

Topological Navigation with Hierarchically Structured Spiking Neural Networks

Master Thesis in Physics

Supervisor: Prof. Dr. Mallot, Cognitive Neuroscience
University of Tübingen

Second Assessor: Prof. Dr. Kley, Computational Physics
University of Tübingen

Matthias M. Brucklacher

September 21, 2020

Erklärung

Hiermit versichere ich, dass ich diese schriftliche Abschlussarbeit selbstständig verfasst habe, keine anderen als die angegebenen Hilfsmittel und Quellen benutzt habe, sowie alle wörtlich oder sinngemäß aus anderen Werken übernommenen Aussagen als solche gekennzeichnet habe. Ich versichere, dass die Arbeit weder vollständig noch in wesentlichen Teilen Gegenstand eines anderen Prüfungsverfahrens gewesen ist, dass ich die Arbeit weder vollständig noch in wesentlichen Teilen bereits veröffentlicht habe, sowie dass die in Dateiform eingereichten Exemplare mit den eingereichten, gebundenen Exemplaren übereinstimmen.

Matthias Brucklacher, Tübingen am 21. September 2020

Abstract

The hippocampus is well understood as representational apparatus of the environment. Since the world is its own best model, a purely representational function would be wasteful, and it is obvious that it must have some behavioral function as well. Here, a model is proposed in which a neuronal structure inspired by the hippocampal network represents the topological syntax of the world. As in the classical understanding of place cells, the agent's position is encoded in a neuronal activity pattern. In a second step, the network is used to plan goal-directed trajectories. By releasing a wave of activity from an ensemble of cells representing the goal, the network's synaptic weights are temporarily altered via spike-timing-dependent plasticity. In this process, goal-directed weights are strengthened to create a synaptic vector field. Simulation experiments demonstrate the model's navigational ability, with motor execution based on a voting process of active neurons. In doing so, this study shows the potential of synaptic vector fields for biologically plausible planning in state-action representations. In an extension of the model, a second neuronal layer is added to capture a hierarchical structure in the environment. Simulation experiments show that such a hierarchical structure can successfully guide the agent to subgoals in tasks of goal-directed navigation. Subgoal selection by the model resembles the strategy of human subjects in comparable navigational tasks. Since hierarchical planning is suspected to increase planning depth in human cognition, the influence of the additional layer on planning across large distances is investigated. Comparison with a network without hierarchical structure indicates a benefit of the second layer on planning depth.

Acknowledgments

I would first like to thank my supervisor, Professor Hanspeter Mallot, whose continuous support and expertise made this thesis possible. I would also like to thank my colleague Tristan Baumann of the Cognitive Neuroscience Lab at the Department of Biology, University of Tübingen for his invaluable, constructive feedback and critical discussions. Lastly, I want to thank Nora Götz, for all your love and support as well as my wonderful family for their never-ending encouragement.

Contents

1	Introduction	11
1.1	Topological Representations of Space	11
1.2	Hierarchically Structured Behavior	14
1.3	A Neuronal Model for Topological Navigation	18
2	Methods	21
2.1	Neuronal Dynamics	21
2.2	Synaptic Plasticity	23
2.3	Motor Efferences	27
2.4	Neuronal Hierarchy	28
3	Simulation Experiments	31
3.1	Comparison to Human Navigation	31
3.1.1	Entering the Goal Region (HN1)	31
3.1.2	Navigation Across Multiple Region Boundaries (HN2)	41
3.2	Planning Depth	47
4	Discussion	57
4.1	Performance in Wayfinding-Tasks	58
4.2	Increased Planning Depth	61
4.3	General Discussion	62
4.4	Suggestions for Further Research	65
5	Conclusion	69
A	Supplementary Information	71
A.1	Inhibitory System	71
A.2	Network Parameters	72
A.3	Low-level Neurons as Place Cells	74
A.4	Comparison to Chance Level in HN1	74
A.5	Code Library References	75
A.5.1	LEMON	75
A.5.2	OpenCV	75

Chapter 1

Introduction

This chapter motivates the model-based approach to cognitive modeling, in the sense that the agent learns an S-A-S'-model consisting of encountered state S, conducted action A and predicted outcome state S'. Following existing work on the topic, levels of behavior in space are presented. It is argued that in order to achieve a certain level of behavioral complexity, model-based (state-action) representations need to evolve that form the basis of the computational model presented in chapter 2. Neurobiological background is given to stress the importance of population-coding in biological representations which in turn demands the biologically plausible model developed in this thesis to be population-coded. Section 1.2 then stresses the power of task-decomposition and hierarchical representations for cognition. To conclude the introduction, 1.3 presents the central part of this thesis: A neural network combining topological state-action representation with hierarchically structured behavior.

1.1 Topological Representations of Space

A Computational Perspective on Behavior in Space

The most fundamental task of spatial behavior is moving away from a stimulus such as light or a chemical gradient - or towards it. Required for fleeing and feeding, this skill requires only a minimum of information processing as realized in the famous Braitenberg vehicles (Braitenberg 1984). Following the line of thought from Mallot and Basten 2009, behavior in space can be arranged in order, ascending in behavioral complexity and necessary computation. The following descriptive list is to be understood in terms of David Marr's first and second level of inquiry (Marr [1982] 2010) that he developed to analyze information processing systems. Firstly, what is the task the system is solving? And secondly, which representations and algorithms can be used to do so? The third level, how such a system can be implemented in hardware, is discussed later.

1. Recognition-triggered response in selective cases, e.g. feeding or fleeing. Computationally, this requires a mapping of few stimuli to respective responses.
2. Path-integration and landmark guidance to allow homing after excursion. Path-integration requires continuously updating a "home-vector" (Cartwright and Collett 1987, Müller and Wehner 1988).
3. Following previously traveled routes to places of interest (Mallot and Basten 2009). This requires an ability to learn new stimulus-response associations, distinguishing this skill from purely inborn stimulus-response mappings.

4. Combining segments of previously traveled routes in novel ways to form a map of the environment (Tolman 1948, Mallot and Basten 2009). Computationally, this recombination requires adaptive stimulus-response mappings that depend on the current motivational drive. In a re-encountered state, the agent has to choose from multiple actions associated with it, depending on its current motivational drive. As demonstrated by the so-called *successor representation* in Stachenfeld et al. 2017, this can be achieved without a full S-A-S' representation. In the successor representation, states are superpositions of future locations in reach. Without knowing which new state S' will exactly result from an action A, value-function based goal-directed behavior and usage of previously learned representations for novel tasks is still possible (ibid.).
5. Route selection and planning. Humans are able to choose between alternative routes. To choose between alternatives in planning and thus enable truly flexible behavior, self-projection as well as inference of future states by having a model of the world is required (Behrens et al. 2018). This ability to plan is indeed based on a model of the world which is constructed from past experiences according to Buckner and Carroll 2007. In spatial cognition, two schools of thought have addressed the construction of such an internal representation: the map-based and the graph-based approach, where the labeled graph is somewhat in between (for a review see Warren 2019). The model described in this thesis is built on the graph-based (topological) approach.

Mallot and Basten continue this order by adding metric knowledge and lastly communication about space as the most complex tasks (Mallot and Basten 2009). Since the approach presented in this thesis does not aim to solve problems based on large-scale metric or communication, but to model a system for flexible, goal-directed navigation, we will not go into further detail on these more complex skills.

Models for Topological Navigation

Many computational models of spatial navigation have been constructed based on topological representations of the environment (Arbib and Liebleich 1977, Schölkopf and Mallot 1995, Arbib and Bonaiuto 2012, Baumann 2019). Already Tolman's narrow strip maps are qualitative predecessors thereof (Tolman 1948). Later, Kevin Lynch used the concept of a graph to describe spatial navigation in his book "The Image of the City". Lynch describes our mental representation of the city as consisting of five fundamental elements: paths and nodes, as well as edges, landmarks and districts (Lynch 1960, pp. 46–47). One of the first quantitative approaches to model the nature of an animal's (in this case a rat's) representation of space in a graph was conducted by Arbib and Liebleich 1977 who coined the term "World Graph". This graph consists of nodes corresponding to places and edges as the motoric movements to travel between them. Kuipers developed a comprehensive computational theory of spatial cognition, summarized in Kuipers 2000. His modularized approach is illustrated in figure 1.1 showing the information flow needed to construct a topological representation and beyond that a metric one from low-level, sensory input. The causal level (as described above as S-A-S'-representations) leads to a topological representation.

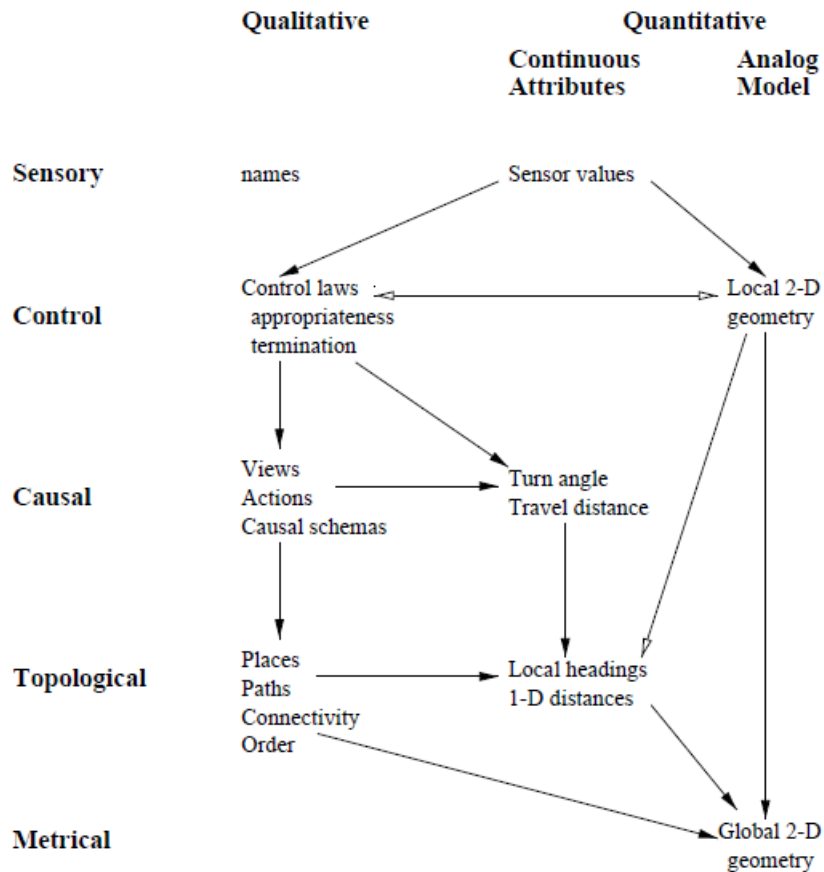


Figure 1.1: Representations of the spatial semantic hierarchy (Kuipers 2000). Open-headed arrows: information flow without dependency, close-headed arrows: dependency. Figure adapted from Kuipers 2000.

Recently, Baumann 2019 has shown that navigation in three-dimensional environments can work with a state-action model in which states are stored as "micro-snapshots". Baumann implemented the model as a graph, with nodes (states) corresponding to SURF image features (Bay and Tuytelaars 2006). These nodes are connected with edges labeled by movement instructions. By allowing multiple such features to be simultaneously perceived by the model, population coding of positions (see below) becomes possible. By computing a bundle of shortest paths based on a graph search algorithm, the agent can navigate in a goal-oriented way. Movement execution is based on a voting process of action labels attached to recognized features preactivated in goal direction by the graph search (Baumann 2019).

Encoding Location in Neuronal Activity

The third and last level of Marr's approach introduced in the beginning of the chapter is concerned with possible hardware implementations of the information flow. For the neurobiological understanding of the implementation of spatial cognition in the brain, the discovery of so-called place cells (O'Keefe and Nadel 1978) was a huge step. Place cells possess firing fields centered at discrete locations in space (ibid.). Over the past decades, research on place cells has led to a theory of the hippocampus

as a representational structure of the environment in which the agent's position in space is encoded. This encoding typically is assumed to be population-coded in the activity of neurons (Wilson and McNaughton 1993). Pfeiffer and Foster 2013 showed that place cells in rats display "preplay": At decision points, rats stop and appear to mentally probe multiple alternatives, sequentially activating place cells corresponding to places along possible trajectories (ibid.). How can such a population-coded representation, be it topologically arranged such as the state-action graph by Bauermann 2019 or a metric map, be used to generate behavior? The model put forward in this thesis addresses this question while aiming for biological plausibility.

1.2 Hierarchically Structured Behavior

Solving Complex Tasks

Studies on human problem solving such as conducted by Huys et al. 2015 have shown a strong tendency of subjects to apply strategic decomposition of complex problem spaces. In addition to approximations and simplifications, humans are able to split problems into subproblems that are easier to solve (ibid.). In the study by Huys et al., subjects were faced with a moderately deep planning problem across six interconnected states in which state-transitions were either rewarded or punished. After learning the rewards associated with the transitions, subjects were asked to plan a path to maximize reward. Statistical analysis of the action selection indicated that while doing so, subjects combined sequences of actions that previously had proven successful to so-called "chunks" and associate those with a joint reward value. This simplification, while leading to suboptimal performance in some cases, saved computational cost associated with evaluating a full decision tree (Huys et al. 2015).

Hierarchical Structure in Space

Such hierarchical representations appear to be important in spatial problem solving as well. Already Kevin Lynch described "districts" in the city, "conceived [...] as having two-dimensional extent, which the observer mentally enters 'inside of,' and which are recognizable as having some common, identifying character" (Lynch 1960, p. 46). Stevens and Coupe 1978 observed a peculiar effect of human large-scale navigation related to representations of such regions. Stevens and Coupe asked subjects to judge the relative position of Reno (Nevada) and San Diego (California). Surprisingly, the majority of subjects judged Reno to be north-east of San Diego, whereas in fact Reno lies to the north-west. Stevens and Coupe concluded that subjects solved the task on a hierarchically super-ordinate level, relying on their knowledge that Nevada lies to the east of California. Many other studies confirm such an influence of regions on spatial cognition, see for example Hirtle and Jonides 1985, McNamara et al. 1989 and recently Schick et al. 2019.

Wiener and Mallot 2003 further investigated the effect of such environmental regions on human navigation in two experiments. Both experiments began with a learning phase, in which subjects explored a virtual environment of discrete places marked

by characteristic landmarks (cf. figures 3.1 and 3.10). In section 3.1, these experiments (referred to as HN1 and HN2, short for human navigation) are used to test the model proposed in this thesis. The virtual environments used in the study by Wiener and Mallot consisted of clearly distinguishable regions. Regions were made obvious by choice of landmarks from the same semantic categories such as "car" or "animal". When asked to navigate to a specific landmark in the testing phase, subjects appeared to follow two interesting heuristics:

1. Experiment 1: The environment consisted of two interconnected islands separated by a river (see figure 3.1.) When choosing between equally long routes to a goal on the other side of the river, subjects preferred routes that led them to the goal region as quickly as possible (Wiener and Mallot 2003).
2. Experiment 2: The environment consisted of places around a hexagonal street, with landmarks from three semantic categories (see figure 3.10). When the subjects could choose between two equally long routes crossing multiple region boundaries, they preferably chose to cross as few region boundaries as possible to reach the goal (Wiener and Mallot 2003).

Based on the presented amount of evidence, it can be concluded that human problem solving deeply relies on such hierarchy-based heuristics and recognition of regions. The findings of Wiener and Mallot 2003 are the motivation for the simulations conducted in section 3.1. By simulating the same navigational tasks with the computational model proposed in this thesis, it is tested if the model architecture is able to replicate the effects of hierarchization on behavior.

Two Types of Hierarchical Structure

In a close, functional investigation of hierarchically structured behavior by Botvinick 2008, the author identified two overlapping structures. In correlational hierarchies, subtasks that frequently occur together, such as picking up a key and unlocking the door when opening a safe, will over time be associated with each other. In addition to action-action associations, such correlational hierarchies can also include associations to the context in which the actions are executed (ibid.). The second type of hierarchical structure determined by Botvinick are instrumental hierarchies. These are built on means-end relations, "pick up the key in order to unlock the door". The end of one action sets up the initial condition for the next. The example cited here already shows the overlap of the two concepts. Nevertheless, it is useful to distinguish between the two. In spatial cognition, correlational hierarchies occur if places share a common functionality, are visually similar or are designated with a regional name that might be purely historical (Schick et al. 2019). Instrumental hierarchies instead are built upon topological representations that were introduced in the previous section. "At intersection X turn right and go 500 meters straight to get to the supermarket", might be an action sequence part of the larger (and hierarchically superior) task of "going shopping". Sutton et al. 1999 use similar examples to introduce the concept of hierarchical reinforcement learning which will be discussed below. Overlap between the two hierarchical concepts can be observed in studies on taxi drivers who prefer to use a system of frequently used routes, a "skeleton in the cognitive map" (Kuipers et al. 2003). Even though these routes are preferred due to frequent usage, the drivers obviously must still know where each road

leads and which to use after that. This hierarchy of routes is closely connected to the chunks of actions observed in the previously described study by Huys et al. 2015.

A last hierarchical concept in spatial cognition worth being introduced here is the importance of decision points. Quoting Kevin Lynch: "[P]eople heighten their attention at such places and perceive nearby elements with more than normal clarity. This tendency was confirmed so repeatedly that elements located at junctions may automatically be assumed to derive special prominence from their location" (Lynch 1960, pp. 12–13). Decision points often are important subgoals in navigation (Spiers and Maguire 2008, Stachenfeld et al. 2017). Even more, they appear to be of importance in the animal kingdom as well: the aforementioned "preplay" in rats occurs particularly frequent at decision points such as crossroads (Pfeiffer and Foster 2013).

Models of Hierarchical Planning

After realizing the advantage of learning a transition model that allows for flexible and adaptive behavior in section 1.1, the question arises, how such a model can be extended to make use of the hierarchical structures described above. The topic of hierarchical planning is of great use in the research field of robotics and artificial intelligence, but since this thesis aims to develop a biologically plausible model, we limit ourselves here to reviewing two approaches inspired by human cognition and refer to the robotics literature for application-oriented systems (see for example Kaelbling and Lozano-Pérez 2011). One approach developed by Wiener and Mallot 2003 to explain the behavior of human subjects in the navigation tasks described above, is the "fine-to-coarse" planning heuristic. It was formulated as an algorithm by Reineking et al. 2008 and describes a heuristic that makes use of the hierarchical structure of the environment. In fine-to-coarse planning, the agent's representation of space consists of a topological state-action graph of discrete places and transitions between them. By adding nodes representing regions as an additional layer and connecting them to nodes of discrete places in the respective region, the representation becomes hierarchical. By creating a so-called *focal representation*, illustrated in figure 1.2 in working memory, the agent plans on multiple levels of resolution simultaneously: Regions further away are not represented in full resolution (figure 1.2b), planning yields only a coarse route to the goal (Wiener and Mallot 2003). In the region currently occupied by the agent, space is represented in full resolution, which enables the agent to take the correct step towards the transition to the next region. At least in this planning phase, the nodes allowing for transition to the next region must be connected to this next region node to allow simultaneous planning on different levels. After taking a step towards the goal, the focal representation is updated (ibid.). Following region transition, places within the newly occupied region are represented in detail in working memory. This results in high computational cost due to frequent replanning, but low memory requirements for fine-to-coarse planning, since the route does not have to be stored in full detail (Wiener and Mallot 2003).

