4

Embodied Neural Dynamics

GREGOR SCHÖNER, CHRISTIAN FAUBEL, EVELINA DINEVA, AND ESTELA BICHO

Over the first three chapters we have seen how neural dynamics goes well beyond the feed-forward processing of sensory input toward motor output. In particular, the neural dynamics of activation variables and fields creates stable states. We have shown that these stable states may instantiate decisions, which are stabilized even as sensory inputs change. Instabilities are the critical points at which states change qualitatively, such as when a decision is first made. Memory traces keep track of past states.

Even though stability was a central postulate in our approach exactly because behavior and cognition may be continuously linked to changing sensory inputs, the organisms we have modeled thus far have been entirely passive and static; they were sitting and waiting for inputs to arrive. Inputs were perhaps structured in space and time, but we did not account for how the spatial and temporal structure of inputs emerges from the behavior of the organism as it moves its body or actuator in a structured environment.

In this chapter we will address how closing the sensory-motor loop affects the neural dynamics and, conversely, how neural dynamics may be embodied in an organism that behaves in a closed sensory-motor loop. This entails recognizing that closing the loop through the outer world creates another level of dynamics, a level we refer to as *behavioral dynamics*. To formalize these dynamics, we will need to introduce new variables that characterize the state of the embodied system within the structured environment in which the organism is situated. Behavioral dynamics is more than control, because it achieves more than keeping the organism aligned with a given goal state. Behavioral dynamics may induce decisions that emerge from the structure of the environment.

Behavioral variables are not the same as activation variables. We will need to understand how neural dynamics with their activation states are linked to behavioral dynamics. That will be a major issue in this chapter. Finally, we will show how the combined neural and behavioral dynamics of an organism brings about autonomy. We will use the A-not-B paradigm of perseverative action to illustrate these ideas.

BEHAVIORAL DYNAMICS IN A BRAITENBERG VEHICLE

Recall the Braitenberg vehicle evoked in the Introduction to Part I of the book. This is a conceptual organism defined by four things: It has sensors, effectors, a body linking the sensors to effectors mechanically, and a nervous system linking the sensor to effectors through activation variables. The organism is also situated in a structured environment, which is critical for organized behavior to emerge.

The activation concept introduced in Chapters 1–3 now makes it easier to make sense of Braitenberg's ideas (Braitenberg, 1984) and helps us understand the difference between neural and behavioral dynamics. We will use the "taxis" vehicle depicted in Figure 4.1 to illustrate (this is Braitenberg's vehicle 3a). The taxis organism has two sensors and two effectors, whose neural connectivity is organized ipsilaterally. The sensors are characterized by a monotonically decreasing sensor characteristic, that is, a decreasing mapping of a physical intensity onto an activation variable. For a light sensor, for instance, this may mean that higher light intensity shining onto the sensor leads to lower activation levels at the output of the sensor. The motor systems of the taxis organism are characterized by a monotonically increasing motor

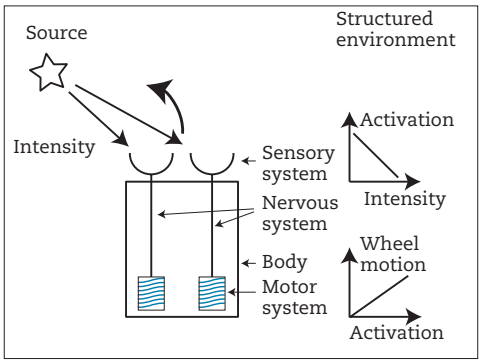

FIGURE 4.1: A "taxis" vehicle of Valentino Braitenberg (center) situated in an environment in which one source (represented by a star) creates a physical intensity pattern (of sound intensity for a sound source, of chemical concentration for a source of some chemical agent, or of temperature for a heat source, for instance). The cups at the front of the vehicle illustrate sensors that are sensitive to the physical intensity at the location in space that they sample. The sensor characteristic shown on the right describes the level of activation at the output of the sensor as a function of the intensity that impinges on the sensory surface. Illustrated is the case where this characteristic is monotonically decreasing. The patterned small squares represent effectors, conceptualized as two self-motorized wheels (seen from above; think of Formula 1 racing tires). Their motor characteristic shown on the right is monotonically increasing, so the wheels turn faster when higher levels of activation are presented at input. The large square represents the body: When the wheels turn, they move the body that the sensors are attached to. The two vertical lines connecting the sensors to the motors are a simple nervous system. They indicate that the activation output by each sensor is passed on as input to the ipsilateral motor system. Taxis behavior, turning toward the source, is hinted at by the curved arrow: Because the sensor on the left is closer to the source, it is assumed to encounter a higher level of intensity than the sensor on the right. As a result, the left sensor sends lower levels of activation to its motor, which thus turns more slowly than the motor on the right, leading the vehicle to drive in a leftward curve.

characteristic, that is, an increasing mapping from an activation level at input to a physical motor action at output. In the conceptual vehicle, physical action is the rotation of the wheel, so larger action output means a higher rate of wheel rotation.

In the Introduction to Part I, we reviewed Braitenberg's verbal account of how this particular arrangement of sensors and motors leads to the behavior of taxis—the orientation to sources of whatever physical intensity the sensors are tuned to. This account was based on the assumption that the environment provides gradient fields of physical intensities, so the two sensors pick up a difference in intensity that indicates the side that the source lies on (higher intensity on the side closer to the source). This difference translates into a difference in activation, with lower activation on the side closer to the source. That difference is handed down to the motors, leading to less wheel rotation on the side closer to the source, which leads the vehicle to turn toward the source.

In this account, activation plays a very limited role; activation merely transduces sensed intensities into motor actions in a one-to-one mapping. We assumed that only positive levels of activation arose and thus did not consider sigmoidal threshold functions. The nervous system of this simple conceptual organism was, therefore, organized in a purely feed-forward fashion. The behavior that emerges, however, closes a sensory-motor loop. At any moment in time, the difference in intensity at the two sensors brings about the turning action of the robot. That turning action then determines how the orientation relative to the source changes as the vehicle advances, leading to a new, reduced value of the difference in intensity sensed on the left and the right. By mentally iterating this closed loop of sensing and acting, we intuitively simulate a dynamical system in which the current orientation of the vehicle relative to the source determines the vehicle's direction and rate of turning.

Braitenberg's goal when he proposed his vehicles was to illustrate how structural principles of neuroanatomy could manifest at the level of function. The taxis vehicle, for instance, served to contrast its ipsilateral neural organization with the contralateral organization of another vehicle (his vehicle 3b) that creates avoidance behavior. He did not formalize the structure–function relationship in his thought experiments with vehicles and did not recognize that dynamical stability plays a critical role in the emergence of behavior (and of cognition in his more advanced vehicles). Our goal now is to formalize the intuitive dynamics implicit in the verbal functional analysis of the vehicles. We will uncover that the function that emerges in the vehicles derives from an implied behavioral dynamics and that stability determines the functions that emerge.

To achieve this goal, we need two things. First, we need a variable that captures the state of the organism within the closed sensory-motor loop. The obvious candidate for such a "behavioral" variable is the orientation of the organism relative to its

environment, the vehicle's heading direction. The heading direction is measured as the angle that the long axis of the vehicle's wheels forms with an arbitrary but fixed world axis. (Later we will address how that world frame is calibrated.) Second, we need a model of the environment that describes the intensities that the sensors will be exposed to given the current orientation of the vehicle relative to the source. In Figure 4.2, we use these two ingredients to derive a dynamical system model of the taxis vehicle for the case in which a single source is present in the environment. To make things simple, we look at the vehicle as it is heading directly toward the source, as sketched at the bottom of the figure. The top panel shows a model of the environment in which the intensity sensed by the vehicle peaks when the vehicle points directly at the source and then falls off as the vehicle turns away from the source. The two sensors on the vehicle point in slightly different directions, so computing the difference between intensities picked up by the left and right sensors amounts to estimating the inverse slope of the intensity profile. To the left of the source, the difference is negative, as the left sensor picks up less intensity than the right sensor. To the right of the source, the difference is positive. When pointing right at the source, the intensities picked up on the left and right sides of the vehicle are identical, so the difference goes through zero at that heading direction. By concatenating the sensor and motor characteristics shown in Figure 4.1, we eliminate activation as a variable and obtain the dependence of each wheel's turning rate on the intensity picked up by the sensor on the same side. That leads to a linear function with a negative slope. The difference between the left and right sides yields the linearly decreasing function shown in the third panel of Figure 4.2. Finally, we may concatenate the functions in the second and third panels of Figure 4.2, eliminating the difference between the intensity sensed on the left and right sides, to directly obtain the difference in turning rate of the left and right wheel as a function of heading direction. The difference in turning rate of the left and right wheel is proportional to the turning rate of the vehicle. This follows from the model of the body to which wheels and sensors are attached. If left and right wheels turn at the same rate, the vehicle moves on a straight path and does not turn. If the left wheel turns more than the right wheel, the vehicle turns to the right, increasing its heading angle. The bottom panel thus shows the functional dependence of the vehicle's turning rate as a function of its heading direction.

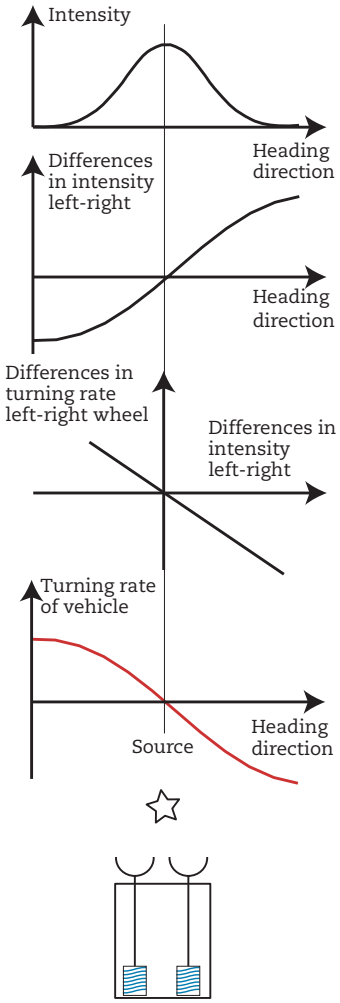

FIGURE 4.2: Model of the taxis vehicle of Braitenberg in the environment sketched at the bottom of the figure, in which the vehicle points toward the single source of intensity. *Top*: The environment is modeled by assuming that the sensed intensity is a bell-shaped function of the heading direction of the vehicle relative to the source, which peaks when the vehicle points exactly at the source. That heading is marked by the thin vertical line that cuts through all panels. *Second from top*: The difference in intensity sensed at the left and right sensors is computed by sampling the model shown in the top panel, at two locations corresponding to the two sensors and computing their difference. *Third from top*: The difference in turning rate of the left and right wheels as a function of the difference in intensity sensed at the left and right sensors is computed by concatenating the sensor and motor characteristic shown in Figure 4.1 and computing the difference, left minus right. *Fourth panel from top*: The difference in turning rate of the left and right wheels is proportional to the turning rate of the vehicle. By concatenating the mappings illustrated in the second and third panels, the turning rate of the vehicle is obtained as a function of its heading.

Note how this derivation makes use of five ingredients: (1) a model of the environment (top); (2) a sensor model (through the sensor characteristic of Figure 4.1); (3) a motor model (through the motor characteristic of Figure 4.1); (4) a model of the body (linking the turning rate of the vehicle to the difference in turning rate of the wheels); and (5) a model of the nervous system (that enabled concatenating sensor and motor characteristics). What we obtain from this derivation is a dynamical system model of the behavior of the taxis vehicle in closed loop: The turning rate of the vehicle is nothing other than the derivative in time, $\dot{\phi}$, of the heading direction, ϕ! Thus, we formally have the functional dependence

$$\dot{\phi} = f(\phi) \tag{4.1}$$

where f is the function depicted in the bottom panel of Figure 4.2. That is a differential equation that mathematically defines a dynamical system. We call this the *behavioral dynamics* of the taxis vehicle.

