2017 IEEE International Conference on Robotics and Automation (ICRA)
Singapore, May 29 - June 3, 2017

Online Optimal Active Sensing Control

Paolo Salaris*, Riccardo Spica†, Paolo Robuffo Giordano† and Patrick Rives*

— This paper deals with the problem of active sensing control for nonlinear differentially flat systems. The objective is to improve the estimation accuracy of an observer by determining the inputs of the system that maximise the amount of information gathered by the outputs over a time horizon. In particular, we use the Observability Gramian (OG) to quantify the richness of the acquired information. First, we define a trajectory for the flat outputs of the system by using B-Spline curves. Then, we exploit an *online* gradient descent strategy to move the control points of the B-Spline in order to actively maximise the smallest eigenvalue of the OG over the whole planning horizon. While the system travels along its planned (optimized) trajectory, an Extended Kalman Filter (EKF) is used to estimate the system state. In order to keep *memory* of the past acquired sensory data for online re-planning, the OG is also computed on the past estimated state trajectories. This is then used for an online replanning of the optimal trajectory during the robot motion which is continuously refined by exploiting the state estimation obtained by the EKF. In order to show the effectiveness of our method we consider a simple but significant case of a planar robot with a single range measurement. The simulation results show that, along the optimal path, the EKF converges faster and provides a more accurate estimate than along other possible (non-optimal) paths.

I. INTRODUCTION

The performance of robots and sensing devices is highly influenced by the quality and amount of sensor information, especially in case of limited sensing capabilities and/or low cost devices. Indeed, such information is often used to estimate the state of the system. As a consequence, the problem of optimal information gathering has been studied in the literature in a variety of contexts.

In the field of optimal sensor placements, in [1] an observability-based procedure is used to characterise the gyroscopic sensing distribution of the wings of the hawkmoth Manduca sexta. In [2] the problem of finding the optimal location of sensors on a wearable sensing glove was addressed in order to minimise the error statistics in the reconstruction of the hand pose.

In the field of localization and exploration for mobile robots, it is important to establish if, for a given system, the observation problem, that consists in finding an estimate of the state of the robot/environment from the knowledge of the inputs and the outputs over a period of time, admits a solution [3]. For instance, in [4] a complete observability analysis of the planar bearing-only localization and mapping problem for a nonholonomic vehicle and for all configurations of landmarks with known (markers) and unknown (targets) positions was studied by using the *Observability Rank Condition* (ORC) tool. The problem is shown to be *locally weakly observable* [5], if and only if the number of markers is equal or greater than two. For nonlinear systems, however, the observability may also depend on the inputs. In particular, for the above example, one can show the existence of *singular inputs* that do not allow the reconstruction of the state: this occurs with the vehicle aiming at the two markers or at the target directly.

The problem of developing intelligent control strategies applied to the data acquisition process [6] in order to maximise the amount of information coming from sensors, i.e. to maximise the "distance" from the singular trajectories, is an active area of research often known as *active sensing control*, *active perception* or *optimal information gathering*. One crucial point in this context is the choice of an appropriate *measure of observability* to optimise.

In [7] the condition number of the Observability Gramian (OG) is used as cost function to find the optimal observability trajectory for an Unmanned Aerial Vehicle (UAV) with GPS measurement and an Autonomous Underwater Vehicle (AUV) with range measurements from a single marker. One peculiarity of this work consists in finding an interesting relationship between the entries of the OG and the geometric properties of the trajectory, which can be useful to characterise the shape of the optimal path. In [8] the authors find optimal observability trajectories for first-order nonholonomic systems with nonholonomic states as outputs by maximising the smallest eigenvalue of the OG.

