

## Article

# Escaping Local Minima via Appraisal Driven Responses

Malte Rørnøse Damgaard \* , Rasmus Pedersen  and Thomas Bak 

Department of Electronic Systems, Automation and Control, Aalborg University, 9220 Aalborg, Denmark

\* Correspondence: mrd@es.aau.dk

**Abstract:** Inspired by the reflective and deliberative control mechanisms used in cognitive architectures such as SOAR and Sigma, we propose an alternative decision mechanism driven by architectural appraisals allowing robots to overcome impasses. The presented work builds on and improves on our previous work on a generally applicable decision mechanism with roots in the Standard Model of the Mind and the Generalized Cognitive Hour-glass Model. The proposed decision mechanism provides automatic context-dependent switching between exploration-oriented, goal-oriented, and backtracking behavior, allowing a robot to overcome impasses. A simulation study of two applications utilizing the proposed decision mechanism is presented demonstrating the applicability of the proposed decision mechanism.

**Keywords:** cognitive robotics; cognitive architecture; appraisals; reflective control; deliberate control; reactive control; variational inference; deadlocks; probabilistic programming idiom; the standard model of the mind



**Citation:** Damgaard, M.R.; Pedersen, R.; Bak, T. Escaping Local Minima via Appraisal Driven Responses. *Robotics* **2022**, *11*, 153. <https://doi.org/10.3390/robotics11060153>

Academic Editor: Thierry Chaminade

Received: 3 November 2022

Accepted: 14 December 2022

Published: 16 December 2022

**Publisher's Note:** MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Copyright:** © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Robotic technology has immense potential to change our daily life. In the industry, human-robot co-working is envisioned to play a key role in the next industrial revolution known as Industry 5.0 [1]. In healthcare robots also see increasing usage e.g., in personalized healthcare for providing assistance to patients, and the elderly [2,3], and during the COVID-19 pandemic, robots were deployed to disinfect common spaces, such as supermarkets and hospitals [4]. Common to the above is the increased need for autonomous robots that can safely and naturally interact with humans while solving different abstractly and/or vaguely defined tasks. Due to the uncertainties in such problems, pure goal-driven problem-solving architectures will often end up in local minima in the problem formulation also known as impasses. I.e., situations where the information or action selection strategy currently available to the robot is insufficient to solve the task. Thus, one core faculty of such robotic systems should be the ability to reflect on the current situation to timely deviate from one action selection strategy to try out other strategies or to retrieve new information about the task.

The next generation of cognitive architectures, based on modern machine learning techniques, has the potential to revolutionize robotics by allowing roboticists to develop such autonomous systems easily. In previous work, we proposed the Generalized Cognitive Hour-glass Model constituting a framework for developing cognitive architectures by composing them from generally applicable probabilistic programming idioms over which powerful general algorithms can perform inference [5]. The idiomatic approach to composing cognitive architectures, encouraged by this framework, allows researchers and practitioners to more easily cooperate by mixing and matching probabilistic programming idioms developed by others while being able to handcraft parts of a system for which current solutions do not suffice.

In another work, we proposed one such probabilistic programming idiom based on the “standard model of the mind” [6] for the task of Active Knowledge Search (AKS) in unknown environments [7]. This idiom defines a probabilistic decision process that

encourages a robot to take actions to discover, i.e., obtain information about, its environment based purely on notions of progress and information gain while avoiding constraint violations. Simulations applying this idiom to the specific problem of active mapping and robot exploration showed promising results. However, limitations were also identified. The main limitation was that in specific situations the simulated robot would get “stuck” in an impasse taking repetitive actions yielding no new information about the environment, thus hindering full exploration of the environment. As we will discuss in more detail in Section 3, this is essentially caused by the fixed strategy for action selection employed by the previous solution.

In the literature related to robot navigation, impasse phenomena are commonly known as “the local minima issue” [8], “deadlocks” [9], “limit cycles” [10], “infinite loops” [11], “dead ends”, “cyclic dead ends”, or “trap-situations” [12]. Like the problem mentioned above, all of these terms refer to situations in which a fixed strategy for action selection results in no meaningful progress towards a goal state or compared to a measure of optimality. To resolve these situations solutions proposed by researchers within robotics usually rely on problem-specific information, e.g., geometric properties, to detect and/or resolve the impasse. As an example consider the approach used in [13] where a grid map is defined over the workspace with a counter attached to each of the cells keeping track of the number of times a given cell has been visited. Whenever this counter reaches a predefined threshold, it is registered as a limit cycle. When a limit cycle is detected a temporary way-point is generated, guiding the robot out of the enclosure causing the limit cycle. Finally, when the robot gets outside the enclosure, a virtual wall is generated, ensuring that the robot does not enter the problematic enclosure again. As another example consider the approach used in [14] where deadlock loops are detected based on the periodicity of the distance to the goal. Similarly, in [15], deadlocks are detected based on a preferred velocity magnitude, the actual velocity magnitude, and the unsigned distance between robots. While the solutions suggested above might work for specific problems, they do not easily generalize to other problems.

In the literature related to cognitive architectures, similar phenomena in which an agent is unable to make progress with the information that is currently available are often referred to as impasses [16,17]. Research in cognitive architectures is mainly focused on developing systems with generally applicable capabilities opposite to the problem-specific solutions commonly proposed in the robotics literature. This is done by taking substantial inspiration from theories about the workings of animal and human cognition developed within cognitive science, based on which computational instantiations are proposed. Nevertheless, solutions to tackle impasse phenomena seem to follow the same pattern as those proposed by researchers in robotics. First, systems are made able to detect impasses. Secondly, systems are induced with some reflective mechanism that based on the detected impasse can choose appropriate temporary decision strategies until the impasse has been resolved. E.g., consider the tri-level control structures implemented by two of the most prominent cognitive architectures, SOAR [16] and Sigma [17]. These tri-level control structures consist of a *reflective control* mechanism, a *deliberative control* mechanism, and a *reactive control* mechanism, each of which activates the control layer above it based on the detection of impasses. These control structures assume that there a priori exists a discrete/symbolic set of operators that re-actively can be either sorted out or proposed for further evaluation, thereby making the detection of impasses straightforward without any problem-specific knowledge. This approach has some clear benefits with respect to attention, i.e., the effective allocation of limited (computational) resources. Since each of the layers focuses computations on the information actually needed to solve a given problem in a specific context/state, a lot of computations are saved. However, the approach also raises some difficulties in robotics, where a lot of the low-level control is more naturally described by means of continuous variables. As an example, consider the position control of a robot. In SOAR and Sigma, such controls are usually abstracted to symbolic representations such as “walk towards target object”, “run towards target object”, “pick up target object”

or “walk towards random object” [17]. These symbolic representations then have to be decoded by an extra module external to the decision process, called the *motor buffer*, before they can be manifested in the environment. This layer of abstraction makes it hard to incorporate uncertainties resulting from low-level control into the decision process, simply due to the inevitably loss of information happening when trying to lump low-level controls, e.g., motor currents or positions, that are continuous in nature together by coarse discretized compound control in the form of symbolic representations. In the end, this will result in less optimal responses being picked since the fine level of control needed within robotics cannot be accounted for as an intrinsic part of the decision process.

