

# AgriColMap: Aerial-Ground Collaborative 3D Mapping for Precision Farming

Ciro Potena<sup>1</sup>, Raghav Khanna<sup>2</sup>, Juan Nieto<sup>2</sup>, Roland Siegwart<sup>2</sup>, Daniele Nardi<sup>1</sup>, and Alberto Pretto<sup>1</sup>

**Abstract**—The combination of aerial survey capabilities of Unmanned Aerial Vehicles with targeted intervention abilities of agricultural Unmanned Ground Vehicles can significantly improve the effectiveness of robotic systems applied to precision agriculture. In this context, building and updating a common map of the field is an essential but challenging task. The maps built using robots of different types show differences in size, resolution and scale, the associated geolocation data may be inaccurate and biased, while the repetitiveness of both visual appearance and geometric structures found within agricultural contexts render classical map merging techniques ineffective.

In this paper we propose AgriColMap, a novel map registration pipeline for that leverages a grid-based multi-modal environment representation which includes a vegetation index map and a Digital Surface Model. We cast the data association problem between maps built from UAVs and UGVs as a multi-modal, large displacement dense optical flow estimation. The dominant, coherent flows, selected using a voting scheme, are used as point-to-point correspondences to infer a preliminary non-rigid alignment between the maps. A final refinement is then performed, by exploiting only meaningful parts of the registered maps.

We evaluate our system using real world data for 3 fields with different crop species. The results show that our method outperforms several state of the art map registration and matching techniques by a large margin, and has a higher tolerance to large initial misalignments. We release an implementation of the proposed approach along with the acquired datasets with this paper.

**Index Terms**—Robotics in Agriculture and Forestry, Mapping, Multi-Robot Systems

## SUPPLEMENTARY MATERIAL

[www.dis.uniroma1.it/~labrococo/fsd/agricolmap\\_sup.pdf](http://www.dis.uniroma1.it/~labrococo/fsd/agricolmap_sup.pdf)

The datasets and our C++ implementation are available at:

[www.dis.uniroma1.it/~labrococo/fsd](http://www.dis.uniroma1.it/~labrococo/fsd)

## I. INTRODUCTION

**C**OOOPERATION between aerial and ground robots undoubtedly offers benefits to many applications, thanks to the complementarity of the characteristics of these robots [1]. This is especially useful in robotic systems applied to precision

Manuscript received: September, 10, 2018; Revised November, 07, 2018; Accepted January, 02, 2019. This paper was recommended for publication by Editor Youngjin Choi upon evaluation of the Associate Editor and Reviewers' comments.

This work was supported by the EC under Grant H2020-ICT-644227-Flourish and by the Swiss State Secretariat for Education, Research and Innovation under contract number 15.0029. <sup>1</sup>Potena, Nardi and Pretto are with Faculty of Department of Computer, Control, and Management Engineering "Antonio Ruberti", Sapienza University of Rome, Italy. Email: {potena, nardi, pretto}@diag.uniroma1.it. <sup>2</sup>Khanna, Nieto and Siegwart are with the Autonomous Systems Lab, ETH Zurich, Switzerland. Email: raghav.khanna@mavt.ethz.ch, jnieto@ethz.ch, r.siegwart@ieee.org




Fig. 1: An overview of AgriColMap. Both the UGV and UAV generate, using data gathered from their onboard cameras, colored point clouds of the cultivated field. The proposed method aims to accurately merge these maps by means of an affine transformation that registers the UGV submap (red rectangular area) into the UAV aerial map (blue rectangular area), taking into account possible scale discrepancies.

agriculture, where the areas of interest are usually vast. A UAV allows rapid inspection of large areas [2], and then share information such as crop health or weeds distribution - indicators of areas of interest with an agricultural UGV. The ground robot can operate for long periods of time, carry high payloads, perform targeted actions, such as fertilizer application or selective weed treatment, on the areas selected by the UAV. The robots can also cooperate to generate 3D maps of the environment, e.g., annotated with parameters, such as crop density and weed pressure, suitable for supporting the farmer's decision making. The UAV can quickly provide a coarse reconstruction of a large area, that can be updated with more detailed and higher resolution map portions generated by the UGV visiting selected areas.

All the above applications assume that both UAVs and UGVs can share information using a unified environment model with centimeter-level accuracy, i.e. an accurate *shared map* of the field. There are two classes of methods designed to generate multi-robot environment representations: (i) multi-robot Simultaneous Localization and Mapping (SLAM) algorithms (e.g., [3], [4]), that concurrently build a single map by fusing raw measurements or small local maps generated from multiple robots; (ii) map registration algorithms (e.g., [5], [6]) that align and merge maps independently generated by each robot into a unified map. On the one hand, the lack of distinctive visual and 3D landmarks in an agricultural field,

along with the difference in the robots' point-of-views (e.g., Fig. 2), prevent direct employment of standard multi-robot SLAM pipelines, either based on visual or geometric features. On the other hand, merging maps independently generated by the UAVs and UGVs in an agricultural environment is also a complex task, since maps are usually composed of similar, repetitive patterns that easily confuse conventional data association methods [7]. Furthermore, due to inaccuracies in the map building process, the merged maps are usually affected by local inconsistencies, missing data, occlusions, and global deformations such as directional scale errors, that negatively affect the performance of standard alignment methods. Geolocation information associated with (i) sensor readings or (ii) maps often can't solve the limitations of conventional methods in agricultural environments, since the location and orientation accuracy provided by standard reference sensors<sup>1</sup> [8] is not suitable to prevent such system from converging towards sub-optimal solutions (see Sec. V)




Fig. 2: Pictures of the same portion of field seen from the UAV point-of-view (left) and from the UGV point-of-view (right). The local crop arrangement geometry, such as the missing crop plants, is generally not visible from the UGV point-of-view. The yellow solid lines represent an example of manually annotated correct point matches. It is important to underline the complexity required in obtaining correct data association, also from a human point-of-view. The fiducial markers on the field have been used to compute the ground truth alignments between the maps.

