Explorative Path Planning

eingereichtes
HAUPTSEMINAR
von

Jasper Friedrichs

geb. am 12.04.1987
wohnhaft in:
Christoph-Probst-Str. 16/611
80805 München
Tel.: 0176 22879091

Lehrstuhl für
STEUERUNGS- und REGELUNGSTECHNIK
Technische Universität München

Univ.-Prof. Dr.-Ing./Univ. Tokio Martin Buss
Univ.-Prof. Dr.-Ing. Sandra Hirche

Betreuer: Dipl.-Inf. Nicolai Waniek, M.Sc. Christian Denk
Beginn: 18.04.2013
Abgabe: 18.06.2013

Abstract

An essential prerequisite for effective interaction with the environment is knowledge about the environment. For animals and autonomous robots alike this means that unknown environments need to be explored. An efficient exploration strategy is thus indispensable for both. In this advanced seminar report we do not develop any such strategy but give an overview about solutions to the explorative path planning problem in robotics as well as in biology. This allows us to contrast the two. The overview of exploration strategies should enable the reader to find a strategy suitable for his requirements.

Zusammenfassung

Effektive Interaktion mit der eigenen Umgebung setzt Wissen über diese Umgebung voraus. Sowohl für Tiere als auch Autonome Roboter heißt das, dass unbekannte Umgebungen erkundet werden müssen. Effiziente Erkundungsstrategien sind daher für Beide unabdingbar. In diesem Hauptseminar entwickeln wir keine solche Strategie, sondern geben eine Übersicht über Lösungen aus der Robotik und der Biologie und vergleichen Diese. Die Übersicht über Erkundungsstrategien soll es dem Leser ermöglichen eine für seine Zwecke geeignete Strategie auszuwählen.

Contents

Chapter 1

Introduction

In order to allow a robot to act autonomously in an unknown environment it needs to perceive and model its surroundings. The task of having a robot simultaneously localize itself in the perceived environment while extending its knowledge about the environment has been extensively researched as the problem of simultaneous localization and mapping (SLAM). A central aspect for optimizing mapping of an unknown environment is choosing where to go next in order to extend the knowledge of the robots surroundings. The selection of future vantage points and the resulting path is considered explorative path planning. In this paper we present a representative selection of various exploration strategies designed for robotic applications. Similar to autonomous robots, animals need to explore unknown environment as well, for example as a prerequisite to foraging or finding shelter. We will therefore also present exploration strategies from biological research. Then the final step is to put the biological exploration strategies and the approaches in the field of robotics in relation to each other.

Chapter 2

Explorative Path Planning in Robotics

The task of exploring an unknown environment with a robot can be considered as the task of guiding a robot through the environment in a way that allows the robot to perceive it. Numerous approaches to this task are presented in literature. We will present a significant selection that covers the different approaches. The exploration strategies are ordered by complexity, which basically means how much information they evaluate. Sim and Dudek have examined several exploration strategies where the format of the explorative path is fixed and independent from the observed environment [1]. Multiple other strategies actually evaluate the already obtained knowledge of the environment in order to consider where to go next. The Frontier-Based Approach [2] only differentiates between currently unexplored and explored areas, guiding the robot to the closest unexplored area. The exploration strategy by [3] actually estimate the information gain obtained when traveling to another location while further approaches evaluate multiple criteria besides the information gain [4]. In the following we will reproduce the referenced exploration strategies.

2.1 Fixed Path Policies

Sim and Dudek examine multiple exploration strategies where the path is based on a policy and independent from the already observed environment [1]. They evaluate results from simulations as well as robotic experiments with regard to accuracy of localization and coverage of unknown space.

The policies are SeedSpreader, Concentric, FigureEigth, Random, Triangel and Star. Example trajectories are depicted in Figure 2.1. The SeedSpreader policy covers an area through parallel zig-zag paths. The concentric traces concetric circles from the starting point while alternating direction with every circle. The FigureEigth policy traces growing figure eight paths, thus moving through the starting point repeatedly. The Random strategy guides the robot in a random direction at each step.