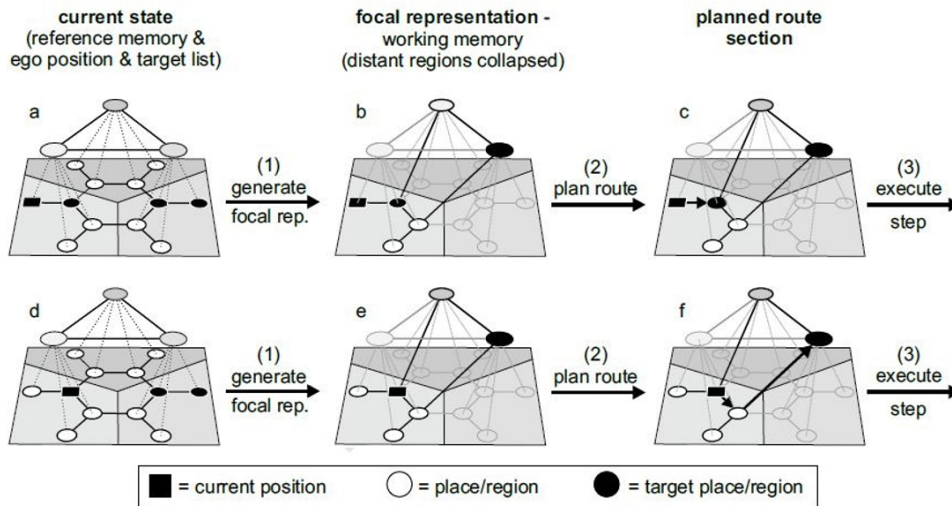


Figure 1.2: Fine-to-coarse route planning as described by Wiener and Mallot 2003. Three regions represented by the nodes on top comprise four places each. When planning the next step towards the goal on the right, the agent creates a focal representation in working memory. (Adapted from Wiener and Mallot 2003)

A second approach to hierarchical planning and action-execution is *hierarchical reinforcement learning* (HRL, Sutton et al. 1999). To tackle problems of traditional reinforcement learning (Sutton and Barto 1998) such as lack of generalization and problems in large task spaces, Sutton et al. 1999 introduced the options framework of HRL (figure 1.3). Whereas an agent in traditional reinforcement learning has the same set of actions to choose from in each state, the HRL agent is equipped with so-called *options*, that correspond to setting subgoals many timesteps ahead Botvinick 2008. In the typical example of the four room problem (figure 1.3a), an option can be to reach the door leading to the room closest to the goal. This option, denoted as "o1" in figure 1.3b then determines the next action choices ("a2-a4"). Botvinick and Weinstein 2014 describe HRL in the context of model-based vs. model-free learning and provide an overview over neurobiological findings concerning a neuronal implementation of HRL (see also Botvinick 2008 and Botvinick et al. 2009). Concerning the selection of subgoals, Sutton et al. 1999 assume that "the subgoals are given and do not address the larger question of the source of the subgoals". 15 years later, Botvinick and Weinstein 2014 still discuss it as a focus of HRL research. This problem of learning appropriate subgoals, termed the *option discovery problem* is fundamental challenge for hierarchical reinforcement learning (Botvinick and Weinstein 2014). With the model introduced in this thesis, we hope to shed some light on how a hierarchically structured neural network as representational apparatus of the spatial task domain might be of help in discovering suitable subgoals.

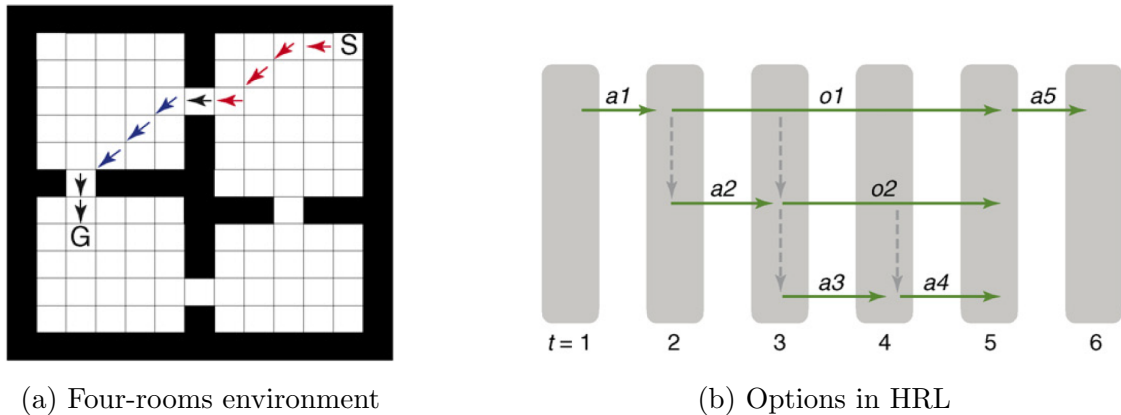


Figure 1.3: Hierarchical Reinforcement learning. In the four-room environment shown on the left, the agent can execute primitive actions (transition to neighboring places). To make reinforcement learning more efficient, Sutton et al. 1999 introduced so called options (shown on the right as $o1$ and $o2$), constituting a mid-term policy that influences action selection ($a1$ - $a5$) and is comparable to setting subgoals. In the example on the left, such an option could be to reach the nearest doorway. (Figure adapted from Botvinick 2008)

1.3 A Neuronal Model for Topological Navigation

As indicated in the previous sections, the goal of this thesis is to implement biologically plausible navigation based on a topological representation of space as described by Baumann 2019. The model is then extended to represent a hierarchy of places and regions in the environment. Conditions for biological plausibility are realistic timing properties of the neurons as well as adaption of synaptic weights based on local information available at the synapse alone, i.e. depending solely on pre- and postsynaptic neuronal activity. Furthermore, mechanisms contained in the model are based on neurobiological findings. Since the aforementioned discovery of hippocampal neurons with spatially selective firing-fields (O’Keefe and Nadel 1978), similarly functioning cells have been the starting point of many neural network models of spatial cognition (Burgess and Recce 1994, Voicu and Schmajuk 2002, Bicanski and Burgess 2018, Li et al. 2020). In these models, the environment is encoded in a layer of place cell-like neurons. A frequent approach in literature is to perform goal-directed navigation by spreading a gradient for hill-climbing from the goal throughout the network (Voicu and Schmajuk 2002). To do this in a biologically more plausible way, Ponulak and Hopfield implemented a spiking neural network in which a wave of activity is released from a set of cells representing the goal (Ponulak and Hopfield 2013, see also Hopfield 2010). The passing wavefront then alters the synaptic weights so that goal-directed weights are strengthened temporarily, which results in a so-called *synaptic vector field* (Ponulak and Hopfield 2013). The model proposed in this thesis builds on the approach of Ponulak and Hopfield, but with some important differences: Instead of encoding a metric representation, the network represents environmental topology: Significant features or places become represented by individual cells. In combination with synaptic connections corresponding to actions, this results in a state-action neural network based

on salient features of the environment, a neuronal version of a state-action graph as described by Baumann 2019. In addition, the developed framework is extended by a hierarchical structure, in which additional cells connect to low-level features belonging to a common region. To summarize, the relation of the neural representation to the external world is threefold:

- A topology of neural connections is learned that corresponds to the topography of the underlying features in reality. Through higher-level neurons, the network is able to encode regions in space and thus represent a hierarchical structure in the environment.
- The agent's position is encoded by a population vector of neuronal activity in the network. This activity typically forms an attractor state Amari 1971.
- Goal-directed movements are planned and then executed by temporarily altering the synaptic weights in the network so that goal-directed weights are strengthened. Then, following the resulting vector field by executing the learned state-action model, the agent navigates to the goal.

This thesis thus aims to investigate the behavior of a model combining the state-action approach and voting-based action selection of Baumann 2019 with the synaptic vector field-based, neurobiologically motivated planning from Ponulak and Hopfield 2013. Navigational behavior of the model is probed in multiple simulated scenarios. Beforehand, it is hypothesized that this hierarchical structure increases planning depth and helps the agent to navigate to suitable subgoals.

Chapter 2

Methods

2.1 Neuronal Dynamics

After the short introduction in section 1.3, this chapter goes into detail on the proposed model and its biological foundation. First, the mechanics of the spiking neurons used to represent the environment are explained, followed by the mechanism of synaptic plasticity to generate synaptic vector fields. Then, an explanation of how these vector fields can be used to generate goal-directed movements together with a demonstration thereof follows. The chapter concludes with a detailed description of the implementation of hierarchy in the neuronal model. A list of all parameters is attached in appendix A.2.

Neuron Model

The model described in this thesis implements a state-action representation of the agent’s environment (Baumann 2019) as a neural network. To do so, leaky integrate and fire neurons with little modification from the approach in Ponulak and Hopfield 2013 take the place of graph nodes. Directed synaptic weights replace the graph edges and give the connection strength between the respective neurons. The weights of these excitatory connections can be learned in a Hebbian manner (Hebb 1949), requiring a dedicated learning phase (as demonstrated in Ponulak and Hopfield 2013). Since the focus of this thesis lies on planning and goal-directed navigation, weights are given beforehand, depending on the effort to get from one feature to the other. Long distance and difficult terrain corresponds to low (or even nonexistent) weights. The exact network architecture depends on the simulation and is described in chapter 3. A representation sufficient for navigation is achieved as soon as each position the agent occupies in the environment yields a stable bump of activity in the network. The membrane potential u_m of each neuron is governed by a differential equation:

$$\tau_m \frac{du_m(t)}{dt} = -(u_m(t) - u_r) + R_m(i_{sens}(t) + i_{syn}(t) + i_{ns}(t) - i_{inh}(t)) \quad (2.1)$$

Here, τ_m denotes the membrane time constant and u_r the membrane potential at rest. R_m is the membrane resistance, i_{sens} the sensory input described in detail later, i_{syn} the incoming synaptic currents from other cells in the network and i_{ns} an optional, neuron-specific, Gaussian noise current used only in the planning depth simulation (section 3.2). i_{inh} denotes the global inhibitory current described in more detail in appendix A.1. Simulated in discretized time, the membrane potential is updated each timestep, following an Euler integration of equation 2.1.

Sensory Input

One of the advantages of equation 2.1 is that sensory input can be integrated in various forms. Using visual input, each cell can be modeled as receptive to a particular feature of the environment, such as in the aforementioned feature-graph of Baumann 2019. This is in line with the discovery of cells in the medial temporal lobe being sensitive to visual features on different levels of abstraction (Steinmetz et al. 2011). Development of these feature sensitive cells is imaginable via competitive learning as described by Rumelhart and Zipser 1985 and applied in Krotov and Hopfield 2019. Upon recognition of a feature, sensory input current to the respective cell is high. Network parameters, especially the inhibitory current, are set so that a stable peak of activity is present in the navigation phase, resembling an attractor network (Amari 1971, Wilson and Cowan 1973, Hopfield 1982). When neurons are modeled as sensitive to recognized features, the center of gravity of the sensory input shifts as the agent moves through the environment. This change in input current shifts the activity bump through the network. Bicanski and Burgess 2018 implemented such a shift of an activity bump by an external signal, albeit not with spiking neurons. Figures 2.3c-f illustrate an activity bump, indicated by active neurons in yellow, moving through the network. Which synaptic input is used in the navigation phase depends on the simulated scenario, and combination of multiple sensory inputs is possible by summation. In chapter 3 for example, place recognition is used, and single cells represent discrete positions in the discretized environment and receive sensory input if the agent is at that specific position.

Synaptic Input

In addition to the sensory input, each neuron k integrates incoming currents from all connected neurons j in the network. These currents are summed up in a supralinear fashion, again based on Ponulak and Hopfield 2013:

$$i_{syn,k}(t) = a_{syn} \cdot \tanh \left(b_{syn} \sum_j H(i_{out,j}(t)) \right) \sum_j w_{jk} i_{out,j}(t) \quad (2.2)$$

Here, a_{syn} and b_{syn} are positive constants and $H(i)$ is a step function which equals one for positive arguments and zero otherwise. $i_{out,i}$ is the current going out from neuron j and w_{jk} the synaptic weight between the two neurons. This process of synaptic input summation is repeated for each neuron k . In combination with the global inhibitory system, the supralinear summation (Nettleton and Spain 2000, Urakubo et al. 2004) limits the spread of activity from randomly spiking neurons by preferring multiple incoming currents over a single one of equal strength. To maintain a stable activity bump created by sensory and synaptic input, the global inhibitory current i_{inh} described in appendix A.1 is tuned accordingly. Such a stable activity bump is necessary for the network to function as a localization system in the environment. If multiple activity bumps were present, the network does not provide clear information to the agent. The same issue occurs if there is no such bump present.

Action Potential

As soon as the membrane potential of a neuron (equation 2.1) surpasses the threshold of u_{thresh} , an action potential is released. This process is modeled as the release of the resulting synaptic current (spike) that immediately reaches its maximum value i_{spike} and then decays exponentially to zero, governed by the time constant $\tau_{spike} = 25$ ms (Ponulak and Hopfield 2013):

$$i_{out,k}(t) = i_{spike} \cdot \exp\left(-\frac{t}{\tau_{spike}}\right) \quad (2.3)$$

where t is the time since the last spike. After the release of an action potential, the neuron goes through a dead time in which the membrane potential is kept at zero. Instead of a neuron specific inhibitory current as a mechanism of spike-frequency adaptation (Ponulak and Hopfield 2013), a longer refractory period until the wavefront has spread through the network is chosen, as implemented in the same paper. Both approaches facilitate the wavefront propagation (ibid.) that is explained in more detail in the next section. To keep the model simpler and because the additional neuron-specific inhibitory current does not alter the network's behavior in principle, the decision was made to use only an extended refractory period.

2.2 Synaptic Plasticity

Synaptic Weight Updates

Route planning in the sense of preparing the network for goal-directed navigation is started by activating the neurons associated with the goal (Ponulak and Hopfield 2013). A goal for navigation thus is defined by a set of features and their corresponding cells. The activation might either be the result of upstream processing or of a motivational drive, as implemented in many computational models (e.g. Voicu and Schmajuk 2002, Arbib and Bonaiuto 2012). In such models, the agent learns to associate drive satisfaction with a set of neurons representing the future goal. Here, this process is not modeled in more detail. Following a sufficiently strong activation of these cells, a wave of spiking neuronal activity throughout the network is released (Ponulak and Hopfield 2013). Ponulak and Hopfield relate such an activity wave to rhythmic activity in the brain, specifically theta-oscillations (Lubenov and Siapas 2009, Patel et al. 2012) and sharp wave ripples (Ellender et al. 2010). During the planning phase that prepares the representation for goal-directed navigation, sensory input is turned off to prevent an activity wave spreading from the cells representing the current position of the agent (Ponulak and Hopfield 2013). Other possible causes of error in wavefront propagation are backward flowing waves due to synaptic currents exciting cells around the current wavefront towards the goal. These are prevented by the refractory period described in the previous paragraph. As the wave propagates and the neurons in the network spike consecutively, synaptic strength is adapted based on a spike-timing-dependent plasticity rule (adapted from Kempster et al. 1999, see also Gerstner et al. 2014, chapter 19.2). The updates occur in incremental fashion each time either the post- or the presynaptic neuron spikes.

This discretization of increasing weights in quanta, only when spikes occur, is expressed in the continuous notation through delta functions:

$$\Delta w_{jk}(t) = a_{Dec} + d \cdot (\delta_k(t) a^+(s_j) + \delta_j(t) a^-(s_k)) \quad (2.4)$$

Here, s_j is the time since neuron j 's last spike. a_{Dec} is a constant that governs the weight decay and d is the parameter controlling the polarity of the STDP process as a globally released neurotransmitter (Seol et al. 2007) depending on the network mode. $d = -1$ is used for wave propagation and $d = 0$ for goal-directed navigation. $d = 1$ is for topology learning, which is not used in the simulations included in this thesis. In equation 2.4, $a^+(s_j)$ and $a^-(s_k)$ define the strength of the STDP/Anti-STDP process, which is why they are preceded by $\delta_i(t)$ and $\delta_j(t)$ respectively, indicating if the post- or presynaptic neuron has spiked.

$$a^+(s_j) = -A^+ \cdot \exp\left\{-\frac{s_j}{\tau_{STDP}}\right\} \quad (2.5)$$

$$a^-(s_k) = -A^- \cdot \exp\left\{-\frac{s_k}{\tau_{STDP}}\right\} \quad (2.6)$$

A^+ is a constant regulating the strength of the STDP process and $A^- > A^+$ controls the Anti-STDP process. This implementation of STDP is based on Ponulak and Hopfield 2013, but with more convenient notation. The STDP process described above allows simulating an exploratory phase as well. Since simulations in chapter 3 only require wavefront propagation and consecutive navigation, a simplified update rule is used:

$$\Delta w_{jk}(t) = d \cdot (\delta_k(t) - 1) \delta_j(t) \cdot a^-(s_k) \quad (2.7)$$

Here, weight updates are conducted only if the presynaptic neuron j spikes, yielding the same net behavior for reverse STDP. Another difference to Ponulak and Hopfield 2013 is that only on the last previous spike from the respective neuron is considered in $a^-(s_k)$ and $a^+(s_j)$, not the whole spike train. This implementation detail should not influence the results, since the decay parameter $\tau_{STDP} = 20$ ms (Morrison et al. 2007) of the STDP process is way smaller than the time between passing wavefronts. Thus, interactions between spike trains collapse to nearest-neighbor interactions between spike-pairs (Morrison et al. 2007, Morrison et al. 2008). As the wavefront passes, Anti-Hebbian updating as observed by Roberts and Leen 2004 leads to goal-directed synapses becoming stronger than their opposing counterparts. In equations 2.4 to 2.6 this is expressed by A^- being larger than A^+ (Ponulak and Hopfield 2013) and in equation 2.7 only the reverse process is left. The process of reverse STDP weight updates is illustrated for two cases in figure 2.1. Neuron j spikes at time t_1 , followed by a spike of a connected neuron k . Depending on the time difference between the two spikes, the reverse synaptic weight w_{kj} increases by a large factor (short time difference, on the left) or a small factor (large time difference, on the right).

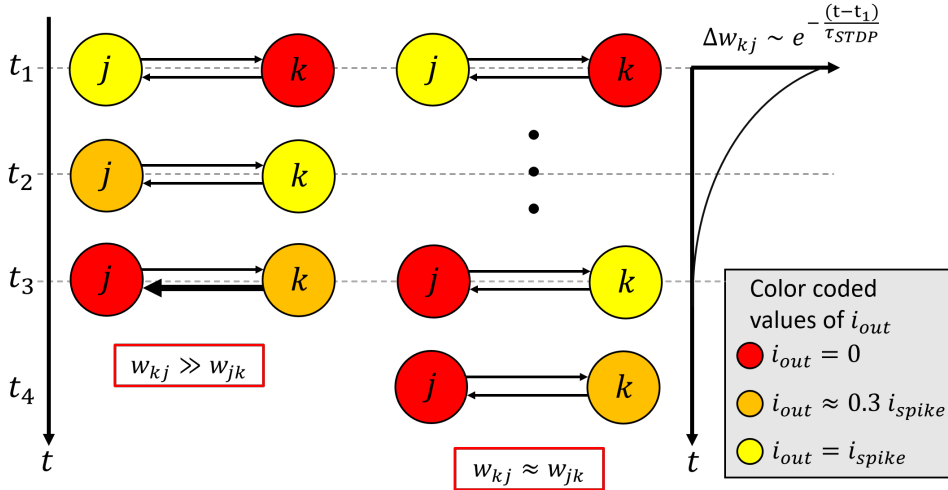


Figure 2.1: Updating of synaptic weights through STDP. Neuron j spikes at time t_1 . Left: Neuron k spikes shortly after j , at time t_2 , leading to temporary strengthening of the synapse leading from neuron k to j indicated by the bold arrow. Center: Neuron k spikes at time $t_3 \gg t_2$, long after neuron j , leaving the synaptic weights unchanged. Right: Exponential decay of weight adaption depending on spike-timing.

Note that during wavefront propagation, sensory input is down-regulated to prevent an activity wave from the current position. Furthermore, the inhibitory system is turned off (Ponulak and Hopfield 2013). When planning a path towards a new goal, the temporary change in synaptic strength is reset to the value before planning. There is experimental evidence that this might be realized with sharp wave ripples in the brain (Lubenov and Siapas 2008).

Displaying the Synaptic Vector Field

To illustrate the representational network, it is useful to depict each cell at its *preferred location*, denoted by \mathbf{x}_j . This position, where the cell is most active, can then be embedded in the environment to yield an interpretable map of the network, such as in figure 3.3a. The network is fine-tuned so that a bump of activity does not move without change, either in the afferent sensory signal or in the network connectivity. This is a fundamental property the network has to fulfill in order to allow navigation guided by a synaptic vector field (SVF). Recalling the threefold relationship of the neural network to the environment in which the agent moves (learning of the topology, localization, goal-directed navigation, cf. section 1.3), the question remains, how the created vector field can be transformed into actual motor commands. Following the idea of a state-action graph in which edges are associated with actions such as a turning angle to change the direction of forward movement (Baumann 2019), the implementation of motor signals in this model is fundamentally different from the one described in Ponulak and Hopfield 2013 and Hopfield 2010. In both these models, the bump of activity is shifted by the SVF. The agent then tries to align the current sensory input with the proposed position from the moved bump in the network (Hopfield 2010) or simply is drawn to the new (metric) position of the bump by a mechanism not described in more detail in Ponulak and Hopfield 2013.

