Behavioural Funnelling

John Pisokas
Computer Science Department
University of Essex
Wivenhoe Park
Colchester, CO4 3SQ, UK
Email: john.pisokas@gmail.com

***Abstract*— The process of planning ahead the future actions of a robot faces a major problem in real word operation: the actual result of each action is not always predictable; that is, there is uncertainty in the interaction between the robot and its environment. Historically, there have been two approaches regarding this kind of uncertainty, the first is to assume that the robot-world interaction is deterministic while the second is to assume that the robot-world interaction is stochastic. Neither of those two assumptions is realistic in real world because they attempt to address only the two extremes of the spectrum of possible cases. Another problem is that historically the assumption that robot-world interaction is stochastic comes with the assumption attached that the observed stochasticity in this interaction is something inherent and unavoidable without the use of clever algorithms.**

This article takes another approach; it sets out a design methodology which leads to naturally quasi-deterministic robot-world interaction. In this way none of the above assumptions need to be made and the robots are more efficient using available resources.[1]

I. INTRODUCTION

Our aim is to build robots capable of deciding their next action and possibly planning ahead their actions. There is a major problem in this process: the actual result of an action is not always predictable because there is uncertainty in the interaction between the robot and its environment. Historically, there have been two approaches regarding uncertainty in real world operation, the first is to assume that the robot-world interaction is deterministic and the second to assume that the robot-world interaction is stochastic.

None of those two assumptions is realistic in the real world, they rather attempt to address only the two extremes of a spectrum of possible cases. Another problem is that historically the assumption that robot-world interaction is stochastic comes with the attached assumption that the observed stochasticity in this interaction is something inherent and unavoidable without clever algorithms. As result, several attempts have been made in addressing the apparent stochasticity of the interaction in software with various kinds of probabilistic modeling and processing techniques. Some examples are provided by [1], [4].

This article does not propose an alternative which promises better performance than the previously mentioned design approaches; it rather proposes a rethinking for wiser use of resources. This article holds the position that building robots in ways that make their interaction with the world stochastic and then building stochastic models to undo the original "mistake" is not clever, is not efficient engineering, and abuses resources which render the robots energy voracious and thus not as efficient as they could be. This, from an engineering stand point, is not an appropriate start and it is eminent that this approach will lead to bottlenecks to what our robots will be able to do in the future when more than moving around and taking pictures will be required. The approach used is to set out a building methodology which leads to naturally deterministic (or quasi-deterministic) robot-world interaction. In this way none of the above assumptions need to be made and the robots are efficient using resources.this article we address the uncertainty issue only in the domain of spatial motion tasks.

II. DETERMINISM AND STOCHASTICITY

In this article the result of a behavior will be called *deterministic* when the outcome of this behavior can be securely predicted before activating that behavior. The accuracy of the prediction is arbitrary, in other words can vary without affecting this definition, as far as it is specified. The result of a behavior will be called *stochastic* when this behavior's outcome cannot be uniquely predicted but rather can only be expressed by a probability distribution of various outcomes. These definitions capture the phenomenon only to its extremes; we call all other cases *quasi-deterministic*. The *level of quasi-determinism* is a human constructed concept for substantiating, measuring, and studying an observed phenomenon. The observed phenomenon is that it appears as if different organisms and their various behaviors have different levels of probability to give a specific outcome.

Provided that one is able to quantize the spectrum of possible outcomes of a behavior into S discrete outcomes (states), and assuming that one can obtain the probability $P(s)$ of each state $s \in S$ occurring, we can use the well known concept of *entropy* for measuring the level of quasi-determinism. Due to practical limitations in measuring multitudes of probabilities, the practice of measuring the probability of occurrence of only the *desired outcome* will suffice. To this end one should first define the *desired outcome* of a behavior as well as an appropriate measure of success in achieving this desired outcome. For example "the robot transverses a corridor to its end within 5*min*". Then several runs of the behavior under

---

[1] Published as: John Pisokas, Behavioural Funnelling, University of Essex, September 2006, CSM- .

