Multiple Objective vs. 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}@cs.usc.edu

1 Introduction

In the *behavior-based* paradigm [4, 16], the control of a robot is shared between a set of purposive processes, called *behaviors*. Based on selective sensory information, each behavior produces/proposes reactions to control the robot with respect to a particular *objective*, i.e., a narrow aspect of the robot's overall task. Behaviors with different and possibly incommensurable objectives may produce conflicting actions that are seemingly irreconcilable. Thus, the design of effective *behavior coordination mechanisms* is a major research issue.

We focus in particular on fuzzy approaches to behavior coordination and argue that conflict resolution using fuzzy inferencing and defuzzification can generate undesirable results. As an alternative we present *multiple objective behavior coordination* (MOBC), which provides useful formal tools for describing, characterizing, and solving behavioral conflicts in a principled manner. This approach is based on multiple objective decision making (MODM) theory [6], which fits nicely into the behavior-based paradigm because it provides approaches for making decision in situations where multiple, potentially conflicting objectives should be considered. We show that MOBC can be used to formulate effective mechanisms that can guarantee a suitable trade-off between a multitude of conflicting behavioral objectives. In addition, we show that multiple objective behavior coordination can be plugged into a fuzzy behavior-based system to replace fuzzy inferencing and defuzzification.

In order to put the fuzzy and multiple objective approaches into a larger taxonomic perspective, in Section 2 we give an overview of classes of behavior coordination mechanisms. Section 3 describes the main pitfalls of fuzzy behavior coordination and motivates multiple objective behavior coordination as an alternative approach. In Section 4, we give an introduction to multiple objective decision making and establish its links to behavior-based control and describe how it can be applied to behavior coordination. The introduction is following by Section 5, where we describe the multiple objective behavior coordination approach in some detail. Rather than providing rigorous empirical results, in this paper we concentrate on the formal and theoretical results. Thus, in Section

6 we provide a mobile robot navigation example task for the purpose of characterizing the behavior coordination mechanisms that we can formulate using the presented multiple objective approach. We conclude with a discussion of results and directions for future work.

2 Overview of behavior coordination mechanisms

In behavior-based control, a complex control problem is divided into a set of simpler controllers, known as behaviors, that collectively solve the original problem. To do this, it is thus necessary to coordinate the activities of the behaviors so to satisfy the initial complex system's control objectives. This is known as the behavior coordination problem, also refereed to as the action selection or behavior fusion problem.

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. For a detailed overview, discussion, and comparison of behavior coordination mechanisms see [23]. Behavior coordination mechanisms can be divided into two main classes: *arbitration* and *command fusion* (Figure 1).

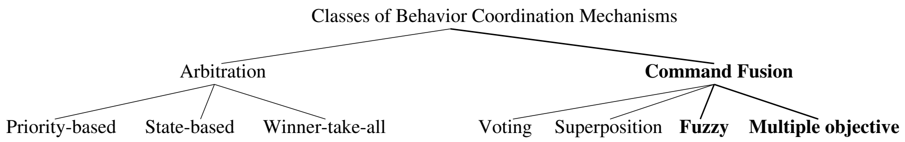

Fig. 1. Classes of behavior coordination mechanisms. Bold indicates those treated in this paper.

Arbitration mechanisms select one behavior, from a group of competing ones, and give it 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. It can focus the use of scarce system resources (sensory, computational, etc.) on tasks that are considered to be relevant. The Subsumption architecture [5] is a *priority-based* arbitration mechanism, where behaviors with higher priorities are allowed to suppress the output of behaviors with lower priorities. *State-based* arbitration mechanisms include the Discrete Event Systems (DES) formalism [12], a similar method called temporal sequencing [3], and action selection based on Bayesian decision theory [13]. 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, such as in the activation network approach [15].

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 behaviors to simultaneously contribute to the control of the system in a cooperative rather than a competitive manner. *Voting* techniques 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. Voting command fusion mechanisms include the DAMN architecture [25], a similar approach called SAMBA [24], and the Action Voting approach [8]. *Superposition* techniques, which combine behavior recommendations using linear combinations, include potential-fields [10], motor schemas [2], and the dynamic systems approach [14]. *Fuzzy* command fusion mechanisms (see [28, 31, 30]) 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*, introduced to robotics by Saffiotti, Ruspini, and Konolige in [28], which allow for weighted combination of behaviors. Finally, the class of multiple objective approaches, first introduced by Pirjanian [20], provides a formal approach to behavior coordination based on multiple objective decision theory. This approach, described in the paper, is in spirit very much related to fuzzy behavior coordination and is based on a similar conceptual framework. However, as will be evident, its formalism affords certain advantages over fuzzy approaches.

