

Modeling behavior, an action selection approach

Alejandro Guerra Hernández
Departamento de Inteligencia Artificial
Universidad Veracruzana
Sebastián Camacho No 5,
Xalapa, Ver., México, 91000
aguerra@mia.uv.mx

February 23th, 1998

Abstract

Modeling behavior is central in building autonomous agents. An autonomous agent is a system that has perceptual faculties to sense its environment and effectors that can modify it. The behavior of an agent is defined as a mapping from the sensors' input to the effectors' output. Agents behave continually generating actions that satisfy agent's goals in the long term. Choosing the right actions has been approached in Artificial Intelligence, planning included, as a search in problem state space. We believe that behavior based approach emphasizes on situated activity, real time response, and autonomy, concepts that are closely related to the agents' intrinsic characteristics. In this presentation we briefly define the concept of autonomous agent and the action selection problem. Then we introduce behavior networks, a behavior based mechanism for action selection, and finally, we present the monots project (modeling and simulating foraging behavior in howler monkeys) where the ideas discussed here have been applied.

1 Autonomous Agents

An agent is a system that has perceptual faculties to sense its environment and effectors to modify it through actions [9]. The behavior of an agent is defined as a mapping from perceptions to actions. Formally we can define an agent as a tuple $\langle S, A, G, B \rangle$ where S is the set of environment states as perceived by the agent, A is the set of actions that the agent can perform, G is a set of goals, and B is a mapping from S^* (streams of perceptions) to

A. In this way, an agent behaves continually acquiring information about the environment and applying B to generate actions that in the long term satisfy the goals it must achieve.

The environment can be described as a tuple $\langle S, A, W \rangle$ where S is the set of environment possible states, A is the set of actions that an agent can display, and W is a function from $S \times A$ to S representing word transitions. We are interested in agents situated in dynamic, non deterministic environments. An environment is dynamic if it changes while the agent is deciding what action to take. An environment is non deterministic if it is not possible to determine its next state on the basis of current state and the action chosen by the agent. Observe that the environment might be perceived as non deterministic by the agent if it is not able to differentiate all environment states. Non determinism can be modeled with a probabilistic W function.

In order to successfully face such a kind of environment, an agent need to be autonomous [2, 17]. Autonomy is related to decision criteria used by the agent to choose actions. As far as what has been described here, an agent can be conceived as an automata. The n variables that define agent perception create an n dimensional space. Automata has a set of rules to generate the action on the basis of a particular region of this space. So, we can control the automata if we know its state and the rules governing its behavior. Observe that an automata can not deal with situations not included in these rules. An autonomous agent does not follow rules, it evaluates alternatives exhibiting some degree of self-control choosing actions based on state, history, and evaluation criteria. Evaluation criteria is usually a cost function with a non minimal, non unique model, and thus it is not completely observable. So we can influence agent behavior, but not absolutely control it.

As we said, the behavior of an autonomous agent is goal oriented, but other systems are also goal oriented. In artificial intelligence goals have been defined in relation to the problem state space concept, which is composed by the set S of possible states in problem domain, a set O of operators that the system can apply to transform one state into another, and goals. A goal is a pair $G = \{s_C, S_f\}$, where $s_C \in S$ is the current state, S_f is the set of final states, and $s_C \notin S_f$ is not one of the final states. The intention to reach one of the final states applying a sequence of operators is implicit in the system. Because of the problem state space concept, we used to think about goals like achievable states. That is, once the state-goal is achieved, it does not influence behavior any more. Agents must have a second kind of goals known as homeostatic goals. A homeostatic goal is one that is always influencing behavior, for example a state where the agent must stay. An autonomous agent must have an extent set of built-in goals designed

to deal with environment and, in addition, it may acquire goals from other agents, humans included, and from the environment. We say the agent is autonomous if it is able to choose at any given time the goal that must be achieved, and also to generate subgoals to achieve its top level goals.

2 Action Selection

As we can see, one of the key problems in modeling and designing autonomous agents is how choice of alternative actions should be performed. The action selection problem is defined as follows: given a set of actions that an agent can perform, given a set of goals that the agent must archive, given perceptual sequences from environment, how can the agent decide its next action in order to progress toward goals archivement? Three different aspects of this question have been identified [1]:

1. The chosen action should be appropriate to the agent's internal and external circumstances, that is, goals and environment.
2. Conflicts between different actions should be solved in such a way that mutually contradictory actions are not executed at the same time, and switching between actions is effected smoothly and coherently.
3. The behavior of the agent should serve goal archivement, including of course, homeostatic goals.

There are two fundamental questions from the designer perspective: should action selection be a centralized or a decentralized process? and should conflict resolution use a fixed priority or should it be dynamically reconfigurable? Ourselves and others [8] argue that in order to gain robustness, action selection must be decentralized, and that autonomy requires dynamically reconfigurable conflict solution.

