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

Alexandre Albore[1,2]
joint work with N. Peyrard[2], R. Sabbadin[2], F. Teichteil-Königsbuch[1]

[1]DCSD - Onera

[2]UR 875 Mia - INRA

Paris, 28/09/2015

Requirements for UAVs crop fields mapping

An important tool in **precision agriculture** to manage the production in crop fields is a map of pest abundance spatial distribution.

Objective: an AI alternative to the **expensive** human annotators

- **Remote sensing tools**
- **Autonomous**
- **On-demand usage**

Requirements for UAVs crop fields mapping

An important tool in **precision agriculture** to manage the production in crop fields is a map of pest abundance spatial distribution.

Objective: an AI alternative to the **expensive** human annotators

- **Remote sensing tools**: use of UAVs (operate below cloud cover, deployed quickly and repeatedly).
- **Autonomous**: on-board computation habilities to exploit dynamic information and for in-flight optimisation of navigation.
- **On-demand usage**: can be deployed on demand, without heavy computation machinery.

Challenges of an AI sampling-based solution

It is impractical to do a complete and exhaustive mapping.

- **Sampling** plots, and then rebuild the whole map.
  - How to represent knowledge and spatial correlations?

Challenges of an AI sampling-based solution

It is impractical to do a complete and exhaustive mapping.

- **Sampling** plots, and then rebuild the whole map.
  - How to represent knowledge and spatial correlations?
- Selection of the best sites where to sample (**quality**).
- Observe those sites within **time and battery constraints.**

Challenges of an AI sampling-based solution

It is impractical to do a complete and exhaustive mapping.

- **Sampling** plots, and then rebuild the whole map.
  - How to represent knowledge and spatial correlations?
- Selection of the best sites where to sample (**quality**).
- Observe those sites within **time and battery constraints.**
  - How to choose an effective sampling strategy that addresses **both criteria**?

Challenges of an AI sampling-based solution

It is impractical to do a complete and exhaustive mapping.

- **Sampling** plots, and then rebuild the whole map.
  - How to represent knowledge and spatial correlations?
- Selection of the best sites where to sample (**quality**).
- Observe those sites within **time and battery constraints.**
  - How to choose an effective sampling strategy that addresses **both criteria**?

Outline for the rest of the talk

1. Challenges in crop fields mapping with autonomous UAVs
2. Coupling AI Planning and Markov Random Field modelling
3. Online update of the MRF
4. Empirical evaluation
5. Novel applications

Automated Planning: a model-based approach

- AI Planning is the **model-based approach** to decision making
- Produces a **general purpose solver**, given a description of the environment and its behaviour.
  - the solution is an action policy.
- Information about uncertainty **often lacks** in planning models.
- **Spatial phenomena** are well described by Graphical Models
  - Representation of uncertainty about the predictions and of the expected gain.

Automated Planning: a model-based approach

- AI Planning is the **model-based approach** to decision making
- Produces a **general purpose solver**, given a description of the environment and its behaviour.
  - the solution is an action policy.
- Information about uncertainty **often lacks** in planning models.
- **Spatial phenomena** are well described by Graphical Models
  - Representation of uncertainty about the predictions and of the expected gain.

Current approaches

In the context of this adaptive sampling technique for weeds mapping;

- (Peyrard et al. 2013) use a *greedy* approach, updating the more uncertain sites
  - rather fast,
  - **but** does not optimise resources such as remaining flying time.
- (Bonneau et al. 2014) derived a Reinforcement Learning approach
  - considers the full sampling horizon and sampling budget,
  - **but** very long off-line computation time.

Neither approach fits our requirements of **costs management** and **deployment on demand**.

Current approaches

In the context of this adaptive sampling technique for weeds mapping;

- (Peyrard et al. 2013) use a *greedy* approach, updating the more uncertain sites
  - rather fast,
  - **but** does not optimise resources such as remaining flying time.
- (Bonneau et al. 2014) derived a Reinforcement Learning approach
  - considers the full sampling horizon and sampling budget,
  - **but** very long off-line computation time.
- (Krause et al 2011) optimization of sensors placement for Hidden Markov Models
  - considers correlation between sensors locations
  - **but** expects a previous knowledge of available observations

Neither approach fits our requirements of **costs management** and **deployment on demand**.

