TUM

TECHNISCHE UNIVERSITÄT MÜNCHEN
INSTITUT FÜR INFORMATIK

Towards Cognitive Navigation: A Biologically Inspired Calibration Mechanism for the Head Direction Cell Network

Dominik Nitschke, Zhenshan Bing, Alois Knoll

TUM-I2291

DEPARTMENT OF INFORMATICS

TECHNISCHE UNIVERSITÄT MÜNCHEN

Master's Thesis in Informatics: Neuroscience

Towards Cognitive Navigation: A Biologically Inspired Calibration Mechanism for the Head Direction Cell Network

Dominik Nitschke

DEPARTMENT OF INFORMATICS

TECHNISCHE UNIVERSITÄT MÜNCHEN

Master's Thesis in Informatics: Neuroscience

Towards Cognitive Navigation: A Biologically Inspired Calibration Mechanism for the Head Direction Cell Network

Auf dem Weg zur kognitiven Navigation: Ein biologisch inspirierter Kalibrierungsmechanismus für das Head Direction Cell Network

Author:	Dominik Nitschke
Supervisor:	Prof. Dr.-Ing. habil. Alois C. Knoll
Advisor:	Dr. rer. nat. Zhenshan Bing
Submission Date:	September 17th, 2021

I confirm that this master's thesis in informatics: neuroscience is my own work and that I have documented all sources and material used.

Munich, September 17th, 2021 Dominik Nitschke

Abstract

To derive meaningful navigation strategies, for example, to find food, animals have to estimate their position and orientation in the framework of their environment. Accordingly, so-called head direction cells (HD cells) were found in rat brains whose activity encodes the animal's orientation. The integration of self-motion cues, often referred to as path integration (PI), is critical for the generation of the HD cells' activity, but introduces errors into the HD signal in the long term. Therefore, animals can exploit environmental features, like visual landmarks, to regularly correct for HD errors. Nevertheless, it is not fully understood how the HD system of an animal is calibrated using the egocentric perception of visual landmarks without the introduction of parallax errors. Researchers tried to explain this by designing computational HD models that mimic this calibration behavior. State-of-the-art place encoding feedback models associate HD with the egocentric perceived bearing of a visual cue, called egocentric cue direction (ECD), and the position of the agent that observes the cue. Therefore, these models can only restore HD information when a visual landmark is perceived a second time in the same position from the same angle.

This work presents a novel efficient computational model designed in a biologically plausible way based on spiking neurons. It addresses the problem of calibrating HD from egocentrically perceived visual cues by associating the angular shift between the ECD and the HD with positional information. This angular shift is described by the so-called allocentric cue direction (ACD). The continuous computations to either infer ACD from ECD and HD, during exploration, or to infer HD from ECD and restored ACD, for calibration, are facilitated using a neural circuit based two two-dimensional conjunctive cell fields. Additionally, the model offers two alternative calibration modes - first glance learning (FGL) switched off or on. With FGL off, the model can derive HD from a visual landmark when perceiving it a second time from the same position but independent from its viewing angle. With FGL on, the model can even derive HD when perceiving a visual landmark a second time from any position and orientation. As both modes do not require the model to perceive the landmark a second time from the same orientation, the model's efficiency to use an egocentrically perceived visual landmark to infer HD is significantly enhanced compared to previously proposed models that require a match of orientation and position. The results demonstrate that it is sufficient to only associate ACD with positional information to correct for HD errors with the aid of the currently derived ECD. With less exploration time in an environment, the presented model achieves already equal HD accuracy as existing models.

Kurzfassung

Tiere müssen ihre Position und Orientierung relativ zur Umgebung wissen, um sinnvolle Navigationsstrategien herzuleiten, zum Beispiel für die Nahrungssuche. Passend dazu wurden zum Beispiel in Rattengehirnen Populationen von Neuronen gefunden, deren Aktivität die Orientierung des Tieres angibt. Es wird angenommen, dass diese so genannten head direction cells (HD cells) den Orientierungssinn darstellen. Zudem wurde festgestellt, dass das Integrieren der Eigenbewegungsinformationen, auch path integration (PI) genannt, entscheidend für die Erzeugung der Aktivität der HD cells ist. Allerdings summieren sich Fehler im Signal der HD cells, wenn ausschließlich auf PI zu dessen Aktualisierung zurückgegriffen wird. Daher nutzen Tiere visuelle Orientierungspunkte aus ihrer Umwelt, um diese Fehler regelmäßig zu korrigieren. Forscher versuchten diesen Fehlerkalibrierungsmechnanismus mithilfe von Rechenmodellen zu erklären, die dessen Funktion nachahmen sollen. Die exakte Funktionsweise, wie die Aktivität der HD cells mithilfe der egozentrisch wahrgenommenen Winkellage, genannt egocentric cue direction (ECD), von nahen oder fernen visuellen Orientierungspunkten ohne Parallaxenfehler kalibriert werden kann, ist jedoch nicht vollständig geklärt. In dieser Arbeit wird ein neuartiges effizientes Rechenmodell vorgestellt, das auf biologisch plausible Weise auf der Grundlage von gepulsten Neuronen entwickelt wurde. Dafür nutzt das Modell, im Gegensatz zu existierenden Modellen, die sogenannte allocentric cue direction (ACD), die den kostanten Zusammenhang zwischen ECD und HD in einer Position beschreibt. Die kontinuierlichen Berechnungen der ACD aus ECD und HD während der Erkundungsphase, und das Berechnen der HD aus ECD und ACD in der Kalibrierungsphase werden durch einen neuronalen Schaltkreis auf der Grundlage von zwei zweidimensionalen conjunctive cell fields durchgeführt. Das Modell bietet zwei alternative Kalibrierungsmodi: first glance learning (FGL) an oder aus. Bei ausgeschaltetem FGL kann das Modell aus einer visuellen Landmarke seine HD ableiten, wenn diese ein zweites Mal von derselben Position aus wahrgenommen wird, jedoch unabhängig von ihrem Blickwinkel. Bei eingeschalteter FGL kann das Modell seine HD ableiten, wenn es eine visuelle Landmarke ein zweites Mal aus einer beliebigen Position und Ausrichtung wahrnimmt. In beiden Modi ist der Blickwinkel von dem aus die Landmarke ein zweites Mal wahrgenommen wird unerheblich, was die Effizienz im Vergleich zu existierenden Modellen, die eine Übereinstimmung von Orientierung und Position erfordern, deutlich verbessert. Die Ergebnisse zeigen, dass es ausreicht, ACD nur mit Positionsinformationen zu verknüpfen, um HD-Fehler mit Hilfe der aktuell abgeleiteten ECD zu korrigieren. In einer ähnlichen Umgebung erreicht das vorgestellte Modell mit weniger Erkundungszeit bereits die gleiche HD-Genauigkeit wie bestehende Modelle.

Contents

1. Introduction

To find food or a reproduction partner, animals, like rodents, have to navigate efficiently through their environment, even under constantly changing conditions, like the time of the year or the weather. The exploration of new areas in the environment can additionally maximize the chance to fulfill their needs (Hills et al. 2015). A newly discovered place can provide food supply or generally new resources, for example, when known ones are depleted. Therefore, animals need to combine spatial information from the known environment with newly explored areas to enable efficient navigation between known places and newly discovered ones.

Accordingly, it was found that many species can store spatial information about the environment in form of neural representations that could act as a mental map (Poulter et al. 2018) which can continuously be expanded when exploring new environmental areas. This allows certain animals, like for example sea turtles, gray wolfs, bears, and homing pigeons, to return to their home base over hundreds of kilometers (Rogers 1988; Avens 2003; Wallraff 2001; Henshaw and Stephenson 1974; Julier and Uhlmann 2001). The mechanism of creating a map by integrating spatial information, while using it to navigate, was already implemented on mobile robots and is often referred to as simultaneous localization and mapping (SLAM) (Poulter et al. 2018; Leonard and Durrant-Whyte 1991).

One central part of SLAM and for the derivation of meaningful navigation strategies is the estimation of orientation (Riisgaard and Blas 2003). An animal will only know the correct direction to reach a goal location if it knows its orientation within the environment. In line with this, so-called head direction cells (HD cells) were reported in several species like rats, mice, chinchillas, guinea pigs, and monkeys (Taube et al. 1990; Yoder and Taube 2009; Muir et al. 2009; Taube 2007; Robertson et al. 1999). These cells respond according to the animal's current head direction (HD), or rather orientation, with respect to the environment, but independently of the animal's position. These HD cells are believed to state the sense of spatial orientation of animals (Taube 2007). For this purpose, an animal needs to continuously perceive its own movements and the environment to create an internal representation of its orientation (Julier and Uhlmann 2001; Taube 2007). Body-internally derived information from the vestibular system, like an animal's angular velocity, is thought to be critical for the generation of this HD signal (Taube 2007). The process of integrating self-motion information from the vestibular system is often referred to as path integration (PI) (Valerio and Taube 2012). It allows, for example, rats to estimate their orientation even in darkness (Valerio and Taube 2012). However, the HD signal accumulates errors over time when only relying on PI (Goodridge et al. 1998). Experiments revealed that landmark information (from

e.g. olfactory or visual cues) is important for animals to correct HD errors arising from PI (Taube 2007).

Accordingly, this work addresses the problem of how HD cells can be calibrated using directional information derived from egocentrically perceived landmarks. Researchers already tried to explain HD cell calibration by designing computational models that mimic this calibration behavior. These HD calibration models can be divided into two categories. The first category contains simple feedback models that directly infer HD information from egocentric cue direction (ECD) (K. Zhang 1996). The ECD describes the egocentrically perceived angular bearing of a visual landmark. These models are only applicable to environments in which the positional relation between the agent and the visual cue is constant, for example when the visual cue is in infinite distance. In this case, the cue acts as a compass direction. When the agent perceives this cue, it can directly infer its orientation. In contrast, these models introduce parallax errors into the HD signal when the agent can move around a proximal visual cue (Bicanski and Burgess 2016). In this case, the agent can perceive the visual cue at the same egocentric angle, but from different orientations. The second category contains place encoding feedback models. These associate a combination of ECD and position of the agent with HD information (Arleo and Gerstner 2000). Subsequently, these models can infer HD information from perceiving a visual landmark from a known position and orientation in the environment. Therefore, place encoding feedback models can cope with proximal landmarks as well. However, the model can only infer HD when perceiving a landmark a second time on the exact same location and with the same orientation. Hence, the procedure of associating all possible combinations of ECDs and positions with HD makes these models less efficient. Both categories are explained in detail in Section 3.1.

This thesis presents a new biologically plausible computational model that efficiently facilitates the correction of HD errors arising from PI using a visual landmark. In contrast to existing place encoding feedback models from Bicanski and Burgess (2016) and Arleo and Gerstner (2000), the model does not have to perceive a landmark a second time in the same position and with the same orientation to infer HD. During the second perception of the landmark the orientation, and, therefore the viewing angle, of the agent is irrelevant. Therefore, an agent needs less exploration in an environment that contains a landmark to infer its HD using the egocentric perception of this landmark. For this purpose, the proposed calibration model involves three concepts. First, the model adopts a neural circuit from McNaughton et al. (1995) to conduct continuous computations of allocentric cue direction (ACD) information from HD and ECD according to an agent's current position in relation to a landmark. This ACD describes the constant angular offset between the HD and the ECD in one position. Second, the place encoding approach from previous calibration concepts from Arleo and Rondi-Reig (2007) and Bicanski and Burgess (2016) is reused to allow the model to derive directional information for HD calibration from proximal and distal landmarks in the same way, in contrast to simple feedback models. However, instead of associating one specific ECD with one HD for each location, like the existing calibration models,

only the computed offset information, ACD, is associated with positional information. Third, another adjustment of the continuous computation model of McNaughton et al. (1995) is implemented to infer HD from restored ACD and perception-derived ECD information when the agent has a second look at a landmark. The proposed calibration model is superimposed on an existing HD model from El-Sewisy (2020) that is solely estimating orientation based on PI. To mimic HD calibration behaviors found by animal experiments, the proposed model offers two separate calibration modes - first glance learning (FGL) on or off. Running without FGL (off), the model associates ACDs to already visited locations, and reuses them to restore HD when revisiting these locations. In the second mode with FGL on, the model stores only information at the first glance that the agent has on the landmark. From this moment, the model is able to infer HD whenever perceiving the landmark again from any (new) location.

Finally, the model is demonstrated and evaluated using experiments in which an agent conducts exploration phases in different spatial environments that contain one visual landmark. The model successfully learns to use ACD information to correct HD errors when perceiving a stable visual landmark. In similar environments, where the agent recurrently visits certain place fields around the visual landmark, the proposed calibration model achieves better performance to the model of Arleo and Gerstner (2000) and similar performance to the model of Bicanski and Burgess (2016) in both modes with even less exploration time. Further demonstrations are conducted with more complex exploration routes where the agent perceives a landmark exactly two times - one time at the beginning of the exploration and one time when returning from the exploration. In this case, the calibration model without FGL does not reset the HD signal as the agent never perceives the landmark two times from the same position. However, the proposed calibration model with FGL still shows adequate HD error resetting capabilities. To the author's best knowledge, this biologically plausible HD calibration model is the first one to be able to infer HD by only perceiving a landmark a second time without considering the agent's orientation. Thus, the orientation from the first perception can be different from the second one. In future work, the presented model could serve as a basis for the further development of computational models that mimic higher-level navigational functionalities such as route planning. This could lead to a deeper insight into how animals or even humans navigate.

The remainder of this thesis is organized as follows. First, the general background knowledge is presented. This includes different experimental findings about neural mental maps and details about how the HD system is driven. Afterward, work-related computational neural models from other researchers are explained. Then, the thesis research question and the requirements for a proper HD model are presented. Next, the proposed HD calibration model and its concepts are outlined, followed by experimental evaluations and a quantitative performance comparison with state-of-the-art models. The main contributions of this work are summarized as follows:

- We propose a novel HD cell calibration mechanism that efficiently uses the egocentrically perceived angular bearing of a visual cue (ECD) in combination with

positional information of an exploring agent to correct for HD errors introduced by PI.

- For this purpose, this work presents a neural circuit based on two two-dimensional conjunctive cells. This allows the model to continuously compute ACD, the global direction of a visual cue from a certain location, with given ECD and HD, during exploration, and to compute HD with restored ACD and perception-derived ECD for calibration. The model offers two alternative calibration modes.
  - first glance learning "off": place associated ACD is reused, when the agent perceives the visual cue a second time from an already visited place, to calibrate the HD cells in combination with the perceived ECD.
  - first glance learning "on": the ACD, the cue distance, and the agent's position from an agent's first glance at the visual cue are used to calibrate the HD cells when perceiving the visual cue again - independent of orientation and position.
- To demonstrate the accuracy of the HD error correction mechanism, the proposed calibration model was tested using data from a navigating robot in both simulated environments and real-world scenarios. These experiments show the model's ability to derive HD in the same way from distal and proximal landmarks. With less exploration experience in similar experimental environments, the proposed model reaches competitive performance to previously proposed calibration models.

2. Background

This chapter provides background information about biological localization. The first sections describe experimental findings about information sources and reference frames that are used by rats for determining their pose and for navigation (Hinman et al. 2018). The pose describes the combination of both orientation and position of the animal within the environment. Then, a selection of spatially modulated cells is presented that were found in rat brains and are thought to serve as a mental representation of the environment. Afterward, specific HD related findings are summarized that concern how HD can be established and driven depending on the availability of different information sources. Finally, discoveries about how visual landmarks as a particular cue can control the firing of HD cells.

2.1. Navigation Cues

Rodents can integrate multimodal sensory information to construct an internal representation of the environment (Lozano Navarro 2016). This information is also used to monitor self-position and direction with reference to allocentric (world-related) cues (O'keefe and Nadel 1978; Hinman et al. 2018). The sources of information can be categorized into allothetic and idiothetic ones (Redish et al. 1999; Lozano Navarro 2016). Both types are explained and illustrated with examples.

2.1.1. Allothetic Information

The origin of allothetic information lies external to the body of the animal (Jain et al. 2017), like features or properties of the environment (Lozano Navarro 2016). An animal can derive allothetic information from these environmental properties by using different sensory modalities, including the visual, auditory, and olfactory system (Jain et al. 2017). For example, Jacob et al. (2017) used two different olfactory cues (vanilla and lemon) that allowed rats to distinguish between two point symmetrical compartments. Thereby, the animal can localize itself after generating knowledge about the whole environment without prior knowledge about its position (Redish et al. 1999). Regarding Jacob et al. (2017), the rat would need to explore and smell the different olfactory cues in each compartment to be able to know in which compartment it is located depending on the currently perceived olfactory cue. Other examples for visual cues can be landmarks or objects that define the local environment (T. Collett et al. 1986; Clark et al. 2012).

2.1.2. Idiothetic Information

Idiothetic sources are self-referential and often referred to as internal cues (Redish et al. 1999). They are generated when the animal moves through the environment (Jain et al. 2017) and allow the animal to determine its pose in the environment like allothetic information. However, no prior knowledge about the environment is needed as only self-information is integrated. That means that the currently estimated pose is always in reference to a previous pose. The self-motion integration describes the change between the two poses. Idiothetic information can be generated from the vestibular system, the proprioception, and the motor system (C. R. Gallistel 1990; McNaughton et al. 2006). One example of idiothetic derived information is the velocity signal of an animal (Hinman et al. 2018). By integrating the self-generated information about its angular and translational velocity, the next pose in the environment relative to the starting pose can be determined. However, this integration of sensory information is prone to error accumulation due to its additive nature (Etienne et al. 2004; Hardcastle et al. 2015).

2.1.3. Combining Allothetic and Idiothetic Information

It was found that, for example, the brain of rodents improves the estimation of position and orientation by using a combination of different sensory information. It does not only rely on one single type of cue but integrates different information cues across modalities at the same time. The information sources are assessed based on their reliability (Yoder, Clark, Brown, et al. 2011). The higher the reliability, the more weight the information gets. The lower the reliability of a cue, the less weight is assigned to it. Thereby, the different and even contradicting sensory information are merged to create a unified pose estimation (Yoder, Clark, Brown, et al. 2011). Nevertheless, the complex rules of how the sensory cues are combined and how the weighing of cues is performed remain to be fully elucidated (Giocomo 2016).

2.2. Reference Frames for Representing Spatial Information

Several studies showed that animals are able to use different reference frames while navigating towards a goal location (Clark et al. 2018). Animals make use of a variety of stimulus sources that correspond to mainly two frames of reference (Knierim and Hamilton 2011). The allocentric and the egocentric reference frame. In the following, both are described using Figure 2.1 that shows an example setup of an animal navigating in a box environment (grey rectangle).

2.2.1. Allocentric Frame

The allocentric reference frame's coordinates are external to the animal and world-centered (Lozano Navarro 2016). An example of an allocentric coordination system is shown in blue in Figure 2.1. It is anchored in its rotation and location to the environment

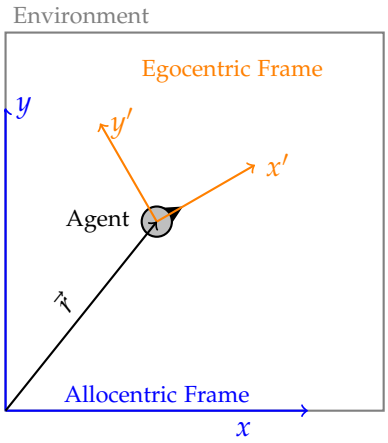

Figure 2.1.: Reference frames: The egocentric reference frame is drawn in orange and is tied to the head of the agent. The allocentric frame (blue) is fixed to the environment.

(Wolbers and Wiener 2014; Arleo and Rondi-Reig 2007). In the example Figure 2.1, the allocentric coordinate system is tied to the bottom left corner of the box environment. However, this could also be any other meaningful position, for example, based on a landmark's position. This allocentric reference frame keeps its position and orientation independently from the animal's pose in the environment. An example of spatial information that is processed in allocentric coordinates, is the animal's orientation within the environment, the so-called HD.

2.2.2. Egocentric Frame

The egocentric reference frame's coordinates are defined relative to the animal's head and, therefore, are body-centered (Arleo and Rondi-Reig 2007). The orange coordinate system in Figure 2.1 illustrates this. The coordinates are anchored to the animal's pose and depend on the animal's movements. It changes its position and orientation together with the rotation and translation of the animal (Lozano Navarro 2016). For instance, the perception of a visual cue is facilitated in egocentric coordinates. That means the perceived bearing of the cue depends on the animal's position and orientation (Arleo and Rondi-Reig 2007; Wolbers and Wiener 2014).

2.2.3. Transformation between Egocentric and Allocentric Reference Frame

As already noted, sensory input is initially represented with reference to the animal's body pose (egocentric), with the internal sense of orientation in an allocentric frame (Page and Jeffery 2018). To derive a unified representation of the environment and for self-localization, complex transformations are required to handle egocentric and allocentric coordinates (Hinman et al. 2018; Byrne et al. 2007). In congruence with that,

it was found that allocentric and egocentric reference frames can be processed together - either in parallel and sequentially (Clark et al. 2018). Yoder, Clark, and Taube (2011) stated that egocentric perceived visual information is transformed into an allocentric frame. This could be achieved using HD information (Hinman et al. 2018). However, the mechanism coordinating spatial information between both reference frames is poorly understood (Clark et al. 2018).

2.3. Spatially Modulated Cells

The firing property of spatially modulated cells represents environmental features and can serve as a mental representation of the environment (Lozano Navarro 2016). In the following, the cells that are thought to be important for navigation and SLAM are shortly described. However, the list of cells is not exhaustive. Only the necessary cells that are useful in the scope of this work are presented. All the cells illustrated are either biologically proven to exist, for example, in brains of rats or are at least proposed by other researchers based on, for instance, computational models.

2.3.1. Head Direction Cells

Head Direction Cells (HD cells) fire according to the animal's head direction, or rather its allocentric orientation, in reference to the external environment regardless of the animal's position (Valerio and Taube 2012; Jacob et al. 2017; Taube et al. 1990; Yoder, Clark, Brown, et al. 2011). Therefore, HD cells function as an internal neural compass and, consequently, as a fundamental base for navigation and SLAM (Lozano Navarro 2016). Transformations between egocentric coordinates (of e.g. sensory feature angles) and allocentric coordinates (location of environmental features) could be facilitated using the HD signal (Bicanski and Burgess 2018; Byrne et al. 2007).

2.3.2. Egocentric Cue Direction Cells

Egocentric cue direction cells (ECD cells) encode the direction of a cue in egocentric coordinates (Wilber et al. 2014) without information about the distance of the cue. K. Zhang (1996) and Skaggs et al. (1995) already hypothesized their existence via computational models. Evidence for cells that resemble view-dependent direction encoding of visual cues was found in Wang et al. (2018) and Wilber et al. (2014). For example, if an animal is positioned in an environment with one visual cue, the ECD cells will encode the direction in which the visual landmark lies with reference to the animal's orientation.

2.3.3. Allocentric Cue Direction Cells

Allocentric cue direction cells (ACD cells) are supposed to fire in relation to the direction of a cue in allocentric coordinates (Wolbers and Wiener 2014). In this case, the direction

encoded by these cells is solely dependent on the animals current position and independent from its orientation. ACD cells have not been reported yet, but are predicted by a computational model from McNaughton et al. (1995). Their firing behavior could be generated by combining information from ECD cells and HD cells (McNaughton et al. 1995).

2.3.4. Conjunctive Cells

Conjunctive cells were discovered by Wilber et al. (2014). These neurons are thought to be aligned in a grid and, therefore, form a two-dimensional conjunctive cell field (CCF) that combines the encoding of allocentric and egocentric reference frames at a time (Wilber et al. 2014) via a two-dimensional activity peak. One dimension codes for HD (allocentric) and the other dimension for ECD (egocentric) (Wilber et al. 2014). They are further explained in Section 5.2.3.4.

2.3.5. Egocentric Boundary Vector Cells

Alexander et al. (2020) reported egocentric boundary vector cells. These appear as a two-dimensional neuronal field. Its neurons respond to the presence of environmental boundaries. Each neuron encodes the position of a boundary relative to the animal's pose (Alexander et al. 2020). For example, if an animal is located in a box environment and faces a corner, the egocentric boundary vector cells whose distance and direction parameter fit to the current distance and perceived direction of the corner wall will be active. Hence, the shape of the perceived environmental walls will be mapped by the activity of the two-dimensional cell field (Hinman et al. 2019). If the animal rotates, the activity of the egocentric boundary vector cells will follow this rotation (Hinman et al. 2019).

2.3.6. Allocentric Boundary Vector Cells

Allocentric boundary vector cells are built as a two-dimensional cell field as well. The firing behavior of the cells in the field depends, like in egocentric boundary vector cells, on the presence of an environmental boundary. In contrast, the allocentric boundary vector cells respond according to a combination of distance and allocentric direction from the animal Solstad et al. (2008). That means that the firing does not depend on the animal's orientation but solely on the animal's allocentric position within the environment (Lever et al. 2009). They are a mental representation of physical environmental boundaries and, thereby, implicitly encode the environmental geometry (Lozano Navarro 2016). If an animal faces a corner that for example lies to the west of the animal's position, the mapping of the corner's walls in the allocentric boundary vector cells will also encode the direction of the corner in allocentric coordinates. In this case, the cells will map the corner also to the west of the animal's position. If the animal stays in this position and just rotates, the firing profile will stay the same. Only if the animal translates, the

firing profile will change. Allocentric boundary vector cells are thought to be a result of consistently integrating the HD signal with information from egocentric boundary vector cells (Evans et al. 2016).

2.3.7. Egocentric Object Vector Cells

Egocentric object vector cells form a two-dimensional field which is encoding the distance and direction of objects, like a visual landmark, in the environment in egocentric coordinates. If an animal perceives a landmark, for example, to its left side at a certain distance, the corresponding neurons to these variables will be active. Thereby, the position of the object is represented by an egocentric map. Egocentric object vector cells have been computationally proposed by Bicanski and Burgess (2018).

2.3.8. Allocentric Object Vector Cells

Allocentric object vector cells (Høydal et al. 2019) or landmark vector cells (McNaughton et al. 1995) code for the position of an object, or rather a landmark, in relation to the animal's position, independent of its orientation, and, thus, in allocentric coordinates, in contrast to egocentric object vector cells (Deshmukh and Knierim 2013). In other words, these cells fire in spatial locations in a relationship to local landmarks (Deshmukh and Knierim 2013). These cells were proposed by McNaughton et al. (1995). Wilber et al. (2014) and Høydal et al. (2019) reported on cells that resemble this firing behavior.

2.3.9. Place Cells

Place cells were reported the first time in the hippocampus of rats by O'Keefe and Dostrovsky (1971). This population of cells is sensitive to an animal's location within the environment (O'Keefe and Dostrovsky 1971). Whenever an animal enters a certain place field the corresponding place cell's firing activity increases. If the animal is located in the center of such a place field, the firing activity will be the highest. If the agent moves away from a place field's center, the firing rate will decrease (Muller et al. 1987). The shape of place fields was found to be circular or slightly elliptical (Lever et al. 2002).

2.4. Head Direction Updating Strategies

This section gives insight into the HD system's modulation. It is capable of integrating different sources of information depending on their availability and reliability (Lozano Navarro 2016). Researchers mainly refer to two strategies to estimate HD (Yoder, Clark, and Taube 2011; Jain et al. 2017; C. R. Gallistel 1990):

- path integration: based on idiothetic cues
- landmark-based updating: based on allothetic cues

Both strategies are described in the next two subsections, followed by combining strategies to optimize the estimation of HD.

2.4.1. Path Integration

path integration, also known as dead reckoning, is a mechanism that uses idiothetic cues to update the representation of one's current pose in relation to a known pose (C. R. Gallistel 1990). Therefore, the animal needs to monitor its trajectory in relation to the known start location (C. R. Gallistel 1990; Whishaw and Wallace 2003). The start can be the nest or the burrow of animals (McNaughton et al. 2006). Then, idiothetic cues which identify the animal's traveled distance and direction of motion and speed are integrated to compute the new pose (C. R. Gallistel 1990; McNaughton et al. 2006). Therefore, no knowledge about the environment is needed (Redish et al. 1999). Furthermore, the integrated information can be used to compute a homing vector to return directly to the start location (McNaughton et al. 2006). The first evidence that mammals use PI was found by M. L. Mittelstaedt and H. Mittelstaedt (1980) who showed that gerbils can return to their nest after retrieving their pups in darkness in a circular arena. Additionally, researchers reported that, for example, rats can maintain their sense of orientation when foraging blindfolded in an environment (Valerio and Taube 2012; Etienne et al. 1998). PI can also be used to distinguish between visually identical targets in a symmetrical environment (Etienne et al. 1998). However, as PI is summing up information, it also accumulates inaccuracies of the idiothetic input which results in a significant drift of the HD signal over time in darkness (Goodridge et al. 1998). The rats can not rely on visual information anymore and must rely on self-motion information only (Giocomo 2016). The resulting HD error strongly correlates with a rat's navigational error when trying to return to a starting position (Valerio and Taube 2012). When errors are accumulated over a longer time the animal even becomes disoriented which results in the degeneration of the HD signal (Valerio and Taube 2012; Goodridge et al. 1998).

2.4.2. Landmark-Based Updating

Landmark-based updating, also known as piloting or landmark navigation (Yoder, Clark, and Taube 2011), is a process by which an animal determines its position and orientation via exploitation of landmarks from the environment (C. R. Gallistel 1990). Yoder, Clark, and Taube (2011) divided landmark-based updating into three parts:

1. Determination of location and directional heading with an environment by perceiving landmarks
2. Computation of a planned route or trajectory
3. Execution of that trajectory

The execution of these steps requires the animal to carry out complex computations in order to recognize and categorize visual features from the environment (Hinman

et al. 2018). Presumably, neuronal mechanisms evolved which can compute spatial position and link this information with the features and location of landmarks (Buzsáki and Moser 2013). Additionally, the animal needs to transform the egocentric perceived environmental features into an allocentric reference frame to get a unified data situation to estimate its pose in the environment (Hinman et al. 2018). For this purpose, an animal needs to collect prior knowledge about the environment but not about its current pose (Redish et al. 1999). Evidence for landmark-based updating was found in many studies. Jacob et al. (2017) reported HD cells in rats that are anchored to local cues in a two compartment experiment. The compartments were connected, and their geometrical layouts were point-symmetrical. Additionally, each compartment contained one local visual cue at opposite walls. Therefore, the visual landmarks are also located in point-symmetric positions. When the rats visited one compartment, HD cells responded in relation to the local perceived visual cue but still in an allocentric manner.

