
DANA WILKINSON, University of Waterloo

Learning Useful Subjective Representations

In a variety of domains it is desirable to learn a representation of an environment defined by a stream of sensori-motor experience. In many cases such a representation is necessary as the observational data is too plentiful to be stored in a computationally feasible way. In other words, the primary feature of a learned representation is that it must be compact, summarizing information in a way that alleviates storage and retrieval demands.

This admits a new way of phrasing the problem: as a variation of dimensionality reduction. There are a variety of well-studied algorithms for the dimensionality reduction problem. We argue that any of these can be useful for learning compact representations as long as additional constraints to the problem are respected, namely that the resulting representation is useful in the context of the actions which generated the observations.

Here, we formalize the problem of learning a subjective representation, clearly articulating solution features that are necessary for a learned representation to be “useful”; the actions must correspond to simple and consistent transformations in the learned representation. Further, we briefly present a possible solution to the newly defined problem and demonstrate its effectiveness for reasoning, planning and localization.