3 Behavior Networks

Behavior Networks [10, 11, 12] were proposed to model action selection as an emergent property of an activation and inhibition dynamics among the different actions an agent can take, avoiding in this way, central forms of control. In this mechanism, an agent is defined as a set of actions represented as competence modules. A competence module i can be described by the tuple $\langle c_i, a_i, d_i, \alpha_i \rangle$. The first three elements are lists of propositions describing

module's preconditions and expected effects: c_i is the list of preconditions that must be observed by the agent in order to activate the module i ; a_i is the list of propositions that may be observed after the module i is activated; and d_i is the list of propositions that may not be observed anymore after module i is activated. The last element α_i , is a positive real variable called activation level. Competence module i is executable if any proposition $p \in c_i$ is observed true. A competence module becomes active if it is executable and its activation level is greater than a threshold.

Competence modules are linked in a network through three different types of links: successor, predecessor, and conflict.

There is a successor link from module i to module j for each proposition $p \in a_i \cap c_j$. There is a predecessor link from module i to module j for each proposition $p \in c_i \cap a_j$. There is a conflict link for each proposition $p \in c_i \cap d_j$.

Observe that even when competence modules definition is very similar to planning operators, such as the ones used in STRIPS [3], their purpose here is different: they are only used to link the network. Competence modules use these links, together with information from goals and environment to activate and inhibit each other through activation level. In this way we obtain dynamically reconfigurable arbitration and relevant behavior to agent's circumstances. This dynamics is as follows:

Current situation. The activation level α_i of a competence module i is increased for each proposition $p \in c_i$ that is observed to be true.

Goals. The activation level α_i of a competence module i is increased for each proposition $p \in a_i \cap G$, where G is the set of goals that the agent must archive, expressed as a list of propositions. It's possible to define two types of goals: *achievable goals*, goals that once archived won't influence action selection process; and *homeostatic goals*, goals that the agent tries to archive permanently.

Protected goals. The activation level α_i of a competence module i is decreased for each propositions $p \in d_i$ that is a satisfied achievable or homeostatic goal.

Successors. The executable competence module i increases the activation levels α_j of its successors j when $p \in a_i \cap c_j$ is true, putting $\alpha(\gamma/\delta)$ on its successors.

Predecessors. The non executable competence module i increases the activation levels α_j of its predecessors j when $p \in c_i \cap a_j$ is not true, putting α on its predecessors.

Conflicts. The competence module i decreases the activation levels α_j of its conflicts j for each $p \in c_i \cap d_j$, taking away $\alpha(\delta/\gamma)$ from its conflicts.

The action selection algorithm running the network performs the next loop:

1. The impact of current situation, goals, and protected goals is computed.
2. The impact of successors, predecessors, and conflicts is computed.
3. A normalization function is applied to keep the sum of activation levels constant over time.
4. The competence module that fulfills the next conditions is activated:
i) the module is executable; ii) Its activation level is greater than the activation threshold; iii) Its activation level is greater than the activation level of other modules fulfilling (i) and (ii). If two or more modules fulfill these conditions, one of them is chosen randomly. If a module is activated its activation level is set to zero. If there is not an active module, threshold is lowered by some percentage.

Five parameters are used to configure a Behavior Network:

1. θ is the activation threshold ($\theta \geq 0$).
2. π is the average activation level. At any time t the sum of all activation levels must be π .
3. γ is the goals and predecessor influence ($\gamma > 0$).
4. σ is the environment and successor influence ($\sigma \geq 0$).
5. δ is the goal protected and conflict influence ($\delta \geq 0$).

The action selection can be made more or less goal oriented varying the ratio from γ to σ . If $\sigma = 0$, traditional backward chaining will be executed. If $\gamma > \sigma > 0$ then an opportunistic, goal oriented behavior is exhibited by the agent. High values for θ avoid local maximum in activation level space, but selection becomes slower. When γ and σ are relatively small in comparison to π spreading of activation through links has more impact, resulting in a bias toward current behavior. If they are close to π the system can switch easier between different behaviors gaining adaptability.

3.1 ABC

ABC [5] is a tool designed to define and simulate action selection in autonomous agents based on Behavior Networks. It is implemented in Tcl/Tk 8.0, and thus, it can run in different platforms, such as Windows, MacOS, and UNIX. ABC has a graphical user interface that enables the user to define agent actions, goals, perceptual state, and network parameters. Perceptual states, goals, and parameters can be changed by the user at run time. ABC also offers analysis tools keeping track of activation patterns, building an ethogram, and registering frequencies of the actions executed. A graphic showing activation patterns has been very useful in tuning the behavior of the agent.

4 Monots Project

The Monots Project studies the foraging behavior of howler monkeys (*Allouatta palliata*) generating models that offer a theoretical framework that helps in explaining observational facts. Foraging behavior is the set of actions aimed to searching for, obtaining and ingesting food. Even when howler monkeys have been widely studied in ethology, there are still open questions related to strategies for foraging. We use data from 1500 hours of observation of howler monkeys in Agaltepec Island (Catemaco, Ver., México) and descriptions of foraging strategies produced by the Neuroethology Institute of Universidad Veracruzana [15, 14, 16]. Stochastic models [13] were proposed as a first attempt to sintesize the main characteristics and regularities in these data. Semi-markovian [4] models were also proposed to introduce temporal dynamics.

