

Decision-Theoretic Planning with Linguistic Terms in GOLOG

Stefan Schiffer

Knowledge-Based Systems Group
RWTH Aachen University
Aachen, Germany
schiffer@cs.rwth-aachen.de

Alexander Ferrein

Aachen University of Applied Sciences
Mobile Autonomous Systems &
Cognitive Robotics Institute
Aachen, Germany
ferrein@fh-aachen.de

Abstract

In this paper we propose an extension of the action language GOLOG that integrates linguistic terms in non-deterministic argument choices and the reward function for decision-theoretic planning. It is often cumbersome to specify the set of values to pick from in the non-deterministic-choice-of-argument statement. Also, specifying a reward function is not always easy, even for domain experts. Instead of providing a finite domain for values in the non-deterministic-choice-of-argument statement in GOLOG, we now allow for stating the argument domain by simply providing a formula over linguistic terms and fuzzy fluents. In GOLOG's forward-search DT planning algorithm, these formulas are evaluated in order to find the agent's optimal policy. We illustrate this in the Diner Domain where the agent needs to calculate the optimal serving order.

1 Introduction

The action language GOLOG [Levesque *et al.*, 1997] has proven useful for encoding the high-level behaviors of a robot or agent (e.g. [Ferrein and Lakemeyer, 2008; Schiffer *et al.*, 2012]). With its foundations in the Situation Calculus [McCarthy, 1963; Reiter, 2001], complex behaviors are described in terms of actions with preconditions and effects. The world evolves from an initial situation due to actions. So-called fluents (predicates with a situation term as the last argument) keep track of changes of the properties of the world. Many extensions to the original GOLOG dialect were proposed, for instance, to deal with continuous change, allow for probabilistic projections, or decision-theoretic planning (e.g. [Grosskreutz, 2000; Grosskreutz and Lakemeyer, 2001; Boutilier *et al.*, 2000]). We build on a variant of GOLOG called READYLOG [Ferrein and Lakemeyer, 2008] which integrates many of the different dialects into an online interpreter that allows to encode high-level behaviors of an agent for dynamic real-time domains. READYLOG has shown its usefulness in applications ranging from robotic soccer to domestic service robots [Schiffer *et al.*, 2006; 2012].

One of the features that we found particularly useful to define the behavior of an agent or robot in a flexible way was

to use decision-theoretic planning. The programmer states the different action alternatives and a reward function in order to select preferred world situations; the optimal policy is then calculated and executed. Besides this non-deterministic choice of actions, GOLOG offers a non-deterministic-choice-of-argument statement. For a finite domain of arguments, an optimal policy is computed. This is in particular helpful when the agent faces incomplete knowledge and particular information has to be sensed at run-time. However, to specify the argument domain is often cumbersome. In this paper, we propose an extension to GOLOG that integrates linguistic terms in non-deterministic argument choices and the reward function for decision-theoretic planning. We demonstrate the extension in the Diner Domain, where a waitron agent has to serve coffee and dishes as hot as possible. The agent has to decide on which order to deliver first as the coffee and meals cool down over time. We show how linguistic terms as defined in the Fuzzy Logic extension of GOLOG [Ferrein *et al.*, 2008; Schiffer *et al.*, 2011] are formally integrated into the forward-search value iteration algorithm used in READYLOG.

The remainder of the paper is organized as follows. In the next section, we briefly introduce the Diner Domain. Then, in Section 3, we review the background of this work, namely READYLOG and our Fuzzy Logic extensions to the Situation Calculus. In Section 4, we bring together the linguistic terms and decision-theoretic planning and define the respective language constructs formally. We conclude with Section 5.