different initial conditions must be performed until a satisfying approximation of the probability of the desired outcome $P(s_1)$ is estimated. Therefore, the probability of the non-desired outcome group occurring is $P(s_2) = 1 - P(s_1)$. Therefore, the set S will have only two states, the desired outcome group s_1 and the undesired outcome group s_2; thus the base of the logarithm in the calculation of the entropy will be 2:

$$H(P) = -\sum_{\forall s \in S} P(s) \cdot log_2 P(s) \tag{1}$$

$H(P)$ expresses the uncertainty about the state of the system, with the minimum possible value, 0, signifying exact knowledge of the outcome state; and the maximum value, 1, signifying that all possible outcome states $s \in S$ are equiprobable and therefore one does not know which is the state of the system. For our purposes from $H(P)$ we derive the quasi-determinism level, $D(P)$:

$$D(P) = 1 - H(P) \tag{2}$$

The value range 0–1 of $D(P)$ is used as measure of quasi-determinism level, with value 0 signifying a behavior with purely random outcome and the value 1 signifying a behavior which has an always reproducible effect — deterministic. We shall mention that this quantitative quasi-determinism measure is useful for the comparison of different behaviors but it is only applicable when the same criterion of accomplishment can be used for all compared behaviors.

III. FUNNELING BEHAVIOR

If a behavior systematically increases the level of quasi-determinism of the robot's position or leaves it as it is we call it a *funneling behavior*. If a behavior systematically decreases the level of quasi-determinism we call it a *divergent behavior*. To determine whether a behavior is funneling or diverging, we take the difference $D_t - D_{t-1}$ where D_{t-1} is the quasi-determinism level before activating the behavior and D_t the quasi-determinism level after the behavior has ended; those quasi-determinism level values are derived by the equation 2. In order to have a value for D_{t-1} either a previous behavior must have finished and then followed by the current behavior without any changes in the state of the robot in the meanwhile, or an appropriate probability distribution of initial states must be derived. In practice the method used for estimating the probability distribution before the activation of the first behavior, is to run experiments with many different initial conditions with equal amounts of runs for each initial condition so that the distribution is uniform and thus $D_{t-1} \rightarrow 0$.

IV. BUILDING BLOCKS OF FUNNELING BEHAVIOR

The preceding paragraphs indicate that for achieving funneling behavior we must select and use only behaviors which have highly repeatable outcome. Since there are a multitude of behaviors that one can envisage and implement, we must short list the ones which have repeatable effect. To this end we short listed those spatial behaviors which are commonly used by animals and have highly repeatable effect.

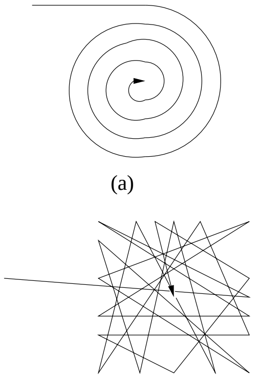
(a)

(b)

Fig. 1. Examples of sensor-less and interoceptive piloting behaviors.

Therefore, we are looking for patterns of behavior which lead the organism to specific desired locations (or states in general). In some cases these locations are not known in advance by the organism, but their attributes are desired; in other words any of several locations with the desired attributes is acceptable. In other cases the location is perceptually unique and the organism can in principle identify this unique desired location; of course perceptual aliasing problems will obstruct an organism from doing so but combination of identification with observation of which other places are around the desired should eventually reveal if the detected place is the desired one, as [3] has illustrated.

Behaviors of various animals — ranging from relatively simple life forms (protozoa, crustaceans) to humans — have been studied in the past and classified in the context of orienting behaviors. From those classes of behavior we distinguish those which have highly repeatable outcome and regroup them in the context of funneling behaviors.

A. Sensor-less and Interoceptive Sensing Behaviors

The first group of funneling behaviors is *sensor-less piloting* and *interoceptive piloting behaviors*. These are behaviors which do not employ exteroceptive sensors; these behaviors only use interoceptive information or no sensory information for their trajectory determination. To be more precise, they still have some means for identifying their goal place, for instance a food source or their nest, which usually involves exteroceptive sensors dedicated to this purpose. Some two-dimensional examples of these behaviors are: opening spiral trajectory in search for the nest, random walk of a restricted span in search for food-source or nest, and up-down flight in a line in search for a food-source or nest; figure 1 gives some examples. This kind of behavior is found in several animals including honeybees, desert isopods, and desert ants.

