RESEARCHERID
THOMSON REUTERS

ISSN: 2277-9655
[Harvey * *et al.*, 7(8): August, 2018] Impact Factor: 5.164
IC™ Value: 3.00 CODEN: IJESS7

IJESRT

INTERNATIONAL JOURNAL OF ENGINEERING SCIENCES & RESEARCH TECHNOLOGY

FREE ENERGY FORAGING IN AN AFFORDANCE LANDSCAPE

Kole Harvey

DOI: 10.5281/zenodo.1345686

ABSTRACT

Introduction

In this article we construct a simulation of a virtual agent which is equipped with a predictive model of its environment and which operates based on the free energy principle to minimize prediction error. The agent is capable of perceiving a landscape of multiple affordances for action in the environment, and selects behaviors towards those affordances based on its internal needs and its relation to certain facets of the environment.

Methods

Through the use of a hierarchical model, the agent is endowed with the ability to choose high level behaviors which stabilize its actions, mid level behaviors which respond appropriately to affordances, and low level actions which modify action online.

Results

We demonstrate the ability of the agent to engage in foraging behavior based on free energy minimization. The agent is shown to balance multiple conflicting needs by responding to appropriate affordances in appropriate contexts. The agent is further shown to dynamically adjust its behavior on the y to respond to obstacles in its path. Overall it demonstrates appropriate behavior over multiple timescales.

Conclusions

The agent we have introduced engages in probablistic inference of the hidden states and causes of the external world, and is led by prediction error minimization to either update its model to be more faithful to the dynamics of the world or to execute actions which cause the world to be more in line with its own predictions. Throughout the course of prediction making, the agent is influenced by the state of the external environment and its internal needs.

Keywords: free energy, affordance competition, predictive processing, adaptive behavior.

I. INTRODUCTION

In this article we give an account of a rational agent capable of foraging in a virtual environment in order to fulfill its internal needs. To do this, we rely on ideas from the free energy principle (FEP) to construct an agent which behaves so as to minimize the surprise of the world, by sampling sensory states in a manner which leads to predictable outcomes. The agent is equipped with a generative model which encodes a probability distribution reflecting the state of certain aspects of the world, as well as a set of `precisions' or inverse variances which encodes the uncertainty of each of those aspects.

The agent adapts to the world by minimizing the prediction errors that its model makes about various modalities of sensory input, adjusting the internal model parameters dynamically down the gradient of free energy. In addition to minimizing prediction errors through perception, the agent can also take actions in the world in order to reduce proprioceptive prediction errors by dynamically selecting action down the gradient of free energy (active inference).

It is a long-standing question how an organism abiding by FEP is capable of engaging in behaviors that are not just `motor reflexes' directly connected to the outside world. Here we rely on the hierarchical aspect of the model to encompass the idea of high-level behaviors which are chosen based on the needs of the agent, and enacted by influencing the direct interactions the agent has with the world. High-level behaviors can essentially act as high-level priors that top-down bias certain actions over others, allowing the agent to focus on managing one need at a time even when other needs are depleting. Further, we will discuss how reflex action can take place on the sub-behavioral level by allowing the agent to make predictions that it will turn away when its

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

sensors detect an obstacle, while still maintaining higher behavior routines such as foraging and mid-level actions such as approaching food.

Another idea we rely on in this article is that of affordances, or possible actions the agent can take in the environment, which the agent can perceive and respond to in a rational way. Each affordance is given a salience and mutually exclusive affordances are competed against one another, eventually leading to a winning affordance or set of affordances which the agent can respond to and thus satisfy its needs. In this way our account is compatible with the affordance competition hypothesis [1], which we believe is a key idea for building intelligent agents capable of traversing complex environments, as we have discussed in previous work [2],[3].

In the rest of this article we will describe the agent in question which builds upon the previous ideas and is responsive to aspects of a virtual environment which are relevant to its needs. We base our model loosely on that given in [4], with the main difference that instead of a preset sequence of salient features, our agent responds dynamically to the actual aspects of the environment, their salience dictated by their relation to the agent's needs and its location in the world. In this way we adopt a relation-based definition of affordances as found in [5], which relies on both aspects of the environment and the agent itself. Through the aforementioned competition process, we ensure that the agent's response to affordances is not completely automatic but rather flexible and situation dependent [6].

The virtual environment in which the agent is placed presents a `landscape of affordances [7]. This landscape can be interacted with by the agent as it makes predictions to which aspects it will respond to. For example, when experiencing thirst the agent will predict it will approach water, and this will lead to the agent taking this action in order to reduce the prediction error that would follow if it did not.

At some levels of the behavior hierarchy of the agent, there is competition between alternate possibilities, whereas at the lowest levels the actions taken are practically deterministic as they have been fully contextualized to the point where there is only a single rational choice. The idea of incremental contextualization of behavior to reduce the myriad choices for action available in an environment to routine or reflex action has been covered previously in [2], and we believe this is a key method for realizing competent behavior in complex environments. The result is an adaptive agent which is capable of fulfilling its needs in a dynamic way based on the state of the changing environment and its relation to it, and which does so based on the underlying principle of free energy minimization.

II. METHODS

2.1 Free energy and active Inference

We develop an agent with a generative model of the world which can predict the dynamics of the hidden state and hidden causes of the external environment. Through reduction of the prediction error of this model over time, the agent essentially adapts itself to the environment. The way this works in practice is by minimizing free energy by either updating the state of the model to better reflect the outside world (perception), or through active inference, enacting proprioceptive prediction errors in the world to better bring it in line with the prediction (action). Free energy itself is an upper bound on the surprise of the encountered sensory states, and minimizing it leads the agent to minimize the amount of possible (probabilistic) states of the world it has to entertain. From an ecological point of view, this prevents the agent from becoming a fish out of water, and allows it to maintain its homeostatic needs' within a viable range.

In the simulations we carry out in this article, we use a hierarchical, dynamic generative model to form predictions and actions which reduce free energy. A key property of this model is that it assumes the probability density function over hidden states is Gaussian, thus only the expectation of the density needs to be manipulated (known as the Laplace assumption). The dynamical change of the internal state and executed action over time can be formalized in terms of a perception and action equation respectively, as shown below. Here we see that the internal state m and action a descend the gradient of free energy From, itself calculated based on sensory input s and the state μ

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

$$\dot{\tilde{\mu}} = D\tilde{\mu} - \frac{\partial}{\partial\tilde{\mu}}F_m(\tilde{s},\tilde{\mu})$$
$$\dot{a} = -\frac{\partial}{\partial a}F_m(\tilde{s},\tilde{\mu}) \quad (1)$$

Here the D operator takes as input a variable in generalized coordinates of motion (noted by the tilde notation) and returns the generalized motion as follows:

$$\tilde{\mu} = [\mu, \mu', \mu'', ...]^T$$
$$D\tilde{\mu} = [\mu', \mu'', \mu''', ...]^T \quad (2)$$

The exact formulation of the equations in (1) depend on the generative model itself, as we describe below

2.2 The generative model

The generative model describes the dynamics of the internal state of the model as defined by the function f , as well as the causal link between levels of the model defined by the function g, which at the lowest level also produces the sensory prediction s(t) at each time step t. ω constitutes noise at each level of the hierarchy and has amplitude of $1/\Pi(i,x)$, $1/\Pi(i,v)$ for hidden state and hidden cause dynamics respectively. Here we can see how the hidden state x is turned into a cause $v(i-1)$ for the previous layer via transformation by the function g(i).

$$s(t) = g^{(1)}(x^{(1)}, v^{(1)}) + \omega^{(1,v)}(x^{(1)}, v^{(1)})$$
$$\dot{x}^{(1)} = f^{(1)}(x^{(1)}, v^{(1)}) + \omega^{(1,x)}(x^{(1)}, v^{(1)})$$
$$\vdots$$
$$v^{(i-1)} = g^{(i)}(x^{(i)}, v^{(i)}) + \omega^{(i,v)}(x^{(i)}, v^{(i)})$$
$$\dot{x}^{(i)} = f^{(i)}(x^{(i)}, v^{(i)}) + \omega^{(i,x)}(x^{(i)}, v^{(i)}) \quad (3)$$

On each time step, the model produces a prediction at each level i about the (expectation of) the hidden state μx and hidden cause μv. These are then compared to the actual conditional expectations μ˜ of the model (or sensory input at the bottom level) to form the following prediction errors:

$$\tilde{\varepsilon}^{(i,v)} = \tilde{\Pi}^{(i,v)}(\tilde{\mu}^{(i-1,v)} - \tilde{g}^{(i)})$$
$$\tilde{\varepsilon}^{(i,x)} = \tilde{\Pi}^{(i,x)}(D\tilde{\mu}^{(i-1,x)} - \tilde{f}^{(i)}) \quad (4)$$

These prediction errors are then used to update the model as follows:

$$\dot{\tilde{\mu}}^{(i,v)} = D\tilde{\mu}^{(i,v)} + \left(\frac{\partial\tilde{g}^{(i)}}{\partial\tilde{\mu}^{(i,v)}}\right)^T \tilde{\varepsilon}^{(i,v)} + \left(\frac{\partial\tilde{f}^{(i)}}{\partial\tilde{\mu}^{(i,v)}}\right)^T \tilde{\varepsilon}^{(i,x)} - \tilde{\varepsilon}^{(i+1,v)}$$
$$\dot{\tilde{\mu}}^{(i,x)} = D\tilde{\mu}^{(i,x)} + \left(\frac{\partial\tilde{g}^{(i)}}{\partial\tilde{\mu}^{(i,x)}}\right)^T \tilde{\varepsilon}^{(i,v)} + \left(\frac{\partial\tilde{f}^{(i)}}{\partial\tilde{\mu}^{(i,x)}}\right)^T \tilde{\varepsilon}^{(i,x)} - D^T\tilde{\varepsilon}^{(i,x)} \quad \textbf{(5)}$$

A full explanation on how free energy minimization and hierarchical message passing works can be found in [8],[9], [10]. The crux of the generative model lies in the functions f and g, which determine the dynamics of the internal state and the sensory predictions generated from this state. The format of the generative model employed in this article lets the agent make predictions which allow it to respond to a subset of the affordances for action in the environment based on internal needs and the current relation of the agent to the environment (e.g. distance to each affordance, current actions it is already performing). The dynamics of the internal state in our model will thus rely on two major features: competing affordances, and response to those winning affordances by choosing behavior.