Similar firing behavior was found in cue manipulation experiments from Clark and Taube (2011) where HD cells maintained their firing behavior anchored to a visual cue even when the cue is manipulated. This was also found by Zugaro et al. (2003) where transitions in the activity of HD cells were recorded 80 ms after a visual cue is rotated. Altogether, these results show that animals use different visual features to aid spatial orientation (Lozano Navarro 2016). However, it is controversial how the visual impression of the environment is processed to finally control the HD signal. Two major models try to explain this behavior (Lozano Navarro 2016). Both are described in the following.

2.4.2.1. Distal Sensory Input Model

The distal sensory input model from Knierim and Hamilton (2011) states that the visual features of landmarks are processed. These features are thought to be relevant to reset HD cells instead of computing the whole visual configuration of the environment. On the one hand, visual cues can anchor the animal's internal coordination system to the environment, and on the other hand, local features like boundaries can serve as a basis for the computation of self-location by providing distance information (Knierim and Hamilton 2011). With the goal to find a specific location, this model suggests that spatial relationships between landmarks need to be computed (Lozano Navarro 2016). Accordingly, Lozano Navarro (2016) investigated if visual landmarks are processed globally as a panoramic scene or based on their individual stimulus features. It was found that animals are differently sensitive to visual features of cue cards, like contrast or orientation, and not to the overall spatial configuration of landmarks (Lozano Navarro 2016). This supports the explanation based on the distal sensory input model.

2.4.2.2. View-Based Model

The view-based model states that a full retina image resets the activity of HD cells instead of explicit landmark features (Zeil 2012). In other words, global positional

and directional information is implicitly included in panoramic visual scenes (Zeil 2012). Therefore, no fine visual details need to be extracted from the environment (Zeil 2012). This would suggest that the animal's HD system is not sensitive to particular environmental features and processes them implicitly based on their spatial relation as a cue array (Lozano Navarro 2016). Cartwright and T. S. Collett (1983) introduced a similar model that was derived from insect experiments. It is based on an image matching strategy to learn visual landmarks using panoramic snapshots. These snapshots can be gathered by a head scanning behavior that was, for example, observed in rats by Cartwright and T. S. Collett (1983). According to the view-based model, this would be enough to set HD and navigate without the need for a cognitive map (Monaco et al. 2014).

Page and Jeffery (2018) supports this model stating that the brain presumably samples panoramic snapshots of the field of vision. Stürzl et al. (2008) evaluated this explanation by investigating the information content of panoramic images taken in an axially symmetrical rectangle box. One goal image was taken in one corner and compared with images taken in other locations in the box by calculating the global image differences (Stürzl et al. 2008). Finally, this showed that the salience of visual cues, like the geometry, corners, or top edges of the arena, within these panoramic snapshots determines the image differences. As a result, this model assumes that no explicit computation of the geometry or visual features from visual input is required (Stürzl et al. 2008).

2.4.3. Combining Head Direction Updating Strategies

Animals were found to combine information from PI and landmark-based updating to create a unified data situation which allows them to estimate their allocentric pose as consistent and accurate as possible (Etienne et al. 2004; Yoder, Clark, Brown, et al. 2011; T. S. Collett and Graham 2004). This combining process includes mainly two mechanisms (Poulter et al. 2018):

1. Update, respectively calibrate, the HD signal to erase introduced errors from PI using visual landmarks

2. Calibrate the underlying mechanisms and parameters of PI to enhance its reliability (e.g. gain modulation)

As the second point is beyond the scope of this work, the focus will stay on the calibration of the HD signal using visual cues.

Evidence for calibrating the HD signal that accumulated errors via PI to restore its accuracy was mentioned in several scientific reports (Taube 2007; Yoder, Clark, and Taube 2011; Lozano Navarro 2016; O'keefe and Nadel 1978). However, the mechanisms how the updating process of the HD system is conducted are not fully understood yet. Nonetheless, two essential parameters seem to influence the combination process significantly. These are the accessibility and the reliability of cues which are explained more in depth.

2.4.3.1. Accessibility of Cues

Idiothetic information, like self-motion, is continuously accessible to an animal which allows the animal to constantly use PI even, for example, in darkness (Taube 2007; Lozano Navarro 2016). In contrast, allothetic information can be inaccessible. For example, visual cues can not be detected in darkness and proprioceptive cues can be unavailable in an open field without boundaries (Taube 2007). Therefore, PI and, for example, landmark-based updating can only be combined to update the HD system if the environmental conditions allow the animal to gather the allothetic information.

2.4.3.2. Reliability of Update Processes

Concurrently working PI and landmark-based updating can indicate conflicting HD. Therefore, the information from both strategies is combined based on the reliability of both processes (Lozano Navarro 2016).

Path Integration. The PI's reliability is affected by the animal's degree of disorientation, or rather degree of disruption of the self-motion signal (Lozano Navarro 2016; Knierim et al. 1995). However, experimental findings investigating the impact of disorientation are not clear and show ambiguities (Poulter et al. 2018). Some experiments showed a strong reliance on landmarks after rats were disorientated. Others found weak reliance on landmarks after conducting the disorientation procedure (Poulter et al. 2018). Lozano Navarro (2016), for example, found strong HD control of visual landmarks in rats which could show that the reliance on self-motion cues was diminished by disorienting them. Additionally, Lozano Navarro (2016) assumed that ambiguous findings of the effect of disorientation can arise from a non-standardized procedure to disorient animals in the experiments.

Landmark-Based Updating. In contrast, the reliability of landmark-based updating seems to depend on more variables. Familiarity with a landmark was found to make landmark-based updating dominant over PI in controlling the HD system (K. Zhang 1996; Goodridge and Taube 1995). Another variable is the salience of an allothetic cue in the environment. A more salient landmark is easier detectable and is less ambiguous (Lozano Navarro 2016). Furthermore, the allothetic cue must be perceived as stable in relation to the environment (Lozano Navarro 2016). Allothetic cues that are perceived as spatially unstable are not feasible to be used as points of orientation. Accordingly, it was shown that rats can use a visual cue to anchor their HD that is perceived as stable but in fact is unstable (Yoder et al. 2015; Knierim et al. 1995). More details about what can serve as a specific visual landmark and which visual features are favorable for animals, like rats, is given in Section 2.5.4.

2.4.3.3. Weighted Average Model

When the HDs derived from idiothetic and allothetic cues are in conflict, a trade-off value is calculated based on the cues' reliabilities. This was shown in several experiments and is often referred to as undershooting (Jeffery et al. 2016; Knight et al. 2014). According to these findings, Fetsch et al. (2013) proposed a model which integrates cues in an optimal way based on Bayesian probability. Therefore, weights are assigned to every cue that needs to be integrated according to its reliability (Page and Jeffery 2018). Subsequently, the weighted average information from all input cues is calculated. Consequently, the corrected HD signal significantly undershoots the indicated HD from, for example, the landmark-based updating process (Knight et al. 2014).

2.4.3.4. Rotational and Global Correction of Head Direction

If the newly estimated HD differs from the currently indicated HD, the HD cells' activity peak is shifted accordingly. This is done via a calibration signal that activates the neurons corresponding to the new HD estimation. Appropriately, two different HD error correction behaviors were observed (Valerio and Taube 2012). If the conflict between the current HD and the HD indicated by the calibration signal is smaller than 45 °, the current HD peak is shifted continuously towards the new position. This is referred to as the rotational correction of the HD peak position (Valerio and Taube 2012). If the conflict is larger than 90 °, the current HD peak degenerates, and afterward, the new HD peak emerges at the position indicated by the calibration signal. This behavior is called global HD correction (Valerio and Taube 2012).

2.4.3.5. Resetting and Remapping

Resetting and remapping are two types of behavioral HD correction or updating processes that are defined depending on the familiarity of an animal with its environment (Valerio and Taube 2012).

Resetting happens when the animal revisits a known (home) location or scene (Poulter et al. 2018). Here, the remembered HD from the last visit to this location is restored and the old known HD signal at this place is re-established (Taube et al. 1990; Valerio and Taube 2012). This means that the HD is already anchored to a memorized environmental cue which associates with the finding of a rat's circling behavior around its home base (Poulter et al. 2018). In this way, the rat gathers information about unexplored areas but keeps its sense of orientation by constantly revisiting the home base to regularly reset HD.

Remapping happens when an animal has to orient itself in an unknown environment (Lozano Navarro 2016) or if a large HD error (disorientation) occurs due to PI (Valerio and Taube 2012). Then, the animal sets a new reference frame for its internal orientation by anchoring the HD to, for example, a newly perceived salient allocentric landmark (Valerio and Taube 2012). Subsequently, the HD remains anchored to this allocentric

cue as the animal becomes familiar with the environment. From this moment on, the resetting process is used. It can be stated that remapping serves as the basis for resetting.

2.5. Visual Landmarks for Updating Head Direction

As the visual system is the most reliable long range-sensing source of information, a visual landmark, as one example for allothetic cues, provides useful data to aid spatial orientation (C. R. Gallistel 1990). In this work, a visual landmark, or a visual cue, is defined as a visual feature of the environment that enables the determination of one's pose within the environment. It was demonstrated that learning of a novel visual landmark can occur very fast (Yoder, Clark, and Taube 2011). According to Goodridge et al. (1998), a single one minute exposure to a new landmark was sufficient to gain control over the HD signal. This section aims to describe how visual landmarks are processed by the visual system of animals (e.g. rodents), what can serve as a landmark, and which visual properties or features are favored to use a landmark as a point of orientation.

2.5.1. Landmark Processing

This work divides the processing of landmarks into three parts. The first part is the perception of the landmark with the visual system. The second and third part describe the dorsal and the ventral visual pathway which are two distinct streams that further process the image data from the primary visual cortex (Mishkin et al. 1983).

2.5.1.1. Visual System

As most of the findings, used in the scope of this work, were retrieved from conducting experiments with rats, the properties of a rat's visual system are shortly summarized. The visual system of rats consists of two laterally placed eyes. This results in a field of vision of 280 °. The overlap between the fields of vision of both eyes lies between 40 ° to 60 ° (Artal et al. 1998; Lozano Navarro 2016). This allows the rat to detect visual cues and to derive egocentric angular bearing and distance information.

2.5.1.2. Dorsal Visual Pathway

In the dorsal visual pathway, spatial aspects of the visually perceived information are processed (Mishkin et al. 1983), inter alia, to transform egocentric perceived information into an allocentric reference frame (Yoder, Clark, and Taube 2011). The following information can be derived from the dorsal stream (Howard and Howard 1982; Yoder, Clark, and Taube 2011):

- distance between viewer and objects
- viewer movement from the optical flow

- spatial relationships between objects
- geometric environmental properties
- up-down orientation of the visual scene

Consistent with this, rats were found to be capable of recognizing shapes and complex features from objects (Tafazoli et al. 2012).

2.5.1.3. Ventral Visual Pathway

The ventral pathway is thought to contain the processes of detection, identification, and classification of visual features, respectively objects (Yoder, Clark, and Taube 2011). These processing steps are necessary before visually perceived information can be used as a landmark (Yoder, Clark, and Taube 2011). For example, visual salient but spatially unstable objects should be classified as useless for deriving one's orientation in space (Yoder, Clark, and Taube 2011).

2.5.2. Examples of Visual Landmarks

In many experiments, it was investigated which visual feature can be used for orientation. It was discovered that rats as an example can rely upon a variety of visual landmarks to anchor their HD, like:

- objects (Zugaro et al. 2001)
- borders and boundaries (Giocomo 2016; Evans et al. 2016; Clark et al. 2012)
- boundary edges (Stürzl et al. 2008)
- corners (Clark et al. 2012)
- cue cards (Goodridge et al. 1998)

In general, any salient visual information (distal or local) derived from the environment that is perceived as stable can serve as a landmark (Knierim and Hamilton 2011; Yoder, Clark, and Taube 2011). Many researchers also reported that geometry or the shape of an environment is highly influential for the rat's sense of orientation (Poulter et al. 2018; Hardcastle et al. 2015; Cheng 1986). However, "geometry" is implicitly provided by boundaries, edges, and corners (Pearce et al. 2004).

2.5.3. Landmark Identity and Spatial Geometry

When an animal orients itself in a bounded environment that contains one landmark, the question arises if an animal relies more on the identity of a single landmark in relation to the geometrical framework of the environment or if the geometrical framework (boundaries, edges, and corners) is solely used for orientation. Both possibilities would

provide enough distance and direction information from the environment to determine one's position exactly. This is similar to the side-angle-side statement of congruence for triangles (Mironitchev and Morales 2019), where two distance information and one angle are enough to fully determine the triangle's form.

It was shown that rats can rely on the identity of a single cylindrical landmark in a black-painted square room to keep their sense of orientation (T. Collett et al. 1986). Poulter et al. (2018) mentioned that rats can use visual features to disambiguate geometrically equivalent locations after training (Cheng 1986). Additionally, Clark et al. (2012) found that multiple salient extra-maze landmarks can override the effect of using environmental boundaries for orientation.

In contrast, other findings suggest that animals prefer geometrical cues, like boundaries or corners, to visual cue cards (Lozano Navarro 2016; Giocomo 2016; Margules and C. Gallistel 1988). This often results in rotational HD errors in symmetrical environments whose geometrical properties offer insufficient data to establish HD (Stürzl et al. 2008). This can even be observed if a colored visual cue would break the symmetry (Stürzl et al. 2008).

It can be concluded that animals, like rodents, are capable of orienting themselves in both ways but have primacy to the geometric arrangement, like the corners in a box environment, which is reported from many mammalian species (Learmonth et al. 2002; Hermer and Spelke 1994; Lozano Navarro 2016). However, they are still able to use the identities of landmarks after training (Greene and Cook 1997).

2.5.4. Impact of Landmark Identity

The perception of a landmark varies depending on the distance and the viewing angle of the observer. Therefore, when using the identity of a landmark for orientation, it is fundamental to extract visual features to create a spatial-contextual representation (Lozano Navarro 2016). Hence, this part deals with what visual features can be extracted from landmarks and which conditions affect their use as points of orientation.

2.5.4.1. Visual Features

Salience and Uniqueness. A unique physical visual appearance of a landmark within the environment is important for orientation (Lozano Navarro 2016). Uniqueness and salience together make a landmark easy to recognize without the risk of confusion compared to when other similar looking features exist (Yoder, Clark, and Taube 2011).

Contrast. Psychometric curves can detect contrast differences in visual cues (Busse et al. 2011). In line with this, Lozano Navarro (2016) discovered that rats prefer a visual cue card with a white bar on a black background to a black bar on white background.

Orientation. Rodents also show a preference for horizontal and vertical orientations of lines to oblique orientations (Roth et al. 2012; Lozano Navarro 2016). This affinity

could arise from natural scenes in which cardinal lines dominate over oblique shapes (Girshick et al. 2011).

2.5.4.2. Positional Aspects

Stability. One fundamental property of a visual cue to be used as a point of orientation is that it is, or at least is perceived as, spatially stable (Lozano Navarro 2016; Yoder, Clark, and Taube 2011; Page and Jeffery 2018). An unstable visual cue cannot give reliable information for orientation. A visual cue is defined as stable when it has a constant distance to other visual cues and a persistent spatial position in the environment (Lozano Navarro 2016).

Location. Visual cues which are placed at certain behavioral important locations become more attractive as points of orientation and are more memorable (Chan et al. 2012). These can be goal locations like food supply, familiar locations, or decision points, like the entrance to a maze (Valerio and Taube 2012).

Distance. Another important property that affects the reliability of visual cues is the distance from the observer. To orientate itself, a rat generally favors distal cues over proximal cues (Knierim and Hamilton 2011). On the one hand, distal cues shift less rapidly as the observer moves. This results in a smaller parallax effect. Therefore, distal cues indicate orientation similar to a compass (Lozano Navarro 2016). On the other hand, a distal cue that has the same physical size as a proximal cue is perceived as smaller and, therefore, more accurate or punctured (Lozano Navarro 2016). Nevertheless, proximal cues can also provide information for the HD system when, for example, no distal cue is present (Lozano Navarro 2016). If both proximal and distal cues are present at a time, proximal cues will be used to derive distance information to find a local goal position (Knierim and Hamilton 2011).

3. Related Work

Researchers already proposed several models to correct for HD errors that arise from PI, using egocentric perceived visual cues. These models and their basic ideas are described in short in the first part of this chapter. Second, this chapter takes a look into computational models that were proposed to process reference frame transformations. The third part deals with computational models that allow combining of conflicting cue information to one unified data situation by weighing the cues according to their reliability.

3.1. Head Direction Calibration Models

Existing HD calibration models can be divided into two categories according to their difference in the use of allothetic information. Figure 3.1 shows the main idea of both categories - simple feedback models and place encoding feedback models. Both categories and related previously proposed models are explained in the following.

3.1.1. Simple Feedback Models

Simple feedback models use the perceived ECD of a visual cue as a single source of allothetic information to directly infer HD. Their general setup is shown in the left part of Figure 3.1. Simple feedback models consist mainly of two rings of neurons. The outer ring consists of ECD cells which directly project onto the inner ring that is composed of HD cells. If the spatial relation between an agent and a visual cue in the environment is stable, these model can sufficiently derive their HD from the egocentric perception of a visual cue.

It must be noted that this model is only working properly under two conditions. Either the visual cue is in infinite distance, or the agent's freedom of movement is limited to a straight line originating from the landmark's position. This ensures that the ECD always has the same offset to the HD. A proximal cue is useless if an agent is allowed to move freely as the agent's position relative to the visual cue affects the offset between the ECD and HD. For example, the simple feedback model shown in Figure 3.1 will introduce an HD error of up to 180 ° if the agent is located in the purple field in the example situation.

K. Zhang (1996), for example, proposed a simple feedback model that uses a continuous attractor network (CAN) consisting of two rings to encode the current estimation of HD by integrating angular velocity input from the agent. A third outer ring, composed of ECD cells, encodes the ECD of a visual landmark. The ECD cells directly project onto

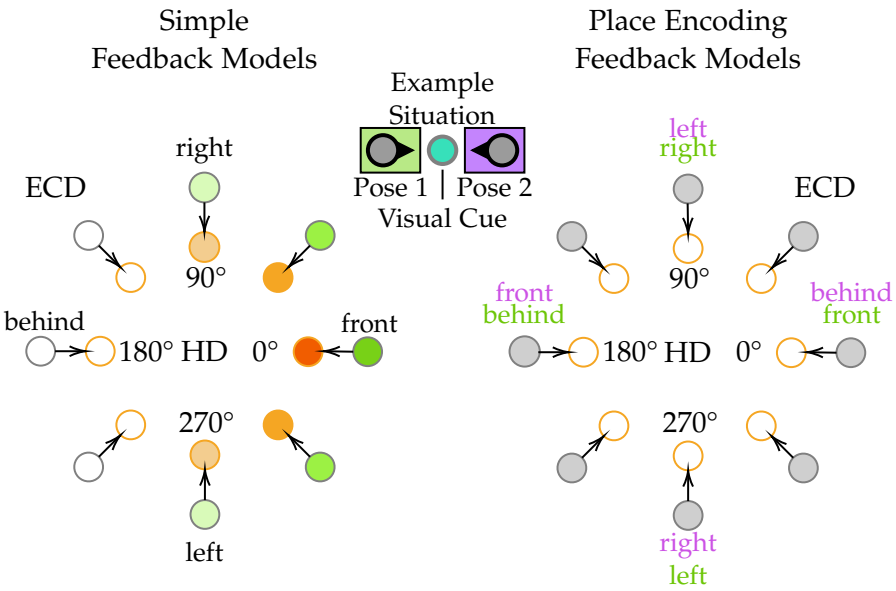

Figure 3.1.: The two main types of HD calibration models: Simple feedback models consist of two rings of neurons. The outer ring encodes the ECD of a perceived visual landmark and directly projects onto the inner ring which encodes HD. In this case, the connections are chosen to suit the spatial relation between the visual cue and the agent when located in the green field. The agent perceives the visual cue to its front. Therefore, the ECD cell representing "front" is the most active and drives the HD cell representing to 0 ° which is in accordance with the agent's orientation. Place encoding feedback models associate the information about the agent's current position with the current perceived ECD to infer HD. Their setup can be imagined as using several ECD rings that directly drive the HD. Depending on the current position of the agent, different ECD cells are active. If the agent is located in the green field, the ECD cells will correspond to the ECDs labeled in green, identical to the setup of the simple feedback model. The agent perceives the visual cue to its front which indicates an HD of 0 °. If the agent is located in the purple field, it perceives the visual cue again to its front. However, its orientation is rotated by 180 °. In this case, the ECD cells responding according to the green labels would introduce a huge HD error. Therefore, the second set of ECD cells corresponding to the purple labels is active and indicates the correct HD of the agent when located in the purple field (180 °).

one of the two rings of HD cells via one-to-one connections. The connection weights are one-shot Hebbian learned at the first glance at the landmark. Subsequently, the ECD cells control the position of the activity profiles in both HD rings to correct for HD errors introduced by PI.

The simple feedback model from Skaggs et al. (1995) is based on a continuous attractor network (CAN) whose activity profile indicates the current HD estimation similar to K. Zhang (1996). It also integrates angular velocity to mimic PI. Another ring of visual feature detectors around the CAN encodes the ECD of a visual landmark. However, the visual feature detectors have constantly Hebbian-learning connections to all neurons of the CAN which allows the model to connect the HD representation to spatially stable landmarks and to disconnect from unstable landmarks.

3.1.2. Place Encoding Feedback Models

The second category, the place encoding feedback models, uses the perceived ECD of a visual cue to infer HD according to the current position of the agent around the landmark. The right part of Figure 3.1 illustrates the general concept of place encoding feedback models. The inner ring consists of HD cells. The outer ring consists of ECD cells that encode the ECD of the visual landmark depending on the agent's current position (green or purple field). The ECD cells project again directly onto the HD cells. For example, the top neuron in the ECD ring represents the ECD "left" when the agent is in the purple field of the example situation and "right" when the agent is in the green field. Thereby, place encoding feedback models can derive HD information from proximal cues as well as from distal cues.

The place encoding feedback model of Bicanski and Burgess (2016) uses sheets of neurons that act as position dependent interface to associate HD with ECD. A PI mechanism accounts for integrating angular velocity to establish the HD signal. The model encodes the ECD of a visual landmark via a visual sensory input, similar to ECD cells. Place cells encode in which place field the agent is located within the environment, and activate exactly one sheet of interface neurons for each place field in the environment. The currently active sheet copies the HD cells' activity. Hebbian-learning connections account for connecting the HD encoded by the currently active sheet of interface neurons with the neurons of the visual sensory input. When perceiving a landmark a second time from the same angle and position, the HD will then be restored to the HD from the first glance at the landmark via the learned connections by the visual sensory input.

The place encoding feedback model of Arleo and Gerstner (2000) uses only one set of interface neurons, called calibration cells, to infer HD from ECD and positional information about the agent. Each calibration cell receives input via learning connections from both the visual sensory input that encodes the ECD of a visual landmark and the position encoding cells. Additionally, the calibration cells have learning connections towards HD cells. When operating in a learning stage, the calibration cells copy the firing activity of HD cells and the three sets of learning connections are adjusted in relation to the place and ECD encoding sensory input using a Hebbian rule. In the

exploitation stage, the learned connections are fixed and, thus, the calibration cells are solely driven by the two sensory input components and, in turn, calibrate the HD cells.

Milford et al. (2004) proposed a place encoding feedback model that uses two three-dimensional attractor cell networks. A pose cell network encodes the position of the agent in two dimensions and the orientation of the agent in the third dimension, such that, a three-dimensional activity packet describes an agent's pose completely. To mimic PI this activity packet is shifted artificially in the network. The second part is called local view cell network and encodes the local perception of landmarks. One dimension accounts for a visual cue's color, one encodes the cue distance, and the last one encodes the ECD. All pose cells network receive Hebbian-learning connections from the local view cells. Thereby, the information about the cues' distances and bearings are linked with the agent's pose. This allows the model to calibrate its HD using several proximal and distal cues.

3.2. Frame Transformation Models

The following models are theoretical and computational proposed mechanisms on how spatial information can be transformed from egocentric coordinates into allocentric coordinates and vice versa. On the one hand, they try to explain how animals, that can collect allothetic information only from an egocentric perspective, can transform this information from an egocentric into an allocentric reference frame. On the other hand, they aim to give an explanation of how animals could use an allocentric representation of their environment to derive navigation strategies that need to be performed from the animal's egocentric perspective. The idea of these reference transformation mechanisms is used to design the presented HD calibration model presented by this work to conduct continuous computations in a biologically plausible way based on spiking neurons (See Section 5.2.1). This allows the HD calibration model to, first, transform ECD into ACD using HD and, second, to derive HD with given ACD and ECD.

The model of Bicanski and Burgess (2018) consists of mainly three components. The first component contains egocentric boundary vector cells and egocentric object vector cells, and the second component contains allocentric boundary vector cells and allocentric object vector cells (See Section 2.3). The third component facilitates the transformation between the neural representations of the first and the second component using the agent's current HD. This shows that the HD signal can be the bridge between cells that encode in egocentric coordinates and cells that use allocentric coordinates.

McNaughton et al. (1995), summarized by Clark et al. (2018), proposed a three-layer model to transform the perceived ECD into ACD. The first layer contains two neural populations: ECD cells and HD cells. The second layer contains conjunctive cells which simultaneously receive ECD and HD information from the first layer. In turn, the conjunctive cells drive the activity of the neural population of the third layer that encodes the allocentric, world-centered, and viewer-independent bearing of the visual cue (ACD). Subsequently, cue distance information could be used to create landmark

vector cells as described in Section 2.3.8.

3.3. Cue Information Weighing Models

The last group to be presented are computational models that aim at mimicking the experimental findings of weighing different cues according to their reliability (Knight et al. 2014). This can, for example, explain the undershooting when animals reset their HD according to a visual cue (Knight et al. 2014). A selection of two models, of Jeffery et al. (2016) and Page et al. (2014), are presented in short. However, several other researchers proposed ideas solving this problem as well, like Campbell et al. (2018), or W.-H. Zhang et al. (2019), whose model is able to integrate and also segregate information. However, these weighing mechanisms are not implemented in the final calibration model presented by this work but offer extension options for further development.

Page et al. (2014) used short-term plasticity connections to connect all ECD cells to all HD cells of their simple feedback model (see Figure 3.1). From start, the connection weights are set such that each ECD cell projects one-by-one onto its corresponding HD cell. When the ECD cells attempt to shift the activity peak in the HD to a new position, a rapid re-weighing of the connections during this shift accounts for setting the new position of the HD cell activity peak to an intermediate position in between its previous position and the position indicated by the ECD cells.

Jeffery et al. (2016) presented another model based on the simple feedback setup. The model combines the short-term plasticity of Page et al. (2014) and the long-term plasticity of Skaggs et al. (1995) to connect the ECD cells to the HD cells. On the one hand, it resets the position of the HD activity packet to an intermediate position between two conflicting indications of HD from two different cues, like Page et al. (2014). On the other hand, the model disconnects the HD ring from spatially unstable cues.

4. Problem Statement

As described in Section 2.4.1, animals can maintain their sense of orientation without allothetic cues by integrating self-motion information (Valerio and Taube 2012). However, any bias or disturbances in the self-motion information are integrated as well. Step by step, the integration of inaccurate information leads to growing errors in the animal's estimation of HD which results in a drift of the HD cell activity over time (Goodridge et al. 1998). In the long run, these effects lead to disorientation of the animal (Valerio and Taube 2012). Therefore, animals exploit directional information provided by stable visual cues to anchor HD (Yoder, Clark, Brown, et al. 2011).

The question arises how animals are able to derive HD from an egocentrically perceived landmark. This motivates the development of a computational HD calibration mechanism to correct for HD errors introduced by PI. In line with this, the aim of the present thesis is to design a neural HD calibration circuit to extend an existing HD model of El-Sewisy (2020) that solely relies on PI to estimate HD and, hence, accumulates errors in the long run. This should allow the extended HD model to connect its HD in relation to the position of a visual landmark and utilize this knowledge to calibrate the HD cells when revisiting the visual landmark after a PI based exploration phase.

Based on the models already described in Section 3.1 and the description of the experimental findings of biological mechanisms for navigation and localization described in Section 2, an optimal biologically plausible HD calibration model should have the following characteristics.

Compatibility. The model has to be compatible with the model of El-Sewisy (2020). El-Sewisy (2020) used spiking neurons as the basic framework to ensure the biological plausibility of his model. Therefore, spiking neurons have to be used as the building blocks for the calibration model as well.

Furthermore, the currents and firing rates of the calibration model must be compatible with the ones appearing in the HD model of El-Sewisy (2020). In El-Sewisy (2020), a firing profile emerges in the HD cells whose peak position indicates the system's current estimation of HD based on PI. A compatible calibration mechanism should generate a calibration signal that, on the one hand, is capable of resetting the peak position of the HD cells' firing activity to any position. On the other hand, this calibration signal should not lead to a large bias of the activity profile's shape (please refer to Section 5.2.2.2 for the descriptions of the activity profile). The shape should still remain symmetric and the peak firing rate should remain about the same size.

Composition. The calibration circuit should only be composed of components that mimic spatially modulated cells that are already proposed or proven to exist. However, this selection is additionally restricted to cells whose firing behavior is not based on HD information. For example, place cells are supposed to be generated based on HD information (Jercog et al. 2019) and, therefore, not useful to reset HD because their accuracy is affected by the accuracy of HD. This restricts the selection to cells that are solely driven by egocentric sensory inputs and independent from HD information.