In this paper, we introduce AgriColMap, an Aerial-Ground Collaborative 3D Mapping pipeline, which provides an effective and robust solution to the cooperative mapping problem with heterogeneous robots, specifically designed for farming scenarios. We address this problem by proposing a non-rigid map registration strategy able to deal with maps with different resolutions, local inconsistencies, global deformations, and relatively large initial misalignments. We assume that both a UAV and a UGV can generate a colored, geotagged point cloud of a target farm environment, e.g., by means of photogrammetry-based 3D reconstruction. (Fig. 1). To solve the data association problem between the input point clouds, we propose to switch from a 3D problem to a 2D one, solved by using a global, 2D dense matching approach. The key intuition behind this choice is that points belonging to a cloud locally share similar displacement vectors that associate such points with points in the other cloud. Therefore, the idea is to employ a regularized 2D matching strategy that penalizes the displacement vectors discontinuities for each point neighborhood<sup>2</sup>. With this formulation, good correspondences are iteratively improved and spread through cooperative search among neighboring points. This approach

<sup>1</sup>Global Positioning Systems (GPSs) and Attitude and Heading Reference Systems (AHRSs)

<sup>2</sup>In other words, a regularized matching enforces the smoothness of the displacement vectors for neighboring points.

has been inspired by the Large displacement Dense Optical Flow (LDOF) problem in computer vision and, actually, we cast our data association problem as a LDOF problem. To this end, we convert the colored point clouds into a more suited, multimodal environment representation that allows one to exploit two-dimensional approaches and to highlight both the semantic and the geometric properties of the target map. The former is represented by a vegetation index map, while the latter through a Digital Surface Model (DSM). More specifically, we transform each input point cloud into a grid representation, where each cell stores (i) the Excess Green index (ExG) and (ii) the local surface height information (e.g., the height of the plants, soil, etc.). Then, we use the data provided by the GPS and the AHRS to extract an initial guess of the relative displacement and rotation between grid maps to match. Hence, we compute a dense set of point-to-point correspondences between matched maps, exploiting a modified, state-of-the-art LDOF system [9], tailored to the precision agriculture context. To adapt this algorithm to our environment representation, we propose to use a different cost function that involves both the ExG information and the local structure geometry around each cell. We select, using a voting scheme, the bigger subset of correspondences with coherent, similar flows, to be used to infer a preliminary alignment transformation between the maps. In order to deal with directional scale errors, we use a non-rigid point-set registration algorithm to estimate an affine transformation. The final registration is obtained by performing a robust point-to-point registration over the input point clouds, pruned from all points that do not belong to vegetation. A schematic overview of the proposed approach is depicted in Fig. 3.

We report results from an exhaustive set of experiments (Sec. V) on data acquired by a UAV and a handheld camera, simulating the UGV, on crop fields in Eschikon, Switzerland. We show that the proposed approach is able to guarantee with a high probability a correct registration for an initial translational error up to 5 meters, an initial heading misalignment up to 11.5 degrees, and a directional scale error of up to 30%. We found similar registration performance across fields with three different crop species, showing that the method generalizes well to across different kinds of farms. We also report a comparison with state-of-the-art point-to-point registration and matching algorithms, showing that our approach outperforms them in all the experiments.

#### A. Related Work

The field of multi-robot cooperative mapping is a recurrent and relevant problem in literature and, as previously introduced, several solutions have been presented by means of either multi-robot SLAM algorithms or map merging/map registration strategies, in both 2D ([5], [10], [11]) and 3D ([6], [12], [13]) settings. Registration of point cloud based maps can also be considered as an instance of the more general point set registration problem [14], [15]. In this work, we mainly review methods based on map registration, since the heterogeneity

of the involved robots and the lack of distinctive visual and geometrical features on an agricultural environment prevent the employment of standard multi-robot SLAM methods; a comprehensive literature review about this class of methods can be found [16].

Map registration is a challenging problem especially when dealing with heterogeneous robots, where data is gathered from different points-of-view and with different noise characteristics. It has been intensively investigated, especially in the context of urban reconstruction with aerial and ground data. In [17], the authors focus on the problem of geo-registering ground-based multi-view stereo models by proposing a novel viewpoint-dependent matching method. Wang *et al.* [18] deal with aligning 3D structure-from-motion point clouds obtained from internet imagery with existing geographic information sources, such as noisy geotags from input Flickr photos and geotagged city models and images collected from Google Street View and Google Earth. Bódis-Szomorú *et al.* [19] propose to merge low detailed airborne point clouds with incomplete street-side point clouds by applying volumetric fusion based on a 3D tetrahedralization (3DT). Früh *et al.* [20] propose to use Digital Surface Models obtained from a laser airborne reconstruction to localize a ground vehicle equipped with 2D laser scanners and a digital camera, detailed ground-based facade models are hence merged with a complementary airborne model. Michael *et al.* [21] propose a collaborative UAV-UGV mapping approach in earthquake-damaged contexts. They merge the point clouds generated by the two robots using a 3D Iterative Closest Point (ICP) algorithm, with an initial guess provided by the (known) UAV takeoff location; the authors make the assumption that the environment is generally described by flat planes and vertical walls: also called the “Manhattan world” assumption. The ICP algorithm has also been exploited in [22] and [23]. Forster *et al.* [22] align dense 3D maps obtained by a UGV equipped with an RGB-D camera and by a UAV running dense monocular reconstruction: they obtain the initial guess alignment between the maps by localizing the UAV with respect to the UGV with a Monte Carlo Localization method applied to height-maps computed by the two robots. Hinzmann *et al.* [23] deal with the registration of dense LiDAR-based point clouds with sparse image-based point clouds by proposing a probabilistic data association approach that specifically takes the individual cloud densities into consideration. In [24], Gawel *et al.* present a registration procedure for matching LiDAR point-cloud maps and sparse vision keypoint maps by using structural descriptors.

