

## HOW SITUATIONS ARE CONSTRUCTED

*Notes for Meeting on 23 May 2001 (4:30 pm)*

John Gero, Rob Saunders, Greg Smith, PakSan Liew and Julie Jupp<sup>1</sup>  
*Key Centre of Design Computing and Cognition  
Department of Architectural and Design Science  
University of Sydney NSW 2006  
Australia*

**Abstract.** This document is an outline of the items discussed during the situatedness discussion group meeting held on 23 May 2001<sup>2</sup>. In particular, this document presents the discussion of how situations are constructed.

### Executive Summary

The following items were covered:

- ☞ A situation is *constructed* as a result of interactions between an agent and its environment;
- ☞ The *situation defines the world*, in which the agent operates, in terms of variables and processes.
- ☞ The variables define aspects of the world that the agent is interested in and the processes operate on these variables.
- ☞ Both variables and processes need to be constructed to form the situation as the agent interacts with the environment;
- ☞ The process of forming a situation is recursive as the agent form meta-heuristic about what it is doing.
- ☞ The starting point of an agent is defined as a single point in time when the agent was first created.
- ☞ After this starting point an agent is always situated although it may become dormant at times.
- ☞ An agent's situation is only constructed from scratch once at the starting point, after that time the agent's situation is modified as it gains experience.
- ☞ A construction strategy for starting up an agent is proposed.

---

<sup>1</sup> Not in any particular order.

<sup>2</sup> Many points and opinions were discussed in the meeting. This document contains only the most central points to the discussion and the general consensus of agreement that was reached. For clarity, some of the points have been expanded to add useful/interesting background information and some may not appear here in the same order that they were raised in the discussion.

- ☞ An agent starts with a predefined architecture that defines the structure and initial state of its processes and internal variables.
- ☞ Enough intuition and biases are given to the agent for it to survive long enough for its learning processes to begin the construction of its situation from early experiences.
- ☞ As the grounded representations learned in the construction of an agent's situation become more useful the initial representations are either modified to ground them in experience, or they are ignored, as their usefulness diminishes with time.
- ☞ A developer can provide instructions that direct an agent to a stopping point or rules can be encoded within the agent about the appropriate stopping criteria.
- ☞ Intrinsic stopping criteria can indicate the complete exploration of solution space while extrinsic stopping criteria can indicate the exhaustion of resources beyond the agent's control.
- ☞ An agent can become a specialist if it can maintain its own design agenda across design projects. The agent's activities are "suspended" after the termination of each design project and the agent is allowed to select design projects in which it can learn.
- ☞ The relative importance of initial conditions and the external environment for the development of situated agents can be determined by two experiments where one experiment exposes two identical agents to different sets of experiences and the other experiment exposes two non-identical agents with the same set of experiences.
- ☞ The notion of the situation may not exist for the agent at all, i.e. an agent may not be able to reflect upon its situation.
- ☞ Any agent is incapable of observing its entire situation because to do so it is necessary for the agent to take itself out of its own situation, which is clearly impossible.
- ☞ An agent can observe aspects of its situation and compare situations with previous situations, and this is what we mean when we say that an agent reflects upon its situation.
- ☞ In general it is very worthwhile for an agent to have knowledge about its situation.
- ☞ Knowledge of the relationship between current and past situations can be especially useful for predicting the consequences of actions and the emergence of unseen phenomena.

### Situations in Design

An agent's situation consists of a set of variables that represent aspects of the world important to the agent and a set of processes that operate on those variables. Similarly, a design situation consists of a set of design variables that define the *domain*<sup>3</sup> that a design agent is working in and a set of design processes that define a *strategy* for operating on those design variables. A situation defines the *world* in which the agent acts. (Interestingly, this implies that the domain of an agent is a subjective notion.)

Importantly, a situation cannot be retrieved like a template from a database; instead a situation is something that must be constructed dynamically by an agent. In other words, an agent must construct the sets of variables and processes that define its situation in response to its world. A

---

<sup>3</sup> The domain defines what the agent will be doing.