Here instead, each synaptic connection is associated with a movement command α_{jk} as illustrated in figure 2.2. This movement command is equivalent to the action to get from the place or feature represented by neuron j to the one represented by neuron k and is interpreted as vector instead of a turning angle as in Baumann 2019.

$$\boldsymbol{\alpha}_{jk} = \mathbf{x}_k - \mathbf{x}_j \quad (2.8)$$

To illustrate where the synapses leading away from neuron j go, a vector is introduced as a weighted average over action commands associated with outgoing synapses:

$$\mathbf{r}'_j = \sum_k w_{jk} \boldsymbol{\alpha}_{jk} = \sum_k w_{jk} (\mathbf{x}_k - \mathbf{x}_j) \quad (2.9)$$

Since finite attractor networks underlie a central bias that draws activity bumps to the center (Muller et al. 1987, Burgess and O'Keefe 1996), it makes sense to consider the change of weights Δw_{jk} as a result of the plasticity process instead of the full weight. The *synaptic vector* then can be written as

$$\mathbf{r}_j = \sum_k \Delta w_{jk} (\mathbf{x}_k - \mathbf{x}_j) \quad (2.10)$$

and can be interpreted as the direction in which the agent will move, controlled by the motor system described in the following section. Due to this interpretability, \mathbf{r}_j will be used to illustrate the SVFs. An exemplary SVF is illustrated in figure 2.3, where a wavefront of activity is spread throughout the network as described above. Again, the grid-like architecture of the network does not reflect the positions of the neurons in the brain tissue, but is an illustrative embedding based on their preferred locations. The resulting vector field in figure 2.3b points towards the goal marked by a circle.

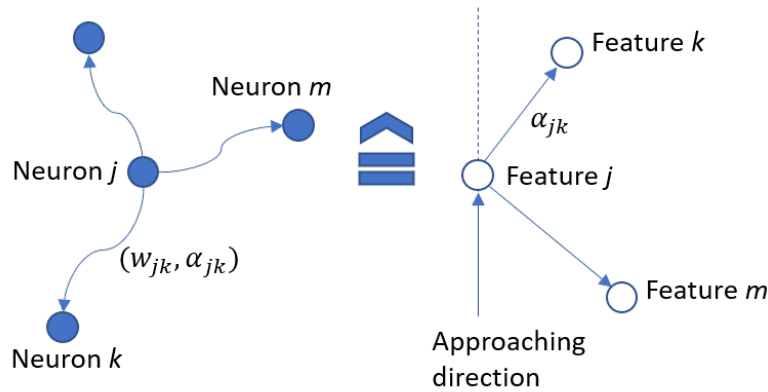


Figure 2.2: Neuronal state-action network as the agent perceives feature j : The action command α_{jk} is associated with the synapse from neuron j to neuron k . α_{jk} corresponds to the action to get from the feature represented by neuron j to the one represented by neuron k , transferring the state-action graph by Baumann 2019 to the neuronal domain.

In such a state-action network, the external topology is reflected in the synaptic connections. Note however, that neurons are not necessarily located next to each other in the brain (as indicated in figure 2.2). In fact this is quite improbable, since neurons representing space such as hippocampal place cells have to undergo frequent remapping and do not seem to be topographically positioned in the brain (O’Keefe and Nadel 1978, Wilson and McNaughton 1993, see also Samsonovich and McNaughton 1997).

2.3 Motor Efferences

To produce motor commands, a voting process is conducted. Since actions are associated with synapses in the manner of a state-action graph, votes are weighted by the synaptic weight. By doing so, less feasible actions with low weights can be overruled. To account for the current position of the agent encoded in the bump of activity, each vote is weighted by the activity of the presynaptic neuron j , given by the outgoing current i_{out} .

$$\dot{\mathbf{x}}_{ag}(t) = \frac{1}{i_{spike}T} \int_{t-T}^t \sum_{j,k} i_{out,j}(t') \boldsymbol{\alpha}_{jk} dt' \quad (2.11)$$

By substituting the synaptic action label $\boldsymbol{\alpha}_{jk}$ with the vector between preferred neuronal positions (cf. equation 2.8) and using the synaptic vector of equation 2.10, the velocity can be written as

$$\dot{\mathbf{x}}_{ag}(t) = \frac{1}{i_{spike}T} \int_{t-T}^t \sum_j i_{out,j}(t') \mathbf{r}_j(t') dt' \quad (2.12)$$

Thus, the result of the voting process over all neurons j determines the direction of movement. Since shifts of direction occur on a much slower timescale (> 100 ms) than neuronal spikes that show (sub-)millisecond timing precision (Nemenman et al. 2008), the integrational time window T is introduced. Figure 2.3 illustrates movement of an agent through a learned 2D-environment. The cells are modeled as place cells, i.e. receive sensory input in the shape of a Gaussian around their preferred location (cf. appendix A.3). After creating the SVF as described above, the agent shifts to the navigation phase, bringing weight updating to a halt. During this phase shift, network activity is reset. Then, a bump of activity builds up due to sensory input (figure 2.3c) representing the agent’s position and navigation commences. Executing motor commands determined by the voting process of equation 2.12 guides the agent correctly towards the goal.

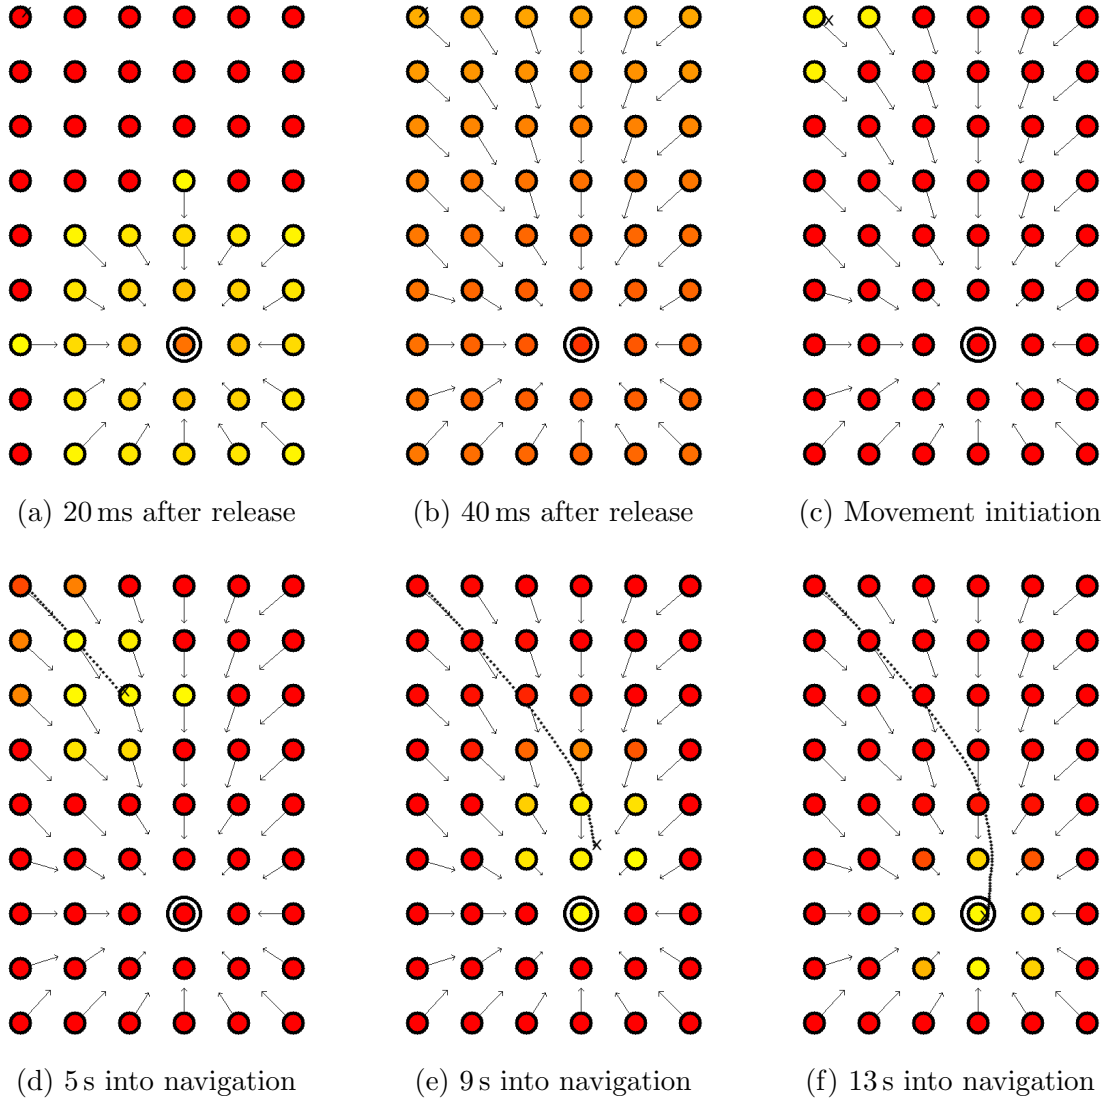


Figure 2.3: Demonstration of model function: (a) and (b): A wavefront of spiking activity is released from the neuron marked by a circle, representing the goal. The inner color of the neurons indicates the outgoing current i_{out} , with red indicating low currents and yellow for high currents. The resulting synaptic vectors (cf. equation 2.10) pointing towards the goal are indicated with arrows. (c): The agent marked by an X starts its movement in the top left corner. Sensory input causes the corresponding neurons to spike. (d)-(f): The agent continues its trajectory to the goal.

2.4 Neuronal Hierarchy

As explained in chapter 1, hierarchical structure is essential to understanding human spatial cognition. To capture what Botvinick terms a correlational, hierarchical structure (Botvinick 2008), this thesis proposes cells disconnected from direct sensory input. This approach is motivated by the discovery of neurons in the medial temporal lobe responding to image properties at multiple levels of abstraction (Steinmetz et al. 2011). Higher-level cells as modeled in this thesis possess connections to multiple cells on the level below them as illustrated in figure 2.4. Apart from

these two properties (lack of direct sensory input and relatively uniform connections to associated lower-level neurons) and their far-reaching connections to neighboring higher-level cells, they are modeled in the same way as the previously described cells, following the equations from section 2.1. The mentioned distinctions define their function in activity transmission. Through these properties, they represent a hierarchically superior level as described by O'Reilly and Frank 2006: they remain active over finer-grained events at levels below and provide a more abstract canvas for planning. Since this thesis is concerned with spatial cognition, the higher-level neurons will be referred to as *region cells*. Further abstractions beyond these two levels shown in figure 2.4 are imaginable. Note that the used type of higher-level cells does not necessarily have to be located in the same part of the brain. It is imaginable that neurons with more abstracted receptive fields are located in other brain areas than the proposed topological system. The prefrontal cortex has been proposed as containing such an abstracted representation for decision making (Martinet et al. 2011). As illustrated in figure 2.4, the regions are overlapping. Neurons in the transition zone between two regions (referred to below as *transition neurons*) are connected to multiple region cells.




Figure 2.4: Hierarchical network structure as implemented in the model. Higher-level cells effectively representing the regions r_1 and r_2 are connected to feature sensitive cells on the lower level via interlevel weights w_{inter} . Neurons on the lower level are connected to nearest neighbors via w_1 and to second-nearest neighbors via w_2 .

Path-Planning in a Hierarchically Structured Network

During the process of wavefront propagation illustrated in figure 2.3, activity quickly spreads across the network of region cells. Figure 2.4 illustrates the overlap of such regions. Neurons in the transition zone are characterized by their affiliation to multiple region neurons. Due to their larger number of connections to higher-level cells than their counterparts further from the border, they receive stronger synaptic inputs when both region neurons are activated. This in turn leads to new waves of spiking activity from these neurons even if they have not yet been reached by the propagating wave of activity on the lower level. The influence of this accelerated signal transmission on navigation with SVFs and on planning depth is investigated in the next chapter.

Chapter 3

Simulation Experiments

3.1 Comparison to Human Navigation

Does the proposed hierarchical structure lead to behavior in line with studies on human navigation? To investigate this question, two experimental studies on human behavior from Wiener and Mallot 2003, introduced in the first chapter as HN1 and HN2, were simulated. The original study by Wiener and Mallot showed an influence of hierarchical structure in the environment on human navigation, described in detail in section 1.2. Experiment HN1 showed a preference of the subjects to take routes that allow access to the target region as soon as possible even when alternative routes were equally short. In experiment HN2, subjects tended to choose routes crossing as few region boundaries as possible.

3.1.1 Entering the Goal Region (HN1)

As in the underlying study from Wiener and Mallot 2003, the virtual agent is placed in a virtual environment of two interconnected islands shown in figure 3.1. Instead of the exploration phase in which subjects learned the topology of the twelve discrete places marked by unique landmarks, the agent’s representation of the environment is given as a neuronal state-action network as explained in chapter 2. In the representational network that is shown in figure 3.2, each neuron corresponds to one of the twelve landmarks distributed across two islands, as illustrated in figure 3.1. This is justified by the fact that the landmarks used in the virtual environment provide by far the most salient structure for feature learning as described by Baumann 2019. To simplify experiment description and discussion, neurons, the places they represent and the positions at which these are located are numbered in the same way: neuron 0 represents place 0 that is located at position 0. Connections between the neurons correspond to transitions between the respective positions.




Figure 3.1: An illustration of the experimental setup from Wiener and Mallot 2003. The twelve connected places shown on the left are distributed on two islands (as can be seen in the middle), each marked by an individual landmark. The rightmost picture depicts one of these landmarks, a monkey. Most subjects preferred to reach the goal region (i.e. island) as quickly as possible in navigation tasks, neglecting alternative, equally short routes. (Figure adapted from Wiener and Mallot 2003)

In the study of Wiener and Mallot, most subjects recognized the hierarchical structure of the environment. Landmarks on each island belonged to different semantic categories: animals such as the monkey in figure 3.1 or cars. To capture such a correlational hierarchy, two region neurons are added, one for each island. These region neurons shown on top of figure 3.2 are connected to all lower-level, landmark-sensitive neurons of that island. Additionally, they are connected to all neurons at transition point (next to a bridge) on the opposite side of the river with synaptic weight $w_{inter} = 10$. Both region neurons are connected with $w_r = 3$. Neurons on the lower level are connected to neighbors corresponding to places one and two edges away in the environment with weights $w_1 = 6$ and $w_2 = 1$. These weights are chosen to yield synaptic currents on the order of a realistic value for synaptic currents, i.e. 0.3 nA (Gerstner et al. 2014, ch. 3.1.2). An overview of the parameters used is given in tables A.1 and A.2 in the appendix.




Figure 3.2: Cross-section of the network architecture in simulation HN1. Of the twelve neurons on the lower level only the foremost are fully shown, the other two rows are indicated by the dotted circles on the right. r_1 and r_2 denote the neurons on the hierarchically superior level. The synaptic weights w_1 , w_2 , w_{inter} and w_r connect the neurons.


Analogously to the human subjects after the learning phase, the virtual agent is placed in one of the twelve locations and asked to navigate to a given target position. As in Wiener and Mallot 2003, the configurations of start and target positions are split into different route types to analyze the resulting route choices dependent on the spatial configuration. These route types are listed in table 3.1.

Route Type	Start Place \rightarrow Target Place
A1, A2	4 \rightarrow 8, 5 \rightarrow 7
B	5 \rightarrow 9
C	2 \rightarrow 6
D	5 \rightarrow 6
E1, E2	4 \rightarrow 11, 5 \rightarrow 10

Table 3.1: Route types for the navigation tasks in the environment displayed in 3.1. Due to symmetry in the network, many route types are equivalent for the model and can be subsumed in one categorie, e.g. the route from place 3 to place 11 is equivalent to 5 \rightarrow 9, hence both belong to category B.

HN1: Vector Field in Small Networks

When given the task to navigate to one of the twelve landmarks, the virtual agent activates the neuron representing the goal (referred to from now on as goal neuron for brevity) to start the planning phase as described in chapter 2. The wavefront of spiking activity spreads throughout the network creating a synaptic vector field as shown in figure 3.3. For better understanding, neurons are plotted at the position they represent, even if they need not be aligned like that in the brain. The synaptic vector illustrated as arrows are calculated following equation 2.10. Before each new navigation task when the agent is placed at the new start position, weights are reset and a new wave of spiking activity is released from the next goal. Both with and without hierarchical structure in the network, the synaptic vectors point towards the goal, indicated by a circle in figure 3.3. To answer the question of the influence of a hierarchical structure, the figure contains the vector fields for both cases. The synaptic vectors from neurons further away from the river separating the two islands appear to have a slightly stronger component towards the center if the network is extended by a hierarchical structure (on the right of figure 3.3). To quantify the difference of the vector fields, the resulting route choices are analyzed below.



(a) SVF without hierarchy

(b) SVF with two region neurons

Figure 3.3: Resulting vector fields after wavefront propagation in the neural network representing the environment of the two islands. The neuron corresponding to the start place is marked by a square, the goal neuron by a circle. The configuration of start and goal position corresponds to route type C from table 3.1. Colors represent the outgoing current of the respective neurons, red indicates low currents, yellow indicates high currents. The network of the left plot consists of only one layer without hierarchical structure. In the right plot, two region neurons are added as a hierarchically superior layer, illustrated as larger circles on the frame of the respective region. Since the synaptic connections of these regions are not associated with actions, they do not possess synaptic vectors.

HN1: Trajectories in Small Networks

After establishing the vector fields in the planning phase, the model shifts to the navigation phase. Since goal-directed navigation in the model is based on a voting process of the currently active neurons (cf. section 2.3), behavior depends on the width of the activity bump. The wider this bump is, the more neurons participate in voting, as defined in equation 2.12. In networks with few neurons that represent important decision points, this leads to erroneous movement signals, because neurons from places far away begin to participate in the voting process. To prevent such an effect, the activity bump is limited to the landmark currently occupied by the agent. An alternative solution that allows for broader bumps of activity would have been to increase the size of the network and have multiple cells representing each place. In the small network, trajectories consist of transitions between the discrete places. The transition probabilities are calculated from the synaptic weights going out from the respective neuron j (cf. figure 2.2):

$$p_{jk} = \Delta w_{jk} \left(\sum_{k'} \Delta w_{ik'} \right)^{-1} \quad (3.1)$$

Instead of using the full value of w_{ij} , the difference to the weight before planning is used, as discussed in section 2.3. The denominator in equation 3.1 ensures that the sum over all outgoing probabilities equals one:

$$\sum_j p_{ij} = 1 \quad (3.2)$$

Following these movement probabilities, the probabilities of the agent taking each route segment lead to the trajectories displayed in figure 3.4 and 3.5. Suboptimal route choices by the agent break the trial and are not shown as trajectories.

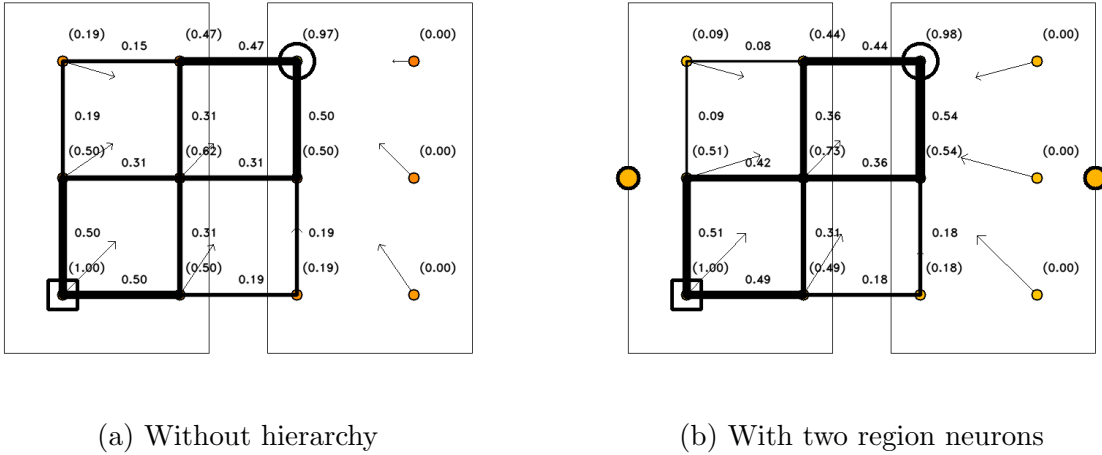


Figure 3.4: Trajectories of the virtual agent in route type C from start position (square) to goal (circle) printed on top of the SVF to illustrate how the vector field leads to goal-directed behavior. Each neuron is shown at the position of the landmark it represents. Numbers in brackets denote the probability of the agent to reach this position, the numbers next to the edges indicate the probability of the agent to travel along this edge during one run.