For these reasons, our intention with this paper is to present our recent efforts toward implementing general reflective mechanisms similar to the ones found in cognitive architectures within the scope of the framework proposed in [5] in a way that is suitable for robotic applications. The main contributions of this paper are:

- A description and implementation of a control structure grounded in stochastic variational inference that is capable of deliberate and reflective control based on architectural appraisals allowing for the incorporation of uncertainties resulting from low-level control.
- A demonstration of how such a control structure overcomes the limitations of the probabilistic programming idiom previously proposed in [7].
- A demonstration of how such general control structure compares to problem-specific approaches commonly used in robotics.
- A discussion of the time complexity of the proposed control structure.

This paper is organized as follows: Section 2 introduces the notation used within this paper. Section 3 shortly describes the previously proposed probabilistic programming idiom in more detail together with the impasse phenomenon observed. Modifications and extensions to the previously proposed idiom are presented in Sections 4 and 5. Simulation results utilizing the modifications are provided in Section 6. Section 7 concludes the paper, and potential future directions are given in this section.

## 2. Preliminaries

As in paper [7] we use the following notation.  $X$  is used to denote observed variables,  $Z$  is used to denote latent variables, and  $C$  is used to denote a collection of both types of variables. A superscript in curly brackets is used to indicate the index of a variable. For time indexes, the set of indexes of future variables is indicated as  $\{t\}^+ = \{t+1, \dots, t+T\}$ . Similarly, the set of indexes of past variables is indicated as  $\{t\}^- = \{t-T, \dots, t\}$ . Furthermore, within this paper the following approximate “probabilistic logic” is used,

$$\begin{aligned} p(z \in \bar{z} \vee y \in \bar{y}) &\stackrel{\text{def}}{=} p(z \in \bar{z}) + p(y \in \bar{y}) - p(z \in \bar{z})p(y \in \bar{y}) \\ p(z \in \bar{z} \wedge y \in \bar{y}) &\stackrel{\text{def}}{=} p(z \in \bar{z}) \cdot p(y \in \bar{y}) \\ p\left(\bigwedge_{i=1}^I z^{\{i\}} \in \bar{z}^{\{i\}}\right) &\stackrel{\text{def}}{=} \prod_{i=1}^I p(z^{\{i\}} \in \bar{z}^{\{i\}}) \end{aligned}$$

where  $\wedge$  and  $\vee$  denotes an approximate *and* and *or* operation, respectively. These approximate “probabilistic logic” rules constitute a probabilistic intersection and union with an independence assumption implied, respectively.

## 3. Related Work

As stated in Section 1, the probabilistic programming idiom proposed in [7] defines a probabilistic decision process for Active Knowledge Search in unknown environments, based on the “standard model of the mind” [6]. This was done by first defining a probabilistic model relating the previous content of working memory,  $Z_{WM}^{\{t\}-}$ , with the future content,  $C_{WM}^{\{t\}+}$ , while taking variables stored in the long-term memory,  $Z_{LTM}$ , into account. In [7],

the working memory was further sub-divided into variables relating to motoric actions i.e., the *motor buffer*,  $Z_{Mb}$ , variables related to the *perceptual buffer*,  $Z_{Pb}$ , *State variables*,  $Z_s$ , representing the state of the agent itself, the environment, and *decision variables*  $C_D^{\{t\}^+}$ . From this a probabilistic decision model with the factorization were derived,

$$p\left(C_{WM \setminus b}^{\{t\}^+}, Z_{Mb}^{\{t-1\}^+}, Z_{Pb}^{\{t\}^+} | Z_{WM \setminus b}^{\{t\}^-}, Z_{LTM}\right) \quad (1)$$

$$\stackrel{def}{=} \left[ \prod_{\tau=t+2}^{t+\bar{T}} p\left(C_D^{\{\tau\}} | Z_s^{\{\tau\}}, Z_{WM \setminus b}^{\{t\}^-}, Z_{LTM}\right) p\left(Z_s^{\{\tau\}} | Z_s^{\{\tau-1\}}, Z_{Mb}^{\{\tau-1\}}\right) p\left(Z_{Mb}^{\{\tau-1\}}\right) \right]$$

$$\cdot p\left(C_D^{\{t+1\}} | Z_s^{\{t+1\}}, Z_{WM \setminus b}^{\{t\}^-}, Z_{LTM}\right) p\left(Z_s^{\{t+1\}} | Z_s^{\{t\}}, Z_{Mb}^{\{t\}}\right) p\left(Z_{Mb}^{\{t\}}\right)$$

where  $Z_{WM \setminus b} = Z_{WM} \setminus \{Z_{Mb}, Z_{Pb}\}$ . Inspired by the work on emotions in [18], a subset of the decision variables,  $x_A$ , was denoted *attention variables*. The purpose of these *attention variables* is to control how the decision process is influenced by the other decision variables: *progress*,  $z_p$ , *information gain*,  $z_i$ , and *constraints*,  $z_c$ , hereafter referred to as *appraisal variables*. In [7] this was done via the fixed relation

$$p\left(x_A^{\{\tau\}} | Z_s^{\{\tau\}}, Z_{WM \setminus b}^{\{t\}^-}, Z_{LTM}\right) \quad (2)$$


$$= \text{Bernoulli}\left(p\left(\left[z_p^{\{\tau\}} = 1 \vee z_i^{\{\tau\}} = 1\right] \wedge z_c^{\{\tau\}} = 1 \mid Z_s^{\{\tau\}}, Z_{WM \setminus b}^{\{t\}^-}, Z_{LTM}\right)\right)$$

which basically states that during the decision process attention should be given to future states that yield progress *or* new information *and* does not violate constraints. Having defined the model in Equation (1) and the relation in Equation (2), Stochastic Variational Inference was used to approximate the posterior over optimal future motoric actions given the attention variables, i.e.,

$$q_{\phi_{Mb}^{\{t-1\}^+}, *}\left(Z_{Mb}^{\{t-1\}^+}\right) \approx p\left(Z_{Mb}^{\{t-1\}^+} | x_A^{\{t\}^+} = 1\right) \quad (3)$$

The above was implemented as an abstract class utilizing the probabilistic programming language Pyro [19], thereby ensuring that the probabilistic programming idiom can be reused in multiple applications by implementing a few abstract methods defined by the abstract class.

To investigate the performance of the idiom, it was used to implement an algorithm for autonomous robot exploration which was simulated on the full HouseExpo dataset [20], containing 35,126 different floor plans. From these simulations, one of the observations was that the robot sometimes would end up taking repetitive actions purely driven by the *progress* appraisal variable. Whereby, the robot would not fully explore its environment as illustrated in Figure 1. In other words, the robot ended up at an impasse. It was further concluded that an alternative to the fixed decision strategy given by Equation (2) would be needed to overcome this problem.



**Figure 1.** A simulated trajectory of using the method presented in [7] for the floor plan with ID “0a1b29dba355df2ab02630133187bfab” from the HouseExpo dataset [20]. The robot keeps driving around in the same room, without exploring the rest of its environment.

#### 4. Overall Idea

Both SOAR and Sigma implement a tri-level control structure, in which the distinction between the *deliberative control* and the *reflective control* mechanisms is architectural rather than conceptual. They both simply comprise a specific architectural response to similar architectural stimuli, i.e., the detection of impasses. When we further consider the statement:

“Work in Sigma on appraisal, and its relationship to attention, has led to the conclusion that the detection of impasses should itself be considered as a form of appraisal” [17].

It hints toward the possibility that similar functionality might be obtained from an architecturally simpler control structure. Instead of treating the detection of and responses to impasses as distinctive architectural mechanisms, we propose that this is treated as *affective responses* to appraisals arising from evaluations of deliberate attention allowing us to incorporate uncertainties from low-level control. As illustrated in Figure 2, our proposal is to have a control structure consisting of a single architectural layer with decisions being the result of three main steps:


- Deliberate Attention Proposal,
- Deliberate Attention Evaluations,
- and Affective Responses.