Despite the extensive literature addressing the problem of map registration for heterogeneous robots, most of the proposed methods make strong context-based assumptions, such as the presence of structural or visual landmarks, “Manhattan world” assumptions, etc. Registering 3D maps in an agricultural setting, in some respects, is even more challenging: the environment is homogeneous, poorly structured and it usually gives rise to strong sensor aliasing. For these reasons, most of the approaches mentioned above cannot directly be applied to an agricultural scenario. Localization and mapping in an agricultural scenario is a topic that is recently gathering great

attention in the robotics community [25], [26], [8]. Most of these systems, however, deal with a single robot, and the problem of fusing maps built from multiple robots is usually not adequately addressed and a little, very recent research exists on this topic. Dong *et al.* [27] propose a spatio-temporal reconstruction framework for precision agriculture that aims to merge multiple 3D field reconstructions of the same field across time. They use single row reconstruction as a starting point for the data association, that is actually performed by using standard visual features. This method uses images acquired by a single UGV that moves in the same field at different times and, being based on visual features, cannot manage drastic viewpoint changes or large misalignments when matching aerial and ground maps. A local feature descriptor designed to deal with large viewpoint changes has been proposed by Chebrolu *et al.* in [28], encoding with such descriptor the almost static geometry of the crop arrangement in the field. Despite the promising results, this method suffers from the presence of occluded areas when switching from the UAV to the UGV point-of-views.

## B. Contributions

Our contributions are the following: (i) A map registration framework specifically designed for heterogeneous robots in an agricultural environment; (ii) To the best of our knowledge, we are the first to apply a LDOF based 3D map alignment; (iii) Extensive performance evaluations that show the effectiveness of our approach; (iv) An open-source implementation of our method and three challenging datasets with different crop species with ground truth.

## II. PROBLEM STATEMENT AND ASSUMPTIONS

Given two 3D colored point clouds of a farmland  $\mathcal{M}_A$  and  $\mathcal{M}_G$  (Fig. 3, first column), built from data gathered from a UAV and a UGV, respectively, our goal is to find a transformation  $F : \mathbb{R}^3 \rightarrow \mathbb{R}^3$  that allows to accurately align them.  $\mathcal{M}_A$  and  $\mathcal{M}_G$  can be generated, for instance, by using an off-the-shelf photogrammetry-based 3D reconstruction software applied to sequences of geotagged images. Our method makes the following assumptions:

- 1) The input maps built from UAVs and UGVs data can have different spatial resolutions but they refer to the same field, with some overlap among them;
- 2) The data used to build the maps were acquired at approximately the same time;
- 3) The maps are roughly geotagged, possibly with noisy locations and orientations;
- 4) They can be affected by local inconsistencies, missing data, and deformations, such as directional scale errors.
- 5)  $\mathcal{M}_A$  is not affected by any scale inconsistencies.

Hypotheses 1, 2, and 3 are the essential data requirements. Hypothesis 4) implies the violation of the typical rigid-body transformation assumption between the two maps: therefore, we represent  $F$  as an affine transformation that allows anisotropic (i.e., non-uniform) scaling between the maps. Hypothesis 5) is an acceptable assumption, since the map created by the UAV is usually wider than  $\mathcal{M}_G$ , and generated by using less noisy

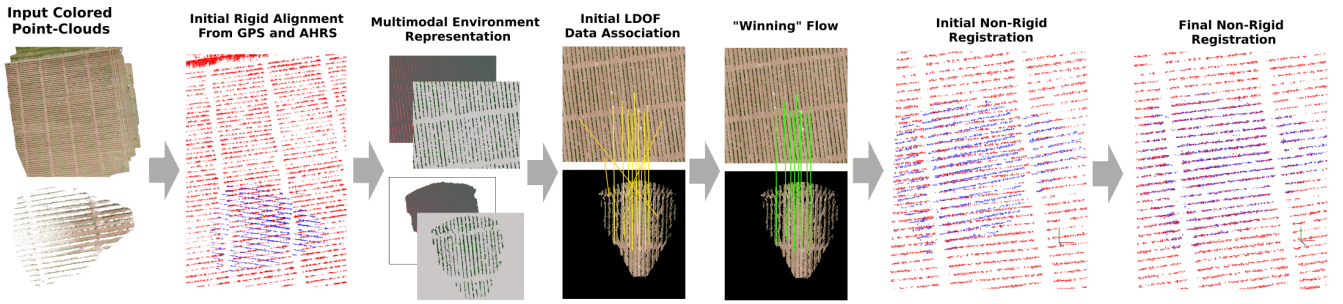


Fig. 3: Overview of the proposed approach. For visualization purposes, in column 2,7 and 8 we colored in blue and red the UGV and UAV point clouds, respectively, pruned from all points that do not belong to vegetation, according to a thresholding operator applied to the ExG index. Starting from the left side, we show: (i) the input colored point clouds gathered by the UAV and UGV; (ii) the initial noisy and biased rigid alignment provided by the GPS and the AHRS; (iii) the generated multimodal grid maps; (iv) the initial LDOF data associations, i.e. the point-to-point correspondences, in yellow; (v) the “winning” data associations (flows), in green, selected by a voting scheme; (vi) the aligned point clouds according to the initial affine transform; (vii) the final non-rigid registration after the refinement step.

GPS readings, so the scale drift effect tends to be canceled: hence, we look for a transformation that aligns  $\mathcal{M}_G$  with  $\mathcal{M}_A$  by correcting the scale errors of  $\mathcal{M}_G$  with respect to  $\mathcal{M}_A$ .

### III. DATA ASSOCIATION

In order to estimate the transformation  $F$  that aligns the two maps, we need to find a set of point correspondences,  $m_{A,G} = \{(p,q) : p \in \mathcal{M}_A, q \in \mathcal{M}_G\}$  between  $\mathcal{M}_A$  and  $\mathcal{M}_G$ , that represent points pairs belonging to the same global 3D position. As introduced before and shown in the experiments (see Sec. V), conventional sparse matching approaches based on local descriptors are unlikely to provide effective results due to the big amount of repetitive and non-distinctive patterns spread over farmlands. Instead, inspired by the fact that when the maps are misaligned, points in  $\mathcal{M}_A$  locally share a coherent “flow” towards corresponding points in  $\mathcal{M}_G$ , our method casts the data association estimation problem as a *dense, regularized*, matching approach. This problem resembles the dense optical flow estimation problem for RGB images: in this context, global methods (e.g., [29]) aim to build correspondences pixel by pixel between a pair of images by minimizing a cost function that, for each pixel, involves a data term that measures the point-wise similarity and a regularization term that fosters smoothness between nearby flows (i.e., nearby pixel to pixel associations).