3 Pitfalls of fuzzy behavior coordination

A significant body of work has demonstrated the advantages of fuzzy approaches to behavior-based mobile robot control [27]. Fuzzy methods describe a behavior (or its output) as a multivalued fuzzy membership function that reflects the (grade of) desirability of each action from that behavior's point of view. Fuzzy behaviors can conveniently be synthesized by a set of easy-to-understand fuzzy if-then rules together with an inference engine (a fuzzy reasoning mechanism). Figure 2 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.

Payton and Rosenblatt [17, 25] demonstrated that multivalued behaviors define a more suitable communication interface because they reduce the loss of information (compared to single-valued behaviors) and hence define a suitable basis for finding consensus and compromize solutions. Furthermore, Payton and Rosenblatt [18] demonstrate that such multivalued representations are suitable for integration of planning and reactivity in concordance with the plan-as-resource view (also known as plan-as-communication) [1]. This view suggests that plans should be used to guide, not control, lower-level behaviors. Even though Payton and Rosenblatt's work was not done in a fuzzy framework, these results are relevant because the behavior semantics are equivalent to fuzzy behaviors. Furthermore, their work was extend into a fuzzy logic approach by Yen and Pfluger [31].

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 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 [27]). This scheme should lead to the selection of an action that somehow represents the consensus among the behaviors and thus best satisfies the decision objectives that they encode. This, however, is not always the case, due to certain shortcomings inherent in standard fuzzy inferencing mechanisms. For example, in Figure 3, the robot is facing a target on the opposite side of a near-by obstacle, thus there is a 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* [27], 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 are used to scale the behaviors' fuzzy outputs accordingly. Another approach, used in the Saphira architecture [11], 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.

While fuzzy behaviors have clear semantics in terms of the desirability over a set of alternative actions, this is not the case for fuzzy behavior coordination due to two main reasons. 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, fuzzy defuzzification operators can generate counterintuitive results for action selection.

Yen and Pfluger [31] use the example illustrated in Figure 4 to demonstrate that center of gravity/area (COG) defuzzification can produce inappropriate results since it selects an action that does not have support from any of the behaviors. Saffiotti [28] uses COG defuzzification but warns about dealing with fuzzy outputs that are not unimodal. Yen and Pfluger propose Centroid of Largest Area (CLA) defuzzification to avoid problems with multi-modal outputs. However, as shown in Figure 5, it is easy to construct an example where CLA defuzzification can produce undesirable results. In this example, the action

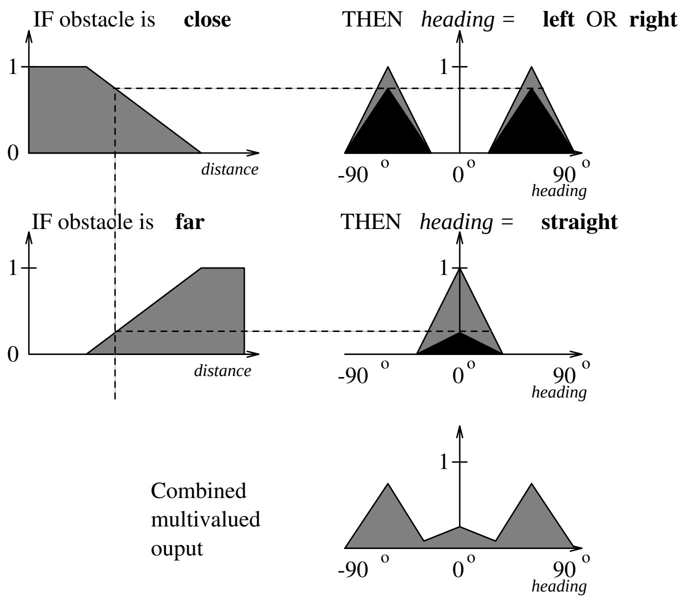

Fig. 2. An example fuzzy rule-base to describe a collision avoidance fuzzy behavior.

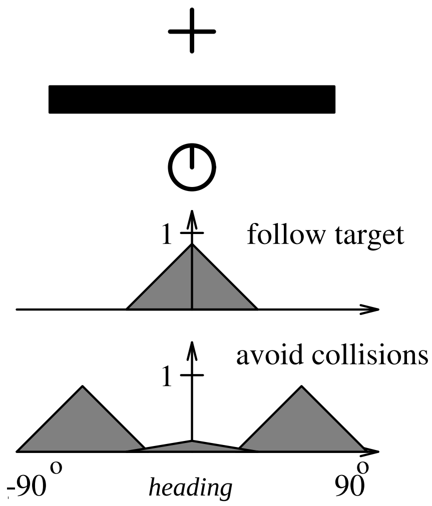

Fig. 3. Conflict between behavior outputs.