Figure 4.3 highlights that behavior emerges from this behavioral dynamics through an attractor state, a stable fixed point of the behavioral variable. In the figure, the vehicle is oriented to the right of the source. The behavioral dynamics is thus sampled at a heading direction to the right of the zero-crossing, generating a negative turning rate of the vehicle. The vehicle will thus reduce its heading direction, turning to the left, until the turning rate becomes zero exactly when the vehicle is oriented toward the source. Analogously, starting out at a heading direction to the left of the source will lead to positive rates of change, increasing heading direction by turning right, again toward the source. As we saw in Chapter 1 for the dynamics of neural activation, a zero-crossing of the dynamics with a negative slope is an attractor, a stable fixed point, now of the behavioral dynamics. That attractor generates the taxis behavior, the behavior of orienting to the source of intensity.

The attractor dynamics thus determines the orientation behavior of the taxis vehicle. This dynamics does not really depend on how the heading direction is measured or calibrated. Contrast, for instance, a calibration in which heading direction is measured relative to the magnetic north with a calibration in which heading direction is measured relative to the magnetic south. The two cases merely differ in how the labels read along the horizontal axis of the dynamics in Figure 4.3. The rate of change is determined by how the vehicle

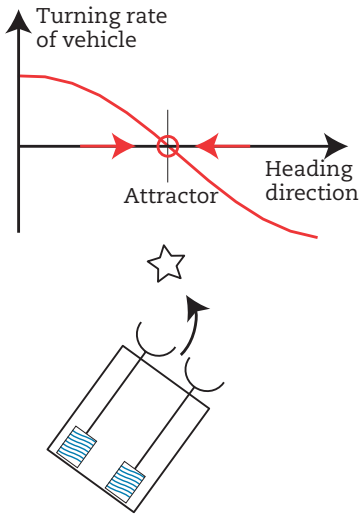

FIGURE 4.3: The dynamics of heading direction has a fixed point at the zero-crossing of the rate of change. When the vehicle's heading corresponds to the fixed point, the rotation rate is zero, so the vehicle remains oriented in that direction. When the vehicle is headed to the right of the fixed point as illustrated at the bottom, the negative turning rate drives the vehicle's heading direction toward the fixed point, as indicated by the red arrow pointing to the left. Similarly, if the vehicle were headed to the left of the fixed point, the positive turning rate would drive the vehicle's heading direction toward the fixed point, as illustrated by the red arrow pointing to the right. The convergence to the fixed point from neighboring states implies that the fixed point is asymptotically stable, a fixed point attractor (marked by a red circle).

is oriented relative to the source, and that relative orientation does not depend on the absolute values of heading direction. Moreover, what determines the movement of the vehicle is the rate of change of its heading direction, which is enacted by sending different commands to the two wheels (based on a simple computation that takes into account the size of the wheels and how far apart they are mounted on the vehicle). The rate of change of heading direction is independent of the reference frame used for heading direction itself. In a sense, the behavioral variable is, therefore, a somewhat abstract concept; it abstracts away from the detailed mechanisms of the sensory and effector systems. The behavioral dynamics provides, however, a process account for movement generation, because it enables generating the modeled behavior using generic sensor or motor models.

In this derivation of the behavioral dynamics from the architecture of the taxis vehicle, we did

not pay attention to the forward movement of the vehicle that is controlled by the average turning rate of the two wheels. Analogous thinking leads to a dynamical system description for that forward velocity that depends on the intensity levels in the environment and how steeply they vary with the distance from the source. Now think of the vehicle as moving forward while it is turning. As a result, the sensors of the vehicle will sample different locations in the environment at different points in time. This may change the intensity profile the vehicle is exposed to. For instance, at a large distance from the source, intensity may fall off more rapidly as heading direction varies than when the vehicle is rather close to the source. (Intuitively, a disk centered at the source describes the spatial range at which intensity has a given level. That disk will loom larger in heading direction when the vehicle is close to the source than when it is far from the source.) A changing intensity profile implies a changing behavioral dynamics! So as the vehicle moves around, the attractor and the negative slope of the dynamics that determines how strongly the heading direction is driven to the attractor may change.

This is not a problem. The attractor itself ensures that the behavioral variable tracks these changes. If the attractor shifts while the vehicle is moving, heading direction is continuously attracted toward the updated attractor. A problem only arises if the shift becomes too rapid for the behavioral variable to follow. That is a real problem that organisms have, too. For instance, you are able to move toward a ball to pick it up. When the ball moves while you try to pick it up, you can update the direction in which you are headed and may be still able to catch the ball. But if the ball moves too fast, like a really fast serve in tennis (or, at the extreme, a cannon ball), then the same updating mechanism will ultimately fail.

The behavioral dynamics changes not only as the vehicle moves around a given environment; different environments create different behavioral dynamics. Figure 4.4 illustrates how an environment with two sources may induce a bimodal intensity profile. By the same logic we used previously, such a profile will give rise to a behavioral dynamics that now has two attractors, one for each local maximum of the intensity profile. The attractors divide the space of possible heading directions into two basins of attraction: One set of initial heading directions leads the vehicle to turn toward one source, the other set leads the vehicle to turn toward the other source. The two basins of attraction are separated by another fixed point, this one with a positive slope of the rate of change, making it a repellor.

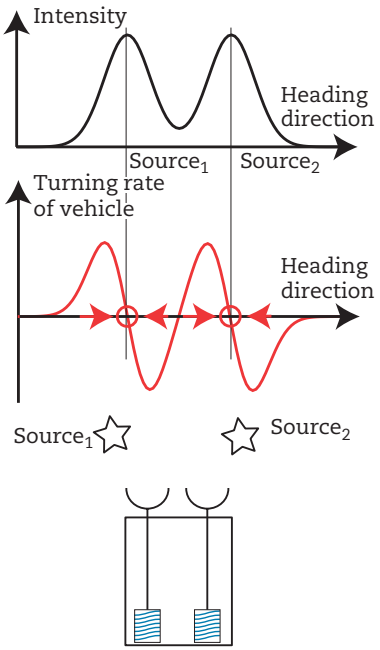

FIGURE 4.4: With two sources of intensity in the environment (assumed of equal strength here), the intensity profile impinging on the vehicle in the symmetric position sketched at the bottom is bimodal, as shown at the top. This leads to a behavioral dynamics of heading direction shown in the middle. This dynamics has two attractors (circles), one at each local maximum of the intensity profile. The third zero-crossing between the two attractors is a repellor. Initial heading directions to the right of the repellor converge to the rightmost attractor, as indicated by the arrows. To see this, consider the sign of the turning rate. Similarly, initial heading directions to the left of the repellor converge to the leftmost attractor.

The coexistence of two attractors, called *bistability*, leads to a selection decision. The initial orientation of the vehicle determines which basin of attraction its heading direction lies in. This determines which attractor the vehicle's heading direction converges to und thus leads to selection of one of the two sources as the target of taxis behavior. If the direction in which the vehicle is headed is initially close to the boundary of the areas of attraction as in Figure 4.4, then the attraction to one of the two stable fixed points pushes the vehicle's direction away from the boundary. In that sense,

the selection decision self-stabilizes. As the vehicle moves, the behavioral dynamics changes and the attractors may shift. Typically, however, the behavioral variable will track the attractor within the basin of attraction in which the variable was initially situated.

Recall that the nervous system of this simple taxis vehicle is purely feed-forward, thus a given input generates a unique output. Even so, when situated in an appropriately structured environment, the behavioral dynamics that emerges from the closed loop makes selection decisions. In this bistable regime, the sensory input no longer uniquely determines the motor behavior. The motor behavior depends instead on the state of the behavioral system. In Chapters 1 and 2 we saw how neural dynamics with strong neural interaction may lead to bistability (and multistability) and how this is a qualitative change from the unique input–output mapping of forward neural networks. In neural dynamics, the internal (recurrent) loops that instantiate neural interaction may break the unique input–output mapping. In behavioral dynamics, the sensory-motor loops through the environment may break unique input–output mappings.

In neural dynamics we saw that instabilities lead to such qualitative change. For instance, detection instability destroys the input-driven activation pattern. Instabilities play an analogous role in behavioral dynamics. This is illustrated in Figure 4.5, where the two sources are closer to each other than in Figure 4.4, so that now the intensity profiles induced by each source fuse to form a monomodal distribution with a single peak located over the averaged heading direction. This leads to a monostable dynamics with a single attractor at that averaged heading direction.

In a mental simulation, imagine an environment that changes continuously, starting out with a single source (as in Figures 4.2 and 4.3), which then splits into two sources that gradually move apart. Up to a critical separation of the two sources, the dynamics is monostable with an attractor at the averaged direction toward the sources (as in Figure 4.5). At a critical separation, the single attractor becomes a repellor, while at the same time two new attractors split off. This is an instability beyond which we find the bistable dynamics of Figure 4.4. The dependence of the fixed points and their stability on the distance between the two sources is illustrated in Figure 4.6. The bifurcation

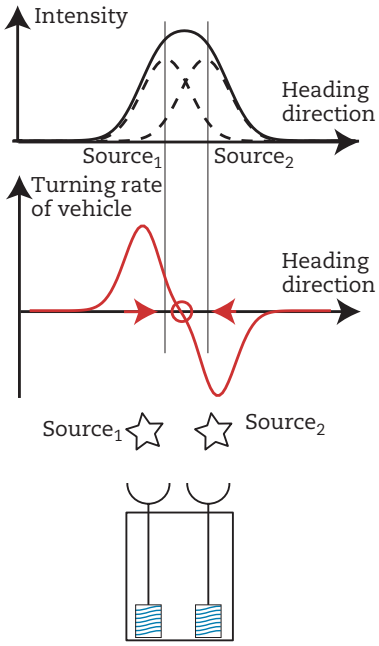

FIGURE 4.5: Shown is the same schema as in Figure 4.4, but now the two sources are at a closer angle. Their individual intensity profiles (top, dashed line) fuse into a monomodal intensity distribution (solid line) that peaks at a heading direction lying near the average of the headings of the two individual sources (marked by thin vertical lines). The behavioral dynamics is monostable, just as the dynamics of a single source is, shown in Figures 4.2 and 4.3.

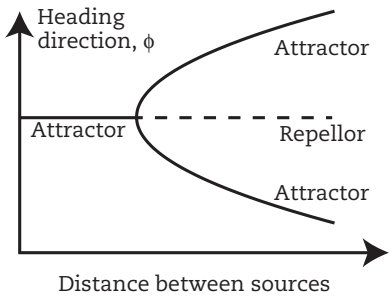

FIGURE 4.6: This bifurcation diagram plots the fixed points of the behavioral dynamics of Figures 4.2 as a function of the distance between two sources. For small distances, a single attractor (solid line on the left) is positioned over the average of the two heading directions under which the two sources are seen. At large distances, two attractors (solid lines on the right) and one repellor (dashed line on the right) exist. The transition occurs at an instability, in which the single fixed point becomes unstable and gives rise to two new stable states and one unstable state.

is called a *pitchfork bifurcation* because of the shape of this bifurcation. In this mental simulation of gradually increasing the separation between two sources of intensity, the capacity of the behavioral dynamics to make selection decisions thus emerges from an instability! This is analogous to how elementary forms of cognition emerged from neural dynamics in the different instabilities discussed in Chapters 1 and 2.

