

Decision-Theoretic Planning with Fuzzy Notions in Golog

Stefan Schiffer

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

Alexander Ferrein

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

In this paper we present an extension of the action language GOLOG that allows for using fuzzy notions in non-deterministic argument choices and the reward function in decision-theoretic planning. Often, in decision-theoretic planning, it is cumbersome to specify the set of values to pick from in the non-deterministic-choice-of-argument statement. Also, even for domain experts, it is not always easy to specify a reward function. 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 GOLOG¹⁴ has proven useful for encoding the high-level behaviors of agents or robots (e.g.^{7,22}). With its foundations in the Situation Calculus,^{15,17} 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 their last argument) are used to keep track of changes of the properties of the world. Many extensions to the original GOLOG dialect have been proposed, for instance, to deal with continuous change, to allow for probabilistic projections, or for decision-theoretic planning (e.g.^{2,11,12}). We build on a variant of GOLOG called READYLOG⁷ 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.^{19,22}

One of the features that we found particularly useful to define the behavior of

an agent or a 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 show how linguistic terms as defined in the Fuzzy Logic extension of GOLOG^{5,20} are formally integrated into the forward-search value iteration algorithm used in READYLOG. Then, 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. This paper is an updated and extended version of a workshop paper.¹⁸ It brings together our previous contributions on using linguistic notions in logic-based agent control^{5,8,20,21} with the idea of integrating these linguistic notions in decision-theoretic planning.

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

2. Background

In this section we briefly introduce the Situation Calculus and READYLOG, showing the forward-search value iteration algorithm in greater detail. Then, we outline how Fuzzy sets can be formalized in the Situation Calculus.

2.1. *Situation Calculus and Readylog*

The Situation Calculus is a second order language with equality which allows for reasoning about actions and their effects. The world evolves from an initial situation due to primitive actions. Possible world histories are represented by sequences of actions. The Situation Calculus distinguishes three different sorts: *actions*, *situations*, and domain *objects*. A special binary function symbol $do : action \times situation \rightarrow situation$ exists, with $do(a, s)$ denoting the situation which arises after performing action a in the situation s . The constant S_0 denotes the initial situation, i.e. the situation where no actions have occurred yet. The state the world is in is characterized by functions and relations with a situation as their last argument. They are called *functional* and *relational fluents*, respectively. Actions in the Situation Calculus are characterized by unique names. For each action one has to specify a *precondition*

axiom stating under which conditions it is possible to perform the respective action and an *effect axiom* formulating how the action changes the world in terms of the specified fluents.

An action precondition axiom has the form $Poss(a(\vec{x}), s) \equiv \Phi(\vec{x}, s)$. The binary predicate $Poss : action \times situation$ defines the precondition of a particular action. If $Poss$ evaluates to \top , the action can be executed. \vec{x} stands for the arguments of action a . For instance, for a robot's move action, the precondition axiom might be $Poss(move(x, y), s) \equiv robotLoc(s) = x$ saying that the robot can only move from x to y if its current location is x .^a After having specified when it is physically possible to perform an action, we need to describe how the action changes the world. In the Situation Calculus, one has to specify negative and positive effects in terms of the relational fluent F , i.e. $\varphi_F^-(\vec{x}, s) \supset \neg F(\vec{x}, do(a, s))$ and $\varphi_F^+(\vec{x}, s) \supset F(\vec{x}, do(a, s))$, respectively. A fluent is a special predicate with a situation term as its last argument. It can vary from situation to situation and describes a property of the world in a specific situation. The effect axiom for a functional fluents f is $\varphi_f(\vec{x}, y, a, s) \supset f(\vec{x}, do(a, s)) = y$. However, describing the positive and negative effect says nothing about those effects which do not change the fluent. The problem of describing the non-effects of an action is referred to as the *frame problem*. The number of frame axioms is very large. For relational fluents there exist in the order of $2 \cdot \mathcal{A} \cdot \mathcal{F}$ frame axioms, where \mathcal{A} is the number of actions, and \mathcal{F} the number of relational fluents. McCarthy & Hayes were the first to mention this problem.¹⁶

A solution to the problem was proposed with so-called *successor state axioms*.¹⁷ The idea behind these axioms is that, if the truth value of F changes from false to true from situation s to situation $do(a, s)$, then $\varphi_F^+(\vec{x}, a, s)$ must have been true. Similarly, for the second axiom. Reiter shows that under consistency assumptions for fluents together with the explanation closure axioms, the normal form axioms for fluent F are logically equivalent to

$$F(\vec{x}, do(a, s)) \equiv \varphi_F^+(\vec{x}, a, s) \vee F(\vec{x}, a, s) \wedge \neg \varphi_F^-(\vec{x}, a, s). \quad (1)$$