Current approaches

In the context of this adaptive sampling technique for weeds mapping;

- (Peyrard et al. 2013) use a *greedy* approach, updating the more uncertain sites
  - rather fast,
  - **but** does not optimise resources such as remaining flying time.
- (Bonneau et al. 2014) derived a Reinforcement Learning approach
  - considers the full sampling horizon and sampling budget,
  - **but** very long off-line computation time.
- (Krause et al 2011) optimization of sensors placement for Hidden Markov Models
  - considers correlation between sensors locations
  - **but** expects a previous knowledge of available observations

Neither approach fits our requirements of **costs management** and **deployment on demand**.

An hybrid approach: coupling MRFs with Planning

- **Non-greedy** approach for sampling in a MRF.
    - A MRF is an **undirected graphical model** with a set of random variables having a Markov property.
    - For us it represents the **probability distribution** of the spatial phenomena.

An hybrid approach: coupling MRFs with Planning

- **Non-greedy** approach for sampling in a MRF.
- We dissociate the problem of **selecting the observation sites** from the one of **planning their visiting order**.

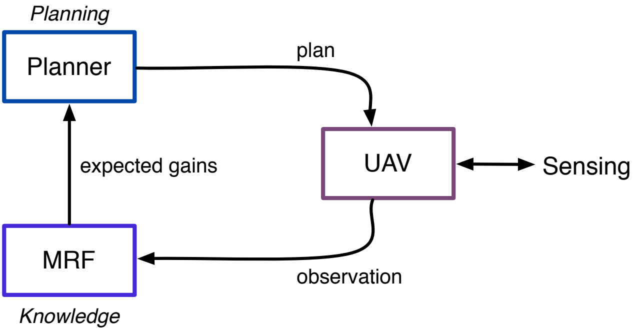

The MRF model for abundance map

The joint distribution of the whole map $X = (X_1, \ldots, X_N)$ is assumed to be expressed as a pairwise MRF:

$$\forall\, \mathbf{x} \in \text{classes}^N,$$

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

- The set E is the set of all pairs of neighbours in the grid,
- Z is a normalising constant,
- f_i and $f_{i,j}$ are non negative functions called order 1 and order 2 **potential functions.**

Choosing the set of sites to sample

Select n sites to sample that maximise the **expected quality gain**:

$$\bar{q}(X_i, x_A) = \max_k \Big(\sum_{i,j} \max_{x_j} \mathbb{P}(x_j | X_i = k, x_A) \Big) - \sum_{i,j} \max_{x_j} \mathbb{P}(x_j | x_A)$$

for a set of observations x_A, and a variable X_i.

Choosing the set of sites to sample

Select n sites to sample that maximise the **expected quality gain**:

$$\bar{q}(X_i, x_A) = \max_k \Big(\sum_{dist(i,j)\leq r} \max_{x_j} \mathbb{P}(x_j | X_i = k, x_A) \Big) - \sum_{dist(i,j)\leq r} \max_{x_j} \mathbb{P}(x_j | x_A)$$

for a set of observations x_A, and a variable X_i.

Our own **implementation** of Loopy Belief Propagation:

- Update only *sensitive perimeter* ($\sim$ 1%)
- Calculated online, depends on the correlation matrix.

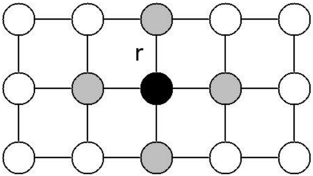

A replanning-based approach

- **Classical planner** to generate a plan, i.e. a sequence of **locations and expected observations.**
- Plan guarantees that the **physical constraints** are not violated.cf. (Ivankovic et al. 2014)
- The plan is executed by the UAV until the number of actual observations that differ from expected ones exceeds a threshold.
- Then, a **replanning** episode is triggered.
- A **partial MDP policy** based on multiple plans.

