

After having detected its goals, an autonomous intelligent agent must continue to process them. The study aims to conduct further research on other aspects of Goal

perceived or inferred facts about the current and future world (data-driven and anticipated data-driven) and when the agent ensures that its desires about the world are reflected in its beliefs (goal-driven).

EXPERIMENTAL ROADMAP

To test the applicability of goal detection concepts elaborated in this work, a simulated household robot that is capable of detecting its own goals was developed

to achieve a goal.

During goal detection, the agent continuously updates its beliefs about the environment and then reasons with these beliefs together with its motives, desires and intentions in order to determine the appropriate goals. Figure 4 illustrates a Goal

knowledge base. Subsequently, the agent creates and inserts a desireIntention tuple.

Anticipatory Decision

Apart from maintaining an updated set of beliefs about the current state of the environment, an agent may also acquire and maintain a set of beliefs regarding the future state of the environment. Anticipated facts about the world are obtained by the agent through communication with an external agent and internal reasoning. For example, the fact that it might rain anytime now can be inferred from the perceived

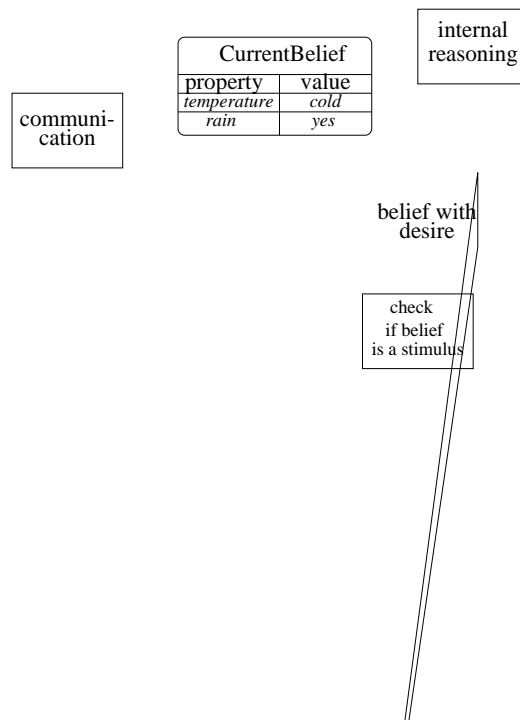
mechanism and condition/reason for detection. Once created, it is *inserted at the top* of the agent's database of desireIntentions. It will be used by the agent when determining the validity and priority as well as assessing the feasibility of the detected goals.

n i n p o o o o n r r n n n o
o r o r r i r v n n r i n o p p i

providing comfort for the house), *and a state that depicts time whether discrete, continuous, interval* (e.g. concept of “always” or “at all times”, morning, etc.).

Current data-driven goal detection is a bottom-up, reactionary process that is activated whenever new data is sensed from the environment. The agent’s purposive behavior is driven by its reactions to the data being sensed. It induces the agent to respond to situations that need immediate attention. Figure 2 shows how the process of current data-driven goal detection is undertaken. Basically, it is a continuous process that: a) compares whether its current beliefs are reflected in its desires or b) checks whether observed data in the environment represent situations that would activate a desire(s). If a current belief does not match a desire (e.g. as shown in Figure 2, the current belief of having a cold temperature does not match the desire of having suitable temperature), a corresponding desire and intention are triggered. Also, if a stimulus is present (e.g. as shown in Figure 2, the presence of rain threatens the motive of protection), a corresponding desire (preventing rain from entering) is detected and subsequently intentions are triggered (closing the door and closing the window) to realize the desire.

A desire can be realized by one or more intentions. For example as shown in Figure 2, to realize the desire of having a suitable temperature inside the house, the agent can either close the window or open and adjust the airConditioner. The agent selects the intention(s) that is applicable under present conditions.



After an intention has been selected, a *desireIntention* tuple is created. A *desireIntention* tuple consists of arguments namely: motive, desire, intention,

also be prescribed by an external agent. For example, a housekeeper can acquire some of its goals from its employer. These types of goals are called communicated desires. They are represented as predicates with desire and intention as arguments. The example depicted in Figure 1 for communicated desires simply states that an external agent requests the agent to turn off the lights for the desire of keeping the lighting suitable.

