Experiments with Staged Learning for a Continually Operating Robot

Gary R. Ireland and Ulrich Nehmzow
Department of Computer Science
University of Essex
Wivenhoe Park
Colchester CO4 3SQ
girela@essex.ac.uk, udfn@essex.ac.uk

Technical Report CSM-366

Department of Computer Science

University of Essex

14th February 2002

1 Introduction

1.1 Motivation

The use of robots is steadily increasing, particularly in manufacturing. Most industrial robots operate in strictly controlled environments. Those which are mobile navigate using beacons, markers or wires embedded into their surroundings([12, pg 8], with pre-installed knowledge of the use and position of the navigational aids.

There are, however, applications for autonomous mobile robots for which continuous, unsupervised operation over long periods of time (several days without interruption) is desirable, for instance surveillance, monitoring or cleaning tasks. These operations would typically be conducted in environments which are designed for human comfort and utility e.g. offices.

For reasons of cost and efficiency it is not practical to modify such an environment to allow robot operation. In such an unmodified environment, using pre-installed knowledge is problematic for the following reasons:

1. Perceptual discrepancy. Robot sensors do not necessarily respond well to the same environmental features a human sees well. A human operator may not be supplying the robot with the most useful data.

2. Inflexibility. Pre-installed competences may not be applicable to situations which were not foreseen at the time of coding.

3. Temporal invariance. Non-learning behaviours may not cope well with environmental changes over time.

4. Spatial invariance. Pre-installed knowledge often enables the robot to operate only in known environments.

5. Brittleness. Inaccurate modelling can produce incorrect operation.

6. Cost. Pre-installing models is time-consuming and costly.

7. Technical complexity. Modelling typically requires skilled operators.

It is, therefore, preferable for the robot to use machine learning techniques to *acquire* competences in interaction with its environment. Learning will allow the robot to exploit salient features of the environment, minimising the need for modifications to the environment. Additionally, if the robot has the ability to learn continuously it can adapt to changes in the environment, an exteremly useful ability when operating in an area in which humans will be working.

1.2 Purpose of the Experiments

Machine learning has regularly been applied to robot control (see [12] for a review), but, to our knowledge, *always* in *condensed* learning scenarios. A *condensed* scenario is a learning situation in which almost all stimuli presented to the robot are relevant to the competence being learned. Such *condensed* learning scenarios are suitable for acquiring competences under controlled conditions, but are likely to fail in real-world scenarios of *continuous* long-term operation, where many perceived stimuli are "irrelevant" to a specific competence.

Staged competence acquisition is an approach which allows the robot to overcome these problems by controlling *when* a competence is learned, and at *what rate* (see [13] for a further discussion).

The first trials with staged competence acquisition were conducted in [8]. During these trials competences for short and middle range navigation were developed which would allow a robot to wander within a limited area and when low on power return to a charging station and re-charge.

The experiments reported here continue these trials, expanding the number of competencies with an obstacle avoidance competence and seeking to establish the repeatability, reliability and utility of the competences. From the results of these experiments, decisions about the use of the competences within a continuous operation scenario are made and whether they can provide a basis from which which the problems with continuous operation and learning can be explored.

2 Method

2.1 Staged Competence Acquisition

Staged competence acquisition ([8]) is a learning process in which the complex behaviour of a robot is decomposed into simpler competencies. The robot acquires each competence when it is in the correct position within the environment and at the correct time of operation. Each competence is acquired until the required functionality can be produced through combination of these simple competencies. Some inspiration for this is drawn from biological systems, such as the stages of child development [17]. During this project staged competence acquisition used a single variable, the operational time of the robot to control the learning of the individual competences.

2.2 Robot Functionality

The robot used in the research, *Strange*, is capable of self re-charging from a specially designed charging station (see section 3.1).

The functionality required from the robot was:

- Wander randomly around the experimental area. When the on-board batteries reach a certain level of discharge return to the charging station, connect and re-charge the batteries. After re-charging resume wandering (figure 1).

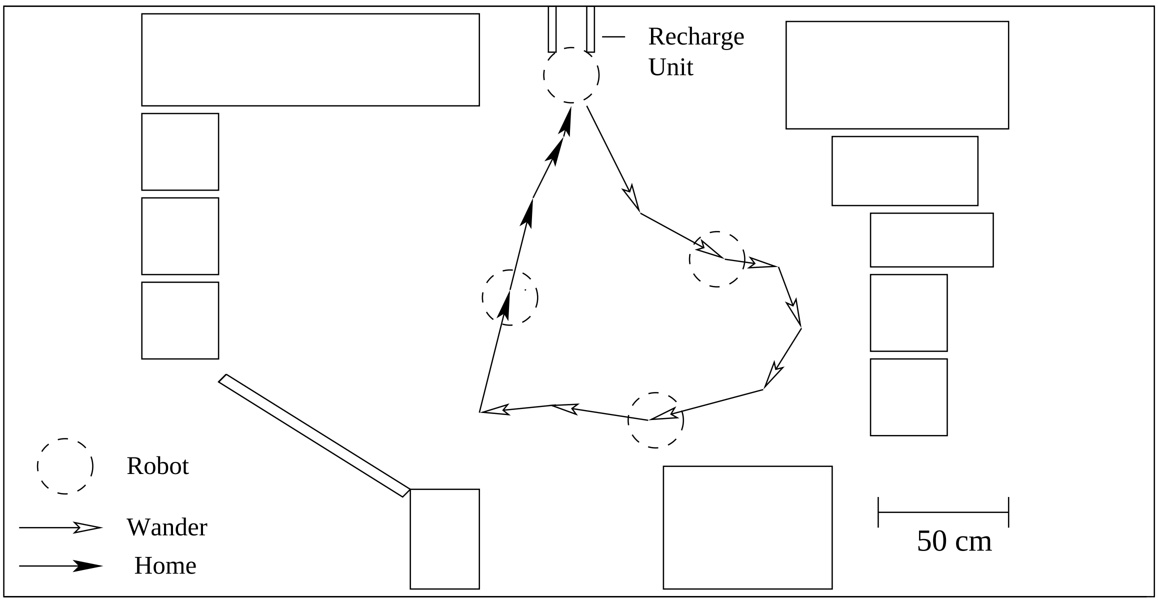

Figure 1: **The functionality required of the robot Strange:** Wander randomly and when batteries are low return to the charging station

For the duration of these experiments, the robot would be confined to the area shown in figure 1. This meant that there was a finite limit to mapping. It also meant that there were effectively two phases of robot operation. A learning phase during which movement and navigation competencies were developed and a wandering phase during which the effectiveness of the learning phase could be evaluated. The robot functionality describes the robots required behaviour during the wandering phase, i.e after all the competences have been acquired.

2.3 Robot Competencies

In order to gain the functionality the robot needed the following competences:

- Connect to charger. This is a reflex-like fixed behaviour which simply moves the robot forward until it makes contact with the charging station, whereupon the robot halts and charges.
- Avoid obstacles. To allow the robot to wander within the experimental environment without becoming trapped, or damaging either the robot or the environment.
- Short range navigation to the immediate proximity of the charger. Required when the robot is near the charging station to orientate the robot towards the charging station and then move forwards. Some performance measure is required to asses when the robot has moved close enough to the charging station for the *connect to charger* function to move the robot into final contact.
- Middle range navigation to the short range navigation area. Produces heading and distance data which allows the robot to find a vector which will return to some point within the short range navigation area (close to the charging station) from anywhere within the experimental area.

The first of these competencies *connect to charger* has no learning component. The other three competencies will use machine learning techniques to acquire their functionality. These competencies have a logical order in which they must be learnt:

- Whilst learning obstacle avoidance the robot will move randomly. The position of the robot will therefore be unknown once the obstacle avoidance competence acquisition has begun. Therefore the competence cannot be learned until the robot is capable of determining its position within the environment, or it will lose the charging station position before mapping has begun and be unable to return.
- Using middle range navigation to home to the general area of the charging station will serve no purpose if the robot is then incapable of assuming the correct heading to charge.

The order in which the competencies should be acquired is therefore:

- Short range navigation to immediate proximity of the charging station;
- Middle range navigation to the short range navigation area;
- Avoid obstacles.

However, if the robot is to map the whole environment for the middle range competence then it needs to start mapping while the short range competence is still being learnt. Therefore middle range navigation competence learning is active concurrently with short range navigation learning.

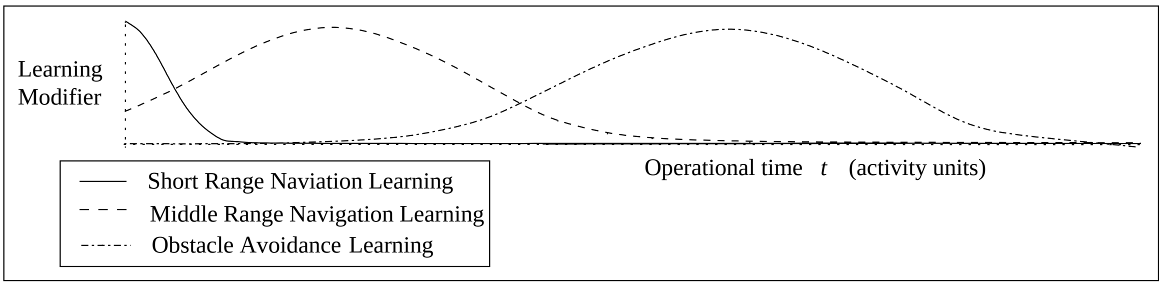

Figure 2: **Competence Staging for the Robot Strange the plot is the learning rate against operational time in activity units**

Figure 2) shows the three competences and when they are trained during the operation of the robot. The training of all three competences will be directly related to the slope of the graph at the operational time. Each competence curve is a function on the operational time. The operational time was not measured in fixed units of time, such as seconds, but in units of robot activity, (henceforth call activity units). An activity unit is a single robot action such as a translational movement, or a rotation. This abstraction away from a direct reading of time allows much greater flexability and robustness, removing the need to factor in absolutes of hardware performance such as different processor speeds or the robot's operational speed. Using activity units the speeds could be changed without needing to recalculate the activation function. The types of function and value of required constants were derived from experimentation. The competences are discussed in greater detail in the following sections.

2.4 Short Range Navigation to the Immediate Proximity of the Charging Station

For the robot to charge successfully it must approach the charging station in the correct orientation so that the robot can engage the contacts and begin charging.

An insect-learning model was used as inspiration for this competence. Honey bees undertake orientation flights, at the start of these flights the insects spend some time hovering close to the nest to ensure that the nest entrance can be recognised from many angles of approach ([5], [4], and [6]). During this hovering period the insect associates images from the compound eyes with the approach vector to the nest entrance.

In a similar manner, the robot would use a neural network to associate environmental information (sonar readings) within the area close to the charging station with vector data which would enable the robot to move towards the charging station.

This was implemented by placing the robot close to the charging station, facing it. The robot then backed out and stopped. Sonar readings were taken and then associated with the heading and distance data (captured through robot odometry) from the vector just travelled. The robot then returned to its starting point.

These vector movements were performed in a fan shaped pattern (figure 3) with the charging station at the focus of the fan and the vector movement along the 'spokes'. This pattern is created by inserting a rotation between translational movements. The pattern was defined as a table of rotational and translational data. The length of the 'spokes' of the pattern increased with operational time.

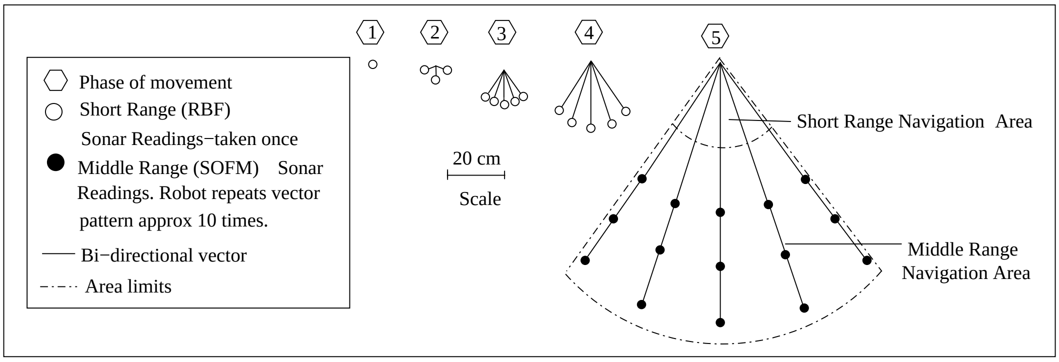

Figure 3: **Robot movement while learning short range and middle range navigation**

This could have been accomplished with random vector movements which would have been more 'natural' but could have lead to the over training of some areas and neglect of others. The selected approach is therefore more robust.

2.4.1 Implementing Short Range Navigation with a Radial Basis Function Network

The short range navigation competence required a network which trained rapidly and could map a continuous function across the area directly in front of the charging station. A Radial Basis Function (RBF) Network was selected because it fulfills both criteria ([15]).

The RBF network is composed of an input layer, a single hidden layer and an output layer. There are two sets of weights, the hidden layer weights between the input and hidden layer, and the output weights between the hidden and output layer.

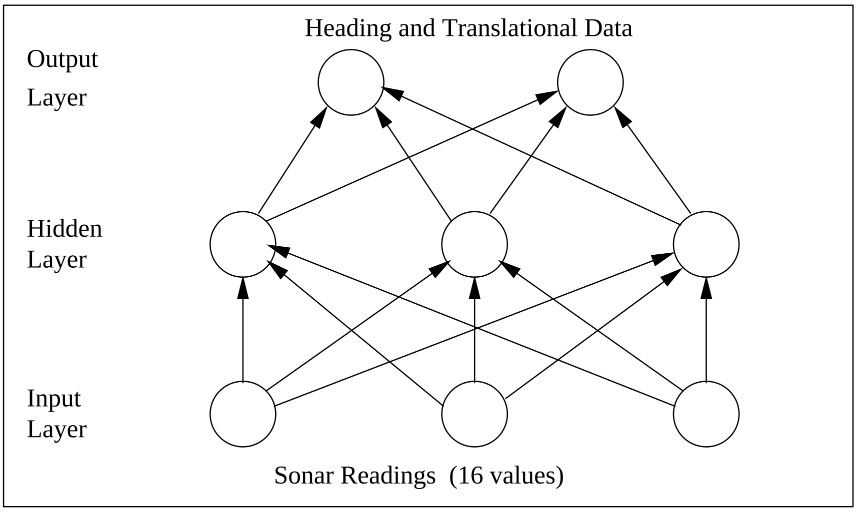

Figure 4: **Radial Basis Function Network: Used for Short Range Navigation**

The network used in the experiment had sixteen input units, one for each sonar sensor of the robot, 14 units in the radial basis layer, and two output units providing distance from, and angle to, the charging station.

While the network was being trained the robot took 14 sonar readings while moving in a fan shaped pattern close to the charging station (see figure 3). These readings were used as the weights of the 14 nodes in the the hidden layer. The readings were also presented to the network through the input layer. For each node in the hidden layer, the Euclidean distance between the input and the hidden node's weights was calculated. A Gaussian function was then applied to the generated value (equation 1). The is the most common radial function used with RBFs (figure 5).

$$h = exp(-\Sigma \frac{||i - w||}{\sigma}), \tag{1}$$

where h is the output of a hidden node, i is the input vector, w is the weight vector and σ is a constant factor.

The result of applying the function is to generate a non linear mapping between the hidden layer and the node.

The output units weights were then trained using a supervised learning schema, the Perceptron Learning Rule (equation 2).

$$\vec{w}(t+1) = \vec{w}(t) + \eta(t)(T - O)\vec{i}, \tag{2}$$

where t is the time at which the following was true, O is the observed output value, T is the target value and $\vec{i}$ is the input vector(in this case the output value from the hidden weight). $\vec{w}(x)$ is the weight vector at time x. η is a learning variable which is expressed as a function of time because it can be varied with time. Indeed, this property of the Perceptron Learning Rule was exploited during the project. The Perceptron Learning Rule is applied to all weights in the output layer.

The output layer is a simple linear mapping. Each node in the output layer multiplies the output vectors from the hidden layer by the output weight vector (equation 3).

$$o = \vec{i} \cdot \vec{w}, \tag{3}$$

where o is the output of the output node, $\vec{i}$ is the vector output from the hidden layer and $\vec{w}$ is the output layer weight vector. It is important to note that the results of all hidden node outputs are used (figure 5), unlike the 'winner takes all' architecture of nets like the SOFM. This is the mechanism by which the RBF is able to form a smooth mapping across the function.

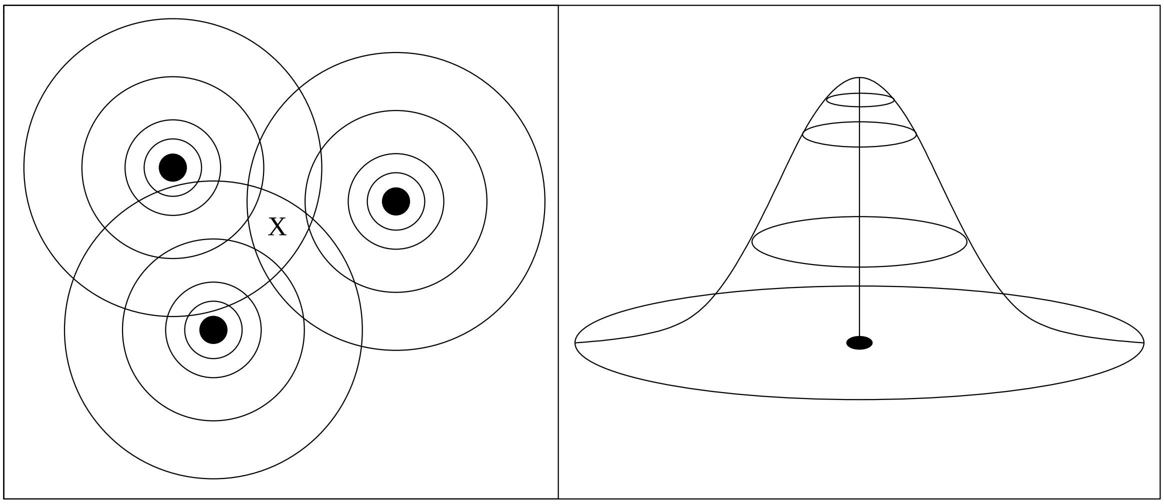

Figure 5: **RBF Continuous Function:** Left:An RBF with a 2 dimensional input space. Each dot is a 'position' which is used as the value of the weights of a hidden layer node. For values generated at the point x the output of the RBF is the 'distance' between the dots and x. One of the group of circles associated with an RBF node showing the Gaussian function shape

When the robot is using the network to provide data to return to the charging station a sonar reading is taken and supplied as input to the network. The output data (distance and heading) is then used to moev the robot toward the charging station.

2.5 Middle Range Navigation to the General Area of the Charging Station

In order to return from some point in the experimental area to the general area of the charging station the robot associates environmental features with a vector to the short range navigation competence's operational area. Several navigation competencies have been identified in social insects which serve as examples of how environmental landmarks are used.

Desert ants (*cataglphyis cursor*) live in an environment which often has few permanent local landmarks, and yet the ant needs to be able to determine a return vector to the nest when foraging. One of the ant's methods of navigation uses a permanent feature of the environment [18]. In addition to their two compound eyes, the ants are equipped with three more primitive eyes called ocelli. These are placed on the 'forehead' between the compound eyes. The ocelli are sensitive to polarised light, each at different rotations. With the ocelli the ant can determine the position it is facing, giving it a 'compass sense'. By combining this information with internally generated distance measurements, the ant can determine both the direction and distance of the nest from any point. Once the ant is within close proximity to the nest it can then use information from the compound eyes to recognise the opening to the nest.

Development of middle range navigation for the robot was based on this model. When the robot is learning how to return to the charging station it will use an odometrically generated heading and path data to determine distance and bearing from the charging station. The robot will then use a neural network to associate sensor readings at a given point with the homing vector information.

Since the middle range navigation competence is similar to the short range competence (but operating over a much larger area) this competence will use a similar fan shaped learning pattern. This is discussed further in the next section.

2.5.1 Implementing Middle Range Navigation with a Self Organising Feature Map

On the basis of previous successful implementations of middle range navigation ([11], [16] and [9]), the middle range navigation competence was implemented using a self-organising feature map (SOFM).

To robot was moved during learning in a fan shaped pattern, similar to the one introduced in section 2.4. The pattern differed because the area covered by the middle range navigation competence is much larger and because training the SOFM is incremental, i.e the robot revisits sonar reading positions. The robot was moved along straight lines from a position close to the charging station, within the short range navigation area, out to the edge of the mapping environment (figure 3). At certain points along the lines the robot halted, took sonar readings and then trained the SOFM by associating the sonar information with odometrically generated data, a homing vector back to the short range navigation area. The robot repeated the fan shaped pattern a number of times (in the case of this particular experiment the robot repeated the pattern 6 times).