The Triangle strategy traces growing concentric equilateral triangles. The Star policy guides the robot along a set of growing rays emerging from the start point.

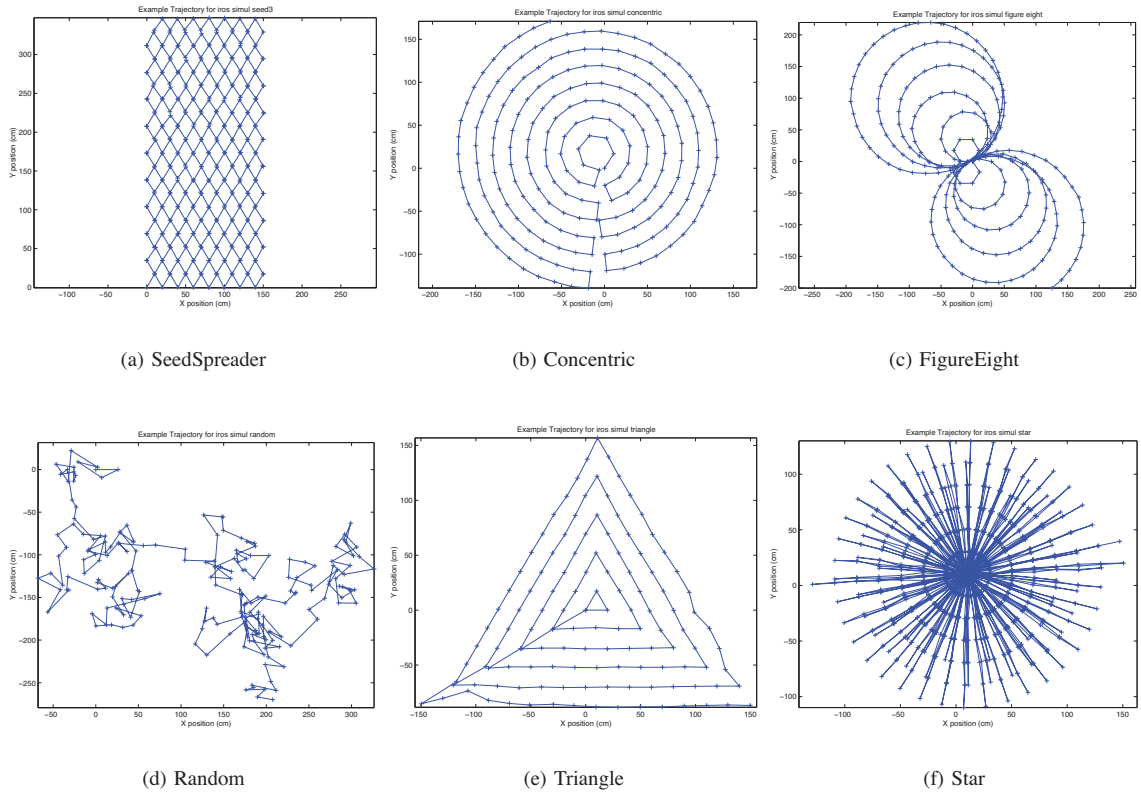

(a) SeedSpreader (b) Concentric (c) FigureEight

(d) Random (e) Triangle (f) Star

Figure 2.1: Example trajectories for fixed path policies [1]

All policies are simulated in an empty rectilinear environment and their position error and exploration efficiency is evaluated. The important results are, that the Star exploration strategy was very accurate due to revisiting known territory, but highly inefficient with respect to obtained data per action for the same reason. The Concentric and FigureEigth strategies proved quite inaccurate because the positioning error accumulated over time, but were more efficient in gathering new information than the Star strategy. As a result Sim and Dudek state that it is difficult to balance accuracy with efficiency as these goals tend to oppose each other.

