

Subjective Mapping

Michael Bowling

Department of Computing Science
University of Alberta
Edmonton, Alberta, Canada
bowling@cs.ualberta.ca

Dana Wilkinson

School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada
d3wilkin@uwaterloo.ca

Ali Ghodsi

Department of Statistics and Actuarial Science
University of Waterloo
Waterloo, Ontario, Canada
aghodsib@uwaterloo.ca

Abstract

Extracting a map from a stream of experience is a key problem in robotics and artificial intelligence in general. We propose a technique, called subjective mapping, that seeks to learn a fully specified predictive model, or map, without the need for expert provided models of the robot's motion and sensor apparatus. We briefly overview the recent advancements presented elsewhere (ICML, IJCAI, and ISRR) that make this possible, examine its significance in relationship to other developments in the field, and outline open issues that remain to be addressed.

Introduction

Mapping is a foundational problem of robotics. It commonly involves estimating the position of obstacles in the environment through the use of a range sensor, or landmarks through a vision-based sensor. In its most general form, though, mapping can be viewed as the construction of a predictive model of the environment. For example, maps containing the positions of obstacles allow a robot to predict future range readings as it navigates. Using a map with landmarks, a robot can predict that a sequence of actions will result in some target landmark being visible with high probability. When viewed in this general form, mapping is a foundational problem for the entire AI enterprise.

Recent advancements in robotic mapping have advocated a probabilistic approach to the problem (Thrun, Burgard, & Fox 2005), where uncertainty is represented explicitly with probability distributions. In this paradigm, a complete predictive model is specified through two probability distributions: a motion model and a sensor model. Let x_t be the robot's state at time t , a_t be the robot's action leading to t , and z_t its observation at t . The motion model specifies the probability distribution $p(x_t|x_{t-1}, a_t)$ and the sensor model $p(z_t|x_t)$. The traditional formulation of mapping has assumed that much of these models are already specified and known a priori. This has proven to be very successful for many robot platforms and environments where rigorous models are well known. For the general AI problem, though, such nearly complete models are not likely to be available.

We propose an approach to mapping which requires no a priori knowledge of models. This short paper will outline this approach and the recently developed techniques that make it possible, as well as examine its significance for the broader artificial intelligence endeavor. We first introduce the basic principles of subjective mapping. We then briefly discuss the recently developed algorithm called action respecting embedding (Bowling, Ghodsi, & Wilkinson 2005), and its extensions (Wilkinson, Bowling, & Ghodsi 2005; Bowling *et al.* 2005), which form the crux of the subjective mapping approach. We then discuss how subjective mapping relates to other recent advances in artificial intelligence. We finally outline what open issues remain to be addressed.

Subjective Mapping

Consider the motion and sensor models described above. These models can be seen as connecting the *subjective* quantities of actions and observations to the robot's state representation, which is traditionally an *objective* quantity. For example, robot state may be defined as the position (x, y) and orientation (θ) from a fixed origin with units in meters and radians. We call such a representation objective because it is independent of the robot itself, *i.e.*, regardless of whether the robot is differential drive, omnidirectional, legged, equipped with sonar, laser, or camera. The motion model, then, provides the connection between the robot's subjective actions, *e.g.*, wheel velocities, and their effect on position and orientation. Likewise, the sensor model provides the connection between the robot's position and its subjective observations, *e.g.*, range readings from a sonar or camera images.

The goal of subjective mapping is to extract fully specified motion and sensor models only from a stream of experience. This is accomplished by removing the requirement of an objective state representation. In fact, a key problem in subjective mapping is extracting an appropriate representation from a stream of subjective experience. It is clear that any extracted representation won't correspond to objective quantities such as meters, since it is the expert provided models, which provided such a translation, that are being replaced. Such a subjective representation, though, can still form the basis of fully specified motion and sensor models and so define a complete predictive model.

Good Representations. Before discussing a technique for extracting a state representation from experience, we first examine the properties of a good representation. Or, why is (x, y, θ) such a commonly chosen objective representation? We propose that three key components are desirable in a representation. First, it should be low dimensional. This is desirable both from a computational viewpoint and that of Occam’s razor, suggesting simpler models are more often correct. Second, it should be a sufficient representation for describing the robot’s observations. Hence, it is a natural representation for defining a sensor model. Third, the robot’s actions in the representation should correspond to simple transformations. Hence, it is a natural representation for the motion model.

Notice that the first two desirable properties are exactly the goals of the well-studied problem of dimensionality reduction. In fact, dimensionality reduction has become a common preprocessing step of high-dimensional sensors in robotics. Dimensionality reduction alone, though, fails to find representations where actions are simple transformations making the result a poor representation for mapping. We have proposed a new technique for dimensionality reduction, which respects actions in order to extract a representation meeting all three properties of good representations. We very briefly describe this result in the next section along with details of how the representation can be used as the basis for constructing motion and sensor models.

