathcal{M}_G\}$ that represents points pairs belonging to the same global 3D position. Inspired by the fact that the points in $\mathcal{M}_A$ locally share a coherent "flow" toward the corresponding points in $\mathcal{M}_G$, we cast the data-association problem as a dense, regularized matching approach. This problem recalls the dense optical flow estimation problem for RGB images: we introduce a multimodal environment representation that allows for the exploiting of such 2D methods on 3D data while enhancing both the semantic and geometrical properties of the maps. We exploit two intuitions:

- A digital surface model (DSM) well highlights the geometrical properties of a cultivated field.
- A vegetation index can highlight the meaningful parts of the field and its visual patterns.

We transform $\mathcal{M}_A$ and $\mathcal{M}_G$ into 2D grid maps $\mathcal{J}_A, \mathcal{J}_G: \mathbb{R}^2 \rightarrow \mathbb{R}^2$, where, for each cell p, we provide the surface height h and the excess green index (ExG), $ExG(p) = 2g_p - r_p - b_p$, with r_p, g_p, b_p being the (average) RGB components of the cell. To estimate the offsets map, we employ a modified version of the LDOF

Figure 9. A visualization interface for the spatiotemporal spectral database showing (a) RGB orthomosaics and (b) the corresponding index maps for a sugar beet field over time. The user can select spectral layers to view a georeferenced reflectance orthomosaic corresponding to a wavelength band, view the color orthomosaic, and toggle through all of the available surveys using the timeline. (Images courtesy of Flourish.)

coarse-to-fine PatchMatch framework (CPM) [30]. We apply the visual descriptor of the original CPM method directly to the ExG channel of $\mathcal{J}_A$ and $\mathcal{J}_G$ while we exploit a 3D descriptor computed over the DSM to extract salient geometric information; the matching cost has been modified accordingly to take into account both descriptors.

The largest set of coherent flows defines a set of matches $m_{A,G}$ that are used to infer a preliminary alignment $\hat{F}$. Finally, we estimate the target affine transformation F by exploiting the coherent point drift registration algorithm [31] over point clouds $\mathcal{M}_A^{\text{veg}}$ and $\mathcal{M}_G^{\text{veg}}$, which are obtained from $\mathcal{M}_A$ and $\mathcal{M}_G$ by extracting only the points that belong to vegetation using an ExG-based thresholding operator.

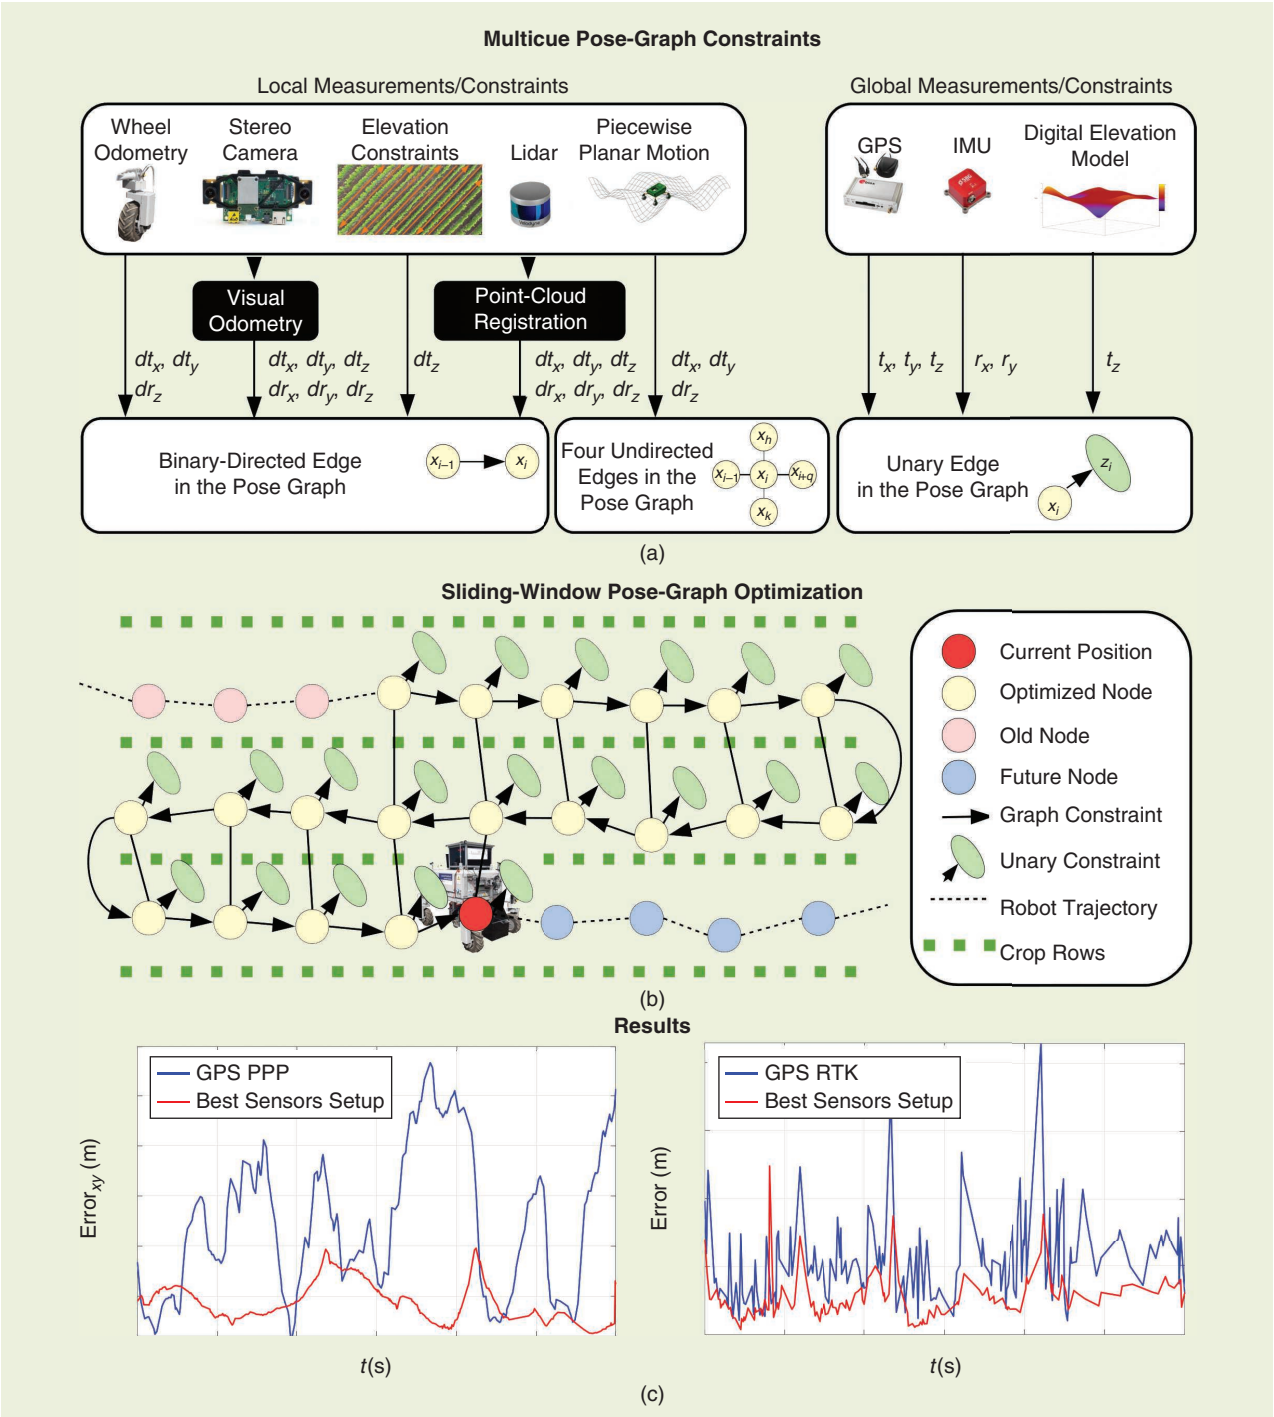

Figure 10. An overview of the proposed graph-based multicue UGV positioning system. (a) The constraints (i.e., the sensors and provided measurements) integrated in the pose graph, (b) the conditional dependencies between nodes, and (c) the absolute error plots for the raw GPS (blue) and optimized trajectories (red), respectively, for two types of GPSs.