In the state at time  $t$ ,  $Z_s^{\{t\}}$ , the *Deliberate Attention Proposal* step proposes one or more relevant attention mechanisms each corresponding to a specific attention variable,  $x_A$ , e.g., as in Equation (2). Then the *Deliberate Attention Evaluations* first follows the steps described in Section 3 to estimate the distribution over specific future motoric actions,  $q_\phi^{\{i\}}(Z_{Mb}^{\{t-1\}^+})$ , corresponding to each of the attention mechanisms proposed by the *Deliberate Attention Proposal*. This is indicated with red, yellow, and blue boxes in Figure 2. After inferring the motoric action posteriors, the *Deliberate Attention Evaluations* step evaluates what the expected appraisals would be from effectuating each of motoric actions,

$$E_{C_D \sim p(C_D | Z_{Mb}^{\{t-1\}^+})} q_\phi^{\{i\}}(Z_{Mb}^{\{t-1\}^+})[C_D].$$

This is indicated by the gray boxes in Figure 2. As it is also indicated in Figure 2, both parts of the *Deliberate Attention Evaluations* step requires access to parts of cognition distal

to the decision process. Based on the expected appraisals of each of the motoric action posteriors the last step in the decision process can initiate different *affective responses*, such as proposing additional attention mechanisms to evaluate or effectuating one of the action posteriors, as indicated by the thin arrows in Figure 2.



**Figure 2.** Illustration of the proposed approach with an indication of where we consider each element of the approach to fit into the approximate timescales at which humans make decisions set forward by Allen Newell in [21].

Considering the tri-level control structure implemented by SOAR and Sigmathis resembles a combination of the deliberate and reflective mechanism. The main difference is that here attention mechanisms for choosing motoric actions similar to Equation (2) are proposed for evaluation rather than operators for which the outcome is known a priori.



The benefit of this is that motoric actions does not need to be represented by symbolic operators a priori. Instead, they are made symbolic implicitly through the choice of deliberate attention mechanism. Thereby, allowing the incorporation of uncertainties of low-level controls. Each of the proposed deliberate attention mechanisms might consider a subset of and/or special combinations and weightings of the appraisal variables available to the robot. Thereby, promoting different behaviors. While the proposed approach conceptually does support deliberate and reflective responses via the *affective responses*, it does not currently have support for reactive responses, since all motoric actions have to be inferred from the deliberate attention mechanisms. However, in Section 7 we will discuss how we imagine that reactive responses could be incorporated into the control structure. Furthermore, modern probabilistic programs such as Pyro [19] can combine stochastic variational inference with enumeration to infer the motoric action posterior. Thereby, making it possible to combine operators represented by both discrete/symbolic and continuous variables in the proposed control structure.

## 5. Idiom Modifications and Extensions

To test the approach proposed in Section 4 several modifications and extensions were made to the probabilistic programming idiom proposed in [7]. This includes additional appraisal variables, the possibility of adding and using additional deliberate attention mechanisms, together with an implementation of simple mechanisms for *deliberate attention proposal* and *affective responses*. In order to make the implementation reusable in the spirit of the framework presented in [5], all of this is implemented as a series of abstract python classes each constituting a probabilistic programming idiom available at [22].

### 5.1. Additional Appraisal Variables

Besides the *progress*,  $z_p$ , *information gain*,  $z_i$ , and *constraints*,  $z_c$ , appraisal variables defined in [7], a couple of new appraisal variables has been implemented. The first was the *accumulated constraints* appraisal,

$$\begin{aligned} p(z_{Ac} = 1 | Z_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM}) \\ = \text{Bernoulli} \left( p \left( \bigwedge_{\tau=t+1}^{\bar{T}} z_c^{\{\tau\}} = 1 \mid Z_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM} \right) \right), \end{aligned} \quad (4)$$

which was implemented due to a need to check constraint violations of a full state trajectory rather than at a single state. The second originated from a need to be able to define desirable/goal states that a robot should seek to attain. First, we approximate the KL-divergence between a desirable state,  $Z_s^*$ , and the state,  $Z_s^{\{\tau\}}$ , after effectuating the motoric action,  $Z_{Mb}^{\{\tau-1\}}$ , as

$$\begin{aligned} D_{KL} [p(Z_s^*) || p(Z_s^{\{\tau\}} | Z_s^{\{\tau-1\}}, Z_{Mb}^{\{\tau-1\}})] \\ = E_{Z_s^{\{\tau\}}} \left[ \log \left( \frac{p(Z_s^{\{\tau\}} = \hat{Z}_s^{\{\tau\}} | Z_s^{\{\tau-1\}}, Z_{Mb}^{\{\tau-1\}})}{p(Z_s^* = \hat{Z}_s^{\{\tau\}})} \right) \right] \\ \approx \frac{1}{I} \sum_{i=1}^I \left( \log(p(Z_s^* = \hat{Z}_s^{\{\tau\}, \{i\}})) - \log(p(Z_s^{\{\tau\}} = \hat{Z}_s^{\{\tau\}, \{i\}} | Z_s^{\{\tau-1\}}, Z_{Mb}^{\{\tau-1\}})) \right) \\ \approx \text{ReLU} \left( \log(p(Z_s^* = \hat{Z}_s^{\{\tau\}})) - \log(p(Z_s^{\{\tau\}} = \hat{Z}_s^{\{\tau\}} | Z_s^{\{\tau-1\}}, Z_{Mb}^{\{\tau-1\}})) \right) \\ \stackrel{\text{def}}{=} D_{Z_s^*}(\hat{Z}_s^{\{\tau\}}). \end{aligned}$$

Inspired by the optimality variable defined in [23] we then define the *desirability* appraisal,  $z_{d, Z_s^*}^{\{\tau\}}$ , as

$$p(z_{d,Z_s^*}^{\{\tau\}} = 1 | Z_s^{\{t+1:\bar{T}\}} = \hat{Z}_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM}) \quad (5)$$

$$= \begin{cases} 0, 0; & \text{if } p(z_{Ac} = 1 | Z_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM}) < 1 \\ \text{Bernoulli}(-e^{-\sigma_d \cdot D_{Z_s^*}(\hat{Z}_s^{\{\tau\}})}); & \text{else} \end{cases}$$

where the subscript  $Z_s^*$  in  $z_{d,Z_s^*}^{\{\tau\}}$  is used to denote the dependency on  $p(Z_s^*)$ , and  $\sigma_d$  is a scaling factor. Equation (5) defines a pseudo probability for which states most similar to the desirable state,  $Z_s^*$ , has the highest probability, and states that are less similar have an exponentially lower probability, while states resulting from trajectories that violate constraints have zero probability. The dependency on the *accumulated constraint* appraisal was introduced to aid in overcoming a small probability of constraint violation observed in [7]. For the same reason the *progress*,  $z_p$ , and *information gain*,  $z_i$ , appraisals has also been modified as follows

$$p(z_i^{\{\tau\}} = 1 | Z_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM})$$

$$= \begin{cases} 0, 0; & \text{if } p(z_{Ac} = 1 | Z_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM}) < 1 \\ p(z_i^{\{\tau\}} = 1 | Z_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM}); & \text{else} \end{cases}$$

$$p(z_p^{\{\tau\}} = 1 | Z_s^{\{t+1:\bar{T}\}} = \hat{Z}_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM})$$

$$= \begin{cases} 0, 0; & \text{if } p(z_{Ac} = 1 | Z_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM}) < 1 \\ p(z_p^{\{\tau\}} = 1 | Z_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM}); & \text{else} \end{cases}$$

where  $z_{\bar{p}}$  and  $z_{\bar{i}}$  are the *progress*,  $z_p$ , and *information gain*,  $z_i$ , appraisals as defined in [7].