By using the concept of entropy, introduced by Shannon, the authors of [9] devised an observability measure and an optimal navigation strategy for a unicycle vehicle with three bearing measurements w.r.t. known markers. The problem of finding informative paths in Gaussian fields has also been tackled in [10], where the authors propose to maximise a cost function based on the concept of mutual information between the state and the measurements. In [11], a Bayesian optimisation approach for determining the most informative path is used. A search based method for planning in information space was developed in [12] to provide an adaptive policy for mobile sensors with non-linear sensing models. In [13], the optimal trajectories for a team of robots moving on a plane and tracking a target by using relative measurements (distances and bearings) were obtained by minimising the trace of the target's position estimate covariance matrix under suitable constraints.

In the field of simultaneous localization and mapping (SLAM), [14] presents sub-optimal paths that minimise an adaptively weighted combination of the uncertainty of the vehicle pose and of the map features. In [15], a trajectory

* are with INRIA Sophia-Antipolis Méditerranée, 2004 Route des Lucioles, 06902 Sophia Antipolis, France, e-mail: psalaris,prives@inria.fr. † are with the CNRS at Irisa and Inria Rennes Bretagne Atlantique, Campus de Beaulieu, 35042 Rennes Cedex, France, e-mail: riccardo.spica,prg@irisa.fr.

optimisation for target localization by using 3D bearing-only measurements from small unmanned aerial vehicles is proposed. The active sensing problem consists in minimising the trace of the inverse of the Fisher Information Matrix (FIM). Similarly, in [16], the objective is to find the best trajectories and camera configurations for a group of aerial Dubins vehicles, equipped with Pan-Tilt-Zoom cameras, that maximise the trace of the information matrix of the Extended Information Filter (EIF) used for estimating the position of a target on the ground.

Because of the difficulty of the problem only few papers tried to tackle active sensing control from an analytic point of view. In [17], the effects of observer motion on estimation accuracy for bearing-only measurements was addressed. In [18], the problem of naximiising the smallest eigenvalue of the OG was casted within the calculus of variation theory ([19]) and solved for a flat 2D system with only one output measurement.

In this paper, we consider a differentiable approximation of the smallest eigenvalue of the OG (i.e. the Schatten norm) as a measure of observability in order to avoid the non-differentiability in case of repeated eigenvalues. The originality of our method is that it combines an *online* gradient-descent optimization strategy with a concurrent estimation scheme (an Extended Kalman Filter in our case) meant to recover an estimation of the true (but unknown) state during motion. The need for an *online* solution is motivated by the fact that, for a nonlinear system, the observability Gramian is a function of the state trajectories, which, in a real scenario, are not assumed available. By using an offline optimisation method that relies on the initial estimation, the resulting trajectory would be sub-optimal – e.g. in a worst-case scenario of a system that admits singular inputs, the optimal trajectory from the estimated initial position may be the singular one from the real initial position.

In order to make the online optimisation problem tractable from an optimisation point of view, we restrict our attention to the case of non-linear differentially flat systems [20] and we represent the flat outputs with a family of curves (B-Spline) function of a finite number of parameters. To check the effectiveness of our method we will consider a planar robot with a single nonlinear output measurement (the squared distance from a marker) for which a partial analytic solution can also be obtained by applying the results in [18] (see the Appendix in [21]) — an analytic analysis that can serve as a 'ground truth' for validating the results of the proposed gradient-based method (which can, instead, be applied to any differentially flat system and output map for which an analytic analysis may not be possible).

We believe that the formulation of our problem is quite general and the computational efficiency and simplicity of the proposed solution allow applying it to more complex systems than the one considered in this paper as a case study, as, e.g. unicycles and quadrotor UAVs. Moreover, the method can be further generalised by including the environment (e.g. targets) as state variables to be estimated and by introducing additional constraints in order to, e.g., avoid obstacles or reach points of interest.

The paper is structured as follows. In Section II the optimal control problem is introduced while in Section III a solution that combines an online gradient-descent optimization strategy with a concurrent estimation scheme is provided. In Section IV we apply our method to a flat 2D system. Finally, the paper ends with some conclusions.