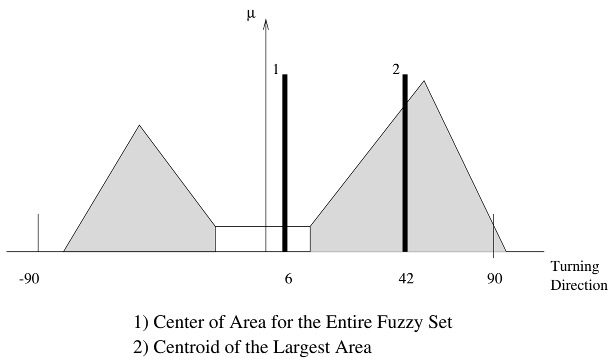

Fig. 4. Examples of two defuzzification techniques [31].

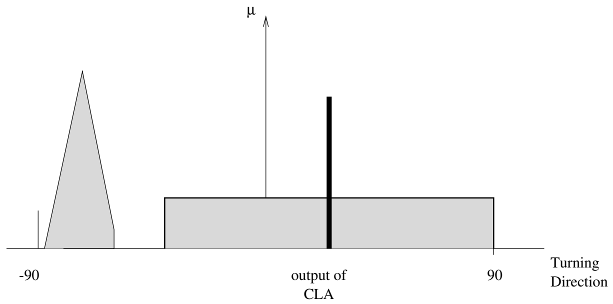

Fig. 5. Examples of the CLA defuzzification technique, where the produced output is not appropriate.

selected by CLA belongs to the large mode on the right with less desirability than the mode on the left. Another approach, used by Pin et al. [19] and by the Saphira architecture [11], is to resolve conflicts by allowing higher priority behaviors to suppress/dominate behaviors with lower priorities. This provides partial solutions to resolving conflict between behaviors in specific situations and systems.

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). MOBC fits well into the framework of fuzzy behavior coordination and can easily replace standard fuzzy inferencing and defuzzification techniques.

4 Multiple objective decision theory applied to behavior coordination

Multiple objective decision theory provides approaches to making decision in complex situations where more than one decision objective should be considered [6]. In this section, we establish the link between multiple objective decision making and behavior-based control. We demonstrate that multiple objective decision theory affords concepts that can be used to describe behavior-based systems and provides techniques that can be used to formulate a class of behavior coordination mechanisms with characterizable properties. We further extend the link from concepts in multiple objective decision theory to their synonymous counterparts in fuzzy logic.

Based on the findings of this section we propose mechanisms for multiple objective behavior coordination in Section 5.

4.1 Objectives and behaviors

Objectives are statements of wants and desires. Thus, 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 alternative's desirability, $o : X \rightarrow R$. Behaviors in this framework are equivalent to fuzzy behaviors (see [28]) because objective functions encode the semantics of fuzzy membership functions and correspond to the notion of desirability. Ruspini [26] establishes the unification of desirability from multivalued logic and preferences from utility theory. Ruspini's work, however, does not address the limitations of fuzzy inferencing for decision making. Multiple objective decision making does. It 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 $[0, 1]$:

$$b : X \rightarrow [0, 1]. \tag{1}$$

The action space, $X = \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n\}$, is defined to be a finite set of alternative actions. The mapping assigns to each action $\mathbf{x} \in X$ a desirability $b(x)$ value, where 1 corresponds to the most desirable one. Note that it is not strictly necessary to confine the range of a behavior to the unit interval.

4.2 Multi-objective decision problem 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:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{argmax} \qquad [o_1(\mathbf{x}), o_2(\mathbf{x}), ..., o_n(\mathbf{x})] \tag{2}$$
$$\text{subject to} \qquad \mathbf{x} \in X,$$

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. 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.

For solving the optimization problem in Equation (2) we do not define a linear order on the set of vector values $v(x) = [o_1(x), ..., o_n(x)]$ so as to find the largest vector and hence optimal solution $\hat{x}$ where $v(x) > v(y), \ \ \forall \ y \in X$. Even if we could specify the most suitable ordering, it could be extremely inefficient for large action spaces ($|X|$). Rather, the philosophy for solving the problem efficiently is based on the notions of *Pareto-optimal* and *satisficing* solutions (explained in the following section).

The solution $\hat{\mathbf{x}}$ to Equation (2) corresponds to the selected action that best satisfies the corresponding behaviors.

4.3 Methodologies for 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 [20, 21, 6]. Rather than merely giving a description of these methods, we highlight the main characteristics of multiple objective decision making (MODM) that make it suitable for behavior coordination. Thus, we start by describing the types of solutions to a multiple objective decision 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 (2) 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 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 [9]. 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 1 (A Satisficing Solution) *A satisficing solution of Simon [29] 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 [29] 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 such 'good enough' solutions. Another useful notion of optimality in the context of multiple objective decision making is *Pareto-optimality.*

Definition 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 non-dominated solution, a non-inferior solution, or an efficient solution.

