

Incremental building of a model of environment in the context of the McCulloch-Craik's functional architecture for mobile robots.

J. Romo, F. de la Paz, and J. Mira

Department of Artificial Intelligence
UNED-Spain
e-mail: {jromo, delapaz, jmira}@dia.uned.es

Abstract. Current robotics is perhaps the most complete paradigm of applied Artificial Intelligence, since it includes generic tasks involving pluri-sensorial integration and internal representation, as well as motor planning and control. In this paper we revise the architecture proposed by Craik and McCulloch and the concept of environment model introduced by K. Craik. Based on this architecture, which links the description in terms of properties with the selection of a mode of action, we study a simple example application in which an incremental procedure is proposed for the construction and use of a model of a structured medium (the interior of a building) using a graph. The type of graph used to store the descriptions of objects and the relations between them is inspired by the work of Hillier and Hanson on the analysis of interiors. The connections between the elements of the environment (graph nodes) are generated in such a way as to facilitate their efficient use for the selection of the most pertinent mode of action at any given moment. The derivation of the graph is carried out autonomously. In the development of this work, we have avoided as far as possible the use of anthropomorphic terms with no causal connection to the symbol level. Posed in this way, the problem of the representation and use of an environment model by a robot reduces to the use of models of generic tasks and methods at the knowledge level together with graphs and finite state machines at the formal level.

1. Problem Statement.

Robotic systems are probably the most complete paradigm of applied Artificial Intelligence. They include the computational counterparts of most of the cognitive tasks used by Neuroscience to describe the interactive behavior of a living being and its environment. It essentially corresponds to what [Newell & Simon, 72] called an "intelligent agent".

The agent interacts with its environment through a set of sensors which physically represent the spatio-temporal configurations of the external and internal environments. Subsequently, a multisensory processing of higher semantics and with reference to the memory contents is carried out. We call it perception. The objective of both processes is to identify the environment according to a model of internal representation which allows the agent to understand the meaning of these spatio-temporal input configuration in order to coordinate elementary actions so as to navigate avoiding obstacles, etc.

Between these two tasks (perception and motion) there is an intermediate set of decision tasks working always between representational spaces of the internal model

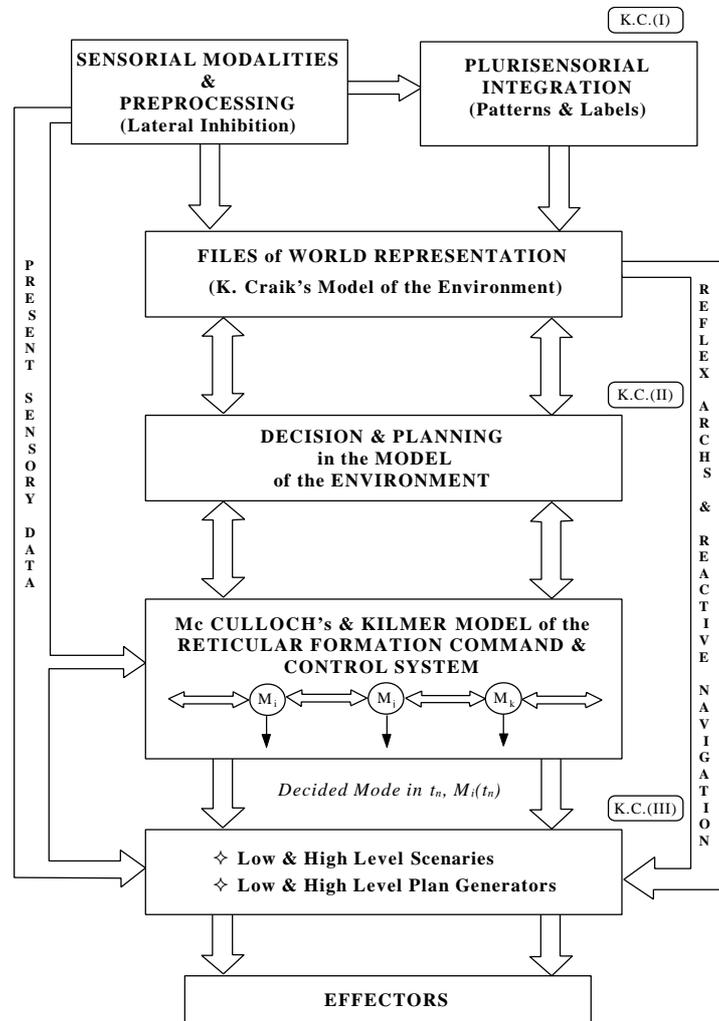


Fig. 1. A simplified version of the structure proposed by W.S. McCulloch for integration of perception and action, based on the Reticular formation in Command and Control system of “modes of behavior”. (Adapted from [Moreno & Mira, 88]. KCI, II, III = K. Craik phases I, II, III.

of the environment but without direct connection with the “real” external environment.

The first part of the ideas underlying this generalized robotic paradigm (the model of the environment), appear in 1943 in the Kenneth Craik’s book “The nature of

explanation”, [Craik, 43], in which, he interprets reasoning in terms of a set of processes geared at building an internal representation of the environment (the model), and using it to predict.

Craik distinguished three essential processes: (I) “Translation” of the external configurations into symbols of the model, (II) Arrival to other symbols by means of a process of inference in the model, (III) “Retraslation” of those new symbols into external processes or “at least recognition of the correspondence between these symbols and external events”.