Behavior Networks were proposed to model foraging behavior [4, 7] because of its properties as a mechanism for action selection. In addition, they have some similarities with models of behavior proposed by ethologists, Fixed-Action Patterns, for example [8]. These similarities might enable us to explain foraging behavior in equivalent terms to those used in ethology.

ABC was used to desing different Behavior Networks to model foraging behavior [6]. The most simple has three actions: resting, eating, and exploring, and three weighted goals or motivations: hunger, lazyness, and curiosity. Weights change over time in function of actions executed and conditions of the environment. The most complex has six actions, adding going-to-food, going-to-shadow, and taking-food, 21 links connecting them, and the same three motivations. Running the networks we observe that activity was divided as follows: resting 57%, eating 28%, and exploring 15%.

Activity in howler monkeys was: resting 60%, eating 25%, and exploring 15%. Other results reproducing statistical parameters associated to foraging will be published soon as a master's thesis.

5 Conclusions

We argue that an agent must be autonomous in order to deal with dynamic, non deterministic environments. That is, an agent must have an extent set of goals and it must be able to choose the goal that must be achieved, at any given time. An agent must, also at any given time, select an action that will contribute to goals archivement in the long term. We present Behavior Networks as an alternative to compute action selection, because they are decentralized mechanisms with dynamically reconfigurable conflict solution. In addtion, Behavior Networks can deal with achievable and homeostatic goals. All these properties make plausible to model agents that exhibit an autonomous behavior using Behavior Networks. Some models of foraging behavior have been proposed in Monots Project to test the viability of our assumptions. We have obtained positive results with simplified models. Future work includes designing more realistic models for foraging, and simulating a group of monkeys to test the influence of social components in foraging behavior.

References

- [1] R. A. Brooks. Intelligence without reason. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pages 569–595, Sydney, Australia, 1991.
- [2] A. A. Covrigaru and R. K. Lindsay. Deterministic autonomous system. *AI Magazine*, pages 110–117, 1991.
- [3] R. E. Fikes and N. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 5(2):189–208, 1971.
- [4] A. García-Vega and et al. Dos modelos de selección de acción ara el estudio de la conducta de forrajeo del mono aullador (allouata paliatta). In *Memorias of 5th Meeting of Internation Biometric Society Network for Central America, The Caribbean, Mexico, Colombia and Venezuela*,

pages 81–92. Instituto Nacional de Estadística, Geografía e Informática, 1997.

- [5] A. Guerra-Hernández. Agentes autónomos y selección de acción, una perspectiva basada en redes de comportamiento. Master's thesis, Maestría en Inteligencia Artificial, Universidad Veracruzana, 1997.
- [6] A. Guerra-Hernández and et al. Abc, una herramienta para implantar modelos basados en redes de comportamiento. In *Memorias Visión Robótica. Primer Encuentro de Computación ENC97*, pages 31–37. Sociedad Mexicana de Ciencias de la Computación, 1997.
- [7] A. Guerra-Hernández and et al. Un modelo de conductas animales basado en redes de comportamiento. In *Memorias of 5th Meeting of Internation Biometric Society Network for Central America, The Caribbean, Mexico, Colombia and Venezuela*, pages 93–104. Instituto Nacional de Estadística, Geografía e Informática, 1997.
- [8] H. Hendriks-Jansen. *Catching Ourselves in the Act, Situated Activity, Interactive Emergence, Evolution, and Human Thought*. MIT Press, 1996.
- [9] L. P. Kaelbling. *Learning in Embedded Systems*. MIT Press, 1993.
- [10] P. Maes. The dynamics of action selection. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, pages 991–997, Detroit, MI, 1989.
- [11] P. Maes, editor. *Designing Autonomous Agents*. MIT Press, 1990.
- [12] P. Maes. Situated agents can have goals. In P. Maes, editor, *Designing Autonomous Agents*, pages 49–70. MIT Press, 1990.
- [13] M. Martínez-Morales. Del modelado estadístico del comportamiento al diseño de agentes autónomos. In *Seminario sobre Vida Artificial*, pages 23–25. Universidad Veracruzana, 1996.
- [14] E. Rodríguez-Luna and et al. Consideraciones en relación a la traslocación de un grupo de monos aulladores (*alouatta palliata*). In *Estudios Primatológicos en México*, pages 81–95. Universidad Veracruzana, 1995.
- [15] E. et al. Rodríguez-Luna. Forrajeo del mono aullador (*alouatta paliatta*) en semilibertad. In *Estudios Primatológicos en México*, pages 133–140. Universidad Veracruzana, 1995.

- [16] J. C. Serio-Silva and et al. Patrón diario de actividades y hábitos alimenticios de *alouatta palliata* en semilibertad. In *Estudios Primatológicos en México*, pages 149–171. Universidad Veracruzana, 1995.
- [17] L. Steels and R. Brooks. *The Artificial Life Route to Artificial Intelligence, Building Embodied, Situated Agents*. Lawrence Erlbaum Associates, Publishers, 1995.