The above formula is called *successor state axiom for the relational fluent F* . The successor state axiom for the functional fluent f has the form:¹⁷

$$f(\vec{x}, do(a, s)) = y \equiv \varphi_f(\vec{x}, y, s) \vee f(\vec{x}, s) = y \wedge \neg \exists y'. \varphi_f(\vec{x}, y', a, s) \quad (2)$$

The background theory for reasoning then is a set of sentences consisting of the effect axioms \mathcal{D}_{ssa} , the action precondition axioms \mathcal{D}_{ap} and axioms characterizing the initial situation \mathcal{D}_{S_0} (along with some foundational axioms Σ and a unique names assumption \mathcal{D}_{una}). It is called a *basic action theory* (BAT) \mathcal{D} . For further details on the Situation Calculus we refer to Reiter.¹⁷ Based on the axioms of the Situation Calculus, the language READYLOG,^{6,7} our variant of GOLOG,¹⁴ is defined and borrows ideas from existing extensions to GOLOG^{2,4,11,12,14} and features the

^aNote that all logical sentences are implicitly universally quantified.

<i>nil</i>	empty program
α	primitive action
$\varphi?$	wait/test action
waitFor (τ)	event-interrupt
$[\sigma_1; \sigma_2]$	sequence
if φ then σ_1 else σ_2 endif	conditional
while φ do σ endwhile	loop
withCtrl φ do σ endwithCtrl	guarded execution
$\sigma_1 \parallel \sigma_2$	prioritized execution
withPol (σ_1, σ_2)	prioritized exec. of σ_2
prob (p, σ_1, σ_2)	probabilistic exec. of σ_1, σ_2
pproj (c, σ)	prob. projection of prog's
{proc $P_1(\vec{\vartheta}_1)\sigma_1$ endproc; ... }	procedures
solve (h, f, σ)	initiate decision-theoretic optimization over σ up to a fixed horizon h
$\sigma_1 \mid \sigma_2$	non-deterministic decision-theoretic choice of prog's
pick (h, \vec{x}, σ)	non-deterministic decision-theoretic choice of arg's

Figure 1. Some of READYLOG's constructs

constructs given in Fig. 1. Besides standard GOLOG constructs, READYLOG also features non-standard constructs such as **pproj**, where a program is probabilistically projected into the future, or the non-deterministic decision-theoretic choices of programs or arguments (“ \mid ” and **pick**, respectively). These constructs are used inside a **solve** statement and leave choices open that are filled by the decision-theoretic forward search algorithm deployed in READYLOG which we will explain next. The other constructs of READYLOG are shown in Fig. 1.

As we are extending decision-theoretic planning (DTP) in GOLOG in this paper, we now have a closer look at the forward-search DTP algorithm that was proposed by Boutilier et al.² The search tree is expanded in a forward direction induced by the basic action theory. Fig. 2 shows the principle. The nodes in the search tree are expanded and the values are propagated back to the root. The path with the highest value represents the optimal policy. As the tree is constructed based on the basic action theory, it is particularly easy to restrict the search.

2.2. An Example in the Maze Domain

In order to illustrate decision-theoretic planning let us consider a simple example in a maze domain. It is basically a grid world with several rooms, where a robot can move around. An example for a maze domain is illustrated in Fig. 3.

Consider an agent in such a maze. To compute a path from its starting position S to a goal position G it could simply use the program shown in Algorithm 1, assuming the robot has the following action set $A = \{go_right, go_left, go_up,$

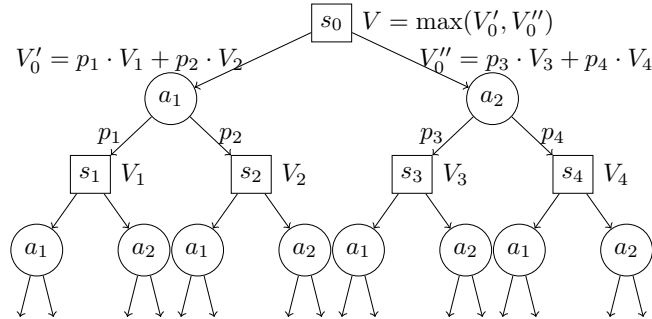


Figure 2. Decision tree search in READYLOG representing a non-deterministic choice of actions. In each situation s_i the agent may choose between a_1 or a_2 . The resulting optimal policy is represented by the branch with maximal value V .




