Appendix A: Motor Schemas

Motor Schema theory has been developed in order to provide a biologically plausible model of the behavior of various animals in certain complex situations [3]; it has also proven to be an useful and interesting framework for applying Distributed Artificial Intelligence techniques to the robotics domain. It consists of perceptual schemas which map sensed data into object representations which in turn are used by motor schema to generate actions. Motor schema are similar to behaviors in many respects; however, rather than using arbitration to select the output of a single behavior, the outputs of all schema are combined in order to produce an action that reflects multiple considerations. This combination of schema is also prioritized, in that each schema has a weight associated with its output command that reflects the importance of the command or the certainty of the information on which it is based.

Arbib and House developed two different schema models which could explain prey-catching and obstacle-detouring behaviors that were observed in frogs [5]. The Vector Field model of path planning associates motor actions with schema objects and combines them; the Orientation Selection model chooses among orientations based on a summation of excitatory and inhibitory influences from each schema. Arbib has pursued the Vector Field model in his subsequent work [4], and this approach has become the robot control paradigm associated with motor schema theory [6]; however, it is not the only possible instantiation of the theory. Although they were not explicitly developed for the purpose, there exist robot control systems which serve as realizations of the Orientation Selection model. One such system is the Vector Field Histogram [11], another is the Distributed Architecture for Mobile Navigation.

A comparison of the two schema paradigms will be made in the context of explaining observed animal behavior, and instantiations of these paradigms will be compared in the context of controlling a mobile robot. In order to understand motor schemas, let us first study the domain for which they were first developed: the observed prey-catching and obstacle-avoiding behavior of frogs.

A.1 Observed Prey-catching Behavior of Frogs

When frogs are presented with stimuli such as worms, the animal selects one of the prey and moves directly towards it, as shown in Figure A-1a; this can be modeled in a classical stimulus-response fashion with each potential prey triggering an excitatory response in the frog. When obstacles such as fence posts are introduced, they influence the decision of which direction to head in, as shown in Figure A-1b; the frog circumvents the fence and then heads toward the worms. However, a simple stimulus-response mechanism cannot explain the behavior observed in Figure A-1c, where the frog moves around the fence posts only to find itself blocked by a solid wall; the frog continues to orient itself toward the chosen prey even though it is no longer visible.

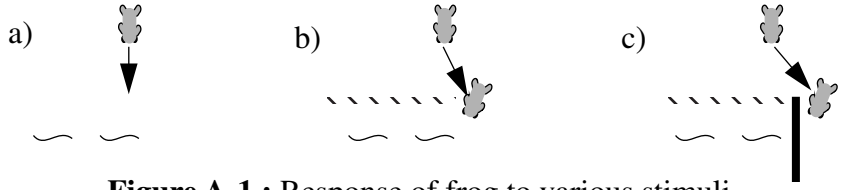

Figure A-1 : Response of frog to various stimuli.
a) prey, b) prey and obstacle, c) occluded prey.

The selection of movement direction as described above is essentially one-dimensional; it depended only on the angles between the frog and its prey and between the frog and the obstacles. However, it is clear from further experimentation that depth information also plays a significant role in the decision-making process. Figure A-2 shows how the behavior of the frog varies under different conditions; the arrows indicate the directions along which the frog moved, along with the percentage of trials for which each direction was chosen. Figure A-2a shows that the frog preferred to go around the fence posts to get to the worm, while Figure A-2b shows that a gap in the fence was taken advantage of when present. In Figure A-2c, the closer fence has a gap in it and the further one does not; the frog chose to go through the gap in the first fence more often than it chose to go around it, yet it did not go through the gap so often as if the second fence were not there. Neither fence in Figure A-2d contained a gap, and so the frog almost always chose to detour around them. In Figure A-2e, a large gap in the closer fence led the frog to often choose to go through rather than detour around, even though the side barriers formed a trap. These observed behaviors indicate that both fences contributed to the decisions made, and that the influence of perceptual information diminishes with the perceived object's distance; the frog is apparently maintaining and using depth information, but it is not processing it in a manner sophisticated enough to account for the more complex topology of the last case. Computational models of the frog's behavior must account for these observed phenomena, as well as the fact that, even under the same initial conditions, the same decision to go through versus detour around the fence is not always made.