Functionality. The calibration model of Arleo and Gerstner (2000) (see Section 3.1) contains three sets of Hebbian-learning connections. Therefore, the model operates in two different stages. During the learning stage, the connections are adjusted to associate place and ECD with HD information. The calibration stage accounts for correcting HD errors via the previously learned connections. In contrast, a more sophisticated model in terms of biological plausibility should facilitate learning and calibration at the same time without the need for artificially switching between the learning and the exploiting stage.

Moreover, an optimal model should also be capable of using distal as well as proximal cues to anchor its HD without introducing parallax errors arising from a shifted egocentric perception of a visual cue during translational movements.

Furthermore, the model should be able to optimally integrate indicated HDs from different cues. For example, the model should learn to only rely on stable perceived visual cues to infer HD and should disconnect from as unstable perceived visual cues. When a conflict between two reliable cues arises, the model should create a unified data situation by computing a weighted average value according to each cue's reliability.

Efficiency. To build an efficient model, the information that is learned from the perception of a visual cue should be meaningful enough to serve the calibration function but as compressed as possible to not store redundant information and to reduce the learning effort. In the model of Bicanski and Burgess (2016) (described in Section 3.1) one sheet of neurons per place is needed to activate the corresponding learning connections that associate HD and ECD. This can be avoided by only associating the combination of offset information, between HD and ECD, with information about the agent's position.

Extensibility The final model should function with a minimum of one visual landmark to orient itself. However, for future work, the model should provide the possibility to be extended to handle more landmarks.

5. Methodology

This chapter explains the newly proposed place encoding feedback model for HD cell calibration and its functionality. The beginning provides an overview of the model architecture. Next, the components of the HD calibration system are explained, followed by the illustration of the connection weights between the components. Lastly, the calibration functionality is presented.

5.1. Model Overview

The overall architecture of the proposed HD calibration model is shown in Figure 5.1. The model consists of two major components.

The first part, the basis HD model, is reused from El-Sewisy (2020). It mimics the PI process by integrating angular velocity input to estimate HD. The position of the emerging activity peak indicates the current HD estimation. This activity peak is shifted clockwise or counterclockwise according to the angular velocity input.

The second component is the newly proposed calibration circuit. The functionality is inspired by Arleo and Gerstner (2000) and Bicanski and Burgess (2016), which are described in Section 3.1. However, instead of associating one ECD with one HD in one specific position of the agent, the idea is to associate information about the agent's position with ACD. The ACD describes the constant relation between HD and ECD. The visual detector stimulates the ECD cells when the agent perceives a visual cue. The CCF 1 facilitates the constant computation of ACD when a visual landmark is in sight according to the three-layer model from McNaughton et al. (1995) (see Section 3.2). ACD cells and ECD cells again project together onto the CCF 2 which in turn calculates the corresponding HD feedback signal. This mechanism is a reverse adaption of the three-layer model of McNaughton et al. (1995). The position encoder (PE) is active when the visual cue is in sight and performs two different tasks. It either reads out the ACD by decoding the activity in the ACD cells or resets the firing activity in the ACD cells corresponding to a desired ACD. When the activity profile of the ACD cells is reset, the ECD cells and the CCF 2 account for resetting the HD cells accordingly.

The PE can process ACDs in two different modes - with or without first glance learning (FGL). In both modes, the PE still accounts for reading out and resetting the activity of the ACD cells. However, the two modes change the way of how the PE processes ACDs that it read out from the ACD cells to compute ACDs for resetting the ACD cells.

Without FGL, the PE associates ACD information to small place fields around a visual

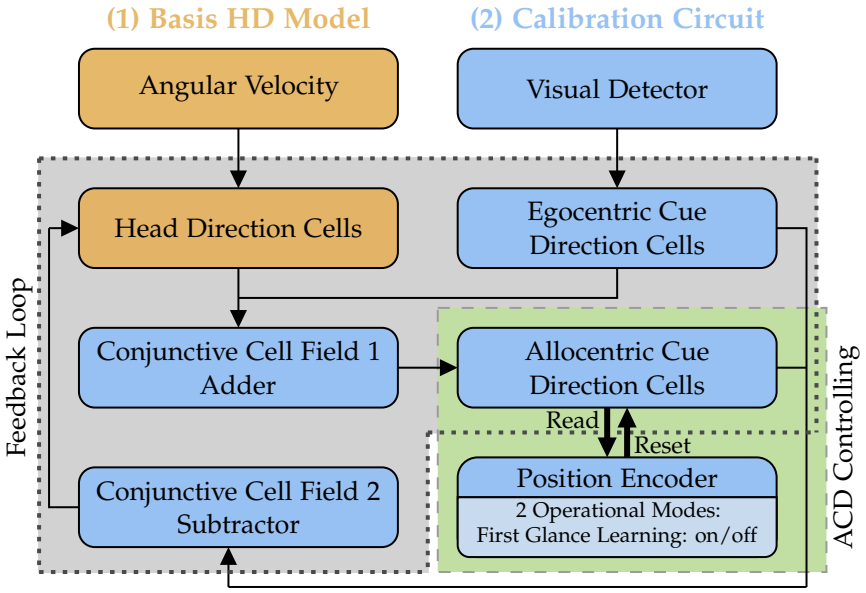

Figure 5.1.: Architecture of the proposed HD calibration model: (1) The basis HD model (orange boxes) indicates the current estimation of HD by integrating angular velocity input via the HD cells. (2) The calibration circuit (blue boxes) facilitates learning and exploiting directional information from visual landmarks. The calibration circuit has two simultaneously running mechanisms: the feedback loop and the ACD controlling. The feedback loop is active when the visual detector input is active. In this case, the feedback loop accounts for creating ACD information by combining HD with ECD information via the CCF 1. Additionally, the feedback loop creates the feedback signal back onto the HD cells by combining ACD and ECD information via CCF 2. The ACD controlling part is composed of the ACD cells and the PE. The PE has two functions. On the one hand, it accounts for reading out and storing the ACD from the ACD cells. On the other hand, the PE resets the ACD cells' activity corresponding to the current position of the agent. The PE offers two separate modes that affect how the read ACD information is processed before resetting: FGL on and off.

landmark. When the agent perceives the landmark, the ACD information encoded by the ACD cells is associated with the place field that is currently occupied by the agent. When revisiting a place field and perceiving the landmark, the associated ACD information is restored to reset the ACD cells' activity. In turn, the HD signal is calibrated via the CCF 2.

When FGL is activated, the PE stores only information at the moment of the first glance the agent has on a visual landmark. The PE stores the encoded ACD, the distance of the cue, and the agent's position. Afterward, the model is able to correct for HD errors whenever perceiving the landmark again, no matter from which position. Therefore, the PE performs a vector addition. The vector describing the direct way from the agent's current position to the stored position from the first glance is added to the vector that is built by combining the stored information about ACD and landmark distance. The direction of the resulting vector describes the estimation of ACD in the agent's current position and is used to reset the activity profile in the ACD cells. Subsequently, the HD cell activity is calibrated by the CCF 2.

Next, the components of the proposed HD calibration model and their operation functionality are described. Afterward, the computation of the connections between these components is illustrated, followed by a detailed explanation of the proposed calibration processes.

5.2. Model Components

At first, the used neuron model is explained followed by the illustration of the calibration circuit components, shown in Figure 5.1, and their operation functionality.

5.2.1. Neuron Model

Spiking neurons serve as the basic building elements for the HD model of El-Sewisy (2020). All components of the calibration circuit are based on the same spiking neuron model.

Networks consisting of spiking neurons closely mimic the behavior of biological neural networks (Wijngaarden et al. 2020). Each neuron's activity is determined by its firing rate. It only depends on the summation of all input firing rates from other neurons. This behavior is referred to as "integration and fire" (Wijngaarden et al. 2020). The concrete computational model to mimic one single neuron is adapted from Pereira and Brunel (2018) that was built on recently recorded neuron data. The differential equation describing a neuron's activity is the following (Pereira and Brunel 2018):

$$\tau \frac{df_i}{dt} = -f_i + \Phi(I_i + \sum_j w_{ij} f_j) \tag{5.1}$$

with:

τ = time constant of firing rate dynamics
f_i = firing rate of neuron i (Hz)
f_j = firing rate of neuron j (Hz)
Φ = single-neuron transfer function
I_i = external inputs to neuron i
w_{ij} = connection (synaptic) weight from neuron j to neuron i

The single neuron transfer function $\Phi(x)$ and its shape defining parameters are reused from El-Sewisy (2020) and Pereira and Brunel (2018). It is defined as follows:

$$\Phi(x) = \frac{r_m}{1 + exp(-\beta_t(x - h_0))} \tag{5.2}$$

with:

$r_m = 76.2\,\text{Hz}$
$\beta_t = 0.82$
$h_0 = 2.46$
$\tau = 20\,\text{ms}$

Pereira and Brunel (2018) chose the parameters so that the sigmoid function matches the recorded data from neurons.

According to Equation 5.1 the firing rate of a neuron always converges to the sigmoid term $\Phi(I_i + \sum_j w_{ij} f_j)$. This is explained by thinking of the following two scenarios:

- If f_i is larger than the sigmoid term, $\frac{df_i}{dt}$ is negative. This results in a decreasing firing rate f_i.
- If f_i is smaller than the sigmoid term, $\frac{df_i}{dt}$ is positive and the firing rate f_i increases.

5.2.2. Basis Head Direction Model

As already mentioned, this thesis is based on the PI HD model of El-Sewisy (2020). Its architecture, components, and functionality will be described in this section. However, only the most important facts are included to provide a general understanding of the model. For more detailed information, it is referred to El-Sewisy (2020).

5.2.2.1. Architecture

The general of the basis HD model's architecture is shown in Figure 5.2. It consists of three components: one HD attractor and two shift layers (SLs). Each of these consists of n neurons. Figure 5.2 shows an exemplary architecture for $n = 8$ neurons. For simplicity, the layers are presented as a line of neurons. Actually, they form a ring so that the first and the last neuron in each line are again neighbored neurons. The position of the activity peak in the HD attractor indicates the current estimation of HD. The SLs receive external stimuli depending on the agent's current angular velocity which causes the activity profile in the HD attractor to shift in line with the agent's change of orientation. In other words, the model integrates angular velocity to infer HD.

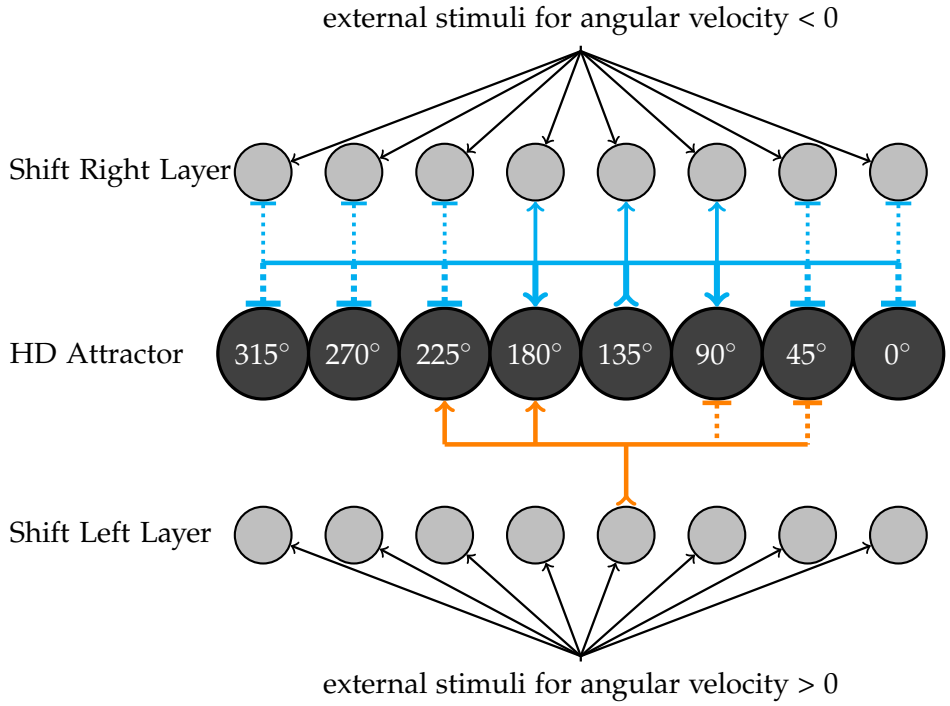

Figure 5.2.: Architecture of the basis HD model: The HD attractor neurons have symmetric connections to each other and towards the shift layers. The blue lines show the example connectivity outgoing from one HD cell representing 135 °. An arrow at the tip of a connection indicates an excitatory connection. A perpendicular line indicates an inhibitive connection. The line thickness indicates the connection strength. The connections to close neighbors in the HD attractor and both shift layers are excitatory (arrowhead). More distant neurons have inhibitory connections (perpendicular bar). The connection weights towards the shift layers are halved. The connections towards the shift left layer are the same as towards the shift right layer but are omitted in the drawing. These connections lead to a Gaussian shaped activity peak in the HD attractor that encodes the current HD. The shift layers have both asymmetric connections towards the HD attractor. The orange lines show the connectivity pattern for one shift left layer neuron. It has excitatory connections to the close left HD cells (arrow) and inhibitive connections to the close right HD cells (perpendicular bar). The shift right layer has excitatory connections to the close right neighbors and inhibitive connections to the close left neighbors. Both shift layers are artificially stimulated according to an agent's current angular velocity. When the external stimuli of the shift layers are the same, their input towards the HD cells cancels out. If the stimuli of one shift layer exceed the other, the input into the HD attractor shifts its activity accordingly to the left or right.

5.2.2.2. Head Direction Attractor

The HD attractor neurons are numbered with $i \in \{0, ..., n-1\}$. Each neuron in the HD attractor corresponds to a certain orientation of the agent within its environment. This is called a neuron's preferred firing direction (PFD) θ. Each neuron's PFD is given by Equation 5.3:

$$\theta_i = \frac{2\pi i}{n} \tag{5.3}$$

Thereby, the set of neurons equally covers the whole circle of 360 °. The closer the orientation of the agent is to a neuron's PFD, the higher is the firing rate of the neuron. The angular distance between the actual HD θ_A of the agent and the PFD θ_i of neuron i can be determined by:

$$\Delta\theta_{iA} = \theta_A - \theta_i \tag{5.4}$$

The firing of each neuron i according to the current orientation of the agent can be inferred by Equation 5.5 which uses the parameters according to El-Sewisy (2020).

$$f(\Delta\theta_{iA}) = A + Be^{Kcos(\Delta\theta_{iA})} = 1.72 + 0.344e^{5.29cos(\Delta\theta_{iA})} \tag{5.5}$$

with:

f = firing rate at angle difference $\Delta\theta_{iA}$ (Hz)
A = lowest firing rate (Hz)
Be^K = peak firing rate (Hz)
K = sharpness parameter of the peak curve

All neurons in the HD attractor are reciprocally connected to each other and, thereby, form a well known CAN. Furthermore, each HD cell is connected to all neurons of the SLs. The connection weights among the HD cells are calculated based on the angular distances between the PFDs of two HD cells. The closer two neurons are in the ring, the smaller is the difference of their PFD and the larger is the connection weight between them. The larger the angular distance between two neurons (up to 180 °) is, the smaller is the weight of their connection. Close neurons have excitatory connections which means that their connection weights are greater than zero. Distal neurons even have inhibitive connections. The weights of these connections are smaller than zero. If there is no activity in any HD cell, the system is in an unstable equilibrium. As soon as a little stimulus is applied to a neuron, a stable activity peak emerges around the excited HD cell. The HD cells' activity profile is maintained via the reciprocal connections within the HD cell layer, even after the external stimulus is switched off. The connections from the HD cells toward the SL neurons are similar. Each HD cell is connected to every shift layer neuron. The only difference is that the weights of the connections are halved compared to the connections among the HD cells. That means that a stable activity peak in the HD attractor leads to activity peaks in both shift layers, halved in their magnitude compared to the activity peak in the HD cells.

5.2.2.3. Shift Layers

The SLs aim to shift the activity peak in the HD attractor according to the rotation movements of the agent so that the activity peak in the HD attractor consistently corresponds to the current HD of the agent. Each neuron in the SLs has one corresponding partner neuron in the HD attractor. Corresponding neurons are drawn on top of each other in Figure 5.2. That means that always two SL neurons are associated with one HD cell. The weights of the connection of both SLs onto HD cells are mirrored. Therefore, the explanation focuses on the connections of the shift left layer (SLL) and, afterward, the idea is transferred for the shift right layer (SRL).

Each neuron in the SLL has connections to all HD cells. The connections are again set up on the angular distance between two neurons like it was done with the angular distance between two PFDs in the HD attractor. The function describing the connection weights from the SLs to the HD cells is the derivative of the function describing the connection weights among the HD cells. That means, that the connection weights also rely on the algebraic sign of the angular distance. As the shift left layer is responsible for moving the activity peak in the HD cells to the left (counterclockwise), the connections from a SLL neuron towards the close HD cells to its left are strongly exciting. In contrast, the connections from a SLL neuron towards HD cells to its right are strongly inhibitive. The further away two neurons are, the weaker their connection is. The weight between a SLL neuron towards the corresponding partner neuron with angular distance zero is set to zero. Whenever a neuron in the SLL is active, it excites the HD cells to its left side and inhibits the ones to the right (see the orange connections in Figure 5.2). This is done the same way for the SRL. The only difference is that the weights are mirrored. The connections from the SRL towards the close neighbored HD cells on the right are exciting. Accordingly, the continuous connection function is the negative derivative of the function describing the connection weights among the HD cells.

5.2.2.4. Functionality

Imagine that the system is initialized with HD $\theta_A = 0\,^\circ$ and no external input into the SLs is applied. In this case, an activity profile emerges in all three rings with the peak at the neurons corresponding to $0\,^\circ$. However, the peak's magnitude in the SLs is half of the peak's height in the HD attractor. In this state, the activity in both SLs is identical. Thus, the input from the SLs to the HD cells cancels out. When the agent starts to rotate, an external input is generated according to the angular velocity of the agent and, subsequently, applied to the SLs. If the agent rotates to the left, the external input for the SRL is set to zero. In contrast, the SLL is excited in relation to the angular velocity of the agent. At this moment, the activity in the SLL exceeds the activity of the shift right layer. This leads to a stronger input to the HD layer from the SLL than from the SRL which results in a (counterclockwise) left shift of the activity peak in the HD layer. Hence, the activity peak in the HD layer still encodes the agent's current orientation within the environment. The same process happens when the agent turns to the right.

The SLL receives zero input. Instead, the SRL is excited according to the agent's rotation speed. The resulting inequality of the SLs' input to the HD layer causes the HD activity profile to shift to the right (clockwise).

5.2.3. Calibration Circuit

In the following part, every single component of the calibration circuit is described. This contains the components' structures and also their functionality, more precisely their firing behavior.

5.2.3.1. Visual Detector

The visual detector is an abstract block governing the whole visual perception process of a visual cue in the environment and is no neural model. It solely represents the visual input signal that drives the ECD cells. This is based on the assumption that animals use extracted visual features of landmarks instead of the whole configuration of a landmark according to the distal sensory input model described in Section 2.4.2.1. This assumption is reused from already proposed calibration models described in Section 3.1.

5.2.3.2. Egocentric Cue Direction Cells

ECD cells represent the egocentric bearing of a visual landmark, as described in Section 2.3.2. Each neuron's activity depends on the bearing of a visual landmark relative to the agent's orientation and position within the environment. n ECD cells are equally distributed around a circle to cover all 360 °. They are numbered with $i \in \{0, ..., n-1\}$. α_i describes each neuron's preferred ECD in the interval $[0, 2\pi)$. Hence, the preferred ECD of neuron i is described by Equation 5.6.

$$\alpha_i = \frac{2\pi i}{n} \tag{5.6}$$

The angular distance $\Delta\alpha_{il}$ between a landmark's angular position in relation to the orientation of the agent α_l and an ECD cell's preferred ECD α_i determines the firing activity $f(\Delta\alpha_{il})$ of neuron i. The angular difference is expressed by Equation 5.7.

$$\Delta\alpha_{il} = \alpha_l - \alpha_i \tag{5.7}$$

The firing activity of every neuron is governed by Equation 5.8 which is adapted from El-Sewisy (2020).

$$f(\Delta\alpha_{il}) = A + Be^{Kcos(\Delta\alpha_{il})} = 1.72 + 0.344e^{5.29cos(\Delta\alpha_{il})} \tag{5.8}$$

with:

f = firing rate at angle difference $\Delta\alpha_{il}$ (Hz)
A = lowest firing rate (Hz)
Be^K = peak firing rate (Hz)
K = sharpness parameter of the peak curve

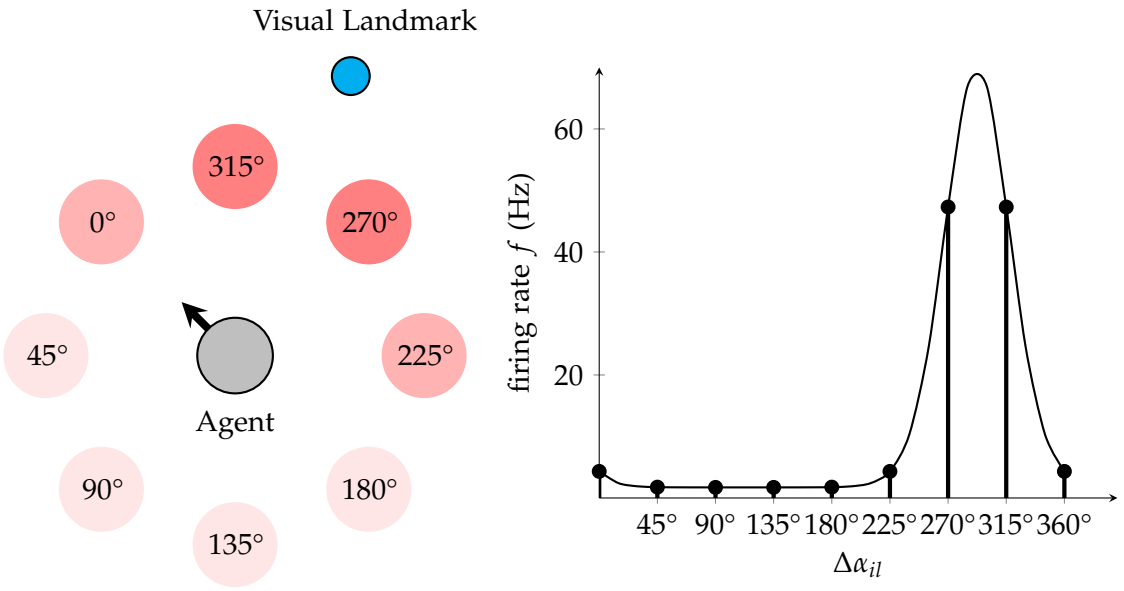

(a) Example spatial relation between agent and visual landmark

(b) Activity peak function

Figure 5.3.: Example of a firing profile of ECD cells: (a) shows a ring of ECD cells and an example relation between an agent and a visual landmark. The corresponding preferred ECD of each ECD cell is fixed to the orientation of the agent. The ECD cells' activities encode the egocentric perceived angle of the visual cue from the agent's pose. (b) shows the exact firing rate of each ECD cell according to its represented angle. In this case, the neurons corresponding to 270 ° and 315 ° are the most active.

Equation 5.8 is the desired continuous target peak function. The peak firing rate and the background firing rate are defined as 70 Hz and 1.72 Hz to mimic the target activity profile of the HD cells of El-Sewisy (2020). The highest firing rate arises at the neuron whose preferred ECD is closest to the angular bearing of the landmark. The lowest activity appears at the ECD cell that is the furthest away from the landmark's position. An example firing behavior of the ECD ring is illustrated in Figure 5.3. In this example, the ECD cell ring consists of $n = 8$ cells that are drawn in a circle around the agent in Figure 5.3a. Every neuron has its fixed position in relation to the orientation of the agent. The neuron $i = 0$ corresponds to "ahead" of the agent (0 °) independent of the agent's orientation. Neuron $i = 2$ represents the agent's left hand side (90 °). An exemplary visual landmark is placed at an egocentric bearing of $\alpha_l = 292.5\,°$ in Figure 5.3a. Therefore, the preferred ECDs of the neurons $i = 6, 7$ are closest to the bearing of the landmark, or rather the angular distances are the smallest. Thus, these neurons are the most active ones. The further the other neurons are away, the less active they are. This relation is shown in Figure 5.3b. The neurons $i = 6, 7$ correspond to 270 ° and 315 ° and fire with ≈ 47 Hz. Hence, the firing behavior of the whole ring encodes the

egocentric bearing of the blue visual cue by discretizing the continuous target peak function. The more neurons are used, the less is the distance between the samples of the function.

5.2.3.3. Allocentric Cue Direction Cells

ACD cells encode the bearing of a visual landmark like ECD cells but in an allocentric way, as described in Section 2.3.3. Each neuron's activity depends on the bearing of a visual landmark solely based on the agent's allocentric position and independent of its orientation. The neuron population consists of n cells that are equally distributed around the circle to cover all 360 ° of directions. They are numbered again with $i \in \{0, ..., n-1\}$. β_i describes each neuron's preferred ACD around the circle in the interval $[0, 2\pi)$. Hence, the preferred ACD of neuron i is described by Equation 5.9.

$$\beta_i = \frac{2\pi i}{n} \tag{5.9}$$

The angular distance $\Delta\beta_{il}$ between a landmark's ACD β_l, in reference to the agent's position, and an ACD cell's preferred ACD β_i determines the firing activity $f(\Delta\beta_{il})$ of neuron i. Therefore, the Equation 5.10 describes the angular distance $\Delta\beta_{il}$ similar to Section 5.7 for the ECD ring.

$$\Delta\beta_{il} = \beta_l - \beta_i \tag{5.10}$$

The functions describing the activity profile for the ACD cells and the ECD cells are similar as well. However, the landmark direction is encoded in allocentric coordinates (see Equation 5.11), which is illustrated in Figure 5.4.

$$f(\Delta\beta_{il}) = A + Be^{Kcos(\Delta\beta_{il})} = 1.72 + 0.344e^{5.29cos(\Delta\beta_{il})} \tag{5.11}$$

with:

f = firing rate at angle difference $\Delta\beta_{il}$ (Hz)
A = lowest firing rate (Hz)
Be^K = peak firing rate (Hz)
K = sharpness parameter of the peak curve

For instance, $n = 8$ ACD cells are aligned in a circle around the agent (see Figure 5.4a). Each ACD cell's preferred ACD is not tied to the agent's orientation but fixed to the world coordinates. The neuron $i = 0$ always corresponds to $\beta_i = 0\,°$ which can be referred to as the compass direction "west" of the agent's position. $\beta_i = 180\,°$ represents the direction "east". Therefore, the visual landmark is located at $\beta_l = 67.5\,°$ in allocentric coordinates. Hence, the neurons $i = 1, 2$, which correspond to 90 ° and 45 °, have the smallest angular distance to the visual cue's angular position and are the most active neurons with a firing rate of ≈ 47 Hz (see Figure 5.4b). The other neurons have a larger angular distance to the landmark and are less active.

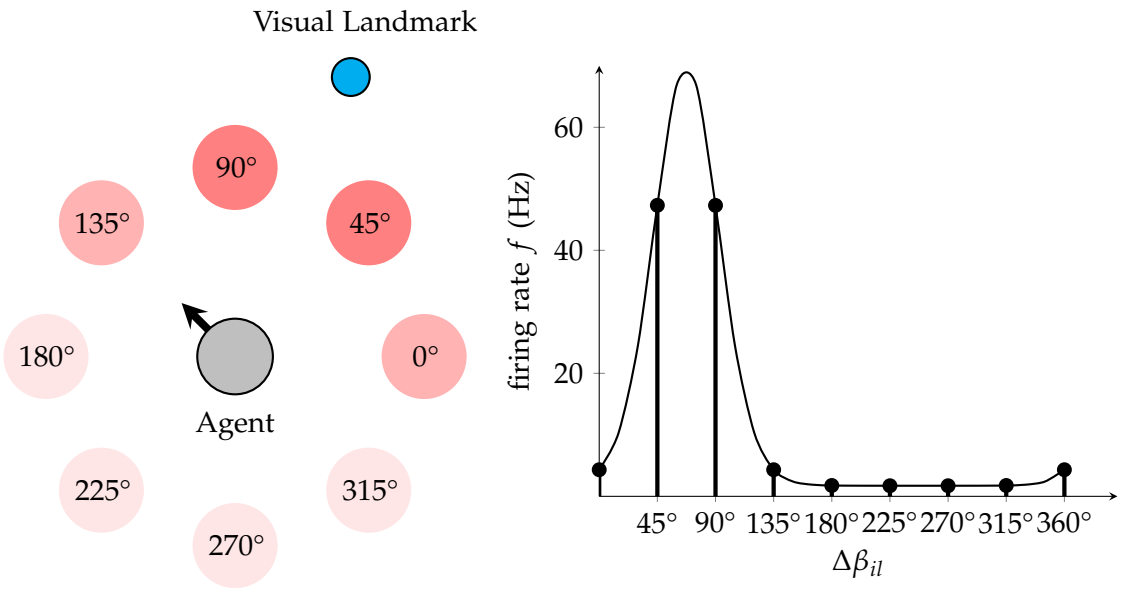

(a) Example: spatial relation between agent and visual landmark

(b) Activity peak function

Figure 5.4.: Example of a firing profile of ACD cells: (a) shows a ring of ACD cells and an example relation between an agent and a visual landmark. The preferred ACDs of the ACD cells are fixed to the world coordinates and, therefore, independent of the orientation of the agent. The ACD cells' activities encode the allocentric angle of the visual cue from the agent's position. (b) shows the exact firing rate of each ACD cell according to its represented angle. In this case, the neurons corresponding to 45 ° and 90 ° are the most active ones.