Figure 3. The Maze66 domain from Hauskrecht et al.¹³

Algorithm 1: Decision-theoretic path planning in READYLOG

```

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

```

go_down }. The actions are stochastic, that is to say that there exists a probability distribution over the effects of the action. Each of these actions takes the robot to the intended direction with a high probability, with a low probability it will end up in an adjacent location. The goal state, position G , has a positive reward while each other field has a negative reward. READYLOG now computes the optimal policy (shortest path) from the start to the goal employing forward-search DTP. The forward-search DTP algorithm is implemented in terms of a number of *BestDo* predicates. For non-deterministic choices of actions, its formal definition following

Boutilier et al.² 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:

$$\begin{aligned}
& \text{BestDo}(\mathbf{pick}(x, \tau, p; p'), s, h, \pi, v, pr) \stackrel{\text{def}}{=} \\
& \text{BestDo}((p|_{c_1}^x \mid \dots \mid p|_{c_n}^x); p', s, h, \pi, v, pr)
\end{aligned} \tag{3}$$

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.

2.3. Fuzzy Golog

We introduced the notion of fuzzy fluents in GOLOG earlier.^{5,20} 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 order to make this paper self-contained, in the following we restate the definitions from our earlier work.^{5,20}

2.3.1. Fuzzy Fluents

We introduce two new sorts to the situation calculus: *real* and *linguistic*. We assume the standard interpretation of reals together with the usual operations and ordering relations. *Linguistic* terms are a finite set of constant symbols c_1, \dots, c_k in the language. They refer to qualitative classes; examples are *close* or *far*. We further require a unique names assumption for these linguistic categories.

Now, having introduced reals and linguistic terms into the language of the situation calculus, we can define the degree of membership of a particular value to a given category. For ease of notation we assume that the domain of a particular category is from the domain of real numbers. In general, the domain can be defined arbitrarily.

Definition 1. Let c_1, \dots, c_k be categories of sort *linguistic*. We introduce a relation $\mathfrak{F} \subseteq \text{linguistic} \times \text{real} \times [0, 1]$ relating each linguistic term c of the domain, a real number, and a degree of membership in the category c as

$$\begin{aligned}
& \forall c, u, \mu_u. \mathfrak{F}(c, u, \mu_u) \equiv \\
& (c = c_1 \supset u = u_{c_1,0} \wedge \mu_u = \mu_{c_1,0} \vee \dots \vee u = u_{c_1,m_1} \wedge \mu_u = \mu_{c_1,m_1}) \vee \dots \vee \\
& (c = c_k \supset u = u_{c_k,0} \wedge \mu_u = \mu_{c_k,0} \vee \dots \vee u = u_{c_k,m_k} \wedge \mu_u = \mu_{c_k,m_k}),
\end{aligned}$$

where all $u_{c_i,j}$ and $\mu_{c_i,j}$ are constants of sort real and $\mu_{c_i,j} \in [0,1]$ respectively, i.e. $\forall c, u, \mu_u. \mathfrak{F}(c, u, \mu_u) \supset 0 \leq \mu_u \leq 1$. To ensure that, for each category, each pair (u, μ_u) is unique, we require unique names for linguistic terms: $\forall c \exists u, \mu_u \forall \mu_{u'} . \mathfrak{F}(c, u, \mu_u) \wedge \mathfrak{F}(c, u, \mu_{u'}) \supset \mu_u = \mu_{u'}$. We further require one of the $u_{c_i,j}$ to equal the centre-of-gravity of the respective category, i.e. $u_{c_i,j} = cog(c_i)$ (cf. Def. 4). ■

Note that the above definition yields a formalization of discrete fuzzy sets. That means that all value-membership pairs belonging to a particular linguistic category are enumerated. While we regard a discrete formalization here, the implementation may make use of a continuous formulation of fuzzy sets. Further note that variables occurring free in the logical sentences are implicitly universally quantified in the following definitions.

Definition 2. A fuzzy fluent f is a functional fluent restricted to take only values from sort *linguistic* or from sort *real*. We write $f(\vec{x}, s)$ to refer to a fuzzy fluent, and $f(\vec{x}, s)$ to refer to a non-fuzzy fluent. ■

2.3.2. Querying a Fuzzy Fluent

To query whether or not a fluent value belongs to a certain category, we introduce, similar to fuzzy control theory, predicates is , $is_{\mathfrak{C}}$, is_{\star} , and is_{\oplus} . These predicates are true if a fuzzy fluent value belongs to the category in question to a non-zero degree.

Definition 3.