Note how the attractors structure the time course of behavior in a manner similar to how attractors of the neural dynamics structure the time course of neural activation. Most of the time, the behavioral variables are in or near an attractor, which they track as the attractors shifts when the vehicle moves around and the bearings of the sources change. Only exceptionally is there a chance to observe a transient where a behavioral variable switches to a new attractor. This happens exactly at bifurcations when a formerly stable state becomes unstable.

The emergence of the capacity to make selection decisions highlights a conceptual difference between the behavioral dynamics introduced here and biological cybernetics. Biological cybernetics is an older approach toward understanding the closed-loop behavior of organisms (e.g., Reichardt & Poggio, 1976). In cybernetic thinking, a sensory signal is coupled into a motor system so as to stabilize a particular goal state, the "set-point" of the closed-loop control system. The sensory signal is often conceived of as an error signal that reflects how the current state of the system deviates from the goal state. The control signal is designed to reduce this error. This view is conceptually not far from information-processing ideas in that the control signal is computed from the error signal in a feed-forward manner, although the closed loop and its stability are taken into account. The selection between two sources is not an obvious cybernetic task, however. It is not clear, for example, how the sensory data could be interpreted as an error signal—relative to which of the two sources should the error be assessed? This conceptual problem notwithstanding, even very simple organisms such as the house fly, studied by Reichardt and Poggio (1976), are able to make such selection decisions. The formalism of biological cybernetics is naturally and easily generalized to the behavioral dynamics introduced here.

The formalization of the Braitenberg vehicle as a behavioral dynamics enables us to make explicit the critical role that the structure of the environment plays in bringing about meaningful behavior. Imagine, for instance, that intensity was a highly irregular function of orientation (perhaps because there would be many sources with a sharp fall-off of intensity compared to the size of the vehicle). The sampling of such an intensity landscape by the two sensors would not lead to a coherent dynamics. The sensed intensity differences would appear to be largely random, and movement behavior would be highly irregular and unpredictable. Only when the environment is appropriately structured do attractor landscapes and consistent behaviors emerge. In the next section we will see how we can use neural fields to re-present environments such that consistent behavioral dynamics emerges.

LINKING DYNAMIC NEURAL FIELDS TO BEHAVIORAL DYNAMICS

In spite of the capacity to make selection decisions, the behavioral dynamics of the taxis vehicle is still very strongly linked to sensory input. Local maxima of the sensed intensity profile induce the different attractors. If, after selecting one local maximum, the vehicle turns due to some other behavior (e.g., driven by obstacle avoidance), it has no way of "remembering" which source it originally chose. It will move toward the attractor whose basin of attraction the heading direction falls into after the distraction. For selection decisions to withstand distraction additional dimensions are required. These serve as inner-state variables that keep track of the initial selection decision as the behavioral variable changes. We know, of course, from Chapters 1–3 about activation and activation fields that provide the dynamic substrate to achieve just that. In this section, we will look at how activation fields might make decisions less dependent on behavioral variables while remaining linked to sensory input.

We will use the robotic vehicle illustrated in Figure 4.7 to develop this point. This vehicle is designed to generate phonotaxis behavior, orientation to sound sources (Bicho, Mallet, & Schöner, 2000). For this purpose, it has an only slightly more complex sensory array than the Braitenberg taxis vehicle: Five microphones are mounted 45° apart so that they roughly sample the angular surroundings of the vehicle. Each microphone is directionally sensitive, with a sensitivity cone approximately 60° wide so that the sensitivity cones of two

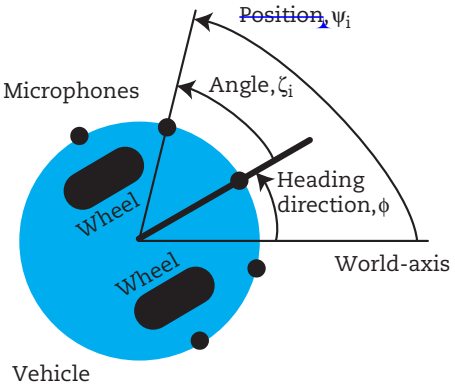

FIGURE 4.7: A robot vehicle seen from above (gray disk) has two active wheels that define an instantaneous heading direction (marked by the thick solid line). The heading direction, ϕ, is measured as the angle between this forward axis and a fixed world axis. Five directionally sensitive microphones (black filled circles) are mounted on the vehicle at fixed angles, ζ_i. The bearings of these sensors, ψ_i, are the directions in the world in which these sensors point.

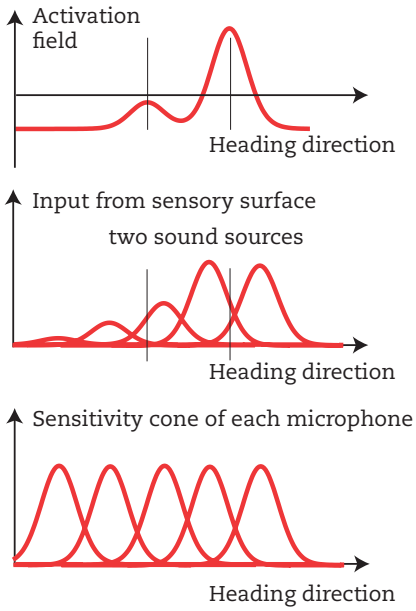

FIGURE 4.8: *Bottom*: Five microphones sample the space of possible bearings of a sound source. The sensitivity cone of the directionally selective microphones is modeled by a Gaussian, centered over the direction in the external frame that the microphone is pointing toward. This sensitivity cone can be conceived of as the tuning curve of the microphone and describes how input from the microphone is distributed within the activation field. *Middle*: This projection occurs by multiplying every tuning curve by the current intensity recorded by each microphone. The curves depicted in this panel result from exposure of the vehicle to two sound sources whose bearing is marked by thin vertical lines. *Top*: The activation field defined over heading direction receives the weighted tuning curves of the middle panel as input and is able to build localized peaks that represent detection and selection decisions about sound sources. In this example, the rightmost source is selected.

neighboring microphones overlap slightly. Figure 4.8 shows how input from these five microphones could drive an activation field, representing the estimated direction in which a sound source lies. The sensitivity cones of the five microphones are modeled as Gaussian functions. In the neural jargon of Chapter 3, these are tuning curves. If the loudness sensed by each microphone at any moment in time is multiplied by its tuning curve (middle panel of Figure 4.8), their superposition yields a sampled representation of the sound intensity profile at the vehicle's location. As for the taxis vehicle, local maxima of this profile could be viewed as estimates of the directions in which a sound source lies. Using the raw readings of the microphones in conjunction with the tuning curves to determine the direction in which a sound source lies has a number of limitations, however. First, the sound coming from the sound sources may vary over time. For instance, if the sound source is a loudspeaker that plays music, the intensity varies as dictated by the music. In the presence of other ambient sources of sound, the direction in which a local maximum of intensity lies may fluctuate wildly. To steer the vehicle toward the loudspeaker, we need to stabilize the estimate of the heading direction in which the loudspeaker lies. When the loudspeaker first comes into acoustic range of the vehicle, we want the vehicle to make a clear decision as to whether a source is present or not, a decision that is then stabilized as the music waxes and wanes. And, of course, we want the nervous system of the vehicle to select the loudspeaker over other distracting sound sources. For instance, maximal sound intensity may be detected in the direction of the loudspeaker, but reflection of the sound from a nearby wall may create an echo, which should be ignored.

We recognize, of course, that these are the typical requirements to which dynamic field theory (DFT) responds. So what we need is an activation field defined over heading direction, in which a

peak of activation represents the detection of a sound source and the location of the peak is an estimate of the direction in which the source lies (as illustrated in the top panel of Figure 4.8). Dynamic fields stabilize detection and selection decisions, as discussed in Chapter 2. The dynamic field takes the weighted tuning curves as inputs. This forward connectivity from the five microphones to the field is analogous to the neural connectivity from a sensory surface to a cortical representation. The forward projection from any location on a sensory surface to the cortical surface is given by the point-spread function in neurophysiology. Conversely, any location in the cortical surface receives input from a range of point-like sources on the sensory surface, mapping out the receptive field or tuning curve. The Gaussians of Figure 4.8 model both of these mappings.

But there is a snag: The direct mapping from the sensory surface to a cortical representation would naturally lie in a reference frame anchored to the sensory surface, here the vehicle itself on which the microphones are mounted at fixed angles, ζ_i ($i=1,\ldots,5$ counting the microphones, compare Figure 4.7). We postulate, however, that the field be defined over the heading direction measured against a fixed world axis. This makes sense, because we want to use the field to steer the vehicle toward sound sources. As the vehicle turns toward a sound source, the direction to the sound source would change if that direction was assessed in a frame attached to the vehicle. In contrast, in an external frame anchored in the world, the direction to the sound source is invariant under any rotation of the vehicle. This difference is critical when the vehicle selects a source and now needs to keep that selection stable even as it reorients under other influences, such as the obstacle avoidance mentioned earlier.

Representing sound sources in a frame anchored in the world while also linking the activation field to sensors mounted on the vehicle requires that we transform vehicle coordinates, ζ_i, into world coordinates, ψ_i. This coordinate transform requires an explicit representation of the vehicle's own heading direction, ϕ. In fact, mathematically, the transformation simply reads:

$$\psi_i = \zeta_i + \phi \tag{4.2}$$

as is obvious from Figure 4.7. For the taxis vehicle, we did not need to know the vehicle's heading direction, only its rate of change. But now, to explicitly represent the direction in the world in which a sensor is pointing we do need an estimate of this behavioral variable.

One simple way of obtaining this estimate is to solve the behavioral dynamics, Equation 4.1, by integrating it in time. This method is called *path integration* or "dead reckoning," based on an analogy with maritime navigation. Sailors used to estimate the position of their ship by integrating in time the speed of their vessel (which they were able to measure by a log). When they did that, the uncertainty about the location of the ship would grow over time, limiting how far they could go without finding some reference landmark. This was because any error in measuring time or speed would remain uncorrected and accumulate in the summation process. This is the problem of calibration, that is, of resetting such an estimate when ground truth is available. Miscalibration of the world frame of the vehicle does not matter, however, because the error in projecting from the sensory surface to the field is cancelled by making the same projection in reverse from the field to a motor command. We will see later that this is what the linking of the dynamic field to the dynamics of heading direction does.

Now, using Equation 4.2 is not really a neural operation. In Chapter 7 we will discuss in depth the neural principles on which such coordinate transformations are based. The upshot is that coordinate transforms are tunable mappings from one neural representation to another. These mappings can be organized to be bidirectional. They can turn less invariant sensory or motor variables into more invariant representations as required here. But they can also be used to predict a variant from an invariant representation, such as when you predict where on your retina a visual object will fall after a planned saccadic eye movement (Schneegans & Schöner, 2012).