#### A. Multimodal Grid Map

Our goal is to estimate  $m_{A,G}$  by computing a “dense flow” that, given an initial, noisy alignment between the maps provided by a GPS and a AHRS (Fig. 3, second column), associates points in  $\mathcal{M}_A$  with points in  $\mathcal{M}_G$ . Unfortunately, conventional methods designed for RGB images are not directly applicable to colored point clouds: we introduce here a multimodal environment representation that allows to exploit such methods while enhancing both the semantic and the geometrical properties of the target map. A cultivated field is basically a globally flat surface populated by plants. A DSM<sup>3</sup> can well approximate the field structure geometry, while a vegetation index can highlight the meaningful parts of the

field and the visual relevant patterns: in our environment representation, we exploit both these intuitions. We generate a DSM from the point cloud; for each cell of the DSM grid, we also provide an ExG index that, starting from the RGB values, highlights the amount of vegetation. More specifically, we transform a colored point cloud  $\mathcal{M}$  into a two dimensional grid map  $\mathcal{J} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$  (Fig. 3, third column), where for each cell we provide the surface height and the ExG index, with the following procedure:

- 1) We select a rectangle that bounds the target area by means of minimum-maximum latitude and longitude;
- 2) The selected area is discretized into a grid map  $\mathcal{J}$  of  $w \times h$  cells, by using a step of  $s$  meters. In practice, each of the  $w \times h$  cells represents a square of  $s \times s$  meters. Each cell is initialized with  $(0,0)$  pairs.
- 3) Remembering that  $\mathcal{M}$  is geotagged (see Sec. II), we can associate each 3D point of  $\mathcal{M}$  to one cell of  $\mathcal{J}$  just using the  $x,y$  and yaw information.
- 4) For each cell with associated at least one 3D point: (a) We compute the height as the *weighted* average of the  $z$  coordinates of the 3D points that belong to such cell; (b) We compute the the ExG index as the *weighted* average of the the ExG indexes of the 3D points that belong to such cell, where for each point  $p$  we have:

$$ExG(p) = 2g_p - r_p - b_p. \quad (1)$$

where  $r_p$ ,  $g_p$  and  $b_p$  the RGB components of the point; (c) we store the 3D global position of the nearest point in the original colored point cloud.

Both the averages use as weighting factor a circular, bivariate Gaussian distribution with standard deviation  $\sigma_{avg}$ : points with  $x,y$  coordinates close to center of the cell get a higher weight.

#### B. Multimodal Large displacement Dense Optical Flow

We generate from both the  $\mathcal{M}_A$  and  $\mathcal{M}_G$  the corresponding multimodal representations  $\mathcal{J}_A$  and  $\mathcal{J}_G$ . In the ideal case, with perfect geotags and no map deformations, a simple geotagged superimposition of the two maps should provide a perfect alignment: the “flow” that associates cells between the two

<sup>3</sup>A DSM is a raster representations of the height of the objects on a surface.

maps should be zero. Unfortunately, in the real case, due to the inaccuracies of both the geotags and the 3D reconstruction, non zero, potentially large displacements are introduced in the associations. These offsets are locally consistent but not constant for each cell, due to the reconstruction errors. To estimate the offsets map, we employ a modified version of the Coarse-to-fine PatchMatch (CPM) framework described in [9]. CPM is a recent LDOF system that provides cutting edge estimation results even in presence of very large displacements, and is more efficient than other state-of-the-art methods with similar accuracy.

For efficiency, CPM looks for the best correspondences of some seeds that are refined by means of a dense, iterative neighborhood propagation: the seeds are a set of points regularly distributed within the image. Given two images  $\mathcal{I}_0, \mathcal{I}_1 \in \mathbb{R}^2$  and a collection of seeds  $S = \{s_n\}$  at position  $\{p(s_n)\}$ , the goal of this framework is to determine the flow of each seed  $f(s_n) = M(p(s_n)) - p(s_n) \in \mathbb{R}^2$ , where  $M(p(s_n))$  is the corresponding matching position in  $\mathcal{I}_1$  for the seed  $s_n$  in  $\mathcal{I}_0$ . The flow computation for each *seed* is performed by a coarse-to-fine random search strategy by minimizing the cost function:

$$f(s_n) = \underset{f_{s_i}}{\operatorname{argmax}}(C(f(s_i))), s_i \in \{s_m\} \quad (2)$$

where  $C(f(s_i))$  denotes the match cost between the patch centered at  $p(s_m)$  in  $\mathcal{I}_0$  and the patch centered in  $p(s_m) + f(\cdot)$  in  $\mathcal{I}_1$ . For a comprehensive description of the flow estimation pipeline, we refer the reader to [9].

Our goal is to use the CPM algorithm to compute the flow between  $\mathcal{J}_A$  and  $\mathcal{J}_G$ . To exploit the full information provided by our grid maps (see Sec. III-A), we modified the CPM matching cost in order to take into account both the height and ExG channels. We split the cost function in two terms:

$$C_{flow}(f(s_i)) = \alpha \cdot C_{DY}(f(s_i)) + \beta \cdot C_{FPFH}(f(s_i)) \quad (3)$$

$C_{DY}(f(s_i))$  is the DAISY [30] based match cost as in the original CPM algorithm: in our case the DAISY descriptors have been computed from the ExG channel of  $\mathcal{J}_A$  and  $\mathcal{J}_G$ .  $C_{FPFH}(f(s_i))$  is a match cost computed using the height channel. We chose the Fast Point Feature Histograms (FPFH) [31] descriptor for this second term: the FPFH descriptors are robust multi-dimensional features which describe the *local* geometry of a point cloud<sup>4</sup>, in our case they are computed from the organized point cloud generated from the height channel of  $\mathcal{J}_A$  and  $\mathcal{J}_G$ . The parameters  $\alpha$  and  $\beta$  are the weighting factors of the two terms. As in [9], the patch-based matching cost is chosen to be the sum of the absolute difference over all the 128 and 32 dimensions of the DAISY and FPFH flows, respectively, at the matching points. The proposed cost function takes into account both the visual appearance and the local 3D structure of the plants.