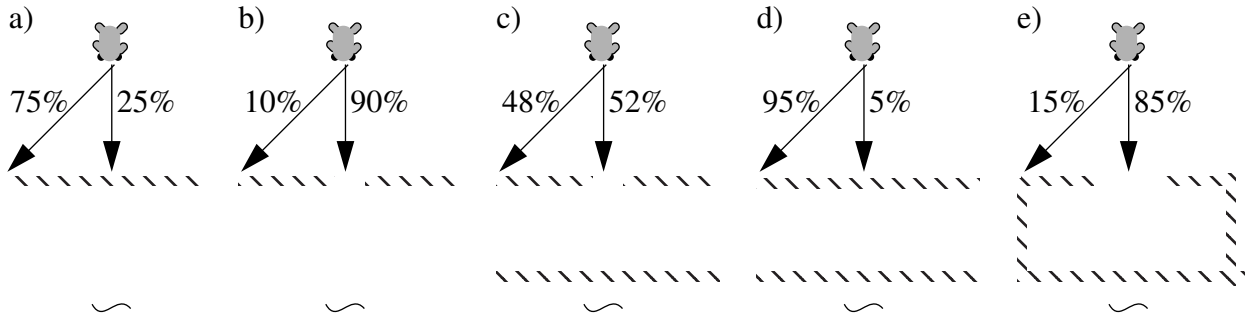

Figure A-2 : Frequency with which frog went through or around a fence.

A.1.1 Schema Models

Schema theory has been used to create models which account for the experimental results described above in a way that is cognitively, neurologically and computationally plausible. Schemas are independent modules which represent perceptual information and contain processes for deciding how to act on that information; together, they constitute a distributed system for reasoning and for acting, either externally to change the state of the physical world or internally to change the state of the system. Each schema has an associated activity level which represents its confidence that the perceptual information it contains is accurate.

In applying schema theory to the domain described above, Arbib and House postulate that two separate depth maps are maintained for prey stimuli and obstacle stimuli, that they are individually processed to generate orientation preferences, and that the motor schemas combine these preferences in order to produce coordinated action [5]. The methodology outlined by schema theory is broad enough that it may be implemented by more than one computational paradigm. Arbib and House postulated both an Orientation Selection model and a Vector Field model. Although only the latter approach was pursued in subsequent work in modelling animal behavior, either model could account for the experimental data on frogs which was presented, and each has their advantages and disadvantages. We present the two models described, and discuss the merits of each.

Orientation Selection Model

In the Orientation Selection model, the schemas operate by convolving input depth maps with masks designed to elicit the appropriate response in a two-dimensional distance/orientation space, as shown in Figure A-3a. The obstacle-avoidance schema convolves the barrier depth map with an inhibitory mask, shown in Figure A-3b, which has the effect of slightly spreading the obstacle laterally (in the orientation dimension); the mask also spreads obstacles backward in distance so that potential goal locations behind the obstacle are also inhibited. The prey-catching schema convolves the prey depth map with an excitatory mask. As can be seen in Figure A-3c, this mask spreads the input stimulus forward along the distance axis so that it may overcome inhibitory influences directly in front of it, and also has a large lateral spread which serves to create secondary goal locations in those cases where the inhibition due to a barrier is not overcome. The resulting arrays output by these two schemas are then summed to create a net excitation array which represents the combination of all excitatory and inhibitory influences. This excitation array is integrated over distance for several orientations to yield an histogram of total excitation as a function of orientation; the orientation corresponding to the highest value in that histogram is selected as the direction in which the frog will move. A depth segmentation module determines the target depth for each orientation, presumably where the excitation reaches its maximum value along that direction; the resulting target depth map can then be indexed by the chosen orientation, so that the frog has now determined the angle in which it will head and the distance it will move in that direction; thus, an action is decided upon by combining the influences of all available inputs.