## HOW SITUATIONS ARE CONSTRUCTED

situation is constructed from what is happening in the environment, what is happening within the agent, what experiences the agent have and the reflections that the agent have on its action. All these are relative to a particular point in time,

### Meta Level Heuristics

Although an agent constructs its situation it cannot be fully aware of its situation in an objective sense. In other words, an agent cannot “step outside” its situation to inspect it. The agent can only reflect at a meta-level, the processes and actions it performs at a specific point in time as it interacts with the environment. This meta-level heuristic is relative to these actions and processes. The act of reflection in itself forms another situation that sits at another meta-level. Thus the process of forming a situation is recursive in nature with one situation enveloping another.

### How does an Agent Start the Construction of its Situation?

*Given that an agent constructs its situation, how does the construction process start?*

The starting point of an agent is defined as the point in time when the agent was first “spawned”. There can only be one starting point for any agent. After this starting point an agent is always situated although it may become dormant at times. An agent’s situation can be thought of as in continual flux and as such it is only constructed from scratch once at the starting point, after that time the agent’s situation is modified as it gains experience. What concerns us here is the problem of the initial construction of an agent’s situation: if situations must be constructed how does the construction process start? In other words, how can an appropriate construction strategy be chosen in the absence of a situation in which to ground the selection process?

To better understand this problem, consider the Rete algorithm<sup>4</sup> that is commonly used in rule-based systems to select rules to fire. The Rete algorithm is popular because it has been shown to be a remarkably efficient way to select rules to fire, more efficient than many other strategies for rule selection by orders of magnitude. The Rete algorithm takes advantage of the empirical fact that, in most rule-based systems, much of the knowledge base, i.e. the domain, is fairly fixed from one rule operation to the next. Although new facts arrive and old ones are removed at all times, the percentage of facts that change per unit time is generally fairly small. Conceptually, the Rete algorithm selects a rule to fire by testing only new added facts against the already selected subset of recently applicable rules, because many of these rules are still likely to be applicable to the new situation.

Unfortunately, if we take the notion of situatedness literally, a newly spawned agent will have no situation from which to select the variables (facts) or processes (rules). Consequently, situated agents must be equipped with some other mechanism than a strategy like the Rete algorithm to determine what to do next. Therefore to allow the construction of an initial situation the designer of an agent must supply some initial variables and processes that can act as an initial situation. An agent is therefore started with a defined architecture and some initial content, we may think of these as a set of innate abilities and instincts that are built into the agent to guide its early development.

---

<sup>4</sup> The principal reference for the algorithm is “Rete: A Fast Algorithm for the Many Pattern/Many Object Pattern Match Problem”, *Artificial Intelligence*, 19, pp 17-37, 1982.

One solution available to the agent designer is to develop a set of processes, or biases, that do not require the pre-existence of internal variables in order to be applicable. We might think of these as being analogous to reflexes that allow a new agent react appropriately to external stimuli without the need for a corresponding internal representation. Like Brook's subsumption-based autonomous agents, design agents produced in this way would have a set of "hard-wired", reactive behaviours. However, if these design agents were given the ability to learn from their experiences they will have the capacity to construct experience-based situations. Unfortunately, the downside to this approach is that the learning process may take considerable time and the abilities of a mature agent may be difficult to predict from the supplied biases.

Alternatively, an agent designer can add an initial set of facts about the world. For example, a designer might provide an agent with a set of design prototypes, i.e. schemas for typical examples of designs within the domain, that have reasonable default values. Care should be taken when starting an agent with predefined representations such as design prototypes because knowledge given to the agent in this way weakens its claim for situatedness. The agent's claim to being situated is weakened because predefined representations are without groundings in the experiences of the agent. Taken to an extreme this means building something like an expert system that dictates the knowledgebase of an agent, typical of early non-situated approaches to artificial intelligence. However, for practical purposes it will most likely be necessary to provide some representations to an agent when it starts to allow it to learn quickly and controllably.

Although the use of an initial representation weakens the situatedness of an agent it need not destroy the notion of situatedness permanently for agents that can learn. Initial representations certainly weaken an agent's claim to situatedness at first but as the agent experiences the world and learns, it will change its internal representations to match the perceived world according to its experiences thereby grounding them and increasingly situating itself. It is assumed that the learning process must be able to change the representations given. Ultimately, it may not matter what the details of the initial representations given to the agent at the start, all that matters to the mature agent would be the initial direction, i.e. bias, that was given to the agent to begin its learning process: the initial representations will either be augmented, or forgotten.