II. PROBLEM STATEMENT

Let us consider a generic nonlinear dynamics

$$\dot{\boldsymbol{q}}(t) = \boldsymbol{f}(\boldsymbol{q}(t), \boldsymbol{u}(t)), \quad \boldsymbol{q}(t_0) = q_0 \tag{1}$$

$$\boldsymbol{z}(t) = \boldsymbol{h}(\boldsymbol{q}(t)) + \boldsymbol{\nu} \tag{2}$$

where $\boldsymbol{q}(t) \in \mathbb{R}^n$ represents the state of the system, $\boldsymbol{u}(t) \in \mathcal{U}$ is the control input ($\mathcal{U}$ is a subset of $\mathbb{R}^m$), $\boldsymbol{z}(t) \in \mathbb{R}^p$ is the sensor output (the measurements available through the onboard sensors), $\boldsymbol{f}$ and $\boldsymbol{h}$ are analytic functions and $\nu \sim \mathcal{N}(0, R(t))$ is a normally-distributed Gaussian output noise with zero mean and covariance matrix $\boldsymbol{R}(t)$ A well-known observability criterion for system (1)–(2), related to the concept of local *indistinguishable* states [3], [18], is the *Observability Gramian* (OG) $\boldsymbol{\mathcal{G}}_o(t_0, t_f) \in \mathbb{R}^{n\times n}$:

$$\boldsymbol{\mathcal{G}}_o(t_0, t_f) \triangleq \int_{t_0}^{t_f} \boldsymbol{\Phi}(\tau, t_0)^T \boldsymbol{H}(\tau)^T \boldsymbol{W}(\tau) \boldsymbol{H}(\tau) \boldsymbol{\Phi}(\tau, t_0)\, \mathrm{d}\tau \tag{3}$$

where $t_f > t_0$, $\boldsymbol{H}(\tau) = \frac{\partial \boldsymbol{h}(\boldsymbol{q}(\tau))}{\partial \boldsymbol{q}(\tau)}$, and $\boldsymbol{W}(\tau) \in \mathbb{R}^{p\times p}$ is a symmetric positive definite weight matrix (a design parameter), that may be used for, e.g. accounting for outputs with different units and/or, as in this paper, for considering the reliability of the outputs with different noise levels. Matrix $\boldsymbol{\Phi}(t, t_0) \in \mathbb{R}^{n\times n}$, also known as *sensitivity matrix*, is given as $\boldsymbol{\Phi}(t, t_0) = \frac{\partial \boldsymbol{q}(t)}{\partial \boldsymbol{q}_o}$ and verifies the following differential equation

$$\dot{\boldsymbol{\Phi}}(t, t_0) = \frac{\partial f(\boldsymbol{q}(t), \boldsymbol{u}(t))}{\partial \boldsymbol{q}(t)} \boldsymbol{\Phi}(t, t_0)\,, \qquad \boldsymbol{\Phi}(t_0, t_0) = I. \tag{4}$$

If the (symmetric and positive definite) matrix $\boldsymbol{\mathcal{G}}_o$ is full rank over the time interval $[t_0,\, t_f]$, then system (1)–(2) is *locally weakly observable*, i.e. it is possible to (locally) recover the state trajectory $\boldsymbol{q}(t)$ from the knowledge of $\boldsymbol{z}(t)$ and $\boldsymbol{u}(t)$, $t \in [t_0,\, t_f]$. The OG, similarly to the well-known Observability Rank (OR) condition [5], can hence be exploited for verifying the observability of a given nonlinear system. However, while the OR condition can only provide a "binary answer" about the observability of the system, the OG also provides a measure of the amount of information gathered by the sensors along the trajectory followed during motion [22]. One can then attempt naximiization of some performance index of the OG (typically a function of its eigenvalues) w.r.t. the system inputs in order to produce a system trajectory with maximum information content over a future time horizon.