Action Respecting Embedding

Suppose we have a stream of experience from a robot: observations, $z_1 \dots, z_n$, and associated actions, a_2, \dots, a_n . In this work, it is assumed that actions come from a small discrete set of uninterpreted labels. The goal is to find a low dimensional representation of z_i , which we’ll write $\Phi(z_i) \equiv x_i$ for which actions correspond to simple transformations. In other words, we want to extract a useful map representation, based on the above principles, strictly from a stream of experience.

Recently, *nonlinear dimensionality reduction* techniques have been successfully used to map a high-dimensional dataset into a smaller dimensional space. Semidefinite Embedding (SDE) is one such technique (Weinberger & Saul 2004). SDE learns a kernel matrix, which represents a nonlinear projection of the input data into a more linearly distributed representation. It then uses Kernel PCA, a generalization of principle components analysis using feature spaces represented by kernels, to extract a low-dimensional representation of the data (Schölkopf & Smola 2002). The kernel matrix K is learned in SDE by solving a semidefinite program with a simple set of constraints. The most important constraints encode the common requirement in dimensionality reduction that the non-linear embedding should preserve local distances. In other words, nearby points in the original input space should remain nearby in the resulting feature representation. The optimization maximizes the trace of K , *i.e.*, the variance of the learned feature representation, which should minimize its dimensionality.

SDE, like other dimensionality reduction techniques, ignores our action labels and so won’t necessarily find a low-

Table 1: Algorithm: Action Respecting Embedding (ARE).

<p>Algorithm: ARE($\ \cdot\ , (z_1, \dots, z_n), (a_1, \dots, a_{n-1})$)</p> <p>Construct neighbors, N, with local metric $\ \cdot\$.</p> <p>Maximize $\text{Tr}(K)$ subject to $K \succeq 0, \sum_{ij} K_{ij} = 0,$ $\forall ij \quad N_{ij} > 0 \vee [N^T N]_{ij} > 0 \Rightarrow$ $K_{ii} - 2K_{ij} + K_{jj} \leq \ z_i - z_j\ ^2$, and $\forall ij \quad a_i = a_j \Rightarrow$ $K_{(i+1)(i+1)} - 2K_{(i+1)(j+1)} + K_{(j+1)(j+1)} =$ $K_{ii} - 2K_{ij} + K_{jj}$</p> <p>Run Kernel PCA with learned kernel, K.</p>

dimensional space where actions have a simple interpretation. The recent Action Respecting Embedding (ARE) algorithm extends SDE to make use of the action labels (Bowling, Ghodsi, & Wilkinson 2005). The ARE approach is to constrain a semidefinite optimization to only consider representations where actions correspond to simple transformations. In particular, the actions in the chosen representation must be distance-preserving transformations: those consisting only of rotation and translation. Therefore, for any two inputs, z_i and z_j , the same action from these inputs must preserve their distance in the learned feature space. For action a , let f_a be the transformation corresponding to that action in the chosen representation. For f_a to be distance preserving in the representation defined by Φ the following must hold,

$$\forall i, j \quad \|f_a(\Phi(z_i)) - f_a(\Phi(z_j))\| = \|\Phi(z_i) - \Phi(z_j)\|. \quad (1)$$

Now, consider the case where $a_i = a_j = a$. Then, $f_a(\Phi(z_i)) = \Phi(z_{i+1})$ and $f_a(\Phi(z_j)) = \Phi(z_{j+1})$, and Constraint 1 becomes,

$$\|\Phi(z_{i+1}) - \Phi(z_{j+1})\| = \|\Phi(z_i) - \Phi(z_j)\|. \quad (2)$$

In terms of the kernel matrix, this can be written as:

$$\forall i, j \quad a_i = a_j \Rightarrow$$

$$K_{(i+1)(i+1)} - 2K_{(i+1)(j+1)} + K_{(j+1)(j+1)} =$$

$$K_{ii} - 2K_{ij} + K_{jj} \quad (3)$$

Notice that this constraint is linear in the entries of the kernel matrix. ARE simply adds Constraint 3 into SDE’s usual constraints to arrive at the optimization and algorithm shown in Table .

Results. Figure 1 shows one example of an ARE discovered representation along with the representation extracted with SDE. The domain used is IMAGEBOT, where a robot is moving an observable image patch around a larger image. The robot’s trajectory in objective coordinates corresponds to an “A” shape, which is largely mirrored in the ARE representation. ARE is also able to determine that pairs of actions are inverses and the pairs are orthogonal, moving along independent dimensions. Dimensionality reduction alone extracted a topologically correct representation, but it’s clear that finding a model of the robot’s actions in this space would be nigh impossible. The original work should be consulted for complete details and further results (Bowling, Ghodsi, & Wilkinson 2005).