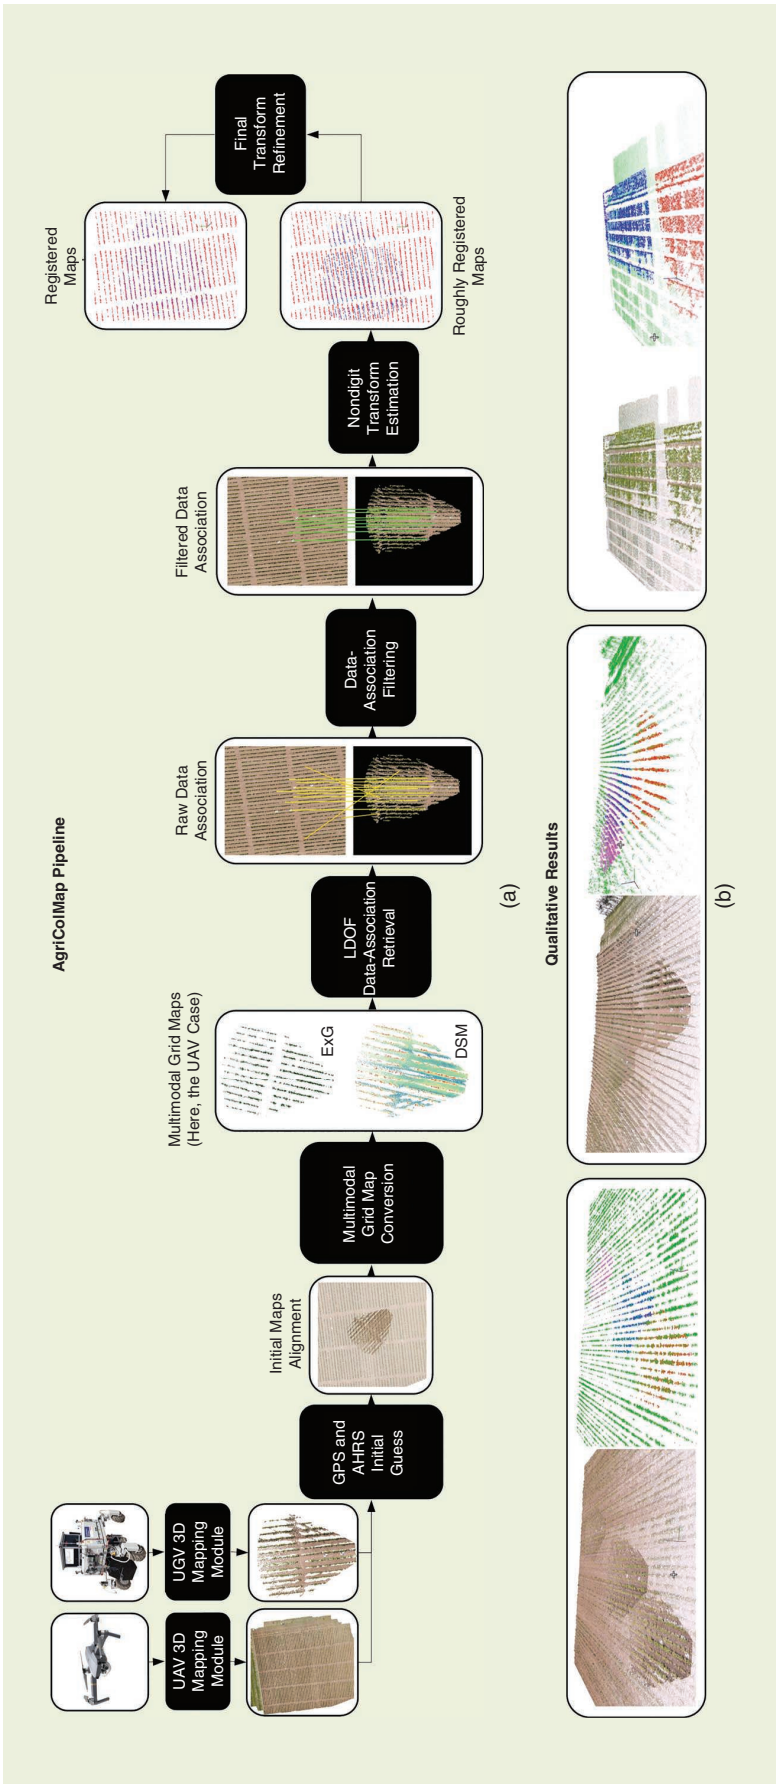

Figure 11. (a) An overview of the AgriColMap method pipeline. (b) Some qualitative registration results (the RGB and ExG-filtered point clouds). The UGV clouds are clearly visible due to their higher points density. AHRS: attitude and heading reference system.

Long-Term Temporal Map Registration

Continuous crop monitoring is an important aspect of phenotyping and requires the registration of sensor data over the entire season. This task is challenging due to significant changes in the visual appearance of the growing crops and the field itself. Conventional image registration based on visual descriptors is typically unable to deal with such drastic changes in appearance. To overcome this challenge, we developed a method for registering temporally separated images by exploiting the inherent geometry of the crop arrangement in the field, which remains relatively invariant over the season. We proposed a scale-invariant geometric feature descriptor that encodes the local plant-arrangement geometry and uses these descriptors to register the images even in the presence of strong visual changes [32]. The registration results allow for the spatiotemporal analysis of data collected over the crop season. This includes applications such as monitoring growth parameters at a per-plant level, as illustrated in Figure 12.

Planning, Navigation, and Coordination

The UAV and UGV have different working areas and roles within each field analysis and targeted-intervention mission. Their action planning and navigation policies should reflect these differences. We introduced an ad hoc UAV navigation module (see the "UAV Mission Planning and Navigation" section) using a planner to effectively perform field-monitoring missions while respecting battery constraints. Crop row localization and safe in-field UGV navigation are addressed in the "UGV Position Tracking and Navigation" section, where the high number of DOF of the UGV is used to improve motion efficiency and smoothness. The inter-robot mission-coordination framework is then introduced in the "UAV–UGV Mission Coordination" section.

UAV Mission Planning and Navigation

A key challenge in agricultural monitoring is developing mission-planning algorithms to define the path for a UAV to optimally survey the field. The planning module needs to maximize mapping accuracy given battery-life constraints, taking into account field coverage and scientifically defined areas of interest. We developed an informative path planning (IPP) framework for adaptive mission planning to meet these requirements [33].

Our framework is suitable for mapping a terrain depending on the type of data received from an onboard sensor, e.g., a depth or multispectral camera. In terms of mapping, the main challenge is fusing the dense visual imagery into a compact probabilistic map in a computationally efficient way. To address this, we present a new method for multiresolution mapping that considers the patterns of the target distributions on the farm. We used Gaussian processes (GPs) to encode the spatial correlations common in biomass distributions. A GP model was exploited as a prior for recursive Bayesian data fusion with probabilistic, variable-resolution sensors. In doing so, our approach enables mapping without the computational burden of standard GP regression, making it suitable for online, on-platform applications.

In terms of planning, a fundamental problem we tackled is trading off image resolution and FoV to find the most useful measurement sites at different flying altitudes. During a mission, the terrain maps built online are used to plan trajectories in continuous 3D space that maximize an information-based objective, e.g., the targeted, high-resolution mapping of areas infested by weeds. Our planning scheme proceeded in a finite-horizon fashion, alternating between replanning and plan execution. This allowed us to create adaptive plans, taking new sensor data into account online to