Once we have computed the dense flow between  $\mathcal{J}_A$  and  $\mathcal{J}_G$  (Fig. 3, fourth column), we extract the largest set of coherent flows by employing a voting scheme inspired by the classical

Hough transform with discretization step  $t_f$ ; these flows define a set of point-to-point matches  $m_{A,G}$  that will be used to infer a preliminary alignment (Fig. 3, fifth column).

#### IV. NON-RIGID REGISTRATION

The estimation of the non-rigid transformation between the maps is addressed in two steps. A preliminary affine transformation  $\hat{F}$  is computed by solving a non-rigid registration problem with known point-to-point correspondences. We compute  $\hat{F} = (\hat{s}\hat{R}|\hat{t})$  by solving an optimization problem with cost function the sum of the squared distances between corresponding points (Fig. 3, sixth column):

$$C_{reg}(\hat{F}) = \sum_{i=0}^N \|p_i - \hat{s}\hat{R}q_i - \hat{t}\|^2 \quad (4)$$

with  $(p_i, q_i) \in m_{A,G}$ ,  $N$  the cardinality of  $m_{A,G}$ ,  $\hat{R}$  and  $\hat{t}$  the rotation matrix and the translation vector, and  $\hat{s}$  is a scaling vector. To estimate the final registration, we firstly select from the input colored point clouds  $\mathcal{M}_A$  and  $\mathcal{M}_G$  two subsets,  $\mathcal{M}_A^{veg}$  and  $\mathcal{M}_G^{veg}$ , that includes only points that belong to vegetation. The selection is performed by using an ExG based thresholding operator over  $\mathcal{M}_A$  and  $\mathcal{M}_G$ . This operation enhances the morphological information of the vegetation, while reducing the size of the point clouds to be registered. We finally estimate the target affine transformation  $F$  by exploiting the coherent point drift (CPD) [33] point set registration algorithm over the point clouds  $\mathcal{M}_A^{veg}$  and  $\mathcal{M}_G^{veg}$ , using  $\hat{F}$  as initial guess transformation.

#### V. EXPERIMENTS

In order to analyze the performance of our system, we acquired datasets on fields of 3 different crop types in Eschikon (Switzerland) - soybean, sugarbeet, and winter wheat. For each crop species we collected: (i) one sequence of GPS-IMU tagged images over the entire field from a UAV flying at 10 meters altitude; (ii) 4-6 sequences of GPS/IMU-tagged images of small portions of the field from a UGV point-of-view. Additionally, for the sugarbeet field, we acquired an additional aerial sequence of images from 20 meters altitude. More comprehensive details regarding the acquired datasets are reported in Table II.

The UAV datasets were acquired using a DJI Mavic Pro UAV equipped with a 12 MP color camera, while the UGV datasets were acquired moving the same camera by hand with a forward-looking point-of-view, simulating data acquisition by a ground robot. The collected images are first converted into 3D colored point clouds using Pix4D Mapper [34], a professional photogrammetry software suite, which are then aligned using the proposed registration approach. To analyze the performance of the proposed approach, we make use of the following error metrics:

$$\delta t = t - \tilde{t} \quad \delta r = \operatorname{Trace}(R^T \cdot \tilde{R}) \quad \delta s = s \oslash \tilde{s} \quad (5)$$

$$e_t = \|\delta t\|_2 \quad e_R = \operatorname{acos}((\delta r - 1)/2) \quad e_s = \|\delta s\|_2 \quad (6)$$

where  $\oslash$  stands for the element-wise division operator and  $(e_t, e_r, e_s)$  are, respectively, the translational, the rotational, and the scale error metrics. We report the AgriColMap related parameters we used in all the experiments in Tab. I.

<sup>4</sup>It is noteworthy to highlight that the FPFH, being a local descriptor, does not embed global displacements along the axes.




Fig. 4: Average success registration rate curves by varying the initial guess and the initial scale error: (i) from left to right, the initial scale error is incrementally increased: 0%, 10%, 20%, 30%; (ii) in each plot within the upper row, the initial heading error  $\delta\psi$  is kept fixed, while the initial translational misalignment  $\delta t$  is incrementally increased until 5 meters. We assume 5 meters to be a reasonable upper bound for the initial GPS translational error. (iii) in the lower row figures  $\delta\psi$  is incrementally increased, while the initial translational misalignment  $\delta t$  is kept constant. It is important to point out that the successful registration rate of the Go-ICP [32] method is only reported for the cases without an initial scale error since this approach only deals with rigid transformations. For AgriColMap, we report the different results obtained in each dataset (sb: Soybean, sg10: Sugabeet 10m, sg20: Sugabeet 20m, ww: Winter Wheat).

TABLE I: Parameter set

Parameter	$\alpha$	$\beta$	$s$	$\sigma_{avg}$	$t_f$
Value	1	.5	0.02 m	0.04 cm	1

TABLE II: Overview of the Datasets: the global scale error is, in general, bigger in the UGV datasets since the camera is carried by hand, and therefore some GPS satellite signals might be not received.

Crop Type	Name	# Images	Crop Size (avg.)	Global Scale Error	Recording Height (approx.)
Soybean	<i>sUGV A</i>	16	6 cm	4%	1 m
	<i>sUGV B</i>	19	6 cm	6%	1 m
	<i>sUGV C</i>	22	6 cm	7%	1 m
	<i>sUAV</i>	89	6 cm	3%	10 m
Sugar Beet	<i>sbUGV A</i>	25	5 cm	6%	1 m
	<i>sbUGV B</i>	26	5 cm	7%	1 m
	<i>sbUGV C</i>	27	5 cm	5%	1 m
	<i>sbUAV A</i>	213	5 cm	3%	10 m
	<i>sbUAV B</i>	96	5 cm	2%	20 m
Winter Wheat	<i>wwUGV A</i>	59	25 cm	9%	1 m
	<i>wwUGV B</i>	61	25 cm	9%	1 m
	<i>wwUAV</i>	108	25 cm	5%	10 m

### A. Performance Under Noisy Initial Guess

This experiment is designed to show the robustness of the proposed approach under different noise conditions affecting the initial guess, and different directional scale discrepancies. For each UGV point cloud, we estimate an accurate ground truth non-rigid transform by manually selecting the correct point-to-point correspondences with the related UAV cloud. We generate random initial alignments between maps by manually adding noise, with different orders of magnitude, to the ground truth alignment heading, translation, and scale. Then, we align the clouds with the sampled initial alignments by using (i) the proposed approach; (ii) a modified version of the proposed approach by moving from the ExG + DSM environment representation to an RGB one (iii) a non-rigid standard ICP, (iv) the coherent point drift (CPD) method [33], (v) a state-of-the art Globally Optimal 3D ICP (Go-ICP) [32], and with standard sparse visual feature matching approaches [35], [36], [37], applied as a data association front-end to our method in place of the proposed LDOF based data association (Sec. III-B): in the last cases, we exploit only the ExG channel

of the grid maps (Sec. III-A). An alignment is considered valid if:  $e_t \leq 0.05$  m,  $e_r \leq 0.1$  rad, and  $e_s \leq 2.5\%$ .