5.2.3.4. Conjunctive Cell Field 1: Adder

The conjunctive cell field (CCF) 1 consists of n conjunctive cells which are aligned in a two-dimensional grid. Figure 5.7 shows an example setup for $n = 36$ neurons. It mimics the cells, which were found by Wilber et al. (2014) and are illustrated in Section 2.3.4. The neurons are numbered with $i \in \{0, ..., n-1\}$. Each row and each column consists of $\sqrt{n}$ neurons. The numbering is done from left to right and row by row, starting from the top row. Each neuron is associated with two input variables: the HD of the agent in the environment and the ECD of a visual landmark. The represented ECD α_i and HD θ_i of neuron i are governed by Equations 5.12 and 5.13.

$$\alpha_i = \frac{2\pi}{\sqrt{n}}(i \mod \sqrt{n}) \tag{5.12}$$

$$\theta_i = \frac{2\pi}{\sqrt{n}}(\sqrt{n} - 1 - (i \text{ div } \sqrt{n})) \tag{5.13}$$

The closer the HD and the ECD are to a neuron's associated values, the more active it is. The activity function for each neuron is given by Equation 5.14, which is adapted from the one-dimensional activity function from El-Sewisy (2020). The angular distances $\Delta\alpha_i$ and $\Delta\theta_i$ are reused from the ECD cells and HD cells defined in Sections 5.2.3.2 and 5.2.2.2.

$$f(\Delta\alpha_{il}, \Delta\theta_{iA}) = A + Be^{0.5K(cos(\Delta\alpha_{il})+cos(\Delta\theta_{iA}))} = 0.0504e^{2.645(cos(\Delta\alpha_{il})+cos(\Delta\theta_{iA}))} \tag{5.14}$$

with:

f = firing rate at HD angle difference $\Delta\theta_{il}$ and ECD angle difference $\Delta\alpha_{il}$ (Hz)
A = lowest firing rate in (Hz)
Be^K = peak firing rate (Hz)
K = sharpness parameter of the peak curve

The lowest and the peak firing rate are set to 0 Hz and 10 Hz according to the measured values of Wilber et al. (2014). The sharpness factor of the peak curve stays the same as used by El-Sewisy (2020). This describes a two-dimensional activity peak that emerges in the CCF and combines information about HD and ECD at the same time. Figure 5.5 shows an exemplary activity peak. Each dimension of neurons is covering 360 ° of the corresponding variable - HD θ_A or ECD α_l. Here, the example values $\theta_A = 90°$ and $\alpha_l = 30°$ are chosen and, thus, define the peak position in the field. The further away a neuron is from this position, the less activity can be observed. The task of the conjunctive cells in the calibration mechanism is to serve as a bridge between egocentric and allocentric reference frame. In this case, the input information, ECD and HD, is used to compute the ACD of a landmark which is subsequently encoded by the ACD cells. Therefore, the CCF is also referred to as "Adder" component. Figure 5.6 shows an example relation between an agent and a visual landmark. The agent's orientation is $\theta_A = 90\,°$. The landmark is perceived at $\alpha_l = 30\,°$ from the agent's pose. The ACD

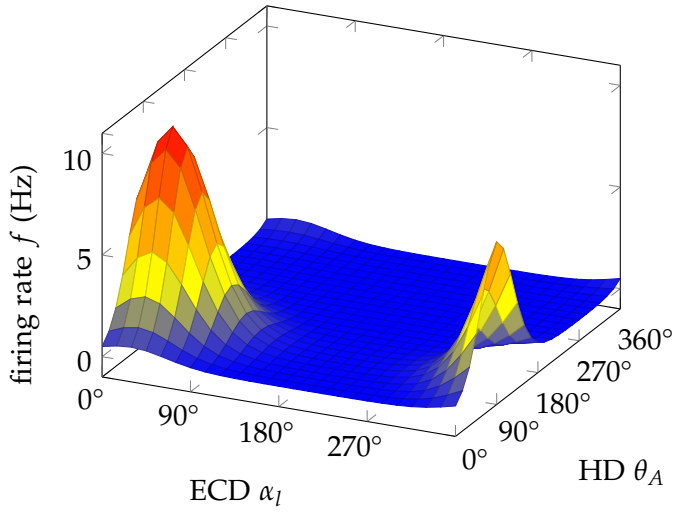

Figure 5.5.: Example of a 2D activity peak function at the CCF 1: The plot continuously shows the firing rate according to the current ECD and the HD of the agent. In this case, the peak position corresponds to $\theta_A = 90°$ and $\alpha_l = 30°$ and shows a maximum firing rate of 10 Hz.

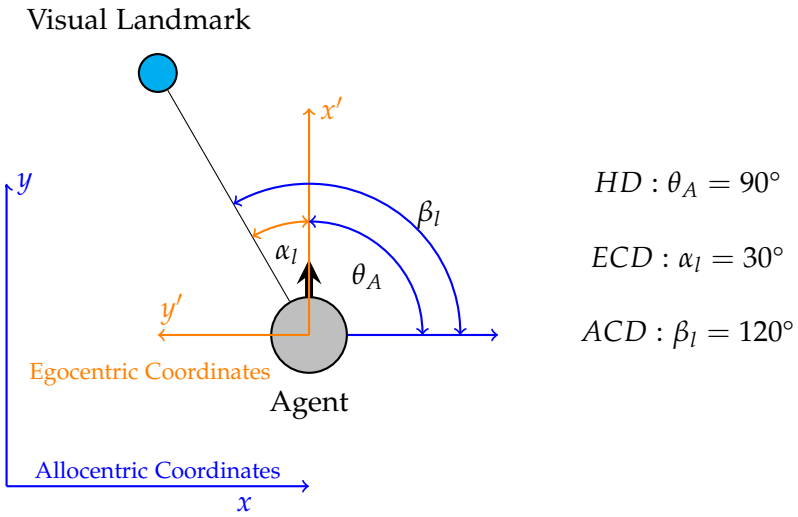

Figure 5.6.: Example spatial relation between agent and visual landmark to show the angular relationships between the pose of an agent and the position of a landmark. The ACD describes the summation of the HD and ECD.

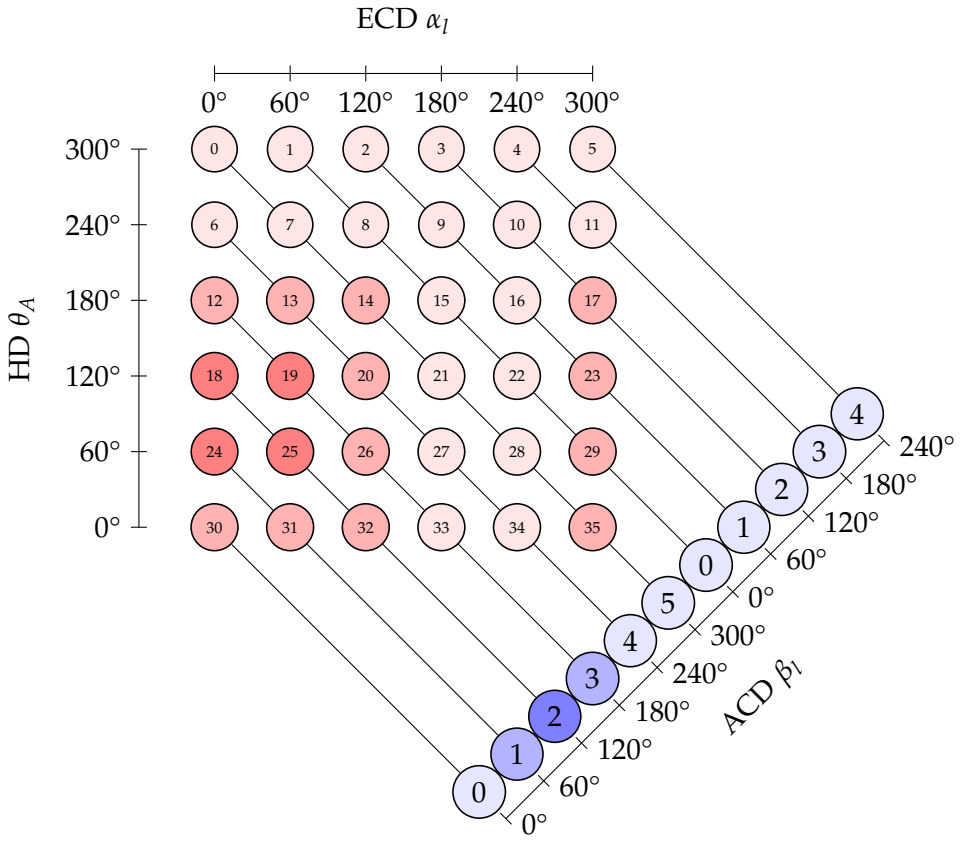

Figure 5.7.: Exemplary setup and activity of CCF 1 - Adder: The conjunctive cells are aligned in a two-dimensional grid. Each neuron's firing activity depends on the current HD and ECD. This leads to a two-dimensional activity peak. In this example case, the most active neurons are numbered with 18, 19, 24, and 25 according to the HD, ECD, and ACD from the example setup from Figure 5.6. The conjunctive cells drive an output ring of neurons that encodes ACD via diagonal connections. The activity peak that emerges in the output ring lies at the diagonal extended position of the two-dimensional activity peak from the conjunctive cells.

from the agent's position is $\beta_l = 120\,°$. The mathematical relation between these three variables is governed by Equation 5.15.

$$\beta_l = (\theta_A + \alpha_l) \mod 360\,° \tag{5.15}$$

This relation can be simply derived from Figure 5.6. Only the modulo operation is needed to restrict the ACD to the interval $[0\,°, 360\,°)$. This adding operation is performed by the CCF 1. Figure 5.7 explains this issue. The ECD ring gives its input via vertical connections onto all conjunctive cells of the CCF 1. The HD ring projects onto the CCF 1 via horizontal connections. Additionally, the CCF 1 drives the ACD ring via diagonal connections. The connection weights towards the ACD cells are described in Section 5.3.2.

In the example setup in Figure 5.7 $n = 6$ ACD cells are used. For simplicity, they are split up. For example, ACD cell number $i = 1$ is only one neuron and receives input from two diagonals. The conjunctive cells show a two-dimensional activity peak with the peak firing rate at the intersection of the ECD $\alpha_l = 30\,°$ and HD $\theta_A = 90\,°$. The red colored neurons illustrate the activity of the neurons. Each ACD cell integrates the accumulated neuron firing rates via the diagonal input connections. For example in Figure 5.7, the neuron number $i = 2$ of the ACD ring receives the largest input as it lies in the diagonal extension of the two-dimensional activity peak. Consequently, the neurons $i = 1, 3$ receive the second largest input. Finally, an activity peak emerges in the ACD ring that encodes the ACD of the landmark. In this case, the resulting activity peak encodes $\beta_l = 120\,°$. This equals the operation from 5.15. If the agent rotates while keeping its position, the two-dimensional activity peak will move only on the diagonal corresponding to $\beta_l = 120\,°$ and, therefore, will not change the encoded ACD. For example, the agent turns 10° to the left. This increases the encoded HD to $\theta_A = 100\,°$. In contrast, the ECD decreases to $\alpha_l = 20\,°$. Consequently, the ACD stays the same ($\beta_l = 120\,°$). If the agent keeps its orientation and moves upwards, the HD will stay constant while the ECD and, in parallel, the ACD will increase equally.

5.2.3.5. Conjunctive Cell Field 2: Subtractor

The CCF 2 consists of n neurons that are aligned in a two-dimensional grid like the Adder (CCF 1). The index $i \in 0, ..., n-1$ numbers all neurons in the field (see Section 5.8). In each row and every column are $\sqrt{n}$ neurons. The neuron field exhibits the same two-dimensional activity profile when stimulated via both dimensions like the Adder (CCF 1) as described in Section 5.2.3.4. In contrast, the purpose of the CCF 2 is to perform a subtraction to obtain a feedback signal for resetting the HD signal. Therefore, the CCF 2 takes ACD and ECD as input to generate the HD signal according to Equation 5.16.

$$\theta_A = (\beta_l - \alpha_l + 360\,°) \mod 360\,° \tag{5.16}$$

Therefore, each conjunctive cell's activity corresponds to a certain ECD α_i and ACD β_i. The represented ECD and ACD of neuron i can be obtained by Equations 5.17 and 5.18.

$$\alpha_i = \frac{2\pi}{\sqrt{n}}(\sqrt{n} - 1 - (i \bmod \sqrt{n})) \tag{5.17}$$

$$\beta_i = \frac{2\pi}{\sqrt{n}}(\sqrt{n} - 1 - (i \operatorname{div} \sqrt{n})) \tag{5.18}$$

For instance in the example setup in Figure 5.8, neuron $i = 1$ corresponds to the ECD $\alpha_i = 240\,^\circ$ and the ACD $\beta_i = 300\,^\circ$. The angular distances from the current ECD α_l and ACD β_l are computed as shown by the Equations 5.7 and 5.10. These angular distances determine the activity of neuron i according to Equation 5.19.

$$f(\Delta\alpha_{il}, \Delta\beta_{il}) = A + Be^{0.5K(cos(\Delta\alpha_{il})+cos(\Delta\beta_{il}))} = 0.0504e^{2.645(cos(\Delta\alpha_{il})+cos(\Delta\beta_{il}))} \tag{5.19}$$

with:

f = firing rate at ACD angle difference $\Delta\beta_{il}$ and ECD angle difference $\Delta\alpha_{il}$ (Hz)
A = lowest firing rate (Hz)
Be^K = peak firing rate (Hz)
K = sharpness parameter of the peak curve

In this example, the neurons $i = 18, 19, 24, 25$ are the most active for an ECD of $\alpha_l = 270°$ and an ACD of $\beta_l = 90°$. Figure 5.8 shows how this operation can be realized with an example setup of $n = 36$ neurons. All neurons of the CCF 2 are diagonally connected to the HD ring. For the explanation of this operation, the HD ring consists of $n = 6$ neurons and only receives input from the conjunctive cells to show which feedback signal arises. In the example (see Figure 5.8), the ECD direction is set to $\alpha_l = 270\,^\circ$ and ACD to $\beta_l = 90\,^\circ$. Thus, the HD cell $i = 3$ receives the largest input as it lies on the diagonal extension of the two-dimensional activity peak of the conjunctive cells and, therefore, is the most active HD cell. In conclusion, this setup of neurons is capable of calculating the HD feedback signal with given ECD and ACD.

5.2.3.6. Position Encoder

The PE encodes the local allocentric position of the agent in relation to the position of the visual landmark and, additionally, interacts with the ACD cells. Firstly, it reads out the encoded ACD from the ACD cells and associates it with the allocentric position of the agent. Secondly, it restores ACD information according to the agent's position and, subsequently, resets the firing activity in the ACD cells. It has two separate operational modes: with and without first glance learning (FGL). Both modes and their interplay with the ACD cells are described in detail in Section 5.4.2.

However, it has to be noted that the PE is not modeled in a biologically plausible way using spiking neurons. Its functionality is only abstractly defined.

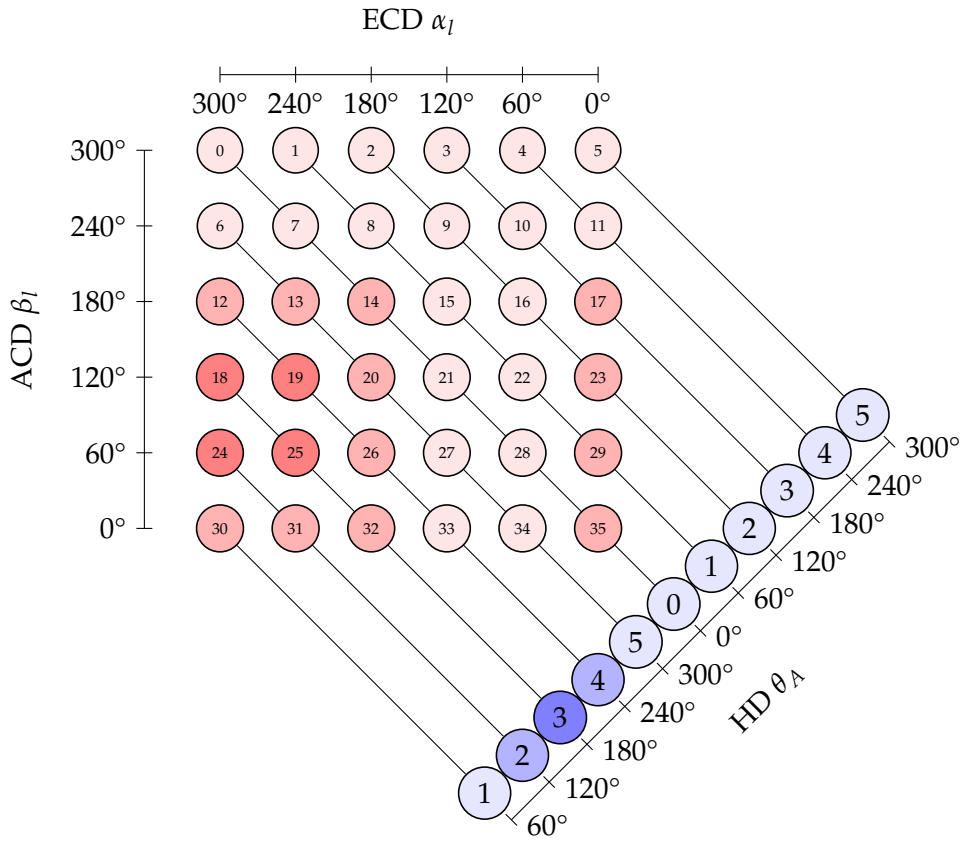

Figure 5.8.: Example activity of the CCF 2 - Subtractor: The conjunctive cells are aligned in a two-dimensional grid. Each neuron's firing activity depends on the current ACD and ECD. In this case, the emerging two-dimensional activity peak's position corresponds again to the example setup shown in Figure 5.6 with ACD $\beta_l = 90°$ and ECD $\alpha_l = 270°$. The conjunctive cells drive an output ring of neurons that encodes HD via diagonal connections. The activity peak that emerges in the output ring lies at the diagonal extended position of the two-dimensional activity peak from the conjunctive cells - in this case at HD $\theta_A = 180\,°$.

5.3. Connection Weights

The connection weights between the components of the calibration circuit and the interface connections between the calibration circuit and the basis HD model are calculated similarly to El-Sewisy (2020). El-Sewisy (2020) aimed to obtain a stable activity peak in the HD cell ring that is shaped as described by Equation 5.5. Therefore, the input current u_i at each post-synaptic neuron i is determined by summing up the input firing rates f_j from the pre-synaptic neurons j multiplied with their corresponding connection weights w_{ij}. Hence, the connection weights need to be adjusted such that the desired

input current u_i is reached. The difference between the desired input current u_i and the actual input current at neuron i is governed by Equation 5.20.

$$\Delta u_i = u_i - \sum_{j=0}^{n-1} w_{ij} f_j \tag{5.20}$$

To evaluate the difference of desired and obtained input currents for a whole ring of reciprocally connected neurons, El-Sewisy (2020) used vectors that discretely describe the input currents $\hat{u}_a$, firing rates $\hat{f}_a$ and weights $\hat{w}_a$ within one ring of neurons in the Fourier domain. Therefore, El-Sewisy (2020) followed the transformation procedure of K. Zhang (1996) that derived a continuous error function in the Fourier domain. The final function describing the difference between the desired input currents and obtained input currents in the whole ring of neurons is shown by Equation 5.21. Accordingly, the optimal Fourier coefficients of the weight distribution can be inferred by minimizing this function. Note that the Equation 5.21 additionally contains a regularization term to control the flatness, via λ, and the smallness of the weight distribution function.

$$E = \sum_{a=0}^{n-1} |\hat{u}_a - \hat{w}_a \hat{f}_a|^2 + \lambda \sum_{a=0}^{n-1} |\hat{w}_a|^2 \tag{5.21}$$

with:

E = error measurement value
$\hat{u}_a$ = a_{th} Fourier coefficient of the input current vector
$\hat{f}_a$ = a_{th} Fourier coefficient of the firing rate vector
$\hat{w}_a$ = a_{th} Fourier coefficient of the connection weight vector
λ = regularization factor controlling the flatness of the weight function

The general solution that minimizes the error function (see Equation 5.21) in the Fourier domain is given by Equation 5.22.

$$\hat{w}_a = \frac{\hat{u}_a \hat{f}_a}{\lambda + |\hat{f}_a|^2} \tag{5.22}$$

Equation 5.22 is again reused from El-Sewisy (2020) that adapted it from K. Zhang (1996). El-Sewisy (2020) used the n Fourier coefficients describing the desired input currents $\hat{u}_a$ and the desired firing rate $\hat{f}_a$ for each HD cells to solve Equation 5.22 for different values of λ. $\hat{u}_a$ and $\hat{f}_a$ can be obtained by computing the Fourier transformation of the desired input currents u_i of the post-synaptic component and the desired firing rates f_j of the pre-synaptic component. As the HD cells form a ring, the post-synaptic and pre-synaptic neurons are the same. Then, El-Sewisy (2020) calculated the error firing rate between the desired activity peak and the obtained activity peak in the HD cells using the resulting weight distributions that were calculated with different values of λ. Subsequently, the weight distribution that yields the smallest summation of error firing rates of the HD cells was chosen.

In this work, Equation 5.22 is used to calculate the weights between neurons of different components of the HD calibration model. Therefore, the pre-synaptic neurons are not anymore in the same neural component as the receiving, post-synaptic, neurons, like for example, the connection from a neuron in the ECD ring towards a neuron in the CCF 1 (see Figure 5.1).

In this case, the firing rate of the pre-synaptic neurons f_j and the desired input currents at the post-synaptic neurons u_i are separately chosen according to the desired activity profiles in both components. According to Equation 5.1, the post-synaptic firing rate f_i of each neuron i always converges to the term $\phi(\sum_{j=0}^{n-1} w_{ij} f_j)$ under the assumption that the external input I_i is zero, and the weights w_{ij} and the pre-synaptic firing rates f_j are constant. Therefore, the desired firing rate f_i of each post-synaptic neuron can be used to calculate the corresponding desired input currents u_i according to $u_i = \phi^{-1}(f_i)$ (Equation 5.2). Afterward, the corresponding Fourier coefficients $\hat{u}_a$ and $\hat{f}_a$ are computed. Then, for one given λ, Equation 5.22 has to be solved for all n values of a to subsequently obtain the resulting weight function by calculating the inverse of the discrete Fourier representation of the n Fourier coefficients $\hat{w}_a$. This procedure is done for a set of flatness parameters λ to obtain several weight functions. Then, the error firing rates between the desired firing activity and the obtained firing activity in the post-synaptic component for each weight function are calculated. $\frac{1}{\lambda}$ is used as regularization factor as small λ values lead to many undesirable wiggles in the weight function. Finally, the weight distribution that gives back the smallest sum of squared errors is used to set the connection weights between the pre- and post-synaptic component. Algorithm 1 shows the general calculation procedure for the weight functions.

Algorithm 1: Calculate Connection Weights

Input:

n: number of neurons
λ: flatness parameter
Λ: vector containing several flatness parameters λ
F: list containing the desired firing rates of n pre-synaptic neurons
U: list containing the desired input currents of n post-synaptic neurons
E: variable storing the smallest summation of squared error for λ

Output:

W_{opt}: list containing the optimal connection weights

initialize E with infinity
$E := \infty$;

compute fourier transforms
$\hat{F} := fft(F)$;
$\hat{U} := fft(U)$;

for λ *in* Λ **do**
    # calculate the fourier coefficients of W
    for $i = 0...n-1$ **do**
        $\hat{W}[i] = \frac{\hat{U}[i]\hat{F}[i]}{\lambda + |\hat{F}[i]|^2}$;
    end

    # compute the inverse fourier transformation
    $W := fft(\hat{W})$;

    # Compute the sum of squared errors and store it
    $X := \sum_{i=0}^{n-1}(U[i] - W[i]F[i])^2 + \frac{1}{\lambda}$;

    if $E > X$ **then**
        $W_{opt} := W$;
    end
end

return the weight list that face the smallest sum of errors
return W_{opt};

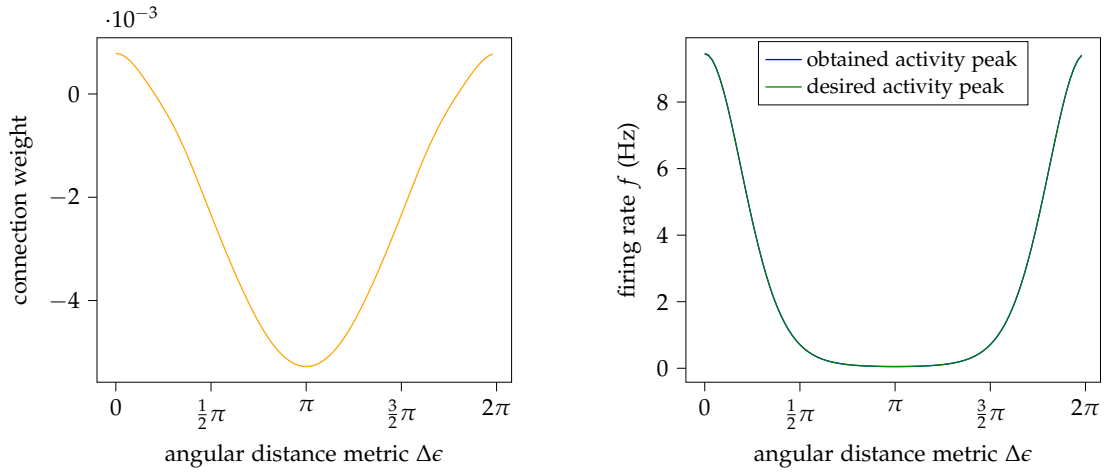

(a) Resulting connection weights

(b) Desired (green) and obtained (blue) peak of the activity band profile in the CCF

Figure 5.9.: Weight calculation results for connecting ECD cells 1 to CCF 1 and 2, and ACD cells to CCF 2. (a) shows the connection weights between two neurons depending on the difference between their preferred ECDs $\Delta\epsilon$. (b) shows the desired activity peak in the CCFs in green and the resulting activity peak when using the obtained connection weights from (a).

5.3.1. Egocentric Cue Direction Cells to Conjunctive Cell Fields 1 and 2, and Allocentric Cue Direction Cells to Conjunctive Cell Field 2

The target activity profiles in the ACD cells and the ECD cells are given by Equations 5.8 and 5.11. As both activity functions are the same, and both components, ECD cells, and ACD cells, project onto neurons in a CCF, the calculation of the weights is the same.

The input vector F for the weights calculation algorithm is obtained by sampling the activity peak function at n (number of neurons per ring) positions starting from 0 with increments of $\frac{2\pi}{n}$. The desired input current vector U for the conjunctive cells could be inferred by sampling the desired two-dimensional activity peak (Equation 5.14) and calculating U via $u = \phi^{-1}(f)$. As the CCF contains more neurons than the ECD cells and the ACD cells, this would lead to a vector U that is much larger than the input vector F. However, the weights calculation algorithm 1 only works with a firing rate vector F and an input current vector U with the same dimension. Therefore, the idea is to exploit the fact that the conjunctive cells receive the same inputs from two components. Hence, the target two-dimensional activity profile is divided into two one-dimensional activity profiles. As a result, if only one input component projects onto a CCF, a band profile will emerge instead of a two-dimensional punctured peak profile (see Figure B.1d). For example, if only the ACD cells project onto the neuron field in Figure 5.8 with a peak activity at 120 °, the most active neurons will be all neurons on the straight line from $i = 18$ to $i = 23$. All others would be less active. If the second input is present

and projecting via the second dimension onto the CCF, the two activity band profiles are super-positioned. This leads to the desired two-dimensional activity profile with the peak at the position where both activity band's maximum firing rates meet. The input currents that correspond to the desired one-dimensional activity band profile are obtained by sampling the two-dimensional activity peak profile described by Equations 5.14 and 5.19 n times along the diagonal cross-section. Therefore, the activity peak that emerges for $\theta_A = 0\,° \,\&\, \alpha_l = 0\,°$ needs to be sampled on the straight line between the point where $\theta_A = 0\,° \,\&\, \alpha_l = 0\,°$ and the point where $\theta_A = 360\,° \,\&\, \alpha_l = 360\,°$. Equation 5.23 describes this diagonal cross-section of the two-dimensional activity peak in one dimension using γ as angle variable.

$$f(\gamma) = Be^{Kcos(\gamma)} = 0.0504e^{5.29cos(\gamma)} \tag{5.23}$$

The function is sampled n times in increments of $\frac{2\pi}{n}$ starting from $\gamma = 0\,°$. Subsequently, the obtained list of firing rates is used to calculate the desired input currents via Equation 5.24. As the conjunctive cells receive input from two components, each input component should account for the half of the input currents for the two-dimensional activity peak in the conjunctive cells. Therefore, the factor 0.5 is used in Equation 5.24.

$$u = 0.5\,\phi^{-1}(f) \tag{5.24}$$

The obtained input currents are put into vector U and handed over to Algorithm 1. The calculated weight function is shown in Figure 5.9a. This one-dimensional function determines all connections from all neurons j from the pre-synaptic ACD cells and ECD cells to all neurons i from the receiving components (CCF 1 and 2). The corresponding optimal flatness parameter that minimizes the error function is $\lambda \approx 713.13$. $\Delta\epsilon$ describes the angular distance metric between the preferred ECD α_j of a pre-synaptic neuron and the post-synaptic conjunctive cell's preferred ECD α_i. If the preferred ECD of a pre-synaptic neuron j (ECD cell) and a post-synaptic neuron i (conjunctive cell) is the same, the connection weight is the highest, with a value of $\approx 0.785 \cdot 10^{-3}$. The same holds for the ACD. If the two neurons represent angles that are π apart from each other the weight value is $\approx -5.28 \cdot 10^{-3}$. The cross-section of the activity band profile that emerges at the CCFs, when only one neuron ring injects its information from one dimension, is shown in Figure 5.9b (blue graph). The green graph is the desired activity profile. Both profiles are almost congruent with a peak firing rate of $\approx 9.45\,\text{Hz}$ and $\approx 0.05\,\text{Hz}$ as the lowest firing rate.