Empirical evaluation: the sampling path [ICAPS15]

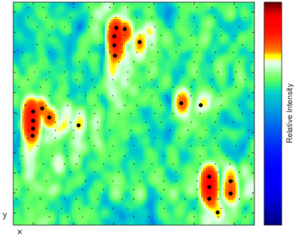

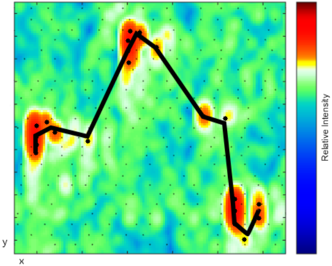

planning-approach

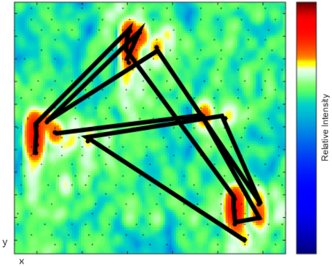

greedy approach

Empirical evaluation in MORSE simulation [ICAPS15]

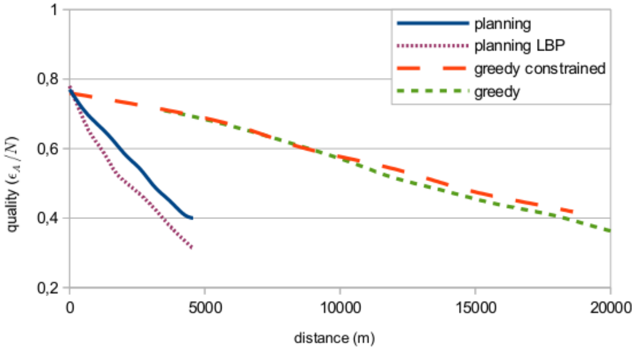

Quality vs. distance for the isotropic model. 4500m distance limit.

Empirical evaluation in MORSE simulation [ICAPS15]

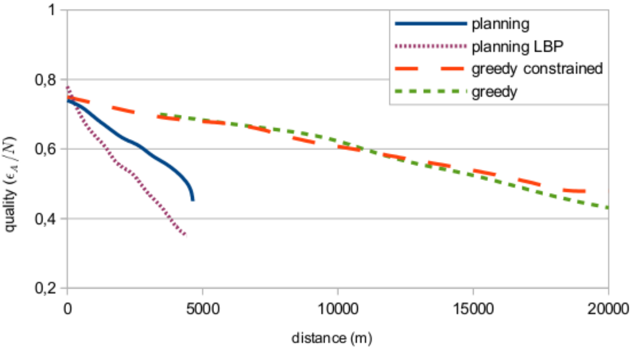

Quality vs. distance for the anisotropic model. 4500m distance limit.

Automated Planning in Precision Agriculture

Good news: There's room for future work!

- **Melissa**: project to map weeds
- Plentiful of other issues that AI Planning can solve:

Automated Planning in Precision Agriculture

Good news: There's room for future work!

- **Melissa**: project to map weeds
- Plentiful of other issues that AI Planning can solve:

Automated Planning in Precision Agriculture

Good news: There's room for future work!

- **Melissa**: project to map weeds
- Plentiful of other issues that AI Planning can solve:
  - Modelling (MDPs)
  - Multicriteria optimization

Automated Planning in Precision Agriculture

Good news: There's room for future work!

- **Melissa**: project to map weeds
- Plentiful of other issues that AI Planning can solve:

- Modelling (MDPs)
- Multicriteria optimization
- Cartography (POMDPs, soft goals)

Automated Planning in Precision Agriculture

Good news: There's room for future work!

- **Melissa**: project to map weeds
- Plentiful of other issues that AI Planning can solve:

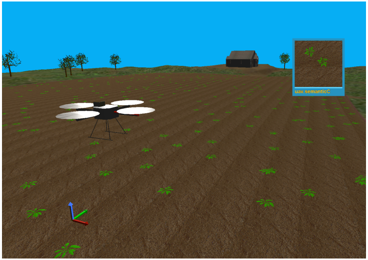

- Modelling (MDPs)
- Multicriteria optimization
- Cartography (POMDPs, soft goals)
- Diseases tracking and spreading modeling

Automated Planning in Precision Agriculture

Good news: There's room for future work!

- **Melissa**: project to map weeds
- Plentiful of other issues that AI Planning can solve:

- Modelling (MDPs)
- Multicriteria optimization
- Cartography (POMDPs, soft goals)
- Diseases tracking and spreading modeling
- Trajectory planning for machines

Automated Planning in Precision Agriculture

Good news: There's room for future work!

- **Melissa**: project to map weeds
- Plentiful of other issues that AI Planning can solve:

- Modelling (MDPs)
- Multicriteria optimization
- Cartography (POMDPs, soft goals)
- Diseases tracking and spreading modeling
- Trajectory planning for machines
- ...

Automated Planning in Precision Agriculture

Good news: There's room for future work!