After this excursion about reference frames, let's return to the dynamics of the activation field that represents the heading directions in which sound sources lie. Here are a few illustrations of the neural dynamics driven directly (via Equation 4.2) by sensors mounted on a robot vehicle (see Bicho et al., 2000, for details). For now, the vehicle is not moving (we disconnected the motors), so we are looking only at the neural, not at the behavioral, dynamics. Figure 4.9 illustrates the detection decision. A loudspeaker playing music is positioned in front of the vehicle. The volume of the music is gradually increased. The evolution over time of the

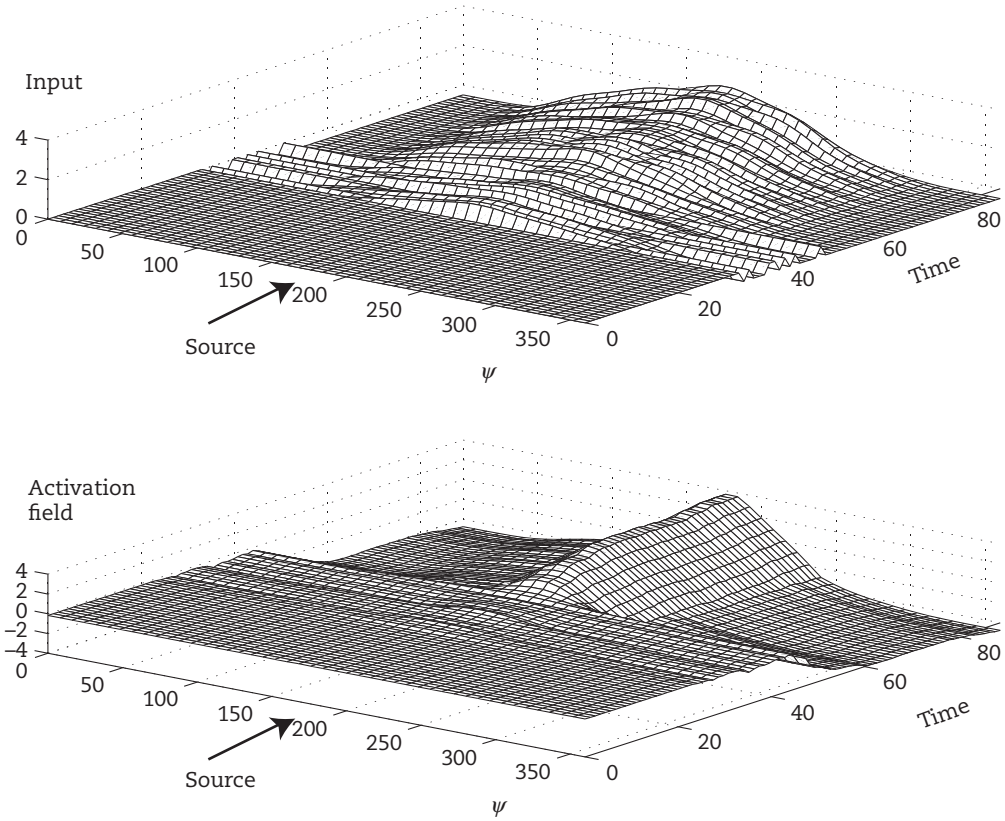

FIGURE 4.9: The weighted tuning curves of the five microphones are summed and plotted as a function of time in the *top* panel as a loudspeaker, whose bearing is marked by an arrow, plays music that increases continuously in loudness. The *bottom* panel shows the activation field driven by this input, which undergoes a detection instability at the moment in time marked as 50 time units.

raw signals from the five microphones, multiplied by the tuning curves of Figure 4.8 and summed, is shown at the top of the figure. These input profiles drive the activation field shown at the bottom of the figure. The detection instability occurs at a particular point in time when input strength reaches a critical level. At that point, the activation pattern switches to a self-stabilized peak that is stabilized from then on, even as input fluctuates. This detection event emerges at a discrete time from a time-continuous change of input.

Selection is illustrated in Figure 4.10, where the same robot is confronted with two loudspeakers. The field initially selects one of the two sources and suppresses the other. This selection takes place at the level of representation rather than at the level of overt motor behavior (discussed in Figure 4.4). Robust estimation is a variant of this form of selection. In Figure 4.11, a single loudspeaker is flanked on one side by a reflecting surface, so that the angular distance from the loudspeaker at which sound is picked up extends further to the right than to the left of the loudspeaker. The field positions an activation peak over the local maximum of input, effectively suppressing the outliers that come from reflected sound. This is a form of robust estimation. That the peak of activation is continuously linked to input is illustrated in Figure 4.12, in which a loudspeaker was moved across the auditory array. The activation peak tracks the moving source.

These demonstrations repeat what we learned in Chapter 2 about fields, now on an embodied system with real sensors placed in a real environment. So far, however, the vehicle is not moving, the loop through the environment is still open. The last outstanding issue then is how to drive a behavioral dynamics from neural dynamics. The idea, of course, is that the behavioral dynamics should control the actual direction in which the vehicle is headed so that the vehicle turns toward any sound source represented by the field. This would generate taxis behavior for sound sources, now implemented

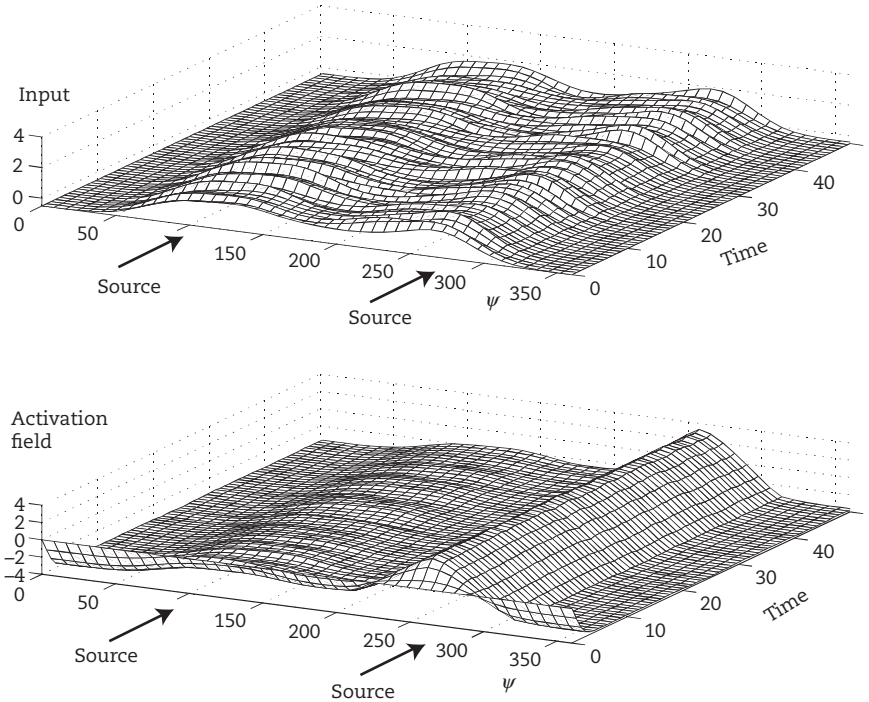

FIGURE 4.10: Shown is the same schema as in Figure 4.9, but with two loudspeakers at the marked bearings. The field on the bottom selects the rightmost source and inhibits activation everywhere else.

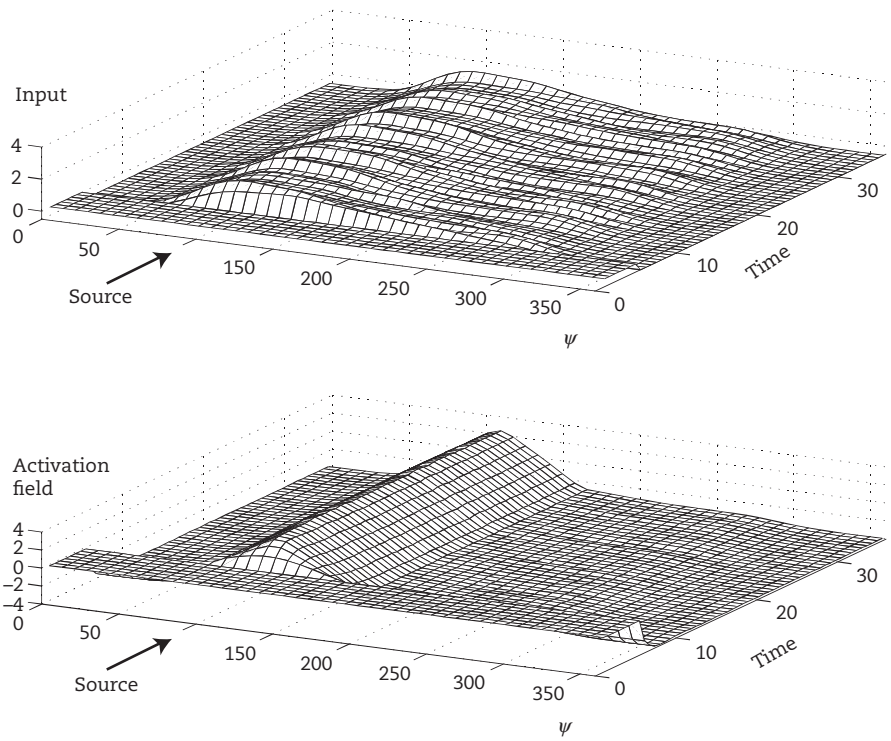

FIGURE 4.11: Shown is the same schema as in Figure 4.9, but with a reflecting surface placed to the right of the loudspeaker, which leads to an input profile with a broad tail on the right. The field below centers its peak on the local maximum of input, effectively suppressing the tail in a form of robust estimation.

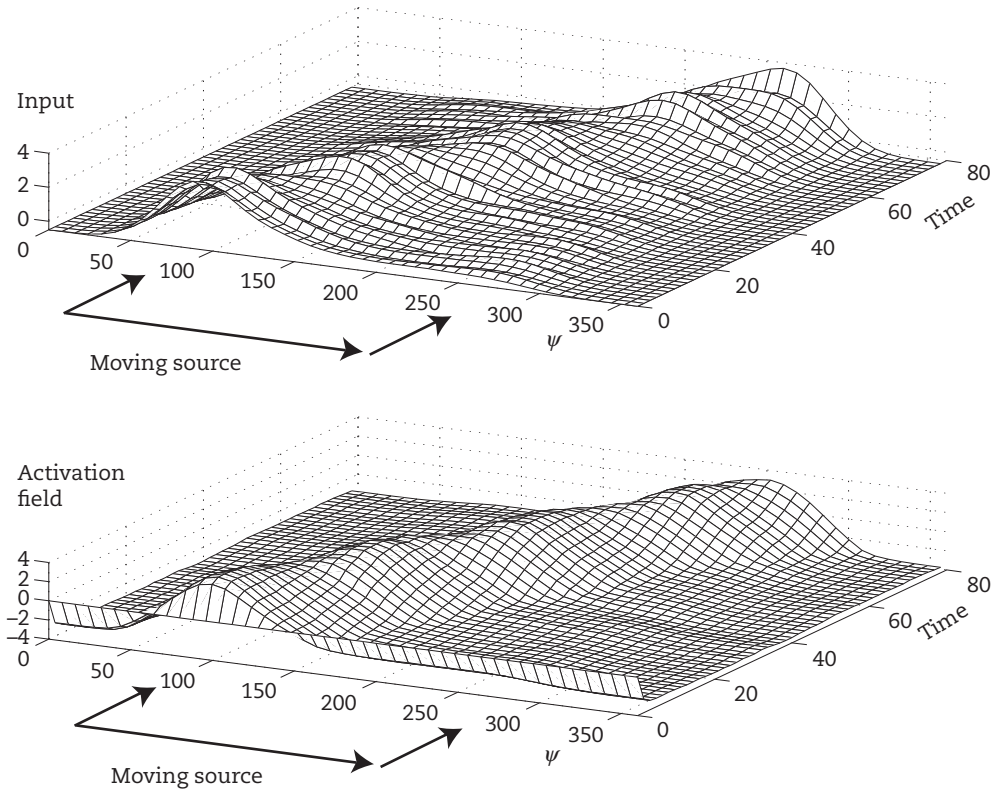

FIGURE 4.12: Shown is the same schema as in Figure 4.9, but the sound source is moved from the leftmost to the rightmost bearing at a constant rate. The peak in the activation field tracks the moving local maximum of the input profile.

with an intermediate neural representation of the bearing of the sound source. To achieve that, we need the peak to induce an attractor in the behavioral dynamics of heading direction that lies in the direction where the peak is located. The attractor would then make the robot turn toward the sound source represented by the activation peak.