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

2.3 Affordance Competition

The affordances which link the skills of the agent to facets of the environment play a large part of the current model. The available affordances at any particular time depend on both the state of the environment and the agent, and are hierarchical in the sense that they can trigger high-level behaviors such as foraging, resting, and finding water, or more immediate actions such as approaching, gripping, and eating. The agent predicts the salience of each affordance and competition between each affordances is done via a softmax function which finds the most salient affordance(s) at any given time, as shown in the following equation. This competition allows the different possibilities for approaching, eating, and drinking to compete against one another. The values that make up this matrix can be further biased when high-level behaviors are engaged, and the predictions for approach action generated as a result of this competition can also be modified on the fly by altering the predicted trajectory to avoid obstacles. As such this competition process lies at the heart of a hierarchy of adaptive behavior.

The model takes two parameters actmin and distmax , which denote the minimum value that the state of an appendage such as grip must reach before it is considered activated, and the maximum distance within which the agent is considered within manipulating range of a part of the environment, respectively. The below matrix essentially allows the agent to approach, eat, or drink relevant targets in an appropriate way based on its needs and position in the environment.

$$salience = softmax\left(\begin{bmatrix} app_{food_1} & \dots & app_{food_nf} & app_{home} & app_{water} \\ eat_{food_1} & \dots & eat_{food_nf} & 0 & 0 \\ 0 & \dots & 0 & 0 & drn_{water} \end{bmatrix}\right)$$

$$app_{food} = \begin{cases} 1 - nutrition & |food - home| < dist_{max}, |agent - food| \geq dist_{max} \\ 6/(min(6, (1 + dist))) & |food - home| \geq dist_{max} \\ 0 & otherwise \end{cases}$$

$$app_{home} = \begin{cases} (1 - x_{rest}) & x_{grip} \geq act_{min}, |agent - home| \geq dist_{max} \\ 0 & otherwise \end{cases}$$

$$app_{water} = \begin{cases} |x_{body} - x_{water}| \geq dist_{max} & hydration < 0.9 \\ 0 & otherwise \end{cases}$$

$$eat_{food} = \begin{cases} (1 - x_{nutrition}) * 4 & |agent - home| \geq dist_{max}, |agent - food| < dist_{max} \\ 0 & otherwise \end{cases}$$

$$drn_{water} = \begin{cases} |x_{body} - x_{water}| < dist_{max} & hydration < 0.9 \\ 0 & otherwise \end{cases}$$

(6)

2.4 Taking action through prediction

Here we explain how the agent can select and modify predictions to take both high and low-level actions based on the winning affordances and the agent's own internal needs.

2.4.1 Top-level: behavior selection

Our agent is attributed 3 primary needs: nutrition, hydration, and rest. These values decrease linearly over time. The second level of the model infers the state of these needs through a causal link with the first level of the model, and uses these inferred values to predict engagement in high-level behaviors depending on which of these needs is the most pressing as follows. These high-level behaviors play a role in stabilizing the activity of the agent so that it can maintain all of its required needs, as will become evident in the experiments to follow. The high-level behavior is given a degree of hysteresis in order to stabilize activity of the agent.

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

$$\dot{x_{behav}} = softmax(0.1(x_{behav}) + 0.9(1 - x_{need})) - x_{behav} \tag{7}$$

2.4.2 Mid-level: action selection

At this level the most complex form of competition is undertaken, as the agent decides which target in the environment to approach, eat, or drink. Once the affordances have been reduced to a winner or small set of winners as shown in equation (6), the agent selects basic actions in response to these most salient affordances, which could also be called 'solicitations' [7]. When engaging in approach behavior, the salience-weighted average location of considered targets, target global, is approached. Along the way the agent will eventually favor one target over the others, thus narrowing down the winners to a single target in an efficient way. This form of incremental decision-making which takes into account multiple targets in parallel reflects the psychological data [11] and can be seen as a way to take optimal and adaptive actions based on indeterminate information which is updated in an online fashion. Further, when choosing to lie down, close its grip, or open its mouth, the agent relies not only on the immediate prediction of salience of the affordance but also the actual state of the environment (e.g. distance to the target). This can be seen as a sort of 'low-level' affordance check, which negates grasping at thin air or lying down in dangerous areas, even if these actions are momentarily entertained at higher levels.

$$\begin{aligned}
target_{global} &= salience_{1,*} \cdot x_{targets} \\
target_{body} &= \frac{target_{global} - x_{body}}{|target_{global} - x_{body}|} \\
\dot{x}_{body} &= target_{body} * 0.05 \\
\dot{x}_{grip} &= \begin{cases} 1 - x_{grip} & |x_{body} - x_{food}| < dist_{max}, |x_{food} - x_{home}| \geq dist_{max} \\ 1 - x_{grip} & x_{grip} \geq act_{min}, |x_{home} - x_{food}| \geq dist_{max} \\ -x_{grip} & x_{grip} \geq act_{min}, |x_{home} - x_{food}| < dist_{max} \\ 0 & otherwise \end{cases} \\
\dot{x}_{mouth} &= \begin{cases} 1 - x_{mouth} & salience_{drn,water} \geq act_{min}, |x_{body} - x_{water}| < dist_{max} \\ 1 - x_{mouth} & salience_{eat,food_i} \geq act_{min}, |x_{body} - x_{food_i}| < dist_{max} \\ -xmouth & otherwise \end{cases} \\
\dot{x}_{lie} &= \begin{cases} 1 - x_{lie} & x_{mouth} < act_{min}, |x_{body} - x_{home}| < dist_{max} \\ -x_{lie} & otherwise \end{cases}
\end{aligned} \tag{8}$$

2.4.3 Low-level: obstacle avoidance

At this level the agent's body is equipped with two proximity sensors on the left and right of its head which can detect nearby obstacles. These values are processed by the agent as it maneuvers toward a goal and change its proprioceptive predictions at the bottom layer of the model in an online fashion by rotating the approach vector away from the sensed obstacle as follows.

$$\begin{aligned}
target_{body} &= R(act_{sens_R} * \tfrac{\pi}{2})R(-act_{sens_L} * \tfrac{\pi}{2})target_{body} \\
act_{sens_L} &= \Sigma_o[max(0, 0.5 - norm(xsens_L - obst_o)) * 3] \\
act_{sens_R} &= \Sigma_o[max(0, 0.5 - norm(xsens_R - obst_o)) * 3] \\
x_{sensL} &= x_{body} + R(-\tfrac{\pi}{8})x_{head} \\
x_{sensR} &= x_{body} + R(\tfrac{\pi}{8})x_{head} \\
x_{head} &= target_{body} * 0.3 \\
R(\theta) &= \begin{bmatrix} cos(\theta) & -sin(\theta) \\ sin(\theta) & cos(\theta) \end{bmatrix}
\end{aligned} \tag{9}$$

Now that we have described the generative model, we will describe the complementary simulation of the external process which acts as the 'external environment and reacts to the agent"s actions, providing sensory feedback at each time step. Further, a depiction of the content of the external process and generative model can be found in Figure 1.

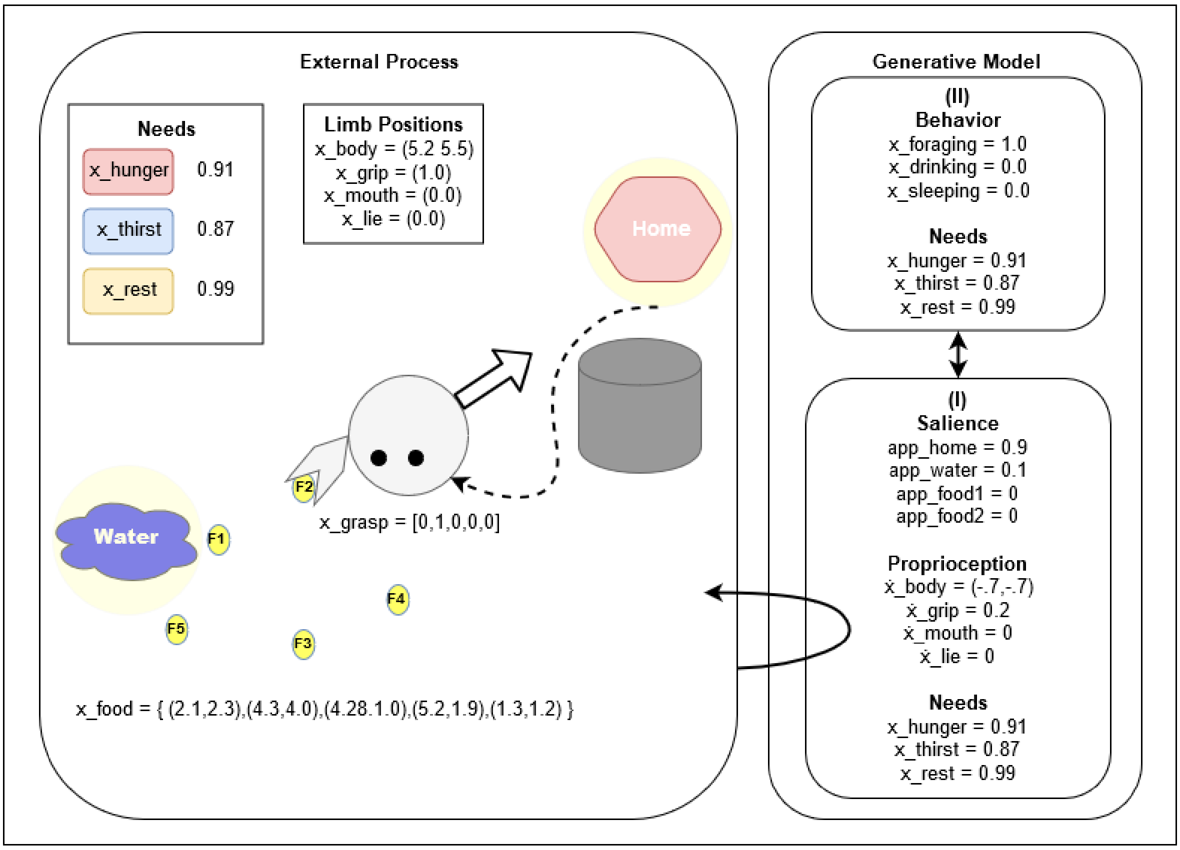

Figure 1. The generative model and external process

(Left) The external process which simulates a real environment with several food pellets, a water pool, a home location, and obstacles, each with a physical location. The agent also has several variables related to its physical apparatus: body position, grip position, mouth position, and whether it is lieing down in a rest position. Further, the external process has an extra variable xgrasp which stores whether the agent is grasping a particular pellet. The physical or 'real' set of needs of the agent are also external, and the agent must infer them along with the other external variables it makes predictions about. (Right) The generative model, consisting of two layers.
The top layer selects a high-level behavior based on the state of the agents needs. The bottom layer predicts approach actions, eating actions, and gripping actions based on the top-level behavior and inferred needs.