- (1) To query if a fuzzy fluent belongs to a given category, we define the predicate $is \subseteq real \times linguistic$ as

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

- (2) Similarly, we define $is_{\mathfrak{C}} \subseteq real \times linguistic$, to know if a fuzzy fluent does *not* belong to a certain category

$$is_{\mathfrak{C}}(f(\vec{t}, \sigma), \gamma) \doteq \neg \exists u, \mu_u. f(\vec{t}, \sigma) = u \wedge \mathfrak{F}(\gamma, u, \mu_u) \vee \\ \exists u, \mu_u. f(\vec{t}, \sigma) = u \wedge \mathfrak{F}(\gamma, u, \mu_u) \wedge \mu_u < 1.$$

A fluent value does not belong to a certain category, if either the value in question is not defined in terms of a fuzzy set, or the value exists and its degree of membership is less than 1. That is because the standard definition of the fuzzy complement is given by $\neg \mu_u = 1 - \mu_u$. This evaluates to a non-zero value for all values of μ_u less than 1.

- (3) For complex queries, for example if a fuzzy fluent value belongs to several overlapping categories at the same time, we define a predicate $is_{\star} \subseteq real \times (linguistic)^n$ for arbitrary n as

$$is_{\star}(f(\vec{t}, \sigma), \gamma_0, \dots, \gamma_n) \doteq \exists u, \mu_{u,0}, \dots, \mu_{u,n}. f(\vec{t}, \sigma) = u \wedge \mathfrak{F}(\gamma_0, u, \mu_{u,0}) \\ \wedge \dots \wedge \mathfrak{F}(\gamma_n, u, \mu_{u,n}) \wedge (\mu_{u,0} \star \dots \star \mu_{u,n} > 0).$$

- (4) Similarly, for asking whether or not a fuzzy fluent value belongs to one category or the other, we introduce the predicate $is_{\oplus} \subseteq real \times (linguistic)^n$

$$is_{\oplus}(\mathfrak{f}(\vec{t}, \sigma), \gamma_0, \dots, \gamma_n) \doteq \exists u, \mu_{u,0}, \dots, \mu_{u,n}. \mathfrak{f}(\vec{t}, \sigma) = u \wedge \mathfrak{F}(\gamma_0, u, \mu_{u,0}) \\ \wedge \dots \wedge \mathfrak{F}(\gamma_n, u, \mu_{u,n}) \wedge (\mu_{u,0} \oplus \dots \oplus \mu_{u,n} > 0).$$

■

Next, we need to define a defuzzifier, a function that computes a single numerical value for a given linguistic category. We need such a defuzzifier because then in the mechanics underlying our situation calculus machinery we can use crisp representatives for our linguistic terms. Note that, while we choose the centre-of-gravity defuzzifier here, our approach is not restricted to this. Instead, any defuzzifier could be used just as well. Then, we define a defuzzifying function that selectively applies the defuzzifier to any linguistic term. This enables us to transparently make use of linguistic terms since they are transformed to numerical representatives whenever necessary.

Definition 4. Let τ be a term of $\mathcal{L}_{sitcalc}$. We define a function $defuzz$ inductively as:

- (1) if τ is an atomic term
 - (a) and τ is of sort *linguistic*, then $defuzz(\tau) = cog(\tau)$
 - (b) otherwise $defuzz(\tau) = \tau$
- (2) if τ is a non-atomic term of the form $f(\vec{t})$ with $\vec{t} = t_1, \dots, t_n$, then $defuzz(\tau) = f(defuzz(t_1), \dots, defuzz(t_n))$

In $defuzz$ we make use of the function cog , which defines the centre-of-gravity defuzzifier. It is defined as:

$$cog(c) = \hat{u} \equiv \\ \exists u_0, \dots, u_k, \mu_{u_0}, \dots, \mu_{u_k}. \mathfrak{F}(c, u_0, \mu_{u_0}) \wedge \dots \wedge \mathfrak{F}(c, u_k, \mu_{u_k}) \wedge \\ u_0 \neq \dots \neq u_k \wedge \forall u^*, \mu^*. (u^* \neq u_0 \wedge \dots \wedge u^* \neq u_k \wedge \\ \mu^* \neq \mu_{u_0} \wedge \dots \wedge \mu^* \neq \mu_{u_k} \supset \neg \mathfrak{F}(c, u^*, \mu^*)) \wedge \\ \hat{u} = \frac{\sum_{i=0}^k u_i \cdot \mu_{u_i}}{\sum_{i=0}^k \mu_{u_i}}$$

■

A fuzzy fluent's function value is eventually substituted by its defuzzified value when applying the defuzzifying function $defuzz$. Note that the number k in the definition of the centre-of-gravity defuzzifier above refers to the number of all value-membership pairs defined in the fuzzy set for the linguistic categories plus one value for each category itself (Def. 1). Further note that the number of value-membership pairs is required to be finite. As the above definition of a defuzzifier is not closed




Figure 4. The one-dimensional domestic robot world.


under division in general, note that the definition is however well-defined. This is because we postulate that the centre-of-gravity for a qualitative category will be added to the set explicitly. In our implementation, where we make use of continuous fuzzy sets, this requirement can be dropped, as the set then is closed under division.