2 The Diner Domain

The example we use in this paper will be from the Diner Domain. In the Diner Domain, a waitron agent has to decide which of its assigned tables it should serve first in order to serve coffee and meals as hot as possible. Of course, the longer the distance for coffee and meals to be served, the cooler the dishes will be when served to the customer. In our example, the waitron was assigned to serve tables T_1 and T_9 . The coffee is regarded as cold if its temperature lies between 0–50 centigrades, it is perceived as luke warm between 45 and 65 degrees, hot between 60 and 80 degrees; above 75 degrees we regard the coffee as veryhot. Despite a negative exponential cooling rate in reality, we assume a linear rate for the sake of simplicity in this example. For every 10 seconds we assume that coffee and meals cool down 1 degree. Traversing a square in the Diner Domain takes the waitron

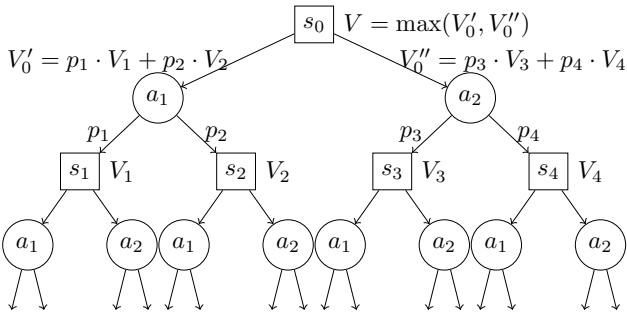


Figure 3: Decision tree search in READYLOG representing a non-deterministic choice of actions. In each situation s_i the agent may choose between a_1 or a_2 . The resulting optimal policy is represented by the branch with maximal value V .

predicates. For non-deterministic choices of actions, its formal definition following [Boutilier *et al.*, 2000] is:

$$\begin{aligned}
 \text{BestDo}((p_1 \mid p_2); p, s, h, \pi, v, pr) &\stackrel{\text{def}}{=} \\
 &\exists \pi_1, v_1, pr_1. \text{BestDo}(p_1; p, s, h, \pi_1, v_1, pr_1) \wedge \\
 &\exists \pi_2, v_2, pr_2. \text{BestDo}(p_2; p, s, h, \pi_2, v_2, pr_2) \wedge \\
 &((v_1, p_1) \geq (v_2, p_2) \wedge \pi = \pi_1 \wedge pr = pr_1 \wedge v = v_1) \vee \\
 &((v_1, p_1) < (v_2, p_2) \wedge \pi = \pi_2 \wedge pr = pr_2 \wedge v = v_2)
 \end{aligned}$$

The non-deterministic choice of action arguments is defined as:

$$\begin{aligned}
 \text{BestDo}(\text{pickBest}(x, \tau, p; p'), s, h, \pi, v, pr) &\stackrel{\text{def}}{=} \\
 \text{BestDo}(p|_{c_1}^x \mid \dots \mid p|_{c_n}^x; p', s, h, \pi, v, pr) &\quad (1)
 \end{aligned}$$

Free variables x in the program p are bound to a finite domain τ ; for each “variable assignment” (denoted by $p|_{c_i}^x$) a new non-deterministic branch in the forward-search DTP is added. The policy is hence optimized for all possible variable assignments leading to the assignment which maximizes the reward function.

3.2 Fuzzy GOLOG

In [Ferrein *et al.*, 2008; Schiffer *et al.*, 2011] we introduced the notion of fuzzy fluents in GOLOG. Fuzzy fluents extend “ordinary” functional fluents in that they have a membership relation that defines, for a number of linguistic fuzzy terms the degree of membership for a particular function value.

In our Diner Domain, we want to serve hot coffee to our customers. The coffee, however, cools down quickly, depending on how long it takes to deliver the coffee. This, in turn, depends on the distance between the counter and the table where the coffee should be served. As an example for a fuzzy set defining the linguistic terms, we look at a distance relation. Distance in our diner example is understood as the Manhattan distance between two positions in the diner. We define distances between 0 and 3 blocks as close, between 3

Algorithm 1: Decision-theoretic path planning in READYLOG

```

1 proc navigate
2   solve( $h, \text{reward}$ , while  $\text{loc} \neq \text{goal}$  do
3     ( $\text{go\_right} \mid \text{go\_left} \mid \text{go\_up} \mid \text{go\_down}$ )
4   endwhile)
5 endproc