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

Approach trajectories can be modified as low-level obstacle avoidance. The proprioceptive predictions made by the bottom layer are transformed into action in the external process. The external process is dynamically updated and returns a sensory input back to the model.

2.5 External process

Sensory input from the external environment is simulated using a single layer based on the following equations, which takes in actions a from the generative model and uses them to update the external process. Essentially, actions are proprioceptive prediction errors which update each part of the agent's physical apparatus such as its body location, mouth, gripper, and whether it is lying down to rest. Further, the external or physical equivalent of the needs of the agent are updated based on their natural decay and consummative activity of the agent. Lastly, the location of food items in the world are updated depending on whether the agent is pulling them. Whether or not the food items are being pulled is determined by the external variable xgrasp , which acts as a placeholder for the dynamics of the food being held in the gripper of the agent. A noise term ω(1, x) is added to the dynamics at each timestep, and finally the actual state of the external environment is returned to the generative model in the form of sensory input s, which for simplicity is just the external state itself plus a noise term, although this could take any form without changing the results drastically as long as the generative model"s g function adequately represents the transformation from external state to sensory input.

Further, the external process is governed by the following rules:

- The agent must take food back home in order to eat it.
- The agent must open its mouth in the range of food to eat it.
- Eaten food cannot be eaten again.
- The agent must go home and lie down to rest.
- The agent must be in range of the water pool and open its mouth to drink from it.

Given the above, we define the equations for the external process as follows:

$$s = x + \omega(1, v)$$

$$\dot{x} = \begin{bmatrix} \dot{x}_{body} \\ \dot{x}_{thirst} \\ \dot{x}_{hunger} \\ \dot{x}_{rest} \\ \dot{x}_{grasp} \\ \dot{x}_{foodloc} \end{bmatrix} + \omega(1, x)$$

$$\dot{x}_{body} = a - x_{body}/8$$

$$\dot{x}_{thirst} = \begin{cases} 0.08 & x_{mouth} \geq act_{min}, |x_{body} - x_{water}| < dist_{max}, x_{thirst} < 1.0 \\ 0 & otherwise \end{cases}$$

$$\dot{x}_{hunger} = \begin{cases} 0.1 & x_{mouth} \geq act_{min}, |x_{body} - x_{food_i}| < dist_{max}, x_{hunger} < 1.0 \\ 0 & otherwise \end{cases}$$

$$\dot{x}_{rest} = \begin{cases} 0.04 & x_{lie} \geq act_{min}, x_{rest} < 1.0 \\ 0 & otherwise \end{cases}$$

$$\dot{x}_{grasp_i} = \begin{cases} 1 - x_{grasp_i} & x_{grip} \geq act_{min}, |x_{body} - x_{food_i}| < dist_{max} \\ -1 - x_{grasp_i} & x_{grasp_i} \geq act_{min}, x_{grip} < act_{min} \\ 0 & otherwise \end{cases}$$

$$\dot{x}_{foodloc_i} = \begin{cases} x_{body} - x_{foodloc_i} & x_{grip} \geq act_{min}, x_{grasp_i} \geq act_{min} \\ 0 & otherwise \end{cases} \tag{10}$$

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

2.6 Summary

In this section we have described the format of the generative model and external process employed in the experiments to follow. We have shown how the physical appendages of the agent are updated in the external process based on actions, which in turn are generated from decisions based on the salience of affordances for action in the environment. These affordances are sensitive to changes in the needs of the agent, which are themselves 'external' properties which the agent must infer in order to react to them and preserve a safe homeostatic range. We have shown how the agent can generate action through the cooperation of three stages:– a high-level behavioral stage (level 2 of the generative model), a mid-level action stage (level 1 of the generative model), and an online reflex stage (manipulation of level 1 predictions based on immediate sensor values). Next we will show the results of several experiments we undertook to demonstrate how the agent adaptively responds to both its internal needs and the state of the external environment by inferring this information and creating proprioceptive predictions that turn into action, and also show how and why each stage of the behavioral hierarchy is necessary for successful adaptive behavior.

III. RESULTS

Simulations for the following experiments were ran in MATLAB using the ADEM foraging.m model which we developed based on the equations as laid out above, and utilizes Karl Friston's spm ADEM.m to simulate the generative model and sensory input generating process. The latter is part of the SPM software suite (Version 12) which can be downloaded from https://www.fil.ion.ucl.ac.uk/spm/software/. All code for the present work can be found at https://github.com/khrv/ADEM_foraging. The simulations were each ran for 450 timesteps, with he following parameter values.

Table 1. Parameters for Experiments

Common Parameters			
act_{min}	0.3	$dist_{min}$	0.8
$\dot{need}_{nutrition}$	-0.005	$\dot{need}_{hydration}$	-0.003
$\dot{need}_{rest}$	-0.002		
Precisions			
Generative Model		External Process	
Π_1^v	e^5	Π_1^v	e^5
Π_1^x	e^5	Π_1^x	e^5
Π_{2needs}^v	e^5	$\Pi_{1proprio}^v$	e^{16}
Π_{2needs}^x	e^5	$\Pi_{1proprio}^x$	e^{16}
Π_{2behav}^v	e^5	Π_2^v	e^5
Π_{2behav}^x	e^5	Π_2^x	e^5

3.1 Experiment 1: Reactionary behavior only

In this experiment we demonstrate the ability of the agent to forage in the virtual environment with solely mid-level behavior which reacts directly to the current state of the agent's needs. Figure 2 shows the results of the experiment.

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

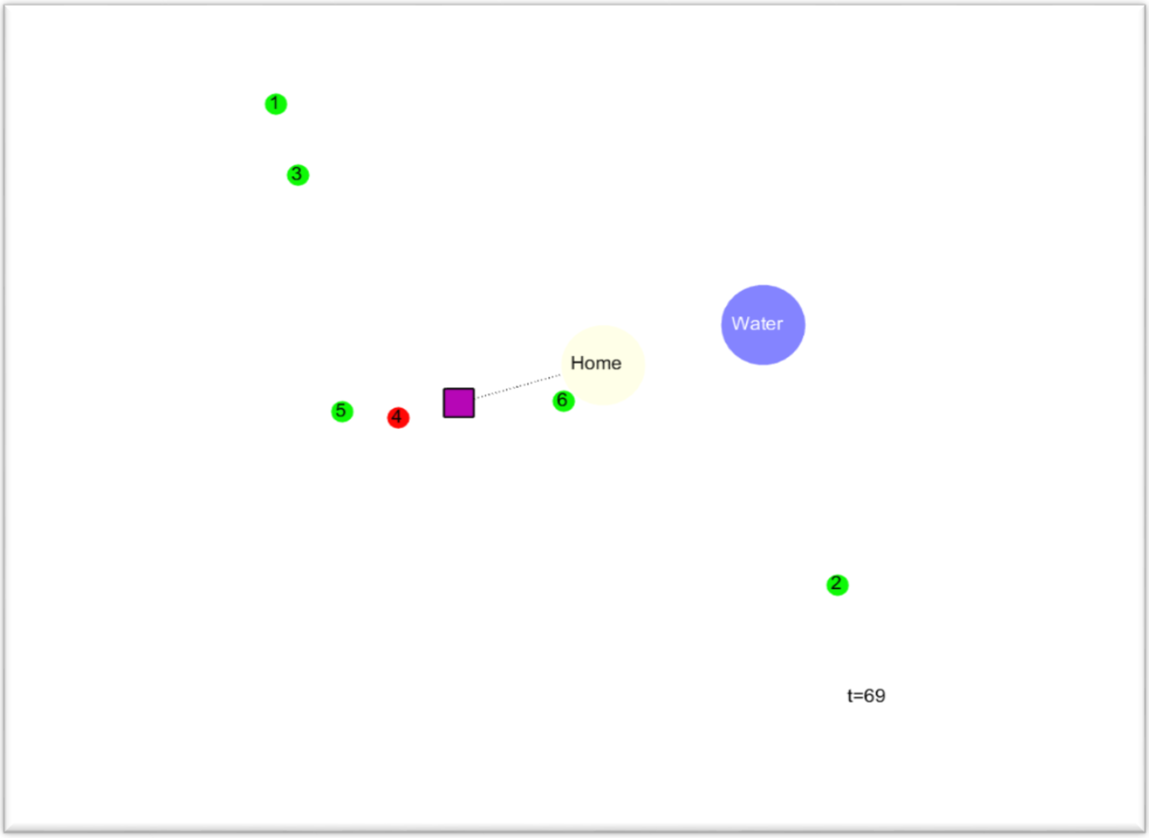

A(I)