Another form of sensor-less trajectory determination is a train running on railway lines until reaching its destination. This train will not need active correction of its trajectory since the confinement provided by the railway lines and the shock absorption by the mechanical dumping system will drive the car. In all these cases only interoceptive or no sensing is required for the trajectory determination while exteroceptive sensing is used only for identification of the goal and consequent ending of the current activity.

B. Kinesis Combined with other Capabilities

The second group of funneling behavioral patterns is *kinesis* orienting behaviors in combination *with particular behavioral or sensing capabilities*. In the example we describe here the *klino-kinesis* type of kinesis is used. In *klino-kinesis* the organism's rate of random turning or angular velocity depends on the intensity of stimulation. For klino-kinesis a single receptor sensing the stimulus intensity is sufficient. Klino-kinesis behavior alone does not produce motion towards a goal location but when combined with a response hysterisis operates as funneling behavior leading the animal systematically to a favorable area of its environment. An example is that of the planarian *Dendrocoelum lacteum*, when there is no suitable stimulus directing the animal it does not move far in a straight line and occasionally turns. In constant lighting conditions its turning frequency is constant. If the light intensity is increased the animal will increase its rate of turning. This is a *klino-kinetic* response. In the sequel the increased rate of turning will fell with time and after approximately 30*min* it will return to the base rate. The combination of the *klino-kinetic* response with the decay time of the rate of turning can lead the animal to the darkest part of an illuminated area. The experimental setting for observing this behavior is a horizontal flat surface illuminated from above in such a manner that a light intensity gradient is formed over the flat surface. If we evenly scatter planarians over the flat surface and the light gradient steepness is suitable, after 1-2 hours the planarians will be aggregated in the darkest end of the surface. The mathematical analysis of this process is a statistical one and is given by Patlak ([6], [5]). Figure 2 focuses on a couple of full turns of the animal; assume that the animal starts at point O where it is fully adapted to the local light intensity and thus the rate of turning is the base one. Initially the animal moves along the line segment OA with light intensity being constant so that at A it turns due to the base rate of turning. Then the animal moves along AB and the light intensity increases and as a result the rate of change of direction increases and the animal turns sooner than normal at point B. Now for simplicity, assume that full adaptation to the local light intensity occurs exactly when turning at B so that the BC distance is the normal one as CD and DE are too. The same events repeat on the trajectory EFGH moving each time the animal more and more towards the right bottom of the area. This simplified explanation is far from accurate but the principle is the same with the real animal's behavior and produced trajectory which leads the animal to the darkest region of the surface. The reported strategy is a funneling behavior because it increases the probability of an organism to be at a desired region of the environment; this strategy can be also used in robotic applications.

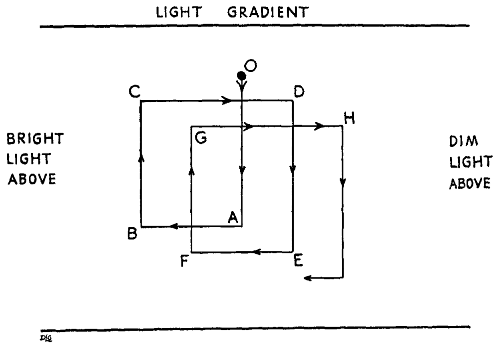

Fig. 2. Trajectory produced by combination of klino-kinesis with stimulus memory.

C. Taxis Behaviors

The third group of funneling behavioral patterns are *taxis* orienting behaviors. These are classified based on the sensory structure and produced trajectories, into three classes: *klino-taxis*, *tropo-taxis*, and *telo-taxis*[2].

In *klino-taxis* the organism turns in response to the change of intensity of the stimulus. Klino-taxis can lead away from or to a stimulus source. To this end a single sensor suffices and it does not need to register the direction towards which the stimulus lies. The commonly referred example of klino-tactic behavior is that of the maggot of certain common flies. These animals after having finished feeding move to a dark place where they pupate. For moving from the food source to the dark place the maggot crawls forward, while turning its head (with a light intensity sensor on it) towards the side with less light intensity. The body follows the orientation of the head and therefore the maggot eventually moves away from light sources. For a more detailed analysis the reader can refer to [2] on pages 59–63. Klino-taxis is commonly found as a reaction to chemical stimuli.