2.2 Frontier-Based Exploration

Yamauchi describes the frontier-based approach as moving to the boundary of open space and unexplored territory [2]. These boundaries are extracted from the currently available representation of the environment. The environment is mapped to

an evidence grid [5], which is a Cartesian grid with cells that store the probability of being occupied. All cells are initialized with an average probability of being occupied. These values are updated based on the sensor data obtained during exploration. The evidence grid allows a classification of the cells into open, occupied or unexplored. A cell is considered open if the probability of being occupied is lower than the initial probability whereas it is considered occupied for a larger value. An unknown cell has the initial probability of being occupied. Consequently a frontier edge cell can be determined as an open cell adjacent to an unknown cell. Edge cells are grouped together as frontiers and their centroid is calculated. Figure 2.2 shows this process. After updating the evidence grid and identifying the frontiers the robot moves to the centroid of the closest frontier. Yamauchi implemented frontier-based exploration on a Nomad 200 mobile robot with laser rangefinder, sonar and infrared sensors and conducted experiments in a real-world office environment as a proof of concept. The following section about information-based exploration compares the performance of both approaches.

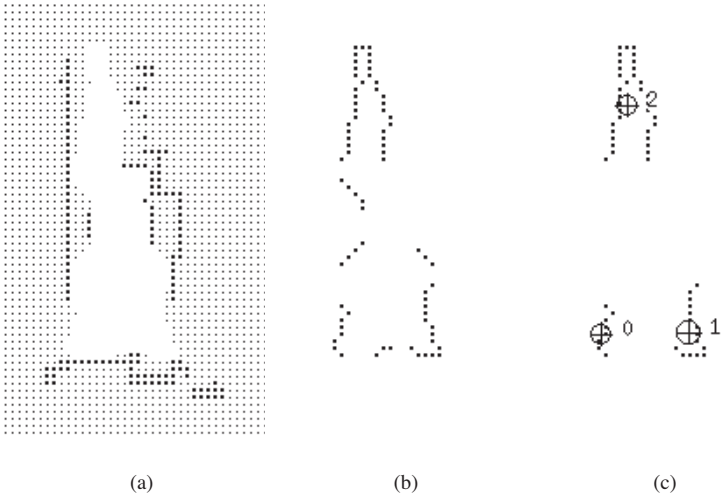

Figure 2.2: Detecting frontier centroids [2]

2.3 Information-Based Exploration

Besides guiding the robot to the closest unexplored position another rather intuitive exploration strategy is to travel to the location that supplies the maximal amount of new information. Stachniss and Burgard introduce one such exploration strategy that aims to maximize the obtained information with each exploration step [3]. Additionally the approach is compared to a frontier-based approach as well as to a combination of the two.

2.3.1 Coverage Map and entropy

For this approach the observed environment is mapped in a coverage map, which poses an extension to the evidence grid mentioned in the previous section. As opposed to regarding an observed grid cell binaurally as either occupied or open as in evidence grids, coverage maps assign a value between one and zero to the cell which expresses the occupied percentile of the cell. As the robot will usually not be able to perceive the ground truth coverage value of a cell, it maintains a probabilistic belief in form of a histogram for each cell. From the histogram we can calculate the entropy as a measurement of uncertainty for each cell. The entropy H of a histogram h consisting of n bins h_i is defined as:

$$H(h) = -\sum_{i=1}^{n} p(h_i) * \log p(h_i). \tag{2.1}$$

A low entropy of a cell corresponds to a high certainty of the cells coverage value. As a result the sum of all cells' entropies can be regarded as the certainty the system has of the map.

2.3.2 Exploration strategies

As stated previously the entropy can be regarded as a measurement of uncertainty. Consequently the information gain of a measurement for a single cell can be calculated as the difference between the entropies before and after the measurments. Resulting in

$$I(h(c_l)|d) = H(h'_d(c_l) - H(h(c_l)) \tag{2.2}$$

for a measurment d into a cell c_l and an expected histogram of $h'_d(c_l)$ after measuring. The overall information gain is the sum over the information gains of all cells perceived by the measurement. However as we do not know the actual measurement we have to integrate over all possible measurement, which Stachniss and Burgard formulate as:

$$E[I(l)] = \sum_{d} p(d|c) * \sum_{c \in C(l,d)} I(h(c_i)|d), \tag{2.3}$$

where $C(l, d)$ are all cells coverd by the measurement and the likelihood of an observation $p(d|c)$ is calculated similar to [5]. Consequently the next viewpoint is obtained by maximizing formula 2.3.