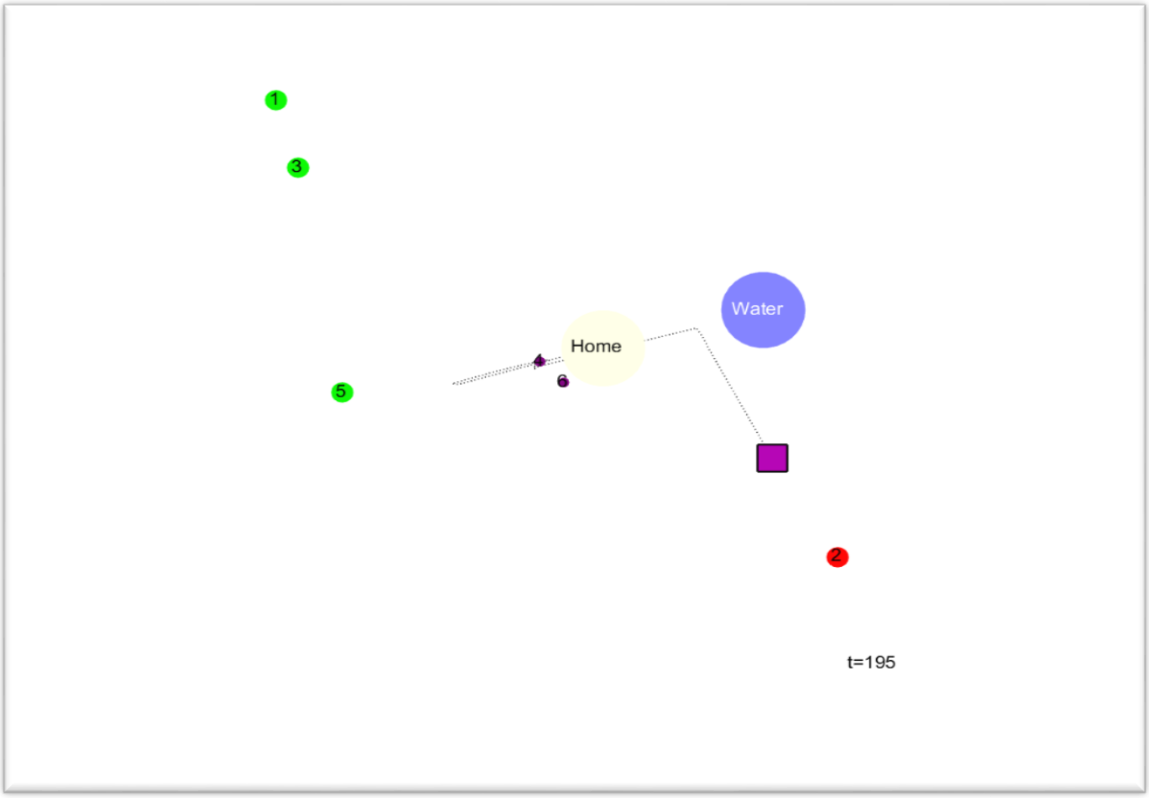

A(II)

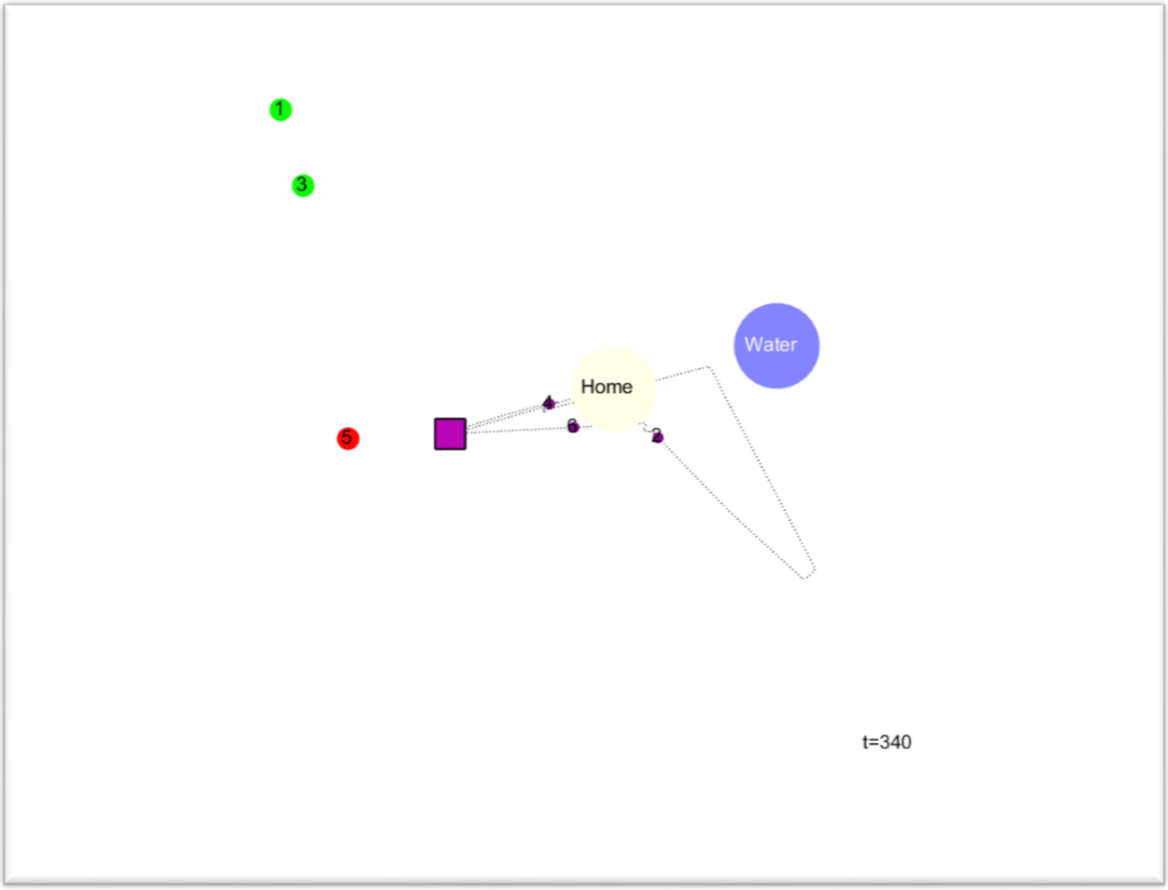

A(III)

Figure 2. Results of Experiment 1

A. Snapshots of agent behavior. (i) Agent (purple square) heads from Home towards Food4 (salient food in red, non-salient food in green). (ii) Agent has taken two food pellets home and eaten them (eaten food shown as small black dots). Agent initially heads for Water but gives up and heads towards Food2 instead. (iii) Agent has brought back and eaten Food2. Agent heads towards Food5, ignoring its thirst. B (Top) Salience of each food item. B (Center) Proprioceptive state of each body part. Grip in blue, Mouth in red, Lie in yellow. B (Bottom) State of each need. Nutrition in red, Hydration in blue, Rest in yellow. C Salience of approaching each target.

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

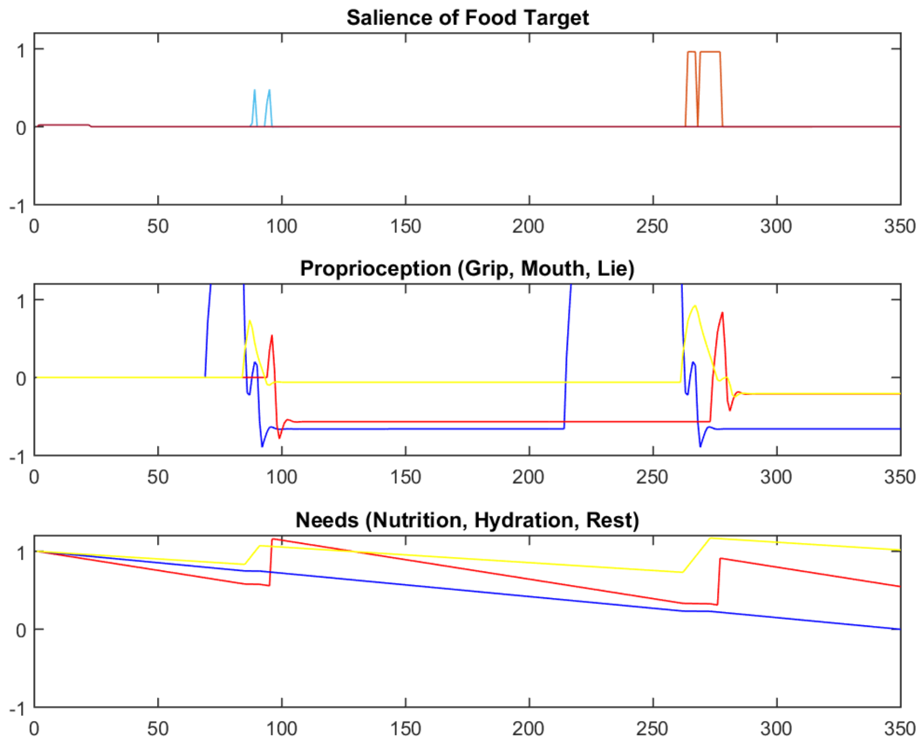

(B)

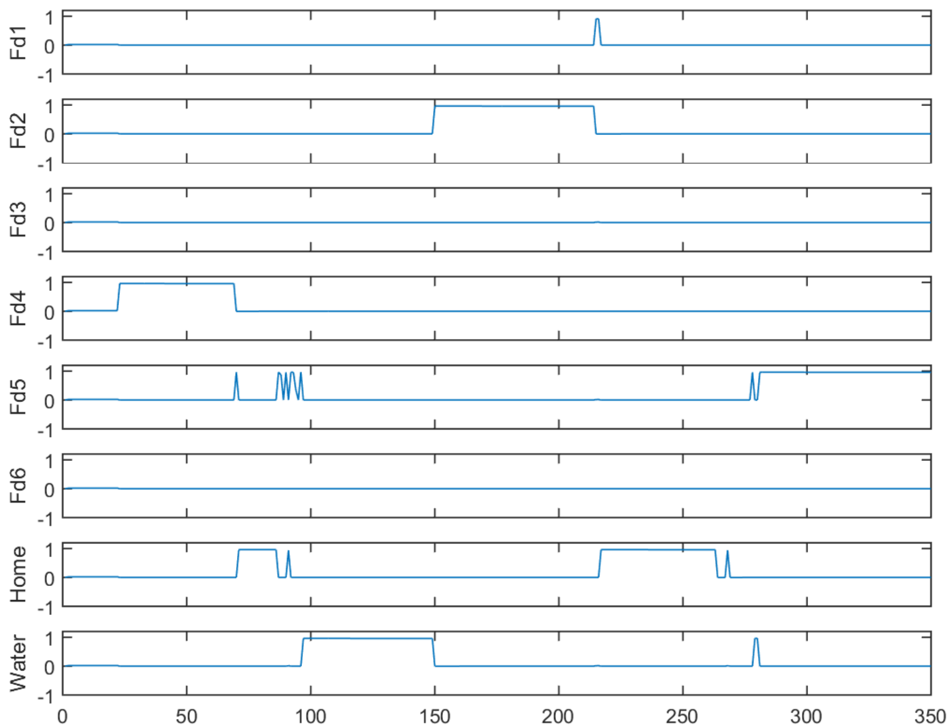

(C)