In the following we outline several methods for generating Pareto-optimal and satisficing solutions.

4.4 Generating Pareto-optimal solutions

Numerous approaches exist to generating Pareto-optimal solutions, however, in the context of behavior coordination we are interested in those that have practical applications. The *weighting* approach is an effective method based on scalar vector optimization and formulated in the following way:

$$\hat{\mathbf{x}} = \underset{x \in X}{argmax} \qquad \sum_{i=1}^{n} w_i o_i(\mathbf{x}) = \mathbf{w}^T \mathbf{o}(\mathbf{x}), \tag{3}$$

where w_i are weights normalized so that $\sum_{i=1}^{n} w_i = 1$. We can claim that $\hat{\mathbf{x}}$ is a Pareto-optimal solution if $\hat{\mathbf{x}}$ is a unique solution of (3) or if $w_i > 0$ for all i. See [6] (theorem 4.6) for details.

4.5 Generating satisficing solutions

In order to go beyond generating Pareto-optimal solutions, some techniques incorporate the decision maker's (DM's) subjective preference structure to identify satisficing alternatives among those that are Pareto-optimal. In the approaches presented here, the DM's preference structure is usually expressed in many forms, including: *weights*, *priorities*, *goals*, and *ideals*. We only describe a few approaches here.

Lexicographic method The lexicographic method requires ranking the importance of all objectives. Assume that the objectives are ranked in decreasing order of importance, $o_1, o_2, ..., o_n$. Then a sequential elimination process is initiated by solving the following sequence of problems until either a unique solution is found or all the problems are solved:

$$\begin{aligned}
&P_1 : \max_{\mathbf{x}\in X} \quad o_1(\mathbf{x}),\\
&P_2 : \max_{\mathbf{x}\in X_1} \quad o_2(\mathbf{x}),\\
&\ldots\\
&P_i : \max_{\mathbf{x}\in X_{i-1}} \quad o_i(\mathbf{x}),\ X_{i-1} = \{\mathbf{x}|\mathbf{x} \text{ solves } P_{i-1}\}, i = 2, ..., n+1,\\
&\ldots
\end{aligned} \tag{4}$$

The basic idea is that the first objective is used to screen the solutions of the second objective, and so on. In some sense, this approach of behavior coordination could be viewed as an extension of Subsumption-style arbitration [5] to command-fusion, where higher-level behaviors inhibit the outputs of lower-levels, in case they disagree (i.e., their recommended sets of actions do not intersect). In cases of agreement (i.e., when their recommended sets of actions do intersect) the selected action represents the consensus of the behaviors.

Goal programming In *goal programming methods* the DM gives his/her preferences in terms of weights, priorities, goals, and ideals. The concept of *best alternative* is then defined in terms of how much the actual achievement of each objective deviates from the desired goals or ideals. Goal programming methods thus *"choose an alternative with the minimum combined deviation from the goal $\hat{o}_1, ..., \hat{o}_n$, given the weights or priorities of the objective functions"*, [6, page 302]. The general formulation of this multi-objective decision problem is known as the surrogate multi-objective decision problem.

$$\begin{aligned}
&\min_{\mathbf{x}\in X} \left\{ d_p(\mathbf{o}(\mathbf{x}), \hat{\mathbf{o}}) \triangleq \left(\textstyle\sum_{j=1}^{n} w_j \left| o_j(\mathbf{x}) - \hat{o}_j \right|^p \right)^{1/p} \right\}\\
&\text{subject to} \quad \mathbf{x} \in X,
\end{aligned} \tag{5}$$

where $1 \le p \le \infty$, $\hat{\mathbf{o}}$ is the goal vector, w_j is the weight or priority given to the jth objective, and $d_p(\mathbf{o}(\mathbf{x}), \hat{\mathbf{o}})$ represents the 'distance' between $\mathbf{o}(\mathbf{x})$ and $\hat{\mathbf{o}}$. Once the goal, $\hat{\mathbf{o}}$, and weights or priorities, $\mathbf{w} \triangleq (w_1, ..., w_n)$, are given, and (5) is formulated for any appropriately chosen value of p, then solutions of the equation determine compromize solutions. The solutions of Equation (5) may not be nondominated. Only when $\hat{\mathbf{o}} = (o_1^*, ..., o_n^*)$, where $o_j^* = max_{\mathbf{x} \in X} f_j(\mathbf{x})$ and either the solution is unique or $w_j > 0$ for each $j = 1, ..., n$ is a solution of Equation (5) nondominated.