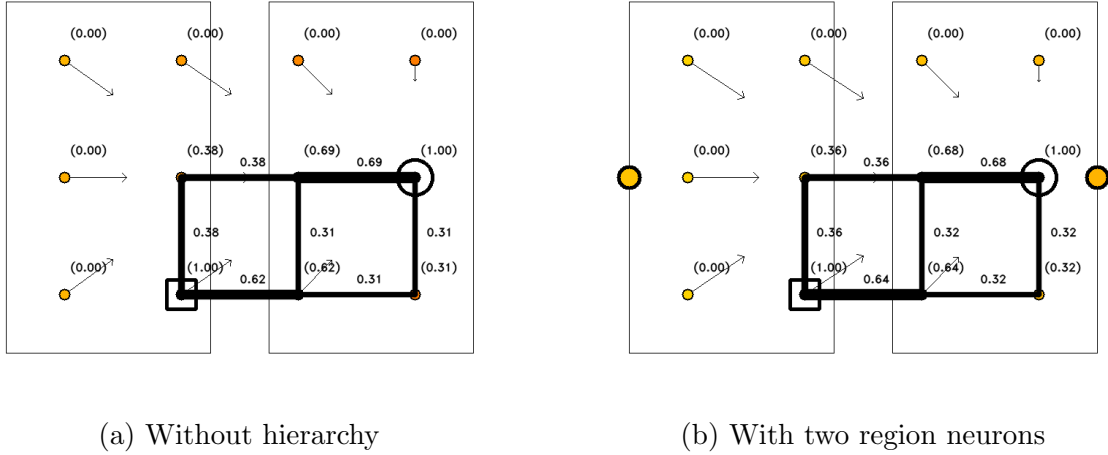


Figure 3.5: Vector field and trajectories of the agent in route type E without and with hierarchical structure in the neuronal architecture. The route allowing fastest possible access to the goal region is the one on the bottom, chosen by the agent in 62% respectively 64% of trials.

In this manner, a probability for reaching the goal and choosing the route to fastest-possible region transition can be obtained for each route type in table 3.1. Figure 3.6a shows the overall success rate for finding an optimal (shortest) route to the goal. Both the hierarchical and the single-level model perform well and have a 97% (single-level) and 96% (two-level) probability for reaching the goal on an optimal route. These numbers are pooled over all route types A-E, and as figure 3.6a shows, variation is low. The worst performance of a single-level, SVF-based agent is a 92% probability for an optimal solution in route type B. The worst performance of the two-level model could also be observed for route type B, at 93%. An agent choosing random, not necessarily goal-directed actions at each position performs considerably worse, at 11% probability. Similarly to the SVF-based agents, the random walk-agent performed worst in route B, reaching only a 6% probability to navigate optimally. Figure 3.6b shows the probability of taking the route with the quickest possible transition to the goal region, together with the chance level. Chance level is determined by the performance of an agent choosing randomly between equally fast, shortest-possible routes to the goal. Detailed calculation of the chance level can be found in appendix A.4. With and without hierarchical structure, the SVF-based model traverses to the goal region at chance level $\pm 8\%$ for all route types with two exceptions in route types C and D.

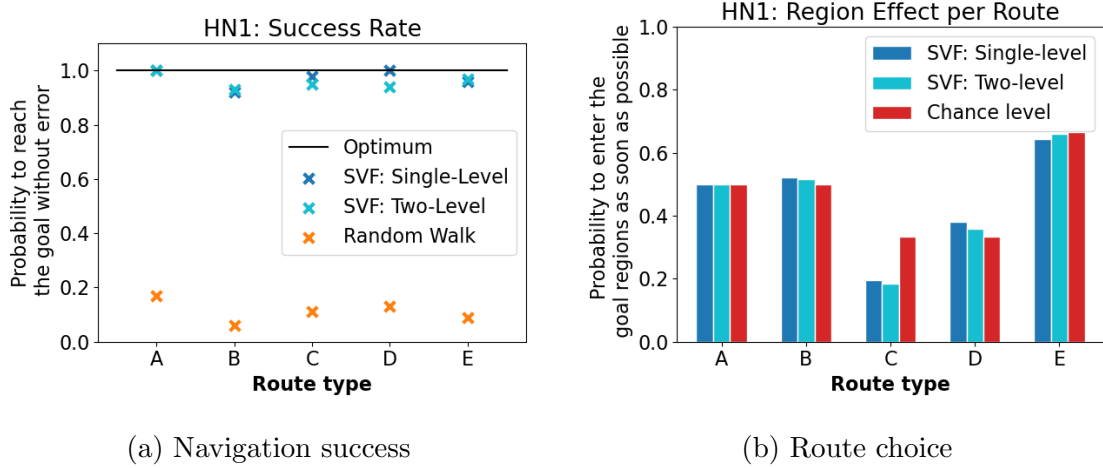


Figure 3.6: Comparison of model behavior with and without hierarchical structure for each of the five route types listed in table 3.1. Additionally, behavior is compared to chance level (choosing randomly between optimal routes) and to a random walk-agent.

In route D, the single-level SVF-based model has a probability of fastest-possible region transition that is 14% higher than chance level, whereas the hierarchical two-level model is only 8% above chance level. The other notable difference occurs in route type C, where performance of the SVF-based models is considerably worse than chance level: 19.6% with one neuronal level, 18.4% with two, compared to 33.3% chance level. Pooled across all five route types, the single-level model chooses fastest-possible region transition with a 44.8% chance, the two-level model with 44.4% and chance level lies at 43%. Note that this must be distinguished from random walk that would achieve a way lower probability since the chance agent chooses only between alternative, optimal routes as explained above.

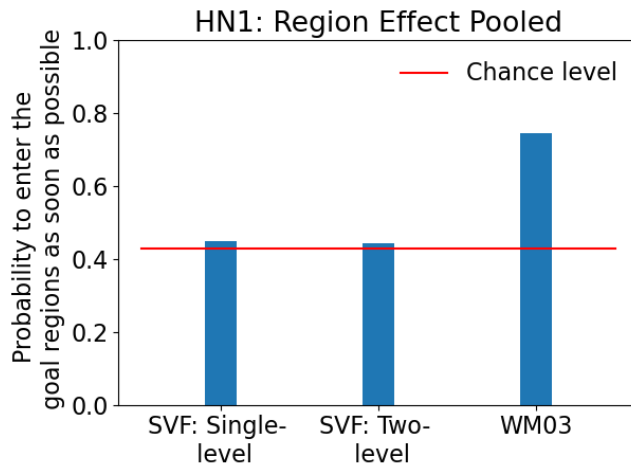


Figure 3.7: Probability for fastest-possible transition to the goal region pooled across route types A-E. Performance of the agent guided by a SVF is compared to the behavior of humans subjects from Wiener and Mallot 2003, indicated here as "WM03", as well as chance level (red).

As can be seen in figure 3.7, the additional hierarchical level does not alter the probability of the agent to take routes leading to the goal region as quickly as possible. Both SVF-based models, consisting of one, or two levels respectively, lead the agent to the goal region at approximately chance level. Human subjects from Wiener and Mallot 2003 instead took the route allowing fastest possible transition to the goal region with an average probability of 67.7% (first test-block) and 74.4% (second test-block).

HN1: Increasing Network Size

The initially suspected effect of the hierarchical structure depends on the activation of neurons in transition areas that are connected to multiple regions (cf. section 1.3). To investigate if the number of low-level neurons connected to one region neuron, i.e. the region size, alters the SVFs, the same experimental setup as above is repeated with a larger network. In theory, this should increase the effect of the hierarchical structure, since activity on the lower level takes longer to traverse the network, while the region neurons already activate the boundary region. The number of neurons on each region is multiplied by a factor of 3 along both axes. Neurons directly by the separating river (next to the bridges connecting the island) are still connected to both regions, leaving the general network architecture from figure 3.2 untouched. By activating the goal neuron, a wave of spiking activity spreads throughout the network and creates an SVF as before.

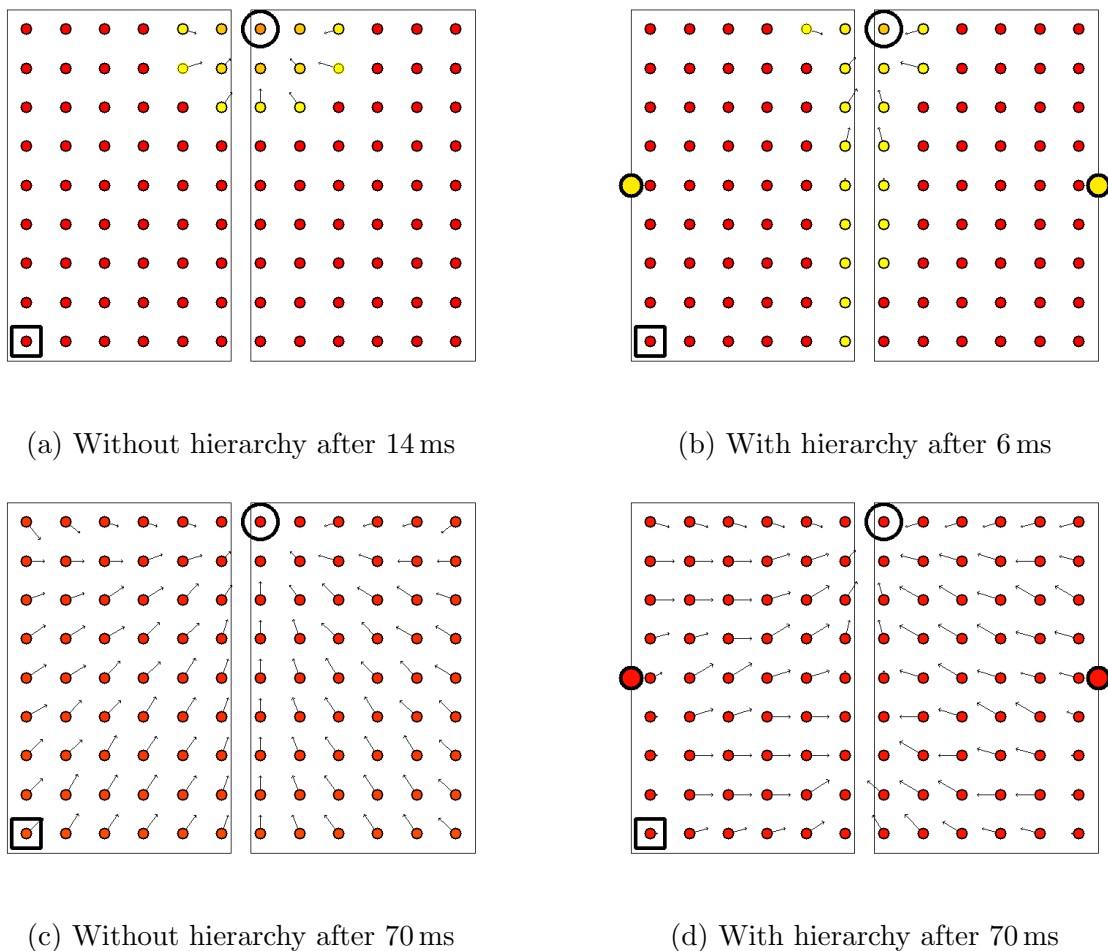


Figure 3.8: Comparison of wavefront propagation between a network without (shown on the left) and with (on the right) hierarchical structure. For the networks with two layers, region neurons are shown in increased size on the rectangle framing the regions. In the bottom figures, activity has spread throughout the whole network, having established an SVF.

Comparing figure 3.8a to figure 3.8b, the hierarchically structured network shows a wavefront with a different shape as the network without hierarchical structure. In the two-level network, neurons in the transition region spike early, after approximately 4 ms, while the single-level network without hierarchical structure shows a wavefront spreading radially from the goal neuron. The SVF of the single-level network displayed in figure 3.8c points towards the goal, while the SVF in the hierarchically structured network (figure 3.8d) shows a bias towards the transition region. Synaptic vectors close to the goal still point towards the goal, but the further away from it, the bigger the tendency of the vectors to be oriented towards the river. Multiple neurons in the boundary region do not possess any synaptic vector. Additionally, synaptic vectors on the river side of the goal region point back towards the other region.

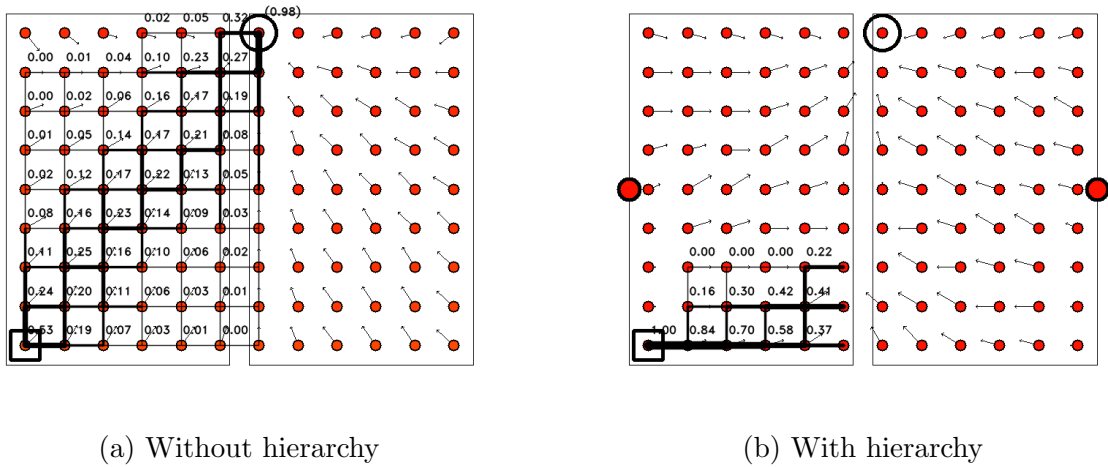


Figure 3.9: Influence of hierarchical structure on agent trajectories in larger networks. The trajectories result from the SVF illustrated in figure 3.8. Numbers next to horizontal arcs denote the probability of the agent travelling along that edge.

Figure 3.9 illustrates the trajectories resulting from the vector fields in figure 3.8. While the single-level network performs well and reaches the goal in 98% of trials, the hierarchical network gets stuck in front of the border and does not cross the river to the other region. The single-level network prefers to navigate the agent directly towards the goal, maneuvering diagonally across the island (cf. figure 3.9a). The agent controlled by the hierarchical model shown in figure 3.9 however, begins by navigating relatively straight towards the boundary region where it then gets stuck.

3.1.2 Navigation Across Multiple Region Boundaries (HN2)

A second experiment from Wiener and Mallot 2003 introduced briefly as HN2 in section 1.2 is simulated to test the navigation behavior of the model across multiple region boundaries and towards multiple goals. To do so, the environment shown in figure 3.10 is represented in a manually constructed network. The virtual park of the study from Wiener and Mallot 2003 consisted of twelve places, each with a unique landmark belonging to one of three semantic categories: cars, animals and buildings. Three regions could be distinguished as indicated by the different shades of grey in the left image of figure 3.10 and places from the same region were connected. The network representing this environment is illustrated in figure 3.11.




Figure 3.10: The environment of experiment HN2. The hexagonal environment layout is shown on the left. In the middle, a birds-eye view of the 3D environment illustrates the virtual environment. The right image depicts one of the six crossroads as encountered by the subjects in Wiener and Mallot 2003. As all twelve places, it is marked with a landmark, in this case a go-kart. (Figure adapted from Wiener and Mallot 2003)




Figure 3.11: Unfolded network representing the hexagonal environment structure of HN2. Neurons on top represent the three regions. r_3 is shown twice to illustrate its connection to both r_1 and r_2 . Lines in black illustrate synaptic connections corresponding to paths between places (state-transitions) in the environment. Neurons marked in yellow represent the three goals for route type A. The neuron for the start position is marked by a square.

As in the previous simulation, each lower-level neuron corresponds to a place, i.e. receives sensoric input if the agent currently perceives this landmark. Neurons on this low level are connected to neighbors corresponding to places one edge away with weight $w_1 = 6$. Neurons in the middle of figure 3.11 represent the six places on the corners of the hexagonal path (cf. figure 3.10). Neurons representing the places at the end of dead-end streets away from the hexagon are shown in the bottom.

HN2: Navigating to Multiple Goals

To compare behavior of the model with the human subjects from Wiener and Mallot 2003, the same wayfinding tasks are set. Both route types consist of three goals each and are illustrated in figure 3.12. Each route type offers two equally long solutions to reach the three goal places. These two options differ in the number of region borders they cross. Wiener and Mallot had added a third, asymmetric route type as a distractor to prevent subjects from learning and exploiting the symmetry of routes A and B instead of planning. Since the model proposed here is obviously incapable of such an understanding, this distractor is omitted in the simulation. To test the influence of a hierarchical structure in the neuronal representation, two different networks architectures are tested. One of these is equipped with a hierarchical structure to capture the regions of the environment. Each of the three regions is represented by a dedicated neuron. These region neurons (illustrated in the upper part of figure 3.11) are connected with weight $w_{inter} = 10$ to all neurons representing places contained in the region, as well as neurons of transition places, i.e. places from which the respective region can be reached. As in HN1, the region neurons themselves are interconnected with weight $w_r = 3$.

In the planning phase, a wave of spiking activity is simultaneously released from all goal positions by increasing the membrane potential of the goal neurons beyond threshold. For route type A, these goals are places 1, 6 and 9 while for route type B they are places 3, 6 and 11, denoted by "G" for "goal" in figure 3.12.




Figure 3.12: The two route types used in Wiener and Mallot 2003 to analyze human behavior in wayfinding tasks across regions are simulated in HN2 with the SVF-based model. Both route types allow for two equally long routes that differ in the number of region boundaries they cross. (Figure adapted from Schick et al. 2019)

HN2: Synaptic Vector Fields

As in the previous simulation, the resulting SVF is illustrated by using the synaptic vectors from equation 2.10. In both route types, the synaptic vectors of each neuron point towards the nearest goal, with the exception of neuron 1. Here, the vector is ambiguously pointing between position 10 and 2 as can be seen in figure 3.13b. Comparing the SVFs in the hierarchical network with the single-level architecture does not show any visible difference, neither for route type A nor B. Consequently, it makes sense to analyze the resulting behavior of the agent. Following equation 3.1 of the previous section, the transition probabilities $p(j, k)$ of the agent are calculated from the synaptic weights. With these transition probabilities, the agent's trajectories can be illustrated (figure 3.14). Each time the agent reaches a subgoal, the synaptic weights must be reset and a new wave must be initialized before commencing navigation. Otherwise, the SVF would draw the agent back to previously visited goals.