The SOFM consists of a multidimensional array of nodes, in this case a two dimensional array of 8 by 8 (determined by experimentation. The weights of each node are equivalent to the size of the input vector.

In order to train a SOFM to create a mapping between two groups, an input is created which consists of two vectors, each vector contains values from one member of each group. In this experiment the vectors consist of data captured during robot movement. The first of these vectors is the sixteen sonar values and the second vector is a homing vector consisting of two values, a heading to, and distance from the short range navigation area.

This is a similar training approach to the one described by Owen and Nehmzow [16].

When executed the SOFM uses a competitive process, a 'winner takes all' mechanism to determine output. The node with the smallest Euclidean difference between its weights and the inputs is detemined to be the winner(equation 4).

$$i = argmin||x - w_j||, j = 1, \cdots, n, \tag{4}$$

where i is the index of the winning node, *arg min* is a function which selects the smallest argument, x is the input vector, w_j is each of the weight vectors for all the nodes in the network.

If the SOFM is still learning, the winning node and the surrounding nodes within a set neighbourhood of that node undergo learning, in order to create clusters.

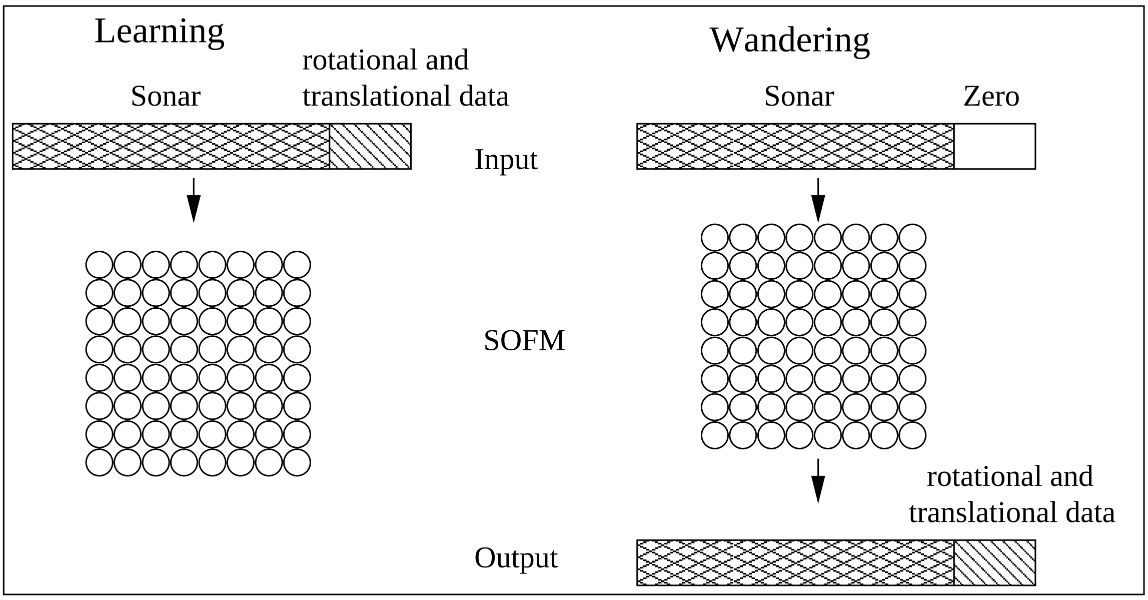

Figure 6: **Using the SOFM as a mapping function between sonar data readings and vector data**

When learning is finished, and the SOFM is to be used for data retrieval, the input vector is the sonar values, read at the robot's current position and zeroes, instead of the homing vector values (see figure 6). The output vector has the same format as the input vector, so the second group values (homing vector) are simply read from that portion of the output vector of the winning node (see [12, pg 121]).

2.6 Avoid Obstacles

Once the robot has mapped the environment it needs to be able to move freely within the environment without colliding, or becoming entangled with objects.

Many experiments have been conducted with insects to show how sensory input can trigger motor actions. An example is cockroach wall following, where directional change is mediated by antennal contact with objects ([3]).

The robot can use sensors to gain information about possible obstacles. A neural network can be used to associate sensory readings with motor actions. Once the neural network is trained the robot takes a sonar reading and uses this as input to the network. The output of the network is the appropriate motor action required to move the robot avoiding obstructions ([10]).

Initially a Radial Basis Function Network had been used to implement the obstacle avoidance competence. However, on very rare occasions the robot collided with the wall. The problem appeared to lie with the way that the Radial Basis Function operated. The output layer had been two 'motor neurons' i.e. each output node drove one of the motors. The output had been thresholded at zero so that the output value was interpreted as forward (zero and positive) or backwards (negative). Since each output node was generated from a summation of the weighted values of the four nodes in the hidden layer occasionally values of 'forward' would occur in both nodes even when there was an obstacle ahead. When this occurred one value would be 'strongly' forward and the other only 'marginly' so. Using a fraction of the output as the motor speed only increased the problem with the robot often following graceful arcs into the wall. Further tests are to be conducted with the RBF in obstacle avoidance to quantify the problem.

Because of the problems with the RBF a pattern associator was used for obstacle avoidance. The pattern associator was implemented in a simple and robust manner. This technique has been used many times before and a description appears in Nehmzow [12, pg. 58].

Basically, if there was no obstruction the robot moved forward. As the robot approached objects they were detected with the robot's five forward facing sonar units. If the range to the object was less than 40 cm the robot then tried different combinations of motor direction combinations until its path was no longer blocked. The network was then trained to associate the successful motor combination with the initial sonar pattern.

2.7 Integrating the Competences

The individual competences now need to be placed into a single framework from which their learning can be controlled. The framework will control when learning for a competence is active and, in the case of the SOFM (middle range navigation competence), what the learning rate is.

When robot operation starts the RBF (short range navigation) learning rate needs to be high and learning needs to be active, decreasing rapidly. The SOFM learning rate needs to be small initially, to increase and then again decrease subsequently. Obstacle avoidance is acquired after the navigation competences have been gained. These requirements have been implemented by controlling the learning rate of all networks using activation functions that take the robot's operational time t as its only parameter. These functions are shown in figure 2.

2.7.1 The Activation Function for the Short Range Navigation Competence

When first switched on the robot acquires the short range navigation competence. This requires the learning rate to be near 1.0 at the start of operation, to stay high for the first two minutes and then to fall rapidly below the activation threshold. An exponential function (equation 5) with negative slope was used to produce this, the activation curve λ_r is given in figure 2 and equation 5.

$$\lambda_r = e^{-t^2/c}. \tag{5}$$

In equation 5, t is the operational time of the robot in units of activity and c is a constant ($c = 500$). Short range navigation learning was only active if λ_r was above 0.3 the activation threshold, which was determined by experimentation.

2.7.2 The Sensitisation Function for the Middle Range Navigation Competence

Unlike the RBF, whose learning is completely inhibited below a value of $\lambda_r = 0.3$, the SOFM learning stayed active throughout the entire operation of the robot. The SOFM learning rate λ_s, however, was modulated by equation 6 (see figure 2).

$$\lambda_s = e^{-(((k-t)^2)/d)} \tag{6}$$

Again, t is the operational time of the robot, k and d are constants (k = 154 and d=70) and λ_s is the learning rate for the SOFM.

As can be seen from figure 2, the SOFM learning rate is initially very low. It peaks after about 2 minutes operational time, and then decreases again.

The constants were chosen so λ_s would be greater than 0.5 when the robot was moving within the short range navigation area over those positions specifically selected for SOFM learning. The value of λ_s would peak at 1.0 when the robot was half way through the SOFM learning period and then drop back to 0.5 as the robot left the short range navigation area. Achieving this required a value for k of 154 (RBF training time plus SOFM training time) and a value for d of 70 (half SOFM training time).

2.7.3 The Activation Function for the Obstacle Avoidance Competence

Once the navigation competencies have finished learning, the obstacle avoidance competence needs to be activated. Like the short range navigation competence the obstacle avoidance activation function is used as a 'switch' to turn learning on and off. When the value of λ_o is above 0.3 Obstacle Avoidance is active (see figure 2).

The obstacle avoidance activation function λ_o used the same equation as the middle range navigation competence, but with different constants, equation 7 (see figure 2).

$$\lambda_o = e^{-(\frac{((k-t)^2)}{d})} \tag{7}$$

Again, t is the operational time of the robot, k and d are constants (k = 280 and d=80) and λ_o is the activation function for obstacle avoidance.

2.7.4 Fusion of the Middle and Short Range Navigation Competence Homing Vector

When the robot is attempting to home to the charging station both the short and middle range navigation competence will generate homing vectors. However, due to the position of the robot within the experimental area, both, one or neither of these homing vectors will prove adequate for homing (see figure 7). In order to arbitrate between the values and decide which are valid, some measure of confidence is required. These are obtained from the competence's networks.

In the short range navigation's RBF network, the output of each unit in the RBF layer is proportional to the Euclidean distance of that unit's weight vector and the currently perceived input vector. In other words: if any RBF unit's output is high, it means a familiar input vector is presented to the network (i.e. the "confidence" value is high). Therefore the integrated activation of *all* RBF units was used as a confidence value in the network's prediction of distance and direction to the charging station.

In the middle range navigation competence's winner-takes-all SOFM network, the Euclidean distance between the winning node and the input vector can be used as a measure of confidence.

To make both confidence values numerically comparable, the SOFM confidence value is modified with the same operations as those applied to an RBF node. A further complication arises because the RBF value is a summation of the distance of the input vector from all nodes, not simply the nearest as in the SOFM. Analysis of the results of some simple experiments suggested that, in order to make the values comparable, the RBF value should be halved. This modification of both confidence values allows the results to be compared directly.

Several tests were conducted. Fourteen of the confidence values from these tests are displayed in figure 7.

At this stage visual inspection of the homing vectors and associated confidence values was used to assess what confidence value would signify an estimated heading within 10 degrees of the true heading. A confidence threshold of 0.03 was chosen. This confidence threshold did allow the occasional homing

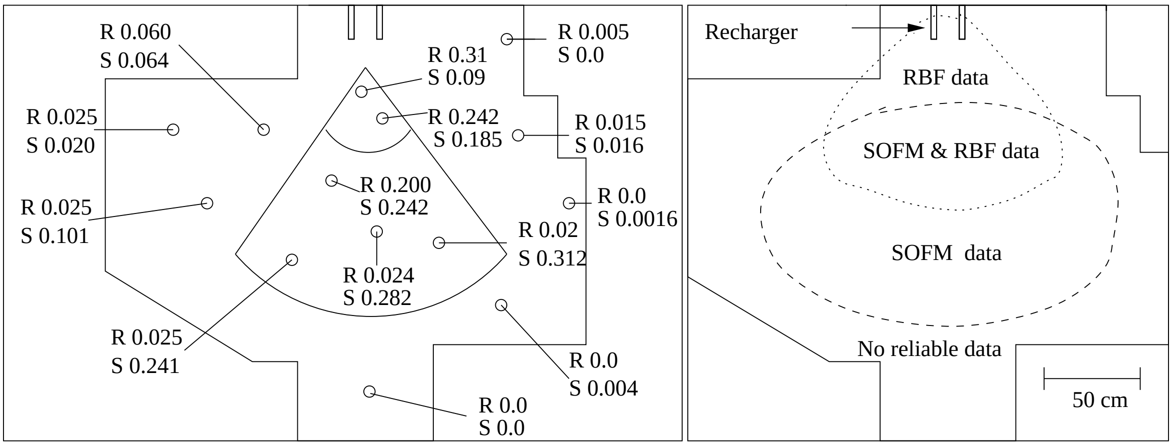

Figure 7: **Left: Confidence values at 14 robot positions. "R" indicates confidence values of the short range navigation competence's RBF network. "S" indicates confidence values of the middle range navigation competence's SOFM network. Right: Approx limits of confidence values $>$ 0.03, the confidence threshold**

vector with a heading error of greater than 10 degrees to be generated, but also maximised the areas covered by the short and middle range navigation competence.

When used during robot homing to the charging station, the confidence values generated by the networks are examined. Confidence values below the confidence threshold of 0.03 are ignored because the associated homing vector will be too inaccurate to enable the robot to home. If only one of the confidence values is above the confidence threshold of 0.03 the homing vector from that network alone is used. If both RBF and SOFM values are acceptable and one confidence value is a factor of 2 greater, then the homing vector from that network alone is used. Otherwise the data from both nets are used proportionally, see equations 8 and 9.

$$rot_{fused} = \frac{(rot_r \cdot c_r)}{(c_r + c_s)} + \frac{(rot_s \cdot c_s)}{(c_r + c_s)}, \tag{8}$$

with rot_{fused} being the final assumed direction to the charger, rot_r and rot_s being the assumed direction to the charger as indicated by the RBF net and the SOFM alone, and c_r and c_s the confidence values of the RBF net and the SOFM respectively.

$$dist_{fused} = \frac{(dist_r \cdot c_r)}{(c_r + c_s)} + \frac{(dist_s \cdot c_s)}{(c_r + c_s)}, \tag{9}$$

with $dist_{fused}$ being the final assumed distance to the charger, $dist_r$ and $dist_s$ being the assumed distance to the charger as indicated by the RBF net and the SOFM alone, and c_r and c_s the confidence values of the RBF net and the SOFM respectively.

3 The Experimental Setup

3.1 Experimental Hardware

The Nomad Scout is a mobile robot designed for research in robotics ([14]). It has a sixteen-sided body with a Polaroid sonar device in each facing (see [14, pg. 47] and [2, pg. 99]). A differential drive mechanism, with built in odometry, provides movement. Low level control of basic hardware functions are provided by a Motorola MC68332 processor running at 16 MHz which has a serial interface for external high level control.

When this robot was used in earlier experiments ([8]), it was using a prototype Amubot robot controller board for high level control. Through the board the robot also gained the ability to accurately measure the level of battery charge. The Amubot board also had a fluxgate compass which provided magnetic heading information. Unfortunately this board proved unreliable and is currently out of service. For the experiments described within this report

High level control was provided by a laptop computer placed on top of *Strange*, connected through the serial interface. With its large storage capacity, the use of the laptop as high level controller enabled the easy capture and analysis of data.

The Nomad Scout robot *Strange* has been modified so that it can connect to a specially designed charging station and recharge without human supervision. The charging station was designed as an MSc project by Roy Henderson [7]. The station is of simple and robust design with a wide capture area, necessities when *Strange*, weighing 23 kg has to periodically perform a controlled collision with the station, possibly with less than perfect precision.

Figure 8: **left to right: the anode contact, Strange contact, cathode contact, two views across the Experimental Area**

The charging station consists of a horizontally orientated board with copper plating on its lower face, this is the station's anode. Beneath it, two sprung cylindrical cathodes are mounted directly to the wall (figure 8). To connect to the charging station, *Strange* was fitted with a cathode and an anode. The cathode is a strip of metal ribbon pinned to a short pole which is attached to the top of the Scout. The anode is a piece of flexible copper plate with an adhesive back stuck along the upper edge of the front half of the robot's body (figure 8).

Although *Strange* can still self-charge without the Amubot board, it cannot accurately measure the charge rate. Also the laptop cannot yet be charged from the robot (a 20 volt line will be added in the future to enable this). For the experiments in this report, charge was undertaken when the robot had halted in contact with the cathodes (not live) of the charging station. The robot and laptop were then connected by hand to their respective charging equipment.

3.2 Experimental Procedure: Sonar Balancing

During the experiments the robot uses its internally generated odometry to provide both heading and a measure of distance moved. As the number of robot movements increases so does the error between the actual movement of the robot and the internally generated measure of movement. While odometry error is quite small with the Nomad Scout, over prolonged operation it can become a serious problem. In the experiments reported here translational odometry only needs to be accurate enough to record a single distance, a movement outwards so that the robot can then move back to the original position. This relative movement is used for both for training and homing purposes and is easily within the capability of the robot. Rotational error is a much greater problem since the robot uses absolute values of rotation to provide the robot with its heading and to rotate the sonar values before they are presented to the neural networks. Since an absolute value is used the cumulative effect of errors in rotation will have a severe effect on the accuracy of the results.

In order to compensate for the rotational error the robot uses an external measure with which to periodically 'reset' the rotational odometry of the robot. When the robot is first activated it is placed 45 cm from the charging station, facing towards it. The sonar is activated and the values of the sonars to the immediate left and right of the robots forward face (sonar numbers 15 and 1) are checked to see if they are equal. If they are not equal the robot is rotated to reduce the difference (figure 9). Once both sonars read the same the robot's internal odometry heading is set to 180 degrees. This will be refered to as sonar balancing in further discussions.

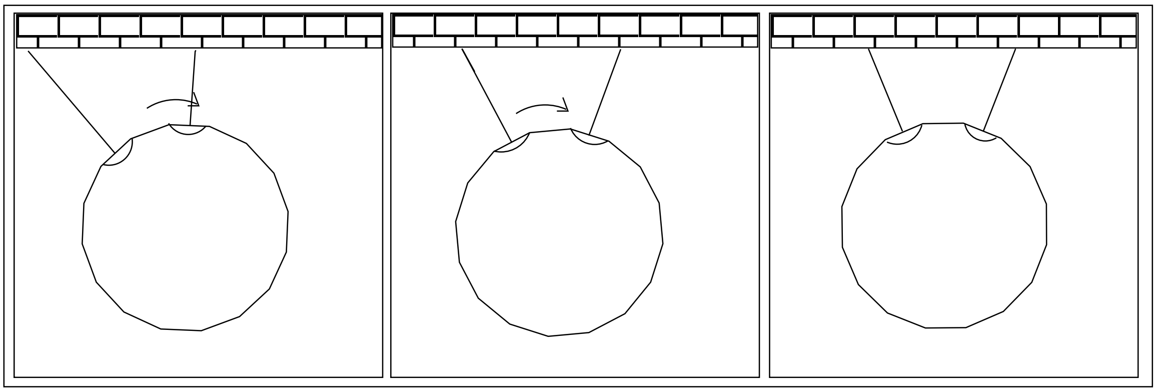

Figure 9: **Sonar Balancing: The robot is rotated until the values of sonar 1 and 15 are equal**

While operating, the robot then relies on the odometry heading alone. When a successful re-charge has occurred the robot is facing the charging station and a wall of known orientation. The robot then carries out the same sequence as before balancing the values of sonar 1 and 15 and then resetting the heading odometry to 180. This means that after each charge the robot's heading is reset, which is more than accurate enough for operation in the experimental area. The major problem with using the odometric heading is loss of bearing due to collision with obstacles. If the robot becomes rotated in a manner which is not recorded by odometry it would be difficult for the robot to regain the correct orientation. For this reason some time has been spent ensuring that emergency stopping mechanisms and obstacle avoidance prevent the robot from striking objects.

4 Experiments with Repeatability of Learning for the Short and Middle Range Navigation Competencies

Before conducting experiments with continuous operation a number of tests were carried out to establish what level of repeatability the learning techniques for short and middle range navigation possessed. If the results obtained in the repeatability experiments were close, this would suggest that the learning technique produces similar results each time it is trained and therefore produces comparable performance. This is important because if the accuracy of the network output varies wildly then it is difficult to make any claims about the usefulness and generalisability of the technique. The networks saved from one of the repeatability experiments (selected at random) were used for the continuous experiment, so that the overall accuracy of the vectors generated during the trial could be directly compared with the results from the repeatability experiment.

4.1 Experimental procedure

This experiment was conducted three times. The networks had the same learning parameters in each trial and each was conducted in the same experimental area, shown in figure 10.

4.1.1 Learning Phase: Short and Middle Range Navigation Learning

At the start of each experiment the robot was placed 45 cm from the charging station, facing the station, aligned so that a forward movement would result in connection between the robot and the station.

The robot initially moved away from the charging station, reversing directly without turning, and then moved back towards the station in the fan-like pattern of ever increasing size (described in section 2.4).