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

Here we immediately see that the experiment lasted until timestep t=350, at which point the agent died of thirst. Let us investigate why this happened. Figure 2A shows the general trajectory taken by the agent, which displays behavior concentrated on foraging of food and rest but with a general neglect of drinking water. C shows the general order in which the agent approaches each target. The agent first heads toward Food4, and proprioceptive results in the second row of B show us that the agent indeed gripped food at around t=80, at which point C shows us that the agent again heads for Home. Upon reaching Home the proprioceptive results show that the agent lies down, and indeed at around t=90 the Rest need is replenished. Next the agent opens its mouth to eat its recovered food pellet, and its Nutrition need is replenished. By this point the agent's Hydration has decreased a lot and sure enough the agent approaches the Water target at t=100. However due to the Nutrition need decreasing at a faster rate than Hydration (see Table 1), the agent becomes hungry mid-approach and gives up fetching water in order to continue foraging. The agent thus continues grabbing food again until dying of thirst at t=350. We thus see that the inability of the agent to focus on the task of obtaining water leads to its demise. Although it is effectively pursuing adaptive behaviors based on its internal needs and the affordances of the environment, it does so in a short-sighted way. It is able to juggle the Rest and Nutrition needs somewhat successfully as these needs both require sustaining at the Home location, however the Water location is out of the way and requires an extra degree of focus in order to coordinate its approach with the other behavior.

3.2 Experiment 2: high-level behavior

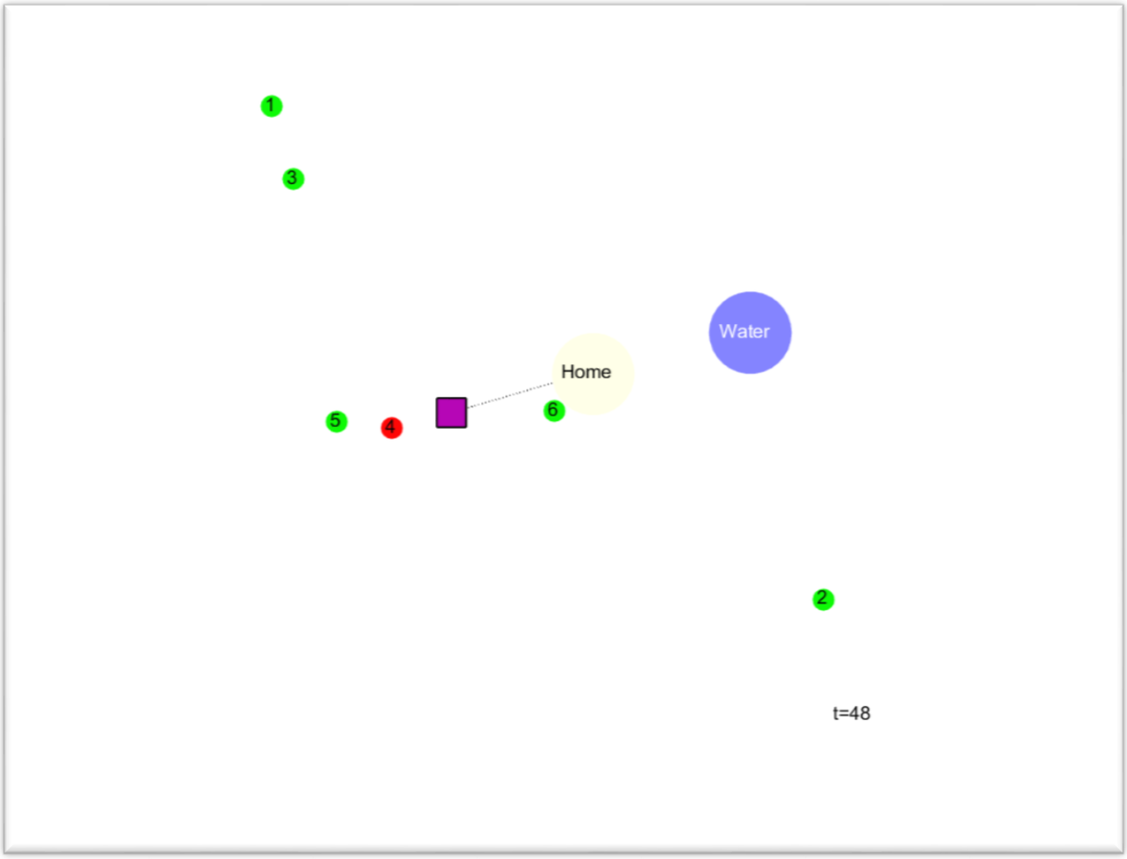

A(I)

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

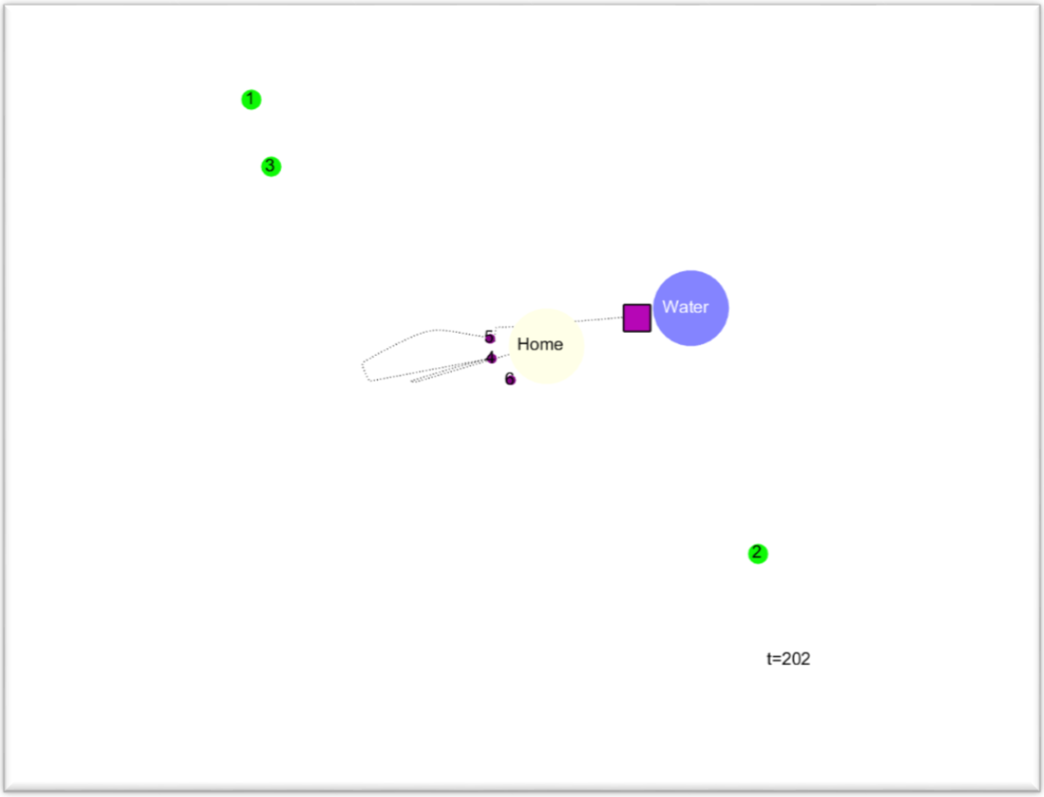

A(II)

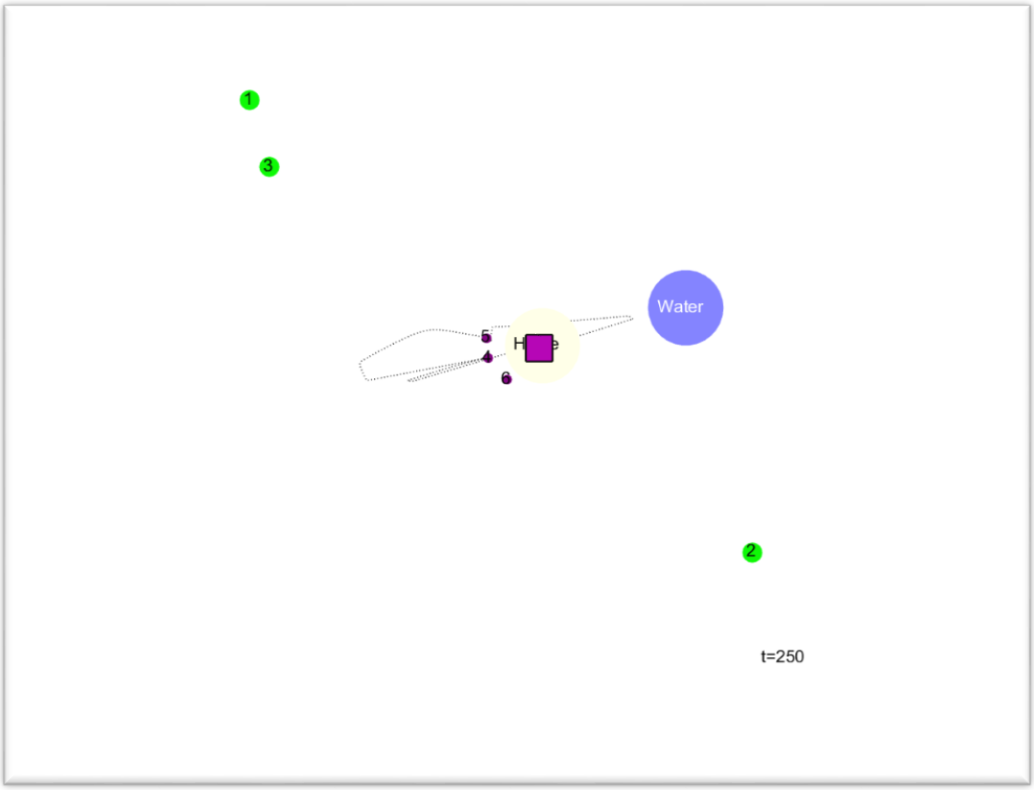

A(III)

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

Figure 3. Results of Experiment 2

A. Snapshots of agent behavior. (i) Agent moves towards Food4. (ii) Agent has recovered and eaten several food pellets and visits the water pool. (iii) Agent has satiated both hunger and thirst and so goes Home to rest. B (top) Salience of each high-level behavior. Foraging in red, Drinking in blue, Sleeping in yellow. B (lower rows), C: See Figure 2.