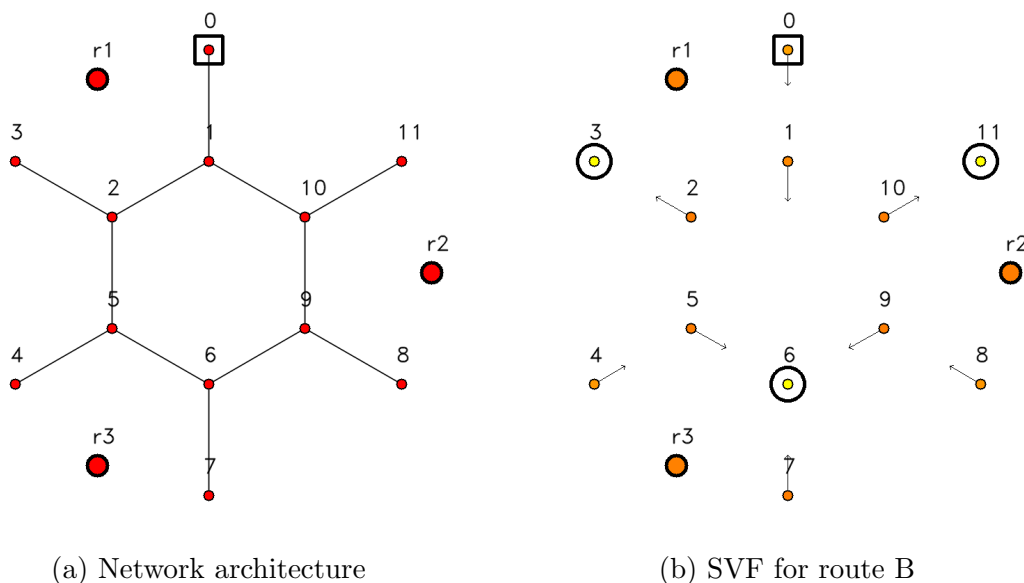
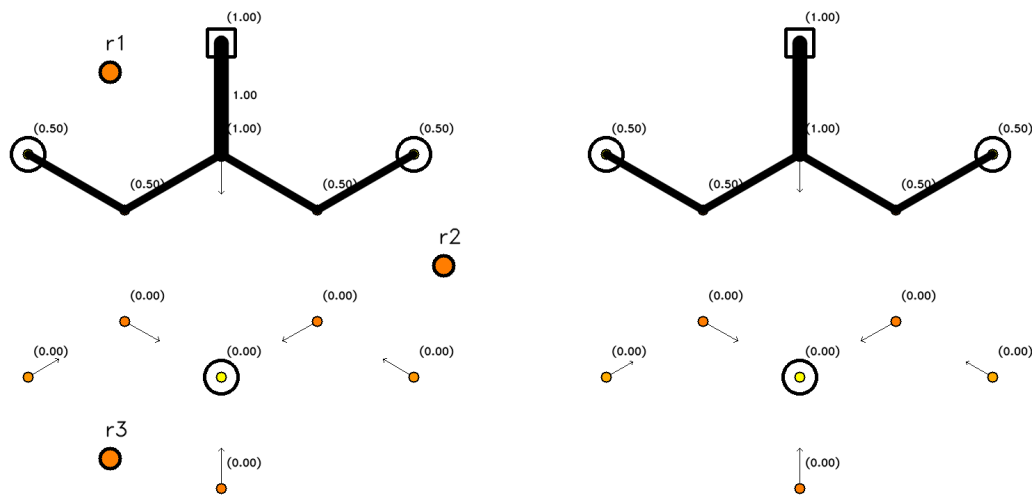


Figure 3.13: Hierarchical network architecture and exemplary SVF for simulation experiment HN2. The three region neurons r1-r3 are illustrated in larger size. Before navigating to the three positions of route B, the agent at position 0 activates the three goal neurons at positions 3, 6 and 11. The resulting vector field is shown on the right.

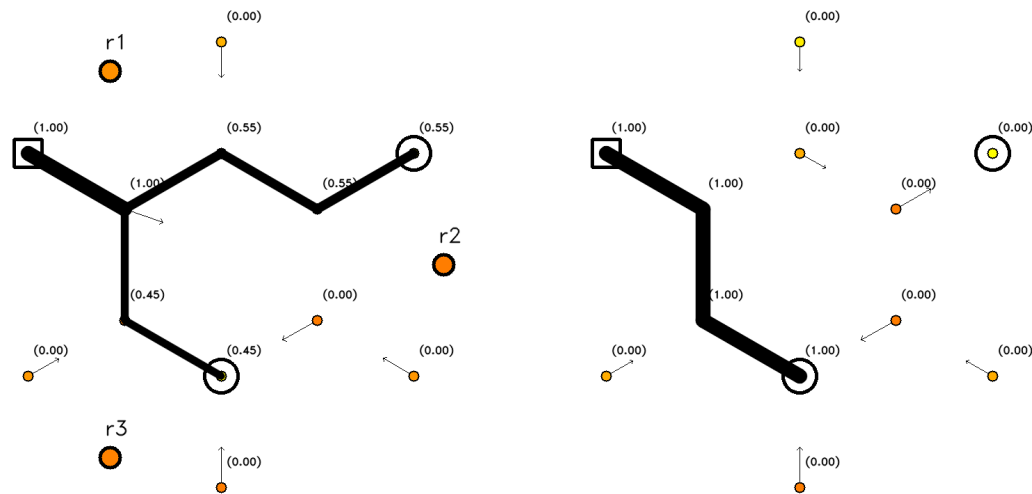
HN2: Trajectories to Multiple Goals

Starting from position 0, both the single-level and the two-level network lead the agent to position 1 with 100% probability, reaching the first subgoal of route A. The discrete nature of the environment then forces the agent to choose one of three possibilities: return (to start position 0), continue left (to 10) or right (to 2). The agent's choice at this position also distinguishes the two shortest routes, illustrated in figure 3.12.



(a) 2-level network, agent at position 0

(b) 1-level network, agent at position 0




(c) 2-level network, agent at position 3

(d) 1-level network, agent at position 3

Figure 3.14: Navigation to the first two subgoals of route B: Comparison of transition probabilities in a hierarchical network (left) with a single-level network (right). Edges with high probabilities are shown in bold. The agent's current position is marked by a square. The probability of the agent to reach a position is denoted in brackets next to it.

For both route types, the agent based on a hierarchically structured network performs similarly to the single-level model: going onward from position 1 to either position 2 or 10 with a 50% probability each. After reaching position 2 or 10 in route type A, both agents (with and without hierarchical structure) reach the other two subgoals with 100% probability. In route type B, the both agents commence to the first subgoal (marked by a circle in figures 3.14a and 3.14b): position 3 or 11, respectively. Figures 3.14a and 3.14b illustrate these trajectories. Then however, performance between the two networks differs. After navigating from position 3 back to position 2, the hierarchical model has a 55% chance of navigating to position 11 - back via position 1 (the upper leg of the split trajectory in figure 3.14c). This corresponds to a suboptimal solution, being one edge longer than going to the other subgoal, position 6, first. With a 45% probability, the hierarchical agent will continue along this optimal route. The single-level model chooses this optimal solution in 100% of cases, as can be seen in figures 3.14c and 3.14d, respectively.



(a) 2-level network, agent at position 6 (b) 1-level network, agent at position 6

Figure 3.15: After reaching the second subgoal at position 6 (marked by a square), the agent commences by replanning towards the remaining targets at positions 6 and 11.

When the agent has reached this second subgoal, replanning by resetting the synaptic weights and releasing a new wave of activity lead it to the last goal with 100% probability. This is shown in the trajectory in figures 3.15a and 3.15b.

HN2: Crossing Region Borders

For comparison to human behavior, the probability of choosing the route crossing less region boundaries, p_{XLESS} , is of interest. Counting only the shortest routes to the goal (as did Wiener and Mallot 2003 with the human subjects), this probability is fully determined by the transition probabilities from position 1 to position 2 and 10 respectively. These probabilities, denoted by $p(1, 2)$ and $p(1, 10)$, are illustrated as labels on the respective edge in figure 3.14a for the hierarchical network and 3.14b for the single-level network.

$$p_{XLESS}(\text{Route A}) = p(1, 10) \quad (3.3)$$

$$p_{XLESS}(\text{Route B}) = p(1, 2) \quad (3.4)$$

Disregarding suboptimal route choices of the two-level model leads to a tie between the single-level and the two-level model. Both models choose the route across less borders with a 50% chance (figure 3.16), showing no influence of the region neurons on model behavior in this regard. Performance of the SVF-based models is equal to chance level (50%) for both route types. As in HN1, chance level refers to choosing randomly between the two optimal solutions illustrated in figure 3.12, not to a random walk. Wiener and Mallot 2003 investigated the same choice in human subjects, counting only subjects who recognized the region structure of the environment. These chose the route crossing fewer boundaries far more frequently, i.e. in 65% of all cases in route type A and in 71% of all cases in route type B (figure 3.16).

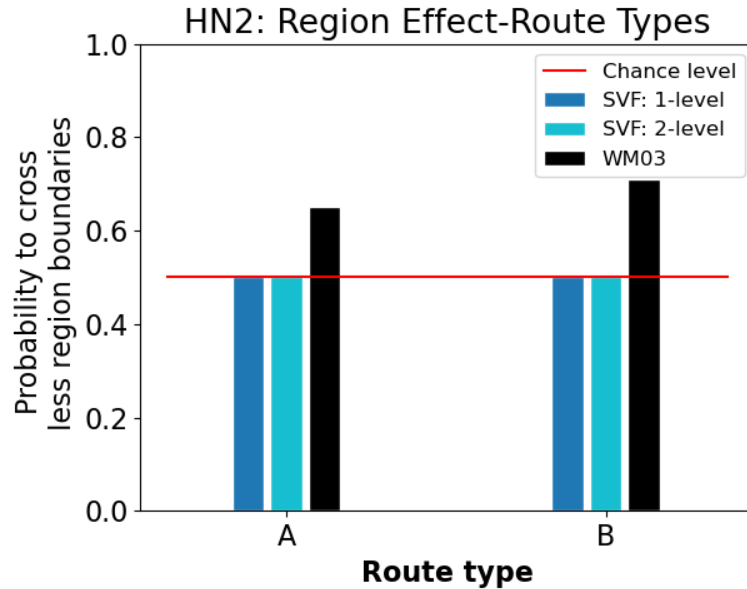


Figure 3.16: Comparison of the chosen routes between the two artificial agents based on SVF-planning with single-level or two-level networks. In addition, the behavior of human subjects from Wiener and Mallot 2003 as well as the chance level of 50% is shown.

3.2 Planning Depth

After investigating the influence of hierarchical structure on the behavior of an artificial agent in the previous section, the goal of this section is to investigate if the neuronal hierarchy can increase the agent’s planning depth. A multitude of studies underline the importance of hierarchically structured behavior for human cognition (cf. section 1.2 of the introduction). Huys et al. 2015 link hierarchical organization to the ability to solve more complex problems. In spatial cognition, this corresponds to navigating larger areas. Can the proposed neuronal structure aid in vector field-based navigation across large networks? To investigate this, the navigational ability of the model is tested with and without hierarchical structure in two settings:

1. With a neuron-specific Gaussian noise current i_{ns} of mean 0 and width 1 nA as additional input to each neuron (cf. equation 2.1)
2. With a neuron-specific Gaussian noise current i_{ns} of mean 0 and width 1 nA as well as a global inhibitory current linearly increasing during the wavefront propagation phase with a rate of $\delta_{inh} = 0.015 \text{ ms}^{-1}$

By doing so, it is evaluated if the hierarchical structure makes the network more robust against noise (and therefore increases planning depth). The second setting simulates a decaying wavefront. In this rather crude manner it will be investigated if the hierarchical structure can carry the wavefront and with it the guiding vector field across larger distance, thereby increasing the navigable area.




Figure 3.17: The network of 24x24 neurons is organized (a): without regions, (b): in two by two regions and (c): in four by four regions of core size six by six. Including overlap with other regions, regions in the network center contain eight by eight neurons (the area marked in grey). In the top right corner, the goal neuron is drawn in yellow.

Network and Noise Configuration

A two-dimensional network of 24x24 neurons is set up. The neurons' preferred locations are set on a grid as illustrated in figure 3.18. A wave of activity is then initiated by activating the goal neuron in the upper right corner of the network, illustrated yellow in figure 3.18. The wave of activity then spreads across the network, creating a goal-directed vector field (cf. section 2.2). Each lower-level neuron is connected to neighbors one and two edges away on the grid with weights w_1 and w_2 respectively (cf. figure 2.4). To test the influence of a hierarchical structure, as well as the influence of region size, three networks, illustrated in figure 3.17 are compared:

1. A network without hierarchically superior neurons, consisting of only one level
2. A network with an additional layer of two by two region neurons that split the environment up into four regions
3. A network in which the additional layer of region neurons consists of four by four neurons, splitting the environment up into 16 regions of equal size

As before, regions overlap, i.e. neurons on the border are connected to region neurons of all adjacent regions. In the second case, each region contains 13x13 neurons considering the overlap. In the third case with 16 regions, each consists of up 8x8 depending on the position in the grid and hence the amount of neighboring regions (whose boundary neurons then are attached as well).




Figure 3.18: Illustration of the network model used in the "Planning Depth" experiment. The grid of 24 by 24 neurons is divided into regions. Shown are regions of size six, but the simulation is repeated with smaller regions, i.e. more region neurons. In the upper right corner the goal neuron is illustrated in yellow. The diagonal arrow away from it indicates the axis along which the vector field quality is evaluated.


Resulting Vector Fields

Despite the applied noise, the network without hierarchical structure displays a stable propagation of the wavefront as shown in figure 3.19a, 3.19c and 3.19e. The number of spurious spikes unrelated to the wavefront is very low (< 1 per experimental run) and the wavefront spreads evenly and radially outward from the goal. When the increasing inhibitory current is applied in the second part of the simulation experiment (figure 3.19b, 3.19d and 3.19f), the wave arrives at a halt, traversing only a limited range in every run. In the simulation run shown in 3.19f), only eight neurons are activated along the diagonal axis.


Repetition of the same simulation with a hierarchically structured network of 2x2 regions leads to a much faster spread of activity. Whereas the single-level network takes 80 ms for half of the network in figure 3.19c, the two-level version takes only a quarter of that time to get this far (figure 3.20c). However, the hierarchical network suffers from a large number of spurious spikes at random positions. These spikes, unrelated to the propagation of the wavefront, lead to partially chaotic spread of spiking activity as can be seen in figure 3.20a, 3.20c and 3.20e. When the wavefront is released accompanied by the inhibitory current (figure 3.20b, 3.20d and 3.20f), the activity still spreads throughout the whole network in all twenty runs. Furthermore, the inhibitory current used to simulate decay of the wave appears to decrease the number of spurious spikes. Nevertheless, the inhibitory current slightly slows down the wavefront, an effect that can be observed with smaller regions as well. Comparing the activity in the network with two-by-two regions without the inhibitory current after 30 ms (figure 3.20c) to the same network with inhibitory current after 38 ms (figure 3.20d) clearly shows the slower speed of activity propagation. In both scenarios, the wavefront in the hierarchical network spreads preferably from the transition regions along the region boundaries indicated by black lines (figures 3.20c and 3.20d).

The third network features smaller regions and thus a larger number of regions (4x4) as illustrated by the black lines separating regions in figures 3.21a, 3.21c and 3.21e. The number of spurious spikes at random positions in the network is again much larger than without hierarchical structure. However, the distortion of the circular shape of the wavefront due to activity at the region borders is lower than in the network with larger regions. When simulating a decaying wavefront by the use of incremental inhibitory current, the activity is still able to spread through the whole network, as illustrated in figures 3.21b, 3.21d and 3.21f.


SVF in Single-level Networks




(a) 1-level network with noise after 30 ms




(b) 1-level network with noise and wave decay after 30 ms




(c) 1-level network with noise after 80 ms



(d) 1-level network with noise and wave decay after 80 ms



(e) 1-level network with noise after 160 ms



(f) 1-level network with noise and wave decay after 160 ms

Figure 3.19: SVF-creating wavefront spreading through network without hierarchical structure. Left: with noise, right: with noise and incremental inhibitory current.

SVF in Two-level Networks of Two-by-two Regions

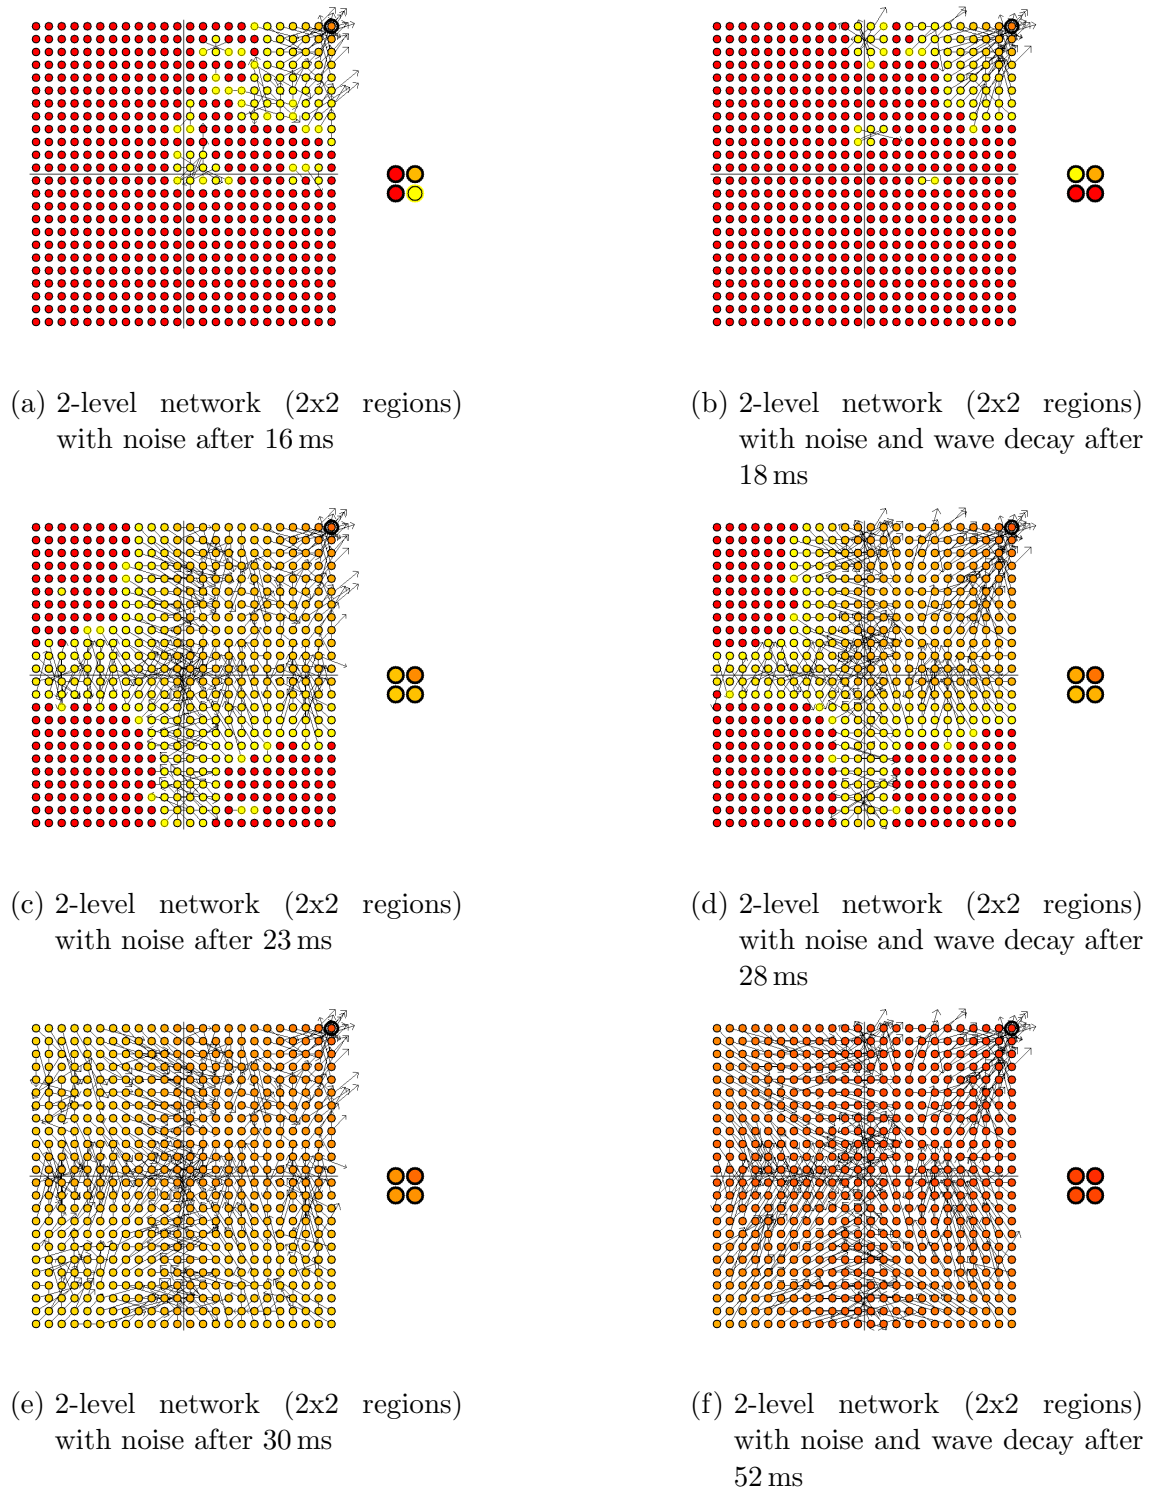


Figure 3.20: SVF-creating wavefront spreading through a network with hierarchical structure of two-by-two superior region neurons. Region neurons are shown on the right. Left: with noise, right: with noise and incremental inhibitory current.