A desire is detected when certain conditions are present in the environment (stimuli). Depending upon the underlying conditions, the agent may choose a method(s) of action that transforms a desired state into reality. Such chosen method represents the agent's intention to achieve its detected desire. It may be represented as a predicate with desire, condition when applicable, and intention as arguments. Figure 1 presents three examples of intentions. The first one states that in order to achieve the desire of having sufficient power supply, the agent must execute an action of recharging its power supply. The next two examples of intention predicates suggest that a single desire may be realized by more than one course of action. For example, to make the temperature suitable (which is currently cold), the agent can either close

descriptions of states that the agent wishes to be achieved (or avoided) when performing its role. They are activated when the conditions that stimulate the agent to satisfy a motive emerge. However, some desires (if not most of them) are not capable of being directly satisfied by the agent. *Intentions* are the desired states that the agent must act upon in order satisfy its desires. An intention is a direct goal (desired state) that an agent can carry out through actions while a desire may just be a causal effect of an agent's accomplished intention.

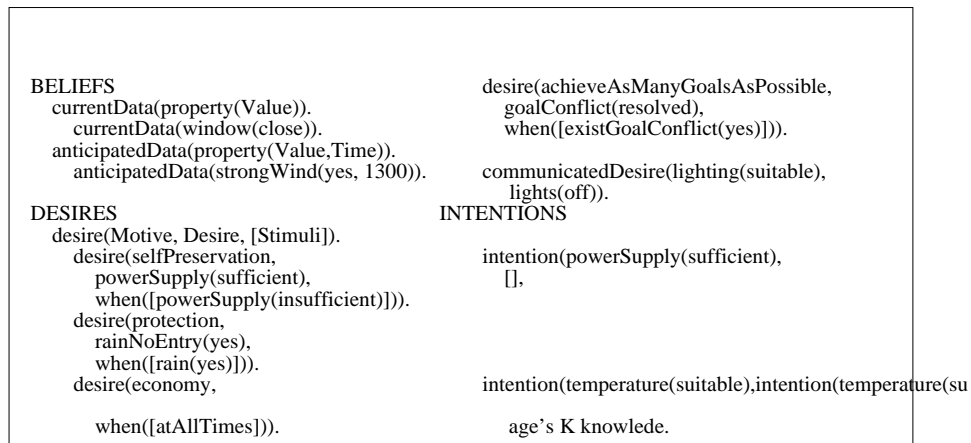
To further clarify the relationship among motive, desire and intention, an example is given. As part of its role, an agent such as a housekeeper is expected to keep the house clean and comfortable. Cleanliness and comfort are examples of a housekeeper's motives. To keep the house comfortable, the housekeeper must see to it that the temperature, air quality, sound, etc. within the house are kept suitable. Keeping the temperature suitable is an example of a housekeeper's desire. It is a state of the world that can be perceived, aspired for and be realized. To fulfill a desire, some actions need to be carried out. For instance, to make the temperature suitable (which is currently cold), the window must be closed or the heater needs to be opened. Closing the window or opening the heater are actions that need to be performed to achieve the desire of making the temperature suitable. Thus, the states of having the window closed or the heater opened are considered to be intentions of the housekeeper. They are achieved directly by the housekeeper through the performance of some actions.

This paper views intentions and desires as the *goals* of the agent. *Goals* are defined to be the end results that an agent would like to achieve or avoid. They are specific and measurable outcomes toward which an agent's efforts are directed within specified time and cost constraints [Roui93]. Some goals are *essential* to the overall mission of the agent (e.g. closing the window to keep the temperature suitable); others are *innovative* (e.g. turning off the lights to economize in the use of electricity) such that their accomplishment improves the agent's performance in fulfilling its mission (i.e., better, faster, cheaper, safer, etc.). However, the agent may also possess some goals (called *meta-goals*) that are not directly related to the agent's mission, but are necessary in guiding or controlling the agent's reasoning and behavior (e.g resolving goal conflicts).