5.3.2. Conjunctive Cell Field 1 to Allocentric Cue Direction Cells

To obtain the firing rate vector F for the calculation of the weights between the CCF 1 and the ACD cells, the two-dimensional activity peak in the CCF 1 needs to be created first. This is done by injecting the desired input current vector U which has been derived in Section 5.3.1 from both dimensions into the CCF. This yields the two-dimensional activity peak in the CCF. Then, the firing rates of all neurons lying on the same diagonal

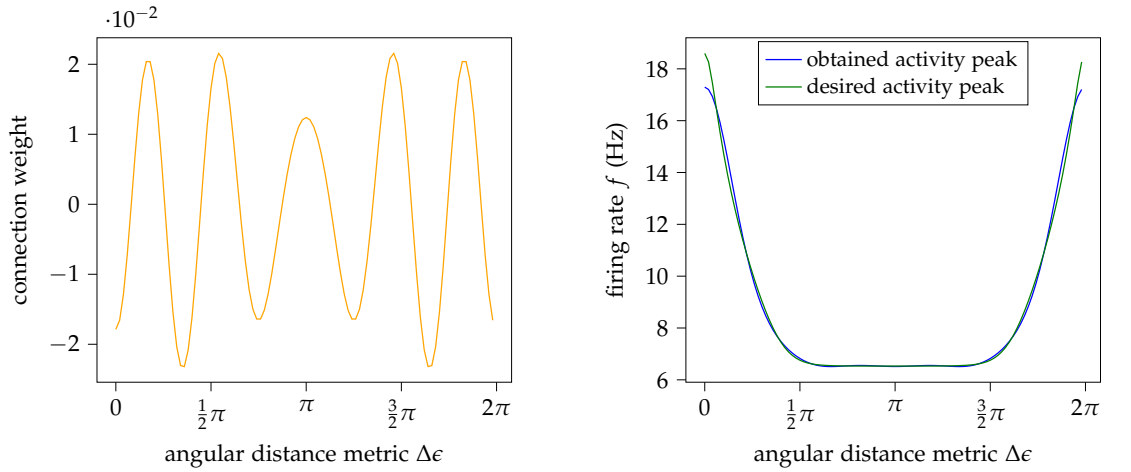

(a) Resulting connection weights

(b) Desired (green) and obtained (blue) activity profile at the ACD cells

Figure 5.10.: Weight calculation results for connecting CCF 1 to ACD cells. (a) shows the connection weights between two neurons depending on the distance of their preferred ACDs. (b) shows the desired activity peak in the ACD cells in green and the resulting activity peak when using the obtained connection weights from (a).

line towards the ACD cells are summed up and stored into the vector F. Hence, each element of F is corresponding to the sum of firing rates of one diagonal. The desired input current vector for the ACD cells U is calculated by, first, computing $u = 0.2\phi^{-1}(f)$ with the activity profile from Equation 5.11 and, afterward, sampling it n times. 0.2 is used as scaling factor to achieve a smaller activity peak in the ACD cells during the exploration phase when the model perceives the landmark, but is not calibrating yet. In this case, the feedback loop is active and, thereby, a calibration feedback signal is produced with a few time steps old HD and ECD information. By lowering the ACD firing in this phase, the feedback calibration signal affects the PI performance of the HD model less, but still exhibits an activity profile that encodes the ACD. The missing rest of $0.8 \times u$ is later artificially injected by the PE when the ACD cells' firing rates are reset. This yields the desired activity profile from 5.11. The obtained optimal weight function corresponds to $\lambda = 0.773$ and is shown in Figure 5.10a. The activity profile that is achieved using this weight distribution is shown by the blue graph in Figure 5.10b. It only slightly differs from the desired one (green graph).

5.3.3. Conjunctive Cell Field 2 to Head Direction Cells

Here, vector F, that contains the sum of firing rates on each diagonal of the conjunctive cells, is reused from the previous weight calculation in Section 5.3.2. The vector U contains the input currents that are calibrating the HD cell ring. To not bias the HD

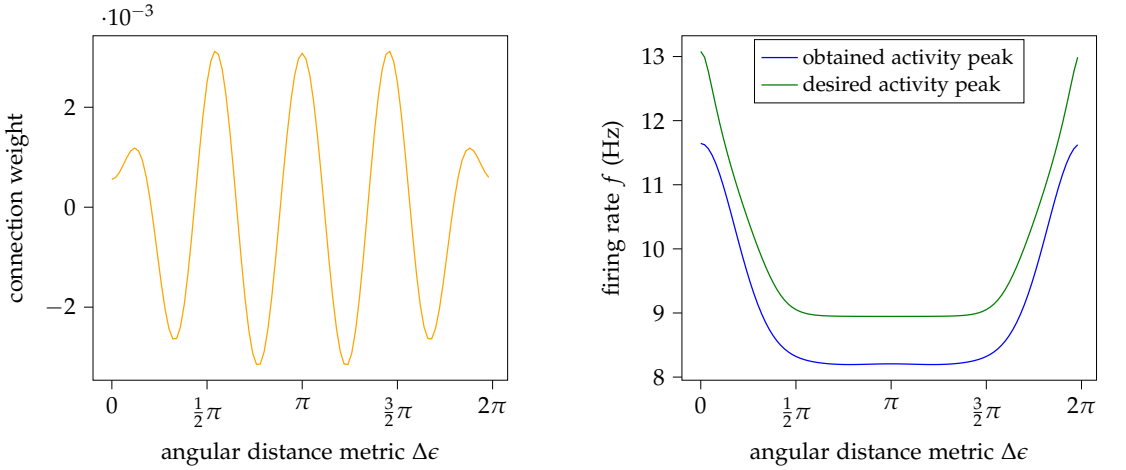

(a) Resulting connection weights

(b) Desired (green) and obtained (blue) activity profile of the feedback signal at the HD cells

Figure 5.11.: Weight calculation results for connecting CCF 2 to HD cells. (a) shows the connection weights between two neurons depending on their angular distance. (b) shows the desired activity peak in the HD cells in green and the resulting activity peak when using the obtained connection weights from (a).

activity profile too much, but to still be able to control the position activity peak, the input current function is set to $u = 0.1\phi^{-1}(f)$ with f as the target firing profile for the HD cell like in El-Sewisy (2020). The weights are then computed with $\lambda = 6.97$. Unfortunately, these weights lead to a faster accumulation of HD errors during PI when no landmark is in sight. This can be explained as follows. If the visual landmark is out of sight, all ECD cells will have constant small activity rates of $\approx 8.5\,\text{Hz}$. In contrast, the HD cells are still driven by PI and, therefore, introduce a persistent input to the calibration circuit. In turn, the ACD cells and CCF 2 show also a small uniform activity of $\approx 8.5\,\text{Hz}$ and $\approx 0.36\,\text{Hz}$) over all neurons. In turn, this leads to a constant uniform input current to the HD ring. This is not significantly affecting the shape of the activity peak in the HD cells. However, it slows down the shifting speed of the HD activity peak during a shift stimulus introduced by the SLs. Thus, the input currents need to be as close to zero as possible, when the visual landmark is out of sight. This is achieved by an offset adjustment of the connection weights. The whole weight distribution function is moved downwards so that the sum of all weights equals zero. If all neurons in the CCF 2 have the same activity, this will lead to zero input to the HD ring. Hence, the PI performance is not affected. The new weight distribution function is shown in Figure 5.11a. The final shape of the HD calibration signal that arises when using the adjusted weights is shown in Figure 5.11b. The obtained firing rate function (blue graph) has a constant negative offset compared to the desired activity from the beginning (green graph) but is still enough to control the firing profile at the HD cells.

5.3.4. Head Direction Cells to Conjunctive Cell Field 1

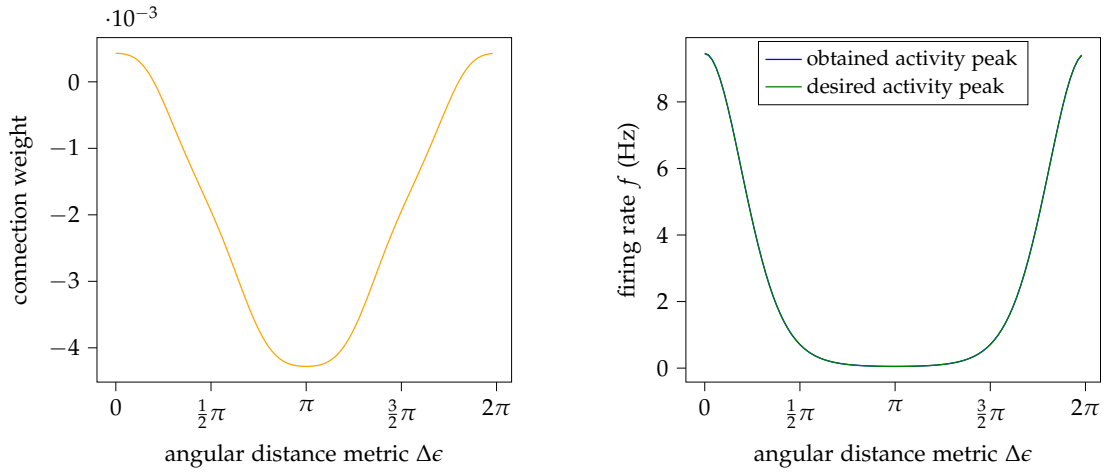

(a) Resulting connection weights

(b) Desired (green) and obtained (blue) cross-section of the activity profile in the CCF 1

Figure 5.12.: Weight calculation results for connecting HD cells to CCF 1. (a) shows the connection weights between two neurons depending on their angular distance. (b) shows the desired activity peak in the CCF 1 in green and the resulting activity peak when using the obtained connection weights from (a).

Lastly, the connections from the HD cell ring towards the CCF 1 are to determine. This connection is the second part that is needed to establish the two-dimensional firing profile at the CCF 1. Similar to El-Sewisy (2020), the HD activity profile without the calibration input from CCF 2 is not matching the desired activity peak function described by Equation 5.5. Consequently, the activity profile can not be inferred by summing the desired currents of the feedback signal to the target activity profile in the HD ring. This issue is caused by the reciprocal connections within the HD attractor network. Hence, the HD signal during the calibration phase is measured from a simulation. Therefore, the HD calibration signal (related to the obtained weights shown in Figure 5.11b) is artificially injected into the HD cells. The resulting activity profile is determined by the calibration input and the inputs via the reciprocal connections. When the activity profile's shape does not change anymore, the firing rates of the HD cells are measured and put into vector F. The target input current vector U for the CCF 1 is reused from Section 5.3.1. The resulting weight distribution function is created with $\lambda = 649.90$ and shown in Figure 5.12a. The cross-section of the activity band profile, when only the HD cell ring projects onto the CCF 1, is presented in Figure 5.12b.

5.4. Calibration Mechanism

The function of the calibration mechanism consists of two mechanisms that work together in parallel: the feedback loop, and the ACD controlling. Both will be described in the following.

5.4.1. Feedback Loop

The feedback loop consists of the HD cells, the ECD cells, the CCFs 1 and 2, and the ACD cells. The feedback loop covers two functionalities. On the one hand, it is able to calculate ACD information from ECD and HD information. On the other hand, it resets the activity of the HD cells according to the activity in the ECD cells and ACD cells. The feedback loop is driven by two inputs. The first input is the angular velocity that is integrated by the HD cells. The HD signal is always present. The second input is the visual detector that is artificially stimulating the ECD cells when a visual landmark is in the agent's field and range of vision. Depending on if a visual landmark is present, the feedback loop will be switched on or off.

Feedback Loop "On" in Presence of Visual Cue. If a visual cue is detected by the visual detector, both the HD cells and the ECD cell project onto the CCF 1 which, in turn, stimulates the ACDs. Then, the ACD cells and the ECD cells project together onto the CCF 2 which is computing the HD feedback signal. This means that the HD cells receive input that is based on the activity in the ACD cells and ECD cells. The ACD cells' activity is determined by CCF 1 and the PE.

Feedback Loop "Off" in Absence in Visual Cue. If there is no landmark visible to the agent (the landmark is neither in range, nor in sight), the visual detector will not project on the ECD cells and no activity peak will emerge at the ECD cells. In turn, the CCF 1 exhibits an activity band profile (see Figure B.1d in the Appendix). Thus, the summation of the firing rates via the diagonal connections towards the ACDs are of equal size. This leads to a small uniform firing profile at the ACD cells. Finally, no feedback signal will be created, as the CCF 2 is driven by the ACD cells and the ECD cells that exhibit no activity peak.

5.4.2. Allocentric Cue Direction Controlling

Whenever, the feedback loop is switched on, the PE interacts with the ACD cells. It can either store ACD information or reset the activity in the ACD cells according to restored ACD information. The processing of the read ACD information and the derivation of the ACD for resetting can be conducted using two alternative modes as already mentioned in Section 5.2.3.6. In both modes, the PE stores ACD, obtained by egocentrically observing a visual landmark. Subsequently, the stored ACD information is used to correct for HD

errors introduced by PI when revisiting a visual landmark. In the following, the two modes are described in detail.

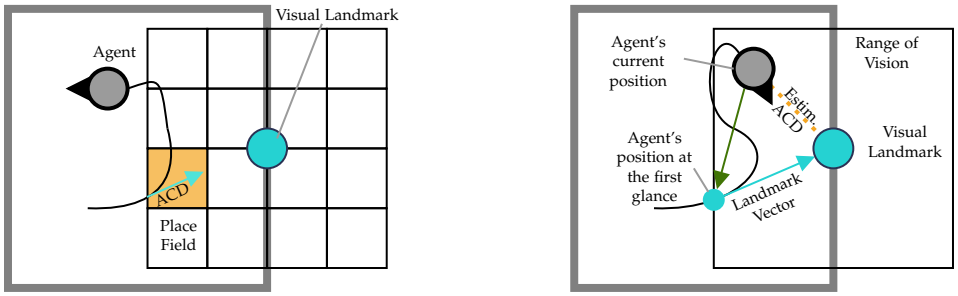

(a) FGL switched off

(b) FGL switched on

Figure 5.13.: Illustration of the PE's operational modes: In Figure 5.13a, the agent associated one ACD to the orange colored place field when the agent perceived the visual landmark the first time in this place field. When revisiting this orange place field, the associated ACD is used to restore the ACD cells' activity. In Figure 5.13b, the PE stored the agent's position from the first glance and the landmark vector which is created with distance and ACD information. The ACD in the agent's current position is estimated via adding the vector to the agent's position from the first glance (green) to the landmark vector.

First Glance Learning "Off". The idea of this first mode is inspired by the place encoding feedback model from Bicanski and Burgess (2016). The position encoding is done via a matrix. Figure 5.13a illustrates an example setup. Every element of the matrix corresponds to a certain squared spatial area around the landmark, called a place field. The central position of the matrix corresponds to the position of the visual cue in the environment. The space that is covered by the whole matrix can be adjusted as well as the amount of the matrix elements. In the example shown in Figure 5.13a, the real space dimension is set to 4. Therefore, the matrix covers a space of $4x4$ units. The edges of this square space are horizontally and vertically 2 units distant from the landmark's position. The visual landmark is in the agent's range of vision when the agent is positioned within the space covered by the matrix. If the granularity of the matrix is set to 4, each place field covered by a matrix element has a size of $1x1$ units. Each matrix element can store a two-dimensional vector. From the beginning, all matrix elements are initialized with zero vectors. Then, the agent explores the environment. If the agent enters a place field and the landmark is in the agent's field of vision, the PE checks the corresponding matrix element. If it contains a zero vector, the current encoded ACD will be transformed into a unit vector, representing the landmark's allocentric bearing, to override the zero vector in the matrix element. If the matrix element already contains a unit vector, the ACD will be restored from this unit vector to subsequently reset the activity of the ACD cells. The

reset signal from the PE towards the ACD cells is present as long as the agent is located in this place field. The ACD cells encode the same ACD as at the time when the agent saw the landmark the first time from this place field.

The overall idea is to exploit the slow error accumulation of the PI process while associating ACDs with positional information. If the agent is in the space covered by the matrix and the landmark is the agent's field of vision, the feedback loop is switched on and ACD information is generated using ECD and HD. Therefore, the current HD error is contained in the ACD signal as well. However, it can be assumed that the size of the error is only slightly changing while the agent moves from one place field to a neighbored place field. This means that the ACDs that are learned while discovering new place fields contain similar sized HDs errors. In this way, the HD can be anchored to a visual landmark with a persistent error.

First Glance Learning "On". In this mode (see Figure 5.13b), the PE stores all necessary information to use a landmark for orientation at once. When a landmark is in the agent's range and field of vision the first time (the moment of the first glance), the PE stores the agent's position, the distance of the landmark, and the ACD encoded by the ACD cells. After that, the agent resets its HD whenever the landmark is again in the agent's field of vision. For this purpose, the real space dimension of the matrix determines if the visual landmark is in the agent's range of vision, similar to when FGL is switched off. The resetting of the HD signal is conducted via a vector addition. Therefore, the agent needs to know its current position and to remember its position from the first glance. Then, the vector from the current position towards the position of the first glance is added up to the vector from the position of the first glance to the landmark. However, this vector addition is not based on neurons and, therefore, is not conducted in a biologically plausible way. The direction information of this newly obtained vector describes the ACD estimation at the current position of the agent. Subsequently, this newly obtained ACD is used to reset the ACD cells' activity. In turn, this resets the HD via the CCF 2.

The idea behind this mode is to allow the model to use a landmark for orientation immediately while only storing information that animals can infer using their senses. The position information of the agent could be inferred using distances of boundaries. Distance information could be derived from binocular visual perception. However in this work, the position information is directly acquired from the simulation environment and the vector calculations are not conducted in a biologically plausible way. Thus, this mode accounts more for mimicking the calibration behavior found from animals than designing a neural vector computing circuit.

6. Experiments

In this chapter, the experimental evaluation of the proposed HD calibration model with its two operational modes is described. It is compared with the HD model from El-Sewisy (2020) that is solely based on PI (without calibration) and with a simple feedback calibration model. This leads to four different HD model setups that will be run in different environments. The simple feedback calibration model is created via direct projections from ECD cells onto the HD cells in line with the model from K. Zhang (1996) described in Section 3.1. However, the connections are fixed from the beginning and are not altered during the agent's exploration phase.

The difference between the encoded HD and the ground truth orientation of the agent, called HD error, serves as metric to evaluate the performance of the HD models. The smaller the HD error, the better the performance. The aim is to investigate if the proposed HD calibration model actually performs better than HD without calibration and the simple feedback model.

6.1. PyBullet Desktop Simulation

The first evaluation experiments are conducted via desktop simulations. The environment for these experiments is based on PyBullet (Coumans and Bai 2016). In each of these simulations, a pioneer robot moves through an environment bounded by walls which contains one visual landmark. The robot setup is reused from (El-Sewisy 2020). The Pioneer P3dx robot has 16 proximity sensors and is guided by a braitenberg controller to move through the environment avoiding obstacles. Furthermore, the agent has a $90\,^\circ$ field of vision, which means $45\,^\circ$ to the left and $45\,^\circ$ to the right. One robot simulation time step is set to 0.05 s. One time step in the neuron model is 0.5 ms. Thus, hundred neuron steps are calculated between two robot steps. The PI input is generated by measuring the angular velocity of the agent. Subsequently, the size of the angular velocity is used to stimulate the SLs which, in turn, shift the activity peak in the HD cells, like in El-Sewisy (2020). The ECD cells are stimulated when the visual landmark is in the agent's field and range of vision.

6.1.1. Box Environment

The first simulation environment is a square box with the upper left corner positioned at $(-1, 1)$ and the lower right corner at $(1, -1)$ (See Figure 6.1a). The cyan object placed in the environment represents the visual landmark. In the beginning, the agent is placed in a lower middle position at $(0, -0.5)$ and oriented to the right (normal vector $(1, 0)$

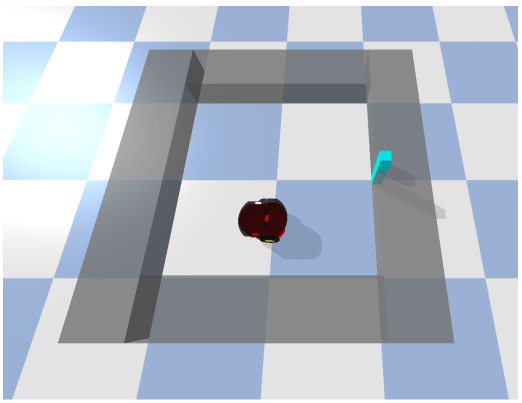

(a) The layout of the box environment with proximal visual cue at $(1.05, 0)$

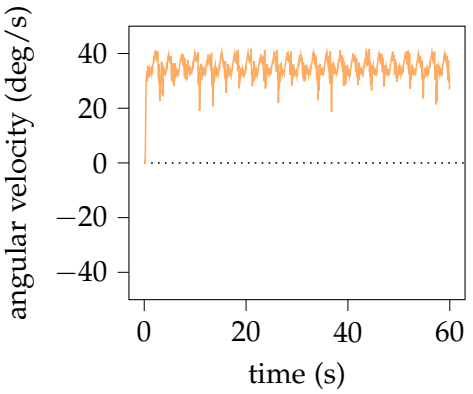

(b) Angular velocity of the agent during a 60 s run

Figure 6.1.: PyBullet box environment: (a) shows the visualization of the simulation environment. The agent starts in a lower middle position and begins to run counterclockwise circles from there. The agent has a field of view of 90 °. In each round, there is a phase in which the agent perceives the landmark. In the example picture, a proximal landmark is shown. For the experiment with a distal landmark, its position is at $(20, 0)$. (b) shows the angular velocity during the run. It is constantly between 20 and 40 deg/s which describes the continuous left turn of the agent.

which corresponds to 0°). Then, the agent runs counterclockwise circles. One episode in the box environment is set to 60 s in which the agent is driving almost six full circles. The agent's angular velocity over time is shown in Figure 6.1b. The angular velocity is mostly between $20\,°/s$ and $40\,°/s$. The dimension of the real space that the matrix (see Section 5.4.2) of the PE covers around the landmark is set to 50 units to ensure that the landmark is recognizable for all four HD model setups, even when the landmark is distal. One unit represents one square field (white or blue) in the chessboard pattern in the PyBullet environment (See Figure 6.1a). The number of matrix elements is set to 150. Hence, every element of the matrix covers a place field of $0.\bar{3}x0.\bar{3}$ units.

6.1.1.1. Box with Distal Landmark

In the first environmental setup the box is used with a distal visual landmark set at position $(20, 0)$ which is ten times more than the dimension of the box. Figure 6.2 shows the performance of all four different HD models.

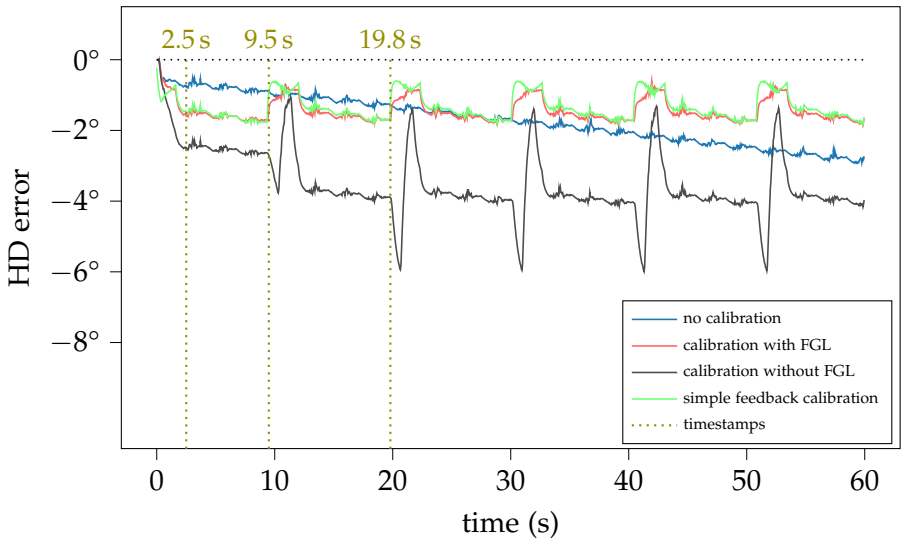

Figure 6.2.: HD error plot for a 60 s run in the box environment with a distal cue at position $(20, 0)$. The timestamps 9.5 s and 19.8 s indicate exemplary moments when the visual landmark is in sight. The timestamp 2.5 s describes the last moment when the visual landmark is in sight during the first circular turn of the agent.

All four HD model setups show negative HD errors. This indicates that the shift of the HD cells' activity peak is too slow to follow the rotation movement of the agent. The blue graph shows the HD error evolving over time with no calibration mechanism. In this case, the HD error increases limitless with a slope of $\approx -0.05\,°/s$. In contrast, all three calibration options limit the HD error to an upper bound. When the landmark is out of sight, the HD errors of all three calibration models increase over time with the

same slope as for the model without calibration. The performance difference lies in the moments when the calibration is active.

The place encoding feedback model without FGL performs the least accurately. It exhibits HD errors up to $\approx -6\,^\circ$. Until 2.5 s, the feedback loop is active as the landmark lies in the agent's range and field of vision. This causes the HD error to increase quicker than when using PI only. The reason is that the current feedback signal is created based on HD and ECD that are already a few time steps old. In this phase, the ACD is associated with the visited place fields around the visual landmark by the PE. Therefore, the smallest error is included in the ACD that is generated right at the start. ACD that is generated later in the run includes a larger error. When the agent perceives the landmark again from already visited place fields (at 9.5 s, 19.8 s, etc.), the ACD corresponding to the currently occupied place field is restored from the matrix and used to reset the ACD cells' activity. Subsequently, the HD signal is calibrated via CCF 2. During the calibration phase (e.g. from 9.5 s until $\approx$ 13 s) the HD error gets even larger in the beginning. Then, the HD error is reset to almost $-1\,^\circ$. Afterward, the HD error increases again.

The reason for this is the granularity of the position encoding matrix. One unit vector describing ACD is associated with one place field. However, an agent can be located in different positions within a place field. The ACDs corresponding to these positions within a place field are not all the same. Hence, the associated unit vector to a place field can not describe all possible ACDs within a place field. Therefore, parallax errors are introduced when the agent revisits a known place field but traverses positions that are different compared to when the agent visited the place field the first time. The closer a place field is to the position of the landmark, the larger the parallax error can be. Nevertheless, there are certain positions within one place field that are accurately fitting to the stored ACD. These positions are placed on a line whose extension exactly intersects the visual landmark's position. If the agent only revisited the positions on this line, the calibration would work optimally. When the agent visits positions apart from this line, parallax errors are introduced into the HD signal. Additionally, the agent is not always revisiting already known places, but also discovers new places, for example, after the first round. This can be seen by the offset between the positions 9.5 s and 19.8 s in the gray HD error graph. In this case, no calibration is conducted, but instead, new place fields are visited the first time and, hence, ACD information is associated with them. From the moment when the agent only revisits place fields, the calibration performance stays constant over time.

The simple feedback calibration and the place encoding feedback calibration model with FGL act on a similar level with slight advantages for the simple feedback calibration. The simple feedback calibration model is resetting the HD via direct projections from the ECD cells onto the HD cells whenever the landmark is in the agent's field of vision. As the landmark is not in infinite distance, small parallax errors are introduced. Hence, both models can not reset the HD error to zero (e.g. at timestamps 9.5 s and 19.8 s). If the cue was in infinite distance, the performance of both models would significantly

increase. Additionally, the HD error increases again at the end of each calibration phase (e.g. at 13 s). This could be explained by degeneration of the firing in the ECD cells. While the firing rate is decreasing as the landmark is out of sight, the ECD cells still give input to the HD. This negatively affects the performance of the PI and can be observed as well for the proposed calibration mechanisms with FGL.

The calibration mechanism with FGL is always a little slower than the simple feedback model. The reason could be that the calibration signal that is introduced via the ACD cells takes twice as long to reach the HD cells because the CCF 2 is in between the ACD cells and the HD cells.

6.1.1.2. Box with Proximal Landmark

In the second simulation environment, the box contains a proximal visual landmark at position $(1.05, 0)$ which is right in the middle of the right border (See Figure 6.1). The resulting HD errors over time are plotted again together for all four HD model setups (see Figure 6.3). All HD error graphs show again negative values. Again, this is caused

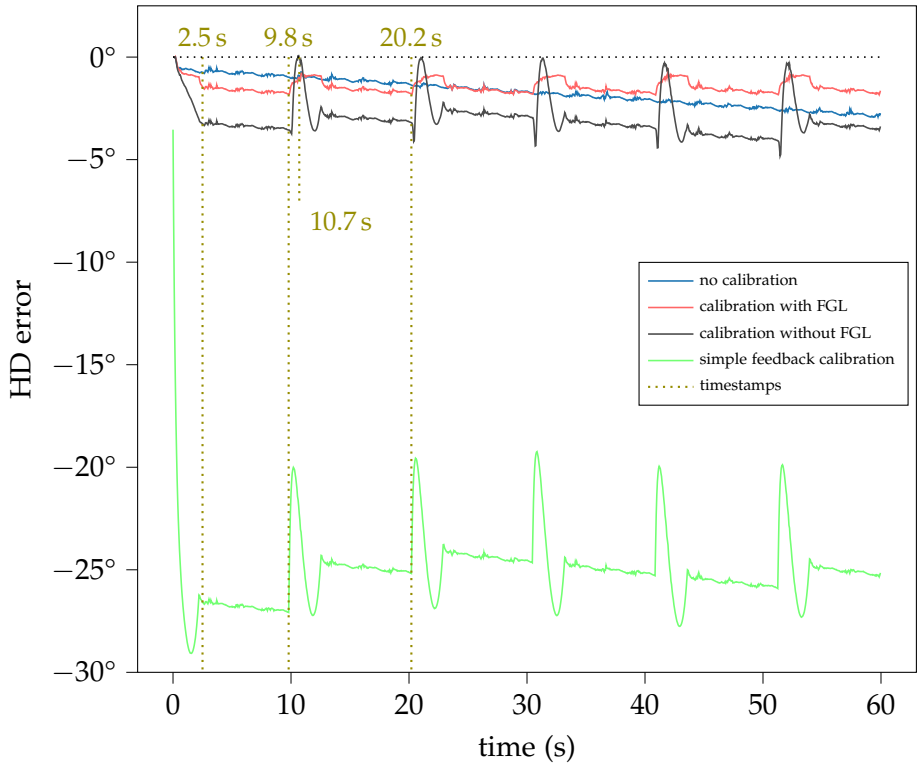

Figure 6.3.: HD error plot for a 60 s run in the box environment with a proximal cue at $(1.05, 0)$: The timestamps 9.8 s and 20.8 s indicate example moments when the visual landmark is in sight. Timestamp 2.5 s describes the last moment when the visual landmark is in sight in the first circular turn of the agent.

by the constant rotation of the agent in one direction and the too slow reaction of the PI mechanism. The calibration with FGL performs on the same level as when using the distal landmark.

