Coupling Markov Random Fields and Automated Planning for Online Decision-Making to Map Spatial Phenomena

Alexandre Albore[†‡] and **Nathalie Peyrard**[†] and **Régis Sabbadin**[†] and **Florent Teichteil-Königsbuch**[‡]
† INRA UR 875 (MIAT), chemin de Borderouge, 31326 Castanet-Tolosan, France - *name.surname*@toulouse.inra.fr
‡ ONERA (DCSD) - 2 avenue É. Belin, 31000 Toulouse, France - *name.surname*@onera.fr

Mapping spatial phenomena (e.g. crop growth, weeds, etc.) has become a central task in precision agronomy for managing production at landscape or field scale. Maps are costly to obtain since they require intensive surveys in the field, most of the time performed by human annotators or with human-controlled Unmanned Aerial Vehicles (UAVs). Data acquisition can also be performed by deploying sensing devices, but given the generally limited available amount of such sensors, and their fixed sensing radius, to decide where to place the sensors constitutes a challenge for mapping spatial phenomena. Many techniques try to optimise their deployment, yet in certain cases it is infeasible in practice, or too costly, to install fixed sensors (e.g. large exploitations).

New technologies of sensors embarked on mobile devices, like autonomous UAVs or wheeled robots, face a similar optimization issue, but can allow to use a single device to cover the whole field. Such a device should be deployed on-demand, reducing as much as possible off-line calculations, and should be usable for a vast array of mapping situations, given the dynamic environment we aim at monitoring in precision agriculture. An adaptive sampling approach for mapping in crop fields, focusing on which plots are more "interesting" to sample, would take advantage of on-line sequential decision-making capabilities of mobile devices.

Gaussian processes and Markov Random Fields (MRF) have been successfully used to model the spatial phenomena in the optimization problem, respectively to select where to place sensor devices on an area, and to perform adaptive sampling in crop fields (Singh et al. 2007; Peyrard et al. 2013). However, past approaches do not take into account the limited battery life that affect mobile devices, which can drastically reduce the amount of gathered data.

Approaches that allow to optimise data collection under limited sampling budget, e.g. (Bonneau et al. 2014), require very long off-line computation time to preliminarily compute the parameters of the sampling policy. This does not match our need for devices to be deployed on-demand, efficiently performing sampling and mapping with light logistics and low-budget computations.

We propose an alternative way to solve these issues, where we couple MRFs and Automated Planning, in order to optimise the sequence of sampling by balancing map quality and cost objectives.

Automated Planning is the branch of Artificial Intelligence that is concerned with selecting the next-to-apply action in order to lead an agent toward a desired goal. Coupling MRFs with Automated Planning techniques offers a methodology to solve on-line tasks for sampling and sensing: on one side, MRFs are used to represent and update probabilistic knowledge about the measured phenomena –that remains uncertain to the observer–, and on the other side, a planner (the solver) can automatically produce decision rules (a controller) from a description of the task, in order to drive an autonomous device toward sampling plots in the field, while minimising the cost of the mission and maximising the quality of the map reconstruction.

In the following, we detail more formally our approach coupling MRFs and planners, and we illustrate it on a UAV application to adaptively sample in a crop field for weeds species mapping (Albore et al. 2015).

Automated Planning

A key element for fully autonomous agents is the ability to select and organize actions over time in order to fulfil some objective (e.g. maximising reward functions, reaching given symbolic system states). This implies deciding on the course of potentially uncertain events and situations that may occur during the execution of the task, as several sources of uncertainty can impact the mission (Ingrand and Ghallab 2014).

Artificial Intelligence planning (Ghallab, Nau, and Traverso 2004) is a model-based and theorem-proving approach to this problem. Given a dynamical model of the available actions, which includes environmental conditions on their applicability and descriptions of their effects on the world, an AI planning solver generates a sequence of actions named *plan*, that can achieve a given objective, in the case such a solution plan exists.

Many variants of AI planning have been studied and applied, ranging from "classical" planning, where effects of actions and environmental observations of a single agent are deterministic, to probabilistic planning where it is asked to the solver to provide a *policy* that associates a state with an action, instead of a plan. The simplest model of planning is where actions are assumed to have deterministic effects and the agent's position is always perfectly known; this classical planning framework has been used to produce ordered sequences of plots to sample in a mapping task with an UAV, while minimizing the travelled distance. The formal model underlying the planning problem can be described by the tuple $\mathcal{P} = \langle S, s_0, S_G, O, \varphi \rangle$ where S is a finite set of states, $s_0 \in S$ is the initial state, and $S_G \subseteq S$ is the set of goal states. Transitions between states are given by O, the set of operators associating to a state s its successor state s' using a state *transition* function $\varphi : S \times O \to S$. A fixed action a is *applicable* in a state s when there exists at least one target state s' such that $\varphi(s, a) = s'$.

AI planning does not provide any means to actually generate the high-level objectives that we call *goals*, but it allows the planning agent to autonomously select relevant actions and organize them over time in order to achieve these objectives. We aim at using graphical models, like Markov Random Fields, to effectively select such goals in the context of mapping crop fields: from the probabilistic representation of the spatial phenomena observed, given by the MRF, we select the list of plots that will provide relevant information to build the map of the field.