## 5.2. Deliberate Attention Mechanisms

Based on the appraisals defined in Section 5.1 five different deliberate attention mechanisms have been implemented. All these can be defined as

$$p(x_A^{\{\tau\}} | Z_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM}) = \text{Bernoulli}(p(\Phi(\tau) | Z_s^{\{t+1:\bar{T}\}}, Z_{WM \setminus b}^{\{t\}-}, Z_{LTM}))$$

where  $\Phi(\tau)$  defines the logic for combining appraisals as in Table 1. Each of these deliberate attention mechanisms promotes different behaviors.


**Table 1.** Definitions of  $\Phi(\tau)$  used for each of the implemented deliberate attention mechanisms.

Deliberate Attention Mechanism	$\Phi(\tau)$ for $\tau \in [t+1; \bar{T}-1]$	$\Phi(\bar{T})$ for $\tau = \bar{T}$
ConstraintAvoidance—CA	$z_{Ac} = 1$	$z_{Ac} = 1$
StateReach—SR( $Z_s^*$ )	$z_{Ac} = 1$	$\begin{cases} z_{Ac} = 1; & P(z_{Ac} = 1) < 1 \\ z_{d,Z_s^*}^{\{\bar{T}\}} = 1; & \text{else} \end{cases}$
StateReachWithProgress—SRP( $Z_s^*$ )	$z_{Ac} = 1$	$\begin{cases} z_{Ac} = 1; & P(z_{Ac} = 1) < 1 \\ z_{d,Z_s^*}^{\{\bar{T}\}} = 1 \wedge z_p^{\{\bar{T}\}} = 1; & \text{else} \end{cases}$
StateReachWithExplore—SRE( $Z_s^*$ )	$z_{Ac} = 1$	$\begin{cases} z_{Ac} = 1; & P(z_{Ac} = 1) < 1 \\ z_{d,Z_s^*}^{\{\bar{T}\}} = 1 \wedge z_i^{\{\bar{T}\}} = 1; & \text{else} \end{cases}$
Explore—E	$z_{Ac} = 1$	$\begin{cases} z_{Ac} = 1; & P(z_{Ac} = 1) < 1 \\ z_i^{\{\bar{T}\}} = 1; & \text{else} \end{cases}$
ExploreWithProgress—EP	$z_{Ac} = 1$	$\begin{cases} z_{Ac} = 1; & P(z_{Ac} = 1) < 1 \\ z_i^{\{\bar{T}\}} = 1 \wedge z_p^{\{\bar{T}\}} = 1; & \text{else} \end{cases}$



### 5.3. Deliberate Attention Proposal and Affective Responses

Based on the *deliberate attention mechanisms* defined in Section 5.2, an *affective response mechanism* has been implemented with the purpose of making a robot effectively explore its environment and possibly navigate towards a goal state,  $Z_s^*$ , if defined. This *affective response mechanism* can be sub-divided into three parts responsible for different types of behaviors with pseudo-code given in Algorithms A1–A3. Algorithm A1 yields behavior that strives for the goal state. Algorithm A2 yields behavior that strives to obtain new information about the environment. Finally, Algorithm A3 yields behavior that strives to backtrack. The combined affective response only depends on the appraisals defined in Section 5.1 and [7] which requires no problem-specific information, thereby, making this affective response mechanism general and reusable. As illustrated in Figure 3, each part of the *affective response mechanism* is activated either as a reflective response to another part of the *affective response mechanism* or based on the *deliberate attention mechanism* that caused a motoric action to be effectuated in the last decision cycle. E.g., if the *deliberate attention mechanism* “Explore (E)” caused the effectuation of  $q_\phi^{\{E\}}(Z_{Mb}^{\{t-1\}})$  at time  $t - 1$ , then Algorithm A2 will be activated first at time  $t$ . In cases where the motoric action was caused by the *deliberate attention mechanism* “ConstraintAvoidance (CA)”, the same algorithm is simply activated again. The intuition behind this affective response mechanism is as follows. If a goal state is known and if it is possible to attain it with the current knowledge directly, this should have first priority. If this is not possible new information should be sought after until the goal state can be attained. Finally, if in a state where new information cannot be obtained, the system should be able to bring itself back to a previous state in which new information can be obtained via backtracking. To support this *affective response mechanism* a *deliberate attention proposal mechanism* has been implemented that simply proposes the *deliberate attention mechanisms* required for each part of the *affective response mechanism*. When combined this exemplifies how both deliberate and reflective responses can be implemented grounded in the appraisals defined in Section 5.2. In particular, notice that something similar to the “no-change” impasse in SOAR and Sigma is obtained on the basis of the “Progress” appraisal in both Algorithms A1 and A2.



**Figure 3.** Overview of the implemented affective response mechanism sub-divided into 3 algorithms. Dotted lines indicate what algorithm is activated based on the *deliberate attention mechanism*,  $DAM^{\{t-1\}}$ , which resulted in the effectuation of a motoric action in the last time step. Solid lines indicate a reflective response resulting in a direct transition between the algorithms. Each direct transition requires a new *deliberate attention proposal* and *deliberate attention evaluation*. A dashed line indicates an indirect transition between the algorithms which takes effect in the next decision cycle.  $Z_s^*$  denotes a goal state, and  $\mathbb{P}_{BT}$  denotes a path of previous states to backtrack.

## 6. Results


To test the effectiveness of the proposed approach two different simulation studies were performed. Both of these were done utilizing the Pseudo-SLAM simulator [20], and using the same implementation of abstract methods for the probabilistic programming idiom that was used in [7]. The exact parameters used for each of these simulations can be found at [22], which also contains scripts to replicate each of the experiments.

### 6.1. Pure Exploration

The first simulation study was done in order to compare with results from [7]. In [7] we tested the exploration capability of the proposed algorithm by simulating it in 35,126 floor plans from the HouseExpo dataset [20]. However, (1) many of these were so small that they were fully discovered in a few iterations, and (2) on the other end of the spectrum some of the floor plans were simply too big to be fully discovered within the maximum of 200 time-steps that was allowed in each simulation. Furthermore, (3) the problems of the previous solution discussed in Section 3 are only noticeable in floor plans with more than one room. Additionally, (4) it was found that for some of the floor plans openings between the rooms were physically too small for the robot to squeeze through. Thus, for the purpose of efficiently testing the approach proposed in this paper we selected a smaller subset of the HouseExpo dataset satisfying the following criteria.

1. The floor plans should have a bounding box larger than 100 m<sup>2</sup> to avoid spending time on simulations redundant due to (1).
2. The floor plans should be fully discovered in the experiment from [7], in order to minimize the influence from (2) and (4).
3. The floor plans should contain more than 3 rooms in order to provoke (3).