The place encoding feedback calibration without FGL is performing slightly better than with the distal landmark. This could occur from the different positions of the landmark. This has two effects. Firstly, the matrix elements cover slightly different place fields when the landmark is placed at $(1.05, 0)$. Secondly, the agent perceives the proximal landmark later than the distal landmark after returning to the starting position after every circle. For example, if both landmarks were present and the agent was in the starting position, the distal landmark would appear as being to the right side of the proximal landmark. This effect can be observed in Figure 6.3 as well. The first calibration phase starts at 9.8 s, which is 0.3 s later than when using the distal landmark. Hence, the agent probably associates ACDs to visited place fields that fit better to the revisiting route through the place fields. For example, at 9.8 s, the HD is calibrated almost perfectly to 0 °. That means that the restored ACD from the currently revisited place field fits perfectly to the real ACD at the agent's position. However, the next calibration step, starting at 10.7 s, increases the HD error again. That means that the ACD does not match the actual ACD corresponding to the agent's position. One reason could be that the associated ACD was stored when the HD signal contained already a larger error (e.g. at 2.5 s). Another reason could be the slightly different route that the agent could take through the known place field as described above.

The simple feedback calibration model introduces a huge HD error. This occurs due to the direct projections from the egocentric perception of the landmark onto the HD system. For example, if the landmark is perceived as being in front of the agent, the simple feedback calibration model will calibrate the HD to 0 °. However, if the agent heads towards 45 °, the calibration mechanism introduces a 45 ° error. This shows that the simple feedback model is not useful when it comes to using a proximal cue as a point of orientation. Concluded, the calibration model with FGL performs best, followed by the calibration without FGL. The simple feedback model performs the worst in this experiment. However, if the agent runs infinitely long in circles, the simple feedback model will be at least performing better than the model without calibration.

6.1.2. Maze Environment with Proximal Landmark

As the simulation in the box environment shows the benchmark test for a good calibration mechanism is a proximal visual cue. Therefore, the maze environment shown in Figure 6.4a contains a proximal cue to evaluate the performances of the calibration mechanisms in a more realistic experiment adapted from El-Sewisy (2020). The visual landmark (cyan object) is positioned at $(8.5, -7.7)$. The real space that the position encoder covers around the landmark is set to 4. The number of matrix elements is set to 12. The agent starts in the lower right corner of the maze at position $(8, -10)$. It is oriented to the top (normal vector $(0, 1)$) which corresponds to an HD of 90 °. From there, the agent drives on the straight way into the middle of the maze and leaves the maze on

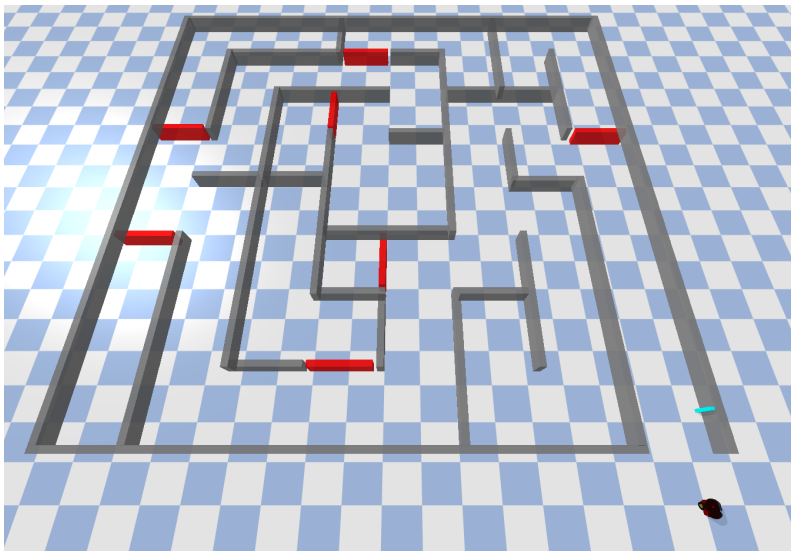

(a) The PyBullet maze environment with proximal visual cue (cyan object)

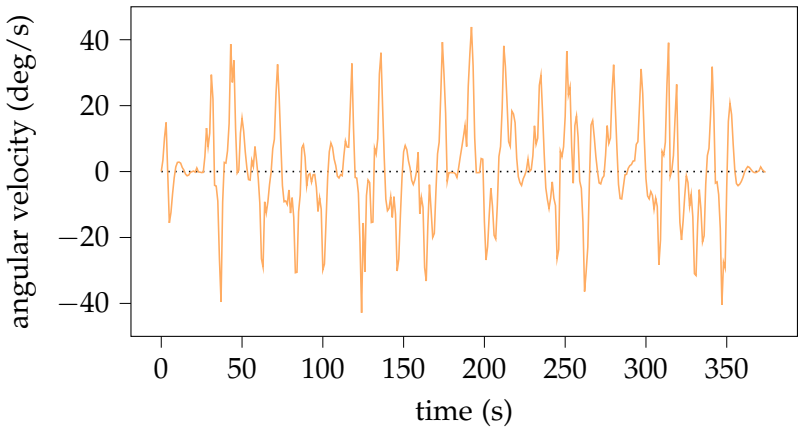

(b) Angular velocity of the agent during a 375 s run in the maze

Figure 6.4.: PyBullet maze environment adapted from El-Sewisy (2020): (a) shows the simulation environment. The agent starts in the lower right corner and is oriented towards the top. The agent enters the maze and takes the direct route into the middle of the maze, where it turns around and takes the direct way out of the maze. The visual cue is placed at the entrance of the maze. There are two phases in which the agent perceives the visual cue during the experiment - when entering the maze and when leaving it. (b) shows the angular velocity of the agent while navigating through the maze. The graph shows that there is an almost equal amount of right and left turns.

the same way, after turning around in the middle. One episode in the maze environment takes 375 s. The angular velocity that the agent has during this run is shown in Figure 6.4b. In contrast to the box environment, the agent turns in both directions in an almost equal amount. In this experiment, the agent visits the landmark exactly two times: one time when entering the maze and one time when exiting the maze. This run simulates the so-called "home base behavior" that was found, for example, in experiments with animals (Eilam and Golani 1989). The agent anchors its internal sense of orientation to a prominent cue that has behavioral importance to the animal. The place of this cue is referred to as the "home base". In this case, this is the visual landmark placed at the entrance of the maze. Then, the agent explores the unknown environment solely based on PI to estimate its orientation. Hence, the HD signal accumulates errors during the exploration phase right until the agent returns to the home position (around the visual landmark). Finally, the visual landmark is used to reset the HD activity to the size of the HD error from the start of the exploration phase. Figures 6.5a and 6.5b show the result for one run in the maze. The result can be divided into three parts: the starting part from 0 s to 2.75 s, the middle part from 2.75 s to 359.5 s and the last part, the calibration part, until 375 s. All models perform equally in the middle part. The landmark is out of the agent's field of vision in this phase. Therefore, all models purely rely on PI to update HD. The performances only differ at the beginning and the end.

The HD model without any calibration feedback ends up with an accumulated error of $\approx -0.31\,^\circ$. A good calibration model should outperform this value. The simple calibration feedback model is not able to reach that goal. It shows already a huge error from the beginning of more than $-60\,^\circ$. This can be explained by the starting position of the agent relative to the landmark. The agent heads towards $\approx 90\,^\circ$ and the landmark lies at the front right of the agent. In egocentric coordinates, this refers to $\approx 340\,^\circ$. Therefore, the ECD ring calibrates the HD system to $\approx 20\,^\circ$. The difference of $\approx 70\,^\circ$ fits the introduced error in HD at the start of the episode. During the middle part, the agent relies on PI. In the end, when the agent leaves the maze, the simple feedback calibration setup has again a similar effect as in the beginning. This results in an HD error of up to $60\,^\circ$.

Compared to pure PI, the place encoding feedback model without FGL introduces a larger error of around $\approx -0.5\,^\circ$ in the starting phase. In the middle part, the offset to the blue graph is constant. However, in the third phase of the episode, the HD error becomes even larger. The reason is that the model is not calibrating at this moment, but rather associates new ACDs to newly discovered place fields. That means that the visual landmark is the first time in the agent's field of vision when the agent visits these place fields. This causes the HD error to increase up to $-0.98\,^\circ$. Then, the agent passes the landmark and relies again on PI, which describes the last wiggle in the HD error graph.

The place encoding feedback model, with FGL switched on, introduces an error of around $-0.38\,^\circ$ in the starting phase. This is larger compared to pure PI but smaller than without FGL. The agent only needs one shot to learn to use the visual landmark to calibrate. Therefore, the error increased only at the very beginning. Then, the visual

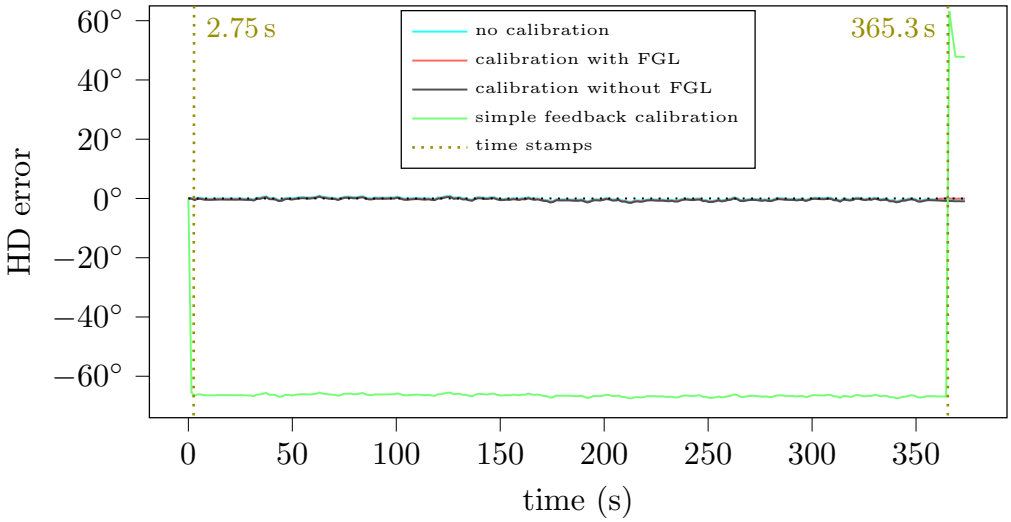

(a) HD error plot including the simple feedback setup (green graph)

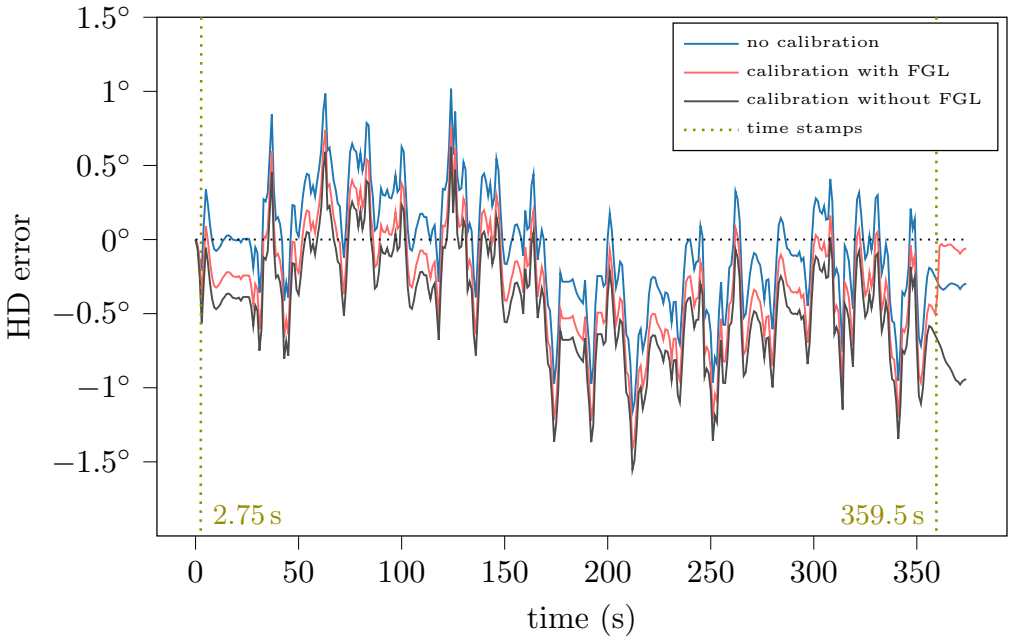

(b) Zoomed HD error plot excluding the simple feedback setup

Figure 6.5.: HD error plots for a 375 s run in the maze environment with a proximal cue at position $(8.5, -7.7)$: the timestamps indicate the moments when the visual landmark is in sight.

landmark is already used to control the HD in the beginning phase before 2.75 s. In the third part, the HD error decreases to $\approx -0.032\,^\circ$ as the model can use the landmark from any position after it was perceived at least once. Thus, the agent does not have to have a view at the landmark in an already discovered place field. It can use the landmark whenever it is in the agent's range and field of vision, independent from the agent's position around the landmark. The last wiggle of the graph arises because the landmark is out of sight again and the HD is driven purely by PI.

Altogether, it can be stated that the calibration with FGL performs best in the "homing" experiment followed by the HD model without calibration which performs even better than the HD calibration mode without FGL. The simple feedback model performs the worst under these conditions.

6.2. Real World Experiment

The second evaluation is done using data recorded from a real world experiment in which a remote-controlled car drives a route in an open field. During the run, the car's position and orientation are recorded. The robotic car is shown in Figure 6.6a.

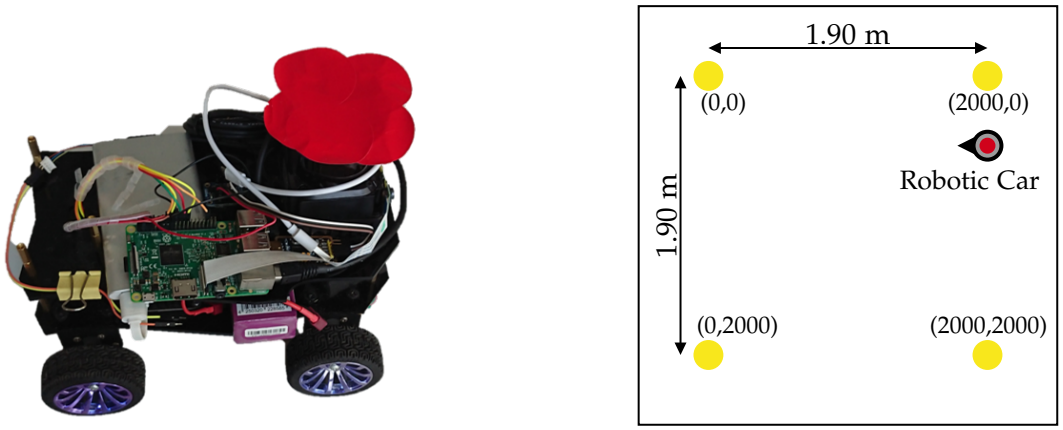

(a) Robotic car

(b) Camera view on the real world setup

Figure 6.6.: Setup of the real world experiment: (a) shows the robotic car (agent). Its position is tracked via detecting the red stickers on its top with a camera. The angular velocity is monitored and recorded using an inertia measurement unit placed in the middle of the car. Figure (b) shows the view of the camera on the experimental field and the agent. The camera is mounted above the field. The four yellow stickers are used as referential points to infer the agent's position (red sticker) during a run.

It contains a raspberry pi that receives drive commands from a remote-control and executes them by controlling the wheels of the car via general purpose inputs/outputs. The remote-controlling is facilitated via the left stick of an Xbox controller. Thereby, the

car is able to drive straight forward, backward, and to rotate left and right. Furthermore, the car contains an inertial measurement unit (Bosch Sensortec BNO055 (Sensortec 2020)) to track the agent's orientation during the run. An RGB camera of an Xbox 360 Kinect sensor (Microsoft Corporation 2010) is used to track the position of the agent. The camera is mounted above the open field in which the car moves. The camera's view on the experimental setup is shown in Figure 6.6b. The robot car has a red sticker on the top. The RGB Kinect camera tracks the position of this red sticker, and thereby the car's position, in relation to the coordination system that is set up by the four yellow stickers on the floor.

The orientation and position data is recorded using ROS (Stanford Artificial Intelligence Laboratory et al. 2018). After recording, the orientation and position data is synchronized by interpolating and sampling the data at time steps of 0.05 s. The finally obtained data set is used to run the HD models. The visual landmark's position is artificially set in the implementation of the model environment and is not part of the data recorded from the real world experiment. Only proximal landmarks are used in the simulation. The positions are set separately for each run to ensure that the agent perceives the landmark from different angles and distinctive often in different runs. The dimension of the square space around the visual landmark that indicates when the agent is close enough to the landmark is set to 1000 for all following experimental setups. This means the visual landmark is only in the agent's range of vision if the agent is located within this square space. The number of matrix elements that are used to encode place fields in the calibration mode without FGL is set to 3000. In the following, the three different runs that have been performed with the robot car are described and the results of the model evaluation are illustrated.

6.2.1. Run 1

In the first run, the agent drives one single circle. The exact positions during this run are shown in Figure 6.7a. The landmark is positioned at $(1850, 650)$. The agent starts and ends up in the upper right corner, where the landmark is positioned as well. Therefore, the agent perceives the visual landmark two times. The angular velocity during this run is plotted in Figure 6.7b. As the agent runs in a left circle, the angular velocity has mostly positive values up to 120 °/s. Figure 6.8 shows the HD errors for the three different HD model setups. At the beginning, the agent stands still up to 6 s. Then, the agent turns slightly and perceives the visual landmark the first time from 6.7 s to 7.8 s. In the beginning of this interval, the calibration models with and without FGL introduce a larger error than without calibration due to the information delay in the feedback loop as explained in Section 6.1.1.1. The calibration model with FGL manages to reduce this error right afterward to reach almost the same error level as without calibration at 7.8 s. It stores all necessary information to use the landmark for orientation at the first glance. At this moment the error is increased. Subsequently, any other view on the visual landmark is used to reset HD. In contrast, the calibration model without FGL keeps a constant offset to the graph describing the HD error of the model without

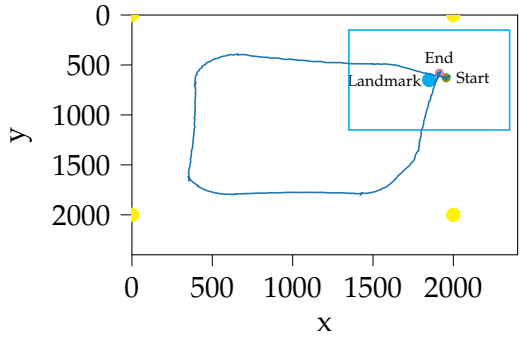

(a) Run 1: agent's route and landmark position

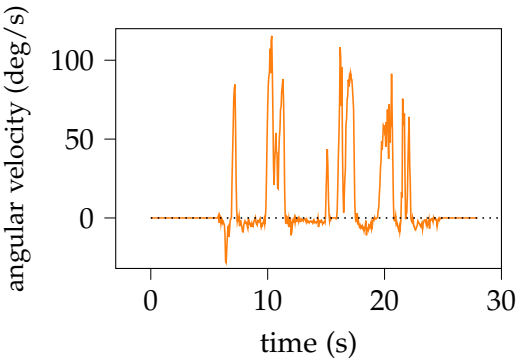

(b) Run 1: angular velocity during the run

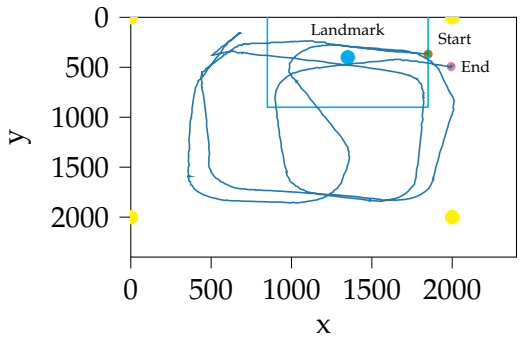

(c) Run 2: agent's route and landmark position

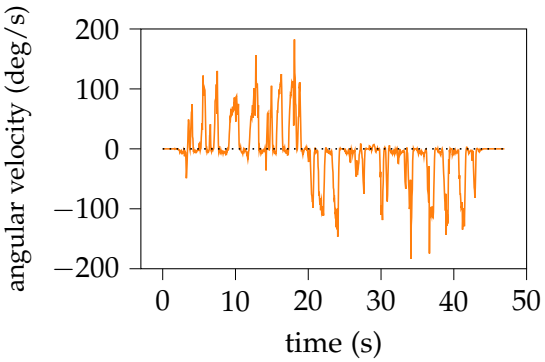

(d) Run 2: angular velocity during the run

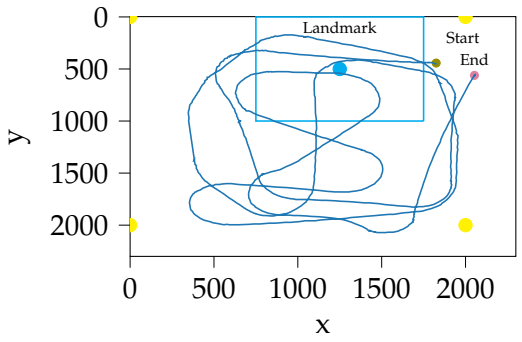

(e) Run 3: agent's route and landmark position

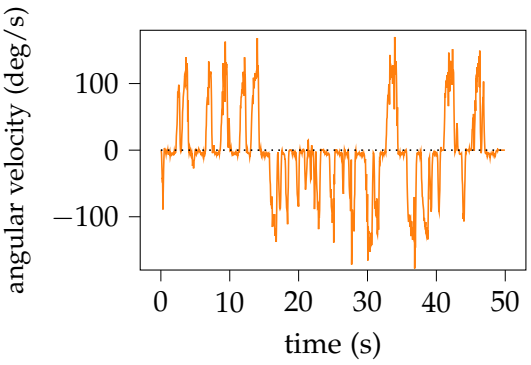

(f) Run 3: angular velocity during the run

Figure 6.7.: Position and angular velocity during the real world experiment: Figures (a), (c), and (e): The blue graph shows the route of the agent tracked in reference to the four yellow stickers for each run. The olive and the purple dots show the agent's starting and final positions. The cyan dot shows the landmark's positions. The cyan blue rectangles with dimensions 1000×1000 describe the maximum distance from which the agent can perceive the landmark. (b), (d) and (f) show the angular velocity over time while the agent moves.

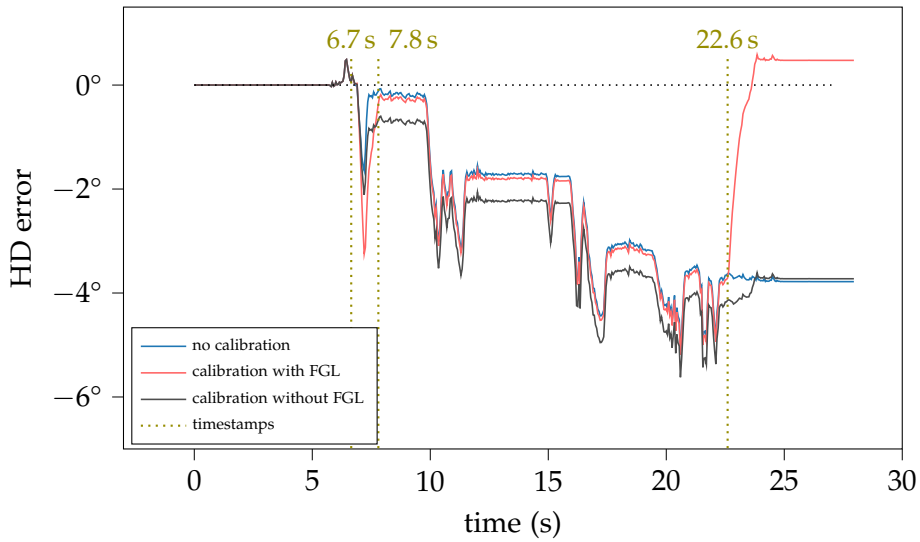

Figure 6.8.: HD error plot for the first run using the real world data: The interval between the timestamps 6.7 s and 7.8 s indicate a phase when the landmark is in sight of the agent. Timestamp 22.6 s indicates the moment that the landmark is in sight when the agent returns to the top right corner.

calibration. It is not able to use the landmark to calibrate directly. It needs a second view at the landmark in the same place field. After 7.8 s, the landmark is out of sight and the agent starts to run the circle. The model perceives the landmark a second time at 22.6 s. The calibration model with FGL resets the HD error to 0.48 °. The reason for the small error at the end could arise from a parallax error introduced while one-shot learning the landmark vector at the beginning of the episode. If then the agent comes very close to the position of the visual landmark, the introduced parallax error is the largest. However, it limits the HD error to a maximum and, therefore, improves the performance compared to the model with no calibration and to the calibration without FGL. The blue and the gray graph show almost the same HD error at the end. As the agent does not revisit one place field where the agent has another view on the landmark, the calibration model without FGL never restores ACD to calibrate. It rather associates ACD to place fields where the landmark is the first time in sight. Therefore, the model introduces parallax errors into the HD signal. At 22.6 s the parallax error even helps the model without FGL to lower the HD error.

It can be concluded that the calibration model using FGL performs best as it resets the HD error when the landmark is seen a second time. The calibration model without FGL performs on a similar level than without calibration because the agent never perceives the landmark a second time from the same position in this run.

6.2.2. Run 2

The second run is used to evaluate the different HD models when an agent revisits a landmark several times and, additionally, performs one phase of left turns as well as one phase with right turns. Figure 6.7d shows the angular velocity over time. In the beginning, the agent is mostly turning left until 20 s. From this moment, the agent starts to do right turns, which results in mostly negative angular velocity. The route that the agent takes and the landmark's position $(1350, 400)$ in this run are shown in Figure 6.7c. The agent starts and ends up again in the top right corner of the environment. The

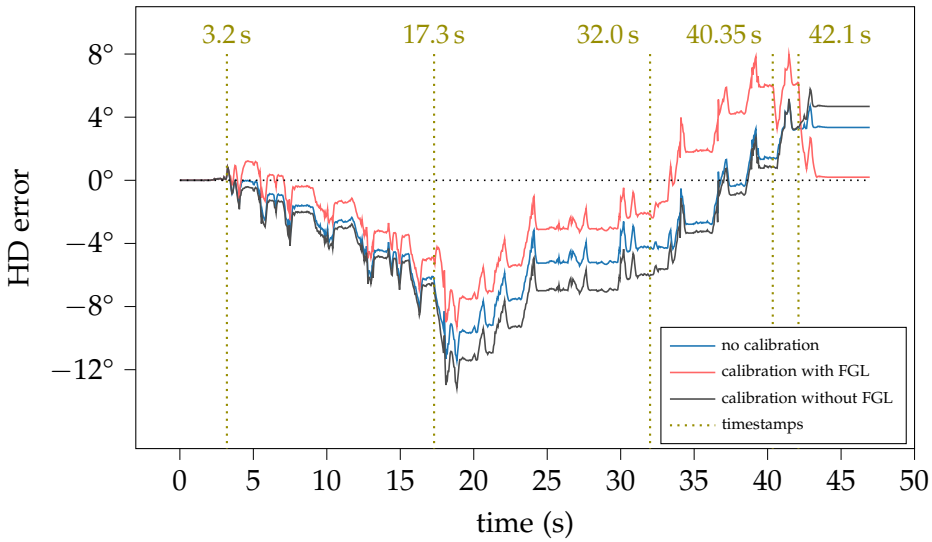

Figure 6.9.: HD error plot for the second run using the real world data: The timestamps indicate the moments when the visual landmark starts to be in sight of the agent.

HD error results for the three different HD models are shown in Figure 6.9. At 3.3 s the agent perceives the landmark the first time. In this phase, both calibration models are learning. Accordingly, they introduce a parallax error in the HD signal. From this moment, the offset between all curves stays almost constant through all the PI phase until 41.35 s. Here, the calibration without FGL increases the HD error as the agent perceives the visual landmark in a new place field. In contrast, the calibration model with FGL is able to reduce the HD error at this point. Exactly the same happens when the agent perceives the visual landmark at 42.1 s.

In line with the results from the first run, the calibration model using FGL achieves the lowest HD error at the end. However, the model without calibration performs the second best and, thereby, better than the calibration model without FGL. The reason is again that the agent is not revisiting a place field in which it already had a look at the visual landmark.

6.2.3. Run 3

The third run's purpose is to investigate the different HD models' performances when the agent perceives the landmark several times while driving an exploration route, similar to the second run. However, in this case, the agent performs multiple changes of direction. Figure 6.7f shows the change of angular velocity during this run. The agent's positions and the location of the landmark $(1250, 500)$ during the last run are shown in Figure 6.7e. The plot of the HD errors (see Figure 6.10) shows again the same results as in the two runs described before. When the agent perceives the visual landmark the first

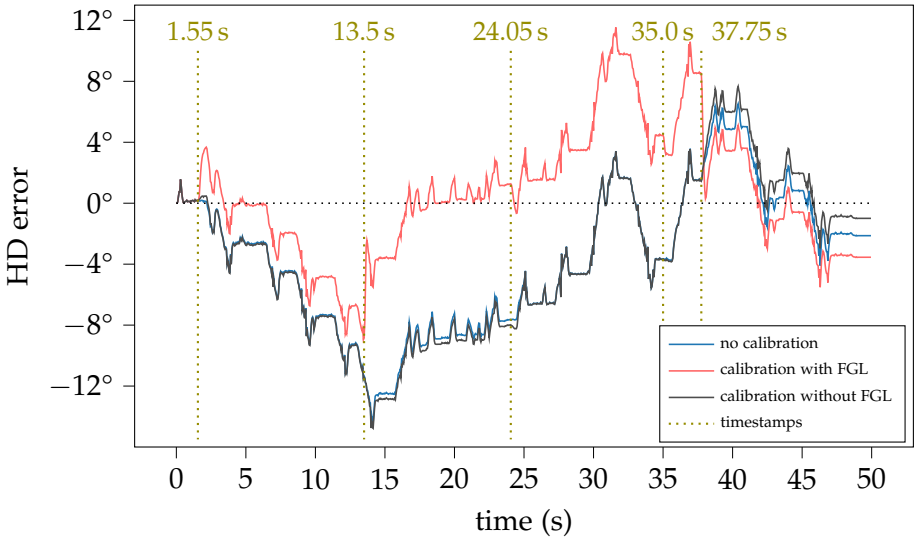

Figure 6.10.: HD error plot for the third run using the real world data: The timestamps indicate the moments when the visual landmark falls into the agent's field of vision.

time (see timestamp 1.55 s), both calibration models introduce a parallax error. Then, the calibration model with FGL is able to use any view on the landmark to calibrate.