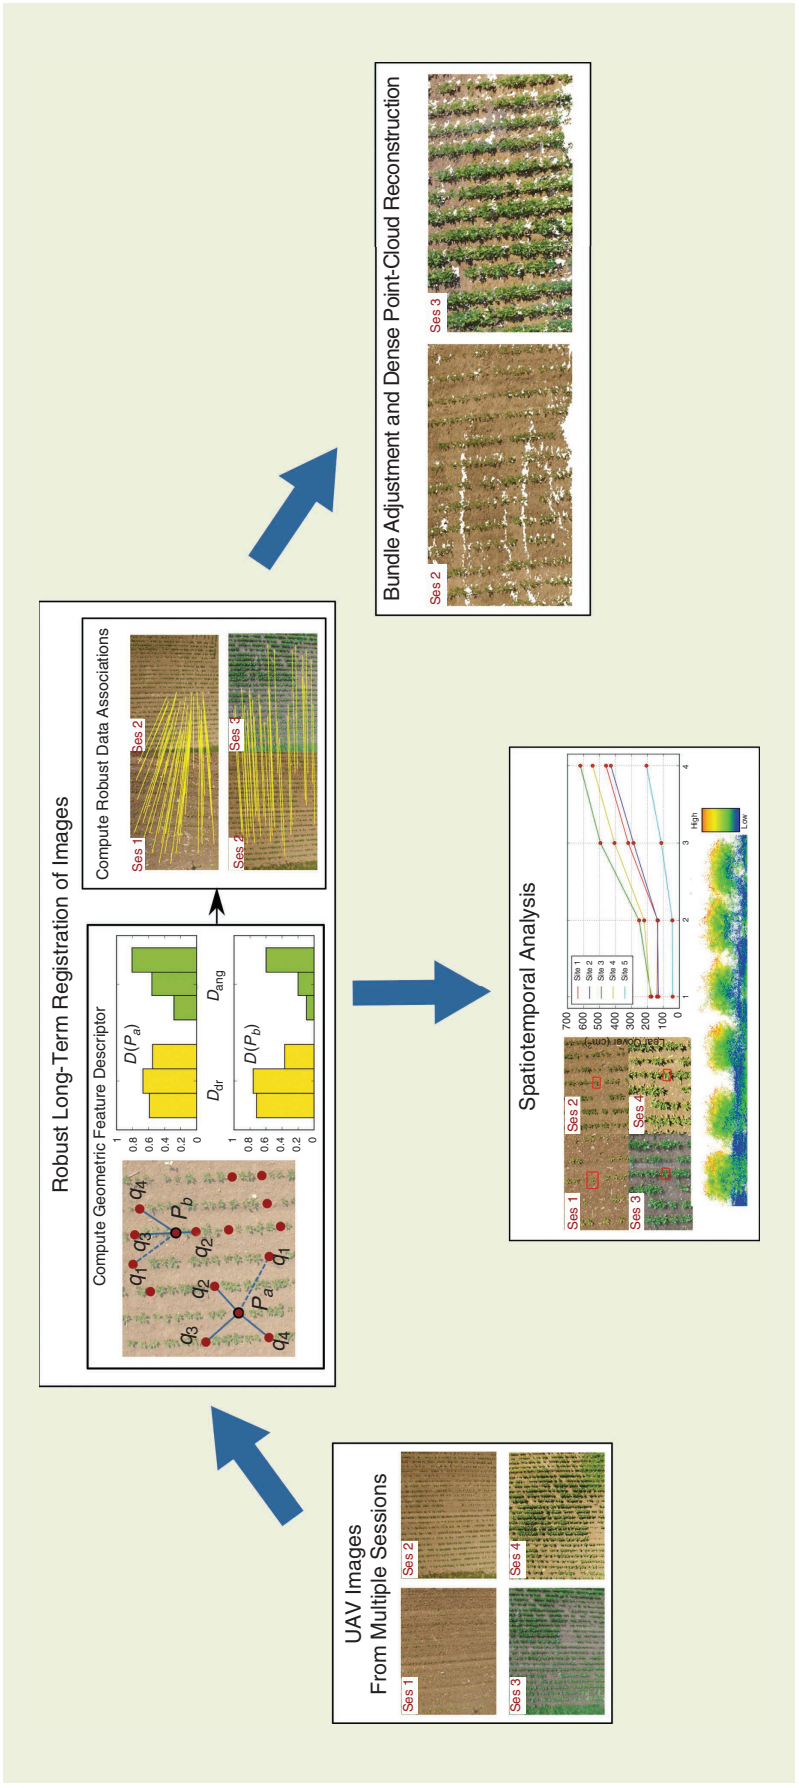

Figure 12. The pipeline used for performing long-term registration of crop-field UAV images. The registration approach computes robust data associations between temporally separated images based on a scale-invariant geometric descriptor. These associations allow for registering the data in a common reference frame, which then forms the basis for performing spatiotemporal analysis. Ses: session.

focus on areas of interest as they were discovered. For replanning, we leveraged an evolutionary technique, i.e., the covariance matrix-adaptation evolution strategy, to optimize the initial trajectory solutions obtained from a coarse 3D grid search in the UAV workspace. Our approach was evaluated extensively in simulation, where it was shown to outperform existing methods (Figure 13) and was validated on the field.

UGV Position Tracking and Navigation

For autonomous navigation on fields, the BoniRob UGV needs to accurately steer along the crop rows without crushing any of the value crops. Moreover, to transition between crop rows, performing tight and accurate turns at the end of the field is essential. There are three key requirements to achieve this: a pose estimate relative to the rows, a path along the crop rows through the field, and smooth velocity commands to precisely follow this path. To this end, we developed a crop row detection algorithm, the pattern Hough transform [34]. We first processed the input from vision or lidar data by extracting the plant features and projecting them onto a feature grid map in the local robot frame [see Figure 14(a)]. Then, our pattern Hough transform determined the pattern of parallel and equidistant lines that is best supported by the feature map, as shown in Figure 14(c). Such a pattern is defined by the orientation θ, the spacing between adjacent lines s, and the offset of the first line to the origin o, as displayed in Figure 14(b). Because our approach takes into account all of the available data used to detect crop