```

and 6 as medium, and above 6 as far. Formally,

$$\begin{aligned}
 \mathfrak{F}(\text{distance}, u, \mu_u) &\equiv \\
 (\text{distance} = \text{close} \supset (0, 1.0) \vee (1, 1.0) \vee \\
 &(2, 0.75) \vee (3, 0.25) \vee (13/12, 0.5)) \wedge \\
 (\text{distance} = \text{medium} \supset (3, 0.25) \vee (4, 0.75) \vee \\
 &(5, 0.75) \vee (6, 0.25) \vee (9/2, 0.5)) \wedge \\
 (\text{distance} = \text{far} \supset (6, 0.25) \vee (7, 0.75) \vee \\
 &(8, 1.0) \vee (9, 1.0) \vee (95/12, 0.5)),
 \end{aligned}$$

where we use (u_i, μ_{u_i}) as an abbreviation for $u = u_i \wedge \mu_u = \mu_{u_i}$. The fuzzy set for the coffee temperature is shown in Fig. 4. Note that fuzzy categories can overlap. For instance, the coffee temperature 62°C belongs to the category luke as well as to the category hot. To query whether a value belongs to a certain category, one has to check if in the respective fuzzy set the value has a positive membership degree in that particular categorize. This is done with the predicate $\text{is} \subseteq \text{real} \times \text{linguistic}$. It is defined as

$$\text{is}(\mathfrak{f}(\vec{t}, \sigma), \gamma) \stackrel{\text{def}}{=} \exists u, \mu_u. \mathfrak{f}(\vec{t}, \sigma) = u \wedge \mathfrak{F}(\gamma, u, \mu_u) \wedge \mu_u > 0,$$

where \mathfrak{f} having the numeric value u is the fuzzy fluent to be queried, \mathfrak{F} is the respective fuzzy set and μ_u is the degree of membership of value u in the fuzzy set \mathfrak{F} . In our distance example above, for instance, we have $\mathfrak{F}(\text{medium}, 5, 0.75)$ to say that the numerical value 5 has a membership degree of 0.75 for the category medium. The predicate holds if the degree of membership is greater zero. For complex queries (logical formulas with fuzzy fluents), we have to define similar predicates is_\ominus for the complement, is_\star for the conjunction, and is_\oplus for the disjunction of fuzzy fluents. See [Ferrein *et al.*, 2008; Schiffer *et al.*, 2011] for the formal definitions. Further, we need to define a function to defuzzify a linguistic term to a numeric value. As a defuzzifying function, we use the center of gravity (cog) which we formally define in [Ferrein *et al.*, 2008; Schiffer *et al.*, 2011]. In our distance example, the center of gravity for the category close is $13/12$. Note that we have to manually add the center of gravity for this respective category in our Situation Calculus fuzzy set formalization.

In the Diner Domain, we want to refer to positions in a room in a qualitative manner. This is why we introduce linguistic categories for the position in X and Y by the following membership functions:

$$\begin{aligned}
 \mathfrak{F}(\text{pos}X, u, \mu_u) &\equiv \\
 (\text{pos}X = \text{left} \supset (1, 1.0) \vee (2, 1.0) \vee (3, 1.0) \wedge \\
 (\text{pos}X = \text{center} \supset (4, 1.0) \vee (5, 1.0) \vee (6, 1.0) \wedge \\
 (\text{pos}X = \text{right} \supset (7, 1.0) \vee (8, 1.0) \vee (9, 1.0)).
 \end{aligned}$$




Figure 4: Coffee temperature membership function

For the y -coordinate we introduce a fuzzy fluent $posY$ and define the categories front, middle, back, referring to the tables whose ordinate have a distance of close, medium, and far from the *Counter*. In the next section, we propose an extension to DTP integrating those linguistic notions.

4 Extending DT-Planning in GOLOG with Linguistic Terms

One of the convenient features when specifying intelligent agents in GOLOG is that the agent designer can leave choices open that the agent then resolves on its own using an optimization theory. As already mentioned, the choices are the non-deterministic choice of action and the non-deterministic choice of argument. The latter is realized by means of the **pickBest** statement. It allows for specifying a set of possible values for a specific fluent for the program in the body of the statement. That program is evaluated with any of the values from the set.

4.1 Picking from Fuzzy Sets

We now propose to use, instead of a finite set of values, a fuzzy expression to specify the set of possible values for a fuzzy fluent. We introduce a new predicate **pickBestF** which takes a fuzzy expression instead of the regular set of the classical **pickBest**.