Opposed to that, the next viewpoint for the frontier-based approach would be selected as the closest viewpoint that allows the observation of a high entropy cell. The frontier-based approach is also referred to as the closest location (CL) strategy.

Another strategy would be to combine the latter two, resulting in a viewpoint selection depending on distance as well as information gain:

$$l_{next} = \arg\max_{l \in L(c)} \left[\alpha * \frac{E[I(l)]}{\max_{l' \in L(c) E[I(l')]}} - \frac{d_c(l)}{\max_{l' \in d_c(l')}} \right], \tag{2.4}$$

where $d_c(l)$ is the distance to l and α is a parameter to adapt the influence of the information gain on viewpoint selection.

2.3.3 Results

Stachniss and Burgard simulate the frontier-based approach as well as variations of strategies evaluating information gain. The path length and the number of measurements are evaluated. The number of measurements is of interest if there is a cost to perceiving the environment.

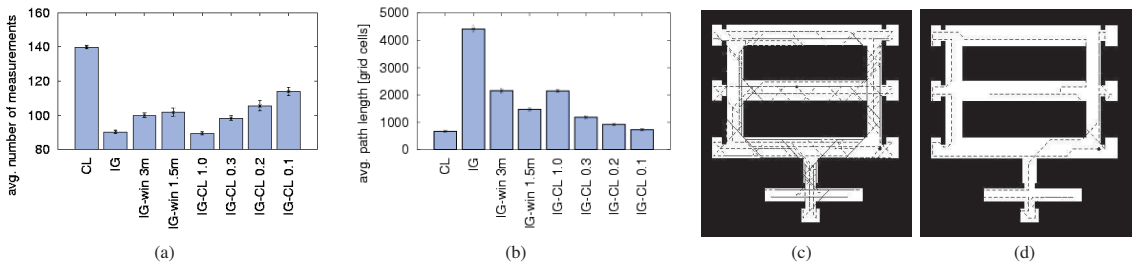

Figure 2.3: Average number of measurements (a), path length (b), and sample paths of "IG" (c) and "CL" (d) strategies [3]

Figure 2.3 shows the experimental results as well as a sample path for the frontier-based and solemnly information based approach. The number behind "IG-win" is the radius of a limited search window that needs to be explored first during the information gain strategy whereas the number behind "IG-CL" is the parameter α (2.4). The comparison shows that the frontier-based approach is optimal if the cost of measurements is neglectable, as the approach has the shortest path as well as the highest number of measurements. Considering both number of measurements and path length "IG-CL 0.3" performs sligthly better that the purely information gain orientated strategies, as these strategies tend to long paths. This result leads to the next section, which describes an approach, where multiple objectives are evaluated to select observation points.

2.4 Multiple Objective Exploration

A common problem in explorative path planning is balancing efficiency and accuracy. The accuracy depends substantially on the localization of the system, while the efficiency depends on path length and information gain. In [4] Makarenko et al introduce an exploration strategy that combines the evaluation of localizability, information gain and navigation cost. Based on the assumption that the robot stationarily performs a 360 ° sensor sweep for maximum information gain, the exploration strategies filter for discrete observation points.

Candidates for the observation points are selected as in the previously discussed frontier-based exploration approach. The candidates are then evaluated by a multi-objective cost function.

One evaluation criterion is the information gain. Similar to the previous information maximizing approach the entropy of cells in the measurement region is evaluated. The information gain utility function is calculated as in [6],

$$U_i^I = -\sum_{C \in W_i} H(s(C)) \tag{2.5}$$

where the entropy H is calculated based on the binary distribution $s(C)$ of a cell C in the region W_i as in [7].