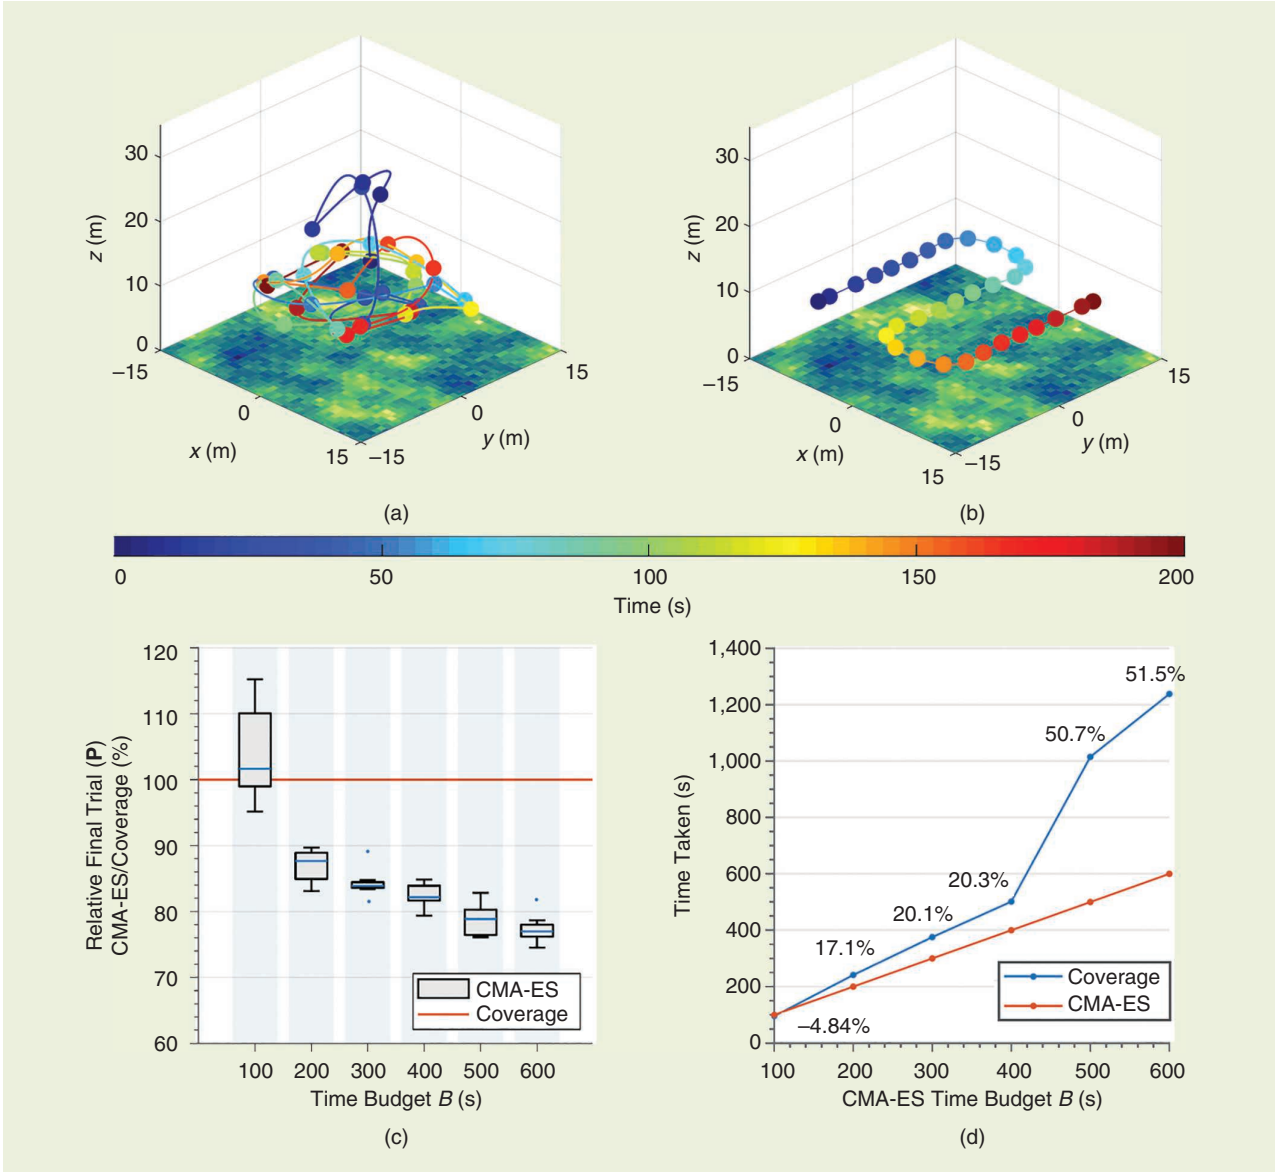

Figure 13. Example comparisons of (a) our covariance matrix adaptation-evolution strategy (CMA-ES)-based approach to (b) "lawnmower" coverage for mapping in 200-s missions. The colored lines and spheres represent the traveled trajectories and measurement sites. Ground-truth maps are rendered. (c) A comparison of the final map uncertainties (measured by the GP covariance matrix trace) for various path budgets. Ten CMA-ES trials were run for each budget. (d) A comparison of the times taken to achieve the same final map uncertainty given a fixed CMA-ES budget.

rows in a single step, it is robust against outliers like weeds growing between the crop rows and yields accurate results during turning, i.e., when the robot is not necessarily aligned with the crop rows.

We integrated the output from our pattern Hough transform into the localization module of the BoniRob. The localization was based on an extended Kalman filter. We fused the odometry and IMU measurements for the prediction. In the correction step, we aligned the detected crop row pattern with a GPS-referenced map of crop rows to correct the pose estimate of the robot relative to the field. Because the crop row pattern provides only lateral and orientation information, i.e., no correction along the crop rows, we corrected the longitudinal position estimate using GPS signals.

We implemented a global planner based on a state lattice planner to ensure that the BoniRob found a path to any reachable pose in the field. The BoniRob can change its track width by adjusting the angles of the lever arms to which the wheels are attached, as shown in Figure 2(a). Thus, whether it can pass through a narrow gap or over an obstacle depends on the wheel positions [see Figure 14(d)]. We developed a path planner that considers the lever angles explicitly [35] by including the arm angles in the state space and adding actions that allow the planner to change them. Adding the arm angles greatly increases the size of the state space, which makes planning with commonly used search algorithms inefficient. Thus, we introduced a novel method to represent the robot state with a reduced cardinality; that is, we tracked valid arm-angle intervals instead of single-arm angles in the robot state.

Our local planner translates a pose path from the global planner into velocities while considering steering constraints. Any robot with slow-turning, independently steerable wheels, such as the BoniRob, has certain steering constraints. The most prominent constraints are limited steering velocity, non-continuous steering, or wheel-angle instabilities when the center of rotation is on a wheel. To avoid violations of these steering constraints, we present a new approach that incorporates steering constraints when generating velocity rollouts [see Figure 14(f)] [36]. Our approach leverages the correspondence between the wheel angles and the instantaneous center of rotation (ICR) of the robot. After projecting the steering constraints into the ICR space, we computed a valid ICR path that satisfies the constraints. From this ICR path, we calculated valid velocity sequences that the robot can execute smoothly. Real-world experiments show that our local planner improves efficiency and leads to smoother execution.

UAV–UGV Mission Coordination

To unlock the potential of the Flourish robotics system, the ability to run coordinated missions between the robots is essential. Because both robots share information via Wi-Fi, this information needs to be kept at a minimal level, and the coordination needs to be ensured even when communication is lost. The only data exchanged are the UAV and UGV poses, the coordinates of the areas of interest, the requests from one robot to the other, and their status messages. Because of the lossy communication, exchanging requests

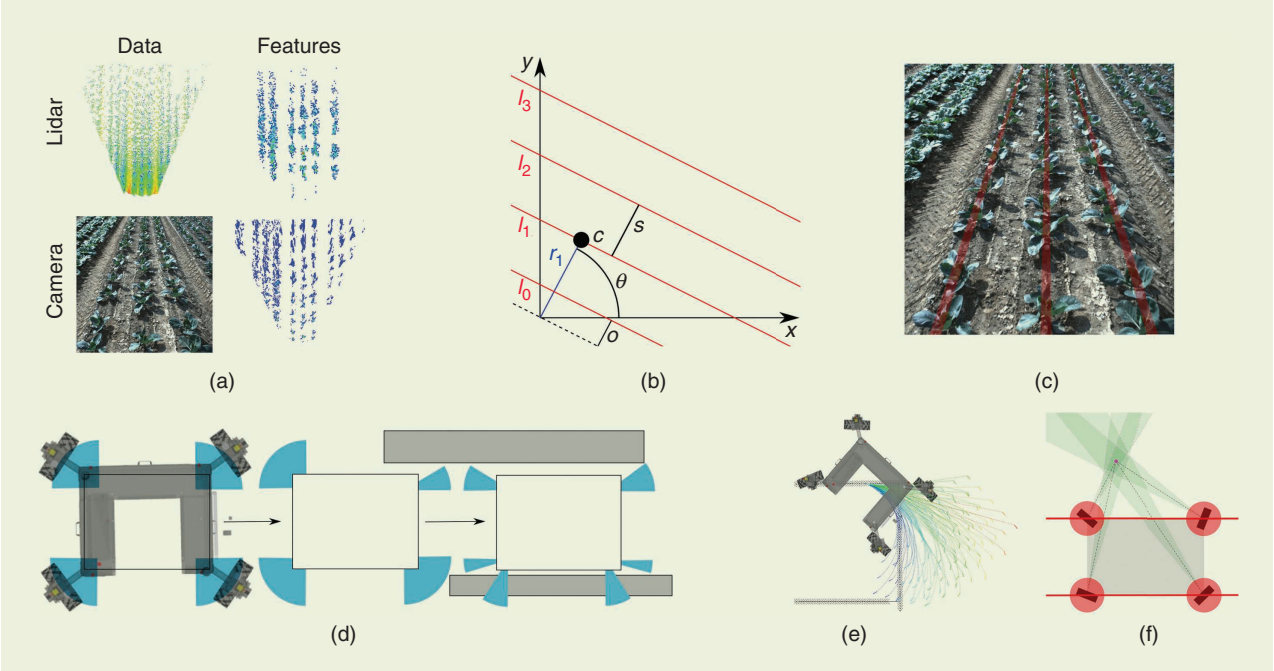

Figure 14. An overview of the UGV navigation system. (a) The pattern Hough transform detects crop rows in lidar or camera data using the extracted plant features. (b) A pattern is defined as a set of parallel and equidistant lines (red) with orientation θ, offset o, and spacing s. (c) The result of the pattern Hough transform (red) on pointed cabbage (≈10 cm). (d) The valid arm-angle intervals when moving close to an obstacle. (e) The velocity rollouts in the local planner. The rollouts are color-coded with their respective costs. (f) The ICR constraints derived from the hardware constraints (red) and the maximum steering velocity (green).

and status is a reliable way to ensure that a message sent by one robot is indeed received by the other.

The mission framework used on both robots is based on ros_task_manager [37], a task scheduler developed for ROS that is particularly easy to use and allows for combining multiple behaviors with elements running in sequence or in parallel, eventually interrupting each other. This framework is based on tasks implemented in C++ that are combined into complex missions implemented in basic Python. Figure 15 is an example of a coordinated mission.

In-Field Intervention: The Collaborative Weeding Use Case

The main use case addressed in the Flourish project is the collaborative weeding application (Figure 1). The UAV flies over the field, running the navigation and planning algorithms discussed in the "UAV Localization and Mapping" and "UAV Mission Planning and Navigation" sections while analyzing the weed pressure using the classification algorithms presented in the "Crop and Weed Detection" section. The UGV is alerted to high weed-pressure areas using the coordination framework described in the "UAV–UGV Mission Coordination" section. Thus, the UGV starts to move toward the selected areas, running the algorithms discussed in the "UGV Global Positioning and Mapping" and "UGV Position Tracking and Navigation" sections. In the following, we describe the tools (see the "UAV Localization and Mapping" section) and methods (see the "Weed Tracking" section) used for actual weed treatment with the possible agronomic impacts reported in the "Agronomic Impacts in Sugar Beet Crops" section. We successfully tested the whole pipeline in a public demonstration during a dissemination event held near Ancona (Italy) in May 2018.