How can we make this transformation from an activation peak to an attractor for a behavioral variable? One might be tempted to think of this transformation as a problem of information processing in which we would first compute the peak's location and provide that information to the behavioral dynamics. The peak's location is something like its "center of mass" if we consider above-threshold activation as mass. This idea is formalized by treating the supra-threshold activation, $g(u(\psi))$ (g is the sigmoid function), as a probability density. The theoretical mean of that probability density is an estimate of the peak location:

$$\psi_{\text{peak}} = \frac{\int \psi g\big(u(\psi)\big) d\psi}{\int g\big(u(\psi)\big) d\psi}. \qquad (4.3)$$

Note that this probabilistic interpretation requires a normalization of the supra-threshold activation by dividing by the total supra-threshold activation. Without such normalization, a less activated peak would lead to a smaller estimate, biasing ψ_{peak} to the left, a more activated peak to a larger estimate, biasing ψ_{peak} to the right, even if the peak location was the same. But this normalization also causes problems. What if there is no peak? That will be the case whenever input is not strong enough to drive the field through detection instability. In that case, the probabilistic interpretation leads to a division by zero, which is not well defined and computationally unstable.

This information-processing view is not useful, nor is it necessary. We do not really need to explicitly compute the peak position. What we need is a behavioral dynamics with an attractor at the right location. When there is no supra-threshold peak in the activation field, the field's contribution to the behavioral dynamics should be zero across all heading directions: The activation field that represents sound sources should not impact the heading

direction of the vehicle at all as long as the field has not yet detected a sound source. There may be other contributions to heading direction from other sensors, for example, to avoid obstacles, and these will then dominate (see Box 4.1).

So let's think strictly dynamically. The activation field must generate a dynamics with an attractor at the location of a peak when such a peak is present, but must make no contribution to the rate of change of heading direction when no peak is present (Figure 4.14). The simple idea is to make the (negative) slope of the contribution that ~~is attracted to~~ the peak location proportional to the strength of the supra-threshold peak. The math goes as follows:

$$\dot{\phi} = -\left[\int g\left(u(\psi)\right) d\psi\right]\left(\phi - \psi_{\text{peak}}\right) \qquad (4.4)$$

BOX 4.1 OBSTACLE AVOIDANCE

Taxis, orienting to sources of stimulation, is one of the most basic behaviors of organisms (see, for instance, the classical treatise by Loeb, 1918, on "tropisms," another word for taxis behavior). Because animals tend to live near interfaces, on a land surface, at the bottom of the sea, or hidden in foliage, they cannot successfully move to sources without at the same time steering clear of the many obstacles such interfaces present. Obstacle avoidance, however, has not been studied anywhere nearly as well as taxis behavior. In fact, how humans avoid obstacles when walking has only recently been studied quantitatively (see Warren, 2006, for a review).

In robotics, by contrast, obstacle avoidance has been a topic from the very start of autonomous-movement generation because it is difficult to move in any natural environment without actively preventing collisions. Among the robotic approaches to obstacle avoidance, the potential field approach comes closest to the ideas we address in this book (Khatib, 1986). In the potential field approach, the position of an effector is the behavioral variable. An attractor is erected in a dynamical system that generates a movement plan as the time course of the behavioral variable. This attractor pulls the effector toward the target. Obstacles are contributions to the dynamical system that repel the behavioral variable. You can think of the movement as being a downhill journey in a potential landscape where the minimum is the target (the attractor) and the obstacles are hills.

We have argued throughout this book that behavior needs to be generated by stable states so that it is robust in the face of competing demands and fluctuating sensory information. The dynamics of heading direction offers a variation of the potential field approach, in which the system is at all times in or near an attractor. An attractor dynamics approach to obstacle avoidance was proposed by Schöner and Dose (1992; see Schöner, Dose, and Engels, 1995, for a comprehensive review). The idea is that the direction, ψ_{obst}, in which an obstacle is detected, adds a contribution to the dynamics of heading direction that repels from that direction. We call this contribution a "force-let," formalized as

$$\dot{\phi} = \ldots + \left(\phi - \psi_{\text{obst}}\right) \exp\left[-\frac{\left(\phi - \psi_{\text{obst}}\right)^2}{2\Delta^2}\right] \qquad (4.9)$$

and plotted in Figure 4.13. This contribution has a zero-crossing at the direction, ψ_{obst}, in which the obstacle lies, and has a positive slope at that point. That leads to repulsion from that direction: If the vehicle is headed to the right of that direction, the turning rate is positive, leading the vehicle to turn even further to the right. If it heads to the left of the obstacle, its turning rate is negative, making the vehicle turn even further to the left. The contribution has limited angular range (hence the term "force-let," a play on "wave-let"), reflecting the fact that an obstacle can be ignored as soon as the vehicle is heading in a direction far enough away from the bearing of the obstacle. This shows that, in a sense, an obstacle contribution of this kind defines

two "half-attractors": Heading direction is attracted toward the boundaries of the repulsive range. If an attractive force is added, these half-attractors can become real attractors, as illustrated in Figure 4.13. Incidentally, the repulsive force-let constructed here is a formalization of Braitenberg's avoidance vehicle (his vehicle 3b), which has a contralateral neural organization. By switching the mapping of sensors to motors, the dynamics we derived in Figure 4.2 is inverted, leading to repulsion from rather than attraction to a source of intensity.

Even though this approach to obstacle avoidance was initially invented to enable robots to autonomously navigate, it turned out to describe in quantitative detail obstacle avoidance by humans. Fajen and Warren (2003) used a virtual reality cave to have humans walk toward a visible target. At defined points in the trajectory, they presented an obstacle at varied angles to and distances from the path and observed how the human walker modified his or her path. It turned out that an ensemble of such obstacle avoidance paths could be captured with only a small number of adjustable parameters from a model similar to Equation 4.9. (The human data required the introduction of an inertial term, so the dynamics was one derivative higher than we have used so far.)

In Bicho, Mallet, and Schöner's (2000) study, they showed that the attractor dynamics of obstacle avoidance could be directly linked to sensory input. There is no need to recognize an object as an obstacle and erect a repellor at the direction in which this object lies. What is needed for obstacle avoidance is distance sensing, for instance, by equipping the vehicle with infrared light-emitting diodes and matching infrared light-sensitive resistors arranged in a similar fashion as the microphones illustrated in Figure 4.7. The further away a reflecting surface is from the vehicle, the less infrared light is reflected from that surface and the smaller the detected signal. Every distance sensor erects a repellor in the heading direction that it is currently pointing toward, its strength modulated by the amount of light detected. Sensors that receive very little reflected light contribute only a weak repulsive force-let. The angular range of repulsion reflects the angular range of the sensor. This is the form of obstacle avoidance used in the various demonstrations in this chapter.

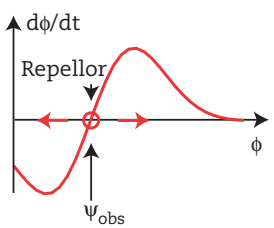

FIGURE 4.13: An obstacle "force-let" is a contribution to the dynamics of heading direction with a zero-crossing at the heading direction, ψ_{obs}, in which an obstacle lies. The positive slope of the force-let at the zero-crossing makes this fixed point a repellor. Heading directions in the vicinity of the fixed point diverge from the repellor, as indicated by the arrows.

where a linear function of ϕ has a zero-crossing at the peak location, ψ_{peak}. The strength of the peak is the integral over its supra-threshold values, which becomes zero if there is no supra-threshold activation. Resolve the equation in parentheses on the right by multiplying by the integral to obtain:

$$\dot{\phi} = -\left[\int g\left(u(\psi)\right) d\psi \phi - \int g\left(u(\psi)\right) d\psi \, \psi_{peak}\right] \quad (4.5)$$

Now insert on the right Equation 4.3 for the theoretical mean, ψ_{peak}. The normalization factor cancels out! This leads to

$$\dot{\phi} = -\left[\int g\left(u(\psi)\right) \phi \, d\psi - \int g\left(u(\psi)\right) \psi \, d\psi\right] \quad (4.6)$$

where we have used the fact that we can move the heading direction, ϕ, under the first integral.

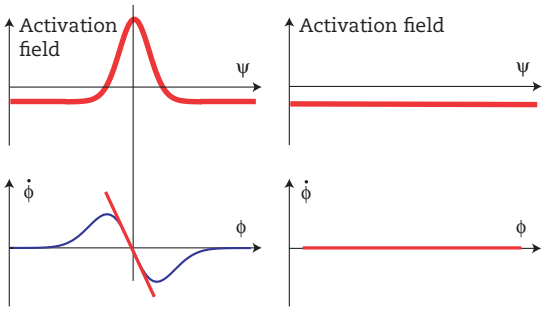

FIGURE 4.14: *Top*: An activation field representing the direction, ψ, in which a sound source lies, is shown with a self-stabilized peak of activation (*left*) and with constant subthreshold activation (*right*). *Bottom left*: Coupling of the activation peak into the behavioral dynamics of heading direction creates an attractor at the peak location (marked by a thin vertical line). The linear dynamics according to Equation 4.7 is shown in red, the range-limited dynamics according to Equation 4.8 is shown in blue. *Bottom right*: The same coupling produces a flat dynamics with rate of change, $\dot{\phi}=0$, when only subthreshold activation is present.

Finally, we pull the common factor in front, now under a single integral, to obtain:

$$\dot{\phi}=-\int g\big(u(\psi)\big)[\phi-\psi]d\psi \qquad (4.7)$$

Because the normalization factor cancels out, there is no longer the problem of division by zero.

Equation 4.7 illustrates consistent dynamical thinking: The activation field ties directly into the dynamics of heading direction. It does so by each field location, ψ, "voting" for a contribution, $-[\phi-\psi]$, to the rate of change of heading direction, which creates an attractor at $\phi=\psi$. The strength of that contribution is proportional to the supra-threshold activation, $g(u(\psi))$, at that field location, ψ. So field sites specify attractors, not computed values.

In practice, variants of Equation 4.7 may be used in which the contributions of each field site to the behavioral dynamics are not necessarily linear. For instance, we may use a range limiting factor, as in

$$\dot{\phi}=-\int g\big(u(\psi)\big)[\phi-\psi]\exp\Big[-(\phi-\psi)^2/2\Delta^2\Big]d\psi \qquad (4.8)$$

with an angular range, Δ, of the attractive "force-let" that each field location specifies. Figure 4.14 illustrates these two forms of coupling, Equations 4.7 and 4.8. As desired, a self-stabilized peak in the activation field induces a dynamics for heading direction that has an attractor (zero-crossing with negative slope) at the heading direction over which the peak is positioned. In the absence of a peak, the dynamics is flat at zero rate of change. Without other contributions to the dynamics of heading direction, all heading directions are fixed points that are marginally stable.

So imagine that the vehicle is driving around, far from sound sources. There is no peak in the field representing sound sources and no contribution from the field to the dynamics of heading direction. Assume we have obstacle avoidance in place, based on a few distance sensors installed on the vehicle. The dynamics of heading direction would then be determined entirely by the contributions of obstacle avoidance, reviewed in Box 4.1. These would push the vehicle's heading direction away from obstacles. A few obstacles scattered throughout the environment would lead the vehicle to turn each time it approaches an obstacle, in effect, exploring the environment. At some point, the vehicle may come sufficiently close to a sound source for the activation field to go through a detection instability. The peak induces a contribution to the dynamics of heading direction that now begins to attract the vehicle to head toward the direction in which the sound source is seen. The peak is self-stabilized, so even if obstacle avoidance forces the vehicle to briefly turn away from the sound source, it will typically turn back once it has cleared an obstacle and may ultimately reach the sound source.