In this experiment we allow the agent to engage in high-level behaviors (level 2 of the model), which act as metastable states on a longer timescale than level 1. The purpose of these behaviors are to focus the agent towards fulfilling goals that require it to temporary disregard competing goals. In this case, the acquisition of water should take priority over the acquisition of food when Hydration is low even though the need for Nutrition is constantly attempting to gain the agent's attention. Figure 3B (top row) shows the salience of each high-level behavior, with the agent choosing to engage in seeking of food, water, and rest in order. Specifically, Foraging behavior is dominant for the first 150 timesteps, followed by Drinking behavior over the next 50 timesteps, and Sleeping behavior over the following 50 timesteps. From the last row of Figure 3B we indeed see that Nutrition, Hydration, and Rest are all balanced for the duration of the experiment. As outlined in Figure 3C we can see that the agent adequately takes turns in approaching food, water, and home targets as necessary to sustain its needs. Further, Figure 3A shows this process as it plays out graphically, starting with 3Ai at timestep t=48, where the agent approaches Food4. Following this in 3Aii, the agent has brought the food back home and eaten it and now proceeds to the water pool to drink at t=202. Finally in 3Aiii at t=250, after having both eaten and drank, the agent returns home to rest. Thus we see a much more balanced set of behaviors that equally sustain each need at viable levels.

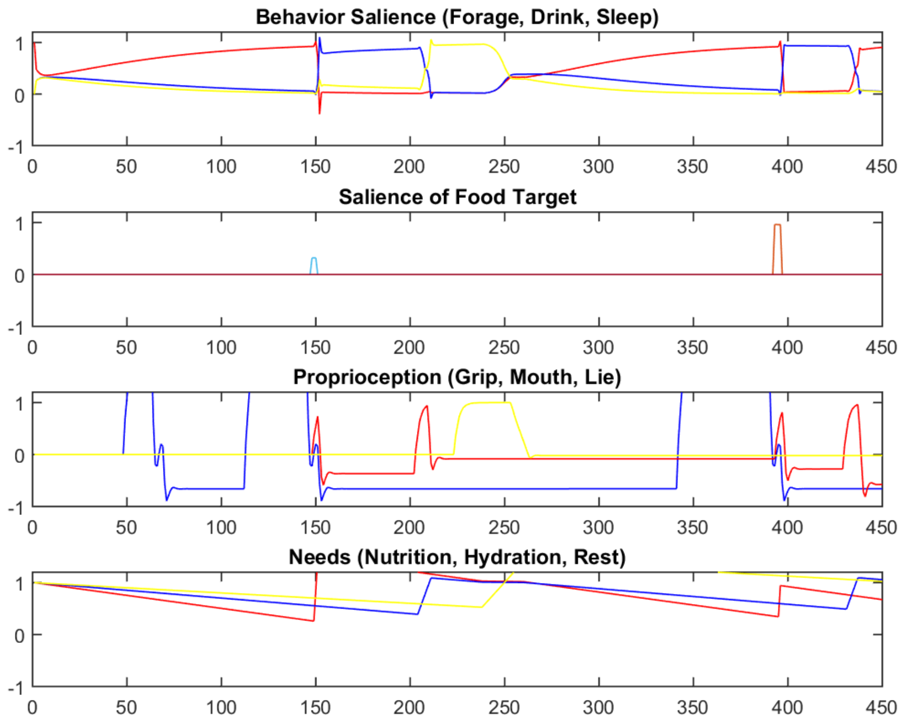

(B)

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

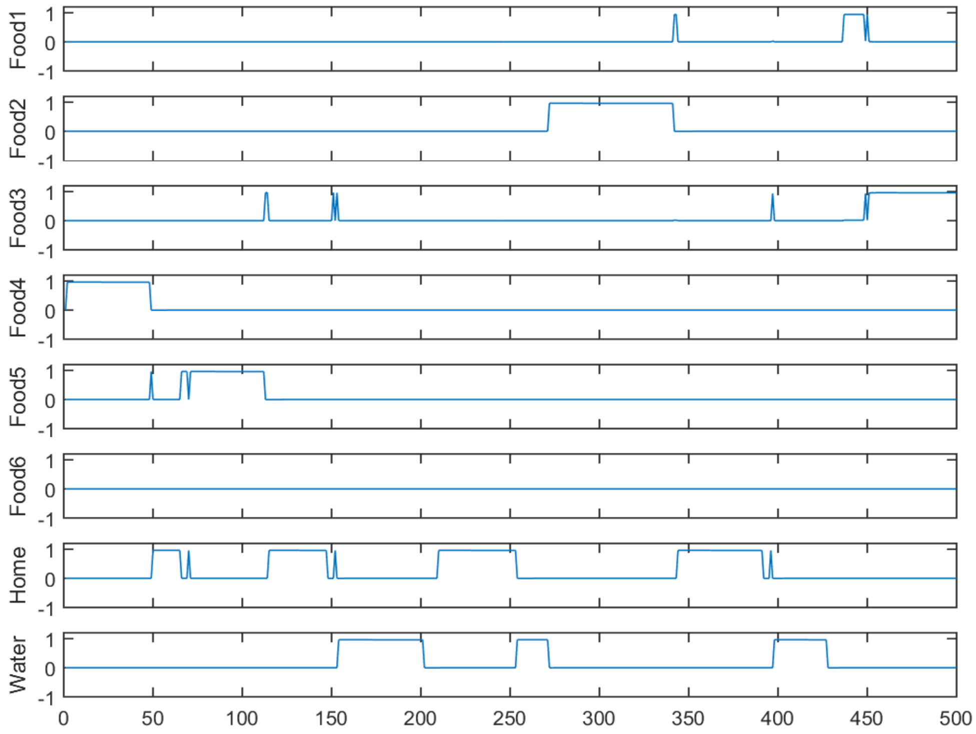

(C)

3.3 Experiment 3: Obstacle Avoidance

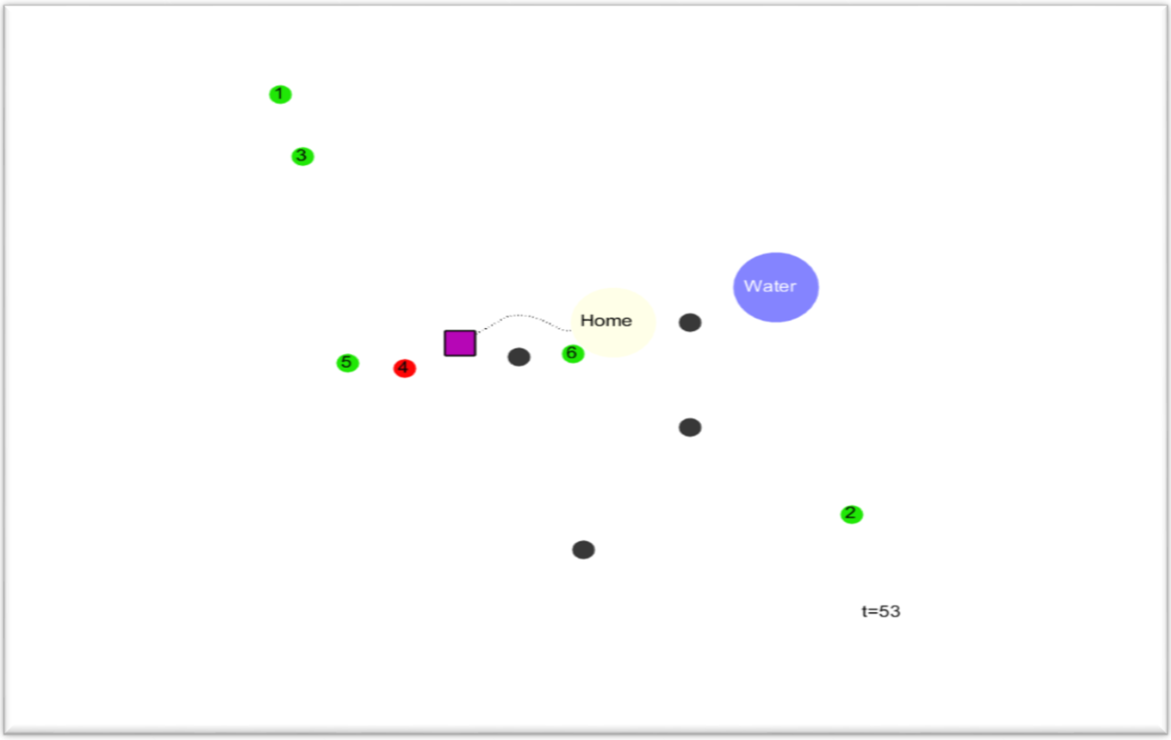

A(I)