At this stage the robot had not developed sensory motor knowledge through learning with the Obstacle Avoidance competence. If an obstacle was detected in the robot's path (360 degree sonars allowing the detection of objects when moving backward or forward) the robot merely stopped and the next

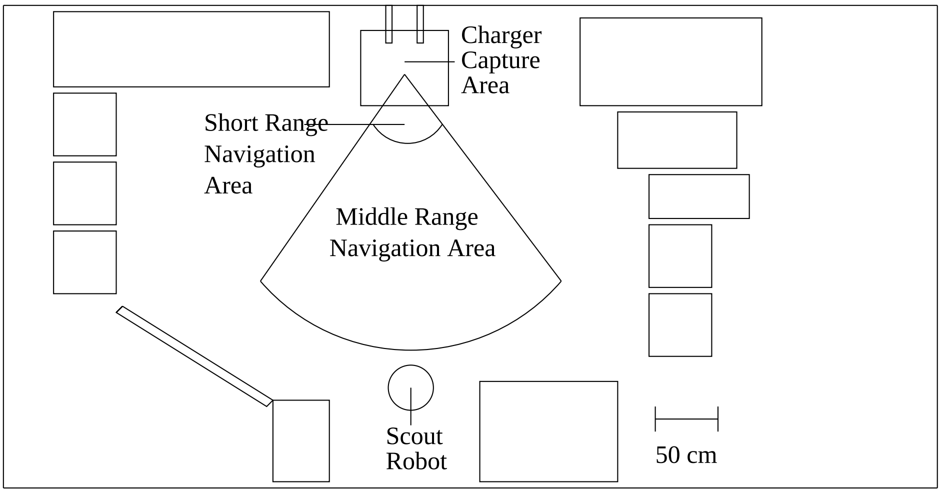

Figure 10: **Areas of operation for the competences**

movement in the fan shaped pattern was attempted. If this movement was in the same direction as the previous attempt this vector movement would be rejected and the next one attempted. This functionality was included to prevent collision with objects in the robot's path and also provides a control method for experimenting with learning problems which may arise when there are objects placed within the environment.

During this first phase short range navigation was learnt, controlled by the activation function λ_r. Middle range navigation learning had also begun, controlled by the sensitisation function λ_s. This phase lasted for about two minutes.

The process continued until the radius of the fan was approximately 50 cm (the short range navigation area in figure 10) .

At the next stage of learning, the robot moved in and out of the charging station at much greater distances (described in section 2.4). Middle range navigation continued to be learnt, with a higher learn rate than in the previous stage. This phase lasted about twenty minutes. The robot's movements during both learning phases are shown in figure 3. At the end of the learning phase the relevant weights and values of the three networks were saved.

4.1.2 Data Collection Phase: Recording the Responses of the Navigation Competences

After the networks had been trained the robot was moved through a fan shaped pattern which was similar to the pattern used during the learning phase (see section md:vm). The robot moved out along seven 'spokes'. Along each spoke the robot stopped 11 times. The first six stopping points were at 12.7 cm intervals, primarily for readings from the short range navigation competence. The last 5 stopping points were at 25.4 cm intervals, for readings from the middle range navigation competence.

At each stopping point the robot halted. Odometrically based readings of translation and heading to the charging station were recorded by the robot. The robot then rotated to a heading of zero degrees. Ten sonar readings were then taken, the robot rotating 36 degrees in an anti-clockwise direction between each reading. During continuous operation the robot could be at any orientation when a sonar reading is taken, the multiple sonar readings simulate this. For each sonar reading short and middle range navigation competence output was generated and collected. This meant that for each network there were 770 recorded homing vectors and confidence values.

4.2 Results

For each of the three experiments with repeatability of learning, the data generated by the robot, including the homing vectors generated by the short and middle range navigation competences, was analysed. Only homing vectors with a confidence value greater than the 0.03 confidence threshold were included

Network	No of Vectors	Mag.Mean	Mag.Median	Mag.sd	Mean	Median	sd
Middle 1	309	10.2	7.2	7.9	5.0	5.8	11.9
Middle 2	446	13.0	9.7	11.3	4.9	6.7	16.5
Middle 3	310	11.0	8.3	7.4	8.0	7.1	10.6
Middle Avrg.	355	11.4	8.4	8.9	6.0	6.5	13.0
Short 1	300	19.1	15.9	14.0	1.5	1.5	23.7
Short 2	304	16.3	15.2	11.3	-0.1	0.1	19.9
Short 3	276	19.6	17.5	13.4	-9.2	-9.3	21.9
Short Avrg.	293	18.3	16.2	12.9	-2.6	-2.6	21.8

Table 1: **Differences between the robot recorded odometric headings and homing vector headings generated by the short and middle range navigation competences during the 3 Repeatability Experiments, with averages for each competence. The units are degrees**

(see section 2.7.4). The differences between the homing vector headings and the robot recorded required headings were analysed. The results of this analysis are shown in table 1 as the mean, median and standard deviation for these differences. The table also contains the mean, median and standard deviation for the magnitude of the differences (i.e. absolute of differences). The magnitude statistics give a good indication of the size of the errors while the actual values any skewing of the values. Figures 12, 13 and 14 each present two heading diagrams which show the short and middle range navigation headings with a confidence value greater than 0.03, the confidence threshold. The circles in the diagrams are based upon those devised for displaying circular statistics in the biological sciences [1]. Figure 11 presents a key to the circles.

There was no analysis of homing vector translation data. Both navigation competences are optimised to provide accurate headings, since the absolute accuracy of the translation data is relatively unimportant.

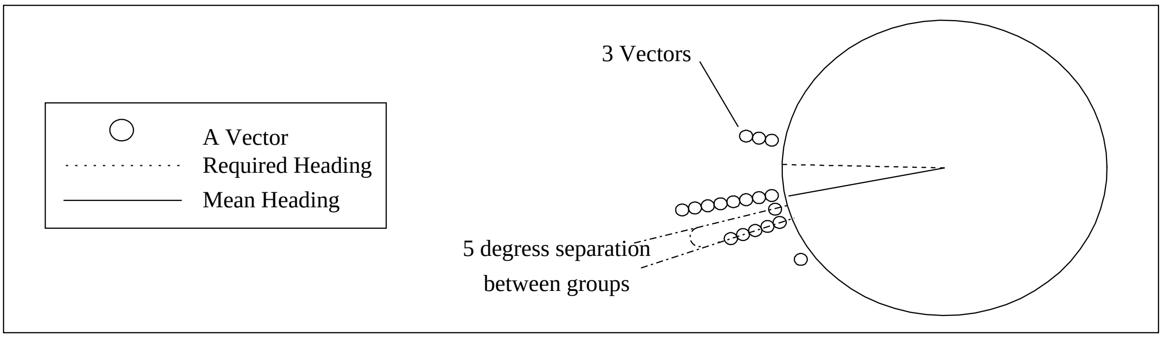

Figure 11: **A Key to the Vector Diagrams: The large circle represents an area within the experimental area. All vectors which were generated within the circular area are grouped into 5 degree wide intervals. The intervals are then displayed on the circumference as small circles, the number of circles in each group corresponding to the number of vectors in the interval. The required heading, i.e the heading to the charging station, is displayed as a dotted radius. The mean of all of the vectors is displayed as a solid radius.**

4.2.1 Results from the 3 Repeatability Experiments for the Short Range Navigation competence

The top diagram in figures 12, 13 and 14 show headings generated by the Radial Basis Function based short range navigation competence in each of the three repeatability experiments. Figures 12 and 13 (top diagrams) show very similar results, however, figure 14, the short range pattern from the third trial, shows a pronounced skew to the right of the charger capture area. This can clearly be seen if the mean

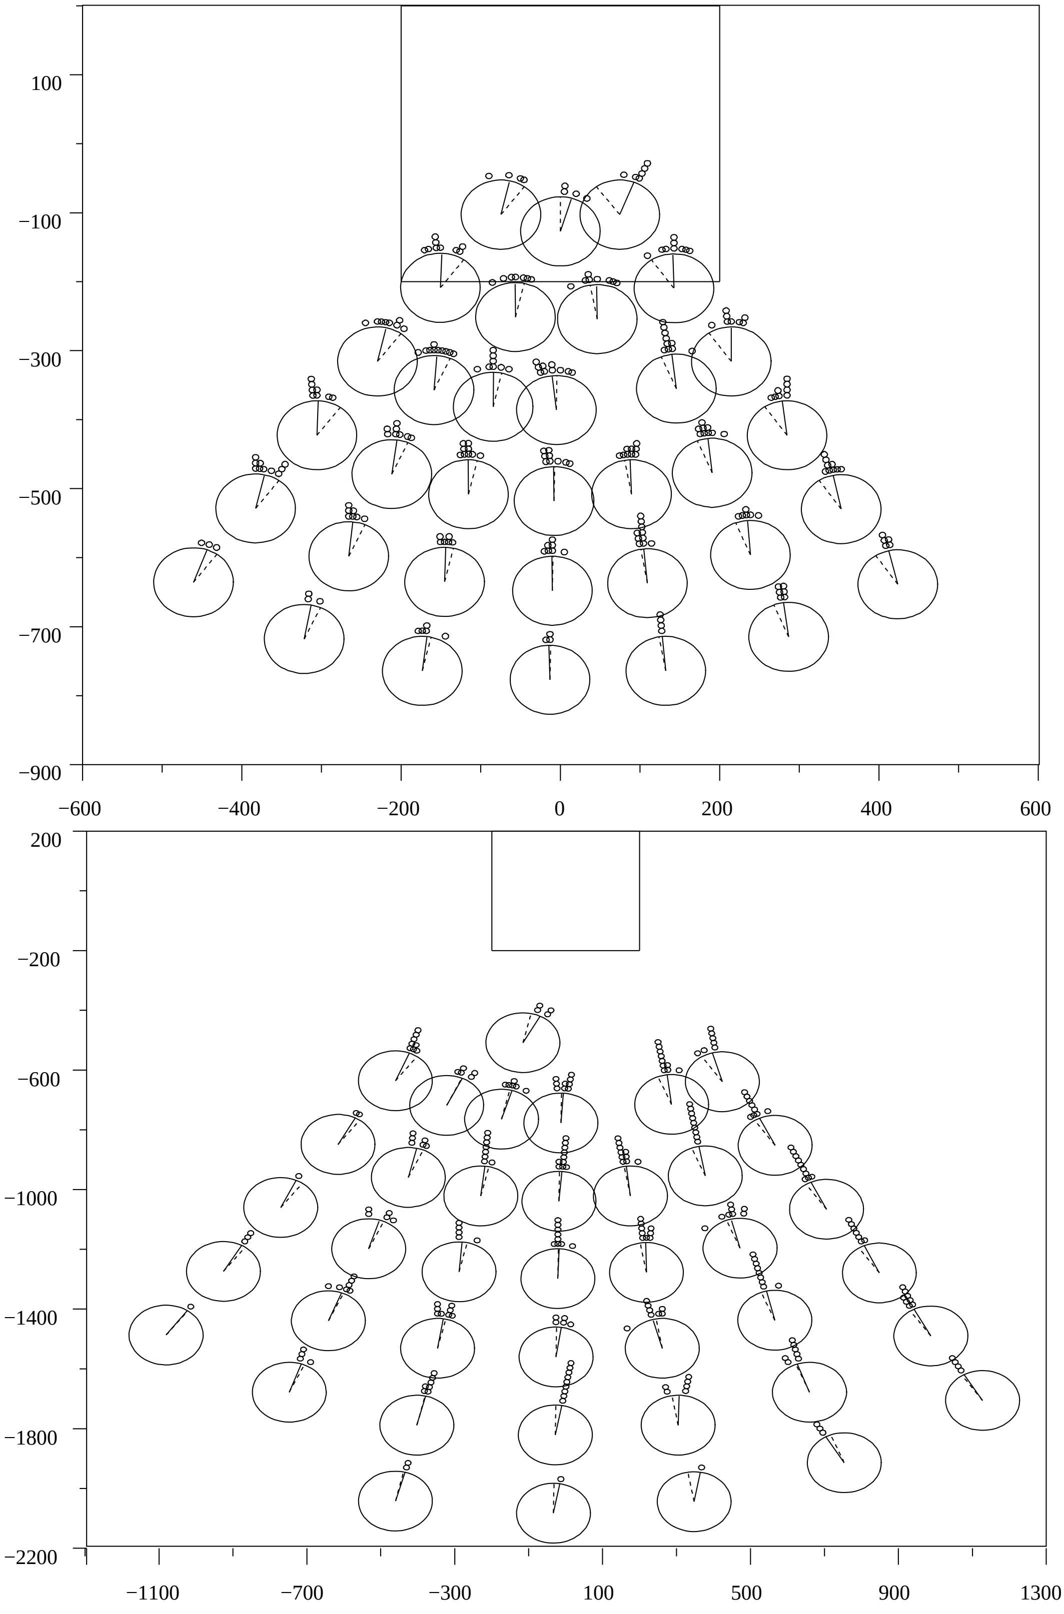

Figure 12: **Short (top) and middle (bottom) range navigation competence generated headings from the first of the 3 Repeatability Experiments: The square is the Charger Capture Area, the circles on in the diagram (short range) are 5 cm radius and in the bottom (middle range) are 10 cm. Axes are in mm.**

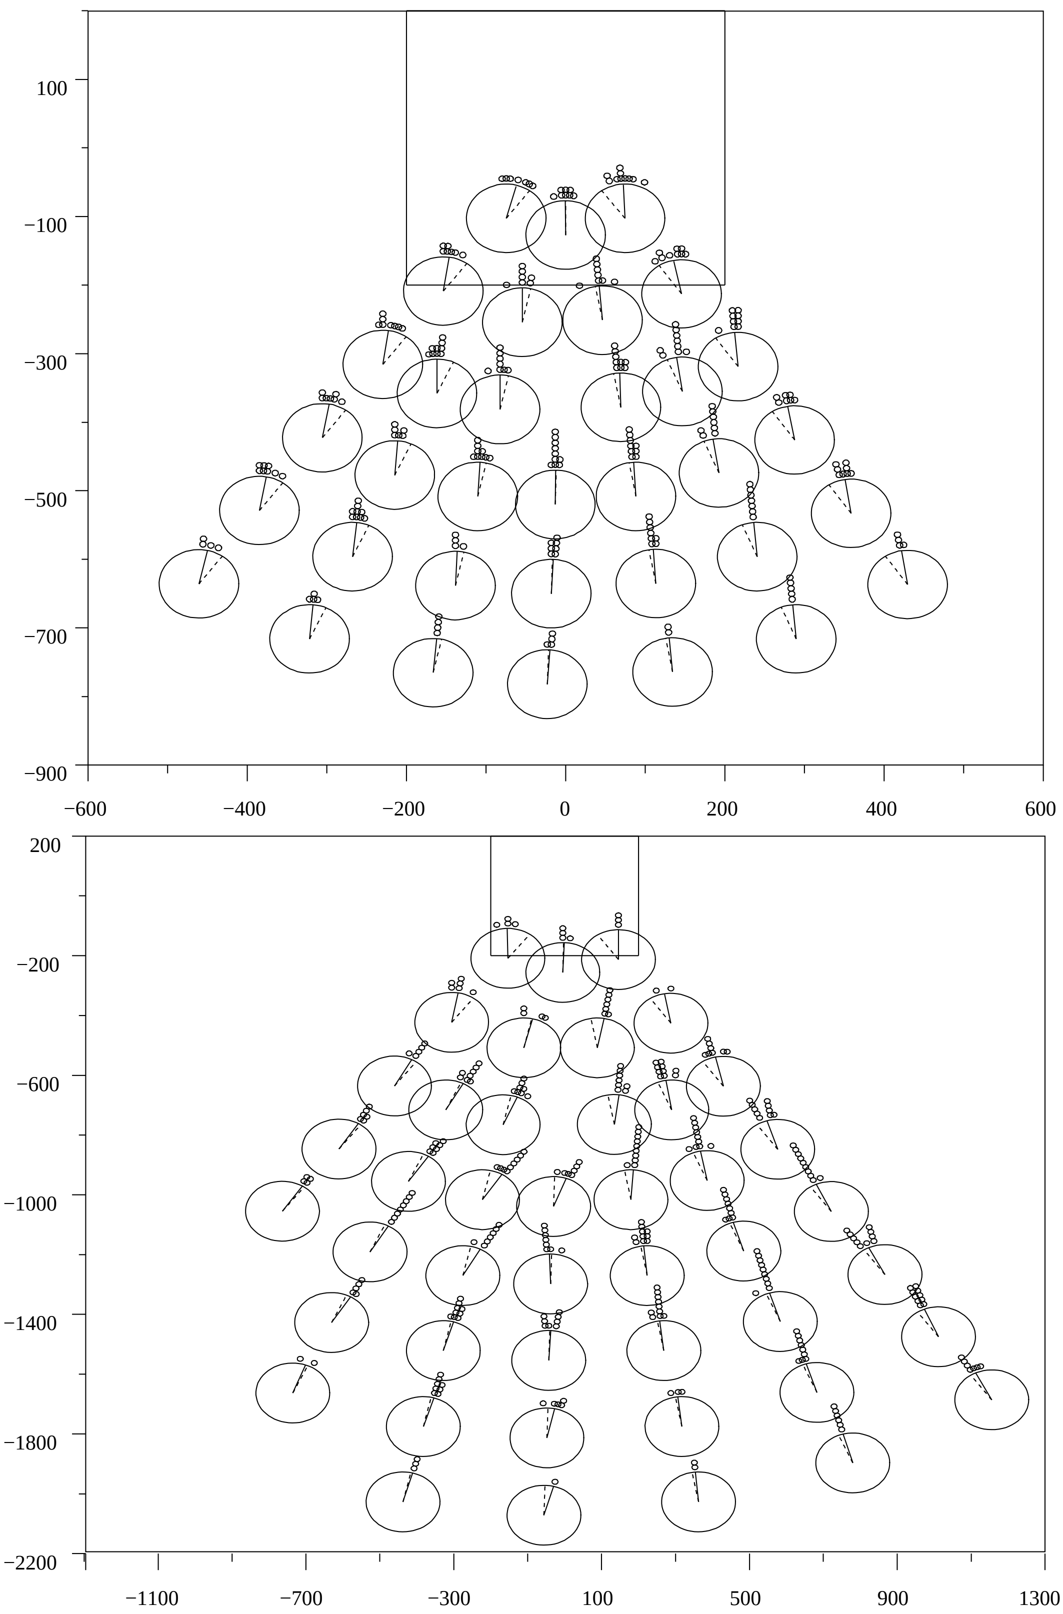

Figure 13: **Short (top) and middle (bottom) range navigation competence generated headings from the second of the 3 Repeatability Experiments: The square is the Charger Capture Area, the circles in the top diagram (short range) are 5 cm radius and in the bottom (middle range) are 10 cm. Axes are in mm.**