In *tropo-taxis* at least two symmetrically placed sensors are used and the organism orients itself so that it receives equal stimulus to both sensors. Depending on whether the organism performs positive or negative taxis it moves towards or away from the stimulus, respectively. Asymmetry in stimulation of the two receptors activates turning reflexes until the two eyes are equally stimulated, in the oriented position . It is experimentally supported that the mechanism operates by simultaneously comparing the activation of the two receptors ([2], [8]).

In *telo-taxis* the organism's receptor(s) can register the direction towards which the stimulus lies and the organism moves directly towards the stimulus source. Of course this kind of taxis is only applicable to directional stimuli, eg light, infrared, ultraviolet electromagnetic radiation. For the production of telo-tactic behavior there is no need for sequential

---

[2] This classification is adhered as collected and presented by Fraenkel and Gunn (1960), without adhering the assumption that these behaviors are pure reactions to stimulus with the nervous system operating as a switchboard. We rather adhere them in the context later developed and outlined by Sche 1984.

in time sensing of intensity or simultaneous comparison of stimulation of two or more receptors but instead sensing of the stimulus direction. An example is that of the isopod, *Aega*, which is performs positive photo-tropo-taxis under certain conditions ([2]).

We should mention that taxis behaviors do not involve shape, object, or place recognition ([2] page 97). Those three cases of taxis behavior can lead an agent systematically and reliably to a goal region, increasing the level of quasi-determinism, thus they are funneling behavioral patterns.

D. Problems

The above behavioral patterns are susceptible to environmental influences which for humans are of minute importance. For example a light reflecting object in the environment will significantly alter the trajectory and end location of an organism which uses kinesis or taxis orientation behaviors. This means that these behaviors often lead to the same location but with “small” changes in the environment they will lead to alternative solutions. This is due to both perceptual and motor limitations of those methods for behavior implementation. Of course such confusing objects and situations might not be frequent in natural environments but a robot should be able to operate in environments constructed by humans; therefore, these behaviors alone will not suffice for all cases. For this reason we list some additional behaviors which are less susceptible to environmental influence.

E. Guarded Motions

Guarded motions are motions which last *until* at least one of a set of conditions is satisfied. An example is the behavioral description: “move along the road *until* your battery charge level runs low, or *until* you reach the end of the road, or *until* you sense an obstacle in proximity, or *until* you overheat”. Satisfaction of any of those conditions will result to ending the current action. Guarded motions/actions are not producing funneling by themselves, but when they are combined with specific behavioral patterns they assist in producing funneling behavior. In practice, guarded motions operate by imposing limits on the possible values of the degrees of freedom of the agent.

F. Path Following

Path following behaviors consist of transversing paths, both semantic and structural ones. These are similar to trail following behaviors (implemented as klino-taxis or tropo-taxis) with the difference that they require more advanced sensory processing, perceptual capability, and behavioral skills. An example of such path following behavior is a pavement following behavior. In this case there is no a chemical trail to follow, but the limits of the pavement. These limits are mostly semantic than structural. That is, it is easy to walk on the street but its conventional use is for cars not for humans, resulting in a higher risk to be harmed if you walk on the street, therefore one walks on the pavement. A pavement following behavior will likely involve some means for detecting the limits of the pavement — eg visual sensing, LASER range finding, or other means. In addition, appropriate processing of these pieces of information together with the internal states of the agent must produce behavior such that the agent proceeds along the pavement without diverting from it. Another example of path following is following a river.

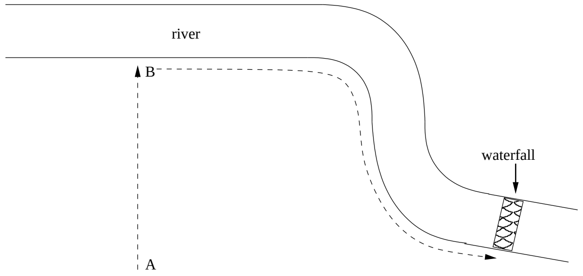

Fig. 3. Direction following behavior example. A person starts at A and moves on the direction indicated by the vector AB *until* he meets “the river”. The river is not visible from A and from most of the path AB. This is a guarded *direction following* behavior; if in the sequel another kind of behavior follows —such as the *path following* behavior “follow the river in the direction of flow”—the exact point B at which the person meets the river does not matter; this is illustrated in the example of figure **??**.