By now, we defined fuzzy fluents as a specialization of functional fluents operating on reals and linguistic terms, introduced qualitative categories as constants of sort linguistic, and defined a fuzzy set in our domain axiomatization which allows for defining which values make up a qualitative category to which degree. We can further query whether or not a fuzzy fluent belongs to a qualitative category. Moreover, we can ask if a fuzzy fluent belongs to several categories at the same time, or if it belongs to the complementary category. In the next subsection, we will show how these fluents can be used for reasoning in the Situation Calculus.


2.4. Reasoning with Linguistic Terms: A One-dimensional Example

To illustrate reasoning with qualitative positional information using linguistic terms and the representations introduced above, consider the following simple example. A robot is situated in a one dimensional room with a length of ten metric units as depicted in Fig. 4. To keep things simple, we restrict ourselves to integer values for positions in the following. We have one single action called $gorel(d)$ denoting the relative movement of d units of the robot in its world. For sake of simplifying the notation in this example, we assume that this action is always possible, i.e. $Poss(gorel(d), s) \equiv \top$. The action has impact on the fluent pos which denotes the absolute position of the robot in the world. The position of the table is defined by the macro $pos_{table} = p \doteq p = 9$. In the initial situation, the robot is located at position 0, i.e. $pos(S_0) = 0$. We partition the distance in categories *close*, *medium*, and *far*, and introduce qualitative categories for the position of the robot as *back*, *middle*, and *front*. We give the (fuzzy) definition of those categories below, where we use (u_i, μ_{u_i}) as an abbreviation for $u = u_i \wedge \mu = \mu_{u_i}$.

For readability reasons, we assume in this example that the robot can only move around in integer steps. Restricting to integers requires to use an altered version $cog'(c)$ of the centre-of-gravity defuzzifier formula: $cog'(c) \doteq \lfloor cog(c) \rfloor$. Of course our function $defuzz$ has to mention cog' instead of cog then. A graphical illustration of the membership functions for position and distance is given in Fig. 5.



(a) Qualitative positions of a robot in a one-dimensional world.



(b) Qualitative distance in the one-dimensional world.

Figure 5. Membership functions for position and distance in our one-dimensional robot domain.

$$\begin{aligned} \mathfrak{F}(\text{position}, u, \mu_u) \equiv & \\ & (\text{position} = \text{back} \supset (0, 0.25) \vee (1, 0.75) \vee (2, 0.75) \vee (3, 0.25) \vee (3/2, 0.5)) \vee \\ & (\text{position} = \text{middle} \supset (3, 0.25) \vee (4, 0.75) \vee (5, 0.75) \vee (6, 0.25) \vee (9/2, 0.5)) \vee \\ & (\text{position} = \text{front} \supset (6, 0.25) \vee (7, 0.75) \vee (8, 0.75) \vee (9, 0.25)), \end{aligned}$$

while the distances can take the values

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

Suppose the robot's position in situation S_0 is characterized by the linguistic term *back* and the table is located at position 9, i.e. $\mathcal{D}_{S_0} = \{\text{pos}(S_0) = \text{back}, \text{dist}(S_0) = 9\}$. Suppose now that the robot travels 4 units to the right. Then we can show that

$$\mathcal{D} \models \text{is}(\text{pos}(\text{do}(\text{gorel}(4), S_0)), \text{middle}) \wedge \text{is}(\text{dist}(\text{do}(\text{gorel}(4), S_0)), \text{medium}).$$

Using regression and the successor state axiom for the fluent *dist* we apply the centre-of-gravity $\text{cog}'(\text{back}) = 1$ if the value of a fuzzy fluent is a linguistic term in the initial situation. It thus holds in S_0 that $\mathcal{D} \models \text{is}(\text{dist}(S_0), \text{far})$. By performing

the action $\text{gorel}(4)$ the robot moves four positions to the right. The proposition holds because $\text{pos}(\text{do}(\text{gorel}(4), S_0)) = 1 + 4 = 5$ and $\mathfrak{F}(\text{middle}, 5, 0.75)$ has a non-zero membership value. The quantitative distance from 5 to 9 equals 4 units or medium distance, as is given by $\mathfrak{F}(\text{medium}, 4, 0.75)$.