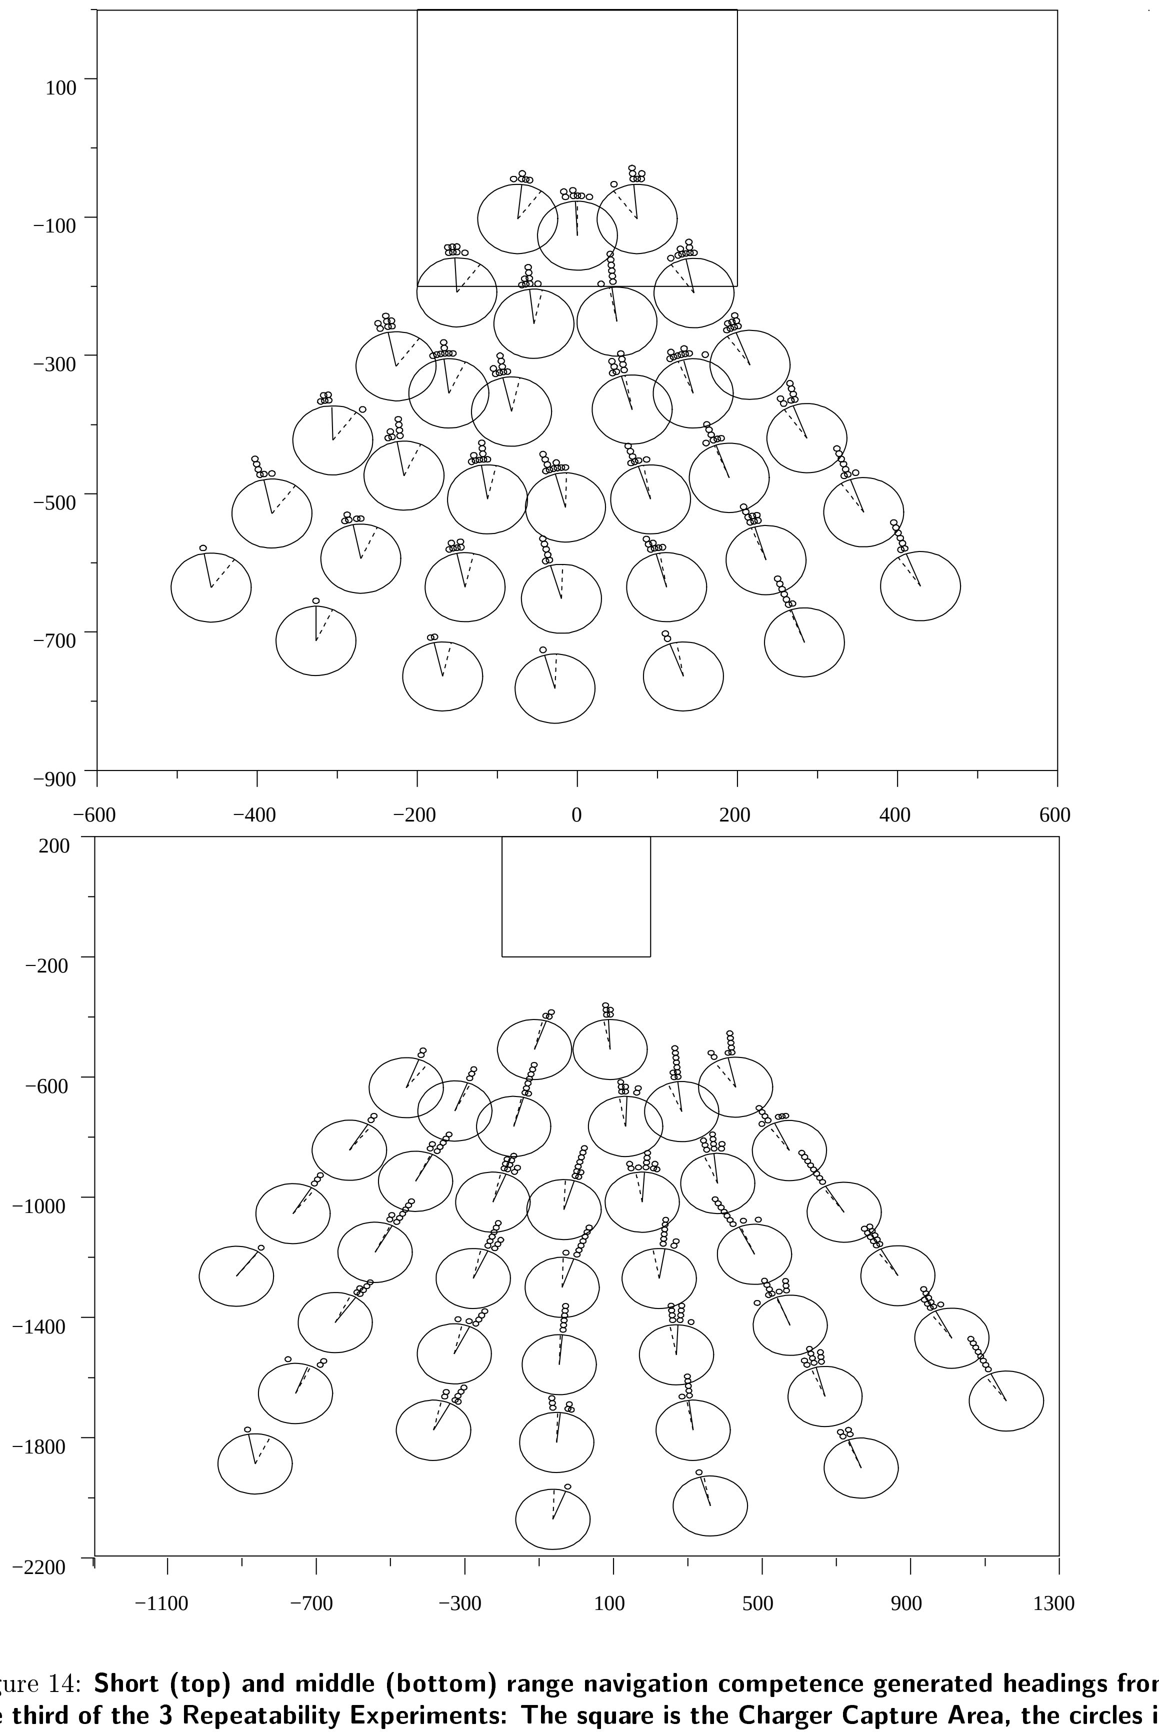

Figure 14: **Short (top) and middle (bottom) range navigation competence generated headings from the third of the 3 Repeatability Experiments: The square is the Charger Capture Area, the circles in the top diagram (short range) are 5 cm radius and in the bottom (middle range) are 10 cm. Axes are in mm.**

Source of Variation	Sums Of Squares	Degrees of Freedom	Mean Squared	F-ratio	Probability (p)	F-distribution Value
Between Groups	18834.51	2	9417.25	19.67	<0.01	4.61
Within Groups	419976.82	877	478.88			
Total	438811.33	879	9896.13			

Table 2: **ANOVA summary table for the difference between network (RBF) generated and required headings for short range navigation competence during the repeatability trial**

of ‘short 3’ (-9.2 degrees) is compared with ‘short 1’ (0.1 degrees) and ‘short 2’ (1.5 degrees) in table 1. The mean of ‘short 3’ of almost -10 degrees from the centre is far higher than the others.

To determine if differences in the results could be due to random variations within the trials, a One-Way Independent Groups Analysis of Variance (One-Way ANOVA) was conducted. ANOVA uses a measure of the mean squared of several samples within a population to determine if differences between the samples could have arisen probabilistically or are the result of some outside factor (the independent variable). In this case the independent variable is membership of a particular trial. The Null hypothesis was that the difference between samples could be explained as random variation. The ANOVA was conducted with a significance level of less than 0.01 (i.e a probability of 99%). The results are displayed in table 2.

The f-ratio of 19.67 was substantially larger than the f-distribution value of 4.61. This means that there is less than a one per cent chance that the three samples are simply random samples from a single population and it therefore appears that there is some influence from an external factor. In order to see which of these samples could be assumed to be part of the same population Tukey’s Honestly Significant Difference (HSD) statistic was then applied to the ANOVA data. The HSD score, of 5.26 is compared with the absolute of difference in mean between pairs of samples. If the difference is greater than the HSD then there is a significant difference. Table 3 shows the results of the comparisons.

Competences	short 1 & short 2	short 1 & short 3	short 2 &short 3
Difference Between Means	1.63	**10.68**	**9.06**

Table 3: **The difference between the means of the network (RBF) generated and required headings for the short range navigation competence during the repeatability trial. Values in bold are above the HSD value of 5.26 and are statistically significant**

This confirmed the earlier comparisons of the means. Based upon the means, the network generated by the short range navigation competence in the third trial is significantly different from the networks generated during the first two trials. Some influence from the environment has affected the nature of learning for the short range navigation competence during the third trial.

It initially appears that the short range competence from the third trial had consistently used an angle with some minus 10 degree deviation from the required heading for training (figure 14, top). But this is not reflected in the middle range navigation competence results from the third trial (figure 14, bottom), so an inaccurate odometric heading is unlikely to be the cause (the middle range navigation competence also uses the odometric heading and there is no form of odometric correction between training the two competencies). Indeed, of the middle scale navigation competence results, the third trial shows the largest deflection in the opposite direction (heading difference mean of 8 degrees)!

Inaccuracies of heading may result from one of, or a combination of, the mechanisms used by the robot, such as, sonar balancing and rotation of sonar readings to a neutral position. The absolute accuracy of these methods has not been individually tested so they may introduce errors of a level greater than those reported here. However, the short range navigation competence is used over a comparatively small distance (approximately 50 cm maximum). If the robot used short range navigation readings, with a constant 10 degree inaccuracy, to travel from the edge of this 50 cm area to the charging station, an error of less than 10 cm offset from the charging station centre would occur. The charging station is 40 cm wide, so this error would not cause a complete failure of the system, but rather would produce a behaviour which ‘favoured’ approaches from one side of the charging station. Hence only if the magnitude

Source of Variation	Sums Of Squares	Degrees of Freedom	Mean Squared	F-ratio	Probability (p)	F-distribution Value
Between groups	1864.55	2	932.28	5.55	¡0.01	4.61
Within Groups	147325.82	877	167.99			
Total	149190.38	879	1100.27			

Table 4: **ANOVA summary table for the magnitudes of the differences between network (RBF) generated and required headings for short range navigation competence during the repeatability trial**

of the error were much greater than 10 degrees would the ability of the robot to home to the charging station be compromised.

Because the system is quite tolerant to small heading deviations, a better test of the similarity of trial results is to look at the magnitude of the differences between the competence generated headings and required headings. This measure looks at the size of the error, not the direction.

The Null hypothesis was that the magnitude of the difference between samples could be explained as random variation. These values were compared across the three trials, again using a One-Way ANOVA test, with a significance level of less than 0.01 (a probability of 99%). The results are presented in figure 4

In this instance, whilst the f-ratio of 5.55 was still greater than the f-value of 4.61, it was only marginally so. This compares with the much larger f-ratio value of 19.67 seen when the direction of the error was included (see table 2), thus suggesting that differences in the direction of errors for the three samples were responsible for the greater part of the f-ratio value. Since the direction of error, ten degrees to the left or ten degrees to the right of the charging station, is relatively unimportant the test without direction of error is a more meaningful measure.

Tukey's HSD test was again applied to see which of the samples could be assumed to be part of the same population. Here the HSD value was 3.12. As table 5 illustrates, differences between the trials are not as marked when comparing mean magnitude of the differences. The third trial is statistically significantly different to the second trial, but the difference is quite small (3.3 degrees).

Competences	short 1 & short 2	short 1 & short 3	short 2 &short 3
Difference Between Means	2.78	0.53	**3.3**

Table 5: **The magnitude of the difference between the means of the network (RBF) generated and required headings for the short range navigation competence during the repeatability trial. Values in bold are above the HSD value of 3.12 and are significant**

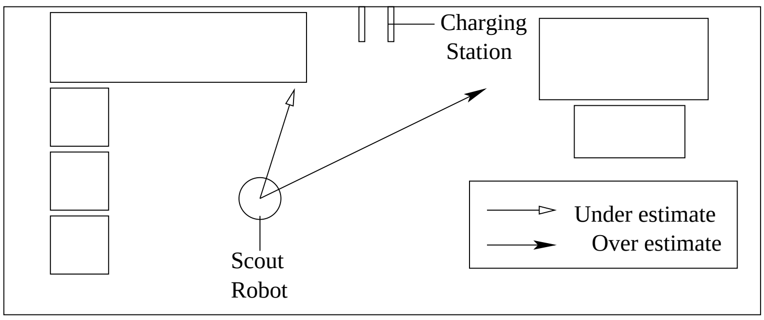

Figure 15: **Robot heading terminology: An underestimated angle occurs if the robot is on one side of the charging station, and the heading would cause the robot to miss the station on that side, an overestimation occurs if the robot would miss the station on the opposite side**

Another potential cause of failure to charge comes from the tendency in trial 1 and trial 2 for the RBF to to underestimate the angle required to place the robot on a heading towards the centre of the charging

Source of Variation	Sums Of Squares	Degrees of Freedom	Mean Squared	F-ratio	Probability (p)	F-distribution Value
Between Groups	1601.81	2	800.91	9.18	<0.01	4.61
Within Groups	92651.64	1062	87.24			
Total	94253.45	1064	888.15			

Table 6: **ANOVA summary table for the magnitude of the difference between network (SOFM) generated headings and required headings for the middle range navigation competence during the repeatability trial**

station. An underestimated angle occurs if the robot is on one side of the charging station, and the heading would cause the robot to miss the station on that side, an overestimation means the robot would miss the station on the opposite side (see figure 15). This becomes more pronounced towards the 'edges' of the group. Note how in figures 12 and 13, top diagrams the mean angles (solid lines) consistently under estimate the required angles (dotted lines). In fact the same tendency can be seen in figure 14 for trial 3 but the minus ten degree mean heading offset 'corrects' the fault on the right edge and accentuates it on the opposite edge. This could lead to the robot consistently failing to charge by missing the charging station when approaching from either side, unless the robot has some strategy for coping with failure to charge.

4.2.2 Results from the 3 Repeatability Experiments for the Middle Range Navigation competence

The bottom diagram in figures 12, 13 and 14 show headings generated with the Self Organising Feature Map-based middle range navigation competence during the three repeatability experiments. The headings all have associated confidence values greater than the confidence threshold of 0.03. While the first (figure 12, bottom) and third (figure 14, bottom) patterns are similar, the second (figure 13, bottom) produces headings much closer to the charging station. This is atypical and the reason for this is unclear, but many of the headings, included at the 'edges' and 'point' of the fan shaped area appear to be of poor accuracy (i.e. the difference between the mean network generated heading and the required heading is typically quite large). The descriptive statistics presented in (table 1, 'middle' entries) verify this. Although the second network (trial 2) produced at least 130 more vectors with confidence values above the 0.03 confidence threshold than either trial 1 or trial 3, the standard deviation of the magnitude was much higher, 11.3 degrees, rather than 7.9 (trial 1) and 7.4 (trial 3) degrees. In fact, the second trial middle range navigation competence has a pattern similar to one produced with a lower confidence threshold (see figure 16). Since the network which the middle range navigation competence generated was preserved it will be possible to conduct further trials to assess what impact the increased inaccuracies, apparent in this network, may have on the viability of the middle range navigation competence.

In order to establish whether observed differences were genuine or had arisen by random chance alone, a One-Way ANOVA was conducted on the magnitude of the differences between network generated headings and required headings for the middle range navigation competence. The magnitude was used because the variance rather than the skew, was being examined. Once more the Null hypothesis was that the difference between samples could be explained as random variation. The results are presented in table 6.

The results of this analysis suggests that there is less than a one percent chance that all three samples could belong to the same population with only random error to account or the differences. Tukey's HSD test was once more used to determine which of the middle range navigation competence trials were significantly different from the others. The HSD score was 2.04, at 99% probability. The HSD score is compared with the absolute of difference in means between pairs of samples. If the difference is equal to or greater than the HSD, then there is a significant difference. Table 7 shows the results of the comparisons.

HSD analysis of the middle range navigation competence data (the 'middle' mean magnitude scores in table 1) show that the second trial is significantly different from the first, and is quite different, from the third, though the difference did not reach the level of statistical significance. This supports the conclusions arrived at from earlier visual analysis of figures 12, 13 and 12 and table 1.

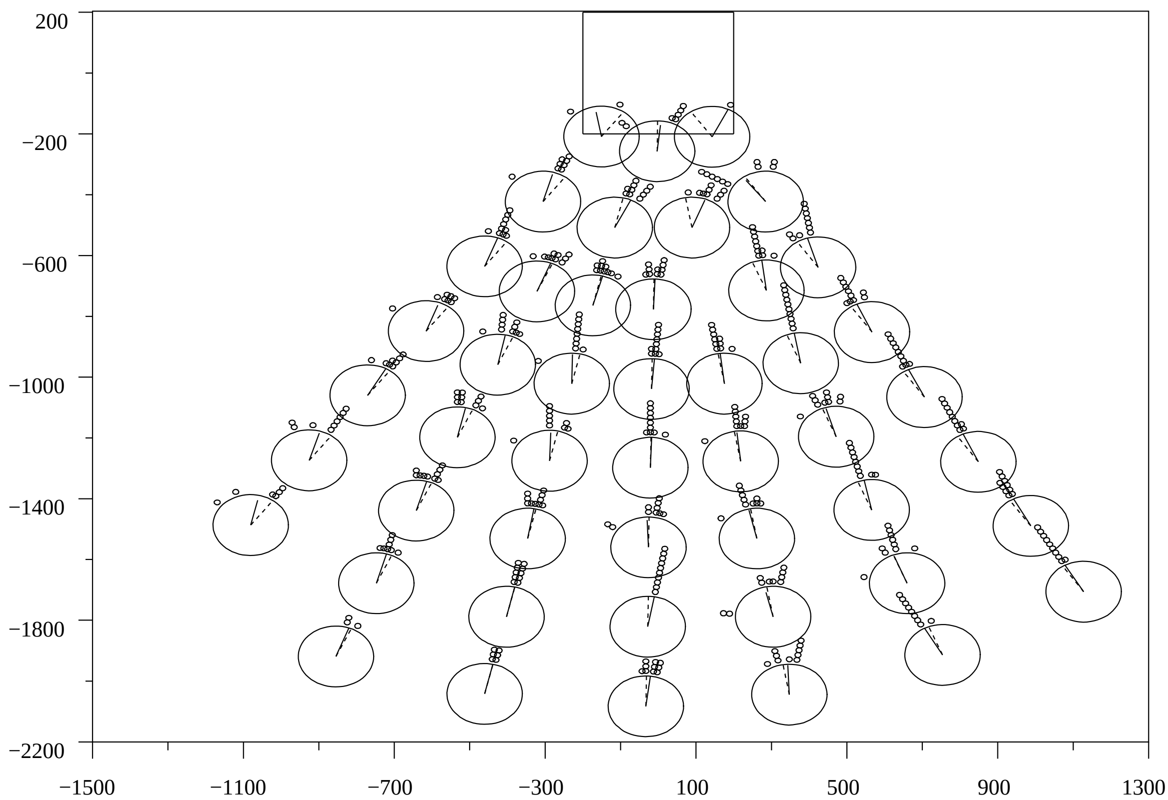

Figure 16: **The middle range navigation competence generated headings from the second of the 3 Repeatability Experiments: This is the same data as figure 12, bottom, but with a much lower confidence threshold, 0.004. Note the similarity with figure 13, bottom, which has a much higher confidence threshold of 0.03 and an atypically large number of homing vectors**

Competences	middle 1 & middle 2	middle 1 & middle 3	middle 2 & middle 3
Difference Between Means	**2.82**	0.83	1.98

Table 7: **The magnitude of the difference between the means of the network (SOFM) generated and required headings for the middle range navigation competence during the repeatability trial. Values in bold are equal to, or above, the HSD value of 2.04 and are significant**

Precisely what could have caused the middle range navigation competence in the second trial to create the observed pattern is unknown. Direct observation of the data generated during the three trials shows far higher SOFM (middle range navigation) confidence values in the second trial for the atypical homing vectors than in the data for the first and second trial (10 to 100 times greater).

Fortunately, should this problem recur, it is not likely to cause operational problems. The short range navigation competence is also active over the area (close to the charger capture area) where the network generated headings show the greatest error (see figure 13, top). Because data between the two navigation competences is fused (see section 2.7.4) the error will be reduced. Further, viewing the network generated data for the area close to the charger capture area it can be seen that the although the middle range navigation competence confidence value is often greater than the confidence threshold, the corresponding short range navigation competence confidence value is often much greater. If one confidence value is twice as large as the other the fusion function will use only the higher value (see section 2.7.4). Therefore quite often only the more accurate short range navigation competence value would be used.

4.2.3 Comparison of the Results of the Navigation Competences by Network Type

When the descriptive statistics for the three trials were grouped by network type (i.e. RBF for the short range competence and SOFM for the middle range competence), it was apparent that the RBF network produced far less accurate headings than the SOFM. If the average values for the two competences over the 3 trial repeatability experiment are compared (table 1), the standard deviation of the magnitude of the difference is 12.9 degrees for the RBF (short range) and 8.9 degrees for the SOFM (middle range).

This was surprising because in the earlier tests with the Amubot board and with a magnetic compass ([8]), the contrary had been found. However problems with hardware prevented extensive capture of information during the Amubot tests, and it was felt that the hardware faults could be causing memory problems. As the largest user of memory, the SOFM would have been most affected. Figure 17 presents the headings from the homing vectors with a confidence value above the 0.03 confidence threshold from the Amubot board experiments ([8]). Note that many of the SOFM (middle range navigation competence) homing vectors are very inaccurate (the vectors are supposed to point towards the charging station). It seems likely that the RBF is no less accurate than during the Amubot tests, but simply that the SOFM is now performing correctly.

The short and middle range navigation competence networks from the first repeatability test were selected randomly for use in the following Continuous Trial.

5 Experiments with Continuous Operation

5.1 Experimental procedure

Rather than using learning to gain the short and middle range navigation competences, the Radial Basis Function (short range) and Self Organising Feature Map (middle range) values from the first repeatability trial were loaded (figure 4). By using a network of known qualities, direct comparisons between the results from the continuous and repeatability experiments could be made.

5.1.1 Learning Obstacle Avoidance

Once the robot had loaded the networks enabling the robot to home to the charging station, the next task was to learn how to avoid objects within the experimental area in a more sophisticated manner than merely stopping (the robot's behaviour during short and middle range navigation competence acquisition).