When a path following behavior is combined with a stopping condition then it is called guarded path following behavior and it leads to funneling of the agent’s possible positions to a group of desired final states, an example will illustrate this in the next section.

G. Guarded Direction Following

In *guarded direction following* behaviors, the agent sets out in a straight line towards a certain initial direction *until* a stopping condition is met. Keeping moving in straight line can be based on either interoceptive sensory information or exteroceptive sensory information or both. In the case of interoceptive sensing the common case is use of dead reckoning while in the case of exteroceptive sensing spatial reference can be provided by celestial cues — such as sun, stars —, other entities — such as trees, mountains, lakes —, visual flow, or human made tools — such as magnetic compasses. Each of these methods will result in different deviations from the straight line trajectory but even the most deviating ones will be successful if the stopping condition can be readily identified — for instance by being a much larger entity than the maximum possible deviation of the agent’s trajectory from the straight line. These facts are well known among mountaineers and navigators who regularly use them in practice. These facts are also successfully used by animals; for instance, honeybees use the sun as reference when they travel in straight line for given distance [7].

An example of *guarded direction following* behavior is illustrated in figure 3. A person starts at A and moves on the direction indicated by the vector AB *until* he meets “the river”. The river is not visible from A and from most of the path AB. This is a *guarded direction following* behavior. We can avoid making assumptions on the exact location of point B along the river by employing a subsequent behavior, such as the *path following* behavior “follow the river in the direction of flow *until* you reach the waterfall”. In this case the exact point B at which the person meets the river does not matter. This is an important consideration since the *guarded direction following* leading the agent from A to B limits the degrees of freedom of the agent so that the agent lies at the bottom coast of the river but no assumption on the exact point can be made.

It should be mentioned that the stopping condition of a *guarded direction following* behavior need not to be a landmark, it can alternatively be the distance traveled by the organism like it is the case with honeybees [7] which travel straight for a certain distance. Concluding, guarded direction following behaviors are primitives which can assist to building funneling behaviors when combined with guarded path following or goal directed behaviors.

H. Goal Directed

Goal directed behaviors have two elements: first the goal is sensed and can be identified from adequate portion of the positions laying on the approach path so that the agent can keep track of it and second the behavior of the agent is such that eventually leads to that goal. We must mention that although it might seem at a first glance that goal directed behaviors are similar to telo-taxis behaviors, in reality they are distinct. Telo-taxis does not involve identification and approach to shapes, objects, or places while goal directed behaviors do. Telo-taxis behaviors can successfully operate even with simple eyes while goal-directed behaviors cannot. In addition goal directed behaviours can keep track of the goal even when there are occasional obstructions of it during the approach process. In contrast telo-tactic behaviors lack the ability to track the goal when it is obstructed. By using one or more stopping conditions for a goal directed behavior we get a guarded goal directed behavior which is also funneling behavior.

I. Turning Behaviors

Turning can be performed in several different ways, with either reference to external sensory cues or to only interoceptive sensory information. Some of the turning methods assist in funneling while others do not. If one turns *until* he faces a particular stable in position entity in the environment — such as mountains, fountains, trees, buildings, or even the sun’s position or stellar cues provided that the agent has a mechanism for correcting for their apparent circadian and seasonal motion — then the turn’s stopping condition will regularly lead the agent to a unique orientation. This is a case of behavioral funneling because regardless the original orientation of the agent the resulting one will be regularly unique.

V. EXPERIMENTAL SUBSTANTIATION

Now, let us see how the above list of funneling behaviors can be used and applied in practice. A simple learning mechanism has been developed for illustrating how funneling behaviors can be used for building maps and planning actions through them, without need for probabilistic models in software. The task of the robot was to move around in a three-roomed environment (Fig. 5) and build an internal graph-like representation of its interaction with the environment. Then this representation was used for planning paths in this environment. Figure 4 depicts a multigraph representation which was constructed by the robot. The sensory perceptions of the robot are represented as discs; each disc corresponding to one view in the environment. Each group of four discs corresponds to a single place, where the robot rotates on the spot and captures four snapshots, 90 degrees apart. Each of these four snapshots is connected to its neighbors with the actions ‘turn left’ and ‘turn right’, which enable the robot to change its point of view.