Imagine in this scenario that there were actually two sound sources in the environment. Figures 4.15 and 4.16 illustrate this scenario. Initially, the vehicle may be closer to one sound source whose input then induces the detection instability. This sound source is selected by the neural dynamics. In Figure 4.15, the vehicle is initially closer to

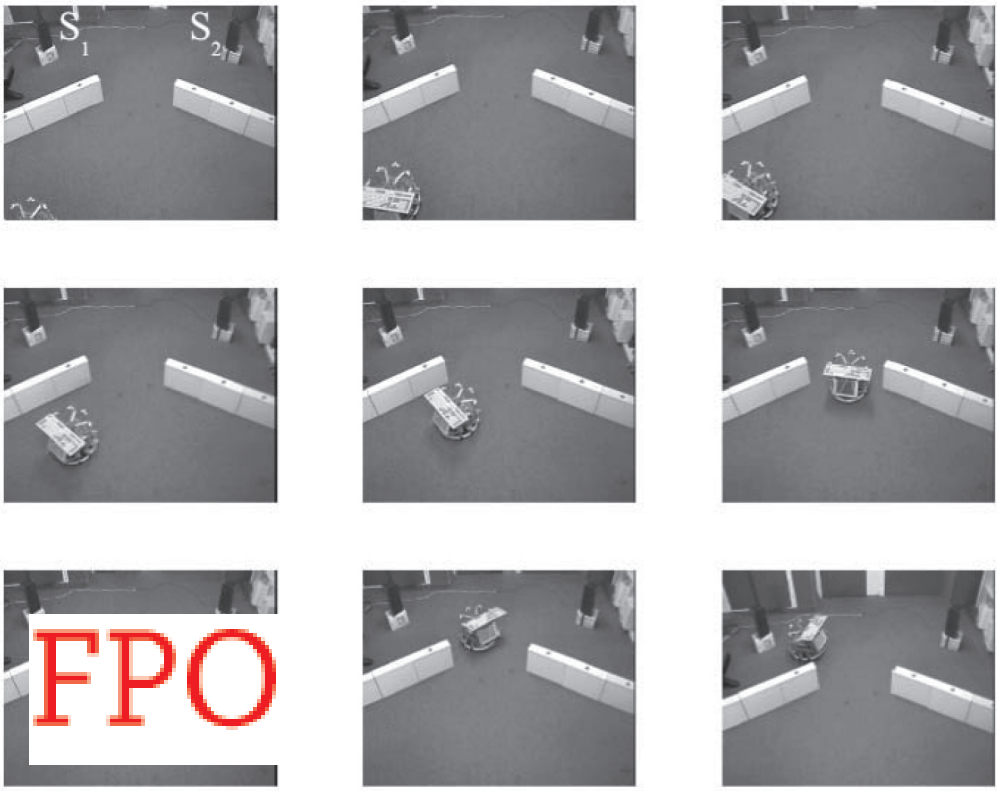

FIGURE 4.15: Series of snapshots from a robotic demonstration, time running from left to right and then from top to bottom. Two loudspeakers are sound sources, marked by S_1 and S_2, top left. Boxes form obstacles in front of the sound source. The robotic vehicle is driven by the dynamics of heading direction described in the text that combines a contribution from a dynamic field representing sound sources and contributions for obstacle avoidance described in Box 4.1. The forward speed of the vehicle is constant. The vehicle approaches the scene from the bottom left and moves to the center under the influence of obstacle avoidance. The fact that it turns to the left once it has passed the obstacles reflects the earlier selection of the leftmost sound source by its neural field.

the leftmost sound source, which wins the selection competition in the dynamic neural field. The arrangement of the obstacles guides the robot toward the central location between the two sound sources where sensed intensities from both sound sources are approximately equal. Because the initial decision is stabilized by the activation field, the robot turns to the left once it has passed the last obstacle. In Figure 4.16, the vehicle starts out on the right, so it selects the rightmost sound source. When it passes through the same central location, it sticks to that decision and turns right once it has moved past the obstacles.

A final note about the dynamics of Equations 4.7 and 4.8: The absolute calibration of the reference frame in which the bearing, ψ, of sound sources is represented and of the estimated heading direction, ϕ, matters because only the difference, $\phi - \psi$, shows up in these equations. That difference lies in the body-centered reference frame of ζ (compare with Equation 4.2)! So really only the body-centered coordinates matter. The reference frames of the bearing angle, ψ, and heading direction, ϕ, are needed only make the dynamics invariant under rotations of the vehicle on the spot. We do not need to calibrate the reference frames of ψ and ϕ as long as we make the same errors in both of them. Equation 4.2 (or its neural implementation in Chapter 7) ensures that any miscalibration of ϕ is copied over to ψ, so this ensures that we make the same errors in both representations. These errors cancel out in Equations 4.7 or 4.8 and thus do not matter.

Although we have used an autonomous robotic vehicle to lay out the ideas, the principles of how neural representations in activation fields can be linked to behavioral dynamics match what is

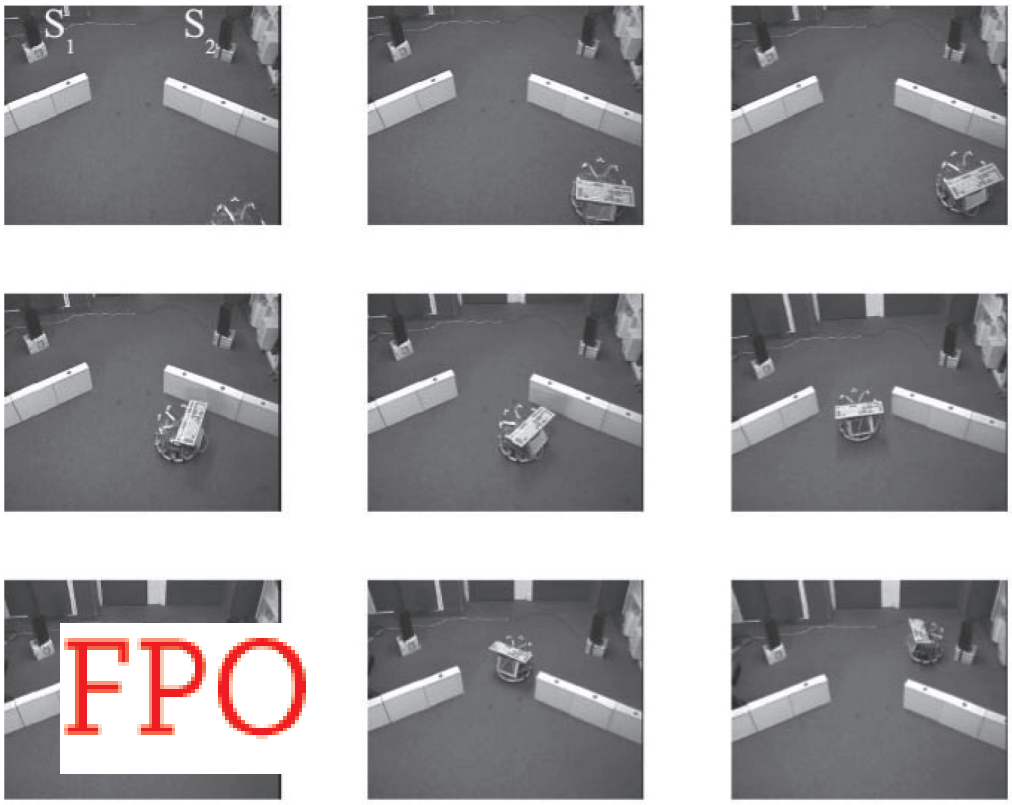

FIGURE 4.16: Same series as in Figure 4.15, but now the vehicle starts out at bottom right. It turns right once it has passed the obstacles, revealing an earlier selection decision of the rightmost sound source.

known about the organization of movement in organisms. In Chapter 3 we reviewed how neural fields capture the way populations of neurons in cortex and subcortical structures like the superior colliculus represent motor parameters. The generation of limb movements based on such motor plans is much more complex than generating movement in the simple vehicle model used here as a metaphor. Even so, limb movement is ultimately brought about by setting attractors for a behavioral dynamics. A brief outline of how that happens is as follows. Limb movements are driven by muscles. The biophysics of muscles, together with the local neural circuitry, including spinal reflex arcs, makes muscles tunable, damped springs (Feldman, 1986). Figure 4.17 provides a simplified illustration of that notion. We have lumped all muscles acting on one particular joint together and described them by a single invariant characteristic that predicts the amount of torque generated by these muscles as a function of the joint angle. Given an external level of torque (e.g., the torque that the weight of the limb creates at that joint), the joint angle will converge to the equilibrium point where the muscles produce the torque that exactly compensates for the external torque. If the joint angle falls short of that equilibrium point, the spinal reflex loops activate extensor muscles and deactivate flexor muscles, decreasing torque generation until the external torque is matched. If the joint extends beyond the equilibrium point, the reflex loops will activate flexor muscles and deactivate extensor muscles, increasing torque generation, again until the external torque is matched. Roughly speaking, the motor periphery acts therefore like an attractor dynamics in which the invariant characteristic together with the external torque sets the attractor state (this requires taking viscosity into account as well, but we will disregard that here for simplicity). Descending input to the motor periphery sets the invariant characteristic to achieve a particular equilibrium point. Movement amounts to shifting the equilibrium point by the descending command, a process that conceptually is analogous

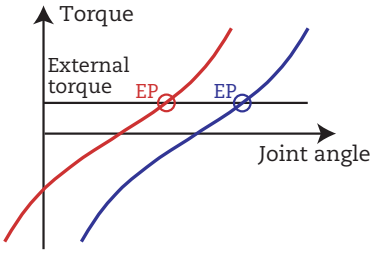

FIGURE 4.17: The invariant characteristic of a muscle-joint system describes the active torque generated by the muscles that converge to a single joint as a function of the joint angle. The invariant characteristic captures the active and passive elastic properties of agonist and antagonist muscles reflected in the monotonic dependence of torque on flexion and extension. The intersection of the invariant characteristic with an external torque (marked by the horizontal line) defines the equilibrium point (EP) to which the joint-muscle system will relax. Movement is induced by shifting the invariant characteristic (e.g., from the instance plotted in red to that plotted in blue). After a shift of the characteristic, the torque induced at the joint through the new invariant characteristic drives the joint to its new attractor posture.

to how peaks in activation fields set attractors for behavioral dynamics. In reality, human voluntary movement generation is much more complex, encompassing issues of movement initiation and termination, movement timing, and inverse kinematics (see, e.g., Martin, Scholz, & Schöner, 2009, for a discussion), but the basic dynamical principles are analogous.

EMBODIED A NOT B

In Chapter 2, we used Piaget's A-not-B paradigm to illustrate the basic instabilities in DFT. Now we will refer back to that paradigm to demonstrate how the DFT account can be embodied—that is, how the neural dynamics in that model can be linked to real sensors and to real motor systems to control a body acting in the world. The robotic demonstration of the DFT account of perseveration uses a video camera as sensory surface. Its visual system filters out those parts of the image that match a particular color (here, an interval of hue values around yellow). The result is a salience image in which only pixels that match the target color have values larger than zero (Figure 4.18). The salience image is summed along the vertical dimension and convolved with an angular kernel for smoothing in a highly simplified account for early visual processing. The angular distribution of salience that results from this operation provides input to an activation field defined again over the heading direction of the vehicle. This entails the same coordinate transform to an allocentric frame detailed previously. The motor system is organized exactly as described for the phonotaxis vehicle. So the A-not-B robot is a taxis vehicle that seeks "yellowness" sources, in a manner of speaking.