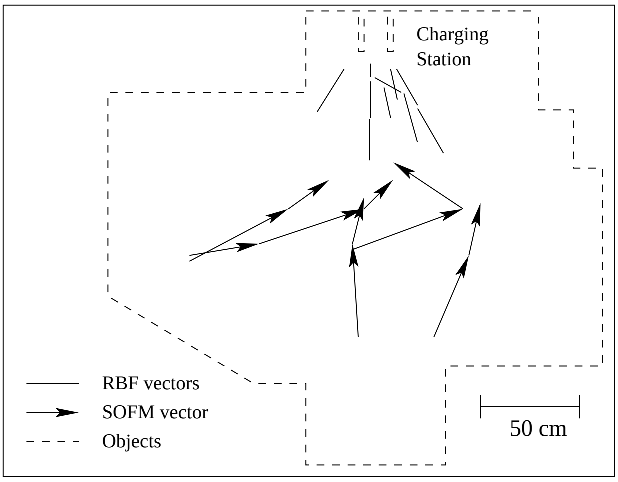

Figure 17: **Short (RBF) and middle (SOFM) range navigation vectors from the earlier Amubot board experiment: The poor accuracy of the SOFM headings is of particular note**

While wandering, the robot needs to be able to navigate without colliding or becoming entrapped. Using the instinct rules described in section 2.6, the robot first learned to move forward when no objects were encountered. As the robot approached objects they were detected with the robot's sonar. The robot then tried different combinations of motor direction combinations until the its path was no longer blocked. The network was then trained to associate the motor combination with the sonar pattern.

5.1.2 Wandering, Homing and Charging

After obstacle avoidance had been acquired the the robot began wandering within the experimental area. The robot's course was changed when an obstacle was encountered.

In 'normal' operation the robot would continue to wander until the batteries required charging, but because the robot's batteries last tens of hours, a quicker 'mock' discharge was simulated to greatly increase the attempted rate of charges. In fact, the 'mock discharge' was set at a rate which would allow the robot sufficient time to wander from the charging station to one of the 'walls' of the experimental area and then use obstacle avoidance to move away from the wall for a few seconds. Once this mock battery level was below a predefined limit the robot behaviour changed. The robot halted and took sonar readings using odometry-based heading data to rotate the sonar readings to a uniform 'zero' direction (this is possible because the 16 sonar sensors are evenly distributed around the circumference of the robot). This was used as input to the RBF (short range) and SOFM (middle range) networks to generate confidence and homing vector values.

If both network-generated confidence values were below or equal to 0.03, the confidence threshold, wandering continued, the robot repeatedly moving, stopping and taking readings until a position was reached where confidence values greater than the confidence threshold were found. When reliable confidence values were obtained, the robot recorded its current heading and then rotated to the network generated heading. It then recorded the generated values and new heading data, and paused whilst external data recording of the robot's position within the experimental area (tape measure) and heading (protractor) was undertaken. The robot then moved forward for 20 percent of the generated translation.

This process was repeated until the RBF generated (short range navigation competence) homing vector's translation value had dropped below 50 cm, or the translation value was less than 100 cm and the robot's sonar detected an object closer than 36 cm ahead in the robot's path (this second rule allowed for inaccuracies in translation values). The connect-to-charger function then assumed control, overseeing connection to the charger.

The connect-to-charger function moved the robot forward 30 cm (though the robot would stop if sensors detected an obstruction at less than 20 cm). Since *Strange* currently has no manner of assessing whether contact has been made, the laptop keyboard is used to enter 'y' for contact and 'n' for no contact. If the cathode of the charging station was touching the re-charge plate on the robot, the charge was deemed successful, the robot and laptop were re-charged (if necessary) and the 'mock' battery level was set to full. Success or failure was recorded along with the displacement between the centre of the robot and the middle of the charging station.

If the charging was deemed successful, the robot moved back from the wall some 45 cm and then used sonar balancing to correct any heading errors in odometry. The robot then rotated by a random angle and moved some 3 metres from the charging station before resuming wandering.

In the case of failure to charge, the robot rotated 180 degrees and moved 1.5 metres before resuming homing behaviour.

If the mock battery level dropped below a value of 20 the robot was deemed to have insufficient power to operate and the robot halted ending the trial. This behaviour never occurred during testing.

This continuous trial lasted some seven hours, including time for collection of external data.

5.2 Results

Although the trial lasted some seven hours, this included robot and laptop re-charge times and halting to record 'real world' robot positional data. By the time the trial was terminated 1012 vectors had been recorded. Along with the results of the repeatability trials, this constitutes a large body of data. Analysis of the continuous trial data began with a review of the accuracy of short and middle range navigation competences and comparisons with the repeatability data.

Network	No of Vectors	Mag.Mean	Mag.Median	Mag.sd	Mean	Median	sd
Middle	477	10.0	7.5	11.3	1.4	3.2	15.1
Short	420	21.7	16.5	17.8	6.5	7.6	27.3
Fused	75	12.3	10.6	8.9	-4.92	3.7	15.2

Table 8: Differences Between Network Generated and Actual Headings from the Continuous Trial

5.2.1 Results from the Continuous Trial for the Short Range Navigation competence

The top diagram in figure 18 shows the short range navigation competence headings, with associated confidence values above the confidence threshold of 0.03, generated during homing towards the charging station during the continuous trial. When compared with the results generated during the repeatability trials, (figures 12, 13 and 14), the continuous trial headings clearly cover a larger area. However most of the headings generated in the 'outer regions' are quite inaccurate. A re-evaluation of the confidence threshold for acceptance of short range navigation competence homing vectors may be beneficial.

By comparing the continuous trial statistics (table 8, 'short' row), with the repeatability trial statistics (table 1, 'short' rows) it can be seen that although the continuous trial mean (6.5 degrees) is obviously very different from those generated during the repeatability trials (1.5, -0.1 and -9.2 degrees), the continuous trial magnitude mean (i.e the mean ignoring the sign of the error) of 21.7 degrees is only a little higher than the repeatability magnitude means (19.1, 16.3, and 19.6 degrees).

The Null hypothesis formulated was that the magnitudes of the difference between the short range navigation competence generated headings and the required headings for the continuous and repeatability trials could be considered to have been drawn from the same population. A One-Way ANOVA was conducted to compare the short range navigation competence generated heading data from all three repeatability trials with data from the continuous trial. The results are presented in table 11. The f-ratio value, 8.0, is greater than f-distribution value, 3.78, this would indicate that the samples cannot be considered to have come from the same population.

Post-hoc Tukey tests were then conducted on the ANOVA data to explore how the samples compared with one another, the HSD score being 3.59 in this instance. The only significant difference was between the third repeatability trial and the continuous trial results. Overall the trial results were very similar.

As in the repeatability trial (section 4.2.1), the RBF (short range navigation) produced headings which suffered quite badly from underestimating the angle to the charging station.

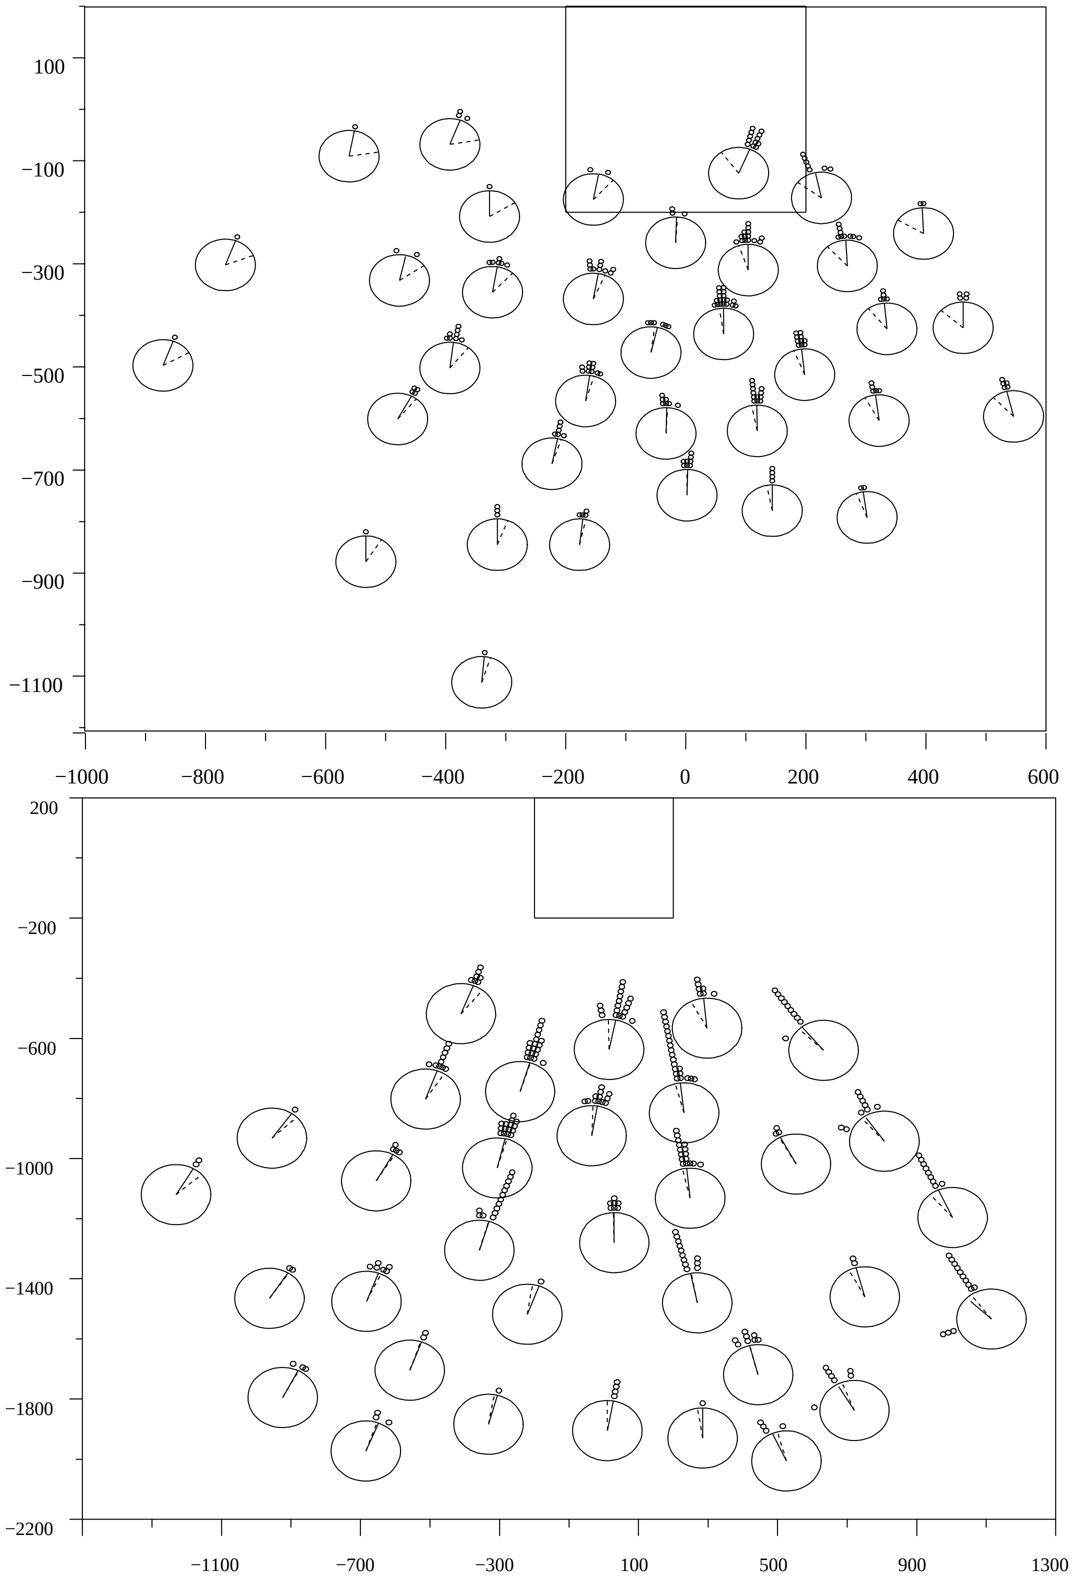

Figure 18: **Short (top) and middle (bottom) range navigation competence generated headings from the Continuous Experiment: The square is the charger capture area, the circles in the top diagram (short range) are 5 cm radius and in the bottom (middle range) are 10 cm. Axes are in mm.**

Source of Variation	Sums Of Squares	Degrees of Freedom	Mean Squared	F-ratio	Probability (p)	F-distribution Value
Between Groups	5191.63	3	1730.54	8.0	>0.01	3.78
Within Groups	280291.69	1296	216.27			
Total	285483.32	1299	1946.82			

Table 9: **ANOVA summary table for the magnitude of the difference between network (RBF) generated headings and required headings for the short range navigation competence during the repeatability and continuous trials**

Competences	short 1 and short 2	short 1 and short 3	short 1 and short C	short 2 and short 3	short 2 and short C	short 3 and short C
Difference Between Means	2.78	0.53	2.63	3.3	**5.4**	2.1

Table 10: **The means of the magnitudes of the difference between the network (RBF) generated and required headings for the short range navigation competence during the repeatability trials and the continuous ('short C') trial. Values in bold are equal to, or above, the HSD value of 3.59 and are statistically significant**

5.2.2 Results from the Continuous Trial for the Middle Range Navigation competence

The differences between the middle range navigation competence generated headings and required headings appear in table 8, 'middle' row. The table presents descriptive statistics for the homing vectors, whose confidence value were greater than the 0.03 confidence threshold. These should be compared with table 1, particularly the row for 'middle 1' which is the repeatability experiment statistics for the network values which were loaded at the start of the continuous experiment. The values for 'middle' and 'middle 1' are similar, but the standard deviations are larger in the continuous experiment. This may be due to the inclusion of homing vectors from positions not covered by the test pattern in the repeatability experiment. Diagrammatic evidence supports this. When compared with figure 12, bottom, there are far more outlying vectors in figure 18, bottom, most of which appear to be quite inaccurate. The mean is actually smaller in the continuous trial data (1.4 degrees) than it was in trial 1 of the repeatability experiment (5.0 degrees). This is notable because it shows that a difference in means can arise between usages of the **same** network.

Obviously, during the continuous trial the robot was not following a set pattern as it did during the repeatability trials. The difference in the form of sampling could be account for the larger errors (the standard deviations for short and middle range navigation competences are larger the continuous trial than in any of the repeatability trials). It certainly accounts for the larger area covered by homing vectors above the 0.03 confidence threshold (compare the area covered in figure 18 with figures 12, 13 and 14. It may be beneficial to replicate the repeatability experiments with the same network values loaded and then compare the results to try and determine if there is an external factor which is responsible for the discrepancy, which, if it exists, could then be further explored. When replicating the repeatability trials it may be useful to broaden the capture area to include outlying homing vectors.

Though the continuous trial mean (1.4 degrees) is obviously different from those generated during the repeatability trials (5.0, 4.9 and 8.0 degrees), the continuous trial magnitude mean (i.e the mean ignoring the sign of the error) of 10.0 degrees is very similar to the repeatability magnitude means (10.2, 13.0 and 11.0 degrees). Additionally the magnitude mean is less sensitive to small heading errors and shows more clearly the level of general error across the data.

The Null hypothesis formulated was that the values observed for magnitude of difference between the short range navigation competence-generated headings and the required headings for the continuous and repeatability trials had arisen from samples which could be considered to have been drawn from the same population. In order to compare the continuous and repeatability trial data further, a One-Way ANOVA was conducted upon the middle range navigation competence generated heading data from all three repeatability trials and the continuous trial. The results of the ANOVA are presented in table 11.

Source of Variation	Sums Of Squares	Degrees of Freedom	Mean Squared	F-ratio	Probability (p)	F-distribution Value
Between Groups	2454.62	3	818.21	8.18	¡0.01	3.78
Within Groups	153911.31	1538	100.07			
Total	156365.93	1541	918.28			

Table 11: **ANOVA summary table for the magnitudes of the differences between network (SOFM) generated headings and required headings for the middle range navigation competence during the repeatability trial and the continuous trial**

The f-ratio value is greater than f-distribution value (substantially so), which indicates that there is less than 1% chance that the samples could have come from the same population.

A post hoc Tukey's HSD test was then conducted to further explore differences between the trials (table 12). The HSD in this case was 2.24.

Competences	middle 1 and middle 2	middle 1 and middle 3	middle 1 and middle C	middle 2 and middle 3	middle 2 and middle C	middle 3 and middle C
Difference Between Means	**2.82**	0.83	0.19	1.98	**3.0**	1.02

Table 12: **The means of the magnitudes of the difference between the network (SOFM) generated and required headings for the middle range navigation competence during the repeatability trials and the continuous ('middle C') trial. Values in bold are equal to, or above, the HSD value of 2.24 and show statistically significant differences**

As can be seen, both instances of a statistically significant difference trial involved 'middle 2', the second repeatability trials middle range navigation competence output. 'middle 2' (trial 2) had actually been found to be significantly different from the other two repeatability test results (section 4.2.2) and had shown atypical learning, so this was to be expected. The relevant observation is the overall comparability of the continuous trial and the repeatability trial.

5.2.3 Results from the Continuous Trial for the Fusion of Homing Vectors

When both the navigation competences produce a confidence value above the confidence threshold and there is less than a factor of two difference between the confidence values, a *fused* homing vector consisting of proportional amounts of both navigation competence homing vectors (see equation 8) is generated. During the continuous trial there were 75 such homing vectors. Statistics for these homing vectors are shown in table 8, the 'fused' row and appear to be closer to the 'middle' row results than the 'short' row, i.e the *fused* homing vectors are more like the middle range navigation competence values.

Figure 19 presents the *fused* homing vector headings diagrammatically. When compared with figure 18, it can clearly be seen that the fused homing vectors overlap the short range navigation competence to quite a large degree. Since the statistics and diagrammatic representations clearly show that the fused homing vectors are more accurate than the short range navigation competence, in this instance, the fusion of data has a positive effect.

5.2.4 Results and Confidence Thresholds

Since there was now a large amount of data captured during a wander-home-charge cycle (the continuous trial) it was decided to determine what effect increasing the confidence threshold would have on the accuracy of the homing vectors produced by the navigation competences. A series of tests was conducted where confidence thresholds of increasing value (in the range 0.03 to 0.2 in steps of 0.01) were applied to the data captured during the continuous trial. At each step, only those homing vectors whose confidence values were higher than the new threshold were selected. As the confidence threshold is increased the number of acceptable vectors is reduced. If there is a correlation between confidence value and accuracy

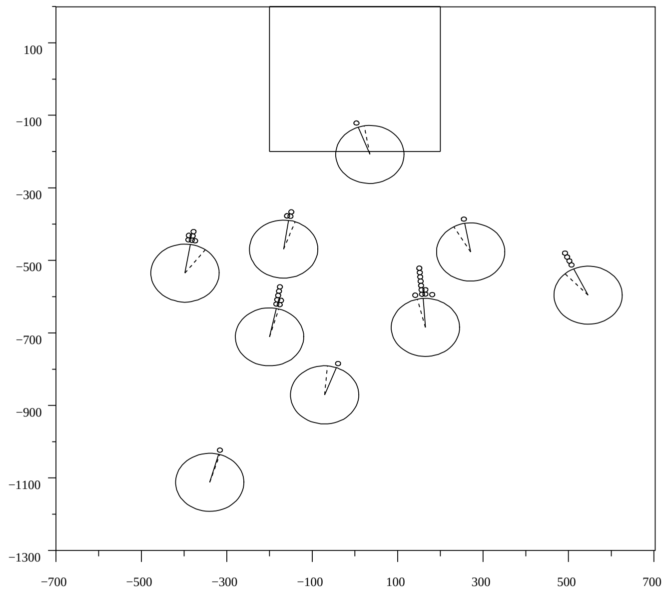

Figure 19: **Fused (middle and short range) homing vector headings from the continuous trial. The square is the Charger Capture Area. The circles are 8 cm radius. The axes are in mm.**

of data then as the confidence threshold increases so should the accuracy of accepted data. The selected homing vectors were analysed using descriptive statistics.

Figure 20 plots the standard deviations of the difference between the competence-generated heading and the required heading against varying confidence thresholds for the short (left box) and middle (right box) range navigation competences. The short range competence shows an almost linear relationship. Increasing the confidence value would reduce errors in the short range competence, but would also make the capture area proportionally smaller. Since the relationship is proportional, the question is simply what is an acceptable acceptable balance between area covered and error arising.

The middle range plot (figure 20, right) was far more useful. There is a sharp increase in accuracy up to a confidence threshold of 0.07, after which there is little change until much higher confidence values (0.17 and above). There is however a large decline in the number of vectors which are acceptable, from 477 at 0.03 confidence threshold to 246 at 0.07 confidence threshold.

To visualise the data, a vector diagram was produced for the short and middle range navigation competences. The diagram uses the data captured during the continuous trial but with a confidence threshold of 0.07. This should be compared with figure 18, also the continuous trial, data but with a confidence threshold of 0.03. For both competences, the only visible effect seems to be the removal of inaccurate values at the edges of the areas. As a result of analysing the statistical data and diagrams, it is intended that in future experiments with these networks the confidence threshold will be raised to 0.07.