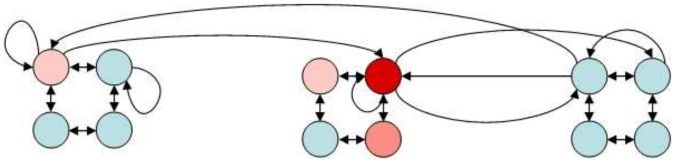

Fig. 4. Example of multigraph created while exploring a flat with three rooms. Each disc represents one view of the environment. The different colors (and intensity) represent the recency of the view’s encounter. Only the first thirty-two actions are depicted, due to the cluttered nature of drawing such data in 2D.

When the robot activates one of its actions interacts with the environment and moves from one place to another or from one direction to another. If the robot moves, its sensory perception (view) will commonly change. In Figure 4, the arrows indicate which funneling actions can be initiated from the current place and direction and what will be their result to the sensory perception (view) of the robot.

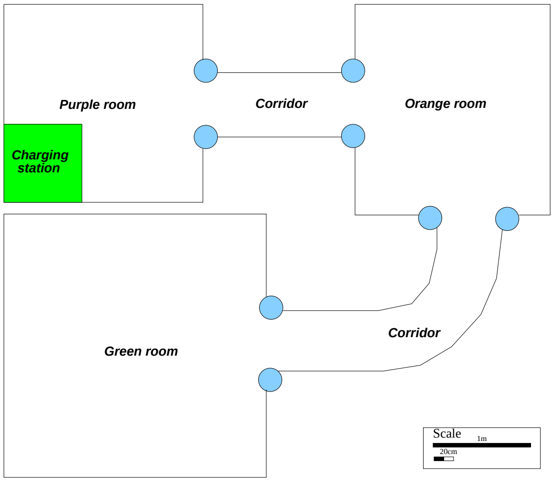

Fig. 5. Top view of the real world environment used for the experiments, in scale.

1) *Experimental Setup:* The experiments were conducted in an indoor experimental arena (Fig. 5), which consists of three rooms connected with corridors. The walls of each room are painted with a different color, the corridor entrances are made of two blue circular pillars placed on the sides and the charging station is a corner which is painted green in room of different color. Color and geometry are

the means for allowing distinguishing among different places due to the perceptual capabilities that the robot is equipped with. A MagellanPro mobile robot has been used during the experiments. A color camera and a LASER scanner mounted on the robot are used, including ultrasonic SONAR distance sensors and bumpers for navigation purposes.

2) Sensor Signals Processing: The sensory feature extraction from the LASER scanner data which is a sequence of 180 angular measurements consisted of a FFT, differentiation, and detection of the angular position of the two maximum and two minimum measurements. These sensory features are used by eight other processing units, four neural networks (ANN) and three rule-based processors for extracting more abstract features: in room?, in corridor?, doorway? and free space?, each of this features is binary. The classification accuracy of each of these classifiers is between 70% and 90%. The sensory feature extraction for the color camera consisted of six colors blob identification and a doorway detector which essentially was looking for two vertical regions of the same color on the sides of the frame with a gap of different colors in between.

3) Behavioral Modules: Motor actions are produced by behavioral modules built on the basis of the previously mentioned funneling behavioral patterns. For achieving reliability and high repeatability of outcome whenever invoked, the behavioral modules were implemented with behavioral redundancy so that for every target competence there were at least two motor pathways driving the robot, so that if one method was failing the other still drives the robot towards the right direction. The behavioral repertoire of the robot was composed by five guarded funneling behaviors (i.e. “go to charger”, “go through the deepest opening (corridor) until you reach a room”, “find corridor and transverse it until it ends”, “turn to doorway” and “turn to blue color”).

4) Exploration: Initially, the robot was left to explore the environment by performing 100 actions randomly selected from those of its behavioral repertoire. Subsequently, the map representation was built off-line, using the acquired data. For building the map for each new view (different than the known ones) a new node on the multigraph was added and the action which lead to that view-node was represented as a link in the graph from the previous view-node to the new. In the sequel the robot was rotating on the spot to take corresponding views on its left, right, and back represented in Fig. 4 as groups of four circles. If in subsequent visits of a node the action taken was not leading to the previously learned next node the behavior was considered not to operate as a funneling one and the link was labeled as unusable. In this manner after 100 actions the resulting map had only links which were repeatedly producing a constant effect. We should mention that in this process the quasi-determinism level was not employed as such since the simpler repeatability we employed had the desired outcome. The quasi-determinism level measure is useful for comparing different implementations of a behavior.