The computational process of inference using the data structures of the model of the environment, produce the output symbols which are equivalent to what the physical causality modelled would have produced in the real external world. Learning, according to Craik, is to incrementally accumulate “knowledge” by updating the model of the environment.

A second legacy of current robotics systems is due to Warren S. McCulloch and his team at the Instrumentation Laboratory of the Massachusetts Institute of Technology [Sutro et al, 64], where Warren McCulloch proposed a general architecture to integrate perception and action. A simplified version of the original proposal is shown in figure 1, adapted from [Moreno & Mira, 88]. The key point in the proposal of McCulloch is that sensorial data are always treated in conjunction with the use of these data to converge into a “mode of action”, by some type of cooperative decision process in which information constitutes power and the decision time is important for survival. [Kilmer & McCulloch, 69] have listed 17 mutually incompatible modes in the general behavior of vertebrates: sleeping, eating, drinking, fighting, hunting, ... in such a way that an animal is in a mode when its principal focus of attention and the goal of its nervous system is centered on the sensory and motor characteristics of the said mode. Many of these characteristics are also present in other modes but with different global significance or included in another “program”.

This proposal to use a repertory with only a few “modes of behavior” has very concrete counterpart in applied Artificial Intelligence and robotics. As is shown in Figure 1, the selection of a particular mode is performed by the command and control system based mostly on present sensorial information and the status of the system in the model of environment (external and internal). Information concerning the selected mode (M_i) is sent to the sensors, which are then to be tuned to optimize the data acquisition strategy pertinent to this mode. It is also sent to the model of the environment, in where the appropriate data is selected to be sent to the decision and planning module, and to the control of execution of the plan. Finally, there are also direct connections between sensors and effectors which are equivalent to reflex paths [Moreno & Mira, 88].

The initial proposals of Kenneth Craik and Warren S. McCulloch concerning the need to find procedures for constructing an environment model and the subsequent use of this model, both for prediction and to act on the environment by selecting one of only a few *incompatible modes of behaviour*, are still of fundamental interest in robotics. The problem lies in trying to make these general principles of robotics more concrete and come up with specific proposals to go from the knowledge level, in the sense of [Newell, 81], to the symbol level, using clear semantic tables.

In this paper we develop a simple example of incremental and autonomous construction of an environment model, in accordance with the sensorial limitations of

the robot and its pre-defined set of behavioural modes. In the first section, we comment on the legacy of Kenneth Craik and Warren S. McCulloch, developing a functional model of the knowledge level and discussing its connection with the most recent proposals which hybridise reactive techniques with other techniques of a higher semantic level. In the second section, we make explicit the suppositions and limitations of the proposed example. That is to say, we specify the type of world, the sensorial limitations of the robot and the result of the processing of the sensor data by a non-recurrent lateral inhibition network which detects spatial contrasts, distinguishing between “sameness” points and “newness” points. This pre-processing in turn limits the different objects the robot is capable of “perceiving”: only those which can be characterised by different patterns after thresholding the output from the contrast detector, this contrast detection being the function of the lateral inhibition network.

In the next section the different incompatible behaviour modes which the robot is equipped with are specified, these being chosen according to the global function which is to be programmed (navigation in the previously specified world) and the perception of this environment. Following this, we present an incremental procedure for constructing a model of this environment, given the sensorial limitations, and for navigating using this concrete set of incompatible behaviour modes, that is, for the set of restrictions specified in the two preceding sections. The paper ends with some reflections on the representational capacity of the type of graph used to construct the environment model and on the possibilities of extending the proposal which, though seeking inspiration to a certain degree in Biology (through the ideas of Kenneth Craik and Warren S. McCulloch) and in Architecture (through the notation of Hillier and Hanson), strives to take a rigorous computational perspective on the problem, thus avoiding anthropomorphism.

The two problems studied by Craik and McCulloch (the computable representation of the environment and the selection and execution of plans in accordance with this environment model) continue to be central to robotics. The existing spatial representations of the world can be classified into three principal categories: **feature-based representations**, **grid-based representations** and **relational representations**.

Feature-based representations model the world as a set of features provided by the robot sensors (laser, cameras, sonars etc.). These features (generally, segments or regions) are used to determine the free space in which the robot can navigate as well as to estimate its position [González et al, 94].

Grid-based representations are based on a tessellation of the space in which the robot must navigate. [Kaiser et al, 95] incrementally construct a map and cover the grid sequentially, [Moravec&Elfes, 85] use the grid to store the probabilistic information concerning the occupation of each cell, while [Borenstein & Koren, 89] use this type of representation of the world to store many different types of data: procedural, geometric, sensorial, etc.

Finally, *relational representations* are used to try to avoid the accumulation of errors characteristic of the two previous representations, storing the relations between signals and markings of the world rather than storing metric information. These models are generally based on graphs; examples can be found in [Kuipers, 78].

There are two aspects of this classification which we would like to emphasise. First, the majority of the spatial representations classified are only concerned with

navigation. Second, the three categories are not exclusive, that is, the representations used may have characteristics from more than one of them. A clear example of this can be found in the work presented in [Fennema et al, 90], in which a hierarchic, relational and feature-based representation of the world is developed.

As