Based on this, a subset of the HouseExpo dataset consisting of 784 floor plans where selected. Figure 4 and Table 2 show the results of simulating our old approach, “AKS”, again as well as the approach proposed within this paper, “AR”. The simulations were performed with the same environmental and robot settings used in [7]. As no goal state was specified, only Algorithm A2 and Algorithm A3 were effectively used to drive the behavior of the “AR” method in these simulations. For each floor plan a random initial position where selected, and this initial position were utilized for both simulations. Notice, that even though one of the criteria for the selection of the subset of floor plans was that it should be fully explored by “AKS” in the experiment in [7], Figure 4 indicates that not all floor plans where fully explored by “AKS” in this new round of simulations. This is simply due to a difference in initial positions between simulations and illustrates a lack of robustness of “AKS”. From Figure 4 it might seem that “AR” does not perform better than “AKS” for small floor plans. By visual inspection of the simulation trajectories, it was found that the reason for “AR” not being able to explore some floor plans fully was due to a lesser willingness to violate constraints compared to “AKS”. This can be verified from the row “Collision pr. Timesteps” in Table 2. This is especially pronounced in small floor plans, where the openings between rooms tend to be smaller. To further verify that the lack of exploration by “AR” is indeed due to its unwillingness to violate constraints, the third series of simulations denoted “AR small” was performed. In these simulations, the size of the robot, the uncertainty in its initial position, and the assumed motion uncertainty was decreased. By changing these parameters, it becomes easier for the robot to take actions through narrow openings without constraint violations. From Figure 4 it is seen that “AR small” fully explores nearly all of the small floor plans, and performs similarly to “AR” for all other maps as expected. As the floor plans get bigger, the ability of all three methods to fully explore the floor plans to a greater extent depends on initial conditions, rather than the ability of the methods to escape local minima. As a result, from Figure 4 it is observed that as the floor plans get larger all methods perform very similarly. Nevertheless, from Table 2 it is evident that “AR” is indeed better for overcoming local minima and making the robot efficiently explore its environment.



**Figure 4.** The area that the robot explored in each of the simulations. The indices of the 784 floorplans have been sorted by the true area of the map in ascending order. “AKS” shows results from using the method presented in [7]. “AR” shows results from using the *affective response* mechanism from Section 5.3. The “smoothed” curves show a moving average with a window size of 100 and shifted 50 indexes. The “not done” scatter shows the exact area explored by the simulations in which the robot did not manage to explore 95% of the map or more.

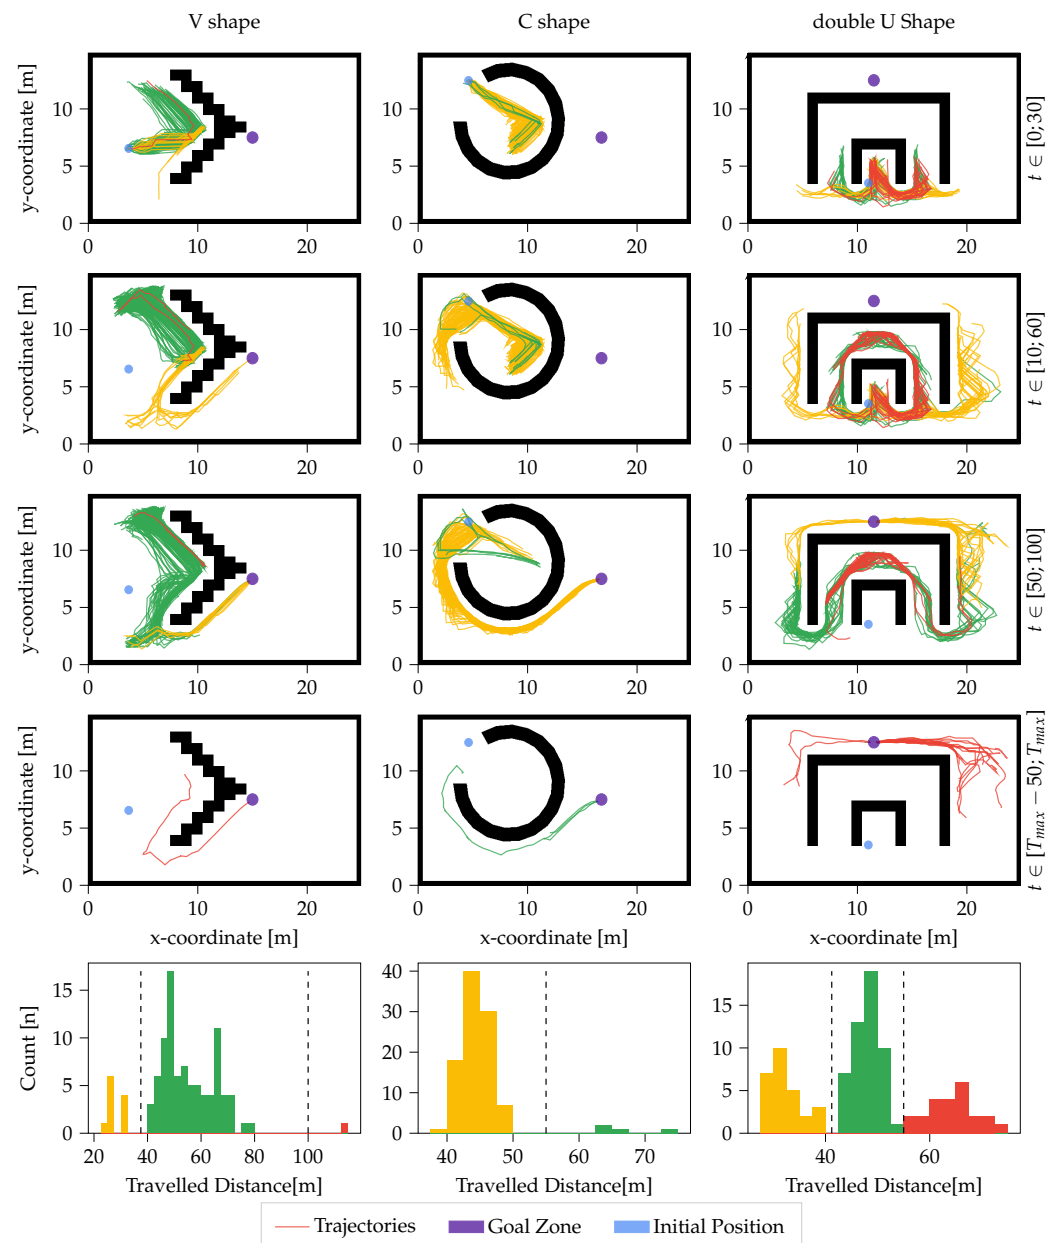
**Table 2.** Comparison of our approach with results for 6 different methods presented in [8].

Metric	AKS	RGS	RGS Small
Maps Not Fully Explored	484	331	314
Mean Exploration Percentage	84.9%	87.3%	90.7%
Mean percentage explored for unfinished maps	78.6%	76.7%	84.4%
Maps with Collisions	18	0	0
Collisions	19	0	0
Collision pr. Timesteps	0.14‰	0.00‰	0.00‰

## 6.2. Goal Seeking

The second series of simulations were performed in order to compare the proposed approach with more problem-specific approaches from [8]. To do so three of the test environments from [8] were recreated in the Pseudo-SLAM simulator as illustrated in Figure 5. These three environments are designed specifically with the purpose of causing local minima, and as such are perfect for testing the proposed approach. Since the approach proposed in this paper is based on probabilistic methods, some degree of variations in results should be expected. Therefore, 100 simulations were performed for each of these environments with the same initial conditions and goal state as in [8]. Since a goal state was specified for these simulations, the full capabilities of the *affective response* mechanism described in Section 5.3 were effectively in use. Table 3 summarizes the results from these simulations and compares them to the results from [8]. In all of the simulations, the robot managed to reach the goal state, thereby substantiating the ability of the proposed approach to escaping local minima. From Table 3 it is furthermore seen that the “AR” method is better than any of the problem-specific methods in all three environments when only considering the minimum traveled distance. However, when considering the average distance for the “V-shape” environment it is nearly twice that of the best method from [8],

i.e., “Reflected Virtual Target”. Considering the first column of Figure 5 it is clear that the robot generally can take the two paths indicated with green and yellow colors. We suspect that the better average performance of the “Reflected Virtual Target” method in the “V-shape” environment is caused by an initial condition that makes the problem-specific methods favor paths similar to the one marked with yellow in Figure 5. The better performance achieved by such preference would not necessarily lead to better performance in general environments/problems, and a more reasonable comparison would probably be obtained by some variations in the initial conditions and/or goal state. As such we do not consider this an inauspicious characteristic of the “AR” approach.