The idea is that instead of giving a finite set of variable or fluent values in the **pickBest** statement, the programmer now can state a formula specifying linguistic categories for a fuzzy fluent. For instance, if we want to optimize the coffee serving temperature, we could simply state to choose a coffee whose temperature is hot. What the **pickBestF** statement does is to translate this into a set of temperatures with positive membership values for the category hot. In our case shown in Fig. 4, this would be translated into the temperatures 60–80 centigrades. For each of the temperatures, the forward-search algorithm would try and optimize the respective program attached, say, $serveCoffee(T_9)$ (serve a coffee at table 9) with the **pickBestF** statement.

For a single linguistic category we define **pickBestF** as

$$\begin{aligned}
 & BestDo(\mathbf{pickBestF}(f : \gamma, p); p', s, h, \pi, v, pr) \stackrel{def}{=} \\
 & \exists u_1, \dots, u_k. \bigvee_{u \in \{u_1, \dots, u_k\}} [f[s] = u \wedge is_{\mathbb{C}}(f[s] = u, \gamma)] \wedge \\
 & \neg \exists u_j. (f[s] = u_j \wedge is_{\mathbb{C}}(f[s], \gamma) \wedge \\
 & \quad u_j \neq u_1 \wedge \dots \wedge u_j \neq u_k) \wedge \\
 & BestDo(p|_{u_1}^f | \dots | p|_{u_k}^f); p', s, h, \pi, v, pr)
 \end{aligned}$$

The intuitive meaning of the above definition is to collect all possible numerical values of a linguistic category as follows: First, we assume that the predicate holds for k numerical values u_i of the category γ . Then, we “check” if these k values are all values for which $is(\cdot, \gamma)$ holds. Lastly, we call *BestDo*, replacing the fluent f in the program p with any such value (denoted by $p|_{u_i}^f$).¹ This is analogous to the definition of *BestDo* for **pickBest** (Eq. 1), where the fluent was replaced with any element of the set τ (cf. Section 3.1). We give the remaining definitions for the complement of a linguistic category and for conjunction and disjunction of several linguistic categories below.

If the expression is the *complement* of a linguistic category we have

$$\begin{aligned}
 & BestDo(\mathbf{pickBestF}(f : \neg\gamma, p); p', s, h, \pi, v, pr) \stackrel{def}{=} \\
 & \exists u_1, \dots, u_k. \bigvee_{u \in \{u_1, \dots, u_k\}} [f[s] = u \wedge is_{\mathbb{C}}(f[s] = u, \gamma)] \wedge \\
 & \neg \exists u_j. (f[s] = u_j \wedge is_{\mathbb{C}}(f[s], \gamma) \wedge \\
 & \quad u_j \neq u_1 \wedge \dots \wedge u_j \neq u_k) \wedge \\
 & BestDo(p|_{u_1}^f | \dots | p|_{u_k}^f); p', s, h, \pi, v, pr)
 \end{aligned}$$

For a *conjunction* of n linguistic categories we have

$$\begin{aligned}
 & BestDo(\mathbf{pickBestF}(f : \Gamma_*, p); p', s, h, \pi, v, pr) \stackrel{def}{=} \\
 & \exists u_1, \dots, u_k. \bigvee_{u \in \{u_1, \dots, u_k\}} [f[s] = u \wedge is_{\star}(f[s] = u, \Gamma_*)] \wedge \\
 & \neg \exists u_j. (f[s] = u_j \wedge is_{\star}(f[s], \Gamma_*) \wedge \\
 & \quad u_j \neq u_1 \wedge \dots \wedge u_j \neq u_k) \wedge \\
 & BestDo(p|_{u_1}^f | \dots | p|_{u_k}^f); p', s, h, \pi, v, pr)
 \end{aligned}$$

where Γ_* is an abbreviation for $\Gamma_* \stackrel{def}{=} \gamma_1 \wedge \dots \wedge \gamma_n$. For the *disjunction* of n linguistic categories we have

$$\begin{aligned}
 & BestDo(\mathbf{pickBestF}(f : \Gamma_{\oplus}, p); p', s, h, \pi, v, pr) \stackrel{def}{=} \\
 & \exists u_1, \dots, u_k. \bigvee_{u \in \{u_1, \dots, u_k\}} [f[s] = u \wedge is_{\oplus}(f[s] = u, \Gamma_{\oplus})] \wedge \\
 & \neg \exists u_j. (f[s] = u_j \wedge is_{\oplus}(f[s], \Gamma_{\oplus}) \wedge \\
 & \quad u_j \neq u_1 \wedge \dots \wedge u_j \neq u_k) \wedge \\
 & BestDo(p|_{u_1}^f | \dots | p|_{u_k}^f); p', s, h, \pi, v, pr)
 \end{aligned}$$