Figure 1 shows how motives, desires and intentions can be represented. Structurally, motives and desires are represented by means of a predicate called "desire", which consists of arguments namely: motive, desire, and the environmental

f 5099(fici5r1"3.02 Tw (must pr)3isfupplathe envireora innd btions thaiiecthe condipr)3isfupplatbe tow

external). Such knowledge depicts a partial description of the total environment. It is acquired through continuous *perception* of the external environment, *communication* with an external agent and *internal reasoning*. Figure 1 illustrates the nature and structural representation of knowledge that is relevant to Goal Detection. To endow an agent with the capability of detecting its own goals, it must possess and continuously reason with: *beliefs* it has about the world, *motives* that drive it to change its current world into a more desirable one, *desires* that need to be achieved (or avoided) as part of its mission and *intentions* that need to be carried out to accomplish the desires and eventually affect the external world.



An agent's set of beliefs is made up of *explicit* and *derived* propositions. Explicit beliefs are acquired by the agent through perception and communication. Derived beliefs are inferred by the agent through reasoning. The agent's beliefs also embody the agent's knowledge about the *current* and *anticipated* states of the world. Current beliefs are characterized as explicit or derived facts about the present status of object properties in the environment. They can be described as a collection of *property-value pairs*, that is, each known property of the environment has a corresponding value associated with it. A property-value pair is also called a *state*. As shown in Figure 1, an example of a current belief (see currentData predicate) is the object property "window" that has a value of "close". This simply means that the agent believes that the window is presently closed. On the other hand, anticipated beliefs are obtained through communication and internal reasoning. They are descriptions of how the world may likely look in the future. They may be represented as a predicate with property, value, and likely time that the belief may occur as arguments. An example (see anticipatedData predicate in Figure 1) is the likely occurrence of a strong wind at 1300.

Aside from possessing a set of beliefs, an agent is also regarded as a *purposive system* that is being behaviorally driven by a set of motives. A *motive* compels the agent to act in order to achieve a desired state. It may be described as a high level form of desire that is satisfied when other lower level desires that are directly related to its fulfillment are met. Associated with motives are desires. *Desires* are

well as different conditions for its accomplishment. Because of some contextual and computational limitations (e.g. limited resources) attributed to the agent, the unpredictability of incessantly changing environmental conditions, and the variations in the requirements of each goal, the agent's goals are impossible to satisfy all at the same time. Thus, in the presence of constraining conditions, the agent must determine wisely which among its goals has to be invoked and pursued.

Goal Detection is the process of identifying the relevant goals that an agent needs to accomplish to affect the external world. It is a process of continuous reasoning about the agent's changing beliefs about the world, motives, desires and intentions. In the past, although a number of AI researchers have addressed some issues relating to Goal Detection, it has not been accorded greater emphasis and elaboration. Wilensky and Chin's work on Goal Detection involves the triggering of a goal by a matching situation through situation-goal pairs, that is when a particular situation has been reached or sensed, its corresponding goal is triggered. [Wile83, Wile90, Chin88] In Luria's work on Knowledge Intensive Planning, the expressed goals (specified by the user) are straightforwardly detected by the KIP. [Luri88]

Georgeff, et al.'s Procedural Reasoning System is endowed with psychological attitudes of beliefs, desires and intentions. It is a reactive system and selects goals based on its current beliefs. [GL87] Beaudoin (whose work is influenced by Georgeff's PRS) theorized the use of goal generactivators (management procedures and programs that execute asynchronously with the management procedures) to generate the agent's goals. [Beau94]

Maes proposed a bottom-up mechanism for behavior selection. This mechanism chooses the appropriate behavior by responding to environmental conditions as well as making use of internal motivation. [Maes91] In the same vein, Norman and Long, and Moffat stressed the importance of motivation in the generation of appropriate behaviors or goals for the agent. Norman and Long's work [NL94] emphasizes the role of motivation in the creation of proactive goals. In Moffat's WILL [Moff94], focus of attention is dictated by the strength of the activation charge.