**Figure 5.** The robot trajectories in each of the 100 simulations for each of the three environments: “V shape”, “C shape”, and “double U shape”. Each of the trajectories is color-coded according to the length of the trajectory.  $T_{\max} = 308$ ,  $T_{\max} = 191$ , and  $T_{\max} = 208$  for “V shape”, “C shape”, and “double U shape”, respectively.


**Table 3.** The traveled distance in 3 different environments utilizing our approach, AR, compared with results for 6 other methods presented in [8]. As AR is based on probabilistic methods we ran 100 simulations and present the mean of the results together with the minimum and maximum values for each of the environments. The best result for each environment is highlighted with bold text.

Environment	Random <sup>1</sup>	Reflected Virtual Target <sup>1</sup>	Global Path Backtrack-ing <sup>1</sup>	Half Path Backtrack-ing <sup>1</sup>	Local Path Backtrack-ing <sup>1</sup>	Wall-Following <sup>1</sup>	AR
C-shaped	55	45	59	101	59	1915	45.34 [39.38–73.97]
Double U-shaped	97	88	100	110	96	466	47.51 [28.07–72.56]
V-shaped	38	27	29	31	28	111	52.91 [22.91–114.98]
Average	63.33	53.33	62.67	80.67	61	830.67	48.59

<sup>1</sup> Results from Table 2 in [8].

### 6.3. Timings

One of the most critical features of any robotics system is the satisfaction of the real-time constraint, i.e., the ability of the system to make decisions on time scales appropriate to the expected behavior of the system. To investigate the computational time required for the proposed approach the average computation times were measured on two different CPUs. The results can be seen in Figure 6. Notice, that these timings are based on a relatively slow python implementation and are uniquely tied to the specific use-case presented in Sections 6.1 and 6.2. As such, they should not be seen as the definitive timings that can be obtained utilizing the method, but rather as indicative of roughly what can be expected by the approach. Nevertheless, the timings given in Figure 6 would probably be too slow or jerky for most real-world robot applications. As should be clear from Figure 6 the most time-consuming part of the approach with the current implementation is the inference part of the *deliberate attention evaluations* step. As such, further optimization of this step would be needed to make the approach usable.



**Figure 6.** Timings for the main steps of the proposed solution running on two different CPUs, for the deliberate attention mechanism and the affective response mechanism presented in Sections 5.2 and 5.3, respectively, and for the abstract method implementations specifically used for the simulations in Sections 6.1 and 6.2. (a) Apple M1. (b) Intel 12600k@5.1 GHz.

## 7. Discussion

The intention of the presented efforts was to implement general reflective mechanisms suitable for robotic applications with an outset in previous work. In Sections 6.1 and 6.2 we demonstrate that the proposed method functionally improves upon our previous proposed probabilistic programming idiom and that it at least can perform as well as, if not better than, problem-specific methods. However, in Section 6.3 it was concluded

that the current implementation would probably be too slow and jerky for real-world robot applications. The approach presented within this paper is supposed to be generally applicable and reusable, making it hard to assess how much the current implementation should be improved to be applicable to real-world robot applications since this would of course depend on each specific use case. One way to assess this anyway could be by comparing it to Allan Newell's analysis of the time scales of human cognition [21]. This is reasonable because the ultimate end goal of our efforts is to make robots as capable as humans.

In a single cycle of *deliberate attention proposal*, *deliberate attention evaluation* and *affective response*, access to parts of cognition distal to the decision process have to have occurred multiple times in order to infer the motor buffer posteriors. This places the proposed approach somewhere above the "biological band" of Newell's analysis said to be on the order of  $\sim 10$  ms. The next step up in Newell's analysis is to the "cognitive band" starting at the level of *deliberate acts* in the order of  $\sim 100$  ms. However, the proposed approach does not merely comprise deliberation, i.e., choosing one known operation over other known operators by bringing available knowledge to bear, since operators are constructed for the to-be-produced response based on the proposed deliberate attention mechanisms. Therefore, the proposed approach also belongs somewhere above the time scales of *deliberate acts*. At the other end of the "cognitive band", we have *unit tasks* in the order of  $\sim 10$  s. At the time scale of *unit tasks*, operations should be composed to deal with tasks. By design, the specific affective response presented in Section 5.3 can only deliver simple responses in one decision cycle and not a plan of responses to solve complete tasks. This leaves us at the time scales of *elementary cognitive operations* or *immediate external cognitive behavior* at  $\sim 1$  s. According to Newell's analysis, such elementary reactions often take  $\sim 2$ – $3$  s, however, with learning from experience, simplification, preparation, and carefully shaped anticipation, it can take less than  $\sim 0.5$  s. By design, the specific affective response presented in Section 5.3 can deliver simple responses within 1 to 3 full cycles of *deliberate attention proposal*, *deliberate attention evaluation* and *affective response*. With the timings in Figure 6a response thus takes anywhere from  $\sim 1.7$  s up to at most  $\sim 5.34$  s in the case of two impasses. Thus, to arrive at the upper end of elementary reactions, the computational times of the current implementation would have to be improved with a factor of  $\sim 2$ – $3$ . Obtaining such improvements does not seem implausible via code optimizations, however, it brings us nowhere near the lower end of  $\sim 0.5$  s. This begs the question: can the proposed approach support the necessary machinery to learn from experience in order to deliver responses at the lower end of  $\sim 0.5$  s?

In Sections 3 and 4 we described the use of stochastic variational inference, as the basis for inferring a parametric approximation of the posterior over the motor buffer,  $q_{\phi}^{\{i\}}(Z_{Mb}^{\{t-1\}^+})$ . It was assumed that this inference process would have to be done from scratch in each decision cycle. However, this need not necessarily be the case. Instead, we could make use of amortized variational inference [24–27]. Thus, instead of making use of a variational distribution with free parameters,  $\phi$ , we would make use of a variational distribution with parameters determined by a parametric function,  $\phi = f_{\phi}^{\{i\}}(Z_{WM \setminus b}^{\{t\}^-, Z_{LTM})$ , e.g., a neural network. When new situations are encountered we would not necessarily gain much by doing so, however, over time this would in principle allow the system to generate proper responses to situations similar to those that the system has previously encountered, without performing any inference. Thereby, removing the need for the most time-consuming step in the decision cycle. Again, when considering the timings in Figure 6a, reducing the inference step to near zero, would bring the total time of a single decision cycle down to around  $\sim 500$  ms with the current implementation. Now if it is possible to improve the other steps with a factor of  $\sim 2$ – $3$  via code optimizations, it would indeed seem plausible to achieve *immediate external cognitive behavior* in around  $\sim 0.5$  s after an initial learning period.