Suppose now that the robot's control program contains the action $\text{gorel}(\text{far})$ mentioning the qualitative term far . At which position will the robot end up in situation $s = \text{do}(\text{gorel}(\text{far}), S_0)$? It follows that

$$\mathcal{D} \models \text{is}(\text{pos}(\text{do}(\text{gorel}(\text{far}), S_0)), \text{front})$$

i.e. the robot ends up in the front part of its world after executing $\text{gorel}(\text{far})$. Determining the robot's position in situation $\text{do}(\text{gorel}(\text{far}), S_0)$ we again use regression. It is sufficient to show that $\mathcal{D}_{S_0} \models \text{is}(\mathcal{R}[\text{pos}(\text{do}(\text{gorel}(\text{far}), S_0))], \text{front})$ which is—according to the successor state axiom above—regressed to $\text{pos}(S_0) = \text{cog}'(\text{back}) \wedge \mathfrak{F}(\text{far}, 7, 0.75) \wedge d' = \text{cog}'(\text{far}) \wedge \text{is}(y = \text{cog}'(\text{back}) + \text{cog}'(\text{far}), \text{front}) \equiv \text{is}(y = 1 + 7, \text{front}) \equiv y = 8 \wedge \mathfrak{F}(\text{front}, 8, 0.75) \wedge 0.75 > 0$. Hence, we can infer that the robot ends up at position front .

Assume that apart from $\text{gorel}(x)$ there is another action $\text{go}(x)$ which makes the robot move directly to position x . The successor state axiom of $\text{go}(x)$ is given as $\text{pos}(\text{do}(a, s)) = y \equiv a = \text{go}(x) \wedge y = x \vee a \neq \text{go}(x) \wedge y = \text{pos}(s)$. What happens if we put in a qualitative category there, i.e. at which position will the robot end up in situation $s = \text{do}(\text{go}(\text{front}), S_0)$? It turns out that we have

$$\mathcal{D} \models \text{is}(\text{pos}(\text{do}(\text{go}(\text{front}), S_0)), \text{front})$$

i.e. the robot ends up in the front part of its world after executing $\text{go}(\text{front})$. When regressing a formula that contains a linguistic term, the defuzzification function (e.g. centre-of-gravity $\text{cog}'(c)$) is applied if the result of a previous successor state axiom assigned a qualitative term to the fuzzy fluent. Then, $\mathcal{D} \models \text{is}(\text{pos}(\text{do}(\text{go}(\text{front}), S_0)), \text{front})$ iff $\mathcal{D}_{S_0} \models \text{is}(\mathcal{R}[\text{pos}(\text{do}(\text{go}(\text{front}), S_0))], \text{front})$ which is regressed to $\text{pos}(S_0) = \text{cog}'(\text{back}) \wedge x = \text{front} \wedge u = \text{cog}'(\text{front}) \wedge \text{is}(u, \text{front}) \equiv \text{is}(u = 7, \text{front}) \equiv u = 7 \wedge \mathfrak{F}(\text{front}, 7, 0.75) \wedge 0.75 > 0$. Thus, it can be inferred that the robot will end up at position front .

3. Extending DT-Planning in Golog with Fuzzy Notions

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 **pick** 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.




Figure 6. The Diner Domain

3.1. 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 **pickF** which takes a fuzzy expression instead of the regular set of the classical **pick**.

The idea is that instead of giving a finite set of variable or fluent values in the **pick** 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 of a waitron agent in a diner cafe (see Fig. 6), we could simply state to choose a coffee whose temperature is *hot*. What the **pickF** statement does is to translate this into a set of temperatures with positive membership values for the category *hot* when served. In our case shown in Fig. 7, 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(T₉)* (serve a coffee at table 9) with the **pickF** statement.

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

$$\begin{aligned}
 & \text{BestDo}(\mathbf{pickF}(f : \gamma, p); p', s, h, \pi, v, pr) \stackrel{\text{def}}{=} \\
 & \exists u_1, \dots, u_k. \\
 & \bigwedge_{u \in \{u_1, \dots, u_k\}} [\text{is}(u, \gamma)] \wedge \bigvee_{u \in \{u_1, \dots, u_k\}} [f[s] = u] \wedge \\
 & \forall u'. (u' \neq u_1 \wedge \dots \wedge u' \neq u_k) \supset \neg(\text{is}(u', \gamma)) \wedge \\
 & \text{BestDo}((p|_{u_1}^f \mid \dots \mid p|_{u_k}^f); p', s, h, \pi, v, pr)
 \end{aligned}$$