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

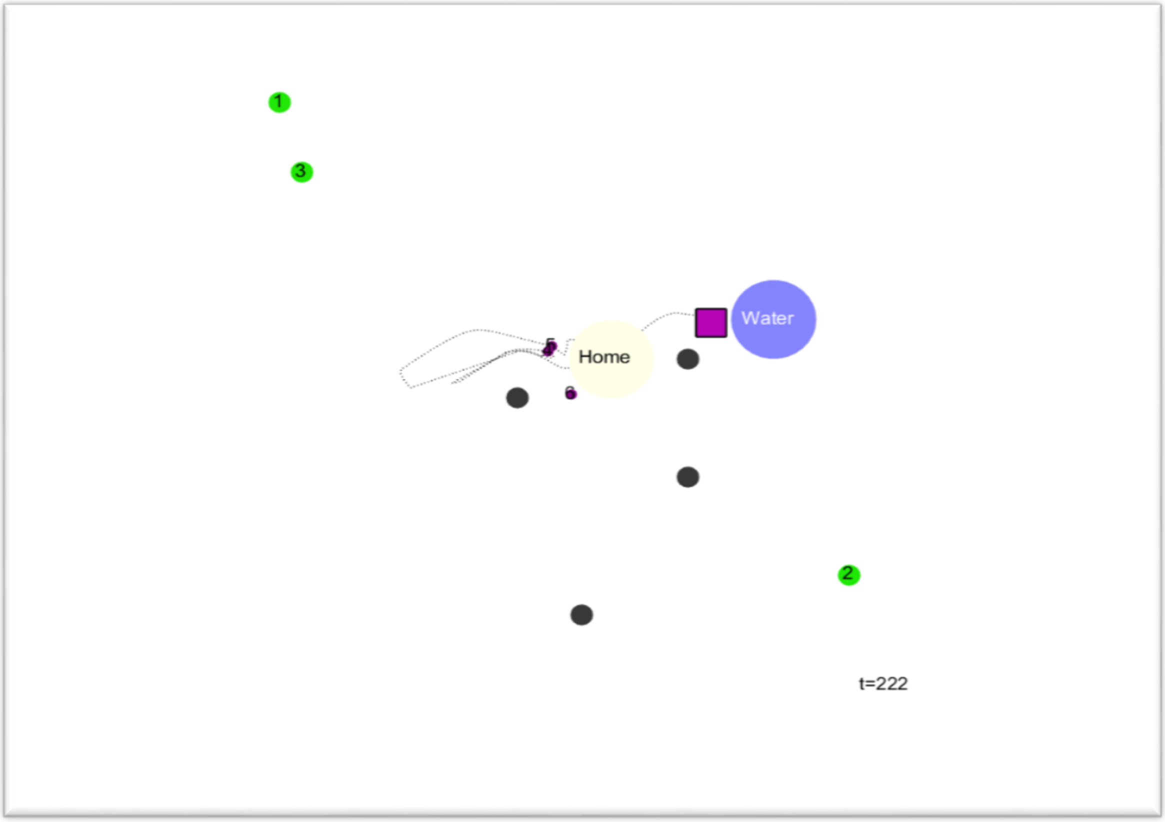

A(II)

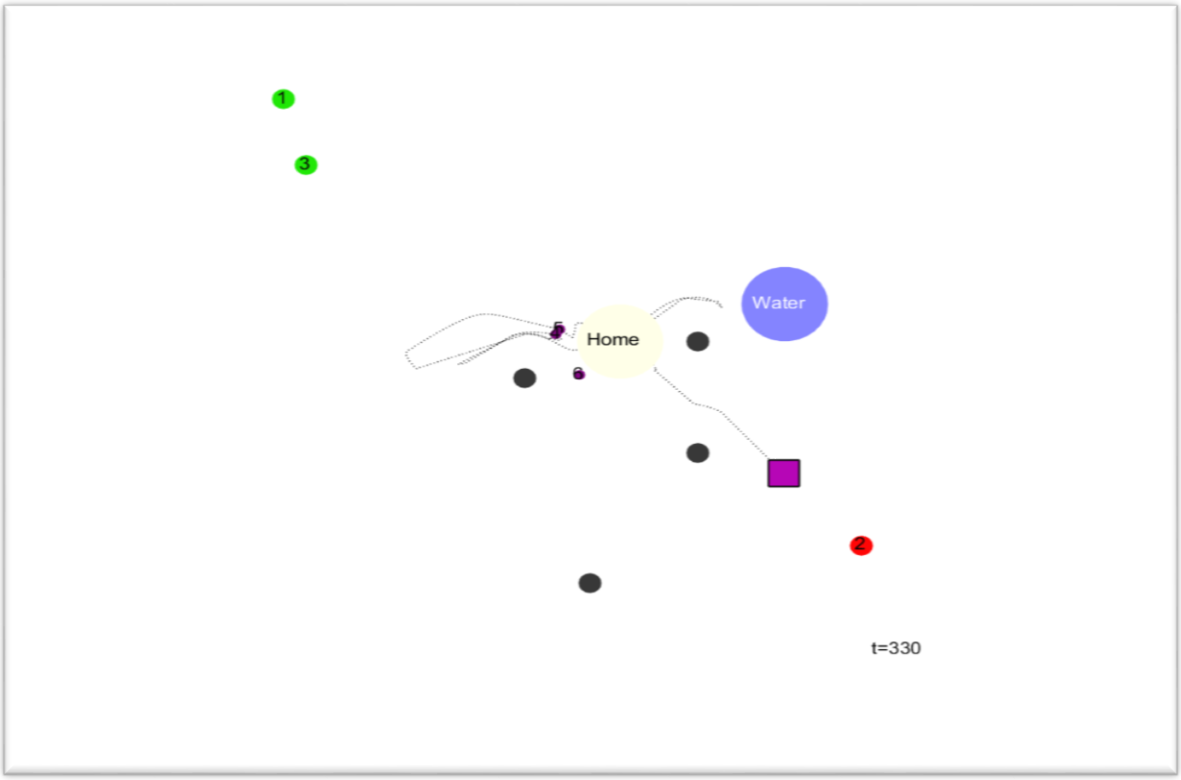

A(III)

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

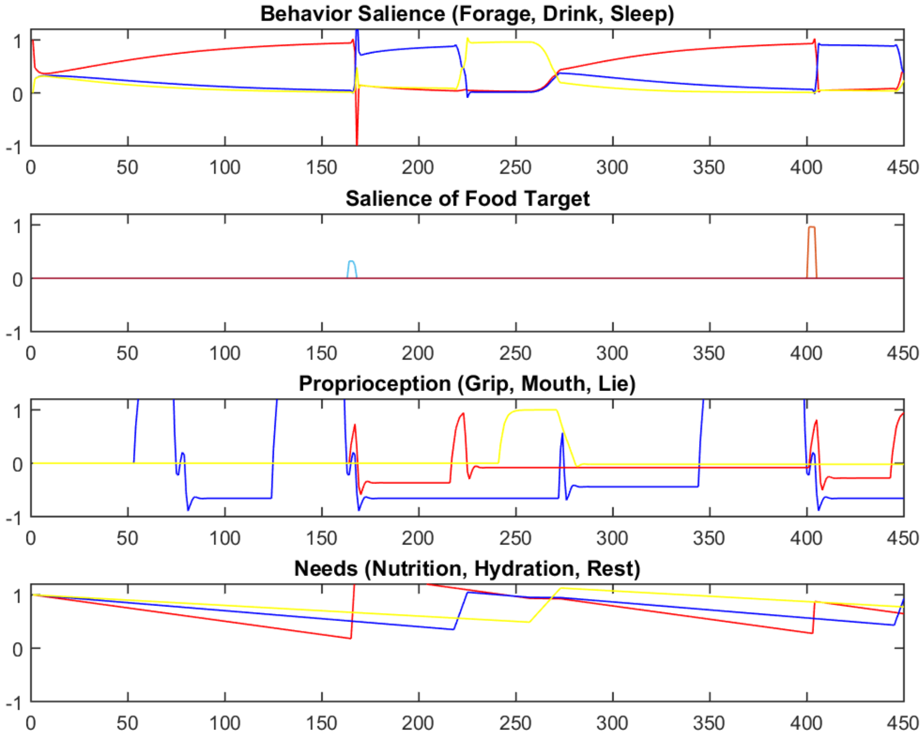

(B)

Figure 4. Results of Experiment 3

A. Snapshots of agent behavior. (i) Agent heads for Food4 while avoiding first obstacle (larger black circle). (ii) Agent has eaten several food pellets and now reaches water after avoiding another obstacle. (iii) Agent avoids obstacle while heading for Food2. B, C: See Figure 2.

In this experiment we again allow the agent to engage in high-level behavior and balance its needs as in Experiment 2, however we now place obstacles in the environment. As seen from Figure 4B and 4C, the agent is able to respond to the affordances in the environment without any noticeable detriment and continue high-level behaviors as if the obstacles were not there. However as seen from Figure 4A, we see that in reality the agent is effectively avoiding obstacles at the motor reflex level. We see the agent maneuver around each obstacle while continuing to approach each appropriate target and engage in foraging, drinking, and resting. It is of note that these micro-actions have not been preprogrammed, but dynamically emerge as the agent responds to affordances for behavior at the top level, approaches and interacts with targets at the mid-level, and engages in immediate navigation around obstacles at the bottom level. As the agent deals with these obstacles in an unreflective and successful way, it never has to change its strategy and so can navigate its environment while at the same time deliberately planning which targets to approach next based on its needs. This shows how seemingly intelligent behavior can come about from a very simple set of behavioral tendencies, organized by metastable states at higher levels and executed by dynamic embodied navigation at lower motor levels.