5.2.5 Results: Charge Attempts

The robot operated continuouslly throughout the trial, charging before the ‘mock’ battery level had fallen to a level which would halt the experiment. However, the robot occassionally missed the charging station and attempted to dock nearby. When this occurred the simple strategy of reversing a short distance was followed. On subsequent charge attempts the robot moved closer to the charging station until finally making contact. Since charge was always successful after, at most, 4 attempts, it appears that the simple strategy (see section 5.1.2) employed on failing to connect with the charging station was quite robust. During the trial there were 69 recharges. Of these 39 (56.5%) were successful on the first attempt. No re-charge took more than four attempts (and there were only two occasions when four attempts were required). Table 13 presents data on the number of charges and attempts. The first row is the number of attempts which a re-charge required, the second row is the proportion of re-charge attempts. As further

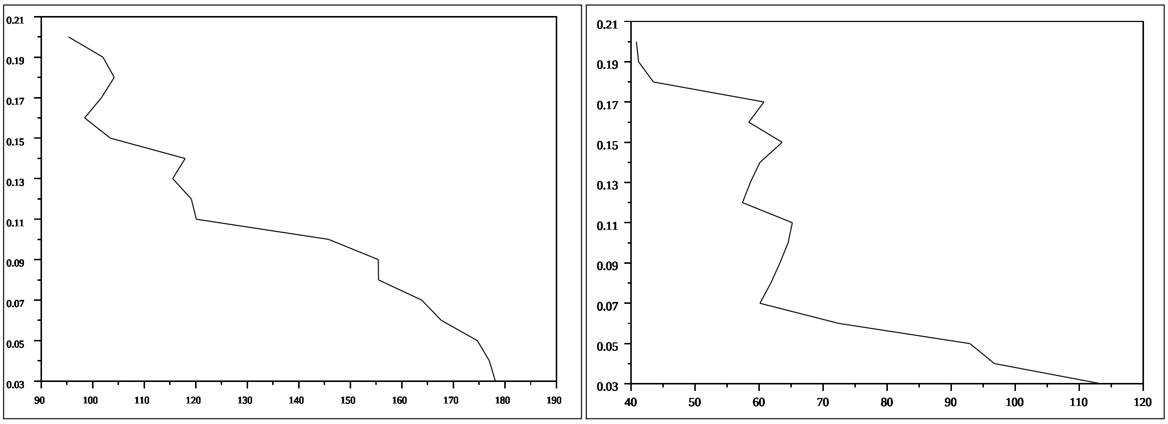

Figure 20: **Standard deviations of the difference between navigation competence generated headings and required headings plotted against confidence thresholds. The left diagram is the short range navigation competence and the right the middle range navigation competence.**

Attempt	1	2	3	4	5	Total	Mean	Median	sd
Number	39	21	7	2	0	69	1.6	1	0.8

Table 13: The number of attempts necessary to reach the charging station after each occassion when the battery level was below the re-charge threshold.

evidence of the robustness of the homing techniques employed, the median is 1, i.e charge on the first attempt, and the value for the standard deviation is 0.8. This information is presented graphically in figure 22.

Analysing several tracks of the robot during homing enabled the production of the simplified profile presented in figure 23. Failure occurred when the heading generated by the navigation competences underestimated the angle required to travel towards the charging station. An underestimated angle occurs if the robot is on one side of the charging station, and this type of heading would cause the robot to miss the station on that side. An overestimation means the robot would miss the station on the opposite side (see figure 15). Overestimations where extremely rare.

However, even if the homing vectors heading were underestimates, the angle of the heading invariably moved the robot towards the charging station. This tendency to underestimate can clearly be seen in figure 18, in both diagrams, particularly towards the 'edges' of the groups. The mean heading (solid radius) is often an underestimate of the required heading (dotted radius). After a charge failure, moving a short distance out, perpendicular to the charging station, placed the robot deeper into the navigation competences areas of operation, allowing a homing vector with a greater confidence value and a more accurate homing vector to be generated than on the previous attempt. The distance from the charging station was, therefore, always reduced on subsequent attempts. This simple strategy works because the error in heading is quite uniform and almost predictable.

5.2.6 Results: Time Required to Reach Charging Station

The robot was using 'mock' battery 'discharge' to maximize the volume of data collected. The total time between two charges does not, therefore, reflect true operational times because in this trial the wandering time is artificially short. However, on a charge attempt, the time from finding the first homing vector with a confidence value greater than the confidence threshold until time of connection at the charging station will be the same as under 'normal' operating conditions. This time was, therefore, examined and results are presented in table 14).

Note that the longest charge time is only 247 seconds, slightly over 4 minutes. Both the mean and median were just over one and a half minutes with a standard deviation from the mean of just 40 seconds.

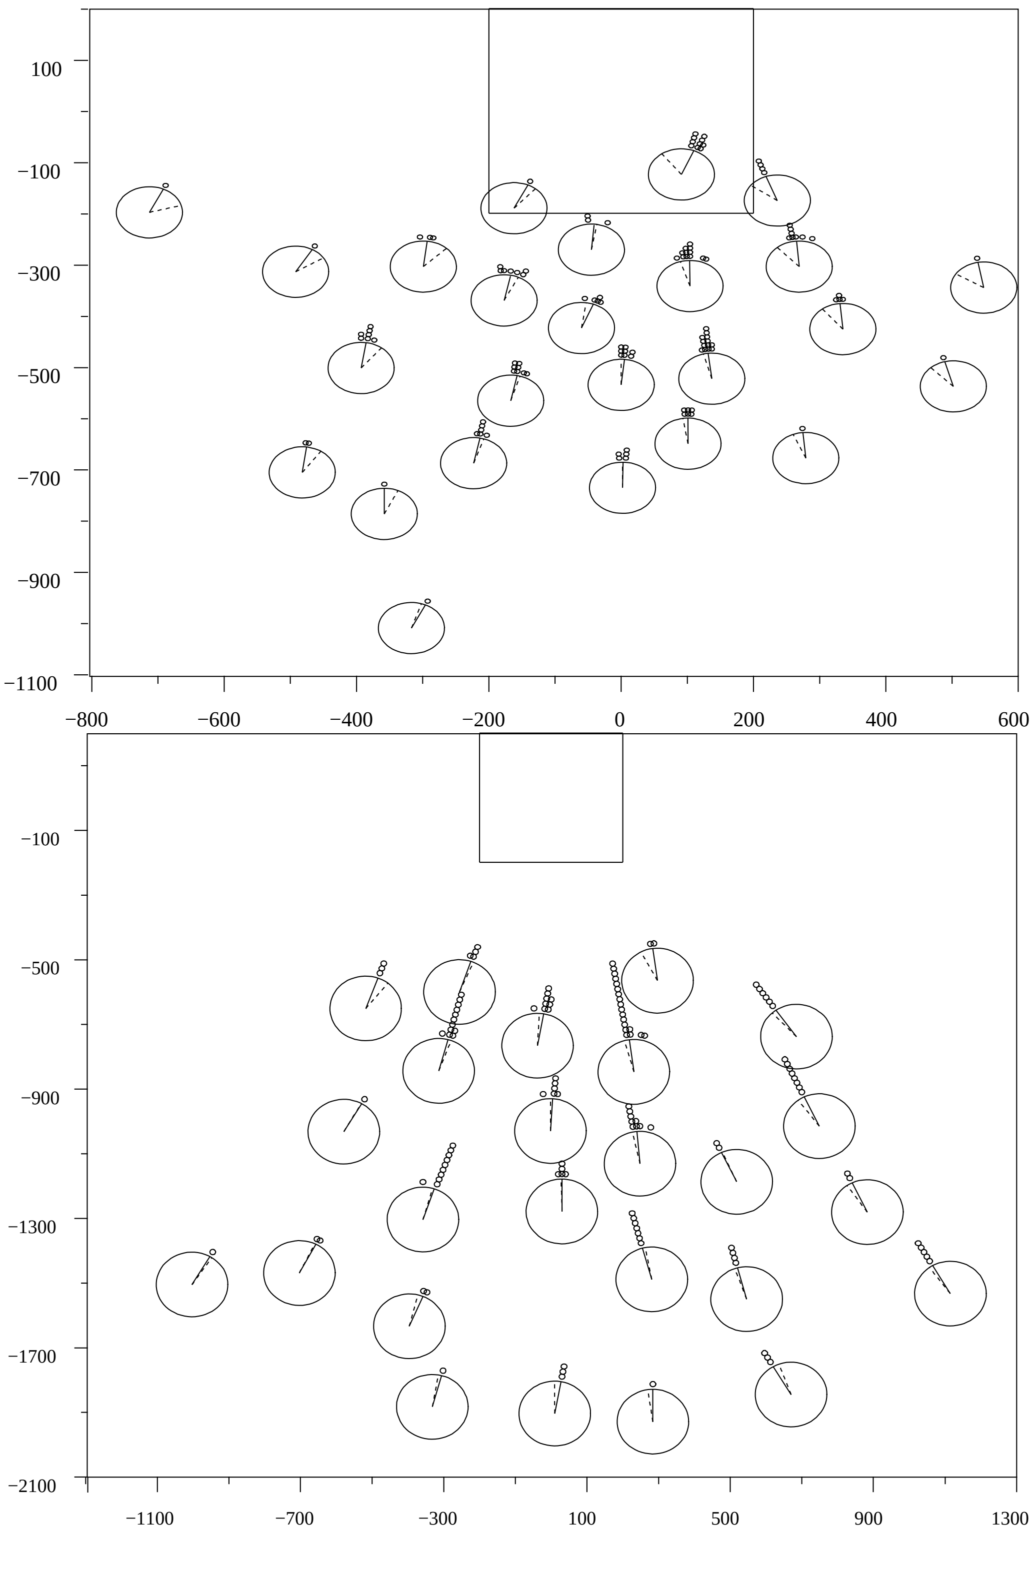

Figure 21: **Short (top) and middle (bottom) range navigation competence generated headings from the continuous trial. Only those headings with confidence values above a confidence threshold of 0.07, are shown (the confidence threshold was 0.03 when the data was captured). The square is the charger capture area. The circles are 5 cm (top) and 10 cm (bottom) radius. The axes are in mm.**

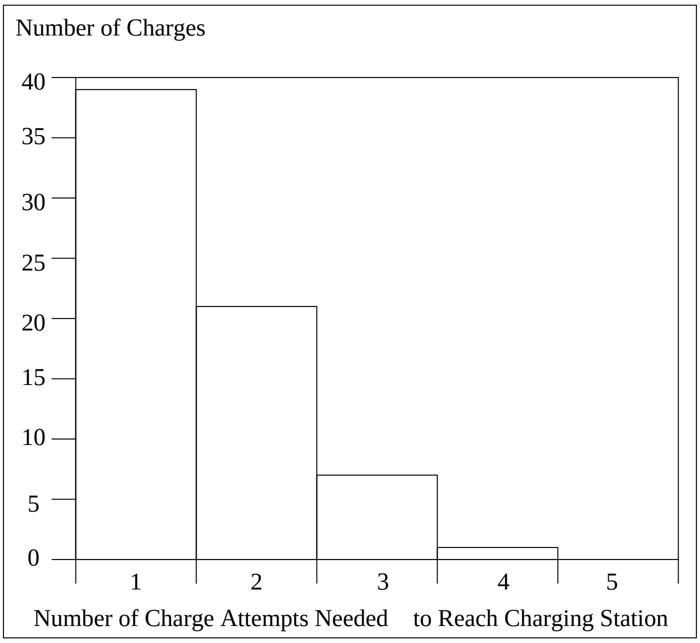

Figure 22: **Charge Attempts: The number of succesful charges plotted against the number of attempts necessary to reach the charging station after each occassion when the battery level was below the re-charge threshold.**

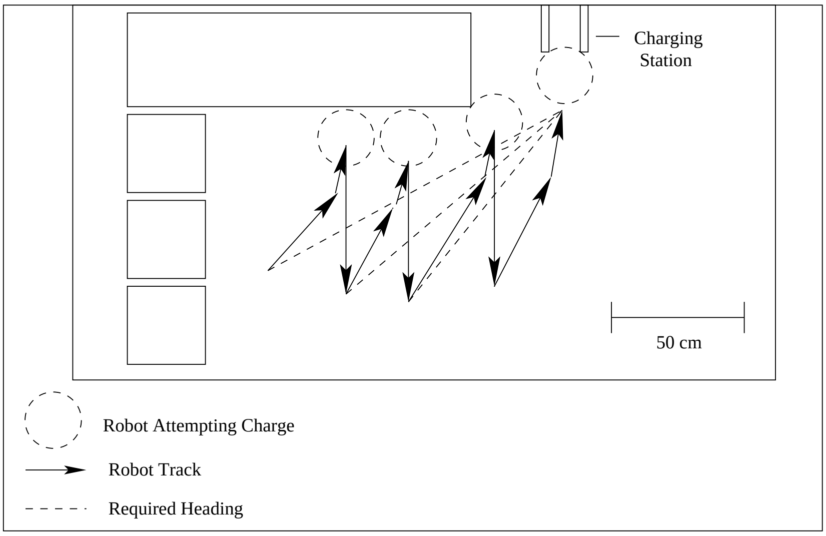

Figure 23: **Simplified representation of robot behaviour when the initial connection attempt has failed**

No of Charges	Mean	Median	sd	Fastest Time	Slowest Time
69	99.8	97	52.7	40	247

Table 14: **On a charge attempt, the time from first homing vector with a confidence value above the confidence threshold to charge (units are in seconds)**

The 4 minute maximum time is much greater than average and occurred in only two instances, which can therefore be considered to be 'outliers'. These findings suggest that when the robot is operating continually, without 'mock' discharge, and once the robot is within range of the middle and short range navigation competences operational area, homing to the charging station will occupy very little of the robot's operational time.

6 Discussion

In order to expand the use of robots beyond the factory floor there is a need for autonomous robots which can operate continuously in unmodified environments. The staged competence acquisition model, loosely based upon animal life cycles, offers a possible approach, albeit on a very small scale, to provide such a system. Complex functionality can be decomposed into simpler competences which can be learned at different stages in the robot's operation, so that at some point the complex overall functionality is acquired.

The experiments presented here investigated how a functionality can be acquired with competences whose staging during learning and utility overlap. A self charging robot wandered within an experimental environment until its batteries were low, homed to a charging station, charged and then resumed wandering. This was achieved by using combinations of three independent competences, short range navigation, middle range navigation and obstacle avoidance. These competences were acquired by the robot using machine learning techniques based on neural networks.

The operational time of the robot was used as a variable upon which each competence based a tailored activation function. The output of the activation function was used to determine when and at what rate learning occurred.

Two types of trial were conducted. The first type, the repeatability trials, sought to provide some measure of how replicable the results of the navigation competences were. The second type was the continuous trial, in which the robot performed the wander-home-charge cycle over an extended period of time. The aim of this trial was to measure the robustness of the mechanism and the survivability of the robot.

During the repeatability and continuous trial data was collected with which to asses the robot's performance. A particularly important measure was the difference between navigation competence generated headings and required headings. Several statistical and diagrammatic techniques were used to analyse the differences in heading data. Analysis showed that for both navigation competences, as expected, there was a gradual deterioration in accuracy of the competence-generated headings towards the 'edges' and 'point' of the competence's areas of operation (see figure 18, 12, 13 and 14). The deteriation was more pronounced for the short range navigation competence. In figures 12 and 13 (top diagrams) the homing vector headings seem to be almost perpendicular to the station, with little deflection towards the required heading. A re-evaluation of the short range navigation competence's Radial Basis Function parameters may be beneficial. The deterioration towards the edges also led to an investigation into the level of the confidence threshold, i.e. the threshold at which a homing vector confidence value is considered to be high enough for use. This led to an adjustment in the value which will be used in later experiments.

Occasionally, (see figure 14, top diagram, for the most extreme case), a competence showed a heading error across the area of operation. This was reflected in the means of the heading differences (tables 1 and 8). These errors were surprisingly uniform. If the homing vector headings were rotated back by the mean the 'fault' was corrected. This error may be due to initial alignment using sonar balancing or some other effect of the interaction between hardware and software. While the reason for these errors will be investigated further, they are not of an order which would threaten the operational survival of the robot.

When the competence heading differences were compared across type, i.e. short range with middle, the Radial Basis Function-based short range navigation competence proved to be less accurate than the Self Organising Feature Map-based middle range navigation competence. The standard deviations were 21.8 for the RBF (short range) and 13.0 for the SOFM (middle range). This was a surprise since earlier tests (see [8]) had shown that the RBF was far more accurate. It is possible that the SOFM was not functioning correctly during the earlier experiments, since there were problems with the equipment, hence the reversal in accuracy of the two types of network is due to an improvement in the performance of the SOFM (middle range navigation), rather than a deterioration in the performance of the RBF. The fact that the RBF is less accurate does not invalidate its use. As used in the trials, the RBF does train much more quickly than the SOFM and the use of the RBF is likely to be extended and methods to enhance its performance explored.

The data collected during the continuous trial for homing vectors based on the fusion of short and middle range competence values yielded some very interesting results (see table 8, 'fused' row). The fused results were close to the middle range navigation competence in accuracy, far better than the short range navigation competence. Since the fusion of values occurred deep into the short range navigation competence area of operation, the fusion clearly helped to increase the accuracy of the robot homing.

The charge attempts were then reviewed. The functionality derived from the interaction of the competences produced a high level of reliability when attempting to charge (figure 22). The maximum number of attempts to reach the charging station was 4 (and this occurred only twice, the median value being much lower at only one attempt) and the maximum length of time 4 minutes (see table 14). When the robot failed to make contact with the charging station, the next attempt was invariably closer to the re-charge station (see figure 23 for a diagrammatic representation of behaviour on failure). At no time did the 'mock' battery level drop to 20, a level which would have signalled the end of the test. These findings were extremely pleasing since they suggest that the functionality derived from the competences would enable the robot to operate for long periods without supervision.

Overall, analysis of the data showed a robust system capable of functioning with a considerable amount of variance. In summary:

- Homing vectors generated by the short range navigation competence were less accurate than anticipated and some tests will be conducted to see if the settings are optimum for the usage.
- The middle range navigation competence proved to be much more accurate than expected.
- During the continuous trial, the fusion of homing vectors from the navigation competences greatly increased the accuracy of many homing vectors.
- The obstacle avoidance competence performed well enough to prevent the robot from striking any object (when not charging) and becoming entangled.
- Using sonar balancing to provide a constant external reference point for odometrically based heading meant that the robot never lost its 'sense of direction' through accumulated error.

The staged competence acquisition approach allowed the robot to gain the required level of functionality through a controlled incremental process. The test area was deliberately confined so that there were two distinct periods where the robot first acquired and then used the functionality. In such a situation basing learning activation and level upon a single variable works perfectly well. If the staged competence acquisition approach is to be used in a less confined, less controlled environment then it will need to be reviewed.

7 Conclusions

The continuous trial was very encouraging. The robot trained correctly and displayed the correct functionality. The robot navigated obstacles successfully while wandering. Homing to the charging station was rapid and robust. At no point did the robot halt through lack of 'power'.

Decomposing a function into several simpler competences and using self learning techniques to acquire the competences in the correct spatial and temporal location certainly possess merit. A robot operating continuously over long periods of time, and learning under such conditions, poses unique problems that are not apparent under "condensed" learning scenarios. The amount of "irrelevant" data (irrelevant with respect to the particular competence being learned at that point in time) will be high, and the effect of noise and contradictory sensory perceptions will be more pronounced than under condensed learning situations. The staged competence acquisition approach has been effective in ensuring that a competence is not exposed to much "irrelevant" data by controlling when learning is active.

Although this approach provided the robot with a robust and repeatable functionality it is still only a step towards providing a system which can truly be called continual. At present the robot can only learn how to navigate over short and middle range distances. There is no provision for longer range mapping. Additionally if changes were made to the experimental environment during a continual trial the robot would not be able to compensate for the change, since there is no mechanism which would allow re-learning. Subsequent sections of this report present a proposal, based upon concepts and measures which arose from this work reported here to explore the issues of both extended exploration and timeliness of knowledge. It is anticipated that the proposed work will give rise to an operating model allowing a self charging robot to function over long periods of time in a large area, in which change can occur. During development, problems due to continuous operation will be encountered. The problems will be investigated and the model adjusted to remove, reduce or utilise the effect of the problems. The development of the model and the investigation of the problems of operating continually will therefore become an iterative process.

8 Future Work

8.1 Issues Arising from current experiments

- The Radial Basis Function Network (RBF), used by the short range navigation competence has been shown to be consistently less accurate than the SOFM, used by the middle range navigation competence (see tables 1 and 14). To ensure that the parameters of the RBF are reasonably close to optimum, a series of tests will be conducted on the robot, varying these parameters. The tests will take the same form as the short range navigation component of the repeatability trial.
- The Radial Basis Function Network, used by the short range navigation competence trains much faster than the SOFM based middle range navigation competence. It would be useful to extend the use of the RBF. Tests could be conducted to see if the RBF can be used to fulfill the role of the middle range navigation competence by increasing the accuracy and range of opereation of the RBF.
- With the capture of large amounts of data during the continuous trial a re-evaluation of the confidence threshold was made. The new value, 0.07, needs to be tested and compared with the previous threshold, 0.03, to see if there is any positive advantage to raising the threshold. This experiment will probably take the form of a shortened version of the continuous trials.