Figure 7. Coffee temperature membership function

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 is predicate holds for k numerical values u_i of the category γ . Also, the value of our fluent f must be one of those k values. Then, we make sure that the k values are all values for which $\text{is}(\cdot, \gamma)$ holds. Lastly, we call BestDo , replacing the fluent f in the program p with any such value (denoted by $p|_{u_i}^f$).^b This is analogous to the definition of BestDo for **pick** (Eq. 3), where the fluent was replaced with any element of the set τ (cf. Section 2.1). 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

$$\begin{aligned} \text{BestDo}(\mathbf{pickF}(f : \neg\gamma, p); p', s, h, \pi, v, pr) &\stackrel{def}{=} \\ \exists u_1, \dots, u_k. & \\ \bigwedge_{u \in \{u_1, \dots, u_k\}} [\text{is}_{\mathbf{C}}(u, \gamma)] \wedge \bigvee_{u \in \{u_1, \dots, u_k\}} [f[s] = u] \wedge & \\ \forall u'. (u' \neq u_1 \wedge \dots \wedge u' \neq u_k) \supset \neg(\text{is}_{\mathbf{C}}(u', \gamma)) \wedge & \\ \text{BestDo}((p|_{u_1}^f \mid \dots \mid p|_{u_k}^f); p', s, h, \pi, v, pr) & \end{aligned}$$

For a *conjunction* of n linguistic categories we have

$$\begin{aligned} \text{BestDo}(\mathbf{pickF}(f : \Gamma_{\star}, p); p', s, h, \pi, v, pr) &\stackrel{def}{=} \\ \exists u_1, \dots, u_k. & \\ \bigwedge_{u \in \{u_1, \dots, u_k\}} [\text{is}_{\star}(u, \Gamma_{\star})] \wedge \bigvee_{u \in \{u_1, \dots, u_k\}} [f[s] = u] \wedge & \\ \forall u'. (u' \neq u_1 \wedge \dots \wedge u' \neq u_k) \supset \neg(\text{is}_{\star}(u', \Gamma_{\star})) \wedge & \\ \text{BestDo}((p|_{u_1}^f \mid \dots \mid p|_{u_k}^f); p', s, h, \pi, v, pr) & \end{aligned}$$

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

^bNote 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.

categories we have

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

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 **pickF** 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.

3.2. 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 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

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

For the complement of a single linguistic category

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

For the conjunction of n linguistic categories

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

For the disjunction of n linguistic categories

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

4. An Application in the Diner Domain

As an application domain for our newly introduced extensions we now consider the Diner Domain, an illustration of which is shown in Fig. 6. 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 .

In the 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 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. 7. 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.

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_{\complement} for the complement, is_{\star} for the conjunction, and is_{\oplus} for the disjunction of fuzzy fluents. See^{5,20} for the formal definitions.

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:

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

For the y -coordinate we introduce a fuzzy fluent $\text{pos}Y$ 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.

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 agent 5 seconds. Fig. 6 shows an example. The agent (*A*) needs 12 actions (*r, r, u, u, u, u, u, u, u, u, u, r*) to get to table T_9 . It takes the agent 60 seconds to reach that table. That means that a hot coffee at 65 centigrades will be lukewarm when served at table T_9 . The `deliverCoffee` action finally delivers the coffee to the customer, once the right table has been reached by the agent.

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 **pickF** 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 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