To summarise, an agent starts with a predefined architecture that defines the structure and initial state of its processes (e.g. sensors, effectors, memory etc.) and internal variables (a.k.a. content, knowledge, or intuitions). An agent should be started with enough intuition, i.e. pre-programmed concepts, and behavioural biases to be able to survive long enough for its learning processes to begin the construction of its situation from early experiences. As the grounded representations learned in the construction of an agent's situation become more useful the representations initially used to bootstrap the process will either be modified to ground them in experience, or be ignored, or forgotten, as their usefulness diminishes with time.

### **Reflex Actions and Intuition**

*Reflex actions* is a better term than *intuition* because "intuition" may imply encoded knowledge when we refer to the initial content to be given to the agent when it was first created. In order for the agent to have the capability to learn such that eventually they become competent in their environment (not to just sit there doing nothing or worse: to die), it is required that the agent be given the following initially: at least some reflex actions, such as babies being able to feed as soon as they are born, and enough machine learning bias that it can learn.

## HOW SITUATIONS ARE CONSTRUCTED

**Machine Learning Bias**

Learning can be defined<sup>5</sup> as generalising from observations such that future novel instances can be handled correctly, and it is inherently inductive. For example:

All men are mortal

Socrates is a man

-----

Socrates is mortal

is a valid deduction. However:

All tigers I have seen are yellow and black

-----

All tigers are yellow and black

is not deductively valid but it is what learning tries to do. Because generalising from a set of observations is not deductively valid, without learning biases a system would have no reason to choose one generalisation over another and learning would not be possible. There are two main kinds of bias:

- ☞ Representation bias: defines the concept space within which any learned concept will be found, e.g. we decide a priori to learn only concepts in conjunctive normal form, or decide to learn a Markov model, or whatever.
- ☞ Search bias: defines the order that the concept space will be searched as training instances are received, e.g. we may start with a concept covering the whole concept space and reduce it as exceptions are found, or start with a null concept and add to it as exceptions are found, and so on.

A good book on machine learning should explain this, such as the one by Langley.

**Termination Point**

*When should the agent terminate its design activities?*

The opposite problem from how to start a design agent is how to stop it. The problem is that, like human designers and unlike standard optimisation algorithms, a suitably complex situated design agent may never come to a definitive stopping place. As developers of situated design agents we must provide some means for an agent to determine when it has reached a 'good place to stop'. As discussed above agent developers can do this either by instructing a learning agent about appropriate stopping criteria or by encoding rules in the agent that embody our common-sense notions of what makes an appropriate time to stop.

Intrinsic stopping criteria for a design agent might include a sense of boredom. Boredom can be modelled as the inverse of the cumulative total of the interestingness of situations over some recent time period. Boredom in relation to a lack of novelty in a situation can be used as an

---

<sup>5</sup> It should be noted that some forms of adaptation can also be considered learning as well as being contingent upon a definition of "observations" that includes "introspections" which isn't a commonly held meaning.

indicator of the amount that an agent is learning as it designs. One criterion for stopping therefore would be a lack of new situations that allow an agent to learn. This would indicate that an agent had exhaustively explored its current design space. To maintain interest in the design, a new design space will be required. This prompts the termination of the current design project with the submission of the best design solutions from the current design space.

Important extrinsic stopping criteria for a design agent include the exhaustion of resources beyond the agent's control. In other words, the designing activities of an agent should be terminated when resources such as time or money run out for that particular project. In general, a design agent will not be in control of the amount of time or money available for a design project, instead these resources are more directly controlled by the agent's client. Consequently, useful design agents must be able to operate successfully within the confines of the resources made available to it by external agents. This suggests that design agents need the ability to plan the allocation of these resources. Allocating resources such as time and money for unpredictable tasks like design is a difficult problem as can be seen in many examples of human design efforts that run over budget and break deadlines. However, if situated design agents are to participate in design projects with external clients that are resource limited then they must be given some way of planning their activity to come to a conclusion at an appropriate time.

An interesting perspective on a design agent's relationship to design tasks comes from considering the end of a design project as merely the suspension of design activities rather than the extinction of the agent. At the end of a design project a design agent's activities are "suspended" and remain dormant until they are needed again for a new design project. However, aspects of the situation constructed during a design project will persist to the next design project when it arrives. This may imply continuity in the design process for the agent not immediately apparent to external observers, as new design tasks will be perceived within the context of previous design tasks by the agent. From an agent's perspective, the commencement of a new design task is a resumption of a continuing design process carried over from the previous project. A situated design agent can thereby maintain its own design agenda across design projects especially if a design agent has some choice over which design briefs it accepts. A situated design agent that has a free choice of the projects it takes on will naturally choose, possibly guided by innate biases, design projects that are interesting to it and that in some way contribute to its learning. Consequently, situated design agents may well become specialists through the selection of design projects without the need for us to pre-program them with specialist abilities.

