A decision-theoretic approach to fuzzy behavior coordination

Paolo Pirjanian and Maja Matarić
USC Robotics Research Laboratory
University of Southern California
Los Angeles, CA 90089-0781
phone: +1 (213) 740-4523, fax: +1 (213) 740-7512
{paolo, mataric}@usc.edu

This work is funded by the Office of Naval Research Grant N0014-99-1-0162.

Abstract

In this paper, we propose a novel approach to behavior-based control, which combines fuzzy logic with multiple objective decision theory. Fuzzy logic provides powerful tools for formulating sophisticated behavior-based systems for robot control. However, it suffers from problems associated with resolving behavioral conflicts, which are mainly due to the deficiencies of fuzzy inferencing techniques. Hence we propose to formulate the inferencing mechanism using a multiple objective decision-theoretic approach, which provides formal tools for behavior coordination and resolving conflicts in a principled manner. Preliminary experimental results are reported.

1 Introduction

A significant body of work has demonstrated the advantages of fuzzy approaches to behavior-based mobile robot control [10]. First, due to its approximate reasoning capabilities, fuzzy logic produces controllers that are robust to uncertainty (sensory noise, perturbations, etc.). Further, fuzzy behaviors can be conveniently synthesized by a set of if-then rules using easy-to-understand linguistic terms to encode expert knowledge. However, as pointed out by several authors [10], fuzzy behavior coordination suffers from problems associated with resolving behavioral conflicts.

Fuzzy behavior coordination is performed by fuzzy inferencing methods (e.g., max-min, prod-sum) to combine the multivalued outputs of the behaviors and then defuzzification (e.g., centroid) is used to select a final crisp action ultimately used for control (see [10]). This scheme should lead to the selection of an action that somehow represents the consensus among the behaviors and thus comprises the action that best satisfies the decision objectives that they encode. This, however, is not true due to certain shortcomings inherent to standard fuzzy inferencing mechanisms. First, combination of fuzzy outputs by various operators (e.g., triangular conorms) has effects on the final action selection in ways that have not been characterized. Second, defuzzification operators (e.g., centroid) can generate counterintuitive results for action selection. In this paper we propose an alternative approach to fuzzy behavior coordination, which is based on *multiple objective decision theory.*

Multiple objective decision theory can be used to formulate effective mechanisms that can guarantee a suitable trade-off between a multitude of conflicting behavioral objectives. We show how MOBC can be plugged into a fuzzy behavior-based system to replace fuzzy behavior coordination, while retaining the main benefits of approximate reasoning and behavior synthesis using a collection of fuzzy rules.

1.1 Related work

Numerous interesting behavior coordination mechanisms have been proposed over the last decade or two. One particular class, based on fuzzy logic and fuzzy control, is subject of detailed investigation in this work. In order to put this class of fuzzy approaches in perspective, we give a brief overview of other alternatives to behavior coordination (or action selection) mechanisms which can be divided into two main classes: *arbitration* and *command fusion.*

Arbitration mechanisms select one behavior, from a group of competing ones, which is given ultimate control of the system (the robot) until the next selection cycle. This approach is suitable for arbitrating between the set of active behaviors in accord with the system's changing objectives and requirements under varying conditions. The subsumption architecture [2] is a *priority-based* arbitration mechanism, where behaviors with higher priorities are allowed to subsume the output of behaviors with lower priorities. *State-based* arbitration mechanisms include the Discrete Event Systems (DES) formalism [5] and a similar method called temporal sequencing [1]. Finally, in *winner-take-all* mechanisms action selection results from the interaction of a set of distributed behaviors that compete until one wins and takes control of the system [6].