Another evaluation criterion is the navigation cost. The utility of navigation

$$U_i^N = -V(x_i) \tag{2.6}$$

is the cost of reaching a point x_i from the current location.

The utility of localizability is important because any position of a feature mapped at an observation point will accumulate the error in the robots position. This metric thus aims to encapsulate the uncertainty of the robots position at a candidate observation point. Makarenko et al define the localizability metric as the minimum vehicle covariance achievable by relocalizing the robot through observations of previously mapped features. The localization uncertainty after k such observations is described by the covariance matrix $\mathbf{P}_{vv}^k$. For more details the interested reader is referred to [4] as well as to the Shannon information measure [7] used to transform the vehile covariance matrix to a scalar. The entropy H for a gaussion distribution on vehicle pose can thus be used as the localization utility,

$$U_i^L = -H(\mathbf{P}_{vv}^k) = -\frac{1}{2}\log((2\pi \exp)^n |P_{vv}^k|). \tag{2.7}$$

It is important to note, that the calculation of this metric is computationally expensive and thus only calculated for a limited number of possible observation points.

Figure 2.4 shows the effect of the different utilities, where darker regions are preffered. As a final step, the total utility of a candidate observation point is calculated as the weighted sum of all utilities. The exploration strategy selects the candidate observation point with the highest total utility.

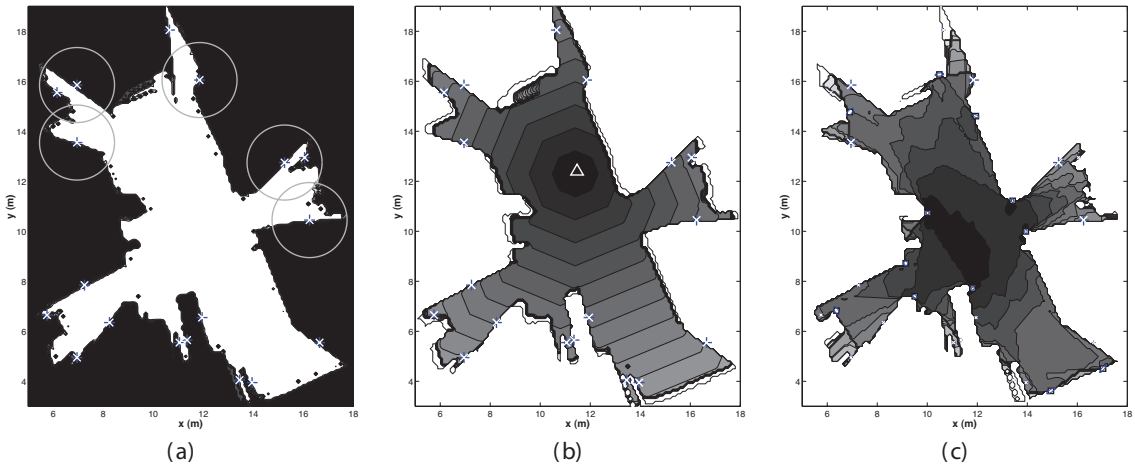

Figure 2.4: Utility of information gain (a), navigation (b) and localization (c) [4]

Chapter 3

Explorative Path Planning in Biology

Biology offers a great model for imitation, copying and learning, and also inspiration for new technologies [8]. Such technologies could be computational systems, algorithms or the explorative path planing problem. However, we have to consider that the complexity of biological cognitive systems draws boundaries that require further research. Moreover considering the explorative path planning problem, difficulties might arise from the differences between spatial understanding and representation by navigation organisms and robotic systems [9, p. 2]. Nevertheless we aim to find inspiration for exploration strategies by regarding research in the field of biology. Animals such as the desert ant genus Cataglyphis rely on efficient searching for surviving in a hostile desert environment. Wehner and Srinivasan not only examine this species searching behavior but also develop an analytical model inspired by it [10]. In this chapter we will reproduce their insights as well as research about the search behavior of rats [11]. This is a prerequisite to the next chapter where we will compare biological to robotic explorative path planning.