Individually, the above cited works do not provide a comprehensive and total approach to goal detection. Although each work proposes a particular means of detecting goals, it does not offer an integrated method of detecting goals. In light of the lack of a total approach to goal detection in autonomous intelligent agents, this work attempts to provide for one. It shall identify the nature and representation of knowledge required of the agent to effectively carry out goal detection, determine and develop the mechanisms that should be employed by the agent during the process of goal detection and develop an integrated architecture for goal detection. In brief, goal detection will be addressed through representation, mechanisms and architecture.

P E A

Broadly speaking, an autonomous intelligent agent can be thought of as a physical or abstract entity that has a representation of the world it inhabits. This representation embodies the agent's knowledge about its environment (internal and

Bulo94b, Moff94, NL94).

Most of the AI planning systems that have been designed and developed so far operate on the basis that a predetermined goal or set of goals already exists. That is, they start with a set of goals that have been previously identified and assigned by the user (or designer). In the end they either accomplish their goals, or when unexpected conditions that go beyond their reasoning capabilities occur, they just decide to completely give up their further pursuit of these goals and then become idle after an attempted search for a solution(s) has failed. [Carb82]

Most of the earlier AI planning systems are not capable of exhibiting “genuine” goal-directed behavior. They are not built with the capability to monitor and analyze the environment and to formulate their own goals. They lack the ability to take advantage of opportunities being presented by changing environmental conditions. They do not have the flexibility of behavior to change their focus of attention or direction when needed most. That is, when most trying circumstances occur unexpectedly (e.g. impending danger), they lack the consciousness to abandon or suspend their current job and to seek for a more pertinent goal(s) to pursue. They are also bereft of the ability to make a rational choice in the light of present circumstances. [Carb82, Geor87, Wile83] To sum it up, as Colby has suggested, most AI planning systems are of “slave-mentality”, that is, they do not have a purpose of their own (KF85 cites Colb78).

To be truly called “intelligent”, an agent must not only be capable of knowing how to achieve its given goals; preferably, it must also have the capability to formulate its own goals. It must reason and decide what goals to pursue and when to achieve them. It must be able to identify its own goals (Goal Detection), assess their feasibility (Goal Feasibility Assessment), prioritize them (Goal Prioritizing), evaluate their validity as to whether they have to be continued, terminated, suspended, or modified (Goal Evaluation/Validation), and modify them in the light of present circumstances (Goal Modification).

The development of this research work is anchored on the above cited rationale for goal formulation. This work also posits the fundamental belief that an intelligent agent’s success in pursuing its goal-directed activities largely depend on the behavior it exhibits during the formulation of goals. However, this paper shall concentrate on the Goal Detection aspect. Specifically, it endeavors to investigate vital issues concerning goals and Goal Detection. To prove the applicability of Goal Detection concepts developed in this work, test run results involving a simulated household robot detecting its own goals will be discussed at the latter part.

2 A DE EC

An autonomous intelligent agent operating in a real world domain has a role to perform (e.g. housekeeper, teacher). Associated with the agent’s role are goals whose fulfillment serves as a barometer to gauge the agent’s overall performance of its role. These goals vary in *importance*, *urgency* and *potential for realization* as the environmental context changes. Each goal may require a different set of resources as

Goal Detection: Representation, Mechanisms, Architecture

Remedios de Dios Bulos
remedios@cogs.susx.ac.uk

School of Cognitive and Computing Sciences
University of Sussex
Falmer, Brighton BN1 9QH
United Kingdom

CSRP 402

January, 1996

ABSTRACT *Today, there appears to be considerable research enthusiasm towards endowing an autonomous intelligent agent with the capability of formulating its own goals. Goal Formulation refers to the intelligent behavior that an agent exhibits when reasoning (and deciding) what goals to pursue and when to pursue them. It may be described as the integration of several processes namely: Goal Detection, Goal Assessment, Goal Prioritizing, Goal Evaluation, and Goal Modification. However, this paper concentrates only on Goal Detection. It seeks to identify the ways by which an autonomous intelligent agent can detect its goals.*

Goal Detection is the process of identifying the goals that an agent needs to accomplish to affect the external world. It is a process of continuous reasoning about the agent's changing beliefs about the world, motives,

INTRODUCTION.uk