MRF modelling of abundance maps

Common practice in field sampling is to divide the crop field in a regular grid of N plots of small area. An observation in a plot is the abundance class of the weed species (discretised in K classes), and we assume no measurement error.

To each plot $i \in V = \{1, \dots, N\}$ of the MRF model, is attached a discrete random variable X_i with domain $D = \{1 \dots K\}$, where K is the number of abundance classes. The joint distribution of the whole map $X = (X_1, \dots, X_N)$ is assumed to be expressed as a pairwise MRF: $\forall x \in D^N$,

$$\mathbb{P}(X = x) = \frac{1}{Z} \prod_{i=1}^{N} f_i(x_i) \prod_{(i,j) \in E} f_{i,j}(x_i, x_j)$$

The set E is the set of all pairs of neighbours in the grid of plots and Z is a normalising constant. The f_i and $f_{i,j}$ are non negative functions called respectively order-1 and order-2 potential functions. Roughly speaking, the order-1 potential functions weight the relative proportions of the K abundance classes while the order-2 potential functions encode spatial correlation between abundance values at different plots. The choice of an appropriate MRF model amounts to the choice of these potential functions.

Sampling in MRF with an objective of map reconstruction is modelled as the problem of finding samples that optimise the Maximum Posterior Marginals (MPM) criterion, classically used in image analysis; we use it to derive an estimator x^* of the hidden map x given the history of past observations $x_A = (x_{A_0}, x_{A_1}, \cdots, x_{A_t})$:

$$x^* = \left\{ x_i^* \;\middle|\; i \in V, \quad x_i^* = \underset{x_i \in D}{\operatorname{argmax}} \, \mathbb{P}(x_i \mid x_A) \right\}.$$

Note that the trajectory from which x_A is obtained can be determined once and for all beforehand, or adaptively chosen on-line assuming that the spatial data don't change during the task; we privilege the latter approach, as more adapted to a light logistics approach to map reconstruction.

The quality of a trajectory is measured as the expected quality of the estimator x^*. In practice, we first define the quality of a trajectory $((A_h, x_{A_h}))_{h=1..H}$ as a function of (A, x_A), where $A = \{A_h\}_1^H \subseteq V$ is the set of the explored sites at steps h and x_{A^h} is the sample output at step h:

$$U\big(A, x_A\big) = \sum_{i \in V} \Big[\mathbb{P}(x_i^* \mid x_A) \Big].$$

$U\big(A, x_A\big)$ can be interpreted as the expectation of the number of well-classified plots, when allocating their values to the modes of the marginals $\mathbb{P}(x_i \mid x_A)$, calculated by Loopy Belief Propagation (Murphy, Weiss, and Jordan 1999).

These values will be used to select a subset of the plots in the field, based on the potential quality gain that an observation can provide to the final map. The planner will then synthesise a plan, i.e. an ordering of the plots that minimises the cost of navigating between them and sensing, which is *de facto* a trajectory.

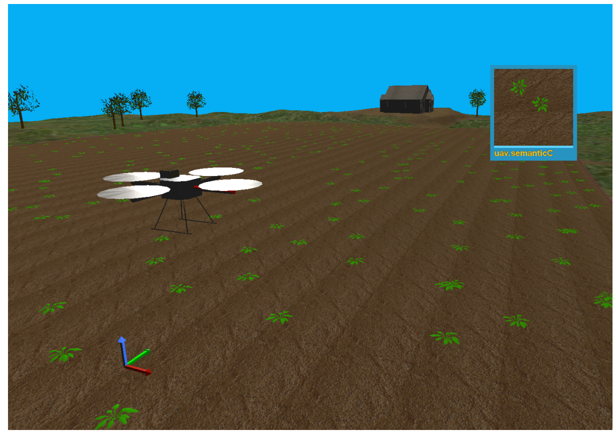

Figure 1: Platform simulation on MORSE. At the upper right corner, the UAV's semantic camera framing weeds.

A replanning-based approach

The planning problem of reconstructing a weed abundance map with an UAV is built and executed in a closed-loop fashion. First, we generate a set of n plots to sample that maximise the expected quality gain derived from the MRF model, given a set of observations (A, x_A). For each variable X_i of the MRF, the *expected quality gain* is an optimistic approximation of the increase of the updated utility $U(A \cup \{i\}, x_A, x_i) - U(A, x_A)$. Given a set of n plots, a planner elaborates a trajectory that connects them while minimising the travelled distance, and considering eventual side effects of observing, such as the uselessness of sampling plots close to the recently visited ones, while maximising the map quality. Since this planning problem is too hard to be exactly solved either off-line or on-line, we propose an approach interleaving planning and execution: from past observations at a given time step, we compute a full plan to find a trajectory that minimises the navigation cost while visiting all the n sampled plots. This plan consists in a sequence of locations *and* expected observations. The plan is then executed by the UAV, and the observations collected. Those observations are then used to update the expected quality gains values from the MRF. We monitor the execution so to stop it and recompute a new trajectory whenever the accumulated differences between the actual observations and the expected ones exceed a given threshold, meaning that the observations gathered so far do not bring enough accretion on the overall reconstructed map quality. In fact, the set of plots in the trajectory is derived from a snapshot of the probability distribution before the execution of the plan; the observations obtained during the navigation update the information about the potential quality gains, and thus a better selection of the plots to visit is always possible.