In this paper, we will consider the smallest eigenvalue of the OG as performance index, i.e. $\lambda_{min}\left(\boldsymbol{\mathcal{G}}_o(t_0, t_f)\right)$, and we will determine the optimal control strategy $\boldsymbol{u}^*(t)$, $t \in [t_0,\, t_f]$ that maximises $\lambda_{min}\left(\boldsymbol{\mathcal{G}}_o(t_0, t_f)\right)$. Indeed, the inverse of the smallest eigenvalue of the OG is proportional to the maximum estimation uncertainty, and hence, its *maximisation* is expected to *minimize* the maximum estimation uncertainty of any estimation strategy that could be used, e.g. a EKF [8]. By increasing the value of the smallest eigenvalue, we also expect the convergence rate of the observer to increase. However, it is well-known that considering the smallest eigenvalue of a matrix $\boldsymbol{A}$ as a cost function can be ill-conditioned from a numerical point of view in case of repeated eigenvalues. For this reason, as also done in [23], we will consider the following cost function (aka Schatten norm),

$$\|\boldsymbol{A}\|_{\mu} = \sqrt[\mu]{\sum_{i=1}^{n} \lambda_i^{\mu}(\boldsymbol{A})} \tag{5}$$

with $\mu \ll -1$, as a differentiable approximation of $\lambda_{min}(\boldsymbol{A})$. Moreover, to ensue well-possness of the optimisation problem, we will constrain the solution to be such that the "control effort" (or energy) needed by the robot for moving along the trajectory from t_0 to t_f is fixed and equal to $\bar{E}$, i.e.

$$E(t_0, t_f) = \int_{t_0}^{t_f} \sqrt{\boldsymbol{u}(\tau)^T \boldsymbol{M} \boldsymbol{u}(\tau)}\, \mathrm{d}\tau = \bar{E}\,.^1 \tag{6}$$

The rest of the paper is hence devoted to propose an *online* solution to our problem that will combine an online gradient-descent optimization strategy with a concurrent estimation scheme (an EKF in our case) meant to recover an estimation $\hat{\boldsymbol{q}}(t)$ of the true (but unknown) state $\boldsymbol{q}(t)$ during motion. The need for an *online* solution is motivated as follows: the Gramian $\boldsymbol{\mathcal{G}}_o$ is a function of the whole state trajectory $\boldsymbol{q}(t)$, $t \in [t_0,\, t_f]$, but the state and, in particular, the initial condition $\boldsymbol{q}(t_0)$, is not assumed available. Therefore, in order to perform an offline optimization of $\boldsymbol{\mathcal{G}}_o$, one would necessarily need to rely on some estimated $\hat{\boldsymbol{q}}(t)$ (for example, by integrating the system dynamics from an initial estimation $\hat{\boldsymbol{q}}(t_0)$). Since $\hat{\boldsymbol{q}}(t)$ is only an approximation of the true state evolution $\boldsymbol{q}(t)$, the resulting optimized path would then represent, in general, a sub-optimal one. On the other hand, during the robot motion, it is possible to exploit a state estimation algorithm, such as a EKF, for improving *online* the current estimation $\hat{\boldsymbol{q}}(t)$ of the true state $\boldsymbol{q}(t)$, with $\hat{\boldsymbol{q}}(t) \to \boldsymbol{q}(t)$ in the limit. Availability of a converging state estimation $\hat{\boldsymbol{q}}(t)$ then makes it possible to continuously refine (online) the previously optimised path by leveraging the newly acquired information during motion.

Before presenting our proposed solution in the next Section, we conclude with three remarks.