The results are illustrated in Fig. 4. The proposed approach significantly outperforms the other approaches, ensuring an almost 100% success registration rate up to a scale error of 25%, and a high probability of succeeding even with a 30% scale error. The ICP-based registration methods [33], [32], due to the absence of structural 3D features on the fields, fall into local minima with high probability. The closest methods, in terms of robustness, are based on local feature matching [35], [36], [37], succeeding in the registration procedure up to a scale error magnitude of 10%. While analyzing the results, however, we verified that, unlike our method, these methods provide a larger number of wrong, incoherent point associations, and such a problem is clearly highlighted for increasing scale deformations above 20% and rotations above 0.1 radians. The superior robustness is also confirmed for noisy initial guesses: unlike the other methods, our approach guarantees a high successful registration rate for a translational error up to 5 meters, and an initial heading error up to 11.5 degrees, enabling it to deal with most errors coming from a GPS or AHRS sensor. Our method generalizes well over the different datasets, showing the capability to deal with different crop species, crop growth stages (i.e., the winter wheat crop is in an advanced growth stage compared to the soybean and sugarbeet), soil conditions, and point cloud resolution (from different UAV altitudes). An additional important outcome is the higher alignment probability obtained with the ExG/DSM environment model than the RGB one.

In Table IV, we report a comparison between the inliers percentages when using the visual (i.e., the ExG or the RGB) and the geometric terms in the cost function of Eq. (3). Most of the information is carried by the visual term, especially by the ExG, while the sole geometric term is not able to provide valid results. Nevertheless, when combined, the latter acts as a strong outlier rejection term, improving the robustness properties of the registration procedure. This is true especially for the sugarbeet dataset, where the inliers percentage increases

TABLE III: Registration accuracy comparison among the proposed approach, the non-rigid ICP, the CPD [33], and the Go-ICP [32] systems. The table reports, for each cell, the average accuracy among all the successful registrations with a specific initial anisotropic scaling error.

crop type	approach	registration err.	registration err.	registration err.	registration err.	registration err.	registration err.	registration err.
		(trans/ros/scale) scale error 0%	(trans/ros/scale) scale error 5%	(trans/ros/scale) scale error 10%	(trans/ros/scale) scale error 15%	(trans/ros/scale) scale error 20%	(trans/ros/scale) scale error 25%	(trans/ros/scale) scale error 30%
Soybean	AgriColMap	0.03m/0.03°/-	0.03m/0.04°/1.9%	0.04m/0.05°/2.0%	0.04m/0.04°/2.1%	0.03m/0.04°/2.2%	0.05m/0.04°/2.2%	0.05m/0.05°/2.3%
	ICP	0.03m/0.07°/-	0.05m/0.08°/2.4%	0.04m/0.09°/2.4%	fail	fail	fail	fail
	CPD [33]	0.02m/0.03°/-	0.04m/0.07°/2.1%	0.03m/0.08°/2.3%	0.03m/0.08°/2.4%	fail	fail	fail
	Go-ICP [32]	0.03m/0.06°/-	-	-	-	-	-	-
	SURF [35]	0.02m/0.04°/-	0.03m/0.04°/2.2%	0.05m/0.06°/2.4%	fail	fail	fail	fail
	[36]	0.02m/0.04°/-	0.04m/0.04°/2.2%	0.05m/0.05°/2.3%	fail	fail	fail	fail
FAST+BRIEF [37]	0.04m/0.05°/-	0.05m/0.06°/2.3%	0.05m/0.07°/2.4%	fail	fail	fail	fail	
Sugabeet 10m	AgriColMap	0.03m/0.04°/-	0.03m/0.04°/2.1%	0.04m/0.04°/2.0%	0.05m/0.06°/2.0%	0.05m/0.07°/2.3%	0.05m/0.1°/2.3%	0.05m/0.1°/2.4%
	ICP	0.04m/0.05°/-	0.05m/0.07°/2.1%	0.05m/0.09°/2.4%	fail	fail	fail	fail
	CPD [33]	0.03m/0.04°/-	0.04m/0.05°/2.1%	0.04m/0.06°/2.2%	0.05m/0.09°/2.4%	fail	fail	fail
	Go-ICP [32]	0.02m/0.05°/-	-	-	-	-	-	-
	SURF [35]	0.03m/0.04°/-	0.03m/0.04°/2.1%	0.04m/0.07°/2.3%	fail	fail	fail	fail
	ORB [36]	0.02m/0.03°/-	0.03m/0.03°/2.2%	0.05m/0.06°/2.4%	fail	fail	fail	fail
FAST+BRIEF [37]	0.02m/0.04°/-	0.02m/0.03°/2.1%	0.05m/0.06°/2.3%	fail	fail	fail	fail	
Sugabeet 20m	AgriColMap	0.03m/0.03°/-	0.04m/0.03°/2.0%	0.04m/0.04°/2.2%	0.05m/0.05°/2.1%	0.05m/0.08°/2.2%	0.05m/0.09°/2.4%	0.05m/0.1°/2.4%
	ICP	0.05m/0.06°/-	0.05m/0.09°/2.3%	fail	fail	fail	fail	fail
	CPD [33]	0.04m/0.05°/-	0.05m/0.07°/2.3%	0.05m/0.08°/2.4%	0.05m/0.1°/2.5%	fail	fail	fail
	Go-ICP [32]	0.04m/0.05°/-	-	-	-	-	-	-
	SURF [35]	0.03m/0.04°/-	0.04m/0.05°/2.1%	0.04m/0.06°/2.4%	fail	fail	fail	fail
	ORB [36]	0.04m/0.05°/-	0.04m/0.05°/2.2%	0.04m/0.05°/2.4%	fail	fail	fail	fail
FAST+BRIEF [37]	0.03m/0.04°/-	0.04m/0.05°/2.1%	0.05m/0.07°/2.4%	fail	fail	fail	fail	
Winter Wheat	AgriColMap	0.04m/0.02°/-	0.04m/0.03°/2.0%	0.04m/0.05°/2.1%	0.04m/0.04°/2.2%	0.05m/0.08°/2.3%	0.05m/0.09°/2.4%	0.05m/0.1°/2.4%
	ICP	0.04m/0.07°/-	0.04m/0.08°/2.2%	0.05m/0.10°/2.5%	fail	fail	fail	fail
	CPD [33]	0.04m/0.05°/-	0.04m/0.05°/1.9%	0.04m/0.05°/2.1%	0.05m/0.09°/2.3%	fail	fail	fail
	Go-ICP [32]	0.03m/0.07°/-	-	-	-	-	-	-
	SURF [35]	0.03m/0.06°/-	0.03m/0.05°/2.2%	0.04m/0.06°/2.4%	fail	fail	fail	fail
	ORB [36]	0.04m/0.05°/-	0.04m/0.04°/2.1%	0.04m/0.06°/2.3%	fail	fail	fail	fail
FAST+BRIEF [37]	0.03m/0.07°/-	0.04m/0.06°/2.3%	0.05m/0.05°/2.4%	fail	fail	fail	fail	