The only new ingredient in this model compared to the phonotaxis robot is a memory trace that was

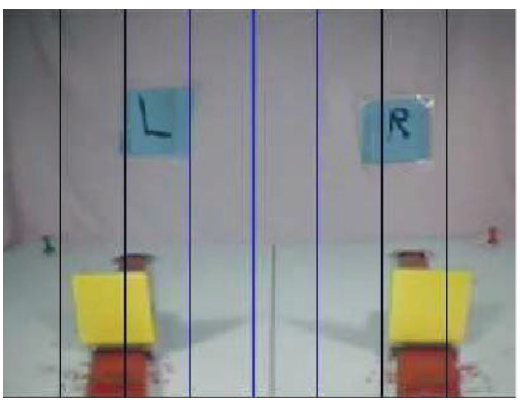

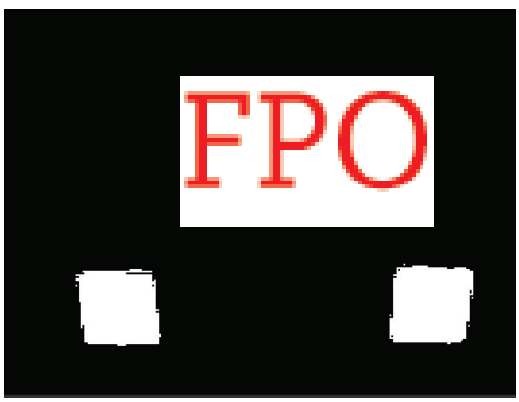

FIGURE 4.18: *Left*: View through the A-not-B robot vehicle's camera of the A-not-B experimental scenario. The yellow cue cards can be moved on two red tracks to move them closer to the robot as needed. *Right*: Salience input into the field is obtained by applying a color filter to the camera image that lets through only those pixels with high enough saturation in a hue interval around yellow. The number of pixels in every vertical bin of the salience image provides input to the activation field that represents the bearning of the visual targets. The count of salient pixels of each column is multiplied by a Gaussian function of heading direction. During the final boost phase of the A-not-B paradigm, a broader Gaussian function is applied.

introduced in Chapter 2 to account for the influence of the motor habit. The memory trace takes the same form as in Chapter 2: Supra-threshold activity in the activation field drives up the memory trace at matching field locations, whereas the memory trace at all other locations decays. When there is no peak in the activation field, however, the memory trace remains unchanged across the entire field, so memory does not decay spontaneously.

We now put the vehicle into an experimental setting, illustrated in Figure 4.19, that mimics the A-not-B paradigm used with infants. Rather than reach toward locations, the robot vehicle turns to orient toward one of two yellow cue cards located at the A and B locations. When a cue card is closer or larger, it subtends a larger angle on the visual array. This generates a stronger salience input (more vertical pixels to sum and more horizontal pixels that overlap). Thus, a cue to the A location is delivered on A trials by presenting a larger cue card closer to the vehicle at the A location for a time interval, followed by a delay. At the end of the delay, a "go" signal is given to the vehicle by moving both cue cards at both locations closer to the vehicle. This input was broadened additionally by applying a broad spatial filter to the salient input at this point. This models the box with the two reaching locations that is being pushed closer to the baby at the end of the delay. After the "go" signal, the robot orients to the selected cue by rotating on the spot—it "reaches" toward A. This is caused by a peak induced by the "go" signal which creates an attractor for heading direction and initiates turning (we will examine this later in Figure 4.21). At the end of a trial, the robot is turned back to its starting orientation. After a small number of A trials (4 or 6, depending on which experiment we model), the same sequence of events occurs but with the cue card presented at the B location.

Figure 4.18 shows what the visual array looks like to the robot vehicle. The two yellow cues cards are picked up by the salience filter, which provides input to the two heading directions in which the yellow cards are seen. Note that both cards are always visible, which models the babies seeing lids at both the A and B locations throughout the experiment. In Chapter 2 this input was called "task input," but it simply comes from the visual array and is as sensory input not different from the "specific input" and "boost" referred to in Chapter 2. All three inputs arise from the visual salience system and are separate only in the sense that the environment is manipulated in the manner of the A-not-B paradigm outlined in Figure 4.19.

Figure 4.20 compares the evolution over time of the activation field and of its memory trace during the first A trial to their time evolution on the first B trial. Initially, the activation field has small bumps

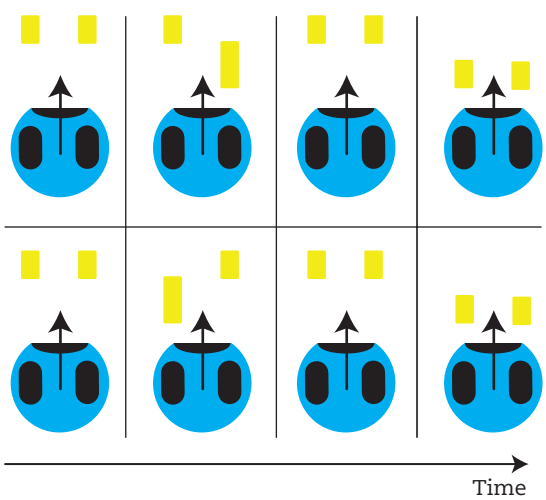

FIGURE 4.19: A-not-B paradigm as implemented for a robot vehicle that embodies the A-not-B dynamic field model. The vehicle, in light blue, is equipped with a camera that faces a scene in which yellow cue cards mark targets. In an A trial (*top row*, time increases from left to right), the cue cards are initially equidistant and of the same size. Then a cue at location A is given by replacing the cue card at that location with a larger copy closer to the robot. During the delay, the cue cards were placed back in their initial position. After the delay, both cue cards are moved closer to the vehicle. The *bottom row* shows the same time steps on a B trial. The only difference is that in the second step, the larger and closer cue card is positioned in the B direction.

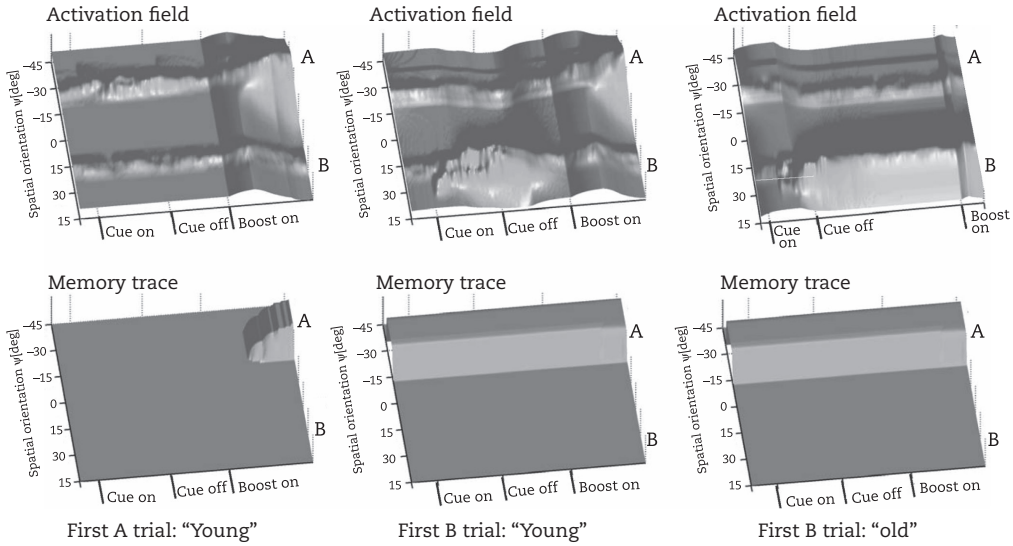

FIGURE 4.20: The activation field (*top*) and the associated memory trace (*bottom*) are shown as functions of time on different trials in the A-not-B paradigm. *Left*: The first A trial for a young robot unable to sustain peaks of activation. *Middle*: The first B trial for a young robot. *Right*: The first B trial for a older robot capable of sustaining peaks.

at both locations that are induced by the salience signals coming from the cue cards in their baseline positions. The specific cue to A generates just a small boost at the A location, which then decays during the delay period. When both cue cards are moved toward the vehicle at the end of the delay period, the broad boost drives the field through the detection instability. A peak forms at the A location, which is still slightly favored from the earlier input. Figure 4.21 shows the activation field at this point together with the dynamics of heading direction. A peak has formed at the A location, which is stable, even though the B location also receives sizeable input at this point. The peak has induced an attractor at the heading that matches the bearing of the A location. The vehicle is still oriented toward the center of the two

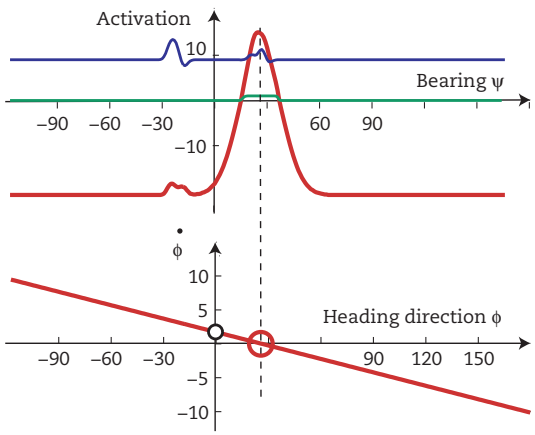

FIGURE 4.21: *Top*: The activation field (top, red solid line) is shown after the end of the delay in the A-not-B paradigm. Input to the field from the visual salience system (blue solid line) has two local maxima over the locations of the two cue cards. At this point, the field has selected the rightmost cue card, even though the current visual input there is weaker. A memory trace has already accumulated and provides input at the selected location (green solid line). *Bottom*: The dynamics of heading direction is driven by the activation peak, which has created an attractor at the heading direction specified by the peak location (dashed vertical line and red circle). The vehicle is still facing toward the center between the two targets, its heading is marked by the black open circle. The vehicle will now start to turn toward the attractor.

locations (zero in the coordinate frame used) but will turn now driven by the attractor. Note in Figure 4.20 that as the peak forms, the memory trace at A builds up. Figure 4.21 registers that memory trace at the beginning of the turning action.

By the time of the first B trial, a sizeable memory trace has built up at the A location, as seen in the middle portion of Figure 4.20. Even though relatively clear input at B drives activation around B close to threshold, this activation decays enough during the delay that the memory trace at A dominates, leading to a peak forming at A and the model making the A-not-B error. The first B trial of an "older" model is shown on the right in Figure 4.20. Here, the field is in the regime in which peaks of activation can be sustained without localized input. The peak induced near B by the cue card is sustained through the delay, so the model responds at B after the delay. The "older" vehicle does not perseverate.

We have used this embodied model of the A-not-B phenomenon to quantitatively account for data from a meta-analysis of the behavior of 400 babies (which is not yet published as this book goes to press). Although the sensory and motor details are different, we were able to reproduce the basic signatures of perseverative behavior as well as a wealth of different conditional probabilities that measure how the history of reaches determines future reaches.

Now that we have an embodied variant of the A-not-B model, we may take the system out of the restricted experimental paradigm to ask more generally what the functional significance of perseveration is in object-oriented action. Why would infants make the "stupid" error? Why give habit so much weight that it may overturn perceptually cued action plans? To address this, we put the robotic vehicle into an arena in which there was a visual target with the yellow color that matched the vehicle's salience filter (Figure 4.22). We added obstacles, which the vehicle was able to avoid with the technique reviewed in Box 4.1. The obstacles merely served to force the robot to turn and thus lose the target from view. So is "out of sight" truly "out of mind" for the young robot? How does the memory trace help? And what does the older robot gain from its capacity to build a working memory of the cued action plan?