Remark 1 *In general, a closed-form expression for the sensitivity matrix* $\boldsymbol{\Phi}$ *may not be available, since finding a solution for* (4) *is as complex as finding a solution for* (1). *However, in some particular cases of interest (e.g. the unicycle) matrix* $\boldsymbol{\Phi}$ *can be found in closed form. For all the other cases, a numerical integration of (4) is required.*

Remark 2 *The* integrand *in* (3) *can have full rank only if* $p \geq n$, *i.e. if the number of available measurements is larger or equal than the number of state variables (such as, for instance, in Structure from Motion (SfM) problems [24]). When* $p < n$ *(as in the case studies reported in Sections IV), maximization of* λ_{min} *(or of its differentiable approximation) is still possible but only in an integral sense (i.e. over the whole time horizon* T*) as shown in the reported results.*

Remark 3 *The method we propose in this paper can be used not only with the OG but also with other measures of information, as e.g. the Kalman filter/smoother covariance matrix, related to the mutual information, or the Fisher information matrix. A sensible conclusion is that all these measures are related and hence the obtained solutions should be similar. Future works will be dedicated to determine an analytical relationship of equivalency between all these measures.*

III. AN ONLINE GRADIENT-BASED SOLUTION TO ACTIVE SENSING CONTROL

We assume, as explained before, that a EKF is run by the robot during its motion for producing an estimation $\hat{\boldsymbol{q}}(t)$ (and associated estimated covariance matrix $\hat{\boldsymbol{P}}(t)$) from the collected measurements and applied inputs. Let $\bar{t} \in [t_0, t_f]$ be a generic time instant during the robot motion and partition (3) as

$$\boldsymbol{\mathcal{G}}_o(t_0, t_f) = \boldsymbol{\mathcal{G}}_o(t_0, \bar{t}) + \boldsymbol{\mathcal{G}}_o(\bar{t}, t_f)\,. \tag{7}$$

The first term $\boldsymbol{\mathcal{G}}_o(t_0, \bar{t})$ represents a "memory" of the past information already collected via the available measurements during $t \in [t_0,\, \bar{t}]$ and can be computed by numerically integrating (3) along the past estimated trajectory. This term is obviously constant and cannot be optimised any longer at time $t = \bar{t}$. On the other hand, the second term $\boldsymbol{\mathcal{G}}_o(\bar{t}, t_f)$ represents the information yet to be collected during $t \in [\bar{t},\, t_f]$ that can, instead, still be optimised.

We note that, although the term $\boldsymbol{\mathcal{G}}_o(t_0, \bar{t})$ is a constant parameter w.r.t. the optimization variables, one still needs to include it in the cost function since in general $\lambda_{min}(\boldsymbol{A} + \boldsymbol{B}) \geq \lambda_{min}(\boldsymbol{A}) + \lambda_{min}(\boldsymbol{B})$ (Weyl's inequality). Furthermore, as explained, both terms in (7) are function of the state evolution $\boldsymbol{q}(t)$ which is assumed unknown: therefore, $\boldsymbol{\mathcal{G}}_o(t_0, \bar{t})$ must be evaluated on the estimated state trajectory $\hat{\boldsymbol{q}}(t)$, $t \in [t_0,\, \bar{t}]$, and $\boldsymbol{\mathcal{G}}_o(\bar{t}, t_f)$ on a "predicted" state trajectory generated from the current $\hat{\boldsymbol{q}}(\bar{t})$.

In order to make the problem more tractable from an optimization point of view (and, thus, to better cope with the real-time constraint of an online implementation) we now make two simplifying working assumptions. First, we restrict our attention to the case of non-linear differentially flat[2] systems [20]: as well-known, for these systems one

[1]Indeed, in general $\lambda_{min}\left(\boldsymbol{\mathcal{G}}_o(t_0, t_f)\right)$ could be unbounded from above without the control effort E.

[2]The class of flat systems includes some of the most common robotic platforms such as, e.g. unicycles, cars with trailers and quadrotor UAVs, and in general any system which can be (dynamically) feedback linearized [25].

以上内容仅为本文档的试下载部分，为可阅读页数的一半内容。如要下载或阅读全文，请访问：https://d.book118.com/618016060052007036