9 A Proposal for Developing a Model to Explore the Problems with Continual Operation

9.1 Motivation

Most experiments with machine learning mechanisms for robot operation occur in a laboratory with strict spatial and temporal limits. The robot is only exposed to data of relevance to the operation being learned or performed. Using such techniques in 'real world' environments, which may not have such limits, is problematic. When attempting to learn, the robot is supplied with much "irrelevant" data, i.e. from the wrong place or the wrong time. Additionally the environment may be prone to change and therefore re-learning also becomes neccesary.

9.2 Purpose

The following sections propose development of a learning model which can be used to explore the problems associated with continual operation. It is intended that an iterative process will be used during development. When a problem occurs which is created by some aspect of continual operation the model will be re-evaluated and modified to eliminate, reduce, or employ the problem's effect. The model will be based on the staged competence acquisition technique (see the preceding sections and [8], which has yielded good results in continual operation, but in a constrained environment.

9.3 Method

9.3.1 Staged Competence Acquisition Revisited

The staged competence acquisition model provided a robot with a certain functionality by first decomposing that functionality into simpler competences and then controlling the acquisition of those competences using the operational time of the robot. In the constrained environment where these tests were conducted this model produced a robust system.

In order to extend this work modifications will be needed. The robot was required to: **wander within an experimental area; when low on batteries home to a charging station; charge and then resume wandering**.

The robot this gained this functionality by acquiring three competences:

- Short range navigation to the immediate area of the charging station.
- Middle range navigation to the short range navigation operational area.
- Obstacle Avoidance, to avoid objects within the environment when wandering.

All three competences used sonar sensors to provide readings from the environment and neural networks to associate these readings with data which could be used to control robot movements.

In the proposed project, the robot will be operating in an environment with no guaranteed limits. The previous robot functionality may, therefore, never be reached. The proposed functionality is to: **explore and map an experimental area; when low on batteries home to a charging station; charge and then resume mapping**.

The functionality from the previous project now describes the robot behaviour when, or if, mapping is completed.

Note that the new functionality is available while learning continues, not when learning is finished. At any time during exploring when charge is low the robot is required to return to the charging station and after charging return to mapping.

The new functionality could be broken down into the following competences:

- Short range navigation to the immediate area of the charging station.
- Middle range navigation to the short range navigation operational area.
- Long range navigation towards the middle range navigation operational area.
- (Obstacle Avoidance, to avoid objects within the environment when wandering.)

9.3.2 Obstacle Avoidance: A Required Competence?

Obstacle avoidance appears in brackets on the competence list above because its role to the new functionality is less well apparent. In the previous experimental setup obstacle avoidance is learned after mapping has occurred, this works because the area is deliberately limited. If mapping is ongoing when can Obstacle Avoidance be learned? Presuming that Obstacle Avoidance is still required, one possible mechanism is to pause mapping after the short and middle range navigation competence have been acquired and use vector homing to repeatedly move the robot close to the charging station (and wall), Obstacle Avoidance can then be learnt as the robot avoids the charging station (and wall). Thereafter vector homing can be used to return the robot to the charging station and mapping continued. However obstacle avoidance is not really needed until mapping is finished. Initially it will be assumed that Obstacle Avoidance will be learnt once all other learning is finished, i.e. the area is fully explored. This may, however require revision according to the findings of future experiments with the proposed functionality.

9.3.3 Navigation Competences

Short and middle range navigation can be based upon the same mechanism as the previous experiment, this produced robust results and there is no reason why it should not be used here. The long range navigation competence is different. Since the area to be explored may change, operational time can no longer control when, and to what extent, learning occurs. The long range navigation mechanism, therefore, needs to be 'open ended', it needs some internal means of assessing when the area is explored and, as a problem for the future, when it needs to be re-explored.

One approach, well researched at Manchester University (CITES), using numerous mechanisms is to provide the robot with route creation and following mechanisms based on machine learning techniques using neural networks.

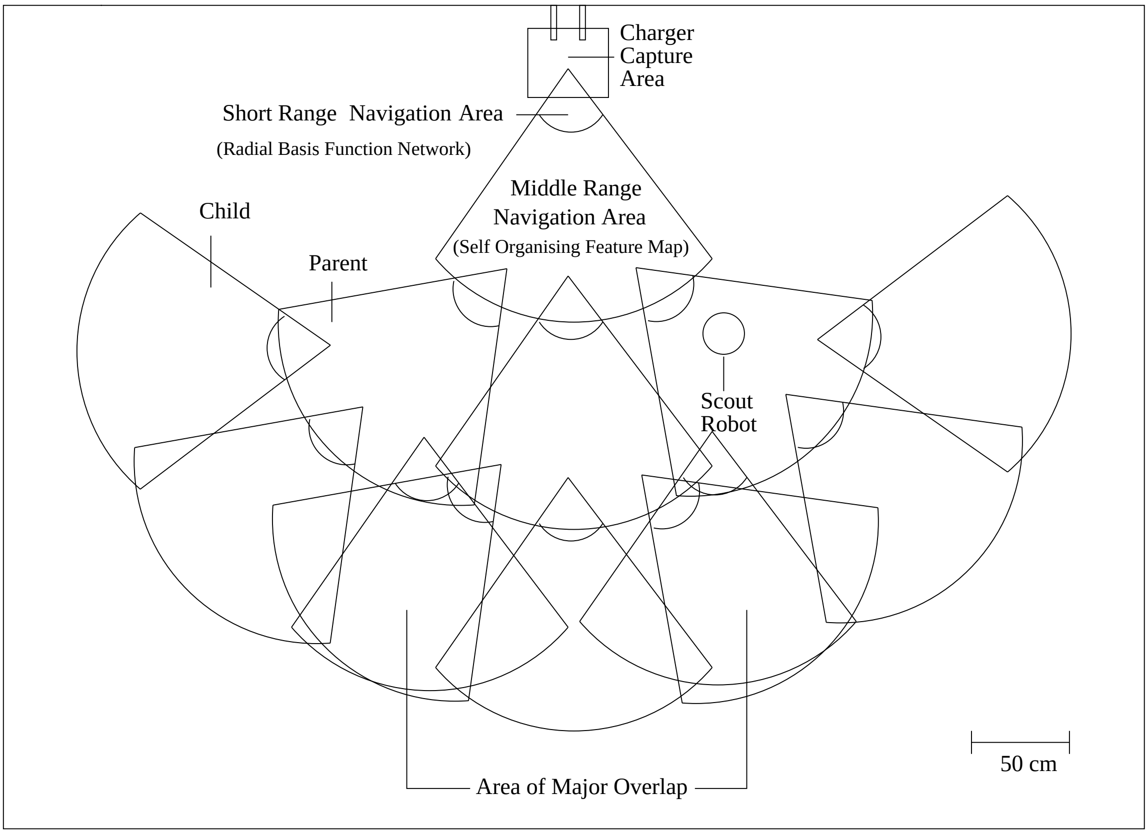

Figure 24: **The proposed long range navigation competence: Effectively the short and middle range navigation competence repeated iteratively. Each epoch of the long range navigation competence maps a fan shaped area. Each fan shaped area becomes the parent of three child areas.**

However, the association of heading vectors with sensor readings though fan like patterns has been shown to provide a robust method for local homing. So another approach is to use the vector mechanism developed for short and middle range navigation iteratively (see figure 24). This means that the long range navigation competence would consist of a number of discrete epochs, each of which is equivalent to the short and middle range navigation competence.

The fan shaped middle range navigation operational area will become the first *parent* area. Near to the edge of each *parent* area are the focal points of three *child* areas begin. Each *child* area is a 'repeat' of the short and middle range navigation competence. The fan shaped learning pattern is thereafter repeated, each *child* becoming a *parent* to further areas, extending the mapped area to the edges of the experimental area. One such possible configuration is shown in figure 24.

Initial experiments will use the navigation competence configurations described in this paper. A separate RBF and SOFM network will be used for each fan shaped area.

An hierarchical linked list of the networks used by the navigation competences will be generated, a trinary tree, each node connected to three lower nodes (see figure 25).

At the top of the tree will be the RBF and SOFM used by the middle and short range navigation competence. At any time one node in the tree will be the active node. This node will point to an RBF and SOFM and these will become the 'active' networks. If a homing vector is required, the 'active' RBF

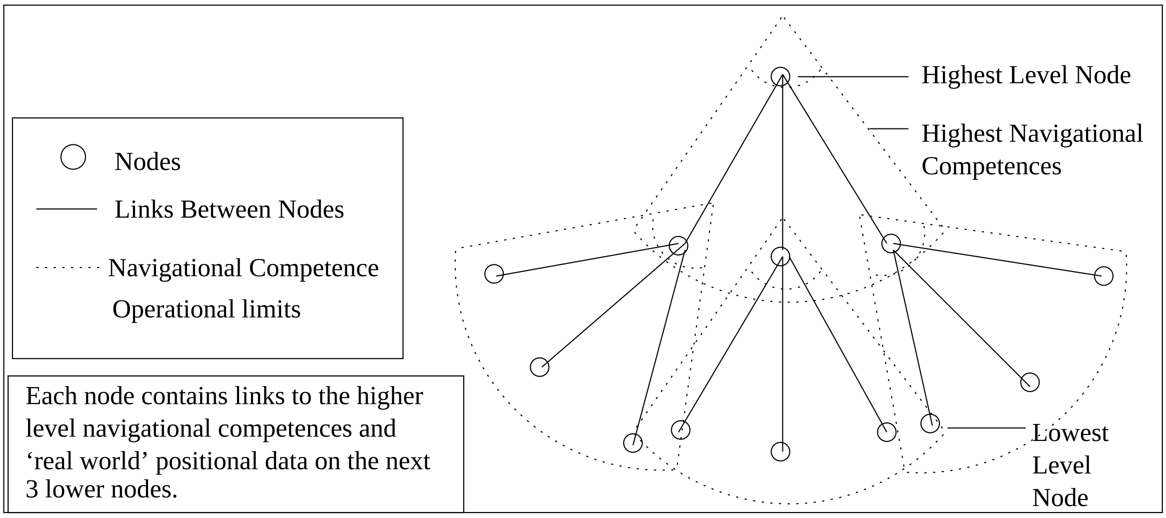

Figure 25: **The hierarchical trinary tree: a list used to link the short, middle and iterative long range navigation competences**

and SOFM will be used to generate it. Initially the active node will be the 'root node'. This node is a special case and points at the connect-to-charger function as well as the short (RBF) and middle (SOFM) range navigation competence networks. In the earlier experiments, the translation distance generated by the RBF (short range navigation) was used to assess when the robot was close to the focus of the fan, i.e. the charger capture area, and then to pass control to the connect-to-charger function. The same mechanism will be used, but (except for the short and middle range competences) closeness to the focus of the fan will cause the parent of the current active node to become the active node. The RBF and SOFM networks linked to the new active node i.e at the next level, will then up be used when generating homing vectors (see figure 26).

This should provide an adequate mechanism for homing to the charging station. However, after charging the robot then needs to return to its mapping duties, and it would be helpful if it did not need to restart mapping from the charging station after each charge. In order to provide a mechanism which will allow the robot to return to its previous mapping location each node will also contain a list of the headings and distance from the 'real world' position of that node to each of the *child* nodes. When homing to the charging station, the robot would build up a list of third child, second child, second child, etc. The robot would simply reverse this list after charging to return to it's previous mapping position. A sonar reading will be recorded at each nodes position. This is included so that the robot can perform some form of checking or course error correction using the reading as the hidden values of a neural network node structure, such as the Restricted Coulomb Energy network. The precise details of these operations are not yet defined.

9.3.4 Controlling Learning

In the Staged Competence Acquisition model, operational time was used as the single variable for controlling when a network was learning. Now, however, the size of the environment will control when learning should occur, i.e. more fan shaped patterns are added until the edges of the environment are encountered. Initially time will still be used as a single factor. The short and middle range competences will be learnt with the activation functions presented in sections 2.7.1 and 2.7.2. Since the long range navigation competence is merely the iterative application of the short and middle range competences the competences activation curves will be used, but in a relative fashion, i.e. not at some absolute time in the robot's operational 'life' but when mapping of a fan shaped pattern begins. At the present time, it will be assumed that the robot can only react to low battery level, and hence begin homing behaviour, between mapping patterns. The activation curves for this adjusted learning control are presented in figure 27. This shows the short and middle range navigation competences and two 'epochs' of long range navigation

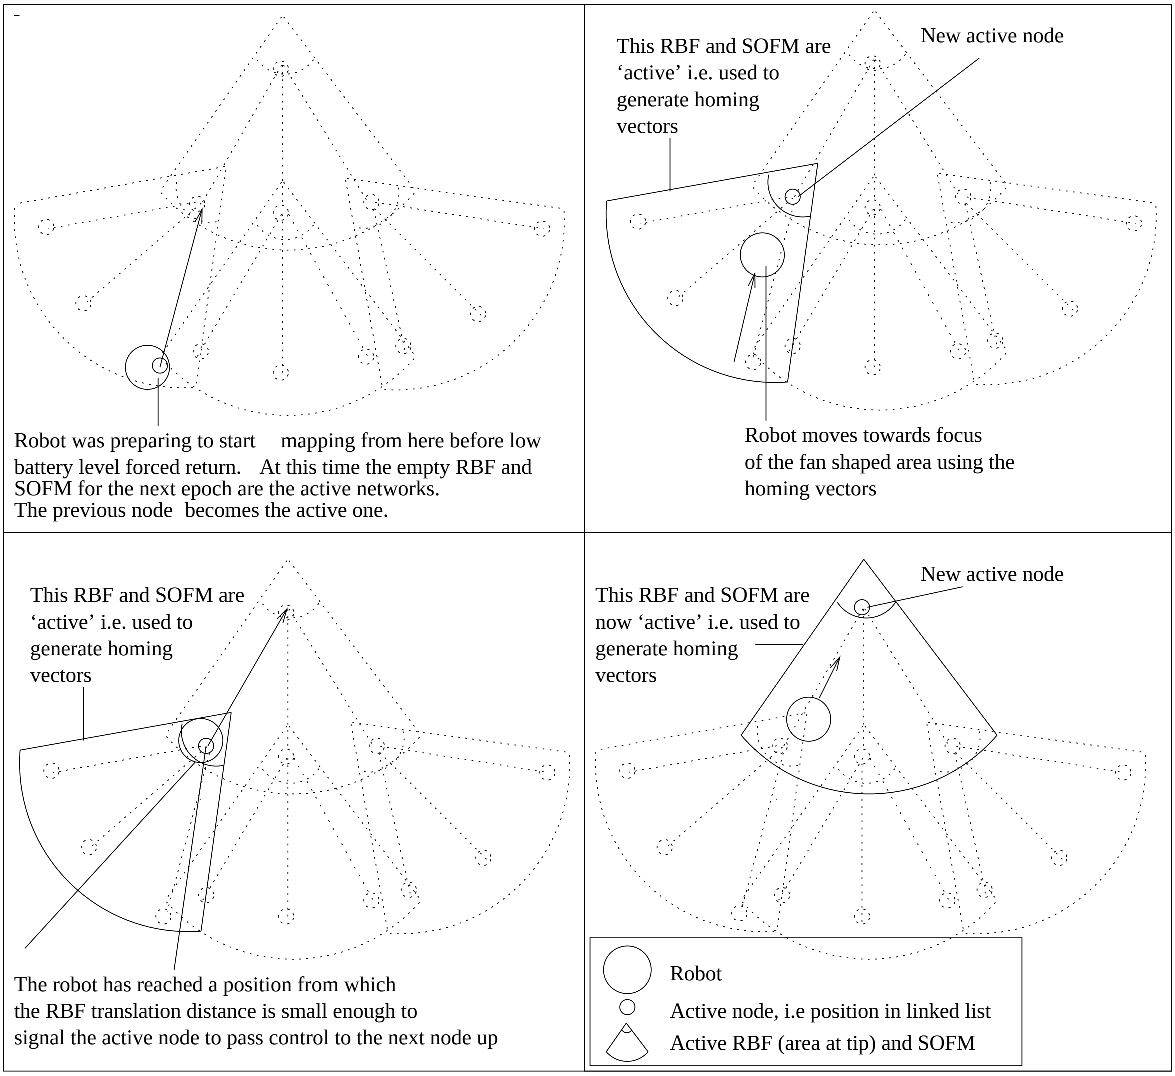

Figure 26: **Homing using the hierarchical trinary tree. The active node selects which RBF and SOFM are to be used for homing vector generation. When the RBF generates a translation value below a 'closeness' threshold control is passed to the next node up the trinary tree**

learning.

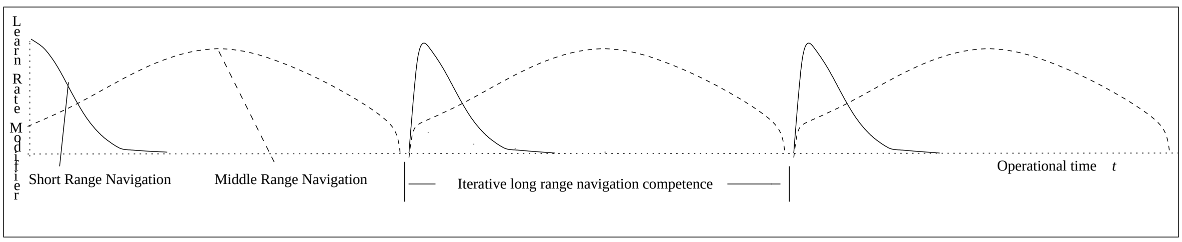

Figure 27: **Competence Staging for the Robot Strange: Short and middle range navigation learning are presented along with two epochs of long range navigation learning. The long range navigation epochs would be repeated until the area is fully mapped**

9.3.5 Repeating the Long Range Navigation Competence

Particularly when using sonar balancing to correct heading (i.e. without the magnetic compass), the robot is quite reliant on odometry. The long range navigation competence is actually a repeat of short and middle range navigation competence learning. The middle range component of this mapping technique (SOFM based) takes about 20 minutes to train and there is always some amount of odometric drift by the end. If this pattern is to be repeated many times then the odometric error will become quite large. One way of avoiding this is to extend the use of the fast training RBF from the short range component to the middle scale and possibly merge the short and middle range components of the iterative procedure used by the long range competence. Some experimentation may be carried out with the RBF prior to the initial experiment (see section 9.4).

9.3.6 Initial Parameter Settings

Since this is a system which relies to some extent on odometry, there are several operational parameters which need to be set and revised after experimentation. As the project progresses it may be possible to make many of these self setting based upon sensor readings.

- In order to provide a definite *link* between *parent* and *child* how 'deep' into the parent fan does the focus point of the child fan need to be? The greater the overlap, the higher the confidence level in the parent, but so is the number of areas required and therefore the training time and memory requirements.
- How much redundancy, overlap at the same level, is desirable or useful? The simple configuration figure 24 has many overlaps, including two areas of almost complete overlap. The same training time and memory considerations as above apply.
- Should the size of all fans remain the same or should the robot use its sensors to decide whether to increase/decrease the size of the fan depending upon the position within the environment (both the radius and the arc of the fan could be adjusted)? If the sizes are adjusted then would a network model such as the Grow When Required network (CITES) be more appropriate than the Self Organising Feature Map?
- Should the number of *child* areas remain the same or should the robot use its sensors to decide whether to increase/decrease the number depending upon it's position within the environment?
- Most importantly will different shapes and sizes of area require specific configurations to optimise mapping? If this is the case then having some self determination in choice of number of children, width of arc or length of area becomes even more important.

9.3.7 Initial Issues

Apart from optimising parameters there are also mechanistic issues:

- The long range navigation competence will be composed of the iterative application of the short and middle range navigation competences in an hierarchical structure (initially a trinary tree). Will the structure be constructed using a breadth first strategy or a depth first strategy? This has a profound effect on how the robot will map the area. A breadth first strategy produces all the possible epochs at one level before producing epochs at the next. Depth first produces all possible descendants of node before producing the next node at the same level (see figure 28). While operation in a corridor environment would be fulfilled best with depth first, but operation in an open area could be better with a breadth first strategy. The robot will use the breadth first strategy in the initial test since the test will be conducted in an 'open area'. A future consideration is whether the robot could be enabled to decide which strategy to use, based upon sensory input.
- As an alternative to using the trinary tree, and having only one pair of networks active, the sensor input could be presented to all of the networks and the vectors from those networks with the highest confidence values used. The major problem with this is that networks from many different places within the environment would be tested and this could increase the chance of perceptual aliasing.