Selective Weed Removal

The weed-intervention module [Figure 3(a)], whose perception system was introduced in the "Ground Vehicle" section, includes further tools designed to address the targeted weed treatment: a weed stamping tool and a selective spraying tool [Figure 3(b)]. The stamping tool is composed of 18 pneumatic stamps arranged in two ranks. All of the stamps are individually controllable, and highly precise positioning is ensured by allowing only 1 DoF for the positioning across to the driving direction. The spraying tool is positioned in the back. It is assembled out of nine nozzles, individually controlled by off-the-shelf magnetic valves.

Both weeding tools are controlled using a scalable programmable logic controller. Modules requiring more computational resources, i.e., weed detection and tracking, are implemented on a computer dedicated to the weed control running Linux and ROS.

The bolt of the stamps has a 10-mm diameter, whereas the footprint of a sprayer is 30 mm when set in the lowest position, as in our experiments. To actually treat a weed while the robot is moving is a time-critical part of the process because a small delay can lead to a position error at the centimeter-level, which is large enough to miss a small weed. In our experiments, the decision of which tool to use on which weed is based only on a size criterion: large weeds are sprayed while small weeds are stamped.

Weed Tracking

The main challenge in weed tracking with nonoverlapping multicamera systems [Figure 3(b)] (see the "Weed-Intervention Module" section) is to deal with the high variance delay between the instant when the image of the first camera is acquired and when a target is sensed by the detection system. To address this issue, a novel tracking system was developed. The inputs were the images and the coordinates of the targets given by the classifier (see the "Crop and Weed Detection" section) in the images of the detection camera (see Figure 3). The outputs were the trigger time and position for the actuators. The main steps are illustrated in Figure 3:

1) The intracamera tracking module estimates the camera pose and the 3D scene map using VO direct methods.
2) After receiving the delayed classification results and scene structures, the object initializer and updater module creates the templates of the received objects, propagates their updated poses, and accumulates their labels.

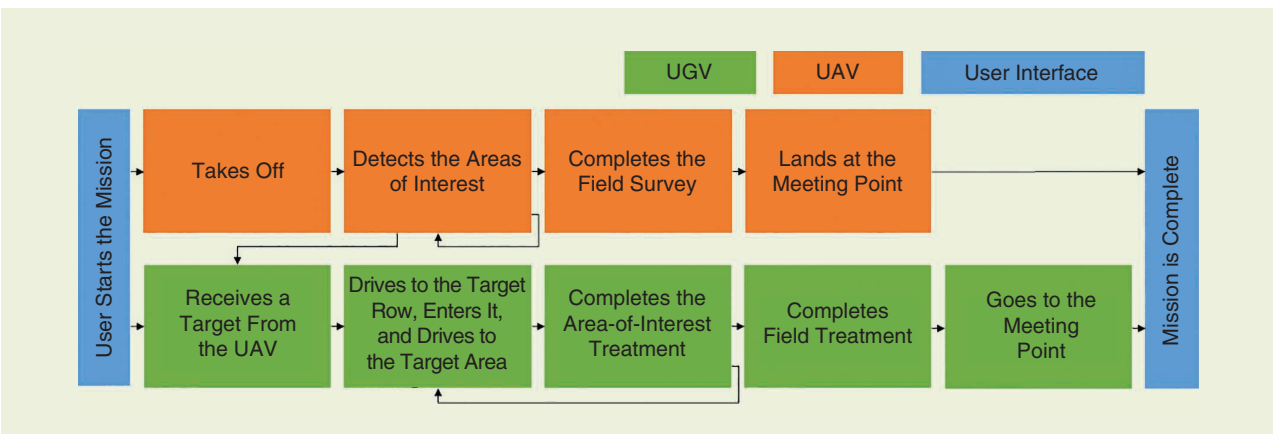

Figure 15. An example of a coordinated mission.

3) To prevent destroying a misclassified crop, a naive Bayes classifier validates its classification based on the accumulated labels.
4) Once a new weed object moves into the tracking camera's FoV, intercamera tracking performs illumination-robust direct tracking to find its new pose and create a new template for intracamera tracking.
5) After repeated intracamera tracking, updating, and intercamera tracking, the weed finally approaches the end effector, where the control algorithm predicts the trigger time and position of actuation for intervention.

The novelty in this module resides in the intra- and intercamera tracking [38] (Figure 16).

Intracamera Tracking

Unlike conventional multiobject tracking algorithms, our proposed VO approach recovers the 3D scene structure before obtaining object information, then formulates each template as a combination of trimmed image and inverse depth map for later tracking upon the arrival of the classification results. This strategy guarantees a constant-time operation despite the change of the number of tracked objects.

Intercamera Tracking

Taking advantage of the fact that only the 2D positions of weeds in the image space were of interest, we extracted the small-frame template of each weed combined with a global illumination-invariant cost to perform a local-image alignment. Then, the weed center and its template boundary were transformed into the current frame using the pose estimate to generate a new template for intracamera tracking. To be robust to changes of viewpoint, the retrieval of weeds objects was achieved using 2D/3D direct template-based matching.

To evaluate the mechanical weed removal, real leaves with an average radius of 10 mm were chosen as targets, including the successfully stamped ones. To evaluate selective spraying, we set up a webcam to monitor the targets after spraying. These experiments and their results are illustrated in Figure 17, where we can observe that the successful treatment rate is almost invariant with the speed in both flat and rough field ground.

Agronomic Impacts in Sugar Beet Crops

The potential impacts of the Flourish methodologies have been investigated through a three-year field campaign