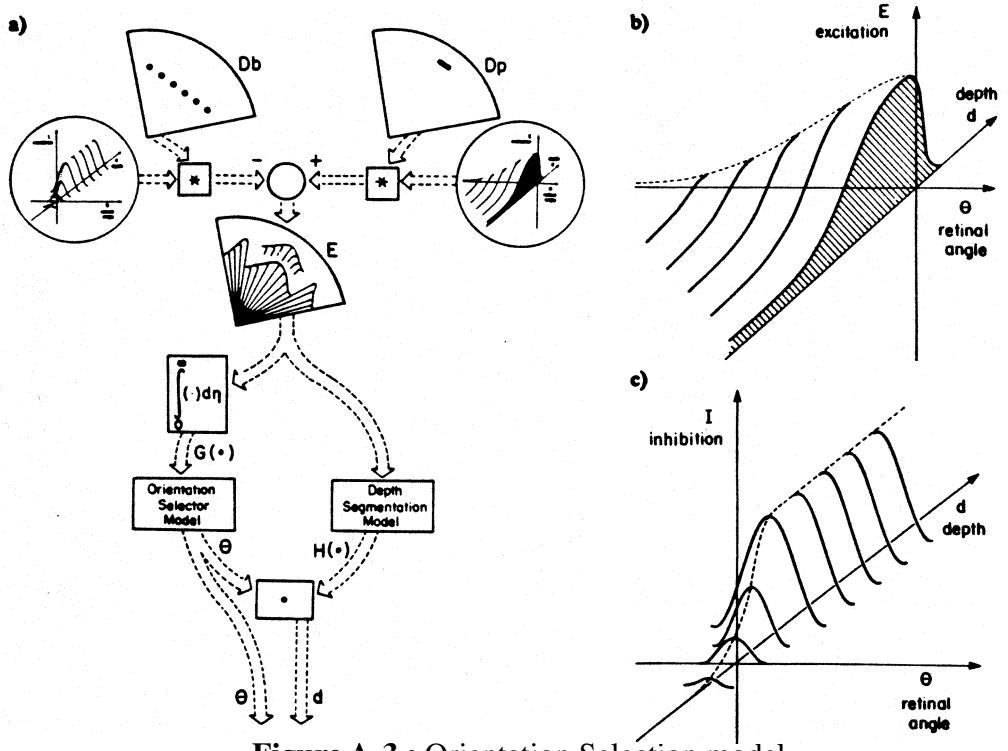

Figure A-3 : Orientation Selection model.
(Reproduced from [5], with permission)

Vector Field Model

In the Orientation Selection model, a level of excitation was determined for each point in a local depth map; the Vector Field model seeks to associate with each point not a scalar value, but rather a vector indicating which direction the animal should head in were it to pass through that point. The magnitude of the vector indicates the strength of the choice for moving in that particular direction, i.e., a small vector at a point means that heading in a different direction would have a relatively small effect, while a large vector indicates that the choice of orientation at that position is an important one. Thus, the prey-catching schema would set up an attractive field of vectors oriented towards it, as shown in Figure A-4a, and the obstacle-avoidance schema would associate with each obstacle a repulsive field like that shown in Figure A-4b. Unlike the prey schema's attractive field, the vector field generated for an obstacle is not symmetric, but instead the vectors point more strongly in the lateral direction relative to the frog's position; this is analogous to the excitation mask used for prey in the Orientation Selection model, in that allows a target to be chosen when it is immediately behind an obstacle. A row of obstacles, such as a series of fence posts, would be combined to create the vector field in Figure A-4c, which would tend to move an object towards either end of the fence. The frog also has a vector field associated with itself which serve to propel the frog forward, as shown in Figure A-4d. For each of these types of fields, the strength of the vectors diminishes with distance from the object creating the field.