- **Melissa**: project to map weeds
- Plentiful of other issues that AI Planning can solve:
  - Modelling (MDPs)
  - Multicriteria optimization
  - Cartography (POMDPs, soft goals)
  - Diseases tracking and spreading modeling
  - Trajectory planning for machines
  - ...

All good things come to an End

Thank you!

Sampling and planning

- We dissociate the problem of **selecting the observation sites** from the one of **planning their visiting order.**

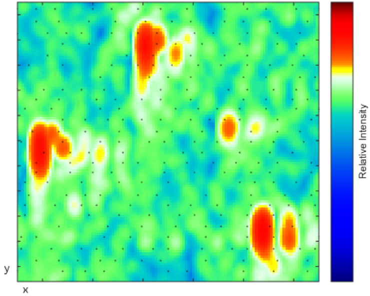

update of the marginals

Sampling and planning

- We dissociate the problem of **selecting the observation sites** from the one of **planning their visiting order.**

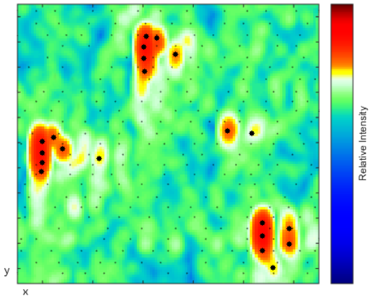

"interesting" points selection

Sampling and planning

- We dissociate the problem of **selecting the observation sites** from the one of **planning their visiting order.**

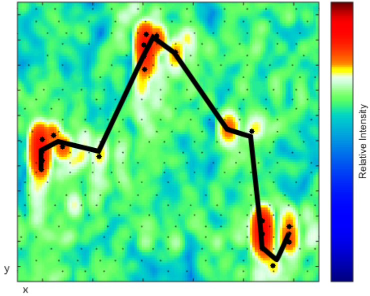

generation of a path (almost) joining them

Path planning is not TSP

Causal links:

An **observation** in a site has influences on the overall quality,

and particularly on **neighbouring sites**,

by **lowering** considerably the expected quality gain after sensing.

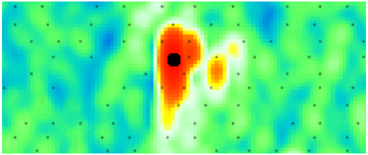

Path planning is not TSP

Causal links:

An **observation** in a site has influences on the overall quality,

and particularly on **neighbouring sites**,

by **lowering** considerably the expected quality gain after sensing.

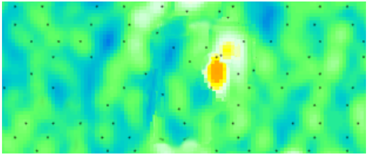

Path planning is not TSP

Causal links:

An **observation** in a site has influences on the overall quality,

and particularly on **neighbouring sites**,

by **lowering** considerably the expected quality gain after sensing.

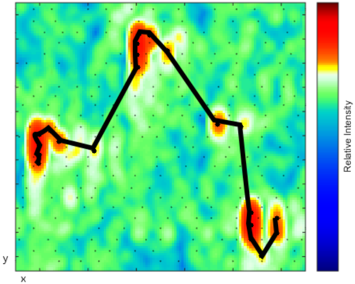

TSP

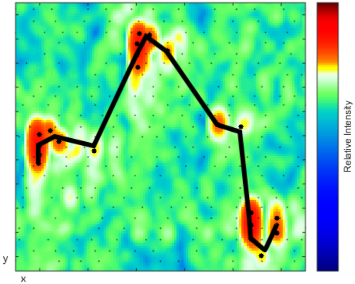

planning

AI Automated (Classical) Planning

Planning is the model-based approach to autonomous behaviour.

Classical planning assume complete information about the environment:

Classical problem: state model $\mathcal{S}(P)$

1. finite and discrete state space S
2. a known initial state $s_0 \in S$
3. a set $S_G \subseteq S$ of goal states
4. actions $A(s) \subseteq A$ applicable in each $s \in S$
5. a deterministic transition function $s' = f(a, s)$ for $a \in A(s)$
6. positive action costs $c(a, s)$

Strips: Basic Language for Classical Planning

Models described in suitable **planning languages** (Strips, PDDL, PPDDL, ...) where states represent interpretations over the language.