This interleaving planning and execution schema is widely used in those frameworks where planning is applied in highly dynamic domains (e.g. robotics), as it allows to reason on a simplified version of the model, while being able to recover from the execution whenever the initial assumptions do not hold any more.

This planning approach to the problem of weeds mapping in crop field compares favourably to a greedy approach, the only other approach that can be used on-line. The latter selects the next plot to be visited by the UAV as the one carrying the biggest amount of uncertainty, without accounting for the future flight duration (Peyrard et al. 2013).

We implemented the previously described replanning algorithm applied to the weeds mapping problem within the Robot Operating System (ROS) framework (Quigley et al. 2009), a robotic meta-operating system. The evaluation of marginals in the MRF and the planner are integrated on the same platform, taking advantage of our implementation of the LBP algorithm, and the (re)planning loop uses a general purpose planner with the Serialized Iterated Width algorithm (Lipovetzky and Geffner 2012), that we adapted as a ROS independent planning package. Tests have been run in the MORSE simulator developed in academic robotics (Echeverria et al. 2011), which enables to perform software architecture-in-the-loop (SAIL) realistic simulations, i.e. to test the exact same functional architecture as the one that will be implemented on-board the real UAV, but replacing the physical sensors and actuators by simulated data (Cf. Fig. 1). In this environment, the planning approach leads to results of similar quality but at a much less cost (measured as the distance covered during the flight) (Albore et al. 2015); this means that if the same distance is allocated to the two approaches, the planner will sample more plots, and therefore provide better quality estimated maps.

Conclusions

Automated Planning, and in particular planning under uncertainty, while efficient in finding policies to navigate toward a desired goal even in domains with incomplete information or exogenous events, does not update the probabilities or the rewards from acting models. Graphical models have the advantage to easily represent knowledge about processes correlated in space or time, and to quickly update information from real-world observations. Even if optimal solutions are difficult to obtain, automated planning can use suboptimal information to elaborate an effective plan to reach the desired goals while minimising resources consumption. Such a coupling has been proven to be well adapted in reactive platforms to adaptive sampling, while interleaving planning, execution, and information update.

Further possible extensions go in the same applicative direction. Adopting hierarchical MRFs would allow to consider in the planning loop the cost of observing with the UAV at different heights, reasoning on the contribution of multi-resolution images. We expect to adopt a trade-off between low resolution images, but that cover a larger area of the field, and high resolution ones, more informative but more costly in terms of distance covered.

As future work, we also consider to extend the use of graphical models to other forms of planning under uncertainty (MDPs, POMDPs, contingent and conformant planning), as they are closer to real-world applications than classical planning, and they explicitly embed in their models uncertainty about the application domain, which is not always easy to evaluate and update without a dedicated mathematical framework.

References

Albore, A.; Peyrard, N.; Sabbadin, R.; and Teichteil-Königsbuch, F. 2015. An online replanning approach for crop fields mapping with autonomous UAVs. In *Proc. of Int. Conf. on Automated Planning and Scheduling (ICAPS-15)*.

Bonneau, M.; Gaba, S.; Peyrard, N.; and Sabbadin, R. 2014. Reinforcement learning-based design of sampling policies under cost constraints in markov random fields: Application to weed map reconstruction. *Computational Statistics and Data Analysis* 72:30–44.

Echeverria, G.; Lassabe, N.; Degroote, A.; and Lemaignan, S. 2011. Modular openrobots simulation engine: Morse. In *Proceedings of the IEEE ICRA*.

Ghallab, M.; Nau, D. S.; and Traverso, P. 2004. *Automated planning - theory and practice*. Elsevier.

Ingrand, F., and Ghallab, M. 2014. Robotics and artificial intelligence: A perspective on deliberation functions. *AI Communications* 27(1):63–80.

Lipovetzky, N., and Geffner, H. 2012. Width and serialization of classical planning problems. In *Proc. European Conference on Artificial Intelligence (ECAI-12)*, 540–545.

Murphy, K. P.; Weiss, Y.; and Jordan, M. I. 1999. Loopy belief propagation for approximate inference: An empirical study. In *Proc. of the conference on Uncertainty in Artificial Intelligence (UAI)*, 467–475. Morgan Kaufmann Publishers.

Peyrard, N.; Sabbadin, R.; Spring, D.; Brook, B.; and Mac Nally, R. 2013. Model-based adaptive spatial sampling for occurrence map construction. *Statistics and Computing* 23(1):29–42.

Quigley, M.; Conley, K.; Gerkey, B.; Faust, J.; Foote, T.; Leibs, J.; Wheeler, R.; and Ng, A. Y. 2009. ROS: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3.

Singh, A.; Krause, A.; Guestrin, C.; Kaiser, W. J.; and Batalin, M. A. 2007. Efficient planning of informative paths for multiple robots. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, 2204–2211.