Behavior coordination using this approach corresponds to selecting actions that deviate the least from the desires (or recommendations) of the behaviors. This approach seeks to satisfy the behaviors as much as possible (i.e., it seeks to *satisfice*).

Interval criterion weights Recall that in the weighting approach the set of Pareto-optimal solutions, X^*, can be generated by varying the weight values, w_i, to assume values in the interval $(0; 1]$ and then solving Equation (3). Rather than letting each w_i assume values in the interval $(0; 1]$, a narrower interval can be used to reflect the importance of each objective function. This approach is known as *interval criterion weights method* and is used to generate a reduced set of X^*. Apart from reducing the computational burden, this method produces the 'best' part of X^*. Here 'best' reflects the DM's preferences, expressed as weight values.

5 Multiple objective behavior coordination

Established the link between a multiple objective decision problem and the action selection (or behavior coordination) problem is the key to formulating the class of multiple objective behavior coordination mechanisms. Objectives correspond to behaviors and solving the multiple objective decision problem in Equation (2) corresponds to action selection. Based on these findings, we propose to formulate the behavior coordination problem as a multiple objective decision making problem. The proposed methodology proceeds in the following steps:

1. **Design of behavior repertoire**: Given a system specification and task requirements, formulate a set of specific objectives, $o_i, ..., o_n$ and associated measurable attributes that are indicative of the degree of achievement of the corresponding objectives. Each objective is implemented by a behavior and hence the objectives define the behavior repertoire necessary for the given system (task).
2. **Design of behaviors**: Each objective function o_i is modeled as a behavior b_i that assesses the attribute values of that objective. Thus behavior b_i is responsible for maintaining the ith objective, o_i. The assessment of the attribute values by a behavior can be based on any appropriate method.
3. **Action selection**: From the set of alternative actions, X, we select actions that are Pareto-optimal **and** satisficing solutions to the vector optimization

problem in Equation (2), where the objective functions $o_1, ..., o_n$ are the ones identified in step 1 and implemented by behaviors $b_1, ..., b_n$ in step 2.

From a system synthesis point of view, this approach specifies a divide-and-conquer methodology (in the same spirit as the behavior-based approach), where a complex system is divided into a set of simple, comprehensible units which are the 'reassembled' to implement the system's overall task. 'Reassembling' is done using techniques for multiple objective decision making.

In behavior-based systems, behaviors constitute the building blocks of a complex control system and thus it is crucial that the semantics of behaviors be well defined. In this approach, each behavior implements a given system objective. Thus, the objective defines the semantics of a behavior and is directly related to the overall system task. Further, action selection corresponds to choosing the best alternative according to the values of the objective functions $o_1, ..., o_n$.

In this approach, it is not so much the methods for generating solutions but the notions of *Pareto-optimality* and *satisficing* that play a significant role in formulating MOBC mechanisms. These allow MOBC to formalize 'good enough' performance and produce actions that represent the consensus among the behaviors. Based on these definitions (see Section 4), we propose action selection proceed in the following step-wise fashion:

1. **Feasible actions**. Based on physical and other hard constraints find the set of feasible and permissible alternatives;
2. **Pareto-optimal actions**. Based on objective criteria, use Equation (3) to find the set of Pareto-optimal actions and reject others to significantly reduce the set of alternatives;
3. **Satisficing actions**. Incorporate subjective knowledge (weights, priorities, goals, aspirations levels) to find the set of satisficing alternatives, using any method in Equations (4), (5) etc;
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 subjective knowledge. First, in step 1, 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, the set of inferior or dominated solutions is excluded, in step 2, 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., the set of satisficing (i.e., sufficiently satisfying) solutions is selected in step 3. Finally, in step 4, other criteria or further subjective knowledge can be included to make an ultimate selection. In this way alternatives are evaluated against the behavioral objectives and ranked in terms of permissible, Pareto-optimal, satisficing and preferred actions.

6 Mobile robot navigation example

In this section we demonstrate the principles of the proposed approach by considering a fairly simple example of mobile robot navigation consisting of three objectives/behaviors: 1) avoid obstacles; 2) maintain (target) heading; and 3) move fast forward. In order to simplify the implementation of the behaviors, we assume that the trajectory of the robot can be approximated with piecewise-circular arcs [7] defined as the ratio r of the translational velocity v to the angular velocity ω, i.e., $r = v/\omega$.

Move fast forward This is the simplest of the behaviors and has the objective to make the robot move forward quickly. The following is the objective function of this behavior:

$$b_{mff}(v, w) = a\frac{v}{v_{max}} + b, \tag{6}$$

where a and b in our case are chosen to be 1 and 0, respectively. The objective function is further normalized and its values are in the interval $[0, 1]$.