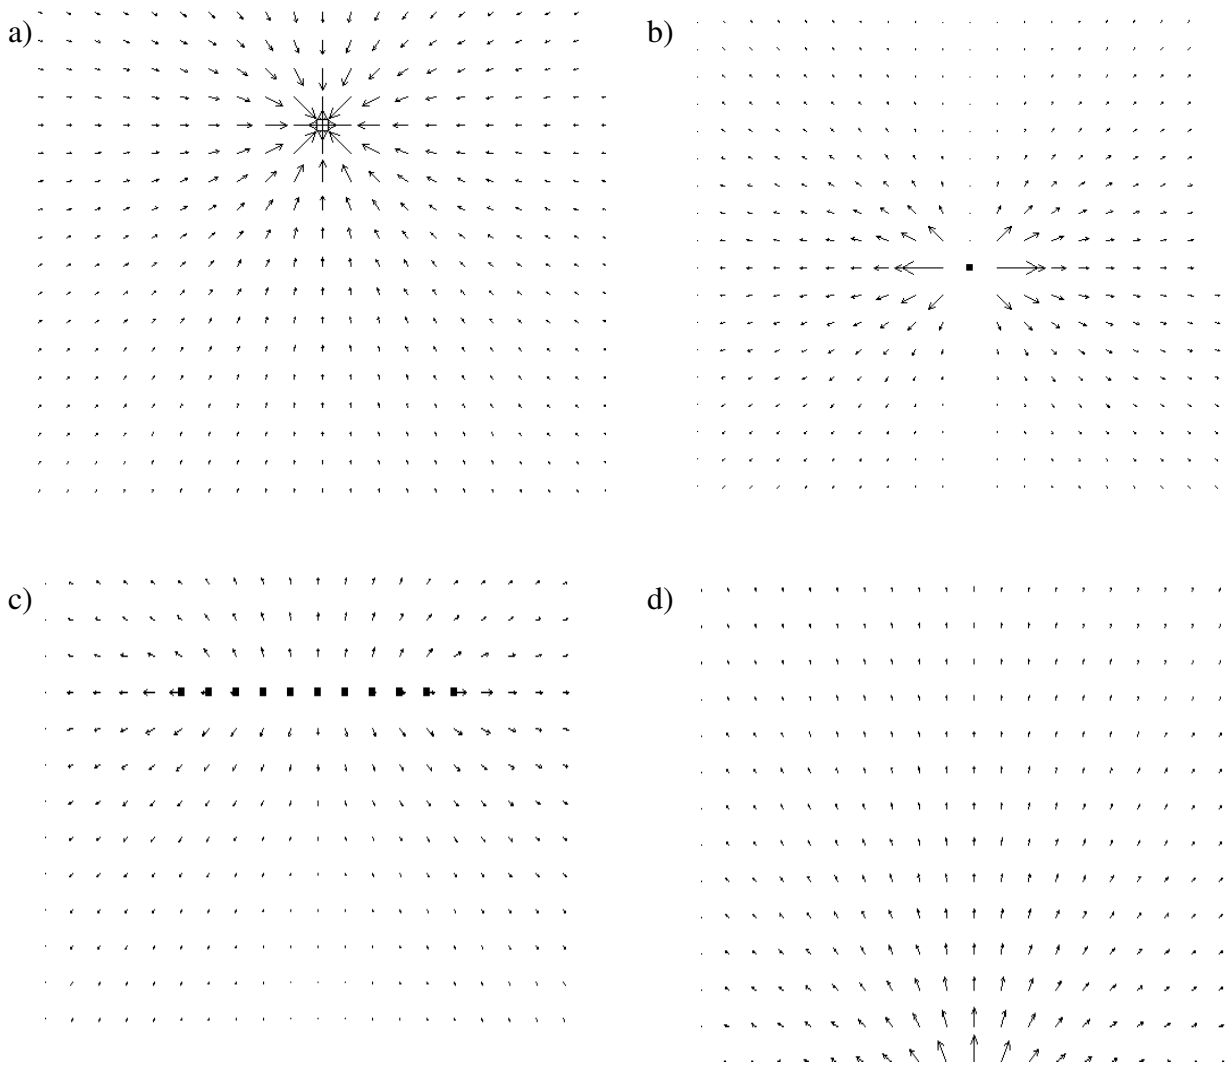

Figure A-4 : Vector fields.
a) attractive toward prey; b) repulsive from obstacle;
c) leading around a series of obstacles; d) self-induced forward propulsion.

The schemas are combined by summing the vector fields. Figure A-5 shows the vector field resulting from the combination of schemas representing the frog facing a prey object behind a fence. Arbib and House hypothesize that a "tracking" creature such as a gerbil would then perform some sort of integration on these vector fields in order to plan a path, and that a field such as that in Figure A-5 would yield two "bundles" of trajectories, one leading around the fence to the left and the other to the right. In the case of a "ballistic" animal like the frog, however, they suggest that the vector field is instead processed to create a map of target locations.