Γ_{\oplus} is an abbreviation for $\Gamma_{\oplus} \stackrel{def}{=} \gamma_1 \vee \dots \vee \gamma_n$. By the above definitions we provide our new **pickBestF** allowing to specify the argument choice in terms of a fuzzy expressions for a single fuzzy fluent based on the existing *BestDo* statements for standard sets. The idea is to branch over all fluent values for which the fuzzy expression holds in the resulting non-deterministic choice of action statement having all occurrences of f replaced by the respective value from the respective fuzzy set.

4.2 Fuzzy Expressions in the Reward Function

As a second way to increase the naturalness of specifying programs for decision-theoretic planning we introduce a means

¹Note that for readability we only use f to refer to a fuzzy fluent. $f[s]$ denotes the fluent f with its situation argument being restored which we need to determine its value in a particular situation.

to use fuzzy expression in the reward function. We propose a statement **fcase** that modifies the reward according to a fuzzy expression, i.e., a single fuzzy category, the complement of a single fuzzy category, the conjunction of several fuzzy categories and the disjunction of multiple categories (all for the same fuzzy fluent f).

The **fcase** statement distinguishes the four above cases and handles them according to the following definitions. For a single linguistic category

$$\begin{aligned} \mathbf{fcase}(f, \gamma, r) &= \mathit{reward} \stackrel{\text{def}}{=} \\ &\text{is}(f[s], \gamma) \wedge (\mathit{reward} = r) \vee \\ &\neg \text{is}(f[s], \gamma) \wedge (\mathit{reward} = 0) \end{aligned}$$

For the complement of a single linguistic category

$$\begin{aligned} \mathbf{fcase}(f, \neg\gamma, r) &= \mathit{reward} \stackrel{\text{def}}{=} \\ &\text{is}_{\mathbb{C}}(f[s], \gamma) \wedge (\mathit{reward} = r) \vee \\ &\neg \text{is}_{\mathbb{C}}(f[s], \gamma) \wedge (\mathit{reward} = 0) \end{aligned}$$

For the conjunction of n linguistic categories

$$\begin{aligned} \mathbf{fcase}(f, \gamma_1 \wedge \dots \wedge \gamma_n, r) &= \mathit{reward} \stackrel{\text{def}}{=} \\ &\text{is}_{\star}(f[s], \gamma_1, \dots, \gamma_n) \wedge (\mathit{reward} = r) \vee \\ &\neg \text{is}_{\star}(f[s], \gamma_1, \dots, \gamma_n) \wedge (\mathit{reward} = 0) \end{aligned}$$

For the disjunction of n linguistic categories

$$\begin{aligned} \mathbf{fcase}(f, \gamma_1 \vee \dots \vee \gamma_n, r) &= \mathit{reward} \stackrel{\text{def}}{=} \\ &\text{is}_{\oplus}(f[s], \gamma_1, \dots, \gamma_n) \wedge (\mathit{reward} = r) \vee \\ &\neg \text{is}_{\oplus}(f[s], \gamma_1, \dots, \gamma_n) \wedge (\mathit{reward} = 0) \end{aligned}$$

4.3 Walking through an Example

To illustrate the extensions proposed above, consider the following example in our Diner Domain: The restaurant has several waitron robots and we need to specify the control program for one of them. Assume the robot is responsible for the tables located in the corners of the room. Let the position of the tables be composed of their x and y cell-coordinates. In terms of a linguistic description, we might then say that the robot needs to serve tables that are in the left or the right part of the room and that are in the front or the back part of the room. The robot can take orders for coffee or meals from any of the tables it needs to serve. Assume the robot has a (finite) list of orders in its world model, each with a number, the table it came from and the temperature the meal was served with. The individual properties of those orders can be retrieved via respective functions, where order_i is used to refer to the order number i . The serving temperature is zero for as long as an order has not been served.