Figure 1: Results of Action Respecting Embedding (ARE) and Semidefinite Embedding (SDE) on extracting a representation from a single “A” shaped trajectory.

Extensions. ARE takes a stream of subjective experience and extracts a representation appropriate for mapping, which is implicitly encoded in the extracted trajectory, x_1, \dots, x_n . A fully specified model requires motion and sensor models that make use of this representation. In an extension of the original work, we showed how the action transformations could be extracted by solving an extended orthonormal Procrustes problem (Wilkinson, Bowling, & Ghodsi 2005). We also showed how the action models were accurate enough to allow for planning in the extracted representation. In a second extension, we described a procedure for using the trajectory and action transformations to construct a fully specified probabilistic model, which could then be used for localization (Bowling *et al.* 2005). These past publications should be consulted for more details and results.

Related Work

We now explore the relationship of this work to other recent developments in robotics and learning. We also briefly discuss the potential implications for the broader field of AI.

SLAM. Simultaneous localization and mapping (SLAM) is the problem of tracking both uncertainty in the robot’s pose as well as uncertainty in the location of obstacles and landmarks (Thrun, Burgard, & Fox 2005). It is the quintessential example of objective mapping where the robot’s state representation is given, along with partially specified models, *i.e.*, the entire motion model and the sensor’s noise model is commonly considered known. In many cases, where the robot and environment is well studied and predictable, this has proven to be a very effective approach. If the robot’s locomotion or sensing apparatus, or even the environment, is not known a priori there can be no partially specified models. Simultaneous localization and mapping with a completely unknown robot is perhaps the “Grand SLAM” challenge, and it’s not clear how current objective mapping approaches could handle the complete absence of models and remain objective. As the models are providing the critical connection between subjective and objective quantities it may, in fact, be impossible.

Predictive Representations of State. New models of dynamical systems have recently been proposed in the rein-

forcement learning community (Littman, Sutton, & Singh 2002). Predictive state representations (PSRs) are one such model for discrete action and observations, where state is represented as predictions on possible future action-observation sequences. It has been shown that such representations are capable of compactly representing a large range of systems (Singh, James, & Rudary 2004) and also can be learned from a stream of interaction with the system (Wolfe, James, & Singh 2005).

The goals of PSRs and subjective mapping are very similar. Both strive to discover an appropriate state representation from only subjective experience. Both learn a fully specified predictive model using the discovered representation. PSRs, though, are addressing a much more challenging problem due to an impoverished set of observations: possibly only two discrete observations for a complicated multi-dimensional state space. In subjective mapping, and robotics in general, robot sensations are not lacking for distinguishable features. Even two seemingly “identical” hallways may have subtle identifying features in a high-dimensional sensor like a camera (*e.g.*, office numbers, lighting differences, or wall scuffs). The problem is often more of identifying salient features in the high-dimensional space, rather than discovering new features based on future predictions. This is not to say that state aliasing cannot be handled with the described subjective mapping approach, which will be discussed later, but just that it may not be as critical.

Other Areas of AI. Although inspired by the problem of reducing the amount of expert knowledge required for robot mapping, the subjective mapping approach is addressing a fundamental problem of AI. How can an agent extract a model of the world from only subjective experience? The reformulation of the problem to one of dimensionality reduction opens up a number of possible connections with other fields of artificial intelligence. In particular, adding constraints that encode temporal interrelationships is a powerful and general idea. It can break the locality assumptions which has been observed to be a limitation with existing non-linear dimensionality reduction techniques (Bengio & Monperrus 2005). Many fields that already make use of techniques for reducing dimensionality, such as natural language, bionformatics, or computer vision, may also be able to benefit from introducing appropriate “action respecting constraints” on the low dimensional representation.

It would be amiss to neglect mentioning the insightful work of Pierce and Kuipers’ (1997), which aimed to achieve a similar goal of mapping without the requirement of expert provided models. At the lowest level they use an application of dimensionality reduction to identify the principal components of action. Learning in these local spaces are then used to find homing sequences and path following behaviors that allow the construction of a global topological map. The subjective mapping approach seeks instead to construct a global geometric map. Not only is this more aligned with the highly successful probabilistic robotics paradigm, it also removes the need for homing sequences to reacquire the robot’s position to keep location uncertainty local. On the other hand, the low-level decomposition of the action space might com-

plement the subjective mapping approach very well.

Open Issues

There are a number of open issues that make this work, while highly promising, admittedly preliminary.