SVF in Two-level Networks of Four-by-four Regions

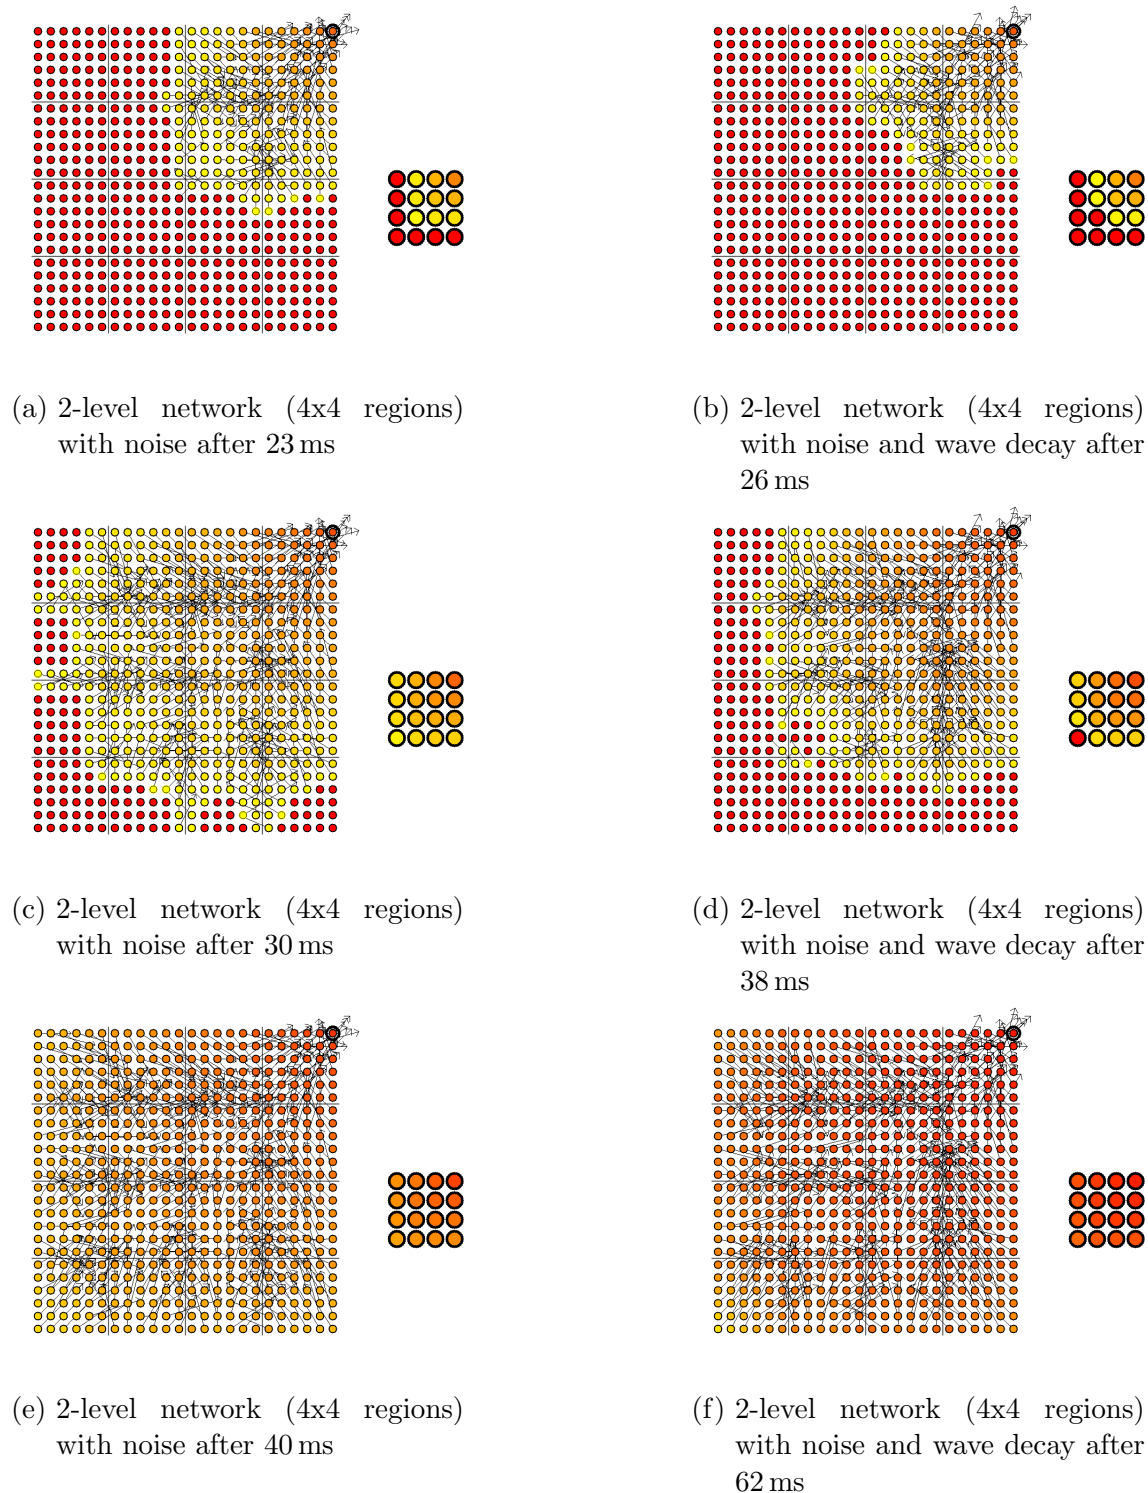


Figure 3.21: SVF-creating wavefront spreading through a network with hierarchical structure of four-by-four superior region neurons. Region neurons are shown on the right. Left: with noise, right: with noise and incremental inhibitory current.

Vector Field Quality

Since the vector field guides the agent by means of the voting process described in section 2.3, successful navigation requires the vectors to point towards the goal. As a quantitative measure of goal-directedness, the vector field quality of neuron j is introduced as the average correlation of synaptic vectors in the neighborhood with the vector towards the goal.

$$S_N(j) = \frac{1}{n_j} \sum_{k \in N(j)} s_j(k) \quad (3.5)$$

where $N(j)$ is the neighborhood of neuron j , i.e. all neurons connected to it with non-zero weights. n_j counts the number of such neighbors for normalization. $s_j(k)$ denotes the correlation of the synaptic vector of neuron k with the goal-directed vector \mathbf{g} from neuron j :

$$s_j(k) = \begin{cases} 0 & \text{if } |\mathbf{r}_k| \cdot |\mathbf{g} - \mathbf{x}_j| = 0 \\ \frac{\mathbf{r}_k}{|\mathbf{r}_k|} \cdot \frac{(\mathbf{g} - \mathbf{x}_j)}{|\mathbf{g} - \mathbf{x}_j|} & \text{otherwise} \end{cases} \quad (3.6)$$

where \mathbf{r}_k is the synaptic vector of neuron k as in equation 2.10, and \mathbf{x}_j denotes the preferred location of neuron j . The scalar product of the normalized vectors in equation 3.6 equals one if they are parallel, minus one in case they are antiparallel and zero if they stand perpendicular. In conclusion, $S_N(j)$ measures the average correlation of synaptic vectors around neuron j with the goal-directed vector.

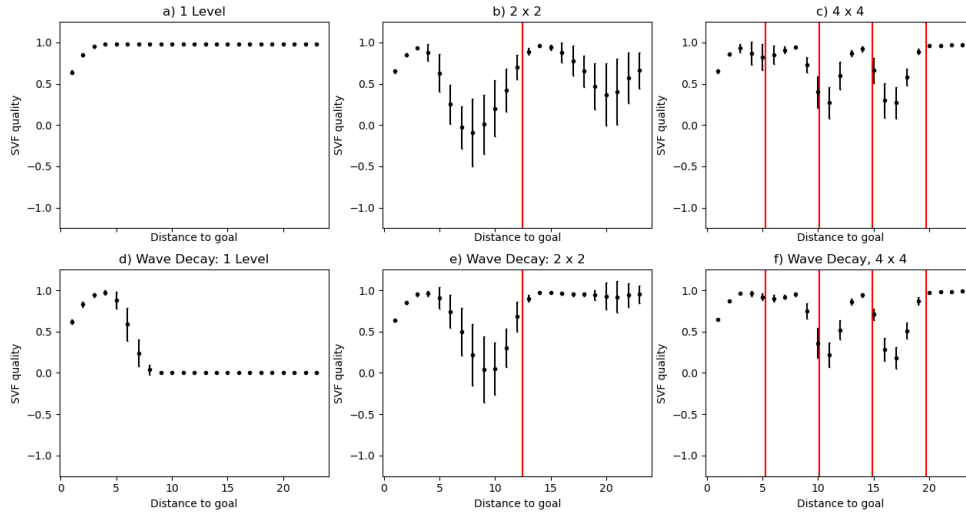


Figure 3.22: Vector field quality along the diagonal axis, $S = 1$ being the optimum: all vectors in the neighborhood point towards the goal. $S = -1$ corresponds to vectors pointing away from the goal. Shown is the mean value across twenty runs with respective standard deviation. Red lines indicate region boundaries. Bottom row: inhibitory current increases each timestep by δ_{inh} during the wave propagation to simulate decay of the wave.

Interpretation of this measure is as follows: The further away from the goal S_N declines, the greater the area that the agent can navigate. The S_N diagrams are calculated for each of the 23 neurons along the diagonal axis (excluding the goal neuron) illustrated in figure 3.18. By averaging over twenty simulatory runs, the mean value as well as the corresponding standard deviation are computed. Figure 3.22 displays the resulting diagrams as an indicator of planning depth. Without superior levels (figure 3.22a), the vector field quality $S_N(d)$ one neuron away from the goal on the diagonal axis is 0.64 ± 0.02 . Moving three neurons further away from the goal, that value rises to 0.98 ± 0.01 where it stays until the bottom left corner.

Adding a hierarchical structure consisting of two by two regions (figure 3.22b), the vector field quality around the goal is not affected. However, four neurons or more away from the goal, the vector field deteriorates and reaches its minimum quality of -0.09 two neurons behind the center of this region, four neurons before the boundary region around the red line in the figure. Noticeably, the standard deviation of S_N increases from initially 0.02 close to the goal up to 0.41 at the minimum of correlation between vector field and vector towards goal. Shortly after the region boundary, S_N almost reaches the optimum again, equaling 0.96 ± 0.02 . Moving further into the lower left corner of the network, vector field quality deteriorates again, however not as badly as before. The local minimum of 0.37 ± 0.38 is reached seven neurons after the boundary region. After that, a slow increase up to 0.66 for the bottom left neuron can be observed. An inverse correlation between standard deviation and vector field quality, i.e. standard deviation being high where the vector field is not pointing towards the goal, can be observed in all configurations with hierarchical networks in figure 3.22. With decreased region size, the vector field quality in the first region does not suffer as much as with larger regions. The indent at the first region boundary (figures 3.22c and 3.22f) is not as deep as for the network with 2x2 regions. Beyond this first region however, the vector field quality shows the same pattern of decrease after region borders (down to 0.18) and simultaneously increased standard deviation. Nevertheless, the standard deviation is smaller compared to the network with 2x2 regions, reaching a maximum value of 0.18.

When the wavefront is assumed to be decaying, the vector field quality of the single-level network is influenced the most. Figure 3.22d shows the rapid decay of S_N to zero. The hierarchical networks (figures 3.22e and 3.22f) show much less of an influence. Even 23 neurons away from the goal, they still reach vector field qualities of 0.95 for 2x2 regions and 0.99 for 4x4 regions.

Chapter 4

Discussion

Topological state-action encodings constitute a useful representation for spatial problem solving (Schölkopf and Mallot 1995, Arbib and Bonaiuto 2012, Baumann 2019). Such representations allow novel combinations of previously traveled route segments in goal-directed planning, a strong evolutionary advantage. Baumann 2019 proposed such a state-action model based on a large number of small image features ("microsnapshots"). However, Baumann's approach lacks a biologically plausible way of route planning and action implementation. Ponulak and Hopfield 2013 proposed a spiking neuron model in which planning was achieved by releasing a wavefront of spiking neural activity through a network of place cells, starting from the neurons representing the goal. The model presented here aims to close the gap between models based on topological representations and biologically plausible planning. Additionally, overwhelming empirical evidence stresses the importance of hierarchically structured behavior - and cognition - in humans (e.g. McNamara et al. 1989, Wiener and Mallot 2003, Botvinick 2008, Huys et al. 2015 and Schick et al. 2019). Thus, functional cognitive models of these structures are of deep importance to understanding cognition. To combine these three core concepts of topological representations, biologically plausible planning and inclusion of hierarchical structures, a neuronal state-action network was implemented. The neurons in the network represent distinctive sensory features (or places) just as the nodes of Baumann's state-action graph. Path-planning is based on SVFs as in Ponulak and Hopfield 2013 and region neurons represent extended areas in space. After the planning phase, the agent executes motor commands associated with synaptic connections based on voting of currently active neurons.

The discussion is built up as follows: Firstly, in the light of the results of the simulation experiments, the model is analyzed with regard to both effectiveness of goal-directed navigation and replication of human performance. Secondly, the question if the implemented neuronal hierarchy was able to increase planning depth is discussed. Thirdly, a general discussion of the model and comparison with other models of hierarchical planning follows. Suggestions for further research as well as applications conclude this chapter.

4.1 Performance in Wayfinding-Tasks

Effectiveness of Goal-Directed Navigation

The proposed implementation of SVF planning for neural state-action representations shows sound, goal-directed behavior. Agents based on a single-level neural network reliably choose the shortest possible route to the goal. In simulation experiment HN1, the agent based on a network of only twelve neurons was able to navigate to the goal on a shortest-possible route with an average probability of 97.2% pooled over all five route types (cf. figure 3.6a). An agent following a random walk performed considerably worse (figure 3.6a), showing the relative difficulty of the task. Simulation experiment HN2 confirmed the reliability of the SVF-based navigation. There, the single-level model reached the goal on optimal routes with a probability of 100% (cf. section 3.1.2). Slightly suboptimal performance in HN1 may be the result of a central bias in the vector field (Muller et al. 1987, Burgess and O’Keefe 1996, cf. section 2.2) that can be observed in figure 3.3: Neurons on the upper end of the island have a small vector component back towards the start position of the agent. Such a central bias could be prevented by adding a padding of cells around the network to achieve homogeneity in outgoing connections across the representational neurons.

Extension of the model by a hierarchical structure partially decreases the probability to reach the goal on an optimal route. In HN1, performance of the two-level network was similar to the single-level network with an average probability to take an optimal route of 96% (figure 3.6a). In route type B of HN2 however, the hierarchical model only chose the optimal route with a probability of 45%. In the majority of trials (55%), the agent based on the hierarchical model took a longer route, as illustrated by the movement probabilities along the splitting trajectory in figure 3.14c. Worse performance than without hierarchical structure (which found the optimal route with 100% probability) stands in contradiction to the better performance of human subjects that recognized the regional structure compared to those that did not (Wiener and Mallot 2003). Due to the process of SVF creation, this tendency to choose the route back across position 1 (cf. figure 3.14c) must be the result of the hierarchical structure speeding up activity propagation from the two remaining subgoals along this pathway. In the single-level network, the optimal route transmitted activity fast enough to "block" the other option. This can be understood when looking at the topology of the network representing the environment from figure 3.10: spiking of neuron 5 before neuron 2 and of neuron 2 before 1 led to the vector field pointing from 2 to 5. Consequently, following this vector field (figure 3.14c) resulted in zero probability for the agent to take the suboptimal route across position 1. The hierarchical structure however, challenged this way of computing optimal routes by speeding up activity transmission on suboptimal pathways, using neurons 11-10-1-2 and the superordinate region neurons. However, this is not in contradiction to literature: hierarchical structure is suspected to make planning more efficient, not necessarily more optimal (Stevens and Coupe 1978, Huys et al. 2015), as described in section 1.2.

In larger networks, a possible pitfall of the hierarchical structure becomes obvious: local minima at region borders can prevent the agent from commencing towards the goal. The trajectories in figure 3.9b illustrate this dilemma: When arriving at the border separating the two regions, the agent encountered a lack of movement commands and got stuck. A similar effect was observable in the "Planning Depth" experiment (section 3.2). The vector field quality measuring the correlation of the vector field with the direction of the goal showed a strong decrease in front of the region boundary, the area left of the red line in figures 3.22b and 3.22e. This can be explained by simultaneous spiking of the neurons in this boundary region due to their high input from two region neurons. Such simultaneous spiking was also observed in HN1 (cf. figure 3.8d) and is in line with the theorized influence of the hierarchical structure from sections 1.3 and 2.4: activation of neurons in boundary regions through strong synaptic input from higher levels. However, this input leads to wavefronts spreading in either direction from the boundary region. These wavefronts in turn create an SVF pointing towards the boundary region, i.e. from either side, thus creating a local minimum (see section 2.2 for the creation of SVFs). A possible solution to this dilemma could be to replan (including a reset of synaptic weights) when in such a local minimum, and inhibit the region node of the previous region. An encoding of which places are behind and which in front could be achieved by spike-timing of cells relative to brain rhythms. This so-called "phase precession" has been related to encoding information about the hippocampal place cells representing places behind or in front of the agent (O'Keefe and Recce 1993).

Replication of Human Navigation

In small networks, addition of a hierarchically superior level does not appear to induce a preference for fast transition to the goal region. As shown in figure 3.7, both the two-level network (44.4%) and the single-level network chose fastest access to the goal region only at chance level (43%) in simulation experiment HN1. This stands in contrast to the clear preference of human subjects for reaching the goal-region as soon as possible. Subjects of Wiener and Mallot 2003, abbreviated as WM03 in figure 3.7, chose such routes far above chance level (first test-block: $p_{asap} \approx 67.7\%$, second block: 74.4%) when faced with the same navigational task. A possible explanation for this behavior can be found when analyzing the wavefront spreading from the goal. Due to the small network size, the wavefront on the lower level reached the boundary region before the neurons there received enough synaptic input (cf. section 2.1) from the region neurons as theorized in section 2.4. Thus, the low-level wavefront prevented an influence of the higher levels on the shape of the wavefront and with it on the SVF (cf. section 2.2). Since navigation of the agent is based on the SVF, the effect of hierarchical structure on the agent's preference to transition to the goal region was very limited in small networks.

In larger networks, the hierarchical model tends to move towards boundary regions first, whereas the single-level version appears to preferably navigate straight towards the goal. The observed difference between the two models is illustrated in figure 3.9 for a network consisting of 108 neurons split up into two regions. The hierarchical network led the agent towards the boundary region, whereas the single-level network led him straight towards the goal. This behavior is in line with the observation of Wiener and Mallot 2003 that human subjects prefer to navigate to the transition re-

gion first. Regarding the network structure, this effect is only possible due to strong input to neurons in the boundary region as a result of overlap between the regions as illustrated in figure 3.2. Exciting neurons in boundary regions beyond threshold released a wavefront from this boundary region as hypothesized in section 1.3. This wavefront then created an SVF leading the agent straight towards the boundary region. Further simulations could test if this tendency to navigate towards the boundary region becomes stronger, the deeper the goal lies in the next region. In theory, the wavefront on the lower level would need longer to reach the agent, and thus increase the effect of the higher level. Indeed, subjects in the island-experiment (replicated here as HN1) of Wiener and Mallot 2003 showed an increased effect for goals further away from the boundary region.

The influence of the proposed hierarchical structure on route planning across multiple region boundaries does not seem to lead to a general preference for routes across less region boundaries. In simulation experiment HN2, both the single-level and the two-level model made the choice between alternative optimal routes at chance-level (figure 3.16). The same figure also shows the choices of human subjects from Wiener and Mallot 2003 that had a far greater tendency (65% in route A, 71% in route B) to choose routes across less region boundaries. As described in section 3.1.2, the crucial decision is made at position 1. Since SVF-based planning strengthens goal-directed weights based on spike-timing (cf. equation 2.7), the decision between walking left to position 10 or right to position 2 depends on which of the respective neurons spikes in shorter temporal proximity to neuron 0. If the activity wave on the lower level of the network (cf. figure 3.11) is fast enough (which was the case here), the hierarchical structure has only a very small influence. This conflict between multiple transition nodes might explain an effect observed in the planning depth simulation (section 3.2) as well: in the center of regions, vector field quality deteriorated (figure 3.22b, 3.22c, 3.22e and 3.22f). In these areas, the agent was drawn to multiple boundary region simultaneously, because all of them received synaptic input from the hierarchically superior neurons.