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     pickF(posX, left ∨ right ) {
5       pickF(posY, front ∨ back ) {
6         tableWithOpenOrder(table, posX, posY);
7         pickF(mealTemp, luke ∨ hot ) {
8           mealWithTempReady(meal, mealTemp);
9           load_meal(meal, tray);
10          bring_meal(tray, table);
11          serve_meal(tray, table); } } }
12    navigate(counter);
13  endwhile
14 endproc

```

as follows. For simplicity, we limit ourselves to a list of only two orders.

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

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

Let us assume, the robot has orders from tables T_0 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 **pickF** 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 **pickF**-statement we collect possible *y*-coordinates 1, 2, 3, 7, 8, 9. Using the *BestDo* definition, we replace in the body of the **pickF** 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 *tableWithOpenOrder*. The only positions

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

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 **pickF**, now for the fluent `mealTemp`. Again, using 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. 4). Hence, we consider to execute the sequence inside the innermost **pickF** 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 **pickF** 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 presented an extension to READYLOG that 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 previous work, we defined a predicate “is” to test this. With this predicate, we can handle negation, conjunction, and disjunction, respectively, of linguistic terms. With these predicates we can check, whether a fluent or its complement belongs to a certain category, if a fluent belongs to one or to another more categories a , or if a fluent is a member of two or more categories at the same time. 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 **pick** statements. To this end, we introduced and defined a statement **pickF** that translates the values for which the given fuzzy fluent formula holds as an argument set for the ordinary **pick** 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.

Designing the reward function of a decision-theoretic agent is not an easy task and requires good domain knowledge. Using fuzzy categories in the reward function allows for a more natural formulation of the reward function and thus alleviates its design. Furthermore, they make it more robust against inevitable variations occurring in the real world as linguistic categories can overlap. This way some form of hysteresis in taking decisions is possible and may help avoiding oscillations in the resulting behavior. Our approach changes nothing about the complexity of decision-theoretic planning per se and hence, does not contribute to making it more scalable. However, with making the fuzzy notions available to the agent designer we facilitate

constraining the search space more conveniently since we are using the forward-search DT planning algorithm.

Our further steps are as follows. So far, we do 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 to planning with preferences is on our agenda. Here we want to investigate how our work can be married with the work of Fritz and McIlraith¹⁰ and Bienvenu et al.¹ who compile modal logic preference formulas into GOLOG programs. Similarly, we will have a look into the works by Finzi and Pirri⁹ and Cesta et al.³ who use temporal interval planning to solve scheduling problems.

References

1. Meghyn Bienvenu, Christian Fritz, and Sheila A McIlraith. Specifying and computing preferred plans. *Artificial Intelligence*, 175(7):1308–1345, 2011.
2. 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.
3. Amedeo Cesta, Simone Fratini, Andrea Orlandini, Alberto Finzi, and Enrico Tronci. Flexible plan verification: Feasibility results. *Fundamenta Informaticae*, 107(2):111–137, 2011.
4. 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.
5. A. Ferrein, S. Schiffer, and G. Lakemeyer. A Fuzzy Set Semantics for Qualitative Fluents in the Situation Calculus. In Caihua Xiong, Honghai Liu, Yongan 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.
6. Alexander Ferrein. Robot controllers for highly dynamic environments with real-time constraints. *Künstliche Intelligenz*, 24(2):175–178, 2010.
7. 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.
8. Alexander Ferrein, Stefan Schiffer, and Gerhard Lakemeyer. Embedding Fuzzy Controllers in Golog. In *Proceedings of the IEEE International Conference on Fuzzy Systems (FUZZ-IEEE'09)*, pages 498–509. IEEE Press, 2009.
9. Alberto Finzi and Fiora Pirri. Flexible interval planning in concurrent temporal Golog. In *Working notes of the 4th international cognitive robotics workshop*, 2004.
10. Christian Fritz and Sheila A. McIlraith. Decision-theoretic golog with qualitative preferences. In *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 153–163, 2006.
11. H. Grosskreutz. Probabilistic projection and belief update in the pgolog framework. In *Proc. of the 2nd International Cognitive Robotics Workshop (CogRob-00) at the European Conference on Artificial Intelligence (ECAI-00)*, pages 34–41. 2000.
12. Henrik Grosskreutz and Gerhard Lakemeyer. On-line execution of cc-Golog plans. In Bernhard Nebel, editor, *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, IJCAI 2001, Seattle, Washington, USA*. Morgan Kaufmann, August 4–10 2001.
13. M. Hauskrecht, N. Meuleau, L. P. Kaelbling, T. Dean, and C. Boutilier. Hierarchical solution of Markov decision processes using macro-actions. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 1998.
14. 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.
15. John McCarthy. Situations, Actions, and Causal Laws. Technical Report Memo 2, AI Lab, Stanford University, California, USA, 3 July 1963.
16. John McCarthy and Patrick J. Hayes. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In Bernard Meltzer and Donald Michie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.
17. R. Reiter. *Knowledge in Action*. MIT Press, 2001.

18. Stefan Schiffer and Alexander Ferrein. Decision-Theoretic Planning with Linguistic Terms in Golog. In Irene Díaz, Anca Ralescu, and Stefan Schiffer, editors, *Proceedings of the Workshop on Fuzzy Logic in AI (FLinAI) 2015, co-located with the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015)*, July 25 2015.
19. 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.
20. 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.
21. Stefan Schiffer, Alexander Ferrein, and Gerhard Lakemeyer. Reasoning with Qualitative Positional Information for Domestic Domains in the Situation Calculus. *Journal of Intelligent and Robotic Systems. Special Issue on Domestic Service Robots in the Real World.*, 66(1–2):273–300, 2012.
22. 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.