Problem in Strips $P = \langle F, A, I, G \rangle$

1. F fluents in P (Boolean variables)
2. A actions
3. I initial situation, **conjunction** of F-literals
4. G goal situation, **conjunction** of F-literals

Operators $a \in A$ are represented by

- the Add list $Add(a) \subseteq F$
- the Delete list $Del(a) \subseteq F$
- the Precondition list $Pre(a) \subseteq F$

From Language to Models

- A Strips problem $P = \langle F, A, I, G \rangle$ defines a **state model** $\mathcal{S}(P)$ where the states $s \in S$ are collections of atoms from F
- (Optimal) solutions for P are (optimal) solutions for $\mathcal{S}(P)$

From Language to Models

- A Strips problem $P = \langle F, A, I, G \rangle$ defines a **state model** $\mathcal{S}(P)$ where the states $s \in S$ are collections of atoms from F
- (Optimal) solutions for P are (optimal) solutions for $\mathcal{S}(P)$
- Planning problem cast as a search problem in the directed graph associated to $\mathcal{S}(P)$

From Language to Models

- A Strips problem $P = \langle F, A, I, G \rangle$ defines a **state model** $\mathcal{S}(P)$ where the states $s \in S$ are collections of atoms from F
- (Optimal) solutions for P are (optimal) solutions for $\mathcal{S}(P)$
- Planning problem cast as a search problem in the directed graph associated to $\mathcal{S}(P)$

From Language to Models

- A Strips problem $P = \langle F, A, I, G \rangle$ defines a **state model** $\mathcal{S}(P)$ where the states $s \in S$ are collections of atoms from F
- (Optimal) solutions for P are (optimal) solutions for $\mathcal{S}(P)$
- Planning problem cast as a search problem in the directed graph associated to $\mathcal{S}(P)$
- Various algorithms: **Blind** search (DFT, BrFS,...) vs. **Informed** search (A*, IDA*, ...)

From Language to Models

- A Strips problem $P = \langle F, A, I, G \rangle$ defines a **state model** $\mathcal{S}(P)$ where the states $s \in S$ are collections of atoms from F
- (Optimal) solutions for P are (optimal) solutions for $\mathcal{S}(P)$
- Planning problem cast as a search problem in the directed graph associated to $\mathcal{S}(P)$

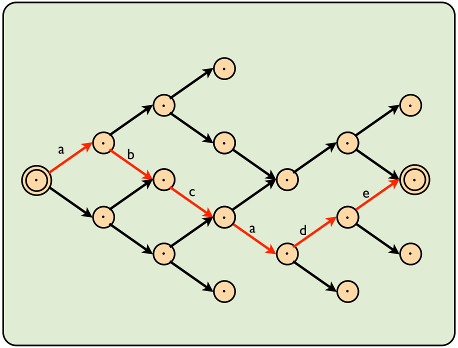

A Classical Planning Model for the navigation task

States: current position, visited sites, battery, accumulated quality

Actions: moving (from $\rightarrow$ to), cost= *dist*
sensing ($\neg$ visited $\rightarrow \bigwedge_{i \in \text{ blanket}}$ visited$_i$), cost= c

Init: initial pose, nothing visited, zero quality

Goal: all sites visited, battery $\geq$ 20%.

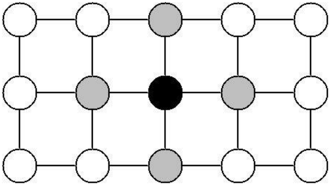

Implementation details

- Two spatial distribution correlationn models: isotropic M1 and anisotropic M2, corresponding to the tillage direction.
- Loopy BP to calculate marginals (approx method) (Murphy, Weiss & Jordan 1999)
- Serialized Iterated Width algorithm (Lipovetzky & Geffner 2012)
- ROS package for the planning module.

Future Work for Melissa

- Real-life tests with an AscTec Firefly UAV.
- A more elaborate image processing module.
- Non-dissociated approach.
- 3D navigation.

Future Work for Melissa

- Real-life tests with an AscTec Firefly UAV.
- A more elaborate image processing module.
- Non-dissociated approach.
- 3D navigation.

Application fields:

- Adaptive methods for weeds mapping
- Monitoring of plant growth
- Mapping of illnesses consequences in crops

Future Work for Melissa

- Real-life tests with an AscTec Firefly UAV.
- A more elaborate image processing module.
- Non-dissociated approach.
- 3D navigation.

Application fields:

- In general on-board decision systems combining:
- Sensing and Planning
- e.g. MORTIMER Project