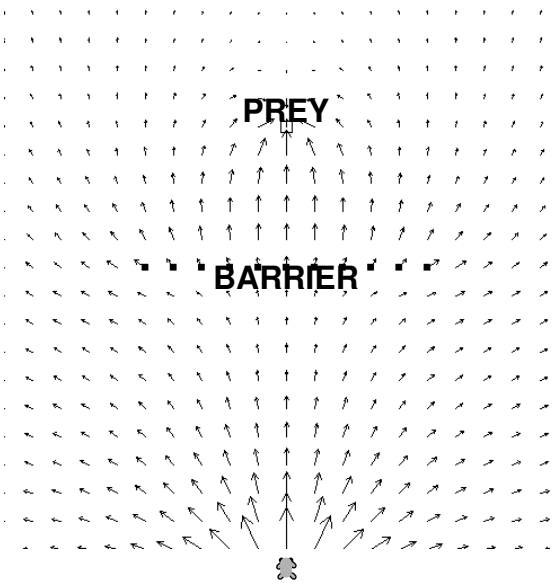

Figure A-5 : Combined vector field for toad, prey, and barrier posts.

Applying the divergence operator to the vector field yields a scalar field; inverted, it affords the animal with a map of convergence such as that shown in Figure A-6. The peaks of this map occur where the vectors in Figure A-5 converge at the two ends of the fence and at the prey position, and they indicate potential goal locations. The frog is presumed to then choose a target from among these peaks, and it may also remember the locations of the other peaks for subsequent decisions; for example, it can decide to sidestep to the left end of the fence and then redirect toward the worm from there.

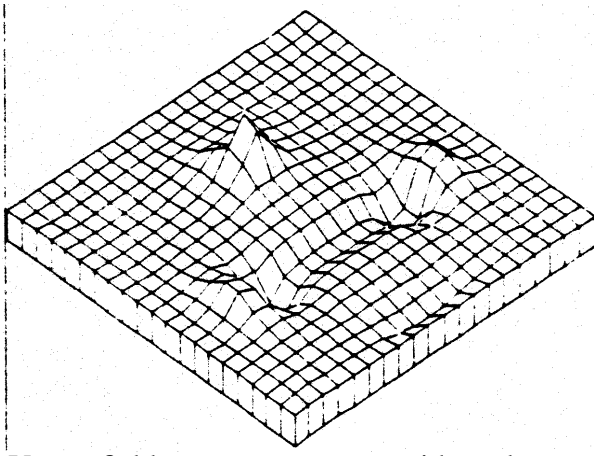

Figure A-6 : Vector field convergence map with peaks at goal locations.
(Reproduced from [5], with permission).

A.1.2 Comparison of Behavior Models

Both the Orientation Selection (OS) model and the Vector Field (VF) model qualitatively reproduce the behavior of frogs observed in the scenarios shown in Figure A-2. The excitation array in the OS model and the convergence map in the VF model both contain peak values only at the prey location and at the ends of the row of obstacles, thus predicting, in agreement with experimental results, that the frog will either move directly through the fence towards the prey or to either end of the fence posts. Figure A-2 also shows the frequency in which the frogs chose to either go forward or detour; the computational models must also be able to explain this apparent non-determinism. Arbib and House were able to reproduce this phenomenon within the OS model by varying the lateral spread of excitation induced by the presence of a prey. A narrow spread diminished the total excitation present at the fence ends, so that the frog would choose to go forward through the fence, while a wider spread used on the same inputs resulted in the frog choosing the detour behavior instead. Variation of this spread parameter within the frog population or even within an individual frog due to motivational factors could explain their probabilistic behavior. No such parameter was identified for the VF model; however, if the schemas were assigned a weight which affected the vector field summation, then variation of these weights might produce the desired result. For example, a higher weight for the prey schema relative to the obstacle schema could result in the frog going directly through the fence, while a lower prey weight could result in the detour behavior.

Arbib and House contend that the VF model is more sophisticated than the OS model and better explains complex behavior, mainly due to the assertion that the OS model creates a "simple decision surface" while the VF model creates a "spatially encoded map of potential motor activity". However, the convergence map in the VF model and the excitation array in the OS model are both two-dimensional surfaces whose peaks correspond to potential targets for the next movement. The OS model uses this map to create a one-dimensional histogram in order to select an orientation, but the VF model must ultimately process the convergence map in a similar manner in order to decide upon a specific target. Additionally, Arbib's evaluation is based on the assumption that an entire path is planned before it is executed, rather than reevaluating the situation as the animal moves along. By continually deciding and choosing a new direction of movement based on the current situation, the OS model can produce arbitrarily complex paths. Arbib and House also claim that, in the OS model, once the frog has selected a target, it cannot remember whether that target corresponds to the location of a prey object or merely an intermediate destination in a detour maneuver; furthermore, if the frog is executing a detour behavior, it will not be able to remember where the original prey target was and therefore cannot orient towards it. However, the preprocessing stages in the orientation selection process generate and maintain essentially the same information as the vector fields; in particular, the input depth maps contain the location of both prey and obstacles, and the output of the depth segmentation module, combined with the orientation histogram, provides goal locations in polar coordinates. Thus, it appears that both the Orientation Selection and the Vector Field models could account for observed animal behavior, and that further experimentation is needed in order to distinguish the two as viable mechanisms. More importantly for our purposes here is the question of how