Scalability. The key limitation of the current formulation of subjective mapping is the computational cost of action respecting embedding, which is needed to discover the state representation. Although the ARE optimization is a semidefinite program and is therefore convex, it is both time and memory intensive to solve even small problems (*i.e.*, with sequences of length one hundred or less) using standard toolbox solvers. Since convex programming is a very active area of research in optimization, we can expect the time complexity to improve in coming years. Even so, it is still likely efficiency improvements are needed to make ARE practical for robotic and general AI applications.

A special purpose solver that takes advantage of the inherent structure in the constraints is one possibility. Since the constraints are growing quadratically with the length of the sequence, a more efficient approach to addressing these constraints would result in huge gains in efficiency. Representation alignment is another possibility making use of a divide-and-conquer approach. Long trajectories could be broken into overlapping shorter trajectories, and a representation discovered for each. A more efficient global optimization problem could then be constructed that found a way to “stitch” the overlapping local representations into a single consistent global representation.

State Aliasing. As mentioned above, state aliasing is not as problematic when working with high-dimensional sensing apparatus as is commonly faced in the POMDP or PSR framework. The subjective mapping optimization, though, still has the power to handle some amounts of aliasing. The distance preserving constraints imply that the same action from the same state will have the same effect, modulo noise. If ARE incorrectly collapses two distinct states together then their future states when taking the same action sequence must also be the same, even if the associated observations differ dramatically. Hence, the optimization is implicitly trading off the penalty in separating two states with similar observations for the penalty of collapsing two states with very different observations. Further investigation is still needed to explore the practicalities of this tradeoff inherent in the optimization. For example, if state aliasing is not given a large enough penalty any data stream can be explained as a single dimensional trajectory through time where actions have no effect and no state is ever revisited.

Continuous Action Spaces. Currently the constraints used in ARE require actions to be discrete labels. A more natural representation of action is as a continuous vector. Re-encoding the distance preserving constraints of ARE for continuous actions is non-trivial, but Pierce and Kuipers’ (1997) work on finding an orthonormal basis for actions is a promising possibility.

Obstacles and Costs. Although we have described how a fully specified predictive model can be constructed from

experience, an open issue is how this model can be used. The common use of a map in robotics is navigation, where a path, *i.e.*, sequence of actions, is planned to reach some goal location. Path planning is made difficult by obstacles or traversal costs for various regions of the state space, which are detected by local sensors on the robot (*e.g.*, a bump sensor for identifying obstacles, or an energy expenditure measurement.) Since an objective representation is no longer maintained, it is not obvious how to incorporate these local traversal costs sensors into the “map”. As with objective mapping it is critical that costs are generalized beyond simple point observations, but subjective spaces will likely have to rely more heavily on adaptive techniques to learn the extent of such generalizations.

Conclusion

This paper has given a brief overview of recent developments toward the goal of subjective mapping, which seeks to construct predictive models of the world from only subjective experience. This is a foundational problem for robotics and artificial intelligence in general. We outlined both connections to other recent developments in the field and open issues that still need to be addressed.

References

- Bengio, Y., and Monperrus, M. 2005. Non-local manifold tangent learning. In *Advances in Neural Information Processing Systems 17*.
- Bowling, M.; Wilkinson, D.; Ghodsi, A.; and Milstein, A. 2005. Subjective localization with action respecting embedding. In *Proceedings of the International Symposium of Robotics Research*.
- Bowling, M.; Ghodsi, A.; and Wilkinson, D. 2005. Action respecting embedding. In *Proceedings of the Twenty-Second International Conference on Machine Learning*, 65–72.
- Littman, M.; Sutton, R. S.; and Singh, S. 2002. Predictive representations of state. In *Advances in Neural Information Processing Systems 14 (NIPS)*, 1555–1561. MIT Press.
- Pierce, D., and Kuipers, B. 1997. Map learning with uninterpreted sensors and effectors. *Artificial Intelligence* 92:169–229.
- Schölkopf, B., and Smola, A. 2002. *Learning with Kernels*. Cambridge, Massachusetts: MIT Press.
- Singh, S.; James, M. R.; and Rudary, M. R. 2004. Predictive state representations: A new theory for modeling dynamical systems. In *Uncertainty in Artificial Intelligence: Proceedings of the Twentieth Conference (UAI)*, 512–519.
- Thrun, S.; Burgard, W.; and Fox, D. 2005. *Probabilistic Robotics*. MIT Press.
- Weinberger, K., and Saul, L. 2004. Unsupervised learning of image manifolds by semidefinite programming. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 988–995.
- Wilkinson, D.; Bowling, M.; and Ghodsi, A. 2005. Learning subjective representations for planning. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, 889–894.
- Wolfe, B.; James, M. R.; and Singh, S. 2005. Learning predictive state representations in dynamical systems without reset. In *Proceedings of the 22nd International Conference on Machine Learning (ICML)*, 985–992.