quite significantly.

TABLE IV: Inliers percentage comparison when changing data terms in the LDOF cost function.

Crop Type	Descriptor Type (% inliers)			
	RGB	ExG	Depth	ExG + Depth
Soybean	11.7% ± 4.3%	53.2% ± 14.9%	0.2% ± 0.1%	<b>54.5% ± 13.2%</b>
Sugarbeet	49.2% ± 11.9%	64.1% ± 12.8%	0.4% ± 0.2%	<b>68.1% ± 13.6%</b>
Winter Wheat	22.9% ± 9.7%	51.8% ± 17.4%	0.1% ± 0.1%	<b>52.4% ± 16.7%</b>

### B. Accuracy Evaluation

To evaluate the accuracy of the proposed registration approach, we compare our results with the ground truth parameters and, by using all the successful registrations, we compute the average accuracy for each crop type and approach. The results are summarized in Tab. III, and are sorted in increasing order of scale error.

On average, our method results in a lower registration error as compared to all the other evaluated methods for the same scale error. The difference in the registration error is even more pronounced when comparing the *Sugabeet 10m* against *Sugabeet 20m* datasets. Indeed, due to the higher sparseness of the points in the latter, all the other methods tend to perform slightly worse than they perform with the *Sugabeet 10m*. Conversely, our method results in almost the same registration error magnitudes, showing that it correctly deals with different density of the initial colored point clouds. We also report some qualitative results in Fig. 5

### C. Runtime Evaluation

We recorded the average, maximum, and minimum computational time for all tested methods over 100 successful

registrations, reporting these values in Tab. V. The method requiring the biggest computational effort is Go-ICP. The proposed approach requires half the computational time as compared to Go-ICP, but turns out to be quite slow compared to the custom-built ICP, and, in general, to all the other matching approaches. Fig. 6 shows the runtime percentages for the proposed approach. The biggest component of the computational effort is required to extract the geometric features (i.e., the FPFH features), meaning that the total computational time might be reduced by switching to a less time consuming 3D feature or by solely using the visual term.




Fig. 5: Qualitative registration results seen from aerial (left) and ground point-of-views. In the former, the UGV clouds are indistinguishable from the UAV, proving the correctness of the registration. Conversely, in the latter, the UGV clouds are clearly visible due to their higher points density.

## VI. CONCLUSIONS

We addressed the cooperative UAV-UGV environment reconstruction problem in agricultural scenarios by proposing an effective way to align heterogeneous 3D maps. Our approach




Fig. 6: Average percentage of total runtime for different parts of the AgriColMap pipeline.

is built upon a multimodal environment representation that utilises the semantics and the geometry of the target field, and a data association strategy solved as a LDOF problem. We reported a comprehensive set of experiments, proving the superior robustness of our approach against standard methods. An open-source implementation of our system and the acquired datasets are made publicly available with this paper.

## VII. ACKNOWLEDGEMENTS

The authors would like to thank Hansueli Zellweger from the ETH Plant Research Station in Eschikon, Switzerland for preparing the fields, managing the plant life-cycle and treatments during the entire growing season. The authors would also like to thank Dr. Frank Liebisch from the Crop Science Group at ETH Zürich for helpful discussions.

## REFERENCES

- [1] R. Käslin, P. Fankhauser, E. Stumm, Z. Taylor, E. Mueggler, J. Delmerico, D. Scaramuzza, R. Siegwart, and M. Hutter, "Collaborative localization of aerial and ground robots through elevation maps," in *Proc. of the IEEE SSRR*, 2016.
- [2] R. Khanna, M. Möller, J. Pfeifer, F. Liebisch, A. Walter, and R. Siegwart, "Beyond point clouds-3d mapping and field parameter measurements using uavs," in *IEEE ETFA*, 2015.
- [3] A. Howard, "Multi-robot simultaneous localization and mapping using particle filters," *Intl. Journal of Robotics Research (IJRR)*, vol. 25, no. 12, pp. 1243–1256, 2006.
- [4] A. Gil, Ó. Reinoso, M. Ballesta, and M. Juliá, "Multi-robot visual SLAM using a rao-blackwellized particle filter," *Robotics and Autonomous Systems*, vol. 58, no. 1, pp. 68 – 80, 2010.
- [5] A. Birk and S. Carpin, "Merging occupancy grid maps from multiple robots," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1384–1397, 2006.
- [6] T. M. Bonanni, B. Della Corte, and G. Grisetti, "3-D map merging on pose graphs," *IEEE Robotics and Automation Letters (RA-L)*, vol. 2, no. 2, pp. 1031–1038, 2017.
- [7] A. Gawel, R. Dubé, A. Surmann, J. Nieto, R. Siegwart, and C. Cadena, "3d registration of aerial and ground robots for disaster response: An evaluation of features, descriptors, and transformation estimation," in *Proc. of the IEEE SSRR*, 2017.
- [8] M. Imperoli, C. Potena, D. Nardi, G. Grisetti, and A. Pretto, "An effective multi-cue positioning system for agricultural robotics," *IEEE Robotics and Automation Letters (RA-L)*, 2018.
- [9] Y. Hu, R. Song, and Y. Li, "Efficient coarse-to-fine patch match for large displacement optical flow," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [10] J. L. Blanco, J. González, and J.-A. Fernández-Madrugal, "A robust, multi-hypothesis approach to matching occupancy grid maps," *Robotica*, vol. 31, pp. 687–701, 2013.
- [11] S. Saeedi, L. Paull, M. Trentini, and H. Li, "Multiple robot simultaneous localization and mapping," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2011.