Writing a program for such an agent includes letting the robot choose which table to serve in which order. Using decision-theoretic planning, we can specify an optimization theory by means of a reward function. With our newly introduced **pickBestF** statement we can write a control program in a very straight-forward manner as given in Alg. 2.

Let us assume that the reward is computed by giving a negative amount for any open order (i.e. any order that has not

Algorithm 2: Decision-theoretic planning in READYLOG for serving a room with fuzzy argument choice

```

1 proc serve_room
2   navigate(counter);
3   while haveOpenOrder(room) do
4     pickBestF(posX, left  $\vee$  right) {
5       pickBestF(posY, front  $\vee$  back) {
6         tableWithOpenOrder(table, posX, posY);
7
8         pickBestF(mealTemp, luke  $\vee$  hot) {
9           mealWithTempReady(meal, mealTemp);
10          load_meal(meal, tray);
11          bring_meal(tray, table);
12          serve_meal(tray, table); } }
13   endwhile
14 endproc

```

been served yet) and by giving a positive amount for food being served with a high temperature. We can use the newly introduced **fcase** statement to specify such a reward function with linguistic terms as follows. For simplicity, we limit ourselves to a list of only two orders.

$$\begin{aligned} \mathit{reward}(s) &= r \stackrel{\text{def}}{=} \\ &r = \text{numOpenOrders}(s) \cdot (-100) + \\ &\mathbf{fcase}(\text{serveTemp}(\text{order}_1, s), \text{hot}, 100) + \\ &\mathbf{fcase}(\text{serveTemp}(\text{order}_2, s), \text{hot}, 100) + \\ &\mathbf{fcase}(\text{serveTemp}(\text{order}_1, s), \text{luke}, 10) + \\ &\mathbf{fcase}(\text{serveTemp}(\text{order}_2, s), \text{luke}, 10) \end{aligned}$$

The fuzzy fluent serveTemp returns the temperature at which a meal was served.

Let us assume, the robot has orders from tables T_9 and T_1 . For simplicity we assume there is no coffee and only one meal to order hence both tables may be served with the same meal. The robot finds two meals M_1 and M_2 ready to serve on the counter with temperatures of 54 and 74 centigrades respectively. If the robot uses the above program it yields an execution trace as follows.

The first **pickBestF** statement has a disjunction as its fuzzy expression. Hence, we apply the corresponding *BestDo* definition. That is, by means of the existential quantifiers we collect those u_i (and only those!) for which the predicate $\text{is}_{\oplus}(u_i, \text{left}, \text{right})$ is true. Using the \mathfrak{F} definition for the posX fuzzy fluent we find six values, namely 1, 2, 3, 7, 8, 9. Similarly for the second **pickBestF**-statement we collect possible y -coordinates 1, 2, 3, 7, 8, 9. Using the *BestDo* definition, we replace in the body of the **pickBestF** statement the variables posX and posY by any of the available values. For each combination we check whether there is a table with an open order at that position with the predicate $\text{tableWithOpenOrder}$. The only positions for which this is true are T_1 and T_9 . For those two tables we continue with the program, i.e. we do another **pickBestF**, now for the fluent mealTemp . Again, us-

Table 1: Table of possible courses of action with their corresponding reward.

serving order	serve-temp1	serve-temp2	total reward
$(M_1, T_1), (M_2, T_9)$	51	62	120
$(M_1, T_9), (M_2, T_1)$	48	59	20
$(M_2, T_1), (M_1, T_9)$	71	42	110
$(M_2, T_9), (M_1, T_1)$	68	39	110

ing the *BestDo* definition for a disjunctive fuzzy expression (hot \vee luke) we collect a set of values to replace the fluent mealTemp in the remaining program. In our example this is the set $\{45, \dots, 80\}$ as per our specification of the fuzzy sets \mathfrak{F} for luke and hot (cf. Sect. 2). Hence, we consider to execute the sequence inside the innermost **pickBestF** for any combination of existing table positions with open orders in the areas that our robot has to serve, each with any of the meals available with temperatures from the set $\{45, \dots, 80\}$ that we have from our *BestDo* definition. First we check whether a meal with a given temperature is ready on the counter by the predicate mealWithTempReady. Only if this is the case, we attempt to load the meal, bring it to a table and serve it. Starting from the initial situation as given above, from the sets constructed by our *BestDo* definitions for **pickBestF** by means of the two predicates tableWithOpenOrder and mealWithTempReady what remains for the innermost program part are for the position (8, 8) and (2, 2), each in combination with a meal of either 54 or 74 centigrades temperature.