3.1 Explorative Behavior of Desert Ants

In [10] Wehner and Sinivasan examine desert ants returning to their nest. This behavior, called homing, triggers a search behavior if the nest is not found. The developed analytical model for the search strategy and the discussion of efficiency are of special interest to us.

The search path can be described as a system of ever increasing loops precisely center around an origin. Figure 3.1 exemplary shows an ant's search path for a period of one hour, where the first 21 minutes are show separately to show the increase in covered area. What makes this searching strategy efficient is that the ant spends more time searching closer to where it suspects its goal than farther away.

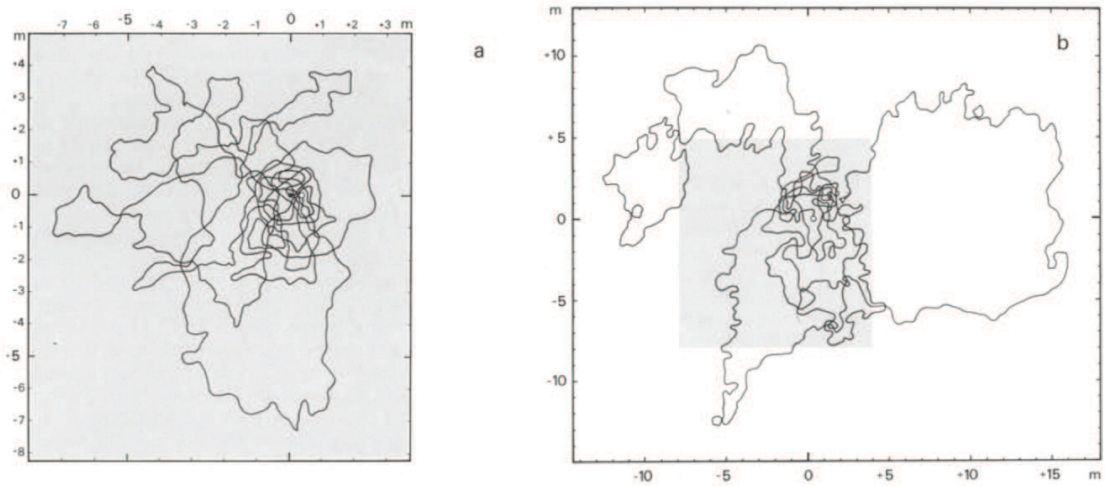

Figure 3.1: A desert ant's search path over a period of 21 minutes (a) as well as 1 hour (b) [10]

In their research Wehner and Sinivasan try to develop an optimal search strategy assuming the likelihood of finding the target is a Gaussian profile around the point of origin, where the target is suspected. They consider a strategy as optimal if it explores the region were the probability of finding the target is highest. A spiraling search strategy is dismissed as it might never find the target if missed once due to imperfect navigation or observation systems. A search pattern with oscillating radial and random tangential movement is developed and refined to match the search pattern of the desert ant. The interested reader is referred to the original derivation in [10]. The analytical model expresses the ants search path as a relation between tangential and radial movement,

$$Q(p) = \frac{T(p)}{R(p)} = \frac{d}{n} * k * V(p) \quad \text{with} \quad T(p)^2 + R(p)^2 = 1, \tag{3.1}$$

where p is the current step, n is the number of steps, d is the distance to the origin, k is a scaling factor and V is a random signal explained in the Appendix of [10]. Figure 3.2 shows an example path of 500 steps calculated by 3.1.

Based on their derivation Wehner and Sinivasan claim that the search pattern is suited to minimizes the time to find a target with visual and navigation systems that are not 100% accurate. Interestringly they also mention several sources that observed similar search patterns in other insect species.