The calibration model without FGL does not calibrate HD as the agent never revisits perceives the landmark a second time in the same place field. It rather associates ACDs to new place fields. This can exactly be seen at the timestamps 13.5 s, 24.05 s, 35.0 s and 37.75 s. The calibration model with FGL always reduces the HD errors at these points. The calibration model without FGL introduces even larger errors at these moments due to the closed feedback loop.

The last point in time when the agent perceives the landmark is right after 37.75 s. At this moment, the calibration model with FGL shows the smallest HD error. The calibration without FGL exhibits the largest error. From this moment on the agent is not perceiving the landmark anymore. This means all models only rely on PI. Therefore, the final HD at 50 s is not meaningful anymore to evaluate the calibration performance of the models.

In conclusion, it can be seen that the calibration with FGL works best. The calibration without FGL and the model without calibration perform on a similar level until 37.75 s which is consistent with the findings from the previous runs.

7. Discussion

In this chapter, the calibration model proposed in this work is critically reflected. In the first place, the experimental results are compared with the results of the models from Bicanski and Burgess (2016) and Arleo and Gerstner (2000). In the second place, the general limits and drawbacks of the proposed model are evaluated and illustrated followed by the possibilities and the advantages that the calibration model offers. Subsequently, the two separate proposed calibration modes, with either FGL on or off, are evaluated.

7.1. Comparison with Existing Place Encoding Feedback Models

Bicanski and Burgess (2016) evaluated their place encoding feedback model with data recorded from an agent that explores a circular arena with one visual landmark (see Figure 7.1). The visual landmark is a visual cue card placed at the wall at the north of the arena. The space in the arena is divided into four equal quadrants. The place

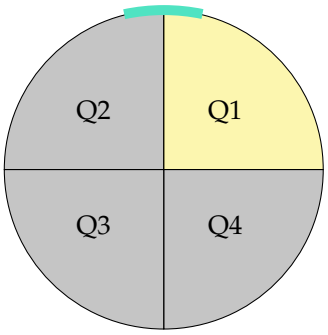

Figure 7.1.: Environmental test setup of Bicanski and Burgess (2016): The circular arena contains one visual cue card to the north. The arena is divided into four equal sized quadrants. Quadrant 1 is used for a comparison with the proposed model of this thesis.

cells' activities encode in which quarter the agent is located. Accordingly, the place cells activate the corresponding interface sheet of neurons that is used to learn the shift between ECD and HD in each quadrant as described in Section 3.1. The agent has a 360 ° field of vision. Therefore, the agent can perceive the bearing of the visual landmark in any pose. At the beginning of the experiment, the agent is oriented to the

north. Then, random translation and rotation movements are conducted by the agent to explore the environment. After each of these movements, the agent dwells for 4 s in its position. One session lasts a few minutes. While the agent explores, the current estimation of HD is compared with the ground truth orientation in each of the four quadrants. Subsequently, Bicanski and Burgess (2016) calculated the average HD error in each quadrant to measure the performance of the model. Finally, these obtained average HD errors are used to compare their proposed place encoding feedback model with a simple feedback model in the circular arena. The experiment in the box environment containing a proximal visual landmark that was conducted in the scope of this thesis (see Section 6.1.1.2) exhibits a similar spatial layout. This fact offers a possibility to compare the model from Bicanski and Burgess (2016) and the one proposed in this thesis quantitatively. However, the proposed model from this thesis only perceives the visual landmark when the agent is located in the bottom right corner in the box environment shown in Figure 6.1a. This space corresponds to the first quadrant of the circular arena of Bicanski and Burgess (2016). Therefore, Table 7.1 shows the results of the experiments conducted in this work and by Bicanski and Burgess (2016) with simple feedback models and place encoding feedback models. The simple feedback models

Table 7.1.: Comparison table with Bicanski and Burgess (2016)

	Bicanski and Burgess		Present Work		
	Simple Feedback	Place Encoding Feedback	Simple Feedback	Feedback Model with FGL	Feedback Model without FGL
HD error	$21.8 \pm 3.3\,^\circ$	$-0.8 \pm 5.1\,^\circ$	$-26.5\,^\circ$	$-1.6\,^\circ$	$-3.8\,^\circ$

introduce absolute HD errors of $21.8\,^\circ$ in the first quadrant of the experiment of Bicanski and Burgess (2016) and $26.5\,^\circ$ in the box environment from this work. In both setups, the direct projection of egocentric bearing information onto the HD system causes these errors. The place encoding feedback model from Bicanski and Burgess (2016) limits this error to an average value of $-0.8\,^\circ$. The proposed calibration model is able to limit the HD error to $-1.6\,^\circ$ with FGL and to $-3.8\,^\circ$ without FGL. This is slightly worse than the model from Bicanski and Burgess (2016). However, the simple feedback model already performed slightly worse in the box environment. This could arise from the difference in the models' viewing angles. As Bicanski and Burgess (2016) use a viewing angle of $360\,^\circ$, the model controls the HD firing in any pose in the first quadrant. In the desktop simulation of the present work, the viewing angle is set to $90\,^\circ$. Therefore, the agent can only derive HD information when the cue is in sight. Furthermore, the agent only runs circles in the experiment of the present work. In Bicanski and Burgess (2016) the agent can perceive the landmark from advantageous angles, for example, when the agent is in the first quadrant but very close to the border of the second quadrant. In this place, the ECD introduces the smallest parallax error into the HD signal. Therefore, it can be stated that both models act on a similar level in their similar environments. However,

the model from Bicanski and Burgess (2016) has the advantage of seeing 360 ° of the environment to learn the shift between ECD and HD in each quadrant. It needs a set of interface neurons to activate and deactivate learning connections in each quadrant. Against this background, the present model can be assumed as being more efficient as it comes without interface neurons and reaches a similar or only slightly worse level of HD errors with less exploration time.

Arleo and Gerstner (2000) implemented their model on a robot and conducted experiments using a square environment with dimensions 80×80 cm that exhibits several visual features. For each run, the robot was placed in a certain position in the environment. Then, the robot starts to rotate randomly for 300 time steps with an angular velocity of 100 °/s. During these rotational movements, the robot stays in its position. This trial was conducted ten times. The absolute value of the HD error, the difference between the agent's ground truth orientation and the estimated orientation of the neuron model, was computed over time during the trials. Finally, the average of the absolute HD error over all trials was calculated. This procedure was done one time for the HD model with calibration mechanism and one time without it, solely based on their PI process. Figure 7.2 shows these results. The experimental setups with the box environment used

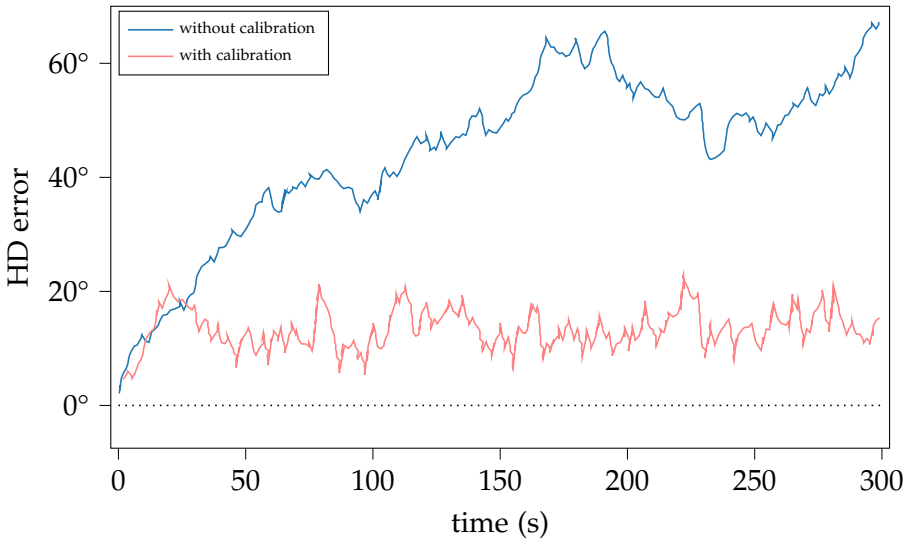

Figure 7.2.: Plot of the average head direction error from Arleo and Gerstner (2000) who conducted $n = 10$ sessions with 300 time steps in a 80×80 cm arena. In one single session, the agent randomly rotates in one fixed position.

in the present thesis are similar. In both experiments, the visual cue that is used for the calibration are seen regularly by the agent. Additionally, the agent has to orient itself using proximal landmarks in both experiments. This allows again a quantitative comparison of the simulation results shown in Figures 6.3 and 7.2. The HD error of both models without calibration is accumulating over time without limit. However, the error introduced in Figure 7.2 rises with 0.2 °/s on average. This is four times more

than in Figure 6.3. This could arise from the higher angular velocity of the agent in the experiment of Arleo and Gerstner (2000). Nevertheless, the overall behavior of both PI mechanisms is similar and can serve as a base to compare the calibration mechanisms.

In both plots, the HD error is limited to a boundary when using the calibration mechanisms. In Arleo and Gerstner (2000) the boundary is around $\approx 21\,^\circ$, whereby in Figure 6.3 the boundary is $\approx 3.5\,^\circ$ without FGL and $\approx 1.8\,^\circ$ with FGL. This difference could again arise from the distinct angular velocity of the agents. However, these numbers show that both calibration models limit the accumulation of errors in a similar way. Additionally, all calibration models introduce an error at the beginning that is larger than when relying purely on PI. In Figure 6.3, it can be seen at timestamp 2.5 s. Figure 7.2 shows this issue at around 20 s. It can be deduced that no model has a clear structural advantage. Differences in the performance can be seen at the time steps when the models are calibrating. As the model has a $360\,^\circ$ field of vision, which allows the model to continuously perceive the visual cue, the calibration time steps can not be derived from the plot in Figure 7.2. This leads to the assumption that the HD is constantly calibrated independently of the agent's orientation due to the $360\,^\circ$ field of vision. Hence, the HD error is consistently kept on the same level around $15\,^\circ$. In contrast, the proposed model from this work only perceives the landmark every turn. Thus, the model calibrates only regularly and not constantly which can also be seen in Figure 6.3.

It can be stated that both calibration models seem to work properly under the given experimental conditions. However, the experimental setup of this thesis (see Section 6.1.1.2) demands more to the calibration model than the experimental setup from Arleo and Gerstner (2000). In the experiment from Arleo and Gerstner (2000), the model only needs to associate different HDs with ECDs in one single place. In the box environment of this work, the agent needs to associate different HDs with ECDs in different places. Hence, the performance of the proposed calibration model of the present work can be estimated as being higher (with and without FGL). If the model from Arleo and Gerstner (2000) was evaluated with the box environment from this thesis, it can be assumed that it would not reach the same performance as the proposed calibration model.

7.2. Possibilities and Advantages of the Proposed Model

In this section, the general advantages and possibilities of the proposed calibration mechanism are evaluated in terms of applicability, biological plausibility, and efficiency without the distinction between calibrating with or without FGL.

7.2.1. Applicability to Environments with Distal and Proximal Landmarks

The results of the experiments in Chapter 6 show that the proposed place encoding HD calibration model is capable of orienting itself using proximal as well as distal landmarks. It limits the HD error to a maximum and outperforms the simple feedback

model adapted from K. Zhang (1996) when using a proximal cue to set HD.

7.2.2. Biological Plausibility

The components of the feedback loop are based on spiking neurons which closely mimic the behavior of biological networks of neurons. Every population of neurons mimics the function of spatially modulated cells that were either proposed or already found to exist, except for the CCF 2. Additionally, no cells have been used that need a correct HD to work properly, like place cells. Furthermore, the calibration feedback signal is not changing the basic shape of the HD activity profile (see Figures B.2a and B.1a). Its magnitude is enhanced when the feedback signal is active but still keeps its overall shape. The feedback signal can shift the HD activity peak to any position that is derived from the landmark, like in K. Zhang (1996) without affecting the performance of the PI process when the feedback loop is switched off. Moreover, the model is able to learn and restore ACDs at the same time similar to animals, that also learn information from the environment to aid their spatial orientation while they navigate in it.

7.2.3. Efficiency of Learning

The feedback loop covers two functionalities to disentangle the relation between ECD, HD, and position of the agent when the landmark is in the field of vision. Firstly, it constantly computes ACD. Secondly, it calibrates the HD signal according to the current activity of the ACD cells. The ACD describes the constant shift between HD and ECD in one certain position. The computation is done in a biologically plausible way based on spiking neurons. This significantly reduces the amount of information that needs to be learned, compared to when every HD needs to be associated with an ECD in each position (like in Bicanski and Burgess (2016), Milford et al. (2004), and Arleo and Gerstner (2000)). Only the relation between ECD and HD has to be learned. This poses the hugest advantage of the calibration model and remarkably increases the model efficiency. In contrast to Bicanski and Burgess (2016), the proposed model does not require any neuron population that acts as an interface to activate learning connections depending on the agent's position in space. Hence, the proposed model is also more scalable than the existing calibration models when it comes to the exploration of larger arenas.

7.3. Limits and Drawbacks of the Proposed Model

In this part, the drawbacks and limits of the proposed calibration model are assessed with again no distinction between calibrating with or without FGL.

7.3.1. Biological Plausibility

Firstly, the PE is not implemented in a biologically plausible way. This means that no neural circuit covers the task of the PE. The tasks of constantly tracking the allocentric position of the agent, calculating the ACDs, filling the storage matrix, and resetting the firing activity of the ACD cells are facilitated by an algorithm. In a more biologically plausible way, the PE should encode the local position in relation to a visual landmark in the given spatial framework of the environment based on spatially modulated cells.

Secondly, ECD information is artificially injected into the ECD cells. A more biologically plausible model could derive visual landmark information from the egocentric image derived from the view an agent has on its environment. However, this was out of the scope of this work.

Furthermore, the proposed model can only deal with one visual landmark. A more biologically plausible mechanism should be able to use more than one visual landmark in larger arenas composed of more compartments.

7.3.2. Optimal Cue Integration

As described in Section 2.4.3, animals can optimally combine cues, which is not covered by the proposed model. The self-motion and the HD information derived from the visual landmark are not evaluated based on their reliability. Instead, the HD information derived from a visual landmark always dominates over PI, even if the landmark was unstable.

7.3.3. Impact on Path Integration Performance

The last drawback concerns the performance of the PI while learning a visual landmark. At the moment when the landmark is seen the first time, the feedback loop is active, but no calibration signal can correct for HD errors at this moment. This causes the feedback loop to control the HD activity peak using information from HD and ECD that is already four time steps old. As a result, the PI performance is significantly decreased in the learning phase.

7.4. Comparison of Calibration with versus without First Glance Learning

In this section the two operational modes of the PE are compared.

Calibration with First Glance Learning With FGL switched on, the calibration mechanism has an equal performance on distal landmarks as the simple feedback model in the experiments. Furthermore, it performs the best of all tested HD calibration models on proximal landmarks. From a behavior analyzing view, the model with FGL is biologically plausible. The agent can orient itself using a visual landmark after perceiving it one

time from one position. Whenever the landmark is perceived again from any position, the knowledge gained from the first glance is used to derive HD information. However, when analyzing the underlying mechanism of FGL, the model needs to perform vector calculations that are not based on a biological mechanism. This poses the question of how this or a similar calibration behavior can be achieved when using a PE that is modeled with spiking neurons.

Calibration without First Glance Learning From the experiments, it was found that the calibration mechanism without FGL is less accurate using a distal cue than the simple feedback model. However, it outperforms the simple feedback model on a proximal cue.

In this mode, ACDs are associated to place fields around the visual landmark. The restoring of ACD can introduce parallax errors into the HD signal, as one ACD can not account for all possible ACDs in one place field.

Moreover, the observed calibration behavior of this mode is biologically less plausible compared to the calibration with FGL. If the agent perceives a visual landmark infinitely long from one place field, this model will still be unable to calibrate its HD when discovering a new place field around the landmark. This is exactly what happened, for example, in the simulations with the real world data. The agent associated ACDs to certain place fields, but as long as the agent only revisits the landmark from different place fields, the agent can not infer HD information. However, the underlying mechanism of this mode is simpler as no vector calculation is needed. Only storing and restoring of ACDs is required. This could be facilitated by Hebbian learning connections between place encoding neurons and the ACD ring. Additionally, the performance of this mode is limited by the viewing angle of 90 °.

8. Conclusion & Outlook

8.1. Conclusion

This thesis presents a biologically plausible HD calibration model which can use a visual landmark as a point of orientation. A proximal landmark can be used as well as a distal landmark. The idea behind the model is to combine information about the local allocentric position of the agent with ECD to generate HD information. This allows the model to reset its internal sense of direction when revisiting a known visual landmark after an exploration phase solely based on PI. It has two operational modes. In the mode with FGL switched on, the model learns a visual landmark on one shot and anchors the HD signal from the first glance at this landmark. When FGL is deactivated, the agent associates ACDs with place fields around the landmark when visiting them the first time. When returning to already visited place fields, the agent uses the learned ACD from the particular place field to reset its HD. The experiments proved that both calibration models can limit the HD error compared to when the agent would only rely on PI. It was further demonstrated that the proposed calibration model using FGL is superior compared to a simple feedback calibration model, like proposed by K. Zhang (1996) and Skaggs et al. (1995). The proposed calibration model without using FGL outperforms the simple feedback calibration model on proximal cues but can introduce parallax errors on distal landmarks. Both modes are more efficient than the state-of-the-art models of Bicanski and Burgess (2016) and Arleo and Gerstner (2000) through the disentanglement of ECD, HD, and position information by calculating ACD using a neural circuit. This significantly reduces the amount of information the model needs to store for calibration. Therefore, less experience is needed in the environment to achieve HD calibration. In both modes, the calibration is independent of the agent's orientation at the second glance compared to state-of-the-art models. Without FGL, the calibration is still position dependent, so that the model needs to perceive a landmark a second time in the same place field to calibrate HD. With FGL, the calibration is even position independent. The model calibrates the HD whenever perceiving the landmark again, independent of position and orientation.

8.2. Outlook

In the outlook, a few ideas are mentioned about how to improve the calibration mechanism in terms of biological plausibility and extending its functionality. In the future, these could help to create a larger neural model that properly mimics, for example, a

rodent's navigation behavior.

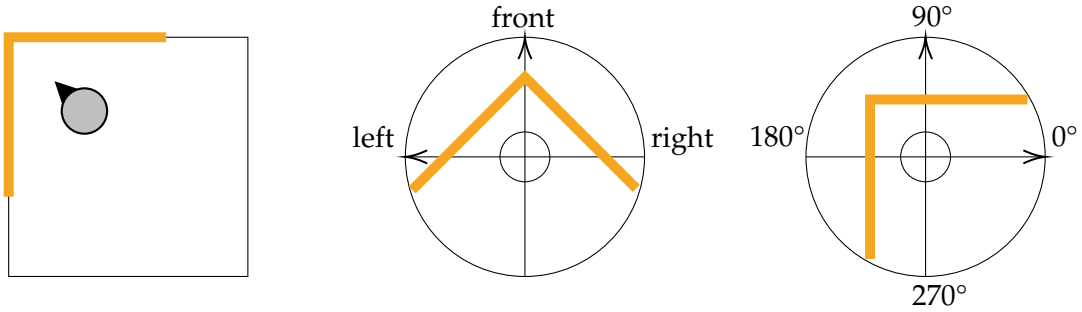

Figure 8.1.: Visualization of boundary vector cells: Left: An agent facing the top left corner in a box environment. Middle: The activity of egocentric boundary vector cells corresponds to the egocentric perception of the boundaries. Right: The activity of allocentric boundary vector cells corresponds to the shape of the environmental boundaries in allocentric coordinates.

Biological Local Position Encoding. The PE could be implemented based on spiking neurons to also make the underlying mechanism more sophisticated. One idea is to use egocentric boundary vector cells (see Figure 8.1). They could represent the local boundaries that the agent perceives. In combination with ECD, this would give enough information to encode the local position in relation to a landmark. An example environment with a straight, infinitely long wall that contains one salient cue card can illustrate this idea. The egocentric perception of the wall allows the agent to determine its orientation in reference to the wall and at what distance it is located from the wall. This encloses the amount of all possible positions to a straight line that is parallel to the wall. The egocentric perception of the landmark lets the agent determine its orientation to the landmark. If both information is combined, the place ambiguity can be resolved and the agent can exactly determine its position in the local environment around the landmark. This place information would be enough to link it with ACD information to set the internal sense of direction in relation to the landmark.

A more advanced idea would be to use allocentric boundary vector cells (see Figure 8.1) to additionally determine the agent's position in an arena with more compartments. These could be created by merging HD information with information from egocentric boundary vector cells. Hence, the allocentric boundary vector cells would contain some error values introduced from the HD system. However, it would be still useful to determine the rough allocentric position of the agent in the larger frame of the environment, which may consist of several similar looking compartments. On top of this rough localization, the local localization could be done using egocentric boundary vector cells as described before.

Another idea for local position determination could be to use the ECDs of two landmarks, for example, the left and right edge of a visual cue card. Two egocentric perceived angles would precisely define the local position of the agent.

A third concept could be to determine the agent's allocentric bearing from the visual landmark in reference to environmental boundaries instead of encoding the agent's position. The landmark's position and the environmental boundary could define an allocentric coordinate system in which the allocentric direction of the agent could be encoded from the landmark's position. The agent could infer this information using egocentric boundary vector cells and egocentric object vector cells, respectively landmark vector cells. Then, the allocentric bearing of the agent could be combined with the computed ACDs from the feedback loop. Thereby, the distance of the landmark from the agent is not considered anymore. Only two angles are associated with each other. As a result, less information has to be learned to use a landmark for orientation.

For further development of the model, Hebbian-learning connections could be introduced to associate position encoding neurons with ACD cells. The ACD cells could serve as post-synaptic neurons and the place encoding neurons serve as pre-synaptic ones. The more often neurons of both populations are active at the same time, the stronger the connection from the place encoding neuron towards the ACD cell becomes. Thereby, the place encoding neurons either gain control over the activity of the ACD cells or will be disconnected over time.

Lastly, the calibration mode using FGL could be implemented in a more biologically plausible way. This requires the model to be able to perform vector additions using neural circuits.

Optimal Cue Integration. Another extension that could improve the performance of the calibration model is the optimal integration of different cues to set the HD. Firstly, this process needs to weigh cues according to their reliability. Secondly, a weighted average of the HD indicated by the different cues has to be computed to subsequently reset the HD. This would allow the model to completely disconnect its HD system from unstable landmarks by a long-term plasticity mechanism as described in section 3.3 (see Figure 8.2). In case two or more cues are equally reliable, optimal cue integration could allow the model to resolve this conflict by resetting its HD to a trade-off value. This value lies in between the HDs indicated by the two separate cues (see Figure 8.3). This behavior could be achieved using short-term learning connections proposed by Page et al. (2014) and described in 3.3. When a landmark causes the HD activity to shift, the short-term plasticity accounts for altering the connection weights during this shifting procedure. Finally, the HD activity peak is positioned between the HDs indicated by the two cues. This capability could be tested by introducing moving cues into the environment.

Visual Stream Implementation. The last proposition to extend the presented model is the implementation of the visual stream of the agent. The egocentric perceived bearing

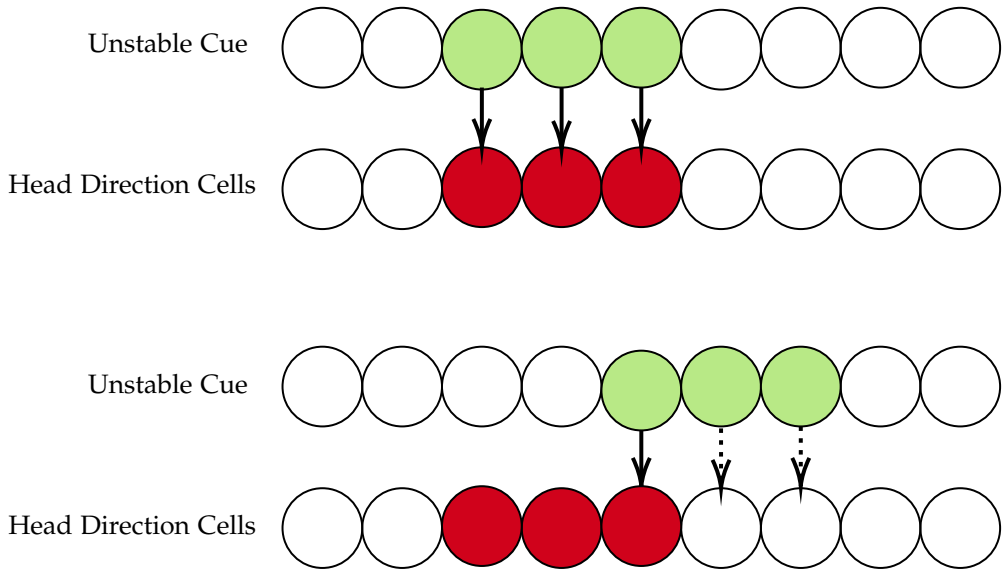

Figure 8.2.: Stability estimation: When the HD indicated from a sensory cue is congruent with the current activity of the HD cells the connections are strengthened. If an unstable cue indicates different inconsistent HDs, the connections are weakened which is shown by the dotted arrows.

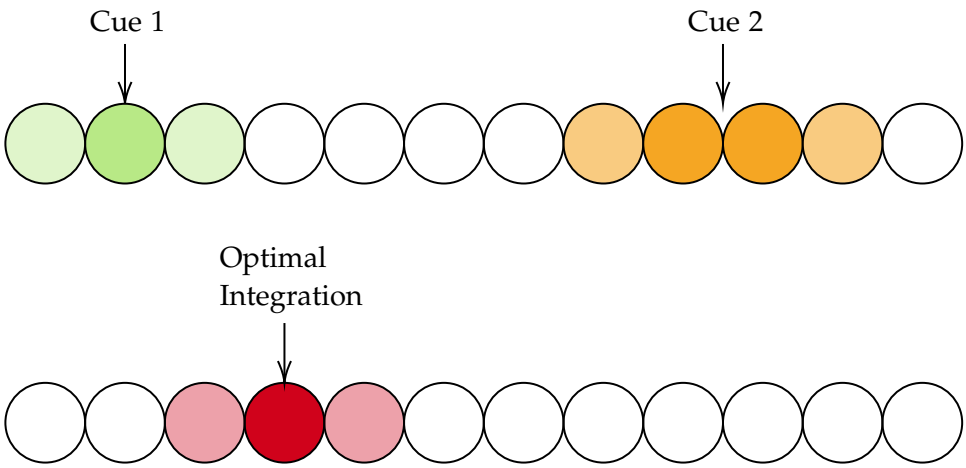

Figure 8.3.: Optimal cue integration: The HD indicated from two conflicting cues are integrated in an optimal way depending on the cues' reliabilities to create an optimal uniform estimation of HD.

of a landmark could be derived from the egocentric view that the agent has on the environment. Several feature detectors could react to different visually salient properties in the environment. Therefore, a visual landmark would be detected as a composition of visual properties and not as a whole configuration. Thereby, the model could deal with more distinct visual landmarks without the need for one visual feature detector per landmark.

A. Implementation Details

The implementation of the calibration model is set up on the existing code from El-Sewisy (2020). The calibration model is mainly implemented in Python (version 3.8). The part for simulating the neural network is written in Python and additionally in CUDA C (NVIDIA et al. 2020). Thereby, the neural steps can be calculated on the graphics card which is adapted from El-Sewisy (2020). The one-dimensional ring components are implemented with $n = 100$ neurons. The two-dimensional ring components are implemented with $n = 10000$ neurons. The connection weights between all the neurons are implemented using matrices. Thereby, the firing rates of a component can be multiplied with the matrix containing the connection weights towards the receiving component to obtain the input currents at the receiving neurons.

The visualization of the firing activity of the components was done using matplotlib (Hunter 2007). When running the implemented model software on the testing computer (Core: Intel Core i5-4210H, Graphics Card: Nvidia GeForce GTX 950M), the following times have been measured for a PyBullet desktop simulation in the box environment (see Section 6.1.1):

- time step of the neural network: 9.3957 ms
- agent time step: 9.6175 ms
- plotting time: 3.002 s

B. Visualization of Components

Figures B.1 and B.2 show the exemplary firing activity of the components of the calibration circuit and the HD cells. Figure B.1 shows the activity when the visual landmark is out of sight. The HD cells are still driven via reciprocal connections and show a maximum firing rate of $\approx 62\,\text{Hz}$. In this example, the HD cells' activity encodes an orientation of $0\,^\circ$ (see Figure B.1a). The ECD cells receive zero input. Accordingly, they show an average firing rate of $8.95\,\text{Hz}$ around the circle corresponding to Equation 5.2 (see Figure B.1b). In turn, the CCF 1 receives only activity input from the HD cells. This manifests itself in an activity band profile shown by Figure B.1d. All conjunctive cells which correspond to an HD of $0\,^\circ$ are slightly active. Therefore, the summations of firing rates of each diagonal yield a uniform firing distribution such that the ACD cells exhibit a uniform firing profile of $8.95\,\text{Hz}$. Hence, the small activity of the ACD cells and ECD cells together leads to almost silent neurons in the CCF 2. In turn, the HD cells receive zero input from the CCF 2. The feedback loop is switched off.