### **Experiments in the Construction of Situations**

*What sorts of tests for situatedness can we conduct on design agents?*

Two interesting experiments could be conducted to investigate environmental effects on the situatedness of agents. In the first experiment two identical agents are given different sets of experiences from which to learn. In the second experiment two non-identical agents are provided with the same set of experiences. The goal of these experiments is to determine the relative importance of initial conditions and the external environment for the development of situated agents. As such, these experiments follow the on-going discussion of the relative importance of nature versus nurture in human, and animal, development.

The first experiment, using two identical agents exposed to different worlds, tests the degree to which early biases are superseded by learned responses. During the course of the experiment it

## HOW SITUATIONS ARE CONSTRUCTED

is assumed that the situations constructed by each agent will diverge with each differing experience. At the end of the experiment, important differences in the situations constructed by each agent can be tested by exposing both to the same world and examining the differences in their responses. For example, two identical design agents could be used, each exposed to a different set of design problems. They are then tested to see what differences have emerged by observing their responses to the same design problem. This experiment would provide an indication of the extent to which the nurturing process has affected each agent's situation.

The second experiment, using two non-identical agents exposed to the same world, tests the degree to which early biases shape the learning of responses. During the course of the experiment it is assumed that the situations constructed by each agent will converge with each identical experience. As with the first experiment, exposing both agents to the same world at the end of the experiment and examining the differences in their responses can be used to test the important differences between the situations constructed. For example, two non-identical design agents could be used, each exposed to the same set of design problems. They are then tested to see what differences remain by observing their responses to the same design problem. This experiment would provide an indication of the extent to which the initial state of the agent has affected the development of its situation.

### **Knowledge of the Situation**

*If an agent constructs a situation can it see it and is it useful for it to see the situation?*

A situation is an emergent property of an agent: its past and present experiences affect the construction of a situation but the situation does not explicitly represent an agent's experiences. Consequently, the notion of the situation may not exist for the agent at all, i.e. an agent may not be able to reflect upon its situation. This is obviously the case for simple agents such as Brook's robots that are situated in the sense that they react to the environment in context sensitive ways, but have no capacity to reflect upon the details of the situation that they construct in the process of reacting.

In fact, the notion of a situation as it has been described here is an abstract concept used to describe the emergent state of an agent by an observer. As a result, any agent is incapable of observing its entire situation because to do so it is necessary for the agent to take itself out of its own situation, which is clearly impossible. Additionally, a situated agent will never be in exactly the same situation twice, the mere fact that an agent has experienced a situation previously will change the situation the second time around. However, this does not mean that an agent cannot observe aspects of its situation and compare situations with previous situations, and this is what we mean when we say that an agent reflects upon its situation.

To reflect upon aspects of a situation an agent must represent those aspects in such a way that they are amenable to appropriate consideration. For example, to reflect upon its situated interactions with the world an agent must first represent the most pertinent aspects of its situation as some form of knowledge about world that it can observe. A reflective agent therefore requires meta-level processes that allow it to reason about its operational processes. Clearly, for an agent to reflect on its situation it requires a limited form of self-awareness.

Why would an agent need to reflect upon its situation? It only makes sense for an agent to have knowledge about its situation when it can do something with that knowledge. In other words, it is only worth having the knowledge about a situation when it can be used to inform some purposeful activity. However, in general it is very worthwhile for an agent to have knowledge about its situation. Knowledge of the relationship between current and past situations can be

especially useful for predicting the consequences of actions and the emergence of unseen phenomena. In particular, this ability can be used to attend to those aspects of a situation that could not be predicted from previous experience. Attending to unexpected phenomena is a useful meta-heuristic for agents that need to autonomously learn in complex environments, as these aspects of the situation are likely to be the most instructive in learning about the nature of the agent's world. We have used this type of meta-heuristic to model curiosity in design agents that allows agents to autonomously explore design spaces with very little predefined domain knowledge.