Figure 4.1: Creation of an ambiguous SVF due to activation of multiple boundary regions (subgoals). The agent is marked as an "X" and region neurons are illustrated on top. In (b) and (c), region neurons are omitted for means of clarity.

Figure 4.1 illustrates this issue: An agent is located in the center of region 1 that overlaps with regions 2 and 3 (figure 4.1a). Upon activation of the region neurons, these boundary regions are activated. This activity wave propagates to the agent (b). The resulting SVF leads the agent to both boundary regions (c). A possible way to reduce this effect would be to add inhibitory connections between transition neurons. Upon activation of the first transition neuron, it could inhibit the other transition neurons on the border between the respective regions, leading to a more selective activation of transition neurons and thus a less ambiguous vector field.

4.2 Increased Planning Depth

The single-level model shows reliable, goal-directed SVFs for navigation across large networks, even when noise currents are applied. The noise current used in section 3.2 was relatively high, drawn from a Gaussian distribution with mean zero and standard deviation of 1 nA, the equivalent of 50% of the synaptic input from a neighboring spiking neuron. As shown in figure 3.19a, 3.19c and 3.19e, the wavefront propagated through the network of 24x24 neurons with very few distortions, and created an SVF of high quality as shown in figure 3.22a. Along the diagonal axis, vector field quality S_N quickly rose to a value of one and stayed there all across the network. The suboptimal value of $S_N = 0.64 \pm 0.02$ of the neuron next to the goal resulted from the neighboring neurons used for the calculation of S_N (cf. equation 3.5). Many of these neurons were located on the outer edge of the network and suffered from the central bias discussed above. Consequently, S_N values decreased in proximity to the goal (3.22). A possible way to prevent this effect could be to add additional neurons on the edges of the network to prevent the goal neuron from being in a corner of the network. An additional negative effect on the S_N value stems from including the goal neuron itself in the computation, since it does not send a movement command to the voting process of neighboring neurons (cf. equation 2.12). This lack of a synaptic vector negatively influences the average vector field quality according to equation 3.5. Despite a lower vector field quality, the lack of movement command might be a desired effect that could lead to slowing down of the agent as soon as the goal neuron becomes active. As in figure 2.3e, the agent stays at the goal without overshooting.

Furthermore, the simulation experiment showed that planning depth can be increased through hierarchical structure if wavefronts are assumed to be decaying. Wave decay was simulated by a global inhibitory current applied to all neurons, thus pulling the membrane potential further and further away from the threshold (cf. section 2.1). Under such circumstances, the single-level network showed a relatively short range of activity propagation (figures 3.20f and 3.22d). Hierarchically structured networks however, achieved to spread activity and with it an SVF throughout the whole network as can be seen in figures 3.20f and 3.21f. This was reflected in the non-zero vector field quality far away from the goal neuron in comparison to the single-level network (figure 3.22) as well. This increased range is in line with the gain of planning depth through hierarchically structured representations (Huys et al. 2015) suspected in section 1.2. The correlational hierarchy (Botvinick 2008) of region neurons allows transmission of activity across a larger number of synapses and with it planning across larger distances.

While not having adversarial effects for planning depth, the additional layer of region neurons leads to partially worse and more fluctuating vector field qualities. In hierarchical networks, a large number of spurious spikes (Ponulak and Hopfield 2013) was observed together with partially worse vector field quality as compared to the single-level network (cf. figure 3.22). This unfavorable effect was probably the result of additional synaptic input, which increased the membrane potential of many neurons. Closer to the threshold (cf. section 2.1), the probability of these neurons to spike increased. Supralinear summation (Ponulak and Hopfield 2013) in combination with lower synaptic weights between levels might reduce the ability of such spurious spikes to release wavefronts and thus mislead the agent.

Navigation with SVFs is not limited in scale. Equipped with enough neurons, the neural network can represent arbitrarily large environments. As established above, a decaying wavefront can be compensated by additional synaptic input, e.g. from a hierarchically superior level. In addition to increasing planning depth, the hierarchical structure speeds up transmission of activity as well. In the simulation experiment described in section 3.2, wavefront propagation was sped up by a factor of approximately four (cf. figures 3.19c and 3.20c). This acceleration resulted from additional synaptic input to all neurons from the activated region neurons in the hierarchically superior level (figure 3.18). Decreased synaptic input by the global inhibitory current in the second part of the "Planning Depth" simulation in turn led to a slightly slower spread of activity. As a result of this increasing inhibition, neurons in this simulation took longer to overcome the threshold for spiking.

4.3 General Discussion

Neural State-Action Networks

The model presented in this thesis is biologically plausible with regard to modeling of the neuronal and synaptic mechanics. It is based on a detailed leaky integrate-and-fire neuron model with experimentally confirmed parameters (listed in appendix A.2). Establishing the SVF is conducted solely with information available at each synapse, i.e. timing of pre- and postsynaptic spikes (Ponulak and Hopfield 2013). SVF-based navigation in networks constructed of topological information and with action voting based on current neuronal activity leads to effective, goal-directed movement, even if multiple goals are present. Goal-directed trajectories of the artificial agent in navigation tasks (e.g. figures 2.3, 3.4 and 3.15) show the capability of this approach. Based on the biological plausibility and the consistently good performance of the non-hierarchical model, it can be used to navigate a simulated three-d environment in future simulations. With the algorithm of Baumann 2019, a graph of features can be constructed. The topology of the graph could then be transferred to a neural network, upon which SVF-based navigation can be performed.

Influence of the width of the activity bump representing the current position of the agent is worthy of further investigation as well. Since active neurons determine the executed movements of the agent according to equation 2.12, it matters how many are active at a time. In simulation experiments HN1 and HN2, the activity bump was assumed to occupy only one neuron to avoid erroneous signals from neurons representing neighboring places. In larger networks, possibly with neurons representing places along the grid edges connecting the places in HN1 (cf. figure 3.1), such a strict limit is unnecessary. Figure 2.3 shows a trajectory with a wider bump of activity, encompassing approximately seven neurons. A very rapidly moving agent can distort the initially circular bump due to the exponential decay of neuronal activity behind the agent. The activity bump then forms a comet tail behind the agent leading to an effect of inertia, since these neurons still send their movement commands to the voting process. This effect does not occur if the timescale at which the activity bump moves through the neural network is much slower than the neuronal timescale of milliseconds (Nemenman et al. 2008), which can generally be assumed to be the case.

A disadvantage of the current model is the inability to represent trajectories and thus choose between alternatives in advance. The fine-to-coarse planning algorithm described in the introduction does not address such representations either (Wiener and Mallot 2003). In another approach to hierarchical planning presented in section 1.2, namely HRL, choosing between midterm strategies is possible through the availability of *options* Sutton et al. 1999.

Initially, representing hierarchical structure was suspected to increase planning depth as put forth by Huys et al. 2015. As discussed above, the simulation experiments partially confirmed this assumption based on the implemented neuronal hierarchy. Largely, this positive effect was due to greater robustness against decay of the wave, a result of additional synaptic input from region neurons. For this benefit it is rather unimportant which places are grouped together, as long as they are in close spatial proximity. Proximity is important, because goal-directed vector fields depend on activation of neurons in succession, the further away they are from the goal, the later they must be activated as the wavefront passes. As long as the topology of region neurons is in compliance with the places they represent, the exact borders of the regions should not matter for this increase in planning depth. Possible pitfalls of the implemented hierarchy have been discussed above, together with ideas for improvement. In an extended model, release of wavefronts by strong input to specific neurons could be used to form a neuronal model of decision points. These have been suspected to play a large role in navigation and constitute possible subgoals (Lynch 1960, Spiers and Maguire 2008 Stachenfeld et al. 2017). Another hierarchical structure in spatial cognition was described by Kuipers 2000: Taxi drivers prefer to stay on a system of frequently used routes as long as possible when navigating across the city and leave it only for the last leg of the route. Learning such a hierarchy of routes can be implemented in the present model through the STDP process: By traveling along a certain route multiple times in the exploration phase, the weights there are strengthened according to equation 2.4. In the wavefront propagation (planning) phase, the wavefront can be expected to spread faster along this route of increased synaptic weights. Rapid spread of activity in the planning phase along a

certain route will later lead the agent preferably along that route, thus replicating the human preference for these routes as observed by Kuipers 2000.

Implementation of Hierarchy

The introduced region neurons are connected to lower-level neurons representing places within this particular region as well as to places from which this region can be immediately reached. Reasons for grouping these places together may include shared functionality, visual similarity or designation with a region name for historical reasons (Schick et al. 2019). All three of these reasons lead to a correlational hierarchy (Botvinick 2008) built upon statistical co-occurrence of the places in the context of the respective functionality, with distinctive visual stimuli or with the common, arbitrary name. Further simulations could be conducted to investigate if correlations between the firing of feature-sensitive neurons can be used to learn the hierarchical structure of the network. The same STDP-based adaption of synaptic weights described in 2.2 could be used to strengthen connections between neurons active in short temporal distance Gerstner et al. 2014, yielding Hebbian learning (Hebb 1949). By doing so, formations of neurons could be tied together through a region node without sensory input (cf. section 2.4) with the purpose of representing the region and with connections to other such region neurons.

As discussed above, the effects of the hierarchical structure implemented in this thesis can be summarized in three points:

1. Increased activation of boundary regions leads to the release of wavefronts from there. These regions then become preferred subgoals for the agent.
2. By adding a second layer of neurons, the strict topology of the single-level network is altered. Thus, wavefront propagation during planning can be faster along suboptimal routes than along shorter alternatives. This can lead to partially suboptimal trajectories.
3. Region neurons provide additional synaptic input to all lower-level neurons, leading to both faster transmission of activity wavefronts in general and activity transmission across larger distances if wavefronts are assumed to be decaying. However, the same additional input leads to increased sensitivity to noise.

Fine-to-Coarse Planning and Hierarchical RL

Agents based on the fine-to-coarse planning heuristic by Wiener and Mallot 2003 replicate human behavior in regionalized environments better than the neuronal hierarchy implemented in this thesis. Nevertheless, the developed neuronal model gives a neurobiological theory for hierarchical navigation and may in the future give an explanation for how regions can be learned through STDP updating of synaptic weights. Implemented as algorithm by Reineking et al. 2008, fine-to-coarse planning was described in section 1.2 as a heuristic to plan in hierarchically structured representations. The topology of the state-action graph used by Wiener and Mallot

is very similar to the neuronal state-action network described in chapter 2: on top of the low-level representation based on discrete entities (whether they are image features or places) lays an abstracted representation of regions. Interlevel connection signal associations of places with regions, while also connecting region nodes to places from which they can be reached. The planning process which has to be repeated after every move does indeed replicate the human preference to reach the goal region as quickly as possible. By solving the planning problem on the superordinate level, the next region to be navigated is selected (*ibid.*). Thus, the preference of human subjects to cross less region boundaries can be explained with the fine-to-coarse heuristic: solving the planning problem first on the highest possible level (Reineking et al. 2008) discards routes across more regions.

The proposed model finds suitable subgoals which has been a challenge for hierarchical reinforcement learning. HRL has been introduced as a second approach to hierarchical planning in section 1.2. It is a promising way to use hierarchical structure for reinforcement learning of complex tasks (Sutton et al. 1999, Botvinick 2008, Botvinick and Weinstein 2014) with increasing neurobiological plausibility (Botvinick et al. 2009). In the framework of reinforcement learning, topological S-A-S' models correspond to the model-based approach (Botvinick and Weinstein 2014) as opposed to the pure updating of value-functions as in traditional Q-learning (Sutton and Barto 1998). In the traditional options framework of HRL (Sutton et al. 1999, Botvinick 2008), subgoals are predefined. This is due to the difficulty of learning meaningful subgoals in complex problem space (Botvinick and Weinstein 2014). The model proposed here provides selection of subgoals in the sense of neurons with a high amount of synaptic inputs from higher levels. Here, these neurons are located in boundary regions where regions overlap. Through the interplay of SVF-based navigation and the implemented neuronal hierarchy, these areas become navigational subgoals.

4.4 Suggestions for Further Research

As discussed above, the overlap of regions in the model can lead to successful selection of subgoals. To investigate if such an overlap indeed is a factor in human navigation, a behavioral experiment could be set up. Subjects explore a virtual environment of regions separated by visual barriers to prevent association of features to multiple regions. Route choices in goal-directed navigation could then be compared to the same setup of regions (albeit maybe marked by different indicators to avoid learning of routes) without the visual barriers and thus allowing for overlap of regions.

Additional Simulations of Interest

Thinking about daily life, sometimes navigational goals are not actual places in the environment, but rather abstract regions (such as when planning a trip to Berlin). This could be replicated in a simulation experiment without altering the model by activating a region neuron alone as a goal. Region neurons themselves have no counterpart in the environment, but represent an abstraction formed by the agent.

Thus, model behavior when navigating to such abstracted entities could be tested.

The association of actions with synaptic connections was described in section 2.3. An extension of the model could investigate learning of these motor commands. Schölkopf and Mallot 1995 built a neuronal, topological model for navigation, in which edges representing state-transitions possess learnable connections to neurons representing actions. By extending the model presented in this thesis with such a mechanism for learning these connections, a further step towards biological plausibility could be made. Learning could be implemented by using the same STDP updating from equation 2.4. As the agent explores the environment, connections between the synapses encoding state-transitions and the motor neurons which are active during the transition would be strengthened. The current model also does not have actions associated to the connections between region neurons. By associating these with an action as well, a task inspired by the study of Stevens and Coupe 1978 could be posed to the model. Activity in the currently active region node would then participate in the voting of actions and thus provide a preference to follow remembered large-scale directions. Replicating the study of Stevens and Coupe, the task posed to the model would be to navigate from San Diego in California to Reno (Nevada). The action associated with the connection from California to Nevada would then be to "go east" and guide the agent there, even if the correct choice would be to navigate westwards, replicating the behavior of human subjects.

Application in Neuromorphic Computing

Spiking neural networks implemented in hardware, such as the Intel Loihi (Davies et al. 2018) or the BrainChip Akida (Posey 2020), promise energy-efficient parallel computation and on-chip learning. In addition to applications in image classification as proposed by BrainChip (Posey 2020), neuromorphic hardware with on-chip learning could also be used to solve navigational tasks. State-action networks with SVFs created by wavefronts of spiking activity could be implemented in future versions of such microchips. With an action-execution system as described in section 2.3, goal-directed navigation becomes possible. Small autonomous vehicles could profit from such small, low-power chips. Especially for airborne devices, power consumption is critical. By generalizing state-action space to problem space, the range of applications is increased beyond the spatial domain. If transmission of spiking activity in such systems is shown to decay, the proposed hierarchical structure provides a possible solution. By increasing the range across which a vector field can be induced, the navigable area is then increased.

Chapter 5

Conclusion

This master thesis presents a computational model for route planning and navigation in hierarchically structured environments. Aiming to be biologically plausible, the mechanics of the model are based on neurobiological findings. As representational apparatus from which later function is produced, a neuronal network represents recognized features. Synaptic connections in this low-level network correspond to transition actions between the respective features in the environment. A second network layer captures regional structures such as places surrounded by water. Route planning is achieved by spreading a wave of neuronal activity from the cells representing the goal throughout the network. By using spiking neurons, it becomes possible to update synaptic weights based on spike-timing. In combination with the spreading wavefront, this synaptic plasticity leads to a temporary strengthening of synapses corresponding to goal-directed actions. This synaptic vector field (SVF) is then used by the agent to navigate to the goal. Movement execution is based on a voting process of synaptic action commands. Each vote is weighted by the corresponding synaptic weight and the activity of the presynaptic neuron. The navigational ability of the model was confirmed in a series of experiments, making further research into SVF-based navigation in state-action representations promising.

Hierarchical structure is assumed to be a cornerstone of planning complex behavior. Evaluation of the implemented neuronal hierarchy however led to ambiguous results in terms of navigational efficiency. While helping to carry the vector field across more synapses if the wavefront was assumed to be decaying, the hierarchical structure also appears to increase the sensitivity of the network against noise. Nevertheless, comparison of model behavior with existing studies on human navigation showed promising results, strengthening biological plausibility of the model. Proposals to overcome malfunctions of the two-level network were made and together with increasing insight the neuronal architecture of the brain could shine more light on the relevance of such neuronal hierarchies for shaping cognition.

Of relevance for the field of reinforcement learning is the promising subgoal selection resulting from the neuronal hierarchy in combination with SVF-based planning. Autonomous subgoal selection is an important problem of hierarchical reinforcement learning. Further simulations could test if the second layer built in by hand here can be learned instead with the same plasticity rules that are used to create the vector field.

Appendix A

Supplementary Information

A.1 Inhibitory System

The inhibitory system based on Ponulak and Hopfield 2013 is switched off in the wavefront propagation phase (except in the planning depth simulation experiment, cf. section 3.2). Its purpose is to ensure the attractor-like behavior of the network during navigation (Ponulak and Hopfield 2013), i.e. a stable bump of activity representing the current position of the agent (ibid.). The global inhibitory current from the inhibitory population is given as in Ponulak and Hopfield 2013:

$$i_{inh}(t) = a_{inh}A_{inh}(t) \quad (\text{A.1})$$

with the gating variable a_{inh} set to one during navigation and to zero in wavefront propagation. This switching is assumed to be governed by globally released neurotransmitter (Ponulak and Hopfield 2013). The activity of inhibitory system is given by

$$A_{inh} = \begin{cases} \alpha_{inh}(i_e(t) - I_{e,0}) & \text{if } i_e(t) - I_{e,0} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.2})$$

with the threshold $I_{e,0}$ above which the inhibitory system becomes active (Ponulak and Hopfield 2013). The input $i_e(t)$ to the inhibitory system is updated based on the following differential equation depending on the total activity in the network (ibid.):

$$\tau_e \frac{di_e(t)}{dt} = -i_e(t) + a_e \cdot \Sigma_{spike}(t) \quad (\text{A.3})$$

where $\Sigma_{spike}(t)$ (pronounced "sigma-spike") counts the number of neuronal spikes over a time window of the last 0.2 ms.

A.2 Network Parameters

Fixed parameters for all simulations

To remain biologically plausible, the parameters were chosen as experimentally measured, if possible. Others were set to optimize network behavior.

Param.	Value	Meaning	Source
a_{Dec}	0 ms^{-1}	Synaptic weight decay	[1]
a_e	20 nA	Increment to inh. system	-
a_{SLS}	10	Factor for supralinear summation	Ponulak and Hopfield 2013
A^+	0.5 ms^{-1}	STDP parameter	Section 2.2
A^-	1 ms^{-1}	Reverse STDP parameter	[2]
α_{inh}	20	Modulation of i_{inh}	-
b_{SLS}	0.1	Factor for supralinear summation	[3]
h	0.2 ms	Timestep size	Ponulak and Hopfield 2013
i_{spike}	0.3 nA	Presynaptic current peak	Gerstner 2014, Ch. 3.1.2
$I_{e,0}$	5 nA	Threshold for inhibition	-
R_m	20 M Ω	Membrane resistance	Ponulak and Hopfield 2013
T	0.2 s	Motonic integration time	-
τ_e	25 ms	Decay of input to inh. system	-
τ_m	20 ms	Decay time of membrane potential	Ponulak and Hopfield 2013
τ_{spike}	25 ms	Decay time of synaptic current	Ponulak and Hopfield 2013
τ_{STDP}	20 ms	Decay time of STDP strength	Morrison et al. 2008
u_r	0 mV	Membrane resting potential	Ponulak and Hopfield 2013
u_{thresh}	10 mV	Threshold for action potential	Ponulak and Hopfield 2013

Table A.1: List of the parameters used in chapter 3

- [1] a_{Dec} is set to zero during wavefront propagation to prevent decay of unused weights.
- [2] To correctly adjust the strength of the STDP process as based on equations 2.5-2.7, it makes sense to regard an exemplary case. Neuron 2 spikes $s = 4$ ms after neuron 1, hence the weight from 2 to 1 (reverse STDP) changes as follows:

$$\Delta w_{21} = h \cdot A^- \cdot \exp\left\{-\frac{4}{20}\right\} = 0.2 \text{ ms} \cdot A^- \cdot 0.82$$

For a realistic temporary change of 20%, A^- can be set accordingly:

$$\Delta w_{12} \stackrel{!}{=} 0.2 \Rightarrow A^- \approx 1 \text{ ms}^{-1} \quad (\text{A.4})$$

- [3] b_{sls} was altered from 0.05 (Ponulak and Hopfield 2013) to 0.1 to profit from supralinear effect (cf. section 2.1) already with two inputs. Comparing the factor $a_{sls} \cdot \tanh b_{sls} \cdot \sum$ from equation (2.2) shows that $10 \cdot \tanh(0.1 \cdot 2) \approx 2$ while $10 \cdot \tanh(0.05 \cdot 2) \approx 1$ yields the same as linear input summation.