Maintain target heading This behavior has the objective to face towards a given target:

$$b_{mh}(v, w) = \frac{1}{1 + \left(\frac{\theta_0(v,\omega) - \theta_{target}}{\beta}\right)^2}, \tag{7}$$

where θ_0 is the predicted robot heading for (v, ω) and θ_{target} is the heading to the target.

Avoid obstacles The objective of this behavior is to keep the robot at a safe distance from obstacles. Range readings, provided by the sonar ring, are used to detect obstacles. For each path determined by the pair of control parameters (v, ω), the obstacle avoidance behavior calculates the shortest distance to the obstacles on its path. Formally the objective function is given by:

$$b_{ao}(v, \omega) = \min\{Far(d) | d = dist(v, \omega, o)\ \forall\ o \in Obs\}, \tag{8}$$

where $Far : R \rightarrow [0, 1]$ is an S-function that maps the distances into the interval $[0, 1]$. Obs is the set of detected obstacle positions $o = (x, y)$ and

$$dist(v, \omega, o) = \begin{cases} \frac{v}{\omega}\gamma & \text{if } \omega \neq 0 \text{ and } (x - x_0)^2 + (y - y_0)^2 = (\frac{v}{\omega})^2; \\ \sqrt{x_o^2 + y_o^2} & \text{if } \omega = 0 \text{ and } x = 0; \\ \infty & \text{otherwise,} \end{cases} \tag{9}$$

$c = (x_0, y_0)$ is the center of the circular path and γ is the angle between the line from c to the robot's position and the line from c to o.

6.1 Behavior coordination

Behavior coordination for the set of behaviors designed in the previous section can be formulated as the following multiple objective decision making problem:

$$\begin{aligned} &\max_{\mathbf{x}} \ [o_{ao}(\mathbf{x}), o_{mh}(\mathbf{x}), o_{mff}(\mathbf{x})] \\ &\text{subject to} \quad \mathbf{x} \in X, \end{aligned} \tag{10}$$

where o_{ao}, o_{mh} and o_{mff} are objective functions estimated by behaviors b_{ao}, b_{mh} and b_{mff} respectively. The equation is subject to constraints (represented by the set X) due to the holonomic constraints of the robot:

$$X = \{(v, \omega) | \ v \in [v_{min}, v_{max}] \ \wedge \ \omega \in [-\omega_{max}, \omega_{max}]\}, \tag{11}$$

where v_{min}, v_{max} and $-\omega_{max}, \omega_{max}$ are the limitations on the translation and angular velocities, respectively.

Action selection consists of solving Equation (10) using the approach described in Section 5.

6.2 Implemented examples

As a proof of concept, we show some simulated and real-world navigation examples produced by the set of behaviors described and specified in the previous section. We used several implementations based on different techniques for generating Pareto-optimal and satisficing solutions. For detailed experimental description see [20, 22]. The following set of parameters were used:

Priorities	:	b_{ao}, b_{mh}, b_{mff}
Goals	: $(g_{ao}, g_{mh}, g_{mff})$	$= (1.0, 0.3, 0.1)$
Weight intervals	: $(w_{ao}, w_{mh}, w_{mff})$	$= \{[0.8; 1.0], [0.1; 0.2], (0.0; 0.1]\}$

Simulation experiments The simulation experiments were carried out on a Nomad200 simulator, which models a synchro-drive robot with a ring of 16 ultrasonic sensors used for obstacle detection, and position encoders used for calculation of robot position (relative to a target). Using the behavior repertoire described in the previous section the robot should move safely to a designated target position (represented with a cross in Figure 6).

Figure 6 depicts two examples with the robot's trail plotted. On the left, the robot reaches the target using a smooth path around the obstacles and towards the target. The example depicted on the right shows that the robot can get stuck in a deadlock situation, which is a common problem for purely reactive systems. The main reasons for this occurrence is attributed to the lack of global spatial information, lack of memory, and absence of temporal reasoning capabilities.

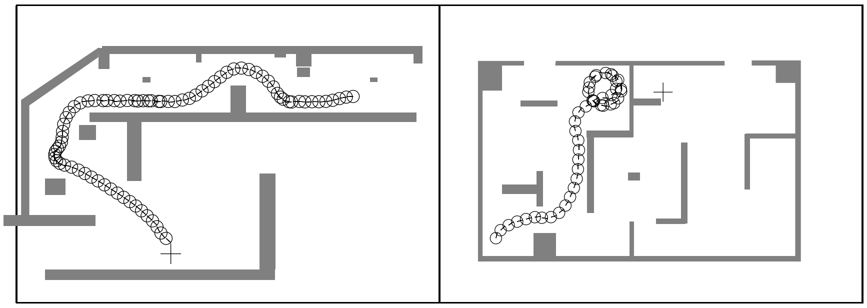

Fig. 6. Simulation examples, showing the results from multiple objective coordination of the **avoid obstacles**, **maintain heading**, and **move fast forward** behaviors for navigation.