	Runtime [sec]		
	Min	Max	Avg
AgriColMap	63.7	118.6	79.8
ICP	2.1	10.6	4.5
CPD	4.9	23.2	8.2
Go-ICP	5.3	689.2	193.1
SURF [35]	4.6	7.2	5.3
ORB [36]	3.9	6.7	4.8
FAST+BRIEF [37]	3.7	6.4	4.5

TABLE V: Runtime comparison.

- [12] C. Frueh and A. Zakhor, "Constructing 3D city models by merging ground-based and airborne views," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2003.
- [13] J. Jessup, S. N. Givigi, and A. Beaulieu, "Robust and efficient multi-robot 3D mapping with octree based occupancy grids," in *Proc. of the IEEE Intl. Conf. on Systems, Man, and Cybernetics (SMC)*, 2014.
- [14] H. Chui and A. Rangarajan, "A new point matching algorithm for non-rigid registration," *Comput. Vis. Image Underst.*, vol. 89, no. 2-3, 2003.
- [15] A. W. Fitzgibbon, "Robust registration of 2D and 3D point sets," in *British Machine Vision Conference*, 2001, pp. 662–670.
- [16] S. Saeedi, M. Trentini, M. Seto, and H. Li, "Multiple-robot simultaneous localization and mapping: A review," *Journal of Field Robotics (JFR)*, vol. 33, no. 1, pp. 3–46.
- [17] Q. Shan, C. Wu, B. Curless, Y. Furukawa, C. Hernandez, and S. M. Seitz, "Accurate geo-registration by ground-to-aerial image matching," in *2nd Intl. Conf. on 3D Vision*, 2014.
- [18] C. Wang, K. Wilson, and N. Snavely, "Accurate georegistration of point clouds using geographic data," in *2013 International Conference on 3D Vision - 3DV 2013*, 2013, pp. 33–40.
- [19] A. Bódis-Szomorú, H. Riemenschneider, and L. V. Gool, "Efficient volumetric fusion of airborne and street-side data for urban reconstruction," in *Proc. of the International Conference on Pattern Recognition (ICPR)*, 2016, pp. 3204–3209.
- [20] C. Fruh and A. Zakhor, "Constructing 3D city models by merging aerial and ground views," *IEEE Computer Graphics and Applications*, vol. 23, no. 6, pp. 52–61, 2003.
- [21] N. M. et al., "Collaborative mapping of an earthquake-damaged building via ground and aerial robots," *Journal of Field Robotics (JFR)*, vol. 29, no. 5, pp. 832–841, Sept 2012.
- [22] C. Forster, M. Pizzoli, and D. Scaramuzza, "Air-ground localization and map augmentation using monocular dense reconstruction," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2013.
- [23] T. Hinzmann, T. Stastny, G. Conte, P. Doherty, P. Rudol, M. Wzorek, E. Galceran, R. Siegwart, and I. Gilitschenski, "Collaborative 3D reconstruction using heterogeneous UAVs: System and experiments," in *Proc. of the Intl. Sym. on Experimental Robotics (ISER)*, pp. 43–56.
- [24] A. Gawel, T. Cieslewski, R. Dubáľ, M. Bosse, R. Siegwart, and J. Nieto, "Structure-based vision-laser matching," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2016.
- [25] U. Weiss and P. Biber, "Plant detection and mapping for agricultural robots using a 3D lidar sensor," *Robotics and autonomous systems*, vol. 59, no. 5, pp. 265–273, 2011.
- [26] A. English, P. Ross, D. Ball, and P. Corke, "Vision based guidance for robot navigation in agriculture," in *IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2014.
- [27] J. Dong, J. G. Burnham, B. Boots, G. Rains, and F. Dellaert, "4D crop monitoring: Spatio-temporal reconstruction for agriculture," in *IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2017.
- [28] N. Chebrolu, T. Läbe, and C. Stachniss, "Robust long-term registration of UAV images of crop fields for precision agriculture," *IEEE Robotics and Automation Letters (RA-L)*, vol. 3, no. 4, pp. 3097–3104, 2018.
- [29] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, no. 1-3, pp. 185–203, 1981.
- [30] E. Tola, V. Lepetit, and P. Fua, "Daisy: An efficient dense descriptor applied to wide-baseline stereo," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 5, pp. 815–830, 2010.
- [31] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (FPFH) for 3D registration," in *IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2009.
- [32] J. Yang, H. Li, D. Campbell, and Y. Jia, "Go-ICP: A globally optimal solution to 3D ICP point-set registration," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 11, pp. 2241–2254, 2016.
- [33] A. Myronenko and X. Song, "Point set registration: Coherent point drift," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 12, pp. 2262–2275, 2010.
- [34] Point clouds generated using Pix4Dmapper by Pix4D. [Online]. Available: <http://www.pix4d.com/>
- [35] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-up robust features (SURF)," *Comput. Vis. Image Underst.*, vol. 110, no. 3, pp. 346–359, 2008.
- [36] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *IEEE Intl. Conf. on Computer Vision (ICCV)*, 2011, pp. 2564–2571.
- [37] M. Calonder, V. Lepetit, C. Strecha, and P. Fua, "Brief: Binary robust independent elementary features," in *Europ. Conf. on Computer Vision (ECCV)*, 2010, pp. 778–792.