Command fusion mechanisms combine recommendations from multiple behaviors to form a control action that represents their consensus. Thus, this approach provides for a coordination scheme that allows all the behaviors to simultaneously contribute to the control of the system in a cooperative rather than a competitive manner. *Voting* techniques ([9, 8]) interpret the output of each behavior as votes for or against possible actions and the action with the maximum weighted sum of votes is selected. *Superposition* techniques, which combine behavior recommendations using linear combinations, include potential-fields, motor schemas [1], and the dynamic systems approach [11]. *Fuzzy* command fusion mechanisms (see [10]) have several contact points with voting techniques, but notions of fuzzy logic and inferencing are used to formalize the action selection processes. In addition, fuzzy approaches enable a new class of coordination mechanisms denoted *context-dependent blending*, which allows for weighted combination of behaviors and was introduced to robotics by Saffiotti, Ruspini, and Konolige in [10]. Finally, the class of multiple objective approaches, introduced in this paper, provide a formal approach to behavior coordination based on multiple objective decision theory. This approach is in spirit very much related to fuzzy behavior coordination and is based on a similar conceptual frame of thought. However, as will be evident, it has certain advantages over fuzzy approaches.

2 Fuzzy behavior-based control

In fuzzy behavior-based control, each behavior is synthesized by a rule-base controlled by an inference engine to produce a multivalued output. Figure 1, illustrates a simple example describing a collision avoidance behavior using only two rules. One rule recommends (expresses the desire) to turn away from a close obstacle and the other suggests to move straight forward if the obstacle is at a safe distance. Using standard fuzzy inferencing, e.g., max-prod in this example, the rules are combined into a multivalued output that encodes the (grade of) desirability of each action from the behavior's point of view.

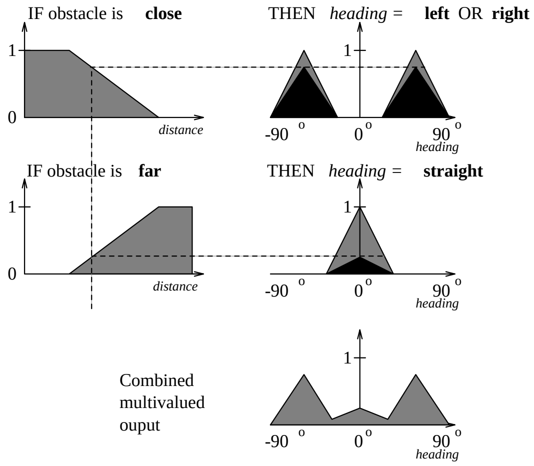

Figure 1: An example fuzzy rule-base to describe a collision avoidance fuzzy behavior.

Behaviors that compete for the control of the robot must be coordinated to resolve potential conflicts. Fuzzy behavior coordination is performed by combining the fuzzy outputs of the behaviors using an appropriate operator such as a triangular conorm (which corresponds to the fuzzy set union), e.g., the max operator. Then defuzzification (e.g., center of gravity, COG) is used to select a final crisp action ultimately used for control (see [10]). This scheme should lead to the selection of an action that somehow represents the consensus among the behaviors and thus comprises the action that best satisfies the decision objectives that they encode. However, in situations where there is conflict between the active behaviors, this approach leads to inappropriate results. For example, in Figure 2, the robot is facing a target on the opposite side of a close-by obstacle, thus there is conflict between the collision avoidance and target following behaviors, which cannot be resolved by fuzzy behavior coordination.

One approach to dealing with conflicts is based on *context-dependent blending* [10], where context dependence is encoded using a number of fuzzy meta rules such as:

```
IF obstacle is close THEN avoid collisions
IF NOT(obstacle is close) THEN follow target
```

If the robot is close to an obstacle, then the contribution from the collision avoidance will be higher than the contribution from target following. This corresponds to weighted decision making, where the weight values are determined by the truth values of the rule antecedents and is used to scale the behaviors' fuzzy outputs accordingly. Another approach, used in the Saphira architecture [4], is to resolve conflicts by allowing higher priority behaviors to suppress/dominate behaviors with lower priorities. These approaches provide partial solutions to resolving conflict between behaviors in specific situations and systems.

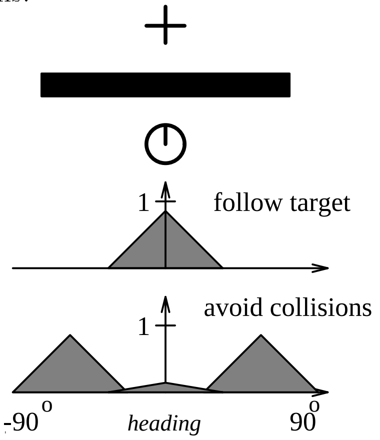

Figure 2: Conflict between behavior outputs.