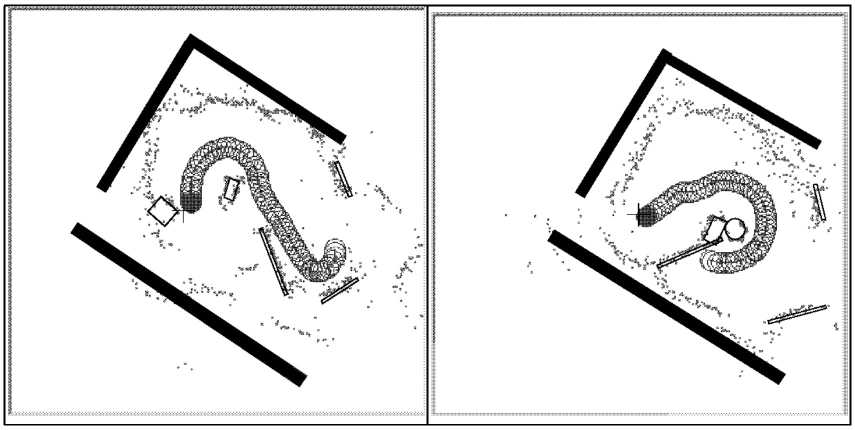

Fig. 7. Example runs with the physical Nomad200 platform. Dots indicate obstacle positions detected by the sonars; obstacles are overlayed as line drawings.

Real-world experiments The simulation experiments described above were repeated on a real Nomad200 robot in a structured indoor laboratory environment (Figure 7).

It can be seen from the examples that the robot successfully negotiates obstacles on its way to the target point. Smooth blending of behaviors is achieved because all the objectives are taken into consideration simultaneously. A detailed analysis of extensive experiments (see [20]) shows that multiple-objective behavior coordination makes suitable trade-offs between behavioral objectives and is able to find solutions that do not severely sacrifice any objective function.

7 Concluding remarks

In this paper, we addressed the main deficiencies of applying fuzzy combination and defuzzification operators for behavior coordination. It is commonly acknowledged that fuzzy approaches can produce undesirable action selection and cannot satisfactorily deal with behavioral conflicts. We introduced multiple objective behavior coordination, which provides formal tools and concepts that can be used to formulate characterizable coordination schemes.

We postulated and argued that multiple objective and fuzzy behavior coordination approaches are compatible due to equivalent semantics of behavior and action selection, and due to their similar mechanics. Hence, they can and should be integrated. This can lead to new approaches that can create a synergy between fuzzy and multiple objective behavior coordination and provide for more powerful tools for devising behavior-based coordination.

7.1 Fuzzy & multiple objective behavior coordination

Note that, as opposed to fuzzy behavior coordination, where there is a clear distinction between behavior combination (e.g., sum-prod) and action selection/decision (defuzzification), in MOBC the combination and action selection processes are very much interleaved. 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 should be evident, the (input/output) mechanics of multiple objective behavior coordination are not different from the fuzzy approaches, hence they can be interchanged.

Further, we use weight values in several of the approaches for generating solutions (Pareto-optimal and satisficing) for multiple objective decision problems (see Sections 4.4 and 4.5). These 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, fuzzy and multiple objective approaches are compatible and thus can be combined. We believe the combination of fuzzy and multiple objective approaches is feasible, and can provide powerful tools for principled and context-dependent behavior coordination.

7.2 Future work

There are three directions we would like to investigate in future work. First, to conduct an empirical study for characterization and comparison of fuzzy and multiple objective behavior coordination (action selection) techniques. Second, to pursue and investigate a combined fuzzy/multiple objective approaches for behavior-based control, which we believe has much potential. Finally, to explore how the proposed approach scales to more complex problems such as team robotics. Intuitively, the coordination of a group of robots seems to have similarities to the action selection problem. The problem of coordination of robots might be formulated in terms of multiple objective decision making where more than one decision maker is present.

Acknowledgments

This work is partially supported by the Office of Naval Research Grant NO014-99-1-0162. We would like to thank the reviewers for providing useful comments.