Further optimization might be achieved by considering when to stop the underlying inference algorithm. In the current implementation, the underlying inference algorithm



With these additions and optimizations of the approach and its implementation, we believe that the approach will be applicable to real-world robot applications, and thereby contribute to the goal of constructing autonomous robots that can safely and naturally interact with humans while solving different abstractly and/or vaguely defined tasks. As such, these optimizations will be the focus of our future work.

**Funding:** This research received no external funding.

**Conflicts of Interest:** The authors declare no conflict of interest.

The following abbreviations are used in this manuscript:

RGS Reflective Goal Search

---

**Algorithm A1:** Affective response for State reach

---

**Require:** *Deliberate Attention Evaluations* of CA, SR( $Z_s^*$ ), SRP( $Z_s^*$ ), and SRE( $Z_s^*$ )

```

1: DAMS  $\leftarrow \{SR(Z_s^*), SRP(Z_s^*), SRE(Z_s^*)\}$ 
2: if  $\nexists \text{DAM} \in \text{DAMS} : p(z_p^{\{\text{DAM}\}} = 1) > P_{z_p, \text{lim}}$  then                                # No-change impasse
3:     Run Algorithm A2                                                                # Reflective Response
4: else
5:     repeat
6:          $\text{DAM}^* \leftarrow \arg \max_{\text{DAM} \in \text{DAMS}} p(z_d^{\{\text{DAM}\}} = 1)$                                 # pick the most desirable DAM
7:         if  $p(z_d^{\{\text{DAM}^*\}} = 1) \geq p(z_d^{\{\text{SR}\}} = 1)$  then
8:             # Only pick DAM* if it is more desirable than SR to avoid motoric
9:             # actions driven mainly by the progress or information gain appraisals
10:            if  $p(z_{\text{Ac}}^{\{\text{DAM}^*\}} = 1) \geq P_{z_c, \text{lim}}$  or  $p(z_{\text{Ac}}^{\{\text{DAM}^*\}} = 1) \geq p(z_{\text{Ac}}^{\{\text{CA}\}} = 1)$  then
11:                effectuate  $q_\phi^{\{\text{DAM}^*\}}(Z_{\text{Mb}}^{\{t\}})$ 
12:            end if
13:        end if
14:        DAMS  $\leftarrow \text{DAMS} \setminus \text{DAM}^*$ 
15:    until DAMS =  $\emptyset$ 
16:    effectuate  $q_\phi^{\{\text{CA}\}}(Z_{\text{Mb}}^{\{t\}})$                                 # pick the constraint avoidance
                                                                    strategy as a backup
17: end if

```

---

**Algorithm A2: Affective response for explore**


---

**Require:** *Deliberate Attention Evaluations* of CA, E, EP, and SR( $Z_s^*$ ) if  $Z_s^* \neq \text{None}$

```

1: DAMS  $\leftarrow \{E, EP\}$ 
2: DAM*  $\leftarrow \arg \max_{DAM \in DAMS} p(z_i^{\{DAM\}} = 1)$  # pick the DAM yielding most
                                                    information
3: if  $\nexists DAM \in DAMS : p(z_p^{\{DAM\}} = 1) > P_{z_p,lim}$  # No-change impasse
   or  $p(z_i^{\{DAM^*\}} = 1) \leq P_{z_i,lim}$  then # No information gain possible
4:    $\mathbb{P}_{BT} \leftarrow \text{set\_path\_to\_backtrack}()$  # Generate the path to backtrack,
                                                     $\mathbb{P}_{BT}$ , from the state tree
                                                    # Reflective Response

5:   Run Algorithm A3
6: else
7:   repeat
8:     DAM*  $\leftarrow \arg \max_{DAM \in DAMS} p(z_d^{\{DAM\}} = 1)$  # pick the most desirable DAM
9:     if  $Z_s^* \neq \text{None}$  and  $p(z_i^{\{SR(Z_s^*)\}} = 1) \geq p(z_i^{\{DAM^*\}} = 1) \cdot (1 - P_{z_i,\Delta})$  # StateReach gives nearly as much
                                                    information
       and  $p(z_p^{\{SR(Z_s^*)\}} = 1) > P_{z_p,lim}$  then # and sufficient progress
10:      DAM*  $\leftarrow SR(Z_s^*)$  # If SR does not violate constraints
                                                    it will be effectuated and Algo-
                                                    rithm A1 will be used in the next
                                                    iteration.

11:    end if
12:    if  $p(z_{Ac}^{\{DAM^*\}} = 1) \geq P_{z_c,lim}$  or  $p(z_{Ac}^{\{DAM^*\}} = 1) \geq p(z_{Ac}^{\{CA\}} = 1)$  then # No constraint violation?
13:      effectuate  $q_\phi^{\{DAM^*\}}(Z_{Mb}^{\{t\}})$ 
14:    else
15:      DAMS  $\leftarrow DAMS \setminus DAM^*$ 
16:    end if
17:  until DAMS  $\neq \emptyset$ 
18:  effectuate  $q_\phi^{\{CA\}}(Z_{Mb}^{\{t\}})$  # pick the constraint avoidance
                                                    strategy as a backup

19: end if

```

---



---

**Algorithm A3: Affective response for backtracking**


---

**Require:** *Deliberate Attention Evaluations* of CA, E, EP, and SR( $\mathbb{P}_{BT}[\tau]$ ) for  $\tau \in [1, \dots, T_{BT}]$

```

1: DAM*  $\leftarrow \arg \max_{DAM \in \{E, EP\}} p(z_i^{\{DAM\}} = 1)$  # pick the DAM yielding most
                                                    information
2: if  $p(z_i^{\{DAM^*\}} = 1) > P_{z_i,lim}$  and  $p(z_p^{\{DAM^*\}} = 1) > P_{z_p,lim}$  then # new information can be obtained
3:   create_new_state_branch() # add branch to state tree
4: else
5:   DAMS  $\leftarrow \emptyset$ 
6:   for DAM*  $\in \bigcup_{\tau=0}^{T_{BT}} \{SR(\mathbb{P}_{BT}[\tau])\}$  do # SR for the first  $T_{BT}$  states in  $\mathbb{P}_{BT}$ 
7:     DAMS  $\leftarrow DAM^* \cup DAMS$ 
8:     if  $p(z_d^{\{DAM^*\}} = 1) > P_{BT,min}$  then # state can be reached
9:       if  $p(z_{Ac}^{\{DAM^*\}} = 1) \geq P_{z_c,lim}$  # without collision
       or  $p(z_{Ac}^{\{DAM^*\}} = 1) \geq p(z_{Ac}^{\{CA\}} = 1)$  then
10:       $\mathbb{P}_{BT} \leftarrow \mathbb{P}_{BT} \setminus DAMS$  # remove the reachable states
                                                    from the path currently being
                                                    backtracked
11:      if  $\mathbb{P}_{BT} = \emptyset$  then # end of path has been reached
12:         $\mathbb{P}_{BT} \leftarrow \text{set\_path\_to\_backtrack}()$  # generate new path to backtrack
13:      end if
14:      effectuate  $q_\phi^{\{DAM^*\}}(Z_{Mb}^{\{t\}})$ 
15:    end if
16:  end for
17: end if
18: end if
19: if  $p(z_{Ac}^{\{DAM^*\}} = 1) \geq P_{z_c,lim}$  or  $p(z_{Ac}^{\{DAM^*\}} = 1) \geq p(z_{Ac}^{\{CA\}} = 1)$  then # DAM* is less risky
20:   effectuate  $q_\phi^{\{DAM^*\}}(Z_{Mb}^{\{t\}})$  # than collision avoidance
21: else
22:   effectuate  $q_\phi^{\{CA\}}(Z_{Mb}^{\{t\}})$ 
23: end if