Starting at the counter, we need 12 steps to reach table T_9 and 6 steps to reach table T_1 . This means, a meal cools down by $12 \cdot 5/10 = 6$ centigrades when it is being delivered to T_9 and $6 \cdot 5/10 = 3$ when it is being delivered to T_1 . With the two orders to serve and two meals to pick from for each we are left with four courses of action, shown with their reward in Tab. 1. The reward for the course of actions essentially depends on the temperature that each meal is being served at. For meals being served with luke temperature the agent receives a reward of 10, for those being served hot it is rewarded with 100. The most rewarding situation is reached with first serving table T_1 with meal M_1 , and then delivering M_2 to T_9 . This yields a total reward of 120. The policy returned for the agent to execute then is navigate(counter), tableWithOpenOrder(T_9 , 8, 8), mealWithTempReady(M_1 , 74), load_meal(M_1 , tray), bring_meal(tray, T_9), serve_meal(tray, T_9), navigate(counter), tableWithOpenOrder(T_1 , 2, 2), mealWithTempReady(M_2 , 54), load_meal(M_2 , tray), bring_meal(tray, T_1), serve_meal(tray, T_1), navigate(counter).

Our newly introduced constructs allow for a seamless integration of linguistic notions in decision-theoretic planning in agent programs. The agent designer can use fuzzy expressions both, to specify the set of values to pick from for the non-deterministic choice of argument and to specify portions of the reward function that is used as the underlying optimization theory in decision-theoretic planning.

5 Discussion

In this paper we proposed an extension to READYLOG which combines fuzzy fluents and decision-theoretic planning. Fuzzy fluents are fluents that have a membership function attached. With this function, it can be checked whether or not a fluent value belongs to the linguistic category in question. In our previous work, we defined a predicate “is” to test this. With this predicate, we can handle negation, conjunction and, disjunction, respectively, of linguistic terms. The decision-theoretic extension of GOLOG implements a forward-search value iteration algorithm and allows to optimize non-deterministic choices of action or arguments w.r.t. a given reward function. The search for an optimal policy can be guided by a GOLOG program to restrict the search space. Our practical experiences with programming robots with READYLOG shows that, in particular, non-deterministic choices of actions and arguments are very useful when specifying the behavior of a robot or agent in a flexible way. For the non-deterministic choice of arguments, the programmer has to give a finite domain from which the program arguments for computing the optimal policy are evaluated. This can be a cumbersome process.

In this work, we extend the non-deterministic argument choice such that it can handle simple fuzzy fluent formulas. This facilitates the specification of the argument set in **pickBest** statements. To this end, we introduced and defined a statement **pickBestF** that translates the values for which the given fuzzy fluent formula holds as an argument set for the ordinary **pickBest** statement. Further, we introduced a statement **fcase** which allows to use simple fuzzy fluent formulas in the reward function of the forward-search value iteration algorithm. The programmer can make use of linguistic terms and fuzzy categories when assigning rewards to preferred world situations. We showed the use of the new constructs by an example from the Diner Domain, where a waitron agent has to find an optimal schedule to serve coffee or dishes to its customers.

Our further steps are as follows. So far, we did not allow arbitrary formulas over fuzzy fluents yet. Enabling such formulas is not as easy because, for example, in fuzzy logic the excluded middle does not always hold. We will look into possible realizations of more complex formulas. Furthermore, broadening the application of linguistic terms in planning with preferences is on our agenda. Here we want to investigate how our work can be married with the work of [Fritz and McIlraith, 2006; Bienvenu *et al.*, 2011] who compile modal logic preference formulas into GOLOG programs. Similarly, we will have a look into [Finzi and Pirri, 2004; Cesta *et al.*, 2011] who use temporal interval planning to solve scheduling problems.