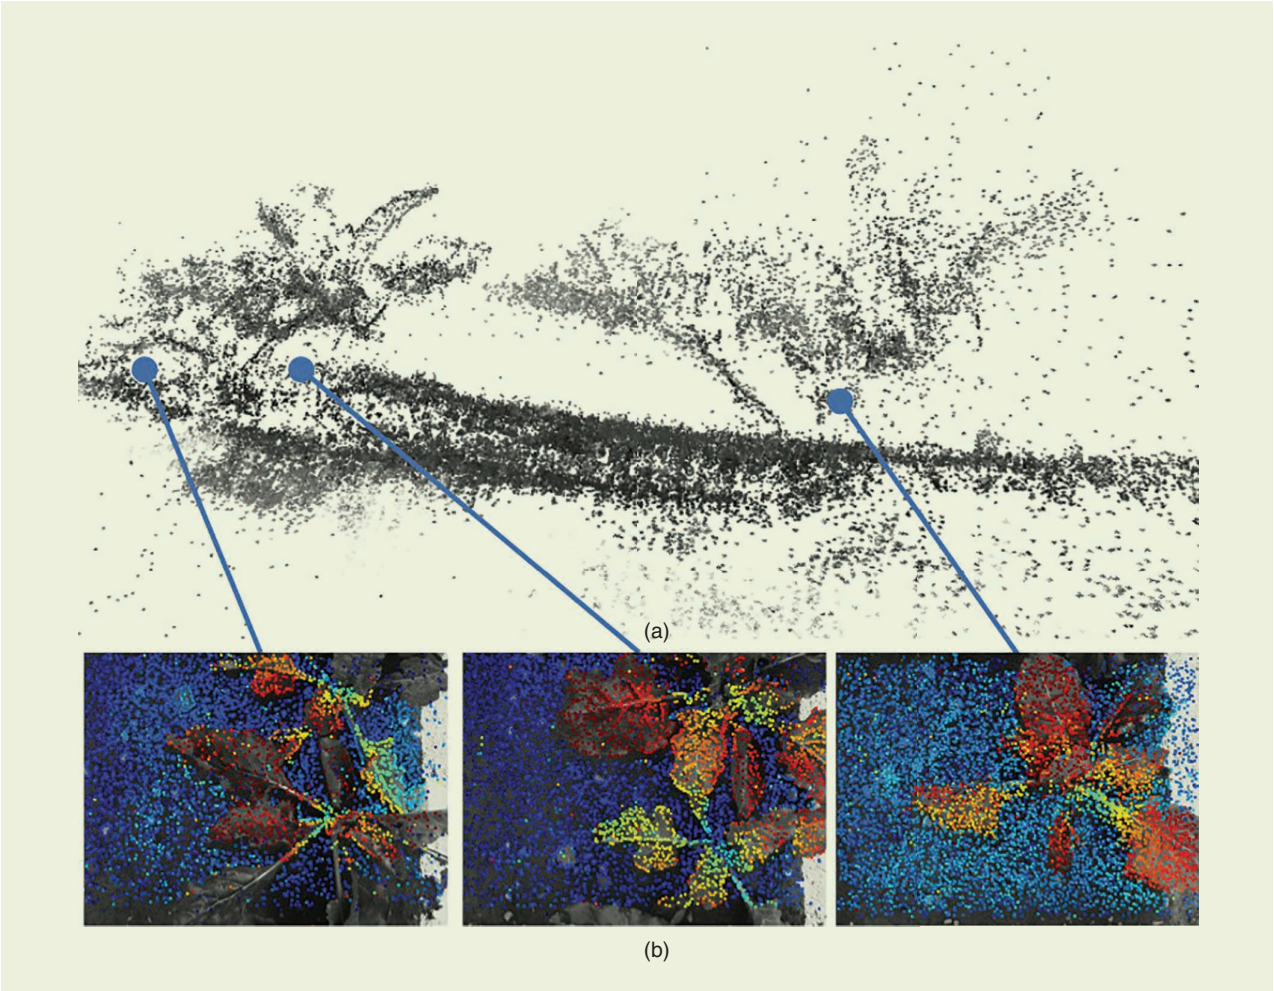

Figure 16. (a) An example of a reconstructed inverse depth map and (b) a 3D point cloud of plants and the ground surface from our proposed intracamera-tracking algorithm.

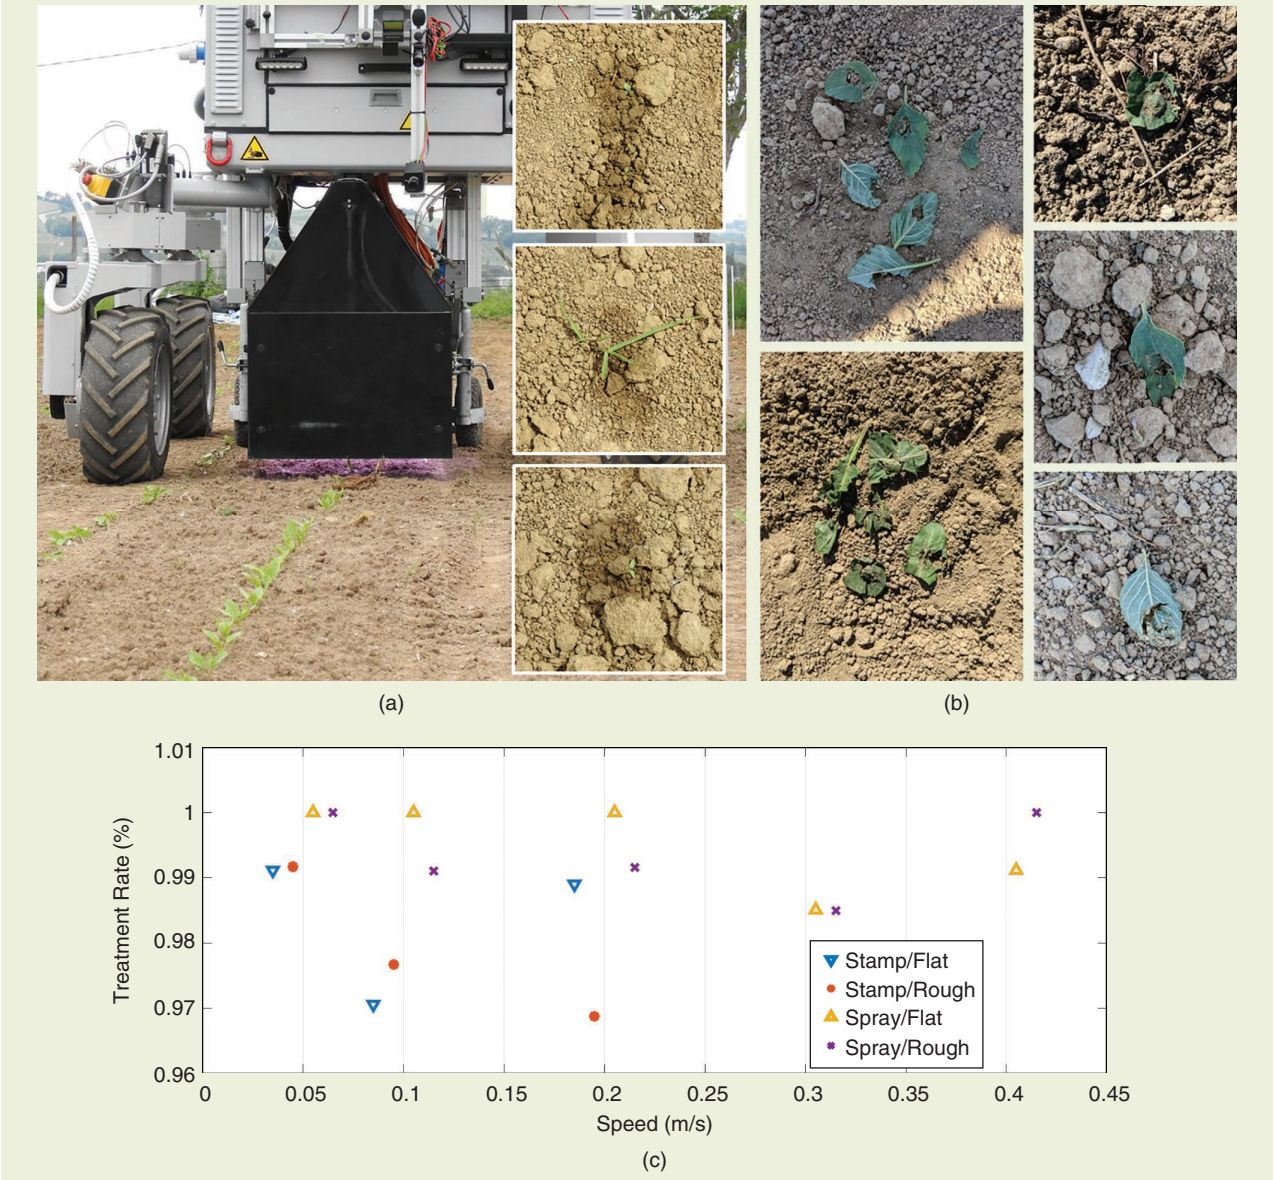

Figure 17. (a) Experiments in a real environment for spraying evaluation and (b) in a real field with fake weeds for stamping evaluation. (c) The treatment rates results for both stamping and spraying in rough and flat environments, respectively.

carried out in Italy. We compared traditional full-field herbicide treatments with a chemical weeding system, targeting the areas with higher weed density and thereby simulating the Flourish selective intervention. We performed four trials on sugar beet experimental plots grown under the same agronomic conditions and on the same site. For each year, four 10-m^2 plots (one plot per trial) were used; in 2017 and 2018, each trial was replicated three times.

For each trial, different preemergence treatments (PRTs) and postemergence treatments (POTs) were performed before and after seedling emergence using dicotyledons/monocotyledons herbicides. Table 1 lists the percentages of plots subjected to treatment (the PRT and POT columns).

At the end of each crop cycle, the sugar beet roots of each plot were harvested and weighed. Representative samples from each plot were delivered to the Agency for Agro-food Sector Services of the Marche Region laboratories for the refractometric estimation of the average root sucrose content (Brix degrees). The amount of sucrose produced in

Table 1. The treatments and tons of sucrose per hectare produced during the 2016–2018 trials.

Trials	PRT	POT	2016	2017	2018
Trial A	100%	100%	15.2	8.4	10.1
Trial B	100%	30%	12.3	8.5	11.5
Trial C	100%	0	11.7	4.9	11.2
Trial D	0	0	11	3.4	6.6

each experimental plot was then computed and related to the agronomic surface unit (the tons of sucrose per hectare). The results are reported in Table 1. The 2017 and 2018 results were calculated for each trial by averaging the results of the three replicates.

The results suggest that sugar beet selective POT (trial B), when associated with an ordinary PRT, represents a sustainable alternative to conventional POT full-field treatments (trial A) because it allows for achieving comparable sucrose production levels while reducing chemical inputs. In fact, compared to trial A, the average sucrose production of trial B was only 3.9% lower over the three years. At the same time, both the selective and full-field POTs turned out to be an effective support to production, especially in the case of crops subjected to marked grass-like weeds pressure. Traditional full-field PRT was essential in controlling weeds (mainly dicotyledons) in a slow-growing sugar beet crop (trial D's average production was 24.5% lower than that of trial C).

