On the Role of Learning in Anchoring
(position paper)

Silvia Coradeschi silvia.coradeschi@aass.oru.se
Alessandro Saffiotti alessandro.saffiotti@aass.oru.se

Center for Applied Autonomous Sensor Systems
Örebro University, S-701 82 Örebro, Sweden

Abstract

A situated agent may include a symbolic subsystem that uses symbols to denote objects in the physical world. *Anchoring* is the problem of connecting these symbols to the perceptual representations of the same objects. Learning is an important, and in some cases essential means to acquire the basic ingredients needed to perform anchoring. In this note, we discuss some issues on the role of learning in anchoring.

The anchoring problem

The focus of this note is the connection between symbol- and sensor-level representations of objects in autonomous robotic systems embedded in a physical environment. We define *anchoring* as the problem of creating, and maintaining in time, this connection. Anchoring can thus be seen as a special case of symbol grounding where the symbols denote physical objects.

Anchoring must necessarily occur in any physically embedded system that comprises a symbolic reasoning component. A typical example is the problem of connecting, inside an autonomous robot, the symbol used by a symbolic planner to refer to a physical object to the data in a perceptual system that pertains to the same object. This connection must be dynamic, since the same symbol must be connected to new percepts when the same object is re-acquired.

All existing embedded systems that comprise a symbolic reasoning component must necessarily incorporate a solution to the anchoring problem. However, this solution is typical given in an implicit form, hidden in the code, and it is developed on a system by system basis for a specific domain. To our knowledge, the first general, although preliminary definition of anchoring was reported in (Saffiotti 1994). More recently, the notion of anchoring has been extended in order to cope with some subtle issues encountered in real applications (Coradeschi & Saffiotti 1999) and (Saffiotti & LeBlanc 2000). Those studies were empirical in nature, leading to a pre-theoretical notion of anchoring. (Coradeschi & Saffiotti 2000) is a first attempt to state the anchoring problem in a computational theory, where the entities and the functionalities involved are formally described. In order to discuss the role of learning in the anchoring problem, we present below an outline of this theory.

A computational theory of anchoring

We consider a situated agent that includes a symbol system and a perceptual system. The basic ingredients of the theory are:

- A *symbol system* Σ, which contains individual symbols (variables and constants), predicate symbols, and an inference mechanism. Our interest is directed to the individual and predicate symbols. Examples of individual symbols are *bottle1, rack2*, while examples of predicates are *cylindrical, big, small, red.*
- A *perceptual system* Π, which includes percepts and attributes. A percept is a structured collection of measurements assumed to originate from the same physical object; an attribute is a measurable property of percepts. Examples of percepts are image regions identified as representing objects. Examples of attributes are *color, shape.*
- A *predicate grounding relation* g, which embodies the correspondence between unary predicates and values of measurable attributes. For instance, g may encode the correspondence between the predicate *red* and the RGB values derived by the sensor data. We do not make any assumption about the origin of g: for instance, g can be hand-coded by the designer of the system, or it can be learned by the system.

The task of anchoring is to create and maintain a correspondence between a symbol used in Σ to denote an *object* in the world, and the percepts generated in Π by observing that object. The pivot to this correspondence is the matching, through the g relation, between the symbolic properties that are predicated of the object, and the observed perceptual attributes. This correspondence is represented by an internal data structure, called an *anchor*, that uniquely represents a given physical object.

We have identified three main functionalities that are necessary to perform the anchoring task for a given symbol x:

Find: create an anchor the first time that the object denoted by x is perceived;

Track: continuously update the anchor while observing the object; and

Reacquire: update the anchor when we need to reacquire the object after some time that it has not been observed.

These are detailed in (Coradeschi & Saffiotti 2000), where we also give examples of their use in two autonomous systems.

The role of learning

Learning can be used to acquire the basic ingredients needed in the anchoring process. We focus here on three aspects:

- learning the predicate grounding relation;
- learning a fusion strategy to combine several component properties; and
- learning the dynamical behavior of a property.

Below, we illustrate these aspects on examples taken from an application developed in our laboratory: an electronic head (Wide, Kalaykov, & Winquist 1999), which, among other sensors, is equipped with an artificial nose and tongue. These sensors try to imitate the corresponding human senses, and they can differentiate one substance (typically, food samples) from another. The task is to endow a mobile platform, equipped with the electronic head, with the ability to identify objects on the basis of a symbolic description including visual, odor, and taste information.

Learning the predicate grounding relation. The predicate grounding relation g can be given a priori to the system and be time invariant. This may not be adequate in many applications; instead, the system should be able to derive its own grounding relation or at least improve the existing one. For instance, in case of the artificial nose and tongue in the electronic head there is no well-established numerical model representing the true function of the sensors. Therefore, the relation between predicates denoting properties that relate to smell and taste and the actual measurements acquired from these sensors has to be learned rather than given a priori. A preliminary experiment in this sense is reported in (Loutfi *et al.* 2001), where an artificial neural network has been trained to recognize vanilla-, lavender-, and yogurt-like aroma using an artificial nose.

Learning to combine several properties. The task of anchoring is to connect symbols denoting objects to the sensory data derived from the corresponding physical objects. The focus on individual *objects* forces us to deal with the co-occurrence of several properties. In fact, while classes can be identified by one predicate, a description of a specific object typically needs the combined use of several predicates, like "a small cup of yogurt with vanilla flavor". The co-occurrence of several properties can influence the sensor data collected, and the g grounding relation for a combined property can be more complex than the superposition of the individual relations for the individual properties. For instance, the recognition of the components yogurt and vanilla in the odor produced by a yogurt with vanilla flavor can require not only the grounding relation for the property of "smelling-like-yogurt" and "smelling-like-vanilla", but also a function indicating how the sensor values tend to change when both substances are present. Such a function would, in most cases, be difficult to define a priori, and it would then have to be learned from data.

Learning property dynamics. The anchoring process presupposes two entities, an object to be identified and an observer. Both entities are located in space and time. In general, the predicate grounding relation g may change as the relative spatial relation, or the time, change. How to maintain and reacquire the anchor while these changes occur is therefore an essential issue in anchoring. For example, the data provided by an odor sensor tend to drift over time, and the sensor can react very differently to the same substance at different moments in time. In order to be able to re-identify an object, a model of this drifting must be used. Since analytical models of the sensor's drift are not available, learning seems to be the only viable alternative here.

References

Coradeschi, S., and Saffiotti, A. 1999. Anchoring symbols to vision data by fuzzy logic. In *Qualitative and Quantitative Approaches to Reasoning with Uncertainty*, LNAI. Berlin, Germany: Springer. 104–115.

Coradeschi, S., and Saffiotti, A. 2000. Anchoring symbols to sensor data: preliminary report. In *Proc. of the AAAI-2000 Conf.*, 129–135.

Loutfi, A.; Coradeschi, S.; Duckett, T.; and Wide, P. 2001. Odor source identification by grounding linguistic descriptions in an artificial nose. In *Proc. of the SPIE conference on Sensor Fusion.*

Saffiotti, A., and LeBlanc. 2000. Active perceptual anchoring of robot behavior in a dynamic environment. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 3796–3802.

Saffiotti, A. 1994. Pick-up what? In Bäckström, C., and Sandewall, E., eds., *Current trends in AI Planning.* Amsterdam, Netherlands: IOS Press. 266–277.

Wide, P.; Kalaykov, I.; and Winquist, F. 1999. The artificial sensor head: A new approach in determining of human based quality. In *Fusion'99, Second International Conference on Multisource-Multisensor Information*, 1144–1149.