References

- [Bienvenu *et al.*, 2011] Meghyn Bienvenu, Christian Fritz, and Sheila A McIlraith. Specifying and computing preferred plans. *Artificial Intelligence*, 175(7):1308–1345, 2011.
- [Boutilier *et al.*, 2000] Craig Boutilier, Ray Reiter, Mikhail Soutchanski, and Sebastian Thrun. Decision-theoretic, high-level agent programming in the situation calculus. In *Proc. AAAI-00*, pages 355–362. AAAI Press, 2000.
- [Cesta *et al.*, 2011] Amedeo Cesta, Simone Fratini, Andrea Orlandini, Alberto Finzi, and Enrico Tronci. Flexible plan verification: Feasibility results. *Fundamenta Informaticae*, 107(2):111–137, 2011.
- [De Giacomo *et al.*, 2000] G. De Giacomo, Y. Lésperance, and H. Levesque. ConGolog, A concurrent programming language based on situation calculus. *Artificial Intelligence*, 121(1–2):109–169, 2000.
- [Ferrein and Lakemeyer, 2008] Alexander Ferrein and Gerhard Lakemeyer. Logic-based robot control in highly dynamic domains. *Robotics and Autonomous Systems, Special Issue on Semantic Knowledge in Robotics*, 56(11):980–991, 2008.
- [Ferrein *et al.*, 2008] A. Ferrein, S. Schiffer, and G. Lakemeyer. A fuzzy set semantics for qualitative fluents in the situation calculus. In Caihua Xiong, Honghai Liu, Yonggan Huang, and Youlun Xiong, editors, *Proceedings of the International Conference on Intelligent Robotics and Applications*, volume 5314 of *Lecture Notes in Computer Science*, pages 498–509. Springer, 2008.
- [Ferrein, 2010] Alexander Ferrein. Robot controllers for highly dynamic environments with real-time constraints. *Künstliche Intelligenz*, 24(2):175–178, 2010.
- [Finzi and Pirri, 2004] Alberto Finzi and Fiorenza Pirri. Flexible interval planning in concurrent temporal golog. In *Working notes of the 4th international cognitive robotics workshop*, 2004.
- [Fritz and McIlraith, 2006] Christian Fritz and Sheila A McIlraith. Decision-theoretic golog with qualitative preferences. In *KR*, pages 153–163, 2006.
- [Grosskreutz and Lakemeyer, 2001] Henrik Grosskreutz and Gerhard Lakemeyer. On-line execution of cc-Golog plans. In Bernhard Nebel, editor, *Proc. IJCAI-01*. Morgan Kaufmann, 2001.
- [Grosskreutz, 2000] H. Grosskreutz. Probabilistic projection and belief update in the pgolog framework. In *CogRob-00*, pages 34–41. ECAI-00, 2000.
- [Levesque *et al.*, 1997] H. Levesque, R. Reiter, Y. Lésperance, F. Lin, and R. Scherl. GOLOG: A logic programming language for dynamic domains. *Journal of Logic Programming*, 31(1-3):59–83, 1997.
- [McCarthy, 1963] J. McCarthy. Situations, actions and causal laws. Technical report, Stanford University, 1963.
- [Reiter, 2001] R. Reiter. *Knowledge in Action*. MIT Press, 2001.
- [Schiffer *et al.*, 2006] Stefan Schiffer, Alexander Ferrein, and Gerhard Lakemeyer. Football is coming home. In *Proceedings of the 2006 International Symposium on Practical Cognitive Agents and Robots (PCAR’06)*, pages 39–50, New York, NY, USA, November 27-28 2006. ACM.
- [Schiffer *et al.*, 2011] Stefan Schiffer, Alexander Ferrein, and Gerhard Lakemeyer. Fuzzy representations and control for domestic service robots in golog. In Sabina Jeschke, Honghai Liu, and Daniel Schilberg, editors, *Proceedings of the 4th International Conference on Intelligent Robotics and Applications (ICIRA 2011)*, volume 7102 of *Lecture Notes in Computer Science*, pages 241–250. Springer, 2011.
- [Schiffer *et al.*, 2012] Stefan Schiffer, Alexander Ferrein, and Gerhard Lakemeyer. CAESAR: An Intelligent Domestic Service Robot. *Journal of Intelligent Service Robotics*, 5(Special Issue on Artificial Intelligence in Robotics: Sensing, Representation and Action):259–273, 2012.