Besides the impacts on production and the environment, the introduction of precise agriculture technologies into ordinary cultural practices could also have a larger-scale effect on the farming sector. To support this intuition, a participatory evaluation (a metaplan model and a strengths, weaknesses, opportunities, and threats analysis) was carried out within the project involving a panel of 17 stakeholders professionally operating in the farming sector. This evaluation highlighted that the application of such technologies would potentially be able to improve the efficiency, efficacy, and safety of farming operations and so reduce labor costs and provide more information on crops, field structure, and meteorological events to increase farming-environmental sustainability and food safety.

Open Source Software and Data Sets

Many of the methods presented in the previous section have been released as open source software with downloadable links reported in the corresponding papers:

- a modified version of the DJI Onboard ROS Software Development Kit [39]
- plant-stress phenotyping data set and analysis software [40] (see the "UAV Localization and Mapping" section)
- the IPP framework (see the "UAV Mission Planning and Navigation" section) used for terrain monitoring [33] (https://github.com/ethz-asl/waypoint_navigator)
- Multi-Cue Agricultural Positioning System (see the "UGV Global Positioning and Mapping" section) [28]
- AgriColMap (see the "Cooperative UAV–UGV Environment Modeling" section) [29]
- algorithms for synchronizing clocks (https://github.com/ethz-asl/cuckoo_time_translator).

We also created and made publicly available several novel data sets:

- Sugar Beets 2016: a novel, vast, long-term data set of a sugar beet field [41]
- Flourish Sapienza data sets [42]: a collection of data sets, with related ground truths, acquired from farming robots
- a data set of four- and five-channel multispectral aerial images dedicated to plant-semantic segmentation [18], [43]
- pixel-wise, ground-truthed sugar beet and weed data sets collected from a controlled field experiment [18] (https://goo.gl/UK2pZq)
- a WeedMap data set [43], which contains high-fidelity, large-scale, and spatiotemporal, multispectral images.

Conclusions

The main goal of the Flourish research project was to develop an adaptable robotic solution for precision farming by combining the aerial survey capabilities of a small, autonomous UAV with a multipurpose agricultural UGV. In this article, we presented an overview of the custom-built hardware solutions and the methods and algorithms developed for the project, which were tailored for cooperation between aerial and ground robots. We demonstrated a successful in-field intervention task integrating the various modules.

We believe that the proposed solutions represent, from several points of view, a step forward in the state of the art of robotic systems applied to precision agriculture, with solutions that are easily applicable to a wide range of robots, farm-management activities, and crop types.

Acknowledgments

This work was supported by the European Commission under grant H2020-ICT-644227-Flourish and by the Swiss State Secretariat for Education, Research and Innovation under contract 15.0029. Alberto Pretto is the corresponding author. Roland Siegwart was the Flourish project coordinator.

References

[1] Flourish Project. Accessed: Aug. 13, 2020. [Online]. Available: http://flourish-project.eu

[2] Rhea Project. Accessed: Aug. 13, 2020. [Online]. Available: http://www.rhea-project.eu/

[3] A. Gasparri, G. Ulivi, N. Bono Rossello, and E. Garone, "The H2020 project Pantheon: Precision farming of hazelnut orchards," in *Proc. Convegno Automatica.it*, Sept. 2018.

[4] P. Astolfi, A. Gabrielli, L. Bascetta, and M. Matteucci, "Vineyard autonomous navigation in the Echord++ grape experiment," *IFAC-PapersOnLine*, vol. 51, no. 11, pp. 704–709, Jan. 2018. doi: 10.1016/j.ifacol.2018.08.401.

[5] Sweeper Project. Accessed: Aug. 13, 2020. [Online]. Available: http://www.sweeper-robot.eu/

[6] W. Dong, P. Roy, and V. Isler, "Semantic mapping for orchard environments by merging two-sides reconstructions of tree rows," *J. Field Robot.*, vol. 37, no. 1, pp. 97–121, 2020. doi: 10.1002/rob.21876.

[7] S. W. Chen et al., "Counting apples and oranges with deep learning: A data-driven approach," *IEEE Robot. Automat. Lett.*, vol. 2, no. 2, pp. 781–788, Apr. 2017. doi: 10.1109/LRA.2017.2651944.

[8] R. Ehsani, D. Wulfsohn, J. Das, I. Lagos, and Z. Ins, "Yield estimation; a low-hanging fruit for application of small UAS," *Resource, Eng. Technol. Sustain. World*, vol. 23, no. 8, pp. 16–18, 2016.

[9] J. Das et al., "Devices, systems, and methods for automated monitoring enabling precision agriculture," in *Proc. IEEE Int. Conf. Automation Science and Engineering (CASE)*, 2015, pp. 462–469. doi: 10.1109/CoASE.2015.7294123.

[10] EcoRobotix. Accessed: Aug. 13, 2020. [Online]. Available: https://www.ecorobotix.com/en/

[11] Blueriver Technology. Accessed: Aug. 13, 2020. [Online]. Available: http://about.bluerivertechnology.com/

[12] Saga Robotics. Accessed: Aug. 13, 2020. [Online]. Available: https://sagarobotics.com/

[13] J. Nikolic et al., "A synchronized visual-inertial sensor system with FPGA pre-processing for accurate real-time SLAM," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2014, pp. 431–437. doi: 10.1109/ICRA.2014.6906892.

[14] P. Lottes, M. Höferlin, S. Sander, and C. Stachniss, "Effective vision-based classification for separating sugar beets and weeds for precision farming," *J. Field Robot.*, vol. 34, no. 6, pp. 1160–1178, 2017. doi: 10.1002/rob.21675.

[15] P. Lottes, R. Khanna, J. Pfeifer, R. Siegwart, and C. Stachniss, "UAV-based crop and weed classification for smart farming," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017, pp. 3024–3031. doi: 10.1109/ICRA.2017.7989347.

[16] P. Lottes, J. Behley, A. Milioto, and C. Stachniss, "Fully convolutional networks with sequential information for robust crop and weed detection in precision farming," *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 3097–3104, 2018. doi: 10.1109/LRA.2018.2846289.

[17] P. Lottes and C. Stachniss, "Semi-supervised online visual crop and weed classification in precision farming exploiting plant arrangement," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2017, pp. 5155–5161. doi: 10.1109/IROS.2017.8206403.

[18] I. Sa et al., "weedNet: Dense semantic weed classification using multispectral images and MAV for smart farming," *IEEE Robot. Automat. Lett.*, vol. 3, no. 1, pp. 588–595, 2018. doi: 10.1109/LRA.2017.2774979.

[19] M. D. Cicco, C. Potena, G. Grisetti, and A. Pretto, "Automatic model based dataset generation for fast and accurate crop and weeds detection," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2017, pp. 5188–5195. doi: 10.1109/IROS.2017.8206408.

[20] C. M. Hoffmann, "Changes in nitrogen composition of sugar beet varieties in response to increasing nitrogen supply," *J. Agronomy Crop Sci.*, vol. 191, no. 2, pp. 138–145, 2005. doi: 10.1111/j.1439-037X.2004.00149.x.

[21] A. Walter et al., "A robotic approach for automation in crop management," in *Proc. Int. Conf. Precision Agriculture*, 2018.

[22] A. Walter, R. Finger, R. Huber, and N. Buchmann, "Opinion: Smart farming is key to developing sustainable agriculture," *Proc. Nat. Acad. Sci.*, vol. 114, no. 24, pp. 6148–6150, 2017. doi: 10.1073/pnas.1707462114.

[23] R. Finger, S. M. Swinton, N. E. Benni, and A. Walter, "Precision farming at the nexus of agricultural production and the environment," *Annu. Rev. Resource Econ.*, vol. 11, no. 1, pp. 313–335, 2019. doi: 10.1146/annurev-resource-100518-093929.

[24] F. Argento, T. Anken, F. Liebisch, and A. Walter, "Crop imaging and soil adjusted variable rate nitrogen application in winter wheat," in *Proc. Precision Agriculture 19*, 2019, pp. 511–517. doi: 10.3920/978-90-8686-888-9_63.

[25] M. Bloesch, S. Omari, M. Hutter, and R. Siegwart, "Robust visual inertial odometry using a direct EKF-based approach," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2015, pp. 298–304. doi: 10.1109/IROS.2015.7353389.

[26] S. Lynen, M. Achtelik, S. Weiss, M. Chli, and R. Siegwart, "A robust and modular multi-sensor fusion approach applied to MAV navigation," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2013, pp. 3923–3929. doi: 10.1109/IROS.2013.6696917.