If the visual landmark is in the agent's field of vision, the peak position in the activity profile of the ECD cells will encode the egocentric bearing of the landmark (see Figure B.2b). In this case, it is $\approx 30\,^\circ$. The HD cells and ECD cells together project with their activity peak profile onto the CCFs 1 which in turn exhibits a two-dimensional activity peak shown by Figure B.2d. Consequently, the ACD cells show an activity encoding the allocentric bearing of the visual landmark of $\approx 30\,^\circ$ (Figure B.2c) as a result of the diagonal connections from the CCFs 1. Finally, the CCFs 2 (Figure B.2e) receives the activity input from ECD cells and ACD cells and computes the feedback signal onto the HD cells. The feedback loop is switched on. The HD cells are driven by their reciprocal connections and, additionally, by the CCFs 2. Thus, the activity peak that emerges in the HD cells is amplified (Figure B.2a) compared to when the feedback loop is switched off. The peak activity is still at $\approx 0\,^\circ$. However, its shape is biased and has a maximum firing rate of $\approx 70\,\text{Hz}$.

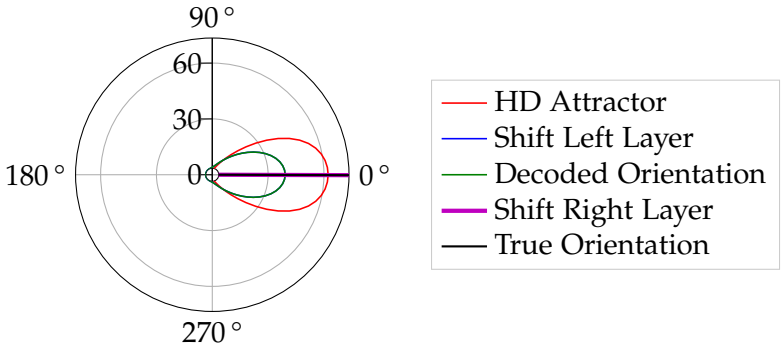

(a) HD cells and SLs: firing rates in Hz

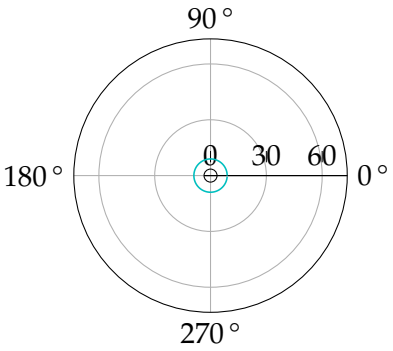

(b) ECD cells: firing rates in Hz

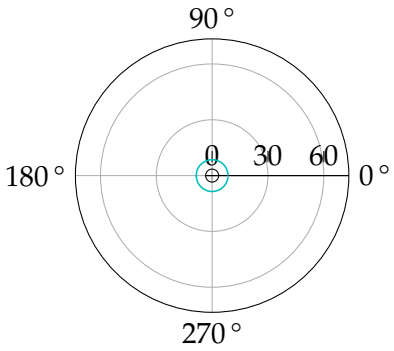

(c) ACD cells: firing rates in Hz

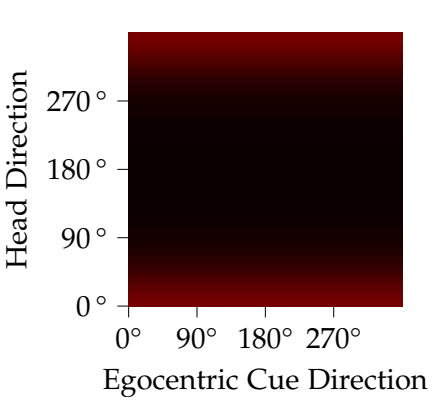

(d) CCF 1: firing rates in Hz

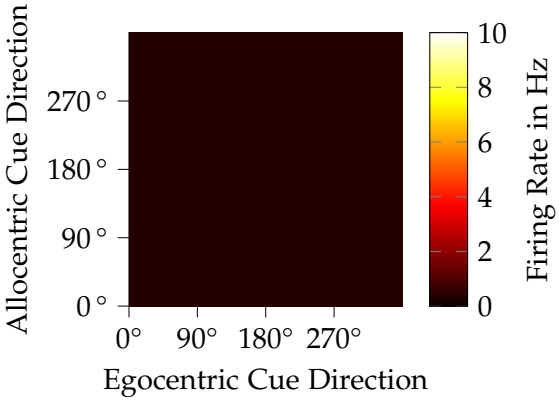

(e) CCF 2: firing rates in Hz

Figure B.1.: Firing rates with feedback loop switched off: Every Figure shows the firing behavior of the corresponding spatially modulated cells from the calibration circuit when the visual landmark is out of the agent's sight.

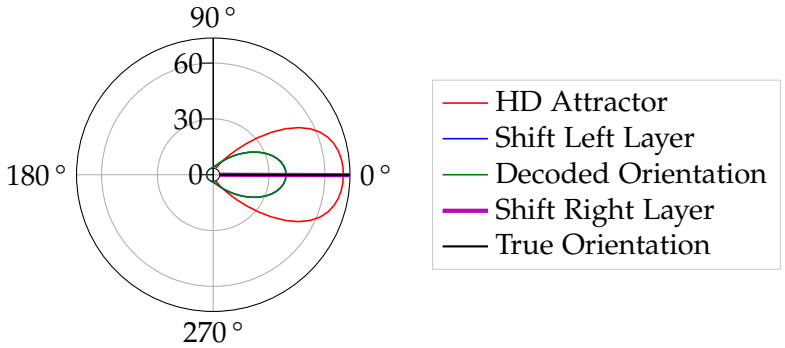

(a) HD cells and SLs: firing rates in Hz

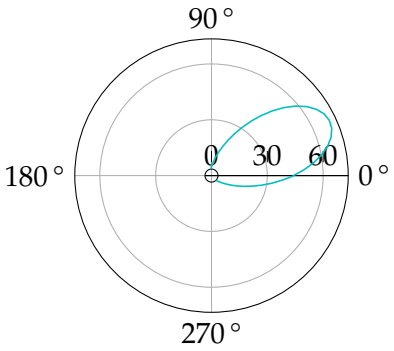

(b) ECD cells: firing rates in Hz

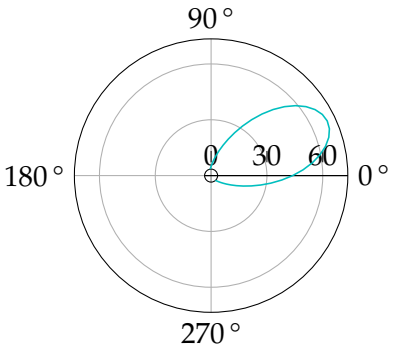

(c) ACD cells: firing rates in Hz

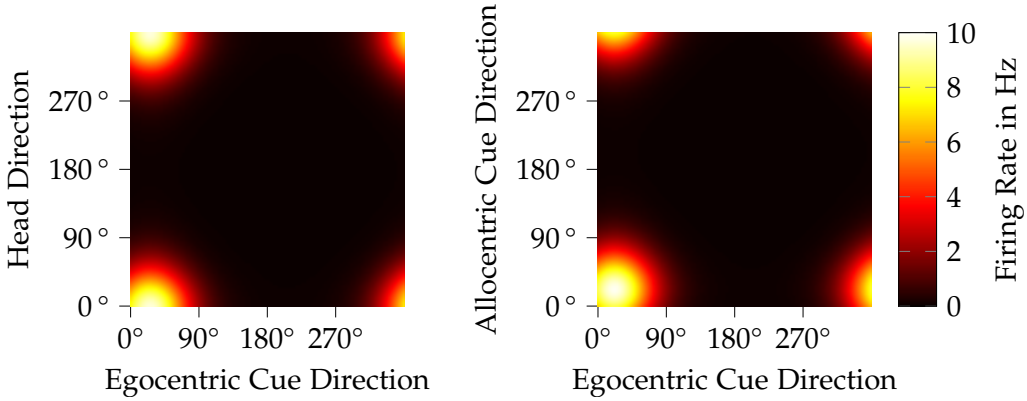

(d) CCF 1: firing rates in Hz

(e) CCF 2: firing rates in Hz

Figure B.2.: Firing rates with feedback loop switched on: Every Figure shows the firing behavior of the corresponding spatially modulated cells from the calibration circuit when the visual landmark is in the agent's field of vision.

List of Figures

Acronyms

ACD allocentric cue direction. v, vii, 2–4, 24, 29–31, 38–45, 50, 51, 54–56, 60, 62, 64, 69, 71, 77–79, 81–83

ACD cell allocentric cue direction cell. 8, 9, 29–31, 38–40, 43, 44, 49–52, 54–56, 60, 61, 77, 78, 83, 89–91, 93

CAN continuous attractor network. 21, 23, 34

CCF conjunctive cell field. 9, 29–31, 40–45, 47, 49–54, 56, 60, 61, 77, 89–91, 93

ECD egocentric cue direction. v, vii, 2–4, 9, 21–24, 28–30, 36–38, 40–45, 47, 49–51, 54, 56, 60, 64, 73–78, 81–83

ECD cell egocentric cue direction cell. 8, 9, 21–25, 29, 36–38, 40, 49, 50, 52, 54, 57, 60, 61, 78, 89–91, 93

FGL first glance learning. v, vii, 3, 4, 29–31, 44, 55, 56, 60–62, 64, 66, 67, 69–74, 76–79, 81, 83

HD head direction. v, vii, 1–5, 7–19, 21–25, 27–36, 40–45, 47, 51–54, 56, 57, 59–62, 64–67, 69–71, 73–79, 81–84, 89, 93

HD cell head direction cell. v, vii, 1–5, 8, 9, 12, 15, 21–25, 27, 29–31, 33–35, 37, 40, 44–46, 51–54, 57, 59–61, 84, 89–91, 93

PE position encoder. 29–31, 44, 51, 54–56, 59, 60, 78, 79, 82, 93

PFD preferred firing direction. 34, 35

PI path integration. v, vii, 1–4, 10, 11, 13–15, 21, 23, 24, 27, 29, 32, 51, 52, 55–57, 60–62, 64, 66, 70, 71, 75–78, 81

SL shift layer. 32, 34–36, 52, 57, 90, 91

SLAM simultaneous localization and mapping. 1, 8

SLL shift left layer. 35, 36

SRL shift right layer. 35, 36

Bibliography

Alexander, A. S., L. C. Carstensen, J. R. Hinman, F. Raudies, G. W. Chapman, and M. E. Hasselmo (2020). "Egocentric boundary vector tuning of the retrosplenial cortex". In: *Science advances* 6.8.

Arleo, A. and W. Gerstner (2000). "Modeling rodent head-direction cells and place cells for spatial learning in bio-mimetic robotics". In: *From Animals to Animats* 6, pp. 236–245.

Arleo, A. and L. Rondi-Reig (2007). "Multimodal sensory integration and concurrent navigation strategies for spatial cognition in real and artificial organisms". In: *Journal of integrative neuroscience* 6.03, pp. 327–366.

Artal, P., P. H. de TEJADA, C. M. TEDÓ, and D. G. GREEN (1998). "Retinal image quality in the rodent eye". In: *Visual neuroscience* 15.4, pp. 597–605.

Avens, L. I. (2003). *Homing behavior, navigation, and orientation of juvenile sea turtles*. The University of North Carolina at Chapel Hill.

Bicanski, A. and N. Burgess (2016). "Environmental anchoring of head direction in a computational model of retrosplenial cortex". In: *Journal of Neuroscience* 36.46, pp. 11601–11618.

– (2018). "A neural-level model of spatial memory and imagery". In: *Elife* 7, e33752.

Busse, L., A. Ayaz, N. T. Dhruv, S. Katzner, A. B. Saleem, M. L. Schölvinck, A. D. Zaharia, and M. Carandini (2011). "The detection of visual contrast in the behaving mouse". In: *Journal of Neuroscience* 31.31, pp. 11351–11361.

Buzsáki, G. and E. I. Moser (2013). "Memory, navigation and theta rhythm in the hippocampal-entorhinal system". In: *Nature neuroscience* 16.2, pp. 130–138.

Byrne, P., S. Becker, and N. Burgess (2007). "Remembering the past and imagining the future: a neural model of spatial memory and imagery." In: *Psychological review* 114.2, p. 340.

Campbell, M. G., S. A. Ocko, C. S. Mallory, I. I. Low, S. Ganguli, and L. M. Giocomo (2018). "Principles governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation". In: *Nature neuroscience* 21.8, pp. 1096–1106.

Cartwright, B. and T. S. Collett (1983). "Landmark learning in bees". In: *Journal of comparative physiology* 151.4, pp. 521–543.

Chan, E., O. Baumann, M. A. Bellgrove, and J. B. Mattingley (2012). "From objects to landmarks: the function of visual location information in spatial navigation". In: *Frontiers in psychology* 3, p. 304.

Cheng, K. (1986). "A purely geometric module in the rat's spatial representation". In: *Cognition* 23.2, pp. 149–178.

Clark, B. J., M. J. Harris, and J. S. Taube (2012). "Control of anterodorsal thalamic head direction cells by environmental boundaries: comparison with conflicting distal landmarks". In: *Hippocampus* 22.2, pp. 172–187.

Clark, B. J., C. M. Simmons, L. E. Berkowitz, and A. A. Wilber (2018). "The retrosplenial-parietal network and reference frame coordination for spatial navigation." In: *Behavioral neuroscience* 132.5, p. 416.

Clark, B. J. and J. S. Taube (2011). "Intact landmark control and angular path integration by head direction cells in the anterodorsal thalamus after lesions of the medial entorhinal cortex". In: *Hippocampus* 21.7, pp. 767–782.

Collett, T. S. and P. Graham (2004). "Animal navigation: path integration, visual landmarks and cognitive maps". In: *Current Biology* 14.12, R475–R477.

Collett, T., B. Cartwright, and B. Smith (1986). "Landmark learning and visuo-spatial memories in gerbils". In: *Journal of Comparative Physiology A* 158.6, pp. 835–851.

Coumans, E. and Y. Bai (2016). *2019. PyBullet, a Python module for physics simulation for games, robotics and machine learning*.

Deshmukh, S. S. and J. J. Knierim (2013). "Influence of local objects on hippocampal representations: Landmark vectors and memory". In: *Hippocampus* 23.4, pp. 253–267.

Eilam, D. and I. Golani (1989). "Home base behavior of rats (Rattus norvegicus) exploring a novel environment". In: *Behavioural brain research* 34.3, pp. 199–211.

Etienne, A. S., J. Berlie, J. Georgakopoulos, and R. Maurer (1998). "Role of dead reckoning in navigation." In:

Etienne, A. S., R. Maurer, V. Boulens, A. Levy, and T. Rowe (2004). "Resetting the path integrator: a basic condition for route-based navigation". In: *Journal of Experimental Biology* 207.9, pp. 1491–1508.

Evans, T., A. Bicanski, D. Bush, and N. Burgess (2016). "How environment and self-motion combine in neural representations of space". In: *The Journal of physiology* 594.22, pp. 6535–6546.

Fetsch, C. R., G. C. DeAngelis, and D. E. Angelaki (2013). "Bridging the gap between theories of sensory cue integration and the physiology of multisensory neurons". In: *Nature Reviews Neuroscience* 14.6, pp. 429–442.

Gallistel, C. R. (1990). *The organization of learning.* The MIT Press.

Giocomo, L. M. (2016). "Environmental boundaries as a mechanism for correcting and anchoring spatial maps". In: *The Journal of physiology* 594.22, pp. 6501–6511.

Girshick, A. R., M. S. Landy, and E. P. Simoncelli (2011). "Cardinal rules: visual orientation perception reflects knowledge of environmental statistics". In: *Nature neuroscience* 14.7, pp. 926–932.

Goodridge, J. P., P. A. Dudchenko, K. A. Worboys, E. J. Golob, and J. S. Taube (1998). "Cue control and head direction cells." In: *Behavioral neuroscience* 112.4, p. 749.

Goodridge, J. P. and J. S. Taube (1995). "Preferential use of the landmark navigational system by head direction cells in rats." In: *Behavioral neuroscience* 109.1, p. 49.

Greene, C. M. and R. G. Cook (1997). "Landmark geometry and identity controls spatial navigation in rats". In: *Animal Learning & Behavior* 25.3, pp. 312–323.

Hardcastle, K., S. Ganguli, and L. M. Giocomo (2015). "Environmental boundaries as an error correction mechanism for grid cells". In: *Neuron* 86.3, pp. 827–839.

Henshaw, R. E. and R. O. Stephenson (1974). "Homing in the gray wolf (Canis lupus)". In: *Journal of Mammalogy* 55.1, pp. 234–237.

Hermer, L. and E. S. Spelke (1994). "A geometric process for spatial reorientation in young children". In: *Nature* 370.6484, pp. 57–59.

Hills, T. T., P. M. Todd, D. Lazer, A. D. Redish, I. D. Couzin, C. S. R. Group, et al. (2015). "Exploration versus exploitation in space, mind, and society". In: *Trends in cognitive sciences* 19.1, pp. 46–54.

Hinman, J. R., G. W. Chapman, and M. E. Hasselmo (2019). "Neuronal representation of environmental boundaries in egocentric coordinates". In: *Nature communications* 10.1, pp. 1–8.

Hinman, J. R., H. Dannenberg, A. S. Alexander, and M. E. Hasselmo (2018). "Neural mechanisms of navigation involving interactions of cortical and subcortical structures". In: *Journal of neurophysiology* 119.6, pp. 2007–2029.

Howard, I. P. and I. P. Howard (1982). *Human visual orientation*. John Wiley & Sons.

Høydal, Ø. A., E. R. Skytøen, S. O. Andersson, M.-B. Moser, and E. I. Moser (2019). "Object-vector coding in the medial entorhinal cortex". In: *Nature* 568.7752, pp. 400–404.

Hunter, J. D. (2007). "Matplotlib: A 2D graphics environment". In: *Computing in science & engineering* 9.3, pp. 90–95.

Jacob, P.-Y., G. Casali, L. Spieser, H. Page, D. Overington, and K. Jeffery (2017). "An independent, landmark-dominated head-direction signal in dysgranular retrosplenial cortex". In: *Nature neuroscience* 20.2, pp. 173–175.

Jain, D., I. R. Jakhalekar, and S. S. Deshmukh (2017). "Navigational strategies and their neural correlates". In: *Journal of the Indian Institute of Science* 97.4, pp. 511–525.

Jeffery, K. J., H. J. Page, and S. M. Stringer (2016). "Optimal cue combination and landmark-stability learning in the head direction system". In: *The Journal of physiology* 594.22, pp. 6527–6534.

Jercog, P. E., Y. Ahmadian, C. Woodruff, R. Deb-Sen, L. F. Abbott, and E. R. Kandel (2019). "Heading direction with respect to a reference point modulates place-cell activity". In: *Nature communications* 10.1, pp. 1–8.

Julier, S. J. and J. K. Uhlmann (2001). "Building a million beacon map". In: *Sensor Fusion and Decentralized Control in Robotic Systems IV*. Vol. 4571. International Society for Optics and Photonics, pp. 10–21.

Knierim, J. J. and D. A. Hamilton (2011). "Framing spatial cognition: neural representations of proximal and distal frames of reference and their roles in navigation". In: *Physiological reviews* 91.4, pp. 1245–1279.

Knierim, J. J., H. S. Kudrimoti, and B. L. McNaughton (1995). "Place cells, head direction cells, and the learning of landmark stability". In: *Journal of Neuroscience* 15.3, pp. 1648–1659.

Knight, R., C. E. Piette, H. Page, D. Walters, E. Marozzi, M. Nardini, S. Stringer, and K. J. Jeffery (2014). "Weighted cue integration in the rodent head direction system". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 369.1635, p. 20120512.

Learmonth, A. E., L. Nadel, and N. S. Newcombe (2002). "Children's use of landmarks: Implications for modularity theory". In: *Psychological Science* 13.4, pp. 337–341.

Leonard, J. J. and H. F. Durrant-Whyte (1991). "Mobile robot localization by tracking geometric beacons". In: *IEEE Transactions on robotics and Automation* 7.3, pp. 376–382.

Lever, C., S. Burton, A. Jeewajee, J. O'Keefe, and N. Burgess (2009). "Boundary vector cells in the subiculum of the hippocampal formation". In: *Journal of Neuroscience* 29.31, pp. 9771–9777.

Lever, C., T. Wills, F. Cacucci, N. Burgess, and J. O'Keefe (2002). "Long-term plasticity in hippocampal place-cell representation of environmental geometry". In: *Nature* 416.6876, pp. 90–94.

Lozano Navarro, Y. (2016). "Landmark processing by head direction cells". PhD thesis. UCL (University College London).

Margules, J. and C. Gallistel (1988). "Heading in the rat: Determination by environmental shape". In: *Animal Learning & Behavior* 16.4, pp. 404–410.

McNaughton, B. L., F. P. Battaglia, O. Jensen, E. I. Moser, and M.-B. Moser (2006). "Path integration and the neural basis of the'cognitive map'". In: *Nature Reviews Neuroscience* 7.8, pp. 663–678.

McNaughton, B. L., J. J. Knierim, and M. A. Wilson (1995). "Vector encoding and the vestibular foundations of spatial cognition: neurophysiological and computational mechanisms." In:

Microsoft Corporation (2010). *Xbox 360 Kinect Sensor*.

Milford, M. J., G. F. Wyeth, and D. Prasser (2004). "RatSLAM: a hippocampal model for simultaneous localization and mapping". In: *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*. Vol. 1. IEEE, pp. 403–408.

Mironitchev, A. and W. Morales (2019). "Trigonometric Proofs for Congruent Triangles". In: *Journal of Universality of Global Education Issues* 5.

Mishkin, M., L. G. Ungerleider, and K. A. Macko (1983). "Object vision and spatial vision: two cortical pathways". In: *Trends in neurosciences* 6, pp. 414–417.

Mittelstaedt, M. L. and H. Mittelstaedt (1980). "Homing by path integration in a mammal". In: *Die Naturwissenschaften* 67.11, pp. 566–567.

Monaco, J. D., G. Rao, E. D. Roth, and J. J. Knierim (2014). "Attentive scanning behavior drives one-trial potentiation of hippocampal place fields". In: *Nature neuroscience* 17.5, pp. 725–731.

Muir, G. M., J. E. Brown, J. P. Carey, T. P. Hirvonen, C. C. Della Santina, L. B. Minor, and J. S. Taube (2009). "Disruption of the head direction cell signal after occlusion of the semicircular canals in the freely moving chinchilla". In: *Journal of Neuroscience* 29.46, pp. 14521–14533.

Muller, R. U., J. L. Kubie, and J. B. Ranck (1987). "Spatial firing patterns of hippocampal complex-spike cells in a fixed environment". In: *Journal of Neuroscience* 7.7, pp. 1935–1950.

NVIDIA, P. Vingelmann, and F. H. Fitzek (2020). *CUDA, release: 10.2.89*. URL: `https://developer.nvidia.com/cuda-toolkit`.

O'Keefe, J. and J. Dostrovsky (1971). "The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat." In: *Brain research*.

O'keefe, J. and L. Nadel (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.

Page, H. J. and K. J. Jeffery (2018). "Landmark-based updating of the head direction system by retrosplenial cortex: a computational model". In: *Frontiers in cellular neuroscience* 12, p. 191.

Page, H. J., D. M. Walters, R. Knight, C. E. Piette, K. J. Jeffery, and S. M. Stringer (2014). "A theoretical account of cue averaging in the rodent head direction system". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 369.1635, p. 20130283.

Pearce, J. M., M. A. Good, P. M. Jones, and A. McGregor (2004). "Transfer of spatial behavior between different environments: implications for theories of spatial learning and for the role of the hippocampus in spatial learning." In: *Journal of Experimental Psychology: Animal Behavior Processes* 30.2, p. 135.

Pereira, U. and N. Brunel (2018). "Attractor dynamics in networks with learning rules inferred from in vivo data". In: *Neuron* 99.1, pp. 227–238.

Poulter, S., T. Hartley, and C. Lever (2018). "The neurobiology of mammalian navigation". In: *Current Biology* 28.17, R1023–R1042.

Redish, A. D. et al. (1999). *Beyond the cognitive map: from place cells to episodic memory*. MIT press.

Riisgaard, S. and M. R. Blas (2003). "SLAM for Dummies". In: *A Tutorial Approach to Simultaneous Localization and Mapping* 22.1-127, p. 126.

Robertson, R. G., E. T. Rolls, P. Georges-François, and S. Panzeri (1999). "Head direction cells in the primate pre-subiculum". In: *Hippocampus* 9.3, pp. 206–219.

Rogers, L. L. (1988). "Homing tendencies of large mammals: a review". In: *Translocation of wild animals. Milwaukee: Wisconsin Humane Society*, pp. 76–92.

Roth, M. M., F. Helmchen, and B. M. Kampa (2012). "Distinct functional properties of primary and posteromedial visual area of mouse neocortex". In: *Journal of Neuroscience* 32.28, pp. 9716–9726.

Sensortec, B. (Nov. 2020). *Intelligent 9-axis absolute orientation sensor*. BNO055. Rev. 17.

El-Sewisy, A. (Sept. 2020). "Biologically Inspired Robotic Mapping, Localization and Navigation using Spiking Neural Networks". MA thesis. Technical University of Munich.

Skaggs, W., J. Knierim, H. Kudrimoti, and B. McNaughton (1995). "A model of the neural basis of the rat's sense of direction." English (US). In: *Advances in neural information processing systems* 7, pp. 173–180. ISSN: 1049-5258.

Solstad, T., C. N. Boccara, E. Kropff, M.-B. Moser, and E. I. Moser (2008). "Representation of geometric borders in the entorhinal cortex". In: *Science* 322.5909, pp. 1865–1868.

Stanford Artificial Intelligence Laboratory et al. (May 23, 2018). "Robotic Operating System". Version ROS Melodic Morenia. In: URL: `https://www.ros.org`.

Stürzl, W., A. Cheung, K. Cheng, and J. Zeil (2008). "The information content of panoramic images I: The rotational errors and the similarity of views in rectangular experimental arenas." In: *Journal of Experimental Psychology: Animal Behavior Processes* 34.1, p. 1.

Tafazoli, S., A. Di Filippo, and D. Zoccolan (2012). "Transformation-tolerant object recognition in rats revealed by visual priming". In: *Journal of Neuroscience* 32.1, pp. 21–34.

Taube, J. S. (2007). "The head direction signal: origins and sensory-motor integration". In: *Annu. Rev. Neurosci.* 30, pp. 181–207.

Taube, J. S., R. U. Muller, and J. B. Ranck (1990). "Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis". In: *Journal of Neuroscience* 10.2, pp. 420–435.

Valerio, S. and J. S. Taube (2012). "Path integration: how the head direction signal maintains and corrects spatial orientation". In: *Nature neuroscience* 15.10, pp. 1445–1453.

Wallraff, H. G. (2001). "Navigation by homing pigeons: updated perspective". In: *Ethology Ecology & Evolution* 13.1, pp. 1–48.

Wang, C., X. Chen, H. Lee, S. S. Deshmukh, D. Yoganarasimha, F. Savelli, and J. J. Knierim (2018). "Egocentric coding of external items in the lateral entorhinal cortex". In: *Science* 362.6417, pp. 945–949.

Whishaw, I. Q. and D. G. Wallace (2003). "On the origins of autobiographical memory". In: *Behavioural Brain Research* 138.2, pp. 113–119.

Wijngaarden, J. B. van, S. S. Babl, and H. T. Ito (2020). "Entorhinal-retrosplenial circuits for allocentric-egocentric transformation of boundary coding". In: *Elife* 9, e59816.

Wilber, A. A., B. J. Clark, T. C. Forster, M. Tatsuno, and B. L. McNaughton (2014). "Interaction of egocentric and world-centered reference frames in the rat posterior parietal cortex". In: *Journal of Neuroscience* 34.16, pp. 5431–5446.

Wolbers, T. and J. M. Wiener (2014). "Challenges for identifying the neural mechanisms that support spatial navigation: the impact of spatial scale". In: *Frontiers in human neuroscience* 8, p. 571.

Yoder, R. M., B. J. Clark, J. E. Brown, M. V. Lamia, S. Valerio, M. E. Shinder, and J. S. Taube (2011). "Both visual and idiothetic cues contribute to head direction cell stability during navigation along complex routes". In: *Journal of Neurophysiology* 105.6, pp. 2989–3001.

Yoder, R. M., B. J. Clark, and J. S. Taube (2011). "Origins of landmark encoding in the brain". In: *Trends in neurosciences* 34.11, pp. 561–571.

Yoder, R. M., J. R. Peck, and J. S. Taube (2015). "Visual landmark information gains control of the head direction signal at the lateral mammillary nuclei". In: *Journal of Neuroscience* 35.4, pp. 1354–1367.

Yoder, R. M. and J. S. Taube (2009). "Head direction cell activity in mice: robust directional signal depends on intact otolith organs". In: *Journal of Neuroscience* 29.4, pp. 1061–1076.

Zeil, J. (2012). "Visual homing: an insect perspective". In: *Current opinion in neurobiology* 22.2, pp. 285–293.

Zhang, K. (1996). "Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory". In: *Journal of Neuroscience* 16.6, pp. 2112–2126.

Zhang, W.-H., H. Wang, A. Chen, Y. Gu, T. S. Lee, K. M. Wong, and S. Wu (2019). "Complementary congruent and opposite neurons achieve concurrent multisensory integration and segregation". In: *Elife* 8, e43753.

Zugaro, M. B., A. Arleo, A. Berthoz, and S. I. Wiener (2003). "Rapid spatial reorientation and head direction cells". In: *Journal of Neuroscience* 23.8, pp. 3478–3482.

Zugaro, M. B., A. Berthoz, and S. I. Wiener (2001). "Background, but not foreground, spatial cues are taken as references for head direction responses by rat anterodorsal thalamus neurons". In: *Journal of Neuroscience* 21.14, RC154–RC154.