Figure 4.22 compares on the left a young and an old robot as they head to a visible target. An obstacle early in the path forces the robot to turn and lose the target from sight. For the young robot, the peak in its target field decays when it is no longer supported by salience input; the robot "forgets" the target, now truly "out of mind," and the robot continues on a straight path past the obstacle, no longer trying to turn back to the original target. The older robot, by contrast, sustains the peak at the direction to the target while it avoids the obstacle. This enables

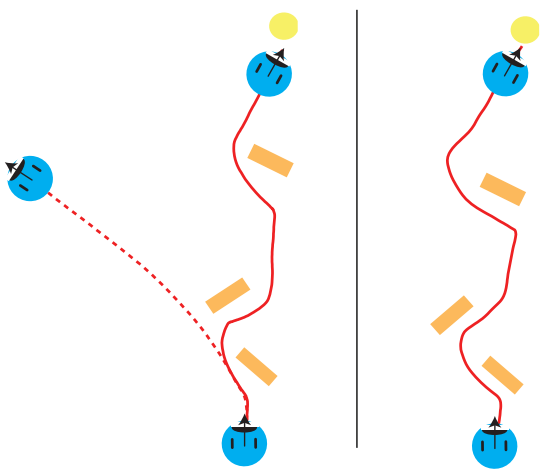

FIGURE 4.22: A-not-B vehicle moving in the presence of a visual target that its salience system is sensitive to (yellow circle on top). Obstacles (brown rectangles) are placed in the scene. The vehicle is initially oriented toward the target (blue circles at the bottom), so its target field builds a peak of activation at the the target's bearing. The *left* panel contrasts the path generated by the "young" robot without a memory trace (dashed red line) with the path generated by the "old" robot (red solid line). The *right* panel shows the path generated by a young robot vehicle with a memory trace. The data were recorded from real robots performing the task. The obstacles were low enough to not occlude the target for the camera mounted high on the vehicle as long as the vehicle was pointing in the direction of the target.

the older robot to turn back toward the target after it has passed the obstacle. When the target thus comes back into view, the peak is updated by current salience input and steers the robot to the target. So clearly, the capacity to sustain peaks through periods when sensory information about the target is no longer available adds stability and enables the older robot to reach the goals under a broader set of environmental conditions than that of the younger robot.

In this demonstration, the memory trace was not active in the two robots. The right panel of Figure 4.22 shows a run of the younger robot, but now with a memory trace in place. Surprisingly, this makes the young robot look like an old robot: It does not lose the target from its mind as it loses it from sight! What happens is that the robot has quickly built a memory trace when it first builds the peak at the initial bearing of the target. This memory trace stabilizes the peak and slows its decay when sensory input from the target is lost. The peak is still there to make the robot turn back toward the source once the first obstacle has been passed. Renewed sensory input then keeps the robot on course. So sustained activation is not the only mechanism for keeping in mind what is out of sight. The memory trace also serves to stabilize movement plans. In light of the sensory-motor challenges to goal-directed movement for young infants, such stabilization is supportive of goal-oriented action. The limitation of the memory trace as a stabilization mechanism is that it is less flexible than sustained activation. A new, sustained peak can be set by sufficiently strong sensory input, and this setting of a new target may overwrite the previous target. This is what happens in the A-not-B paradigm on the switch from the A to the B trials. The memory trace, by contrast, cannot be switched as rapidly by sensory information. A new memory trace only forms as the system "experiences" a new neural activation pattern. So neither the younger nor the older robot loses from mind what is out of sight. The older robot is more flexible in what is on its mind and the A-not-B paradigm is sensitive to that flexibility.

Why would a younger infant or robot rely on the slower memory trace to stabilize decisions rather than on the faster mechanism of sustained activation? The reason is not really known. One possibility is that the kinds of coordinate transforms postulated here are harder to achieve for younger systems. A fast switch may induce categorical errors when coordinate frames become misaligned during the occlusion of a target. The slower memory trace may filter out what has been reliably tracked in spite of a system's difficulty of stabilizing and aligning reference frames. In Part 2 of this book we will provide the theoretical tools to address processes of transforming and aligning reference frames. In Part 3 we will introduce concepts and models that begin to address some of the developmental issues implied in this interpretation.

CONCLUSION

We have seen that the closed loop through which overt motor action controls sensory input to an organism generates a second type of dynamics. This behavioral dynamics can be captured through variables that characterize the state of the physical, embodied system relative to its environment. Attractors are critical to bringing about consistent behavior in the face of fluctuations and distractors and stabilizing simple sensory-motor decisions. Behavioral dynamics is inherently limited in flexibility, however, requiring the continuous availability of sensory inputs. We saw how the neural representations provided by DFT enhance the flexibility of behavior. In fact, returning to the phenomenon of infant perseverative behavior first used in Chapter 2 to illustrate the core concepts of DFT, we were able to get a concrete sense of the developmental trajectory of increasing flexibility as the system goes from more strongly input-driven to interaction-dominated dynamics. This topic will be a central theme of Part 3.

Ultimately, overt behavior always entails both neural and behavioral dynamics. We examined in detail how neural dynamics ties into behavioral dynamics, recognizing that there was no need to "read out" the estimates or decisions generated by activation fields. Instead, peaks of activation directly create attractors of behavioral dynamics. This was a beautiful instance of pervasive dynamical thinking that is useful to keep in mind as we move forward to increasingly abstract, cognitive levels of processing. Having laid the conceptual foundations for both the dynamics of behavior and the dynamics of elementary forms of cognition, we will address in Parts 2 and 3 more ambitious forms of cognitive processing. In Part 2, we will discover new cognitive functions that derive from multidimensional dynamic fields, including a neural-process account for how the coordinate transforms can be achieved that were assumed in the present chapter. This will enable us to provide a neural dynamic foundation

of visual cognition. In Part 3 we will address learning and development and extend our investigation into cognition by looking at cognitive control and sequence generation.

REFERENCES

Bicho, E., Mallet, P., & Schöner, G. (2000). Target representation on an autonomous vehicle with low-level sensors. *The International Journal of Robotics Research, 19*, 424–447.

Braitenberg, V. (1984). *Vehicles. Experiments in synthetic psychology*. Cambridge, MA: MIT Press.

Fajen, B. R., & Warren, W. H. (2003). Behavioral dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance, 29*, 262–343.

Feldman, A. G. (1986). Once more on the equilibrium point hypothesis (λ-model) for motor control. *Journal of Motor Behavior, 18*, 15–54.

Khatib, O. (1986). Real-time obstacle avoidance for manipulators and mobile robots. *International Journal Robotics Research, 5*, 90–98.

Loeb, J. (1918). *Forced movements, tropisms, and animal conduct*. Philadelphia: Lippincott.

Martin, V., Scholz, J. P., & Schöner, G. (2009). Redundancy, self-motion and motor control. *Neural Computation, 21*, 1371–1414.

Reichardt, W., & Poggio, T. (1976). Visual control of orientation behaviour in the fly: I. A quantitative analysis. *Quarterly Reviews in Biophysics, 9*, 311–375.

Schneegans, S., & Schöner, G. (2012). A neural mechanism for coordinate transformation predicts pre-saccadic remapping. *Biological Cybernetics, 106*, 89–109.

Schöner, G., & Dose, M. (1992). A dynamical systems approach to task-level system integration used to plan and control autonomous vehicle motion. *Robotics and Autonomous Systems, 10*, 253–267.

Schöner, G., Dose, M., & Engels, C. (1995). Dynamics of behavior: Theory and applications for autonomous robot architectures. *Robotics and Autonomous Systems, 16*, 213–245.

Warren, W. H. (2006). The dynamics of perception and action. *Psychological Review, 113*, 358–389.

EXERCISES FOR CHAPTER 4

The simulator for this exercise is provided in the file `launcherRobotSimulator`. Running this file will open a main graphical user interface (GUI) window showing a dynamic neural field with a coupled attractor dynamics and control elements, and an additional window showing a top-down view of a simulated robot in a small arena. The robot is depicted as a gray circle with an arrow indicating its heading direction (think of a simple differential-drive robot like the Khepera here). It has nine directional sensors (such as light sensors) placed equidistant along its front half. The noisy outputs of these sensors are shown in the bar plot in the top part of the window. The sensors respond to targets in the arena, with intensity depending on distance to the target. These targets are shown as smaller red circles in the arena plot. You can add or remove targets at any time by clicking on the corresponding button at the bottom of the window and then clicking on a location in the arena.

In the main GUI window, the top plot shows the activation of a one-dimensional field receiving inputs from the robot's sensors. The field is defined over the space of robot orientations in an allocentric reference frame (fixed in the world, not rotating with the robot). Note that the x-axis is flipped to allow a more intuitive mapping to the sensor geometry on the robot while retaining the mathematical conventions for specifying orientations. The field provides input to the attractor dynamics shown in the bottom plot, in the form introduced in this chapter. The red plot gives the turning rate (rate of change in heading direction) for every possible heading direction; the red circle on this plot indicates the actual heading direction and instantaneous turning rate of the simulated robot. You can control field parameters and strength of coupling between field and attractor dynamics via the sliders at the bottom (hover over the slider to get a description of the controlled parameter). In addition, you can control the forward speed of the robot via the slider on the bottom right. Clicking the `Reset` button will reset the field activation and also put the robot back in its initial position.

The goal of this exercise is to explore the role of the detection and selection instabilities for the orientation behavior of the robot.

Exercise 1: Detection Instability

When you start up the interactive simulator the robot environment is created with a single target in the upper left. As it is quite far from the robot, the target affects the activation field only weakly. This is the perfect setup to study the detection instability. You can start with the preset parameter values. Use the slider v_r to set the forward speed of the robot to positive values. The robot will drive until it has reached the target and then automatically sets its speed back to zero. You can re-place the robot at its initial position by clicking the "reset robot" button. At some point during the robot's movement, a detection instability will occur. You can pause the

simulator when this happens and also reproduce this event several times by using the reset button.

1. How does the detection instability manifest itself? Input fluctuates due to noise modeled for the sensors. Does the peak fluctuate with input after the detection instability?
2. What happens to the behavioral dynamics at the detection instability? How does this affect the motor behavior of the robot?
3. After the detection instability, the peak tracks sensory input from the target. What does this do to the behavioral dynamics?
4. If you run through this path at a higher speed, the robot turns later in the path. Why?

Exercise 2: Selection Instability

For selection instability, add a second target and place it to the upper right of the robot. Try to place it at a distance equal to that of the other target. Again, activate the forward velocity and let the robot run.

1. Observe the input profile (green curve) and watch how detection instability occurs. What happens to the alternate peak when one target is selected?
2. Reset the robot and repeat the trial. Can you observe different selection outcomes?
3. By removing and then again adding a target, you can vary its location. Can you manipulate the probability that the target will be selected over the default target?
4. Set up a situation with symmetric targets in which either target can be selected. Turn off the neural interaction by setting all three parameters of the interaction kernel to zero. What happens now when the robot heads for the targets?
5. Go back to the initial setting by quitting and restarting the simulator. Add several targets near the initial target. What does that do to the detection instability?
6. What happens in this case as you approach the target? Do you see a transition from monomodal input to the field to multimodal input? What happens to the self-stabilized peak itself?
7. You can play with the h-level to enable sustained peaks or not. Try removing a target right as the robot is heading toward a target, perhaps in the presence of another target. Can you see the effect of sustained activation?

Exercise 3: Avoidance

You can explore a simple form of avoidance by changing the sign of the coupling from the field into the attractor dynamics to negative (use the "parameters" button to get access to the parameter values for the attractor dynamics). Explore avoidance behavior as a form of obstacle avoidance.