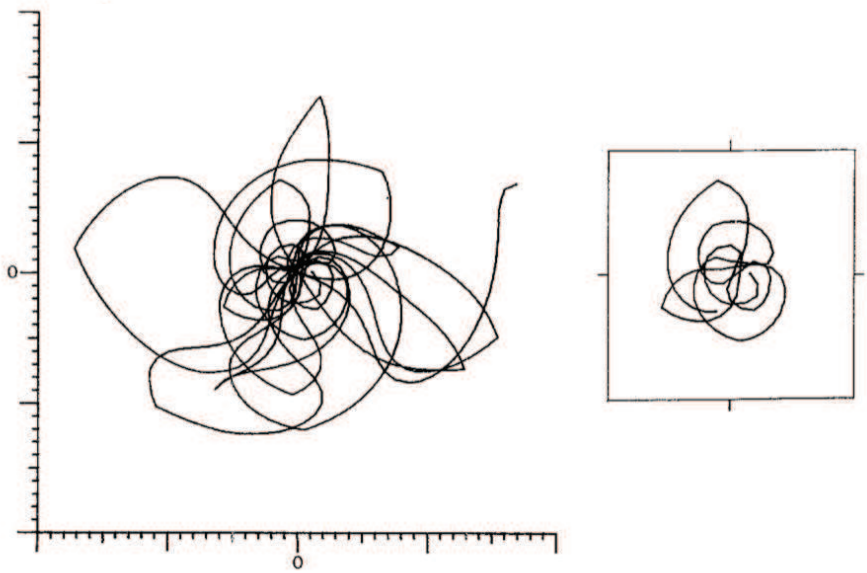

Figure 3.2: Exemplary path derived by analytical model of the desert ant search path [10]

3.2 Explorative Behavior of Rats

Previous research showed that rats show a preference to a certain location, called a home base, and perform explorative excursions from there [12]. In [11] Tchernichovski et al further examine the structure of rat's explorative behavior in unknown environments. They describe an alternating pattern between stopping and excursions and show an increase in excursion length similar to the previously described behavior of desert ants. It is suggested that the increase of excursion length is based on an evaluation criterion of familiarity.

Observations were made with 43-50 day old Long Evans hooded rats in a $6.5\,m$ diameter circular arena. The starting point for exploration was a tube similar to the rats' previous shelter, establishing it as a home base. The rats were exposed to the area for 30 minute sessions. The main findings of [11] include an increase of excursion lengths during sessions as well as from session to session. This behavior is modeled with a change from home base attraction to home base repulsion profile, based on the measured rats' velocities. This basically means, that from an unvisited location the rat is attracted back to the home base while after some visits to that location the rat is repulsed from the home base at that location. Based on statistical analysis of the change of attraction and repulsion Tchernichovski et al suggest that only the familiar portion of an excursion increases while the unfamiliar portion stays constant.

Chapter 4

Comparison of Biological and Robotic Exploration Strategies

After we have presented different approaches to explorative path planning from robotic as well as biological research we will now put the approaches in relation to each other. After a short comparison of the biological strategies we will contrast them to the robotic strategies. Finally we will classify the robotic solutions to explorative path planning.

The search strategy of desert ants and the home based exploration of rats share the expanding path length behavior. Both strategies are thus suited to explore an area where the regions closer to the point of origin are prioritized. A substantial difference between the research about desert ants and rats lies in how the authors try to explain the behavior. Wehner and Sinivasan come to the conclusion that their analytical model describing a randomized explorative path planning policy should be feasible for the ant's nervous system [10]. Tchernichovski et al suggest that the rat's explorative path planning might actually rely on evaluating the familiarity of a location, though they cannot prove that thesis [11]. It is important to note, that if we aim to be inspired by biological solutions to explorative path planning instead of trying to accurately mimic them, behavioral explanations are of interest even if their validity isn't proven. After all even just a new perspective can produce valuable insights. Therefore the next step is to compare the biological to the robotic solutions.

The analytical model of desert ants' search seems comparable to the fixed path policies in explorative path planning for robotics. Sim and Dudek mention in their evaluation of fixed path explorative policies, that balancing accuracy and efficiency is a challenge [1]. It seems that the desert ant strategy could supply a competitive strategy here, as it efficiently explores an area if the objective is to explore regions close to the origin more thoroughly. The strategy also revisits previously visited locations, which can be seen as a compromise to accuracy as it should allow relocalization. It also offers a chance to compensate for inaccurate sensors or suboptimal

viewfields by multiple measurments.