A. Experimental Results

Once trained, the performance was assessed by planning for five different, randomly selected tasks. For each task, the robot had to move from a starting to a goal position and take a picture with its camera. We performed 20 runs for each of the tasks. For planning a path, a search was made on the multigraph representation for finding a complete path from start to goal view.

The testing criterion was the *success rate* of the planner, which was defined as the percentage of successful runs over the number of trials for each task. An unsuccessful run occurs when the robot stops, not being able to make a plan from the current position to the goal, or if the time used for accomplishing the task exceeded *five* minutes. The success rate of each planner for five different tasks is depicted in Fig. 6. The measured overall success rate is 96%.

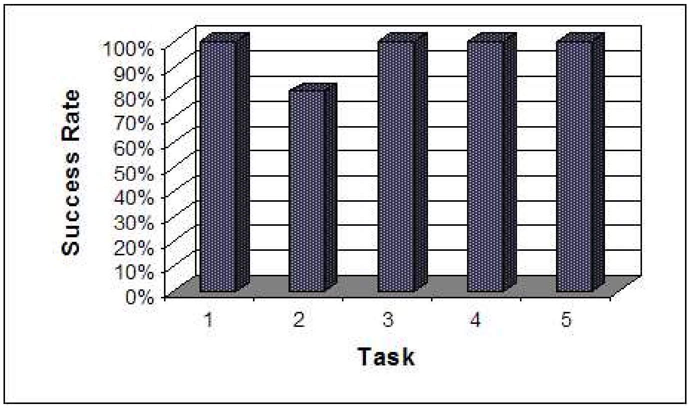

Fig. 6. Experimental results for the mapping and planning mechanism operating in the environment depicted in Fig. 5.

These results illustrate that even a really primitive planner can achieve high success rate in real world operation if the behaviors used are based on specific behavioral patterns which have been successful in nature for multitude of years. Further experiments have been performed in an unmodified furnished flat with three rooms, but the available space does not allow inclusion of these results which will be published in a separate article. In that experiment different sensory processing and funneling actions have been used; the achieved overall performance was 93%.

VI. CONCLUSIONS

The experimental results confirm that even a really primitive planner can achieve high success rate in real world operation if the behaviors used are based on specific behavioral patterns which have been successful in nature for multitude of years.

A list of such behaviors — called funneling behaviors — was given. This is not a full list of all possible funneling behavioral patterns since we do not have currently a method to prove that these are all the possible patterns. It is rather a list of the funneling behavioral patterns that are currently known.

REFERENCES

[1] Anthony R. Cassandra, Leslie Pack Kaelbling, and Michael L. Littman. Acting optimally in partially observable stochastic domains. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 1994.

[2] Gottfried S. Fraenkel and Donald L. Gunn. *The orientation of animals: kineses, taxes and compass reactions*. New York: Dover Publications, 1961.

[3] B. Kuipers and Y. T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Robotics and Autonomous Systems*, 8:46–63, 1991.

[4] L. Mihaylova, T. Lefebvre, H. Bruyninckx, and J. De Schutter. *Series III: Computer and Systems Sciences*, volume 198 of *NATO Science Series ‘Data Fusion for Situation Monitoring, Incident Detection, Alert and Response Management’*, chapter Active Robotic Sensing as Decision Making with Statistical Methods, pages 129–143. IOS Press, 2006.

[5] C. S. Patlak. A mathematical contribution to the study of orientation of organisms. *Bulletin of Mathematical Biophysics*, 15:431–476, 1953.

[6] C. S. Patlak. Random walk with persistence and external bias. *Bulletin of Mathematical Biophysics*, 15:311–338, 1953.

[7] J. R. Riley, U. Greggers, A. D. Smith, D. R. Reynolds, and R. Menzel. The flight paths of honeybees recruited by the waggle dance. *Nature*, 435:205–207, May 2005. I HAVE IT.

[8] Hermann Schöne. *Spatial Orientation: the spatial control of behavior in animals and man*. Princeton Series in Neurobiology and Behavior. Princeton University Press, 1984.