[27] T. Schneider et al., maplab: An open framework for research in visual-inertial mapping and localization, 2018. [Online]. Available: arXiv:1711.10250

[28] M. Imperoli, C. Potena, D. Nardi, G. Grisetti, and A. Pretto, "An effective multi-cue positioning system for agricultural robotics," *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 3685–3692, Oct. 2018. doi: 10.1109/LRA.2018.2855052.

[29] C. Potena, R. Khanna, J. Nieto, R. Siegwart, D. Nardi, and A. Pretto, "AgriColMap: Aerial–ground collaborative 3D mapping for precision farming," *IEEE Robot. Automat. Lett.*, vol. 4, no. 2, pp. 1085–1092, 2019. doi: 10.1109/LRA.2019.2894468.

[30] Y. Hu, R. Song, and Y. Li, "Efficient coarse-to-fine patch match for large displacement optical flow," in *Proc. Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 5704–5712. doi: 10.1109/CVPR.2016.615.

[31] A. Myronenko and X. Song, "Point set registration: Coherent point drift," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 12, pp. 2262–2275, 2010. doi: 10.1109/TPAMI.2010.46.

[32] N. Chebrolu, T. Läbe, and C. Stachniss, "Robust Long-term registration of UAV images of crop fields for precision agriculture," *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 3097–3104, 2018. doi: 10.1109/LRA.2018.2849603.

[33] M. Popović et al., "An informative path planning framework for UAV-based terrain monitoring," *Autonom. Robots*, vol. 44, no. 6, pp. 889–911, 2020. doi: 10.1007/s10514-020-09903-2.

[34] W. Winterhalter, F. Fleckenstein, C. Dornhege, and W. Burgard, "Crop row detection on tiny plants with the pattern through transform," *IEEE Robot. Automat. Lett.*, vol. 3, no. 4, pp. 3394–3401, 2018. doi: 10.1109/LRA.2018.2852841.

[35] F. Fleckenstein, C. Dornhege, and W. Burgard, "Efficient path planning for mobile robots with adjustable wheel positions," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017, pp. 2454–2460. doi: 10.1109/ICRA.2017.7989286.

[36] F. Fleckenstein, W. Winterhalter, C. Dornhege, C. Pradalier, and W. Burgard, "Smooth local planning incorporating steering constraints," in *Proc. 12th Conf. Field and Service Robotics (FSR)*, 2019.

[37] C. Pradalier, "A task scheduler for ROS," Jan. 2017. [Online]. Available: https://hal.archives-ouvertes.fr/hal-01435823

[38] X. Wu, S. Aravecchia, and C. Pradalier, "Design and implementation of computer vision based in-row weeding system," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2019, pp. 4218–4224. doi: 10.1109/ICRA.2019.8793974.

[39] I. Sa et al., "Build your own visual-inertial drone: A cost-effective and open-source autonomous drone," *IEEE Robot. Automat. Mag.*, vol. 25, no. 1, pp. 89–103, 2018. doi: 10.1109/MRA.2017.2771326.

[40] R. Khanna, L. Schmid, A. Walter, J. Nieto, R. Siegwart, and F. Liebisch, "A spatio temporal spectral framework for plant stress phenotyping," *Plant Meth.*, vol. 15, no. 1, 2019. doi: 10.1186/s13007-019-0398-8.

[41] N. Chebrolu, P. Lottes, A. Schaefer, W. Winterhalter, W. Burgard, and C. Stachniss, "Agricultural robot dataset for plant classification, localization and mapping on sugar beet fields," *J. Robot. Res.*, vol. 36, no. 10, pp. 1045–1052, 2017. doi: 10.1177/0278364917720510.

[42] "Flourish Sapienza Datasets," Sapienza University of Rome, Italy. [Online]. Available: https://www.dis.uniroma1.it/~labrococo/fsd/

[43] I. Sa et al., "WeedMap: A large-scale semantic weed mapping framework using aerial multispectral imaging and deep neural network for precision farming," *Remote Sens.*, vol. 10, no. 9, 2018, Art. no. 1423. doi: 10.3390/rs10091423.

Alberto Pretto, Department of Computer, Control, and Management Engineering, Sapienza University of Rome, Italy, and IT+Robotics Srl, Padova, Italy. Email: alberto.pretto@it-robotics.it.

Stéphanie Aravecchia, International Research Lab 2958, Georgia Institute of Technology-National Center for Scientific Research, Metz, France. Email: stephanie.aravecchia@georgiatech-metz.fr.

Wolfram Burgard, Department of Computer Science, University of Freiburg, Germany; Toyota Research Institute, Los Altos, California. Email: burgard@informatik.uni-freiburg.de.

Nived Chebrolu, Photogrammetry and Robotics Lab, University of Bonn, Germany. Email: nivedchebrolu@gmail.com.

Christian Dornhege, Department of Computer Science, University of Freiburg, Germany. Email: dornhege@informatik.uni-freiburg.de.

Tillmann Falck, Robert Bosch GmbH, Corporate Research, Renningen, Germany. Email: tillmann.falck@de.bosch.com.

Freya Fleckenstein, Department of Computer Science, University of Freiburg, Germany. Email: fleckenf@informatik.uni-freiburg.de.

Alessandra Fontenla, Agency for Agrofood Sector Services of the Marche Region, Osimo, Italy. Email: fontenla_alessandra@assam.marche.it.

Marco Imperoli, Department of Computer, Control, and Management Engineering, Sapienza University of Rome, Italy. Email: imperoli@diag.uniroma1.it.

Raghav Khanna, Autonomous Systems Lab, Department of Mechanical and Process Engineering, ETH Zürich, Switzerland. Email: raghav.khanna@mavt.ethz.ch.

Frank Liebisch, Department of Environmental Systems Science, Institute of Agricultural Sciences, ETH Zürich, Switzerland. Email: frank.liebisch@usys.ethz.ch.

Philipp Lottes, Photogrammetry and Robotics Lab, University of Bonn, Germany. Email: philipp.lottes@igg.uni-bonn.de.

Andres Milioto, Photogrammetry and Robotics Lab, University of Bonn, Germany. Email: amilioto@uni-bonn.de.

Daniele Nardi, Department of Computer, Control, and Management Engineering, Sapienza University of Rome, Italy. Email: nardi@diag.uniroma1.it.

Sandro Nardi, Agency for Agrofood Sector Services of the Marche Region, Osimo, Italy. Email: nardi_sandro@assam.marche.it.

Johannes Pfeifer, Department of Environmental Systems Science, Institute of Agricultural Sciences, ETH Zürich, Switzerland. Email: johannes.pfeifer@ble.de.

Marija Popović, Autonomous Systems Lab, Department of Mechanical and Process Engineering, ETH Zürich, Switzerland; Smart Robotics Lab, Department of Computing, Imperial College London, United Kingdom. Email: m.popovic@imperial.ac.uk.

Ciro Potena, Department of Computer, Control, and Management Engineering, Sapienza University of Rome, Italy. Email: potena@diag.uniroma1.it.

Cédric Pradalier, International Research Lab 2958, Georgia Institute of Technology-National Center for Scientific Research, Metz, France. Email: cedric.pradalier@gmail.com.

Elisa Rothacker-Feder, Robert Bosch GmbH, Corporate Research, Renningen, Germany. Email: elisa.rothacker-feder@de.bosch.com.

Inkyu Sa, Autonomous Systems Lab, Department of Mechanical and Process Engineering, ETH Zürich, Switzerland; CSIRO Data61, Australia. Email: enddl22@gmail.com.

Alexander Schaefer, Department of Computer Science, University of Freiburg, Germany. Email: aschaef@cs.uni-freiburg.de.

Roland Siegwart, Autonomous Systems Lab, Department of Mechanical and Process Engineering, ETH Zürich, Switzerland. Email: rsiegwart@ethz.ch.

Cyrill Stachniss, Photogrammetry and Robotics Lab, University of Bonn, Germany. Email: cyrill.stachniss@igg.uni-bonn.de.

Achim Walter, Department of Environmental Systems Science, Institute of Agricultural Sciences, ETH Zürich, Switzerland. Email: achim.walter@usys.ethz.ch.

Wera Winterhalter, Department of Computer Science, University of Freiburg, Germany. Email: winterhw@informatik.uni-freiburg.de.

Xiaolong Wu, International Research Lab 2958 Georgia Institute of Technology-National Center for Scientific Research, Metz, France. Email: xwu@gatech.edu.

Juan Nieto, Autonomous Systems Lab, Department of Mechanical and Process Engineering, ETH Zürich, Switzerland. Email: jnieto@ethz.ch.