Simulation Specific Parameters

Depending on the simulation experiment, different settings were used for the simulation, listed below in table A.2.

Sim. (#Lvls.)	Chapter	SLS	w_1	w_2	w_{inter}	w_r	δ_{inh}
Demo (1)	2.3	yes	3	1	-	-	-
HN1.small (1)	3.1.1	no	6	1	-	-	-
HN1.small (2)	3.1.1	no	6	1	10	3	-
HN1.large (1)	3.1.1	no	6	1	-	-	-
HN1.large (2)	3.1.1	no	6	1	10	3	-
HN2 (1)	3.1.2	no	6	-	-	-	-
HN2 (2)	3.1.2	no	6	-	3	3	-
PD (1)	3.2	no	6	1	-	-	-
PD (2)	3.2	no	6	1	5	3	-
PD-dec (1)	3.2	no	6	1	-	-	0.015 ms^{-1}
PD-dec (2)	3.2	no	6	1	5	3	0.015 ms^{-1}

Table A.2: Parameters and settings used in each simulation. Simulation name (PD = planning depth experiment, PD-dec = wave decay), in brackets: number of hierarchical levels. SLS: yes = supralinear summation as described in section 2.1, no = linear summation of synaptic inputs. w_1 and w_2 : weights between neurons corresponding to places one and two edges away in the environment, w_{inter} : weights between hierarchical levels, w_r weights between higher-level neurons. δ_{inh} increase of inhibitory current in the wavefront propagation phase as described in section 3.2.

A.3 Low-level Neurons as Place Cells

For the demonstration of navigation in chapter 2, sensory input to each cell is implemented by the Euclidean distance of the agent to the preferred location of the cell (Ponulak and Hopfield 2013). The firing field of the cell corresponds to a Gaussian centered at the preferred location $\mathbf{x}_k = (x_k, y_k)$:

$$i_{sens,k} = i_{sens,max} \cdot \exp\left\{-\left(\frac{(x_k - x_{ag})^2 + (y_k - y_{ag})^2}{2\sigma_{sens}^2}\right)\right\} \quad (\text{A.5})$$

where $i_{sens,max}$ is the maximum value of i_{sens} , σ_{sens} determines the size of the cells place field and (x_{ag}, y_{ag}) is the agent's current position. This approach is frequently used in models of position-sensitive cells (Bicanski and Burgess 2018) and can be explained neurobiologically by input of path-integrating grid-cells (Bicanski and Burgess 2018, Li et al. 2020).

A.4 Comparison to Chance Level in HN1

To interpret the results of simulation HN1 better, it makes sense to compare behavior of the artificial agent to an agent choosing goal-directed routes at chance level. We consider an agent taking only the shortest possible routes (called solutions below). This is equivalent to evaluating only the trials in which a solution was achieved as did Wiener and Mallot 2003. To calculate the chance level for quickest possible region transition, the fraction of solutions allowing for fastest region transition over the number of total solutions is calculated for each route type from table 3.1:

- (A) $4 \rightarrow 8$, the only solution leading to the goal region as quickly as possible is 4-7-8, the only alternative solution is 4-5-8, leading to a chance level of 50%.
- (B) $5 \rightarrow 9$, solutions with quickest region transition: 5-8-7-10-9, 5-8-11-10-9 and 5-8-7-6-9, other solutions are 5-4-7-10-9, 5-4-3-6-9 and 5-4-7-10-9 leading to a chance level of 50%.
- (C) $2 \rightarrow 6$, solutions with quickest region transition: 2-5-8-7-6, other solutions: 2-5-4-7-6, 2-5-4-3-6, 2-1-4-7-6 and 2-1-0-3-6, leading to a chance level of 33.3%.
- (D) $5 \rightarrow 6$, solutions with quickest region transition: 5-8-7-6, other solutions: 5-4-7-6 and 5-4-3-6, leading to a chance level of 33.3%.
- (E) $4 \rightarrow 11$, solutions with quickest region transition: 4-7-10-11 and 4-7-8-11, other solution: 4-5-8-11, leading to a chance level of 66.7%.

A.5 Code Library References

A.5.1 LEMON

The spiking neural network was implemented in C++ as a graph in the LEMON graph library.

LEMON Graph Library 1.3.1, (Dezso et al. 2011),
<https://lemon.cs.elte.hu/trac/lemon>

A.5.2 OpenCV

OpenCV was used to create the visual output of the model from the C++ script, i.e. the images of spiking neurons and agent navigation.

OpenCV 3.4.0, 2014 (Bradski and Kaehler 2000),
<https://opencv.org/>

Bibliography

- Amari, S. (1971). “Characteristics of randomly connected threshold-element networks and network systems”. In: *Proceedings of the IEEE* 59.1, pp. 35–47. DOI: 10.1109/PROC.1971.8087.
- Arbib, M. A. and Bonaiuto, J. J. (2012). “Multiple levels of spatial organization: World Graphs and spatial difference learning”. In: *Adaptive Behavior* 20.4, pp. 287–303. DOI: 10.1177/1059712312449545.
- Arbib, M. A. and Liebllich, I. (1977). “Motivational learning of spatial behavior”. In: *Systems Neuroscience*. Ed. by J. Metzler. New York: Academic Press, pp. 221–239.
- Baumann, T. (2019). “Microsnapshot navigation”. Master thesis at the University of Tübingen.
- Bay, H. and Tuytelaars T. and Van Gool, L. (2006). “SURF: Speeded up robust features”. In: *European conference on computer vision*. Springer. Berlin, Heidelberg, pp. 404–417. DOI: 10.1007/11744023_32.
- Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., et al. (2018). “What is a cognitive map? Organizing knowledge for flexible behavior”. In: *Neuron* 100.2, pp. 490–509. DOI: 10.1016/j.neuron.2018.10.002.
- Bicanski, A. and Burgess, N. (2018). “A neural-level model of spatial memory and imagery”. In: *eLife* 7. DOI: 10.7554/eLife.33752.
- Botvinick, M. M. (2008). “Hierarchical models of behavior and prefrontal function”. In: *Trends Cogn. Sci.* 12, pp. 201–208. DOI: 10.1016/j.tics.2008.02.009.
- Botvinick, M. M., Niv, Y., and Barto, A. G. (2009). “Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective”. In: *Cognition* 113.3, pp. 262–280. DOI: doi.org/10.1016/j.cognition.2008.08.011.
- Botvinick, M. M. and Weinstein, A. (2014). “Model-based hierarchical reinforcement learning and human action control”. In: *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 369.1655, p. 20130480. DOI: 10.1098/rstb.2013.0480.
- Bradski, G. and Kaehler, A. (2000). “The OpenCV library”. In: *Dr. Dobb’s Journal of Software Tools* 3.
- Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, Massachusetts: The MIT Press.

- Buckner, R. L. and Carroll, D. C. (2007). “Self-projection and the brain”. In: *Trends in cognitive sciences* 11.2, pp. 49–57. DOI: 10.1016/j.tics.2006.11.004.
- Burgess, N. and O’Keefe, J. (1996). “Neuronal computations underlying the firing of place cells and their role in navigation”. In: *Hippocampus* 6.6, pp. 749–762. DOI: 10.1002/(SICI)1098-1063(1996)6:6<749::AID-HIP016>3.0.CO;2-0.
- Burgess, N. and Recce M. and O’Keefe, J. (1994). “A model of hippocampal function”. In: *Neural networks* 7.6-7, pp. 1065–1081. DOI: 10.1016/S0893-6080(05)80159-5.
- Cartwright, B. A. and Collett, T. S. (1987). “Landmark maps for honeybees”. In: *Biol. Cybern.* 57, pp. 85–93. DOI: 10.1007/BF00318718.
- Davies, M., Srinivasa, Na., Lin, T., et al. (2018). “Loihi: A neuromorphic manycore processor with on-chip learning”. In: *IEEE Micro* 38.1, pp. 82–99. DOI: 10.1109/MM.2018.112130359.
- Dezso, B., Jüttner, A., and Kovacs, P. (2011). “LEMON—an open source C++ graph template library”. In: *Electronic Notes in Theoretical Computer Science* 264.5, pp. 23–45. DOI: 10.1016/j.entcs.2011.06.003.
- Ellender, T. J., Nissen, W., Colgin, I. L., Mann, E. O., and Paulsen, O. (2010). “Priming of hippocampal population bursts by individual perisomatic-targeting interneurons”. In: *J. Neurosci.* 30, pp. 5979–5991. DOI: 10.1523/JNEUROSCI.3962-09.2010.
- Gerstner, W., Kistler, W. M., Naud, R., and Paninski, L. (2014). *Neuronal dynamics. From single neurons to networks and cognition*. Cambridge: Cambridge University Press. DOI: 10.1017/CB09781107447615.
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York: Wiley. DOI: 10.1002/cne.900930310.
- Hirtle, S. C. and Jonides, J. (1985). “Evidence of hierarchies in cognitive maps”. In: *Memory and Cognition* 13, pp. 208–217. DOI: 10.3758/BF03197683.
- Hopfield, J. J. (1982). “Neural networks and physical systems with emergent collective computational abilities”. In: *Proceedings of the national academy of sciences* 79.8, pp. 2554–2558. DOI: 10.1073/pnas.79.8.2554.
- Hopfield, J. J. (2010). “Neurodynamics of mental exploration”. In: *Proc. Natl. Acad. Sci. U.S.A.* 107, pp. 1648–1653. DOI: 10.1073/pnas.0913991107.
- Huys, Q. J. M., Lally, N., Faulkner, P., et al. (2015). “Interplay of approximate planning strategies”. In: *Proc. Natl. Acad. Sci. USA* 112, pp. 3098–3103. DOI: 10.1073/pnas.1414219112.

- Kaelbling, L. P. and Lozano-Pérez, T. (2011). “Hierarchical task and motion planning in the now”. In: *2011 IEEE International Conference on Robotics and Automation*. Shanghai, pp. 1470–1477. DOI: 10.1109/ICRA.2011.5980391.
- Kempster, R., Gerstner, W., and van Hemmen, J. L. (1999). “Hebbian learning and spiking neurons”. In: *Phys. Rev. E* 59, pp. 4498–4514. DOI: 10.1162/neco.2007.19.6.1437.
- Krotov, D. and Hopfield, J. J. (2019). “Unsupervised learning by competing hidden units”. In: *Proceedings of the National Academy of Sciences* 116.16, pp. 7723–7731. DOI: 10.1073/pnas.1820458116.
- Kuipers, B. (2000). “The spatial semantic hierarchy”. In: *Artificial Intelligence* 119, pp. 191–233. DOI: 10.1016/S0004-3702(00)00017-5.
- Kuipers, B., Tecuci, D., and Stankiewicz, B. (2003). “The skeleton in the cognitive map: A computational and empirical exploration”. In: *Environment and Behavior* 35, pp. 81–106. DOI: 10.1177/0013916502238866.
- Li, T., Arleo, A., and Sheynikhovich, D. (2020). “Modeling place cells and grid cells in multi-compartment environments: Entorhinal–hippocampal loop as a multisensory integration circuit”. In: *Neural Networks* 121, pp. 37–51. DOI: 10.1016/j.neunet.2019.09.002.
- Lubenov, E. V. and Siapas, A. G. (2008). “Decoupling through synchrony in neuronal circuits with propagation delay”. In: *Neuron* 58, pp. 118–131. DOI: 10.1016/j.neuron.2008.01.036.
- Lubenov, E. V. and Siapas, A. G. (2009). “Hippocampal theta oscillations are travelling waves”. In: *Nature* 459, pp. 534–539. DOI: 10.1038/nature08010.
- Lynch, Kevin (1960). *The image of the city*. Vol. 11. Cambridge, Massachusetts: The MIT Press.
- Mallot, H. A. and Basten, K. (2009). “Embodied spatial cognition: Biological and artificial systems”. In: *Image and Vision Computing* 27, pp. 1658–1670. DOI: 10.1016/j.imavis.2008.09.001.
- Marr, D. ([1982] 2010). *Vision : A computational investigation into the human representation and processing of visual information*. Cambridge, Massachusetts: The MIT Press.
- Martinet, L.-E., Sheynikhovich, D., Benchenane, K., and Arleo, A. (2011). “Spatial learning and action planning in a prefrontal cortical network model”. In: *PLoS Comput. Biol.* 7. DOI: 10.1371/journal.pcbi.1002045.

- McNamara, T. P., Hardy, J., and Hirtle, S. C. (1989). “Subjective hierarchies in spatial memory”. In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 15, pp. 211–227. DOI: 10.1037/0278-7393.15.2.211.
- Müller, M. and Wehner, R. (1988). “Path integration in desert ants, *Cataglyphis fortis*”. In: *Proceedings of the National Academy of Sciences* 85.14, pp. 5287–5290. DOI: 10.1073/pnas.85.14.5287.
- Morrison, A., Aertsen, A., Diesmann, M., and Gerstner, W. (2007). “Spike-timing dependent plasticity in balanced random networks”. In: *Neural Comput.* 19, pp. 1437–1467. DOI: 10.1162/neco.2007.19.6.1437.
- Morrison, A., Diesmann, M., and Gerstner, W. (2008). “Phenomenological models of synaptic plasticity based on spike timing”. In: *Biol. Cybern.* 98, pp. 459–478. DOI: 10.1007/s00422-008-0233-1.
- Muller, R. U., Kubie, J. L., and Ranck, J. B. (1987). “Spatial firing patterns of hippocampal complex-spike cells in a fixed environment”. In: *Journal of Neuroscience* 7.7, pp. 1935–1950.
- Nemenman, I., Lewen, G. D., Bialek, W., and de Ruyter van Steveninck, R. R. (2008). “Neural coding of natural stimuli: information at sub-millisecond resolution”. In: *PLoS Comput. Biol.* 4.3. DOI: <https://doi.org/10.1371/journal.pcbi.1000025>.
- Nettleton, J. S. and Spain, W. J. (2000). “Linear to supralinear summation of ampa-mediated EPSPs in neocortical pyramidal neurons”. In: *J. Neurophys.* 83, pp. 3310–3322. DOI: 10.1152/jn.2000.83.6.3310.
- O’Keefe, J. and Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Oxford University Press.
- O’Keefe, J. and Recce, M. L. (1993). “Phase relationship between hippocampal place units and the EEG theta rhythm”. In: *Hippocampus* 3.3, pp. 317–330. DOI: 10.1002/hipo.450030307.
- O’Reilly, R. C. and Frank, M. J. (2006). “Making working memory work: a computational model of learning in prefrontal cortex and basal ganglia”. In: *Neural Comput.* 18, pp. 283–328. DOI: 10.1162/089976606775093909.
- Patel, J., Fujisawa, S., Berenyi, A., Royer, S., and Buzsaki, G. (2012). “Traveling theta waves along the entire septotemporal axis of the hippocampus”. In: *Neuron* 75. DOI: 10.1016/j.neuron.2012.07.015.
- Pfeiffer, B. E. and Foster, D. J. (2013). “Hippocampal place-cell sequences depict future paths to remembered goals”. In: *Nature* 497.7447, pp. 74–79. DOI: 10.1038/nature12112.

- Ponulak, F. and Hopfield, J. J. (2013). “Rapid, parallel path planning by propagating wavefronts of spiking neural activity”. In: *Front. Comput. Neurosci.* 7.98. DOI: 10.3389/fncom.2013.00098.
- Posey, B. M. (2020). *What is the Akida event domain neural processor?* URL: <https://brainchipinc.com/what-is-the-akida-event-domain-neural-processor> (visited on 08/22/2020).
- Reineking, T., Kohlhagen, C., and Zetsche, C. (2008). “Efficient wayfinding in hierarchically regionalized spatial environments”. In: *International Conference on Spatial Cognition*. Springer. Berlin, Heidelberg, pp. 56–70. DOI: 10.1007/978-3-540-87601-4_7.
- Roberts, P. D. and Leen, T. K. (2004). “Anti-hebbian spike-timing-dependent plasticity and adaptive sensory processing”. In: *Front. Comput. Neurosci.* 4.156. DOI: 10.3389/fncom.2010.00156.
- Rumelhart, D. E. and Zipser, D. (1985). “Feature discovery by competitive learning”. In: *Cognitive Science* 9, pp. 75–112. DOI: 10.1207/s15516709cog0901_5.
- Samsonovich, A. and McNaughton, B. L. (1997). “Path integration and cognitive mapping in a continuous attractor neural network model”. In: *Journal of Neuroscience* 17.15, pp. 5900–5920. DOI: 10.1523/JNEUROSCI.17-15-05900.1997.
- Schick, W., Halfmann, M., Hardiess, G., Hamm, F., and Mallot, H. A. (2019). “Language cues in the formation of hierarchical representations of space”. In: *Spatial Cognition & Computation* 19.3, pp. 252–281. DOI: 10.1080/13875868.2019.1576692.
- Schölkopf, B. and Mallot, H. A. (1995). “View-based cognitive mapping and path planning”. In: *Adaptive Behavior* 3.3, pp. 311–348. DOI: 10.1177/105971239500300303.
- Seol, G. H., Ziburkus, J., Huang, S., et al. (2007). “Neuromodulators control the polarity of spike-timing-dependent synaptic plasticity”. In: *Neuron* 55.6, pp. 919–929. DOI: 10.1016/j.neuron.2007.08.013.
- Spiers, H. J. and Maguire, E. A. (2008). “The dynamic nature of cognition during wayfinding”. In: *Journal of environmental psychology* 28.3, pp. 232–249. DOI: 10.1016/j.jenvp.2008.02.006.
- Stachenfeld, K. L., Botvinick, M. A., and Gershman, S. J. (2017). “The hippocampus as a predictive map”. In: *Nature Neuroscience* 20.11, p. 1643. DOI: 10.1038/nn.4650.
- Steinmetz, P. N., Cabrales, E., Wilson, M. S., et al. (2011). “Neurons in the human hippocampus and amygdala respond to both low- and high-level image properties”. In: *J. Neurophysiol.* 105, pp. 2874–2884. DOI: 10.1152/jn.00977.2010.

- Stevens, A. and Coupe, P. (1978). “Distortions in judged spatial relations”. In: *Cognitive Psychology* 10.4, pp. 422–437. DOI: 10.1016/0010-0285(78)90006-3.
- Sutton, R. S. and Barto, A. G. (1998). *Introduction to reinforcement learning*. Vol. 135. Cambridge, Massachusetts: The MIT Press. DOI: 10.1016/S1364-6613(99)01331-5.
- Sutton, R. S., Precup, D., and Singh, S. (1999). “Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning”. In: *Artificial intelligence* 112.1-2, pp. 181–211. DOI: 10.1016/S0004-3702(99)00052-1.
- Tolman, E. C. (1948). “Cognitive maps in rats and man”. In: *Psychological Review* 55.4, pp. 189–208. DOI: 10.1037/h0061626.
- Urakubo, H., Aihara, T., Kuroda, S., Watanabe, M., and Kondo, S. (2004). “Spatial localization of synapses required for supralinear summation of action potentials and EPSPs”. In: *J. Comput. Neurosci.* 56, pp. 251–265. DOI: 10.1023/B:JCNS.0000025688.64836.df.
- Voicu, H. and Schmajuk, N. (2002). “Latent learning, shortcuts and detours: a computational model”. In: *Behav. Process.* 59, pp. 67–86. DOI: 10.1016/S0376-6357(02)00060-8.
- Warren, W. H. (2019). “Non-Euclidean navigation”. In: *Journal of Experimental Biology* 222.Suppl 1. DOI: 10.1242/jeb.187971.
- Wiener, J. M. and Mallot, H. A. (2003). “‘Fine-to-coarse’ route planning and navigation in regionalized environments”. In: *Spat. Cogn. Comput.* 3, pp. 331–358. DOI: 10.1207/s15427633scc0304_5.
- Wilson, H. R. and Cowan, J. D. (1973). “A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue”. In: *Kybernetik* 13.2, pp. 55–80. DOI: 10.1007/BF00288786.
- Wilson, M. A. and McNaughton, B. L. (1993). “Dynamics of the hippocampal ensemble code for space”. In: *Science* 261.5124, pp. 1055–1058. DOI: 10.1126/science.8351520.