effectively these schema models may be applied to the robotics domain. Implemented physical systems can provide feedback regarding matters such as the suitability of the representations used, the efficiency of the algorithm, and the applicability to various other situations and environments.

VECTOR FIELD HISTOGRAM

The Vector Field Histogram [11] is a method used for robot control which serves as a realization of the Orientation Selection model, although it was not developed with that intention. The Vector Field Histogram (VFH) method begins by building an egocentric map of observed obstacles. It does so by creating a local histogram grid whose elements represent the number of times the robot's sensors have reported the presence of an obstacle at that location, as illustrated in Figure a. This corresponds to the input obstacle depth map used in the OS model, with the addition of the histogram counts which provide the certainty information which Arbib hypothesized but did not include in his model. As in the OS model, the next stage of processing involves integration along various orientations to yield a one-dimensional polar histogram of obstacle density (Figure b); like the Vector Field model, the significance of a perceived obstacle diminishes with distance. Note that this histogram corresponds only to the inhibitory influence of obstacles in the OS model, rather than a general summation of excitation that includes other factors such as the presence of attractors. The histogram grid is not convolved with a mask as in the OS model, but the same effect of spreading activation is achieved in the VFH method by smoothing the polar histogram so that orientations adjacent to obstacle directions also receive inhibitory influence. Finally, a threshold τ is applied to the polar histogram, and those areas whose obstacle density is below that threshold are considered as candidate directions for movement. The presence of goals such as a prey object now comes into play, as that candidate direction which is closest to the direction of the goal is chosen as the direction of travel.

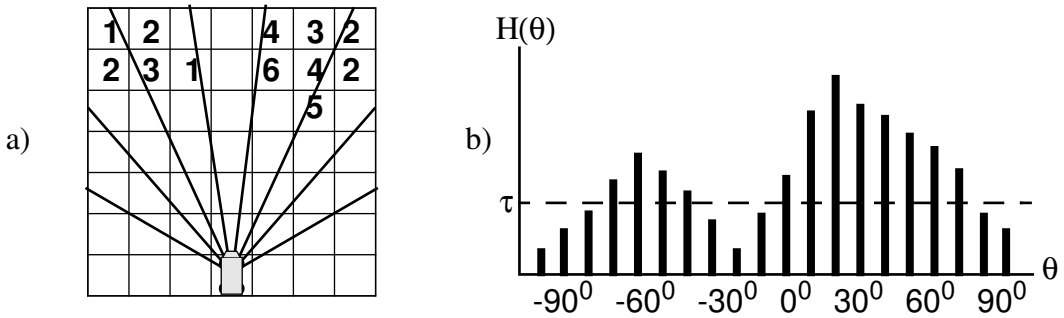

Figure A-7 : Vector Field Histogram.
a) Histogram Grid, b) Polar Histogram

POTENTIAL FIELDS

Potential Fields is a method of planning robot trajectories based on combining the vector fields induced by mapped objects, such as repulsive fields from obstacles and attractive fields from goals [8] [37]. This approach to mobile robot control is very similar to the Vector Field model; as

such, it has become the control paradigm most closely associated with motor schema theory. Work on motor schemas has been directly inspired by the Vector Field model and uses potential fields to implement mobile robot analogs of the prey-attractant and barrier-repellent schemas, as well as motor schemas such as stay-on-path and avoid-moving-obstacle [7]. The result of using these combined fields for the control of a mobile robot is shown in Figure A-8a.