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

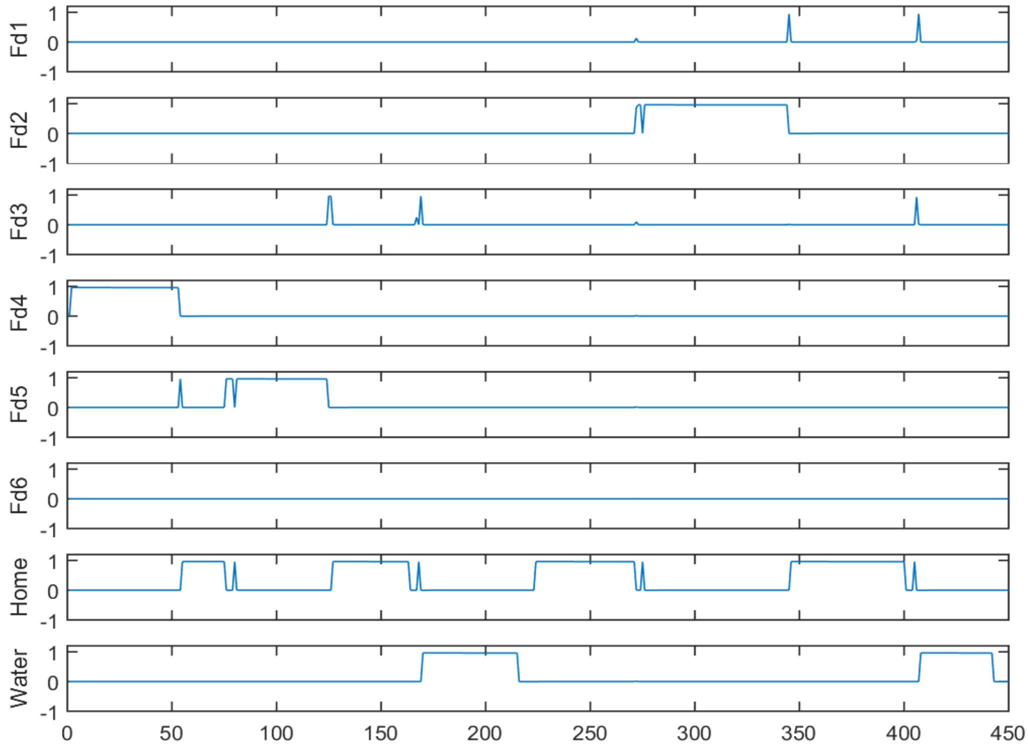

(C)

IV. CONCLUSIONS

In this article we have constructed a foraging agent based on ideas from FEP and affordance competition. The agent we have introduced engages in probablistic inference of the hidden states and causes of the external world, and is led by prediction error minimization to either update its model to be more faithful to the dynamics of the world or to execute actions which cause the world to be more in line with its own predictions. Throughout the course of this prediction making, the agent is influenced further by the state of the external environment and its internal needs. At each step the agent is juggling multiple competing affordances for action, each possibly leading to fulfilment of a different need. The influence of these affordances ultimately leads to predictions about the proprioceptive state of each of the agent's appendages, and the ensuing actions lead to adaptive behavior in the environment.

4.1 Perceiving and responding to affordances

The agent is tasked with not only inferring the personal relevance of each affordance for action available, but also their appropriateness at any paritcular moment. The way this comes about for the agent, although implemented here in an almost rule-like fashion, can be thought of as follows. Phenomenologically speaking the agent is in a sense'directly perceiving' the affordances themselves [12], for example through the Salience matrix in (6). From this viewpoint, the agent is not engaging in a series of cognitive calculations, but is in fact directly making predictions about its own behavior based on the immediate salience of affordances. For example, there is no need to consider the action of grasping when there is no food around, because that affordance is not supported by the present environment. The agent simply has no inclination to grasp at thin air, neither does it have the inclination to grasp at food when it is not hungry, as that is not an adaptive strategy for need reduction. In our implementation these low-level 'appendage' movements are implemented in a rather deterministic fashion, due to the simplicity of the simulation. In a more detailed simulation however, the ideas laid out in this article could be expanded to be more dynamic at the appendage/motor level also, relying more on the low-level physical properties of having a body. This would give credence to the idea of the 'embodied mind' [13], as heeding affordances could be done on a sub-behavior motor level as an 'embodied decision' [14]. We have attempted to implement the idea of sub-behavioral level affordances in our paper through the agent's ability to avoid obstacles, which in practice is implemented as manipulation of low-level trajectory predictions based on sensors that detect the obstacles in the environment. In a future expansion of the present theory we believe we could replace its more deterministic aspects with this kind of dynamic 'prediction alteration' on multiple behavioral levels based on the present affordances and their salience. Such dynamic interaction between environment and predictive model woud likely require a more complex simulated body, as the intelligence of the agent begins to be offloaded into the body and environment in true embodied fashion.

4.2 The behavior hierarchy

Through our experiments we found that it is important for the agent not just to utilize salient affordances in a greedy myopic way, but rather to focus attention on restoring particular needs at a time. The way in which our agent does this is through the use of high-level behaviors which take on a metastable dynamic and change at a slower rate than the underlying actions, thus stabilizing the lower level and ensuring that the agent does not respond to the environment in an erratic, overly reflex-like fashion. We demonstrated how the agent is capable of balancing three separate needs: nutrition, hydration, and rest, each with its own particular affordances that required a particular behavioral response from the agent. Namely, the agent was able to acquire food and bring it back to its nest to consume, approach the water bath in order to drink, and approach the nest and lie down in it in order to sleep. While the set of needs employed here is rather basic, we believe that the same principle of hierarchical behavioral management could apply to more complex agents, allotted with hierarchical generative models split into layers determining the predictions made at each time scale. Additionally, as outlined above, the lower levels could have a more in-depth relation to the agent's body and external environment. In this way we could emulate the sort of low-level 'reflex' to high-level 'cognitive' progression of behavior as is believed to exist in animals, and in particular
human beings.

4.3 Dealing with the unexpected

In the present simulation, the agent does not have to deal with hugely unexpected situations. For example, while it may drop food as it is carrying it and approach it again to pick it back up, the behavior still remains within the realm of what Heidegger would call 'readiness-to-hand ' [15], or purely unreflective behavior. However, when the truly unexpected happens (say, the agent is caught at a wall of obstacles and cannot simply maneuver around them), it would be necessary to send the prediction error which is not dealt with at the action level up the hierarchy so that it may be resolved by changing behavior at a higher level. In this way we see the beginnings of a more detached or 'present-at-hand ' style of behavior, which we believe is key to truly intelligent agents that are capable of forming mid to long term plans and dealing with the unexpected in a reasonable and adaptive way. Nevertheless we believe we have provided an interesting starting point for developing agents based on the FEP which can respond to aspects of the environment and their own physical body in a dynamic and adaptive way, while adeptly managing multiple competing needs in a balanced fashion.

V. COMPLIANCE WITH ETHICAL STANDARDS

5.1 Ethical Approval

This article does not contain any experiments with human or animal participants performed by the author.

5.2 Informed Consent

Does not apply.

5.3 Funding

No funding was received for this study.

REFERENCES

[1] Paul Cisek. "Cortical mechanisms of action selection: the affordance competition hypothesis." In: Philosophical transactions of the Royal Society of London. Series B, Biological sciences 362.1485 (2007), pp. 1585–99. issn: 0962-8436. doi: 10.1098/rstb.2007.2054. url: http://www.ncbi.nlm.nih.gov/pubmed/17428779.

[2] Kole Harvey. "Artificial Dasein: Solving the Frame Problem with Incremental Contextualization". In: OALib 05.04 (2018), pp. 1–17. issn: 2333-9721. doi: 10.4236/oalib.1104535. url: http://www.oalib.com/paper/pdf/5293435.

ISSN: 2277-9655
Impact Factor: 5.164
CODEN: IJESS7

[3] Kole Harvey. "An Open-Ended Approach to Piagetian Development of Adaptive Behavior". In: OALib 05.03 (2018), pp. 1–33. issn: 2333-9721. doi: 10.4236/oalib.1104434. url: http://www.oalib.com/paper/pdf/ 5293221.

[4] Karl Friston. "Embodied inference and spatial cognition". In: Cognitive Processing 13.S1 (2012), pp. 171–177. doi: 10.1007/s10339-012-0519-z. url: http://link.springer.com/10.1007/s10339-012-0519-z.

[5] Erik Rietveld and Julian Kiverstein. "A Rich Landscape of Affordances". In: Ecological Psychology 26.4 (2014), pp. 325–352. issn: 1040-7413. doi: 10.1080/10407413.2014.958035. url: https://www.tandfonline.com/ doi/full/10.1080/10407413.2014.958035.

[6] K´evin Roche and Hanna Chainay. "Is there a competition between functional and situational affordances during action initiation with everyday tools?" In: Frontiers in Psychology 8.JUN (2017). issn: 16641078. doi: 10.3389/fpsyg.2017.01073.

[7] Jelle Bruineberg et al. "A Rich Landscape of Affordances". In: Frontiers in Human Neuroscience 16.3 (2008), pp. 1–48. issn: 1662-5161. doi: 10.3389/fnhum.2015.00237.

[8] Karl Friston. "The free-energy principle: a rough guide to the brain?" In: Trends in cognitive sciences 13.7 (2009), pp. 293–301. issn: 1364-6613. doi: 10.1016/j.tics.2009.04.005. url: http://www.ncbi.nlm.nih. gov/pubmed/19559644.

[9] Karl Friston and Stefan Kiebel. "Cortical circuits for perceptual inference." In: Neural networks : the official journal of the International Neural Network Society 22.8 (2009), pp. 1093–104. issn: 1879-2782. doi: 10.1016/ j.neunet.2009.07.023. url: http://www.ncbi.nlm.nih.gov/pubmed/19635656.

[10] Karl Friston. "Hierarchical Models in the Brain". In: PLoS Computational Biology 4.11 (2008). Ed. by Olaf Sporns, e1000211. issn: 1553-7358. doi: 10.1371/journal.pcbi.1000211. url: http://dx.plos.org/10. 1371/journal.pcbi.1000211.

[11] M. J. Spivey, M. Grosjean, and G. Knoblich. "From The Cover: Continuous attraction toward phonological competitors". In: Proceedings of the National Academy of Sciences 102.29 (2005), pp. 10393–10398. issn: 00278424. doi: 10.1073/pnas.0503903102. url: http://www.pnas.org/cgi/doi/10.1073/pnas.0503903102.

[12] James J. Gibson. The ecological approach to visual perception. Boston: Houghton Mifflin, 1979.

[13] Andreas K. Engel et al. "Where's the action? The pragmatic turn in cognitive science". In: Trends in Cognitive Sciences 17.5 (2013), pp. 202–209. issn: 13646613. doi: 10.1016/j.tics.2013.03.006.

[14] C Burr. "Embodied decisions and the predictive brain". In: Philosophy and predictive processing November (2017), pp. 102–124. doi: 10.15502/9783958573086.

[15] Martin Heidegger. Being and time: A translation of Sein und Zeit. SUNY press, 1996.

CITE AN ARTICLE

Harvey, K. (2018). FREE ENERGY FORAGING IN AN AFFORDANCE LANDSCAPE. *INTERNATIONAL JOURNAL OF ENGINEERING SCIENCES & RESEARCH TECHNOLOGY*, *7*(8), 450-469.