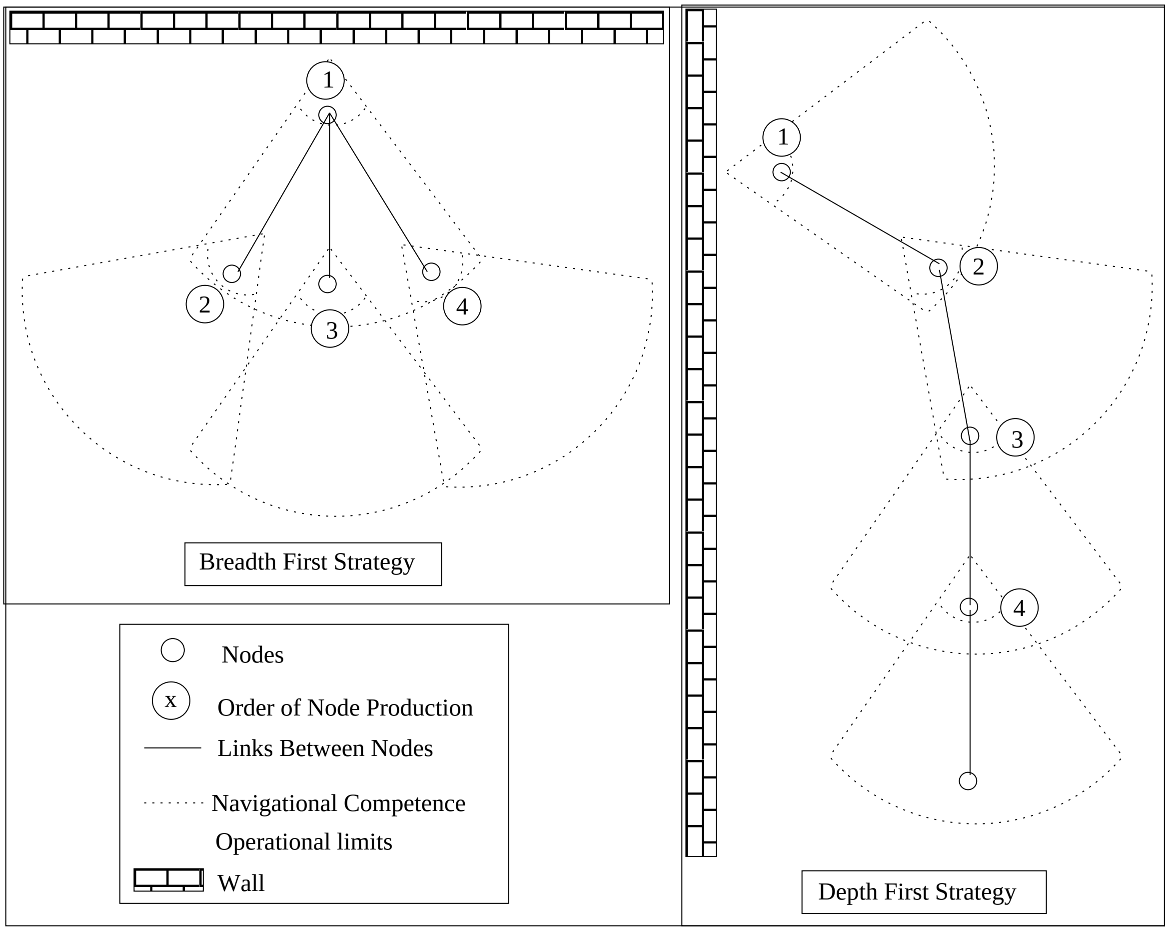

Figure 28: **Breadth and Depth Strategies and how their use in the long range navigation competence would affect mapping**

Even if not of general use, such a mechanism could be useful as an element in a search procedure if the robot becomes lost, or even used as a 'vote' alongside a value generated by using the trinary tree.

- With the greatly increased number of vector movements, will the odometry of the robot remain sufficiently accurate? How many fan patterns can the robot generate before it needs to return to the charging station to correct the heading? If the heading can be adequately corrected, or the magnetic compass is used, will translational errors become significant? Is some form of sonar balancing, or positioning, to correct translation possible (this would probably require the identification of some unique feature in the landscape; the charging station is not at present such a feature). If odometry is a problem does the SOFM need to be trained more rapidly (repeating the fan shaped pattern less) or even dropped in favour of the fast training RBF?
- Could another network, such as the Grow When Required (GWR) or the Restricted Coulomb Energy, be used alongside the other networks, in order to increase the accuracy of the system by allowing 'voting' from a network supplying a different perceptual perspective. In such a configuration there would be a single GWR network growing over all the mapped areas.
- Can the system cope with confined areas, such as a corridor leading off the main learning area? At present the robot uses an emergency braking routine to stop the if it moves too close to a wall while learning; further sensor reading points along that vector are truncated. If there is a corridor connected to the 'main learning area' the current mechanism should be capable of traversing and

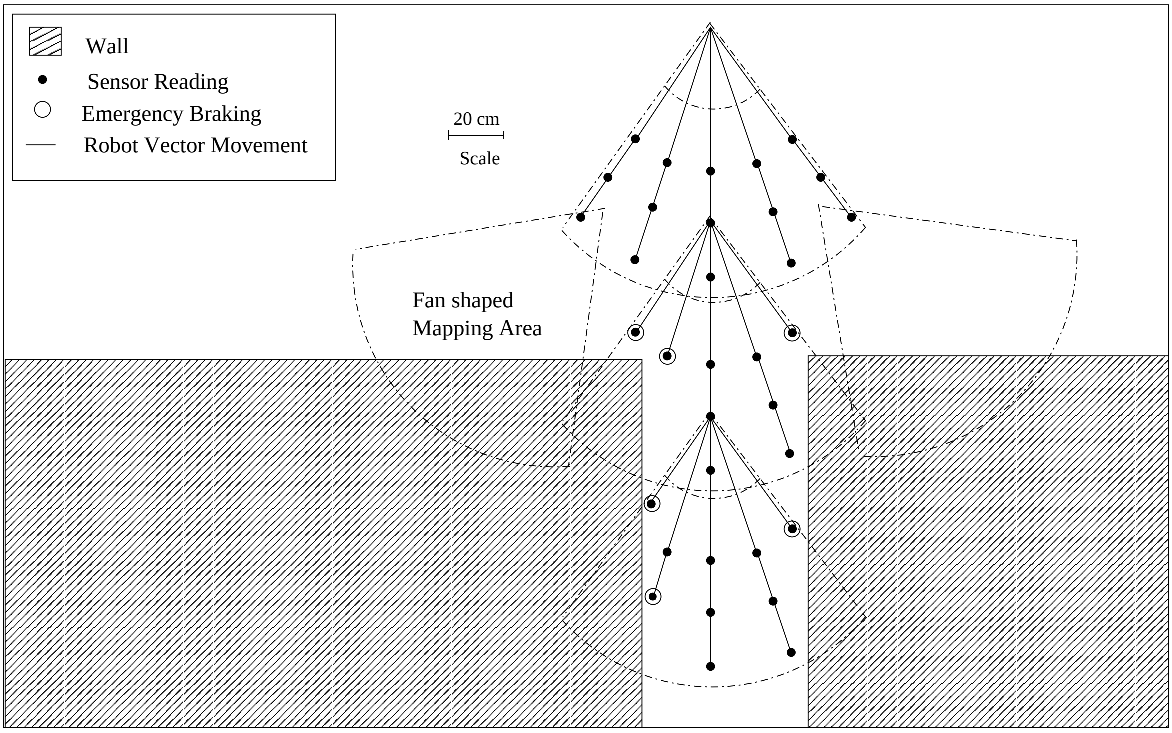

Figure 29: **Mapping Corridors with the Fan Shaped Vector Movement Pattern**

mapping it as graphically represented in figure 29. Once the basic operation of the proposed model is proved to work this will require testing by 'constructing' (or using real) corridors.

9.4 Initial Experiment

Prior to the initial experiment there will be an investigation into extending the use of the Radial Basis Function network. If it is discovered that the RBF can produce accurate enough results over the middle range navigation operational area it may well be used instead of the slow training Self Organising Feature Map. However, since the result of these tests is at present unknown the proposed initial experiment is discussed with the presumption of both networks being present.

9.4.1 Experimental Hardware

The modified self charging Nomad scout *Strange* and the custom built charging station will once more be used. A laptop will provide high level control.

9.4.2 Experimental Setup

The experimental area will no longer be constrained to that of little more than the short and middle range navigation competences. Although the area will not be large enough to use the idealised pattern shown in figure 24, an area large enough for at least one long range navigation epoch (hopefully two) will be used. Figure 30 shows a likely configuration. While not ideal, this set-up can be used to test the basic functionality of the long range navigation competence mechanisms.

9.4.3 Experimental Procedure

The purpose of this initial experiment is to see if the mechanisms developed for the long range navigation competence mechanism can be integrated into a model alongside the short and middle range navigation competence. The success of the experiment (and the model) will be based upon the robots ability to first

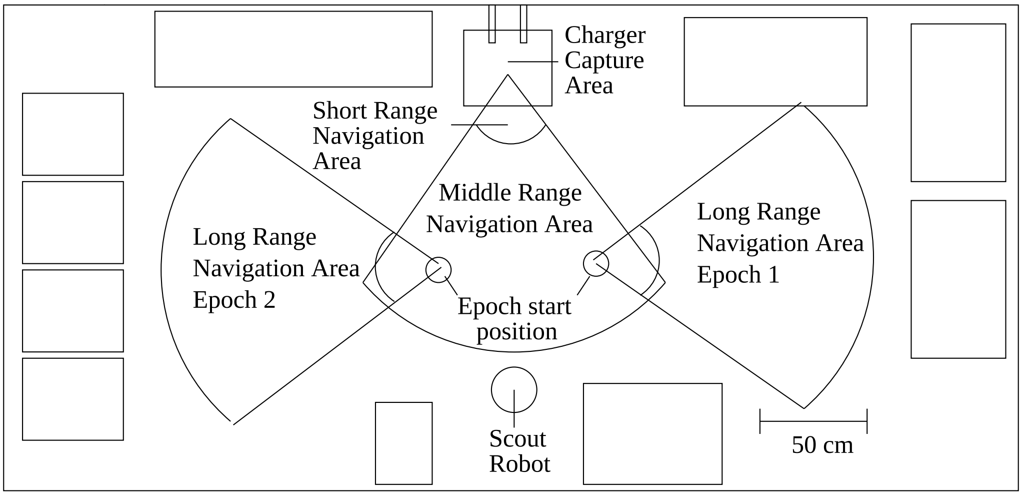

Figure 30: **Possible experimental area configuration for the first experiment**

map the area and then navigate back to the charging station from various points within the experimental area. At the start of the experiment, the robot will be placed 45 cm from the charging station, facing the station, aligned so that a forward movement will result in connection between the robot and the station.

The first node of the tree (in this case a binary) is established pointing at the connect-to-charger function and the empty networks for short and medium range navigation learning.

The robot initially moves away from the charging station, reversing directly without turning, and then moves back towards the station in the fan-like pattern of ever increasing size (described in section 2.4).

If an obstacle is detected in the robot's path, the robot will merely stop and attempt the next movement in the fan shaped pattern. If the new movement is in the same direction as the previous attempt, this movement will be rejected and the next one attempted.

During this first phase, short range navigation will be learnt, controlled by the activation function λ_r. Middle range navigation learning will also begin at the same time, controlled by the sensitisation function λ_s. This phase continues until the radius of the fan is approximately 50 cm (the short range navigation area in figure 30) and lasts for about two minutes.

At the next stage of learning, the robot moves in and out of the charging station at much greater distances (described in section 2.5). Middle range navigation continues to be learnt, with a higher learn rate than in the previous stage. This phase lasts about twenty minutes. The robot's movements during both learning phases are shown in figure 3.

The robot returns to the initial start point. The robot creates a new node for the tree and adds a link from the active node to this new node. The new node then has pointers made to long range navigation competence epoch one's empty RBF and SOFM. The robot then moves out to the long range navigation epoch one start position. The absolute heading to return to the previous position is stored in the new node along with the translational distance moved. The robot now maps the epoch one area by using the procedure described for the short range (epoch one's RBF) and middle range (epoch one's SOFM) navigation learning.

Once this learning phase is over and the robot is once more at epoch one's start position, the robot returns to it's position at the start of the experiment. This could be done using the short and middle range competence to produce homing vectors, but initially will use the absolute heading and translation data stored at the node associated with epoch one.

Once the robot has returned to the original starting position, a new node will be created and linked with the root node. This will be the node for epoch two. The procedure described above for epoch one will then be repeated for epoch two.

At the finish of the long range navigation learning (i.e. the end of epoch two in this experiment), the robot will return to the initial start position. There will be no obstacle avoidance learning at this time. In this test the robot will not be required to wander freely.

The testing and data collection phase will then begin. The robot will be moved out under joystick control and stopped within the short, middle, long epoch one or long epoch two navigation operational area. The robot will be then be given a number, 1-3 where 1 is short or middle range, 2 is long epoch one and 3 is long epoch 2. This will robot which area it is in and which RBF and SOFM networks to use. The robot will then be told to start homing. Homing data will be collected as in the continuous trial. Once the charging station has been reached, or the robot has quite clearly failed, the robot will be guided out under joystick control and the homing restarted. The initial experiment will test the soundness of the basic idea and the functionality of the various mechanisms.

9.5 Temporal Problems: The Timeliness of Data

The initial experiment described above seeks to develop a mechanism to investigate the spatial problems of continuous operation. When operating in an 'open environment' there are also problems with the timeliness of data. To operate continuously the robot requires mechanisms to allow it to determine when a learned competence is no longer valid because of changes in the environment, the robot, or the robot's interaction with the environment. A machine learning technique provides the robot with information or an ability which allows the robot to perform a new function. If this new function is repeatedly executed on the robot there should be some aspect of the robot's performance which can be measured and later used as a gauge to assess how successfully the function is being applied.

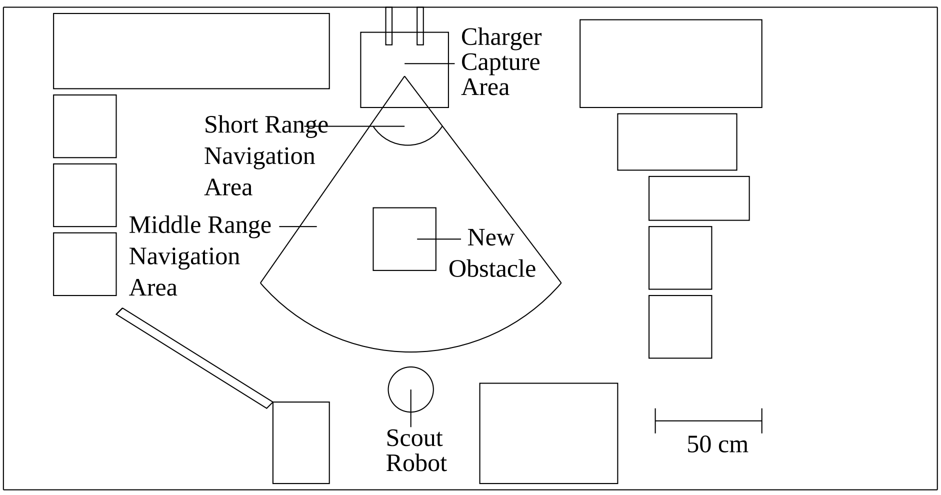

Figure 31: **Introducing an Obstacle into the Middle Range Navigation Area after learning**

As an example, when an obstacle is placed within the middle range navigation area, the middle range navigation competence will no longer be capable of providing accurate homing information. However, as long as the change in the environment is only partial, the high dimensionality of the sonar readings (sixteen sonars) will probably allow the robot to reach the charging station. Initial tests have been conducted with the short and middle range navigation competence. After the 'learning phase' was over, i.e. short range navigation competence, middle range navigation competence and obstacle avoidance had been learnt, a box was placed in the centre of the middle range operational area (seen figure 31). Time from first successful homing vector (i.e one with a confidence value above the confidence threshold) to connection with the charging station was over 10 minutes. The longest time this operation took during the continuous trial was 4 minutes (and this was an unusually large value).

The time to connection could, therefore, be used as a indicator of when the environment had changed.

Initial values for the time to connection could be generated immediately after the 'learning phase' by forcing several charge attempts, building up a 'profile' of times from first successful homing vector to connection with the charging station.

The robot could then estimate the time between generating the first successful homing vector to connection with the charging station. If later times were significantly longer than the estimated value then the network would be forced to retrain.

The performance measure could take several forms:

- A simple maximum value;
- A probabilistic measure;
- Using a neural network to provide an association between the values of the first successful homing vector with the length of time to connection with the charging station.

Experiments will be conducted to determine the advantages and disadvantages of each form.

Investigations into means of generating self-regulation for the competences used in the project would run concurrently with the main tests. Integration into the main project would occur when the techniques displayed some measure of success.

9.6 Proposals: Summary

The work reported in this paper, along with [8] has demonstrated that Staged Competence Acquisition is a robust approach to solving problems associated with continual operation of autonomous mobile robots. Section 8 of this paper has outlined some immediate considerations raised by the work conducted thus far. Section 9 has presented a proposal to extend the work, specifically to address problems in continual operation associated with long-range navigation and changing/unpredictable environments. The proposed work will build upon the strengths already demonstrated by the Staged Competence Acquisition model. It will furthermore address the problems identified with components of the model (eg Radial Basis Function accuracy). Finally, since development of the model and associated techniques will be iterative, the proposed work will enable the Staged Competence Acquisition model to be refined and strengthened.

References

[1] E. Batschelet. *Circular Statistics in Biology.* Academic Press, London, 1981.

[2] H. R. Borenstein, J. Everett and L. Feng. *Where Am I: Sensors and Methods for Mobile Robot Positioning.* University of Michigan, 1996.

[3] J. M. Camhi and E. Johnson. High-frequency steering maneuvers mediated by tactile cues: Antennal wall-following in the cockroach. *Journal of Experimental Biology*, 202:631–643, 1999.

[4] E. A. Capaldi, A. D. Smith, J. L. Osborne, S. E. Fahrbach, S. M. Farris, D. R. Reynolds, A. S. Edwards, A. Martin, G. E. Robinson, G. M. Poppy, and J. R. Riley. Ontogeny of orientation flight in the honeybee revealed by harmonic radar. *Nature (letter)*, 403(6769):537–540, 2000.

[5] T. S. Collett. Insect navigation en route to the goal: Multiple strategies for the use of landmarks. *The Journal of Experimental Biology*, 199:227 – 235, 1996.

[6] T. S. Collett. Survey flights in honeybees. *NATURE (News and Views)*, 403(6769):488–489, 2000.

[7] R. Henderson. *A Continuously Operating Mobile Robot.* Manchester University, 1999.

[8] G. Ireland. Staged learning for a continually operating robot. In *Towards Intelligent Mobile Robots '01.* Department of Computer Science, University of Manchester, 2001.

[9] O. Lemon and U. Nehmzow. The scientific status of mobile robotics: Multi-resolution mapbuilding as a case study. *Robotics and Autonomous Systems*, 24(1–2), 1998.

[10] U. Nehmzow. Autonomous acquisition of sensor-motor couplings in robots. Technical Report UMCS-94-11-1, Department of Computer Science, University of Manchester, 1994.

[11] U. Nehmzow. An episodic mapping algorithm for mobile robot self-localisation: "meaning" through self-organisation. In *International workshop on "Computation for Metaphors, Analogy and Agents", Aizu Wakamatsu*, 1998.

[12] U. Nehmzow. *Mobile Robotics: A Practical Introduction.* Springer Verlag, 1999.

[13] U. Nehmzow. Continuous operation and perpetual learning in mobile robots. In *International Workshop on Recent Advances in Mobile Robots.* DeMonfort University Leicester, 2000.

[14] Nomad Technologies. *Nomad Scout User's Manual.* Nomadic Technologies, Inc., 1999.

[15] M. J. L. Orr. *Introduction to Radial Basis Function Networks.* University of Edinburgh, 1996.

[16] C. Owen and U. Nehmzow. Map interpretation in dynamic environments. In *Proceedings of the 8th International Workshop on Advanced Motion Control.* IEEE Press 1998 (ISBN 0-7803-4484-7), 1998.

[17] J. Piaget. *The Origins of Intelligence in Children.* International Universities Press New York, 1952.

[18] R. Wehner and M. Srinivasan. Searching behaviour of desert ants, genus cataglyphis (fomicidae, hymenopoptera). *Journal of Comparative Physiology*, 142:315–338, 1981.