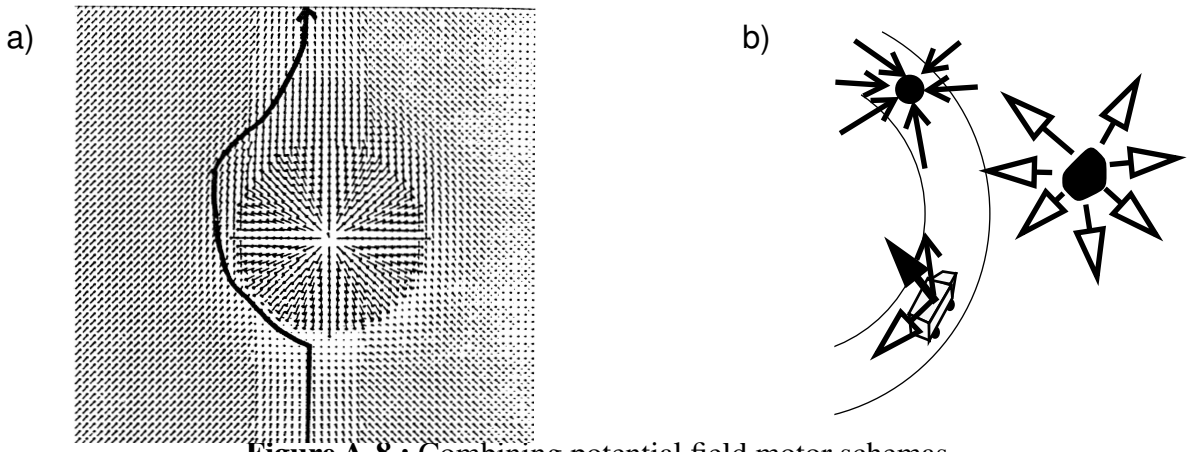

Figure A-8 : Combining potential field motor schemas.
a) combined fields for obstacle and goal, resulting path, b) instantaneous vector addition.
(Figure a reproduced from [7], with permission.)

In the Vector Field model, the entire field is created at once, and then a path is planned through it and followed to the goal. This would be advantageous if the obstacle and goal locations were known with certainty *a priori*; in practice, however, it has been found that it is necessary for a robot to reevaluate its path often in order to maintain reactiveness in a dynamic and uncertain environment. In more recent implementations of Potential Fields, only the force vector for the current robot position is calculated for each schema. The sum of these vectors is then used to control the instantaneous motion of the robot, as shown in Figure A-8b. The process is then repeated from the next robot position to be evaluated, taking into consideration any new perceptual information that may have become available.

Comparison of Implemented Schema Models

The three systems described above have all been implemented and used successfully for mobile robot navigation. The Vector Field Histogram is somewhat limited in that it only deals with obstacle avoidance while pursuing a single goal location. Both the Distributed Architecture for Mobile Navigation and the Motor Schema framework allow multiple behaviors to be instantiated and their outputs combined for a variety of different tasks and environments, and are therefore more general and also more closely mirror the schema paradigm for distributed processing and command fusion. While the Motor Schema framework offers a means of fusing commands from multiple behaviors, it suffers from the well known problem of local minima in potential fields, although there are various means of overcoming this limitation. Another, perhaps more serious

problem, is that arbitration via vector addition can result in a command which is not satisfactory to any of the contributing behaviors. The DAMN system does not average commands, but rather selects the command which has the most votes from the behaviors, and the VFH system selects the direction that is closest to the goal direction and unobstructed by obstacles. A related problem with a potential field implementation is that the robot cannot pass through closely spaced obstacles. This effect, which can be observed in Figure A-4c, was desirable in modelling frog behavior because the frog did not pass through closely spaced fence posts unless there was a prey object immediately behind exerting a large attractive force; however, this is problematic for robots that wish to explore areas that may only be accessible by passing through narrow entrances such as doorways. Borenstein also reported conditions under which potential fields suffer from oscillations and instability [12]. Potential fields have the advantage that the entire path may be computed if the obstacle and goal locations are known *a priori*; however, it has been found in practice that it is necessary for a robot to reevaluate its path often in order to maintain reactiveness in a dynamic and uncertain environment. The most significant lesson learned from these implemented systems is that entire paths cannot be planned before any action is taken, and then blindly followed until the goal is achieved. Both animals and robots must be reactive to changes in their surroundings and to new perceptual information as it becomes available.