These advantages apply to the exploratory path of rats as well, as they too revisit locations closer to their home base. Converted to a robotic exploration strategy the home base could be specifically used for relocalization or even selected as a position where localization is facilitated by multiple observable features. The suggestion that rats evaluate the familiarity of a location seems very similar to the information-based exploration strategies in robotics. Unfortunately Tchernichovski et al only show that the evaluation of familiarity can explain the explorative behavior but do not suggest how this is done. Further research in this area might supply valuable insights for robotic exploration strategies.

Chapter 5

Conclusion

The previously presented solutions to explorative path planning show differing characteristics and are thus suited under different circumstances. The fixed robotic exploration policies pose interesting solutions for explorative path planning in systems with low computational power that need to calculate their explorative path without evaluating the environment. The numerical model of the desert ants search strategy can be viewed as a fixed policy as well and might even outperform the presented fixed robotic policies. The advantage of the fixed policies lies in their simplicity. The frontier-based strategy also performs very well despite its simplicity. It is suitable if a system is required to minimize the traveled path while the number of measurements is neglectable. However, if the number of measurements needs to be minimized an information-based approach should be better suited. This comes with the drawback of a large traveled path, but a combined approach, that tries to find a compromise between reducing path length as well as number of measurements was suggest as well. Along that line multi-objective approaches allow a lot of flexibility when trying to find a compromise between several evaluation criteria. They can be adapted via weights and extended with more criterions. The characteristics of the different strategies for explorative path planning show that an ultimate solution does not exist, because it depends on the system as well as the requirements for the exploration task.

Bibliography

[1] R. Sim and G. Dudek. Effective exploration strategies for the construction of visual maps. In *Intelligent Robots and Systems, 2003. (IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, volume 4, pages 3224–3231 vol.3, 2003.

[2] B. Yamauchi. A frontier-based approach for autonomous exploration. In *Computational Intelligence in Robotics and Automation, 1997. CIRA'97., Proceedings., 1997 IEEE International Symposium on*, pages 146–151, 1997.

[3] C. Stachniss and W. Burgard. Exploring unknown environments with mobile robots using coverage maps. In *Proc. of the International Conference on Artificial Intelligence (IJCAI)*, 2003.

[4] A.A. Makarenko, S.B. Williams, F. Bourgault, and H.F. Durrant-Whyte. An experiment in integrated exploration. In *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, volume 1, pages 534–539 vol.1, 2002.

[5] H.P. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Robotics and Automation. Proceedings. 1985 IEEE International Conference on*, volume 2, pages 116–121, 1985.

[6] Alberto Elfes. Robot navigation: Integrating perception, environmental constraints and task execution within a probabilistic framework. In *Reasoning With Uncertainty in Robotics*, pages 93 – 130, Berlin, Germany, 1995. Springer-Verlag.

[7] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley-Interscience, New York, NY, USA, 1991.

[8] Janine M. Benyus. *Biomimicry: Innovation Inspired by Nature*. William Morrow Paperbacks, 1998.

[9] F.L. Dolins and R.W. Mitchell. *Spatial Cognition, Spatial Perception: Mapping the Self and Space*. Cambridge University Press, 2010.

[10] Rüdiger Wehner and Mandyam V. Srinivasan. Searching behaviour of desert ants, genuscataglyphis (formicidae, hymenoptera). *Journal of comparative physiology*, 142(3):315–338, 1981.

[11] Ofer Tchernichovski, Yoav Benjamini, and Ilan Golani. The dynamics of long-term exploration in the rat. *Biological Cybernetics*, 78(6):423–432, 1998.

[12] D. Eilam and I. Golani. Home base behavior of rats (rattus norvegicus) exploring a novel environment. *Behavioural Brain Research*, 34(3):199 – 211, 1989.