References

1. Philip E. Agre and David Chapman. What Are Plans for? *Robotics and Autonomous Systems*, 6:17–34, 1990.
2. Ronald C. Arkin. Towards the Unification of Navigational Planning and Reactive Control. *AAAI Spring Symposium on Robot Navigation*, March 1989.
3. Ronald C. Arkin. Integrating Behavioral, Perceptual and World Knowledge in Reactive Navigation. *Robotics and Autonomous Systems*, 6:105–122, June 1990.
4. Ronald C. Arkin. *Behavior-Based Robotics.* Intelligent Robotics and Autonomous Agents series. MIT Press, May 1998.
5. Rodney A. Brooks. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, March 1986.
6. Vira Chankong and Yacov Y. Haimes. *Multiobjective Decision Making - Theory and Methodology*, volume 8. North-Holland, 1983.
7. Dieter Fox, Wolfram Burgard, and Sebastian Thrun. The Dynamic Window Approach to Collision Avoidance. *IEEE Robotics and Automation Magazine*, 4(1):23 – 33, March 1997.
8. Joel Hoff and George Bekey. An Architecture for Behavior Coordination Learning. In *IEEE International Conference on Neural Networks, Perth, Australia*, volume 5, pages 2375–2380, November 1995.
9. Peter G. W. Keen. The Evolving Concept of Optimality. In Martin K. Starr and Milan Zeleny, editors, *Multiple Criteria Decision Making*, volume 6 of *Studies in the management sceiences*, pages 31–57. North-Holland Publishing Company, 1977.
10. Omar Khatib. Real-Time Obstacle Avoidance for Manipulators and Mobile Robots. *The International Journal of Robotics Research*, 5(1):90–98, 1986.
11. Kurt Konolige. *Saphira Software Manual.* ActivMedia Inc., version 6.1 edition, 1998.

12. Jana Koŝecká and Ruzena Bajcsy. Discrete Event Systems for Autonomous Mobile Agents. *Proceedings Intelligent Robotic Systems '93 Zakopane*, pages 21–31, July 1993.
13. Steen Kristensen. Sensor Planning with Bayesian Decision Theory. *Technical Report, Laboratory of Image Analysis, Aalborg University*, June 1995.
14. Edward W. Large, Henrik I. Christensen, and Ruzena Bajcsy. Scaling the Dynamic Approach to Path Planning and Control: Competition among Behavioral Constraints. *International Journal of Robotics Research*, 18:37–58, January 1997.
15. 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.
16. Maja Mataric. Behavior-Based Control: Examples from Navigation, Learning, and Group Behavior. *Journal of Experimental and Theoretical Artificial Intelligence, special issue on Software Architectures for Physical Agents*, 9(2-3):323–336, 1997.
17. David W. Payton. Internalized plans: A representation for action resources. In Pattie Maes, editor, *Designing Autonomous Agents*, pages 89–104. MIT/Elsevier, The MIT Press, Cambridge, Massachusetts, London, England, 1990.
18. David W. Payton, J. Kenneth Rosenblatt, and David M. Keirsey. Plan guided reaction. *IEEE transactions on Systems, Man, and Cybernetics*, 20(6):1370–1382, Novermber/December 1990.
19. F. G. Pin and Y. Watanabe. Automatic Generation of Fuzzy Rules Using the Fuzzy Behaviorist Approach: The Case of Sensor-Based Robot Navigation. *Intelligent Automation and Soft Computing*, 1(2):161–178, 1995.
20. 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. *http://www-robotics.usc.edu/~paolo/publications*.
21. Paolo Pirjanian. Multiple objective action selection in behavior-based control. In *The proceedings of the 6th Symposium for Intelligent Robotic Systems, Edinburgh, Scotland*, pages 83–92, July 1998.
22. Paolo Pirjanian. Satisficing action selection. In *Sensor Fusion and Decentralized Control in Robotic Systems*, volume 3523, pages 157–168. SPIE Conference, November 1998.
23. Paolo Pirjanian. Behavior coordination mechanisms – state-of-the-art. Technical Report IRIS-99-375, Institute of Robotics and Intelligent Systems, School of Engineering, University of Southern California, October 1999. http://iris.usc.edu/~irislib.
24. 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.
25. 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.
26. Enrique H. Ruspini. On truth and utility. In R. Kruse and P. Siege, editors, *Symbolic and Quantitative Approaches to Uncertainty*, pages 297–304. Springer-Verlag, 1991.
27. Alessandro Saffiotti. The Uses of Fuzzy Logic in Autonomous Robot Navigation: a catalogue raisonné. Technical Report 2.1, IRIDIA, Universite Libré de Bruxelles, November 1997.

28. Alessandro Saffiotti, Kurt Konolige, and Enrique H. Ruspini. A multivalued logic approach to integrating planning and control. *Artificial Intelligence*, 76:481–526, March 1995.
29. Herbert A. Simon. *The New Science of Management Decision*. Harper and Row, New York, 1960.
30. Eddie Tunstel. Coordination of Distributed Fuzzy Behaviors in Mobile Robot Control. In *IEEE International Conference on Systems, Man,and Cybernetics*, pages 4009–4014, Vancouver, BC, Canada, October 1995.
31. 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.