```

---

## References

- Demir, K.A.; Döven, G.; Sezen, B. Industry 5.0 and Human-Robot Co-working. *Procedia Comput. Sci.* **2019**, *158*, 688–695. [CrossRef]
- Fang, B.; Guo, X.; Wang, Z.; Li, Y.; Elhoseny, M.; Yuan, X. Collaborative task assignment of interconnected, affective robots towards autonomous healthcare assistant. *Future Gener. Comput. Syst.* **2019**, *92*, 241–251. [CrossRef]
- Farid, F.; Elkhodr, M.; Sabrina, F.; Ahamed, F.; Gide, E. A Smart Biometric Identity Management Framework for Personalised IoT and Cloud Computing-Based Healthcare Services. *Sensors* **2021**, *21*, 552. [CrossRef] [PubMed]
- Kaiser, M.S.; Al Mamun, S.; Mahmud, M.; Tania, M.H., Healthcare Robots to Combat COVID-19. In *COVID-19: Prediction, Decision-Making, and Its Impacts*; Santosh, K., Joshi, A., Eds.; Springer: Singapore, 2021; pp. 83–97. [CrossRef]
- Damgaard, M.R.; Pedersen, R.; Bak, T. Toward an idiomatic framework for cognitive robotics. *Patterns* **2022**, *3*, 100533. [CrossRef] [PubMed]
- Laird, J.E.; Lebiere, C.; Rosenbloom, P.S. A Standard Model of the Mind: Toward a Common Computational Framework across Artificial Intelligence, Cognitive Science, Neuroscience, and Robotics. *AI Mag.* **2017**, *38*, 13–26. v38i4.2744. [CrossRef]
- Damgaard, M.R.; Pedersen, R.; Bak, T. A Probabilistic Programming Idiom for Active Knowledge Search. In Proceedings of the 2022 International Joint Conference on Neural Networks (IJCNN), Padua, Italy, 18–23 July 2022; pp. 1–9. [CrossRef]
- Tashtoush, Y.; Haj-Mahmoud, I.; Darwish, O.; Maabreh, M.; Alsinglawi, B.; Elkhodr, M.; Alsaedi, N. Enhancing Robots Navigation in Internet of Things Indoor Systems. *Computers* **2021**, *10*, 153. [CrossRef]
- Grover, J.S.; Liu, C.; Sycara, K. Deadlock Analysis and Resolution for Multi-robot Systems. In *Algorithmic Foundations of Robotics XIV*; LaValle, S.M., Lin, M., Ojala, T., Shell, D., Yu, J., Eds.; Springer International Publishing: Cham, Switzerland, 2021; pp. 294–312. [CrossRef]
- Boldrer, M.; Andreetto, M.; Divan, S.; Palopoli, L.; Fontanelli, D. Socially-Aware Reactive Obstacle Avoidance Strategy Based on Limit Cycle. *IEEE Robot. Autom. Lett.* **2020**, *5*, 3251–3258. [CrossRef]
- Krishna, K.M.; Kalra, P.K. Solving the local minima problem for a mobile robot by classification of spatio-temporal sensory sequences. *J. Robot. Syst.* **2000**, *17*, 549–564. [CrossRef]
- Mohanty, P.K.; Kodapurath, A.A.; Singh, R.K. A Hybrid Artificial Immune System for Mobile Robot Navigation in Unknown Environments. *Iran. J. Sci. Technol. Trans. Electr. Eng.* **2020**, *44*, 1619–1631. [CrossRef]
- Ordonez, C.; Collins, E.G.; Selekwa, M.F.; Dunlap, D.D. The virtual wall approach to limit cycle avoidance for unmanned ground vehicles. *Robot. Auton. Syst.* **2008**, *56*, 645–657. [CrossRef]
- Sanchez, G.M.; Giovanini, L.L. Autonomous navigation with deadlock detection and avoidance. *Intel. Artif.* **2014**, *17*, 1323.
- Alonso-Mora, J.; DeCastro, J.A.; Raman, V.; Rus, D.; Kress-Gazit, H. Reactive mission and motion planning with deadlock resolution avoiding dynamic obstacles. *Autono. Robots* **2018**, *42*, 801–824. [CrossRef]
- Laird, J.E. *The Soar Cognitive Architecture*; The MIT Press: Cambridge, MA, USA, 2012.
- Rosenbloom, P.S.; Demski, A.; Ustun, V. The Sigma Cognitive Architecture and System: Towards Functionally Elegant Grand Unification. *J. Artif. Gen. Intell.* **2017**, *7*, 1–103. [CrossRef]
- Rosenbloom, P.S.; Gratch, J.; Ustun, V. Towards Emotion in Sigma: From Appraisal to Attention. In *International Conference on Artificial General Intelligence*; Bieger, J., Goertzel, B., Potapov, A., Eds.; Springer International Publishing: Cham, Switzerland, 2015; pp. 142–151. [CrossRef]
- Bingham, E.; Chen, J.P.; Jankowiak, M.; Obermeyer, F.; Pradhan, N.; Karaletsos, T.; Singh, R.; Szerlip, P.A.; Horsfall, P.; Goodman, N.D. Pyro: Deep Universal Probabilistic Programming. *J. Mach. Learn. Res.* **2019**, *20*, 28:1–28:6.
- Li, T.; Ho, D.; Li, C.; Zhu, D.; Wang, C.; Meng, M.Q.H. HouseExpo: A Large-scale 2D Indoor Layout Dataset for Learning-based Algorithms on Mobile Robots. In Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, NV, USA, 24 October 2020–24 January 2021; pp. 5839–5846. [CrossRef]
- Newell, A. *Unified Theories of Cognition*; Harvard University Press: Cambridge, MA, USA, 1990.
- Damgaard, M.R. ProbMind. 2022. Available online: <https://github.com/damgaardmr/probMind/tree/ec996e295575c384879b3d72cfc7e64b8085b9a5> (accessed on 3 November 2022)
- Damgaard, M.R.; Pedersen, R.; Bak, T. Study of Variational Inference for Flexible Distributed Probabilistic Robotics. *Robotics* **2022**, *11*, 38. [CrossRef]
- Zhang, C.; Butepage, J.; Kjellstrom, H.; Mandt, S. Advances in Variational Inference. *arXiv* **2017**, arXiv:1711.05597.
- Rezende, D.J.; Mohamed, S.; Wierstra, D. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In Proceedings of the 31st International Conference on Machine Learning, Beijing China, 21–26 June 2014; Xing, E.P., Jebara, T., Eds.; Proceedings of Machine Learning Research (PMLR): Beijing, China, 2014; Volume 32, pp. 1278–1286.

- 
26. Kingma, D.P.; Welling, M. Auto-Encoding Variational Bayes. In Proceedings of the 2nd International Conference on Learning Representations (ICLR 2014), Banff, AB, Canada, 14–16 April 2014.
  27. Shu, R.; Bui, H.H.; Zhao, S.; Kochenderfer, M.J.; Ermon, S. Amortized Inference Regularization. In *Advances in Neural Information Processing Systems*; Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R., Eds.; Curran Associates, Inc.: Sydney, NSW, Australia, 2018; Volume 31.