Defuzzification can also produce inappropriate results since it selects an action that does not have support from any of the behaviors. This is demonstrated in [13] by the example in Figure 3. Yen and Pfluger propose Centroid of Largest Area (CLA) defuzzification to avoid problems with multi-modal outputs. However, it is easy to construct an example where CLA defuzzification can produce undesirable results.

In summary, due to the unclear semantics and unpredictable characteristics of fuzzy behavior coordination, alternative approaches should be considered. One alternative approach, that avoids the common pitfalls of fuzzy behavior coordination, is multiple objective behavior coordination (MOBC) which fits well into the framework of fuzzy behavior coordination and can easily replace standard fuzzy inferencing and defuzzification techniques.

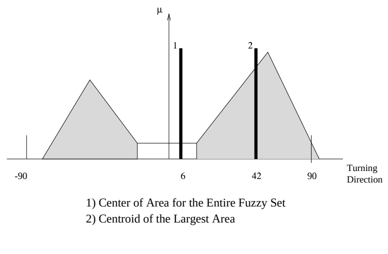

Figure 3: Examples of defuzzification techniques [13].

3 Multiple objective behavior coordination

Multiple objective behavior coordination is based on multiple objective decision theory [3]. Since multiple objective decision theory has many concepts that have synonymous counterparts in fuzzy logic, we only introduce aspects of the theory that support the discussion of MOBC.

3.1 Objectives and behaviors

Objectives define the standards against which the quality of a given alternative may be evaluated. An objective function, o, is thus a mapping that assigns to each alternative, $x \in X$, a value reflecting that objective's desirability. Multiple objective decision making (MODM) provides approaches for making decisions in complex situations where more than one decision objective should be considered.

In multiple objective behavior coordination, a behavior, b, is formalized as an objective function that defines a mapping from an action space, X, to the interval $[0,1]$: $b : X \to [0,1]$. The action space, $X = \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n\}$, is defined to be a finite set of alternative actions. Behaviors in this framework are equivalent to fuzzy behaviors because objective functions encode the semantics of fuzzy membership functions.

3.2 Multi-objective decision problems and behavior coordination

Once the link between behaviors and objective functions is established, we can formulate the coordination of multiple behaviors as a multiple objective decision problem that seeks to *optimize a number of objective functions simultaneously*, where the objective functions correspond to behavior multivalued outputs. From the set of alternative actions X, an action $\hat{\mathbf{x}} \in X$ should be chosen that best satisfies all objective functions. This is formulated as a multiple objective decision problem in the following way:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{argmax} \qquad [o_1(\mathbf{x}), o_2(\mathbf{x}), ..., o_n(\mathbf{x})] \qquad (1)$$
$$\text{subject to} \qquad \mathbf{x} \in X,$$

where $o_i(\mathbf{x})$ for $i = 1, ..., n$ are the objective functions estimated by corresponding behaviors b_i and $\mathbf{x}$ is an N-dimensional control vector $(x_1, x_2, ..., x_N)$.

Further $X \subseteq R^N$ is the constrained action space[1]. The solution $\hat{\mathbf{x}}$ to equation (1) corresponds to the selected action that best satisfies the corresponding behaviors.

3.3 Solving multiple objective decision problems

There are a number of methods described in the literature for solving multi-objective decision problems each with certain advantages and disadvantages [7, 3]. Rather than merely giving a description of these methods, we highlight the main characteristics of multiple objective decision making that make it suitable for behavior coordination. Thus, we start by describing the types of solutions to a MODM problem.

A solution is *optimal* if it results in the maximum value of each of the objective functions simultaneously. That is, $\hat{\mathbf{x}}$ is an optimal solution to equation (1) iff $\mathbf{o}(\hat{\mathbf{x}}) \geq \mathbf{o}(\mathbf{x})$ for all $\mathbf{x} \in X$, where $\mathbf{o}(\mathbf{x}) = [o_1(\mathbf{x}), o_2(\mathbf{x}), ..., o_n(\mathbf{x})]$.

An important characteristic of MODM problems is that the multiple objectives, $o_1, ..., o_n$, can be competitive and conflicting, i.e., the improvement in one can be associated with a deterioration in another. Hence optimal solutions might not exist and thus the concept of optimality should be modified to endow similar and useful concepts in MODM problems. In practice we can only hope for solutions that *satisfy* each of the objectives to a *sufficient* degree. This idea is encompassed in the following definition.

Definition 3.1 (A Satisficing Solution) *A satisficing solution of Simon[12] extended to a multiple objective decision problem is a reduced subset of the feasible set, which exceeds all of the aspiration levels (goals) of each objective.*

Simon points out that it might be impractical to search for an optimal solution and argues that in practice it is more efficient to look for solutions that are 'good enough'. A satisficing solution corresponds to the behavior of a decision maker who looks for 'good enough' solutions. Another useful notion of optimality in the context of multiple objective decision making is *Pareto-optimality.*

Definition 3.2 (A Pareto-Optimal Solution) $\mathbf{x}^*$ *is a Pareto-optimal solution to a multiple objective decision problem iff there does not exist any* $\mathbf{x} \in X$ *such that* $o_i(\mathbf{x}^*) \leq o_i(\mathbf{x})$ *for all* i *and* $o_j(\mathbf{x}^*) < o_j(\mathbf{x})$ *for at least one* j.

Intuitively, a Pareto-optimal solution is one in which no other solution can improve one objective function without a simultaneous deterioration of at least one of the other objectives. A Pareto-optimal solution is also known as a nondominated solution, a non-inferior solution, or an efficient solution.

3.4 Multi-objective behavior coordination

It is not so much the methods for generating solutions but the notions of Pareto-optimal and satisficing that play a significant role in formulating behavior coordination mechanisms with useful properties. See [7] for techniques for generating PO and satisficing solutions for multi-objective decision problems. Based on the definitions of Pareto-optimality and satisficing, we propose a behavior coordination scheme that proceeds in the following step-wise fashion:

1. **Feasible actions**. Based on physical and other hard constraints find the set of feasible and permissable alternatives

2. **Pareto-optimal actions**. Based on objective criteria find the set of Pareto-optimal actions and reject others to reduce significantly the set of alternatives

3. **Satisficing actions**. Incorporate subjective knowledge (weights, priorities, goals, aspirations levels) to find the set of satisficing alternatives

4. **Most preferred action**. Use other criteria to finally select the most preferred action.

This approach systematically narrows down the set of possible solutions, by including increasingly more decision criteria. First, the set of hard constraints given by the nature of the task and physical limitations of the robot (e.g., limited acceleration) are used to exclude the set of non-feasible actions. Then, based only on the objective criteria, we can exclude the set of inferior or dominated solutions simply because they are not efficient when compared to the Pareto-optimal solutions. By incorporating some subjective preference structure in terms of weights, priorities etc. we can then select, out of the Pareto-optimal solutions, the set of satisficing, i.e., sufficiently satisfying, solutions. Finally, other criteria or

[1] The action space is constrained due to various physical and geometrical properties, such as a set of non-holonomic constraints imposed by the mechanical construction of the robot, into a subset of *feasible* (e.g., physically possible) actions.

further subjective knowledge can be included to make a final selection. In this way alternatives are evaluated against the behavioral objectives and ranked in terms of permissable, Pareto-optimal, satisficing and preferred actions.

4 A multi-objective approach to fuzzy behavior coordination

Rather than letting the behavior coordination depend on the operators provided by the underlying fuzzy mechanism, multiple objective behavior coordination is first and foremost concerned with the quality of the selected actions (in terms of Pareto-optimality and Satisficing). However, as it should be evident by now, the (input/output) mechanics of multiple objective behavior coordination are not different from the fuzzy approaches, hence they can be interchanged.

Further, several of the approaches for generating solutions (Pareto-optimal and satisficing) to multiple objective decision problems (see [7]) use weight values . The weight values have the same role as in context-dependent blending, and determine the degree of contribution for each behavior. Thus, also with respect to behavior blending the fuzzy and multiple objective approaches are compatible and thus can be combined. One straight-forward way to combine both approaches is to:

1. Specify each behavior using a fuzzy rule-base, and generate multivalued behavioral outputs using fuzzy inferencing,
2. Use fuzzy context rules to generate weight values, and
3. Coordinate the behaviors by generating Pareto-optimal and satisficing solutions using multiple objective decision making techniques.

In this way we combine the advantages of approximate reasoning and ease of behavior specification of fuzzy logic with the qualitative and formal behavior coordination capabilities of multiple objective decision making.

5 Experimental results

We have conducted some preliminary simulation experiments which clearly indicate that the proposed approach provides better conflict resolution when compared to fuzzy behavior coordination (using COG defuzzification and context-dependent blending[2]). Multiple-objective behavior coordination (MOBC) makes suitable trade-offs between behavioral objectives and is able to find solutions that do not severely sacrifice any behavior's objective.

The experiments are carried out on a Nomad200 simulator using a ring of 16 ultrasonic sensors for obstacle detection, and encoders for calculation of robot position. In the example, we consider the navigation task for a mobile robot placed at some arbitrary location and commanded to go to a goal position. This task involves three behaviors: 1) obstacle avoidance, 2) maintain target heading, and 3) move fast forward. Table 1 summarizes the results for 20 experiments for each of the mechanisms (MOBC and FBC). The metrics used for performance characterization of the MOBC vs. fuzzy command fusion mechanisms are: 1) safety, which is the average distance to obstacles; 2) velocity, which is the average velocity of the robot during the runs; 3) number of successful runs, which is a global metric that measures the number of trials where the robot safely reaches the goal position; and 4) combined deviation from ideals, which is a metric indicating the degree of compromise made between the behaviors.

Combined deviation A powerful characteristic of command fusion mechanisms is that they select actions that benefit not only one but the collection of behaviors that are involved in the action selection process. An attribute that is indicative of the degree of compromise is the combined deviations from ideals. An ideal $\mathbf{o}^*$ is defined as $\mathbf{o}^* = (o_1^*, ..., o_n^*)$, where $o_j^* = max_{\mathbf{x} \in X}\ o_j(\mathbf{x})$. $o_1, ..., o_n$ are the objective functions. Thus, for a selected action $\mathbf{x}'$, the combined deviation for a set of objectives $o_1, ..., o_n$ is given by: $d = \sum_{j=1}^{n} |o_j(\mathbf{x}') - o_j^*|$. The value for d will in general be in the range $[0; n]$ and in the range $[0; 3]$ for three behaviors, with 0 being the best (i.e., no deviation from ideals) and 3 the worst case (i.e., maximum deviation from ideals).

The results show that MOBC is superior to fuzzy behavior coordination (FBC) based on the metrics of compromise and the number of successful runs. MOBC produces twice as many successful trials (12 vs. 6) and nearly twice as good a compromise (0.268 vs. 0.492). Further, looking at the velocity metric we see that MOBC produces an average velocity that is nearly twice as fast as for FBC, yet it has a higher success rate because it makes a suitable trade-off be-

[2]The distance to the closest obstacle is used to determine the weight values for blending the behaviors. Context-dependent blending is also used in the multiple objective behavior coordination mechanism.

	MOBC	Fuzzy	Ideal value
Safety (inch)	39.3	46.1	∞
Velocity, v (inch/s)	19.2	10.6	**20**
# Successful runs	12	6	**20**
Smoothness (deg)	11.6	1.2	**0**
Compromise: Deviation from ideals			
Avoid obstacle	0.147	0.316	**0**
Maintain heading	0.117	0.175	**0**
Move fast	0.004	0.000	**0**
Combined dev.	0.268	0.492	**0**

Table 1: Experimental results for 20 experiments for each of the behavior coordination mechanisms and the ideal values as a point of reference.

tween system objectives. However, the smoothness metric indicates that FBC produces a smoother path, which is due to COG defuzzification. Figure 4 depicts situations where center of gravity defuzzification fails to resolve conflicts between behaviors (mainly between obstacle avoidance and maintain target heading), whereas MOBC provides a suitable trade-off between behavioral objective to resolve conflicts.

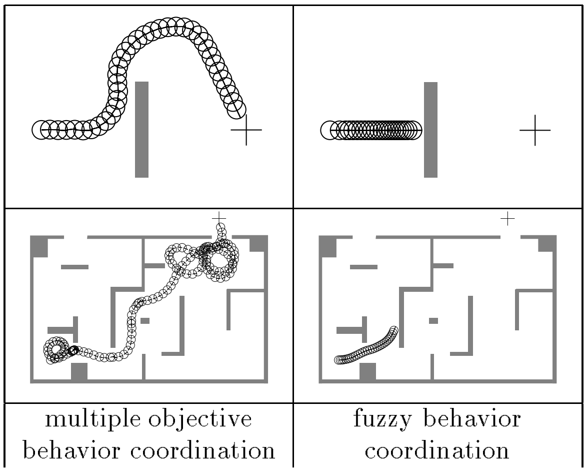

Figure 4: Plot of trails produced from example trials.

6 Concluding remarks

It is commonly known and acknowledged that fuzzy approaches can produce undesirable action selection and cannot satisfactorily deal with behavioral conflicts. Multiple objective behavior coordination provides useful formal tools to formulate characterizable coordination schemes. Thus we proposed a combined approach that maintains the advantages of approximate reasoning and ease of behavior specification of fuzzy logic with the qualitative and formal behavior coordination capabilities of multiple objective decision making. The experimental results support the fact the multiple objective behavior coordination provides better means of resolving conflicts and produces more suitable trade-off between behavioral objectives than fuzzy behavior coordination.

References

[1] Ronald C. Arkin. Towards the Unification of Navigational Planning and Reactive Control. *AAAI Spring Symposium on Robot Navigation*, March 1989.

[2] Rodney A. Brooks. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, March 1986.

[3] Vira Chankong and Yacov Y. Haimes. *Multiobjective Decision Making - Theory and Methodology*, volume 8. North-Holland, 1983.

[4] Kurt Konolige. *Saphira Software Manual.* ActivMedia Inc., version 6.1 edition, 1998.

[5] Jana Košecká and Ruzena Bajcsy. Discrete Event Systems for Autonomous Mobile Agents. . *Proceedings Intelligent Robotic Systems '93 Zakopane*, pages 21–31, July 1993.

[6] Pattie Maes. How To Do the Right Thing. Technical Report NE 43 - 836, AI-Laboratory, Massachusetts Institute of Technology, 545 Technology Square, Cambridge, MA 02139, USA, 1989.

[7] Paolo Pirjanian. *Multiple Objective Action Selection & Behavior Fusion using Voting.* PhD thesis, Department of Medical Informatics and Image Analysis, Institute of Electronic Systems, Aalborg University, Fredrik Bajers Vej 7, DK-9220 Aalborg, Denmark, August 1998. *Available online: http://www-robotics.usc.edu/~paolo/publications.*

[8] Jukka Riekki and Juha Röning. Reactive Task Execution by Combining Action Maps. In *International Conference on Integrated Robots and Systems (IROS)*, pages 224 – 230, Grenoble, France, September 1997.

[9] Julio K. Rosenblatt. DAMN: A Distributed Architecture for Mobile Navigation. In *AAAI Spring Symposium on Lessons Learned from Implemented Software Architectures for Physical Agents, Stanford, CA.* AAAI Press, Menlo Park, CA., March 1995.

[10] Alessandro Saffiotti. The Uses of Fuzzy Logic in Autonomous Robot Navigation: a catalogue raisonné. Technical Report 2.1, IRIDIA, Universite Libré de Bruxelles, 50 av. F. Rossevelt, CP 194/6, B-1050 Brussels, Belgium, November 1997.

[11] Gregor Schöner and M. Dose. A dynamics systems approach to task level systems integration used to plan and control autonomous vehicle motion. *Robotics and Autonomous Systems*, 10:253–267, October 1992.

[12] Herbert A. Simon. *The New Science of Management Decision.* Harper and Row, New York, 1960.

[13] John Yen and Nathan Pfluger. A Fuzzy Logic Based Extension to Payton and Rosenblatt's Command Fusion Method for Mobile Robot Navigation. *IEEE Transactions on Systems, Man, and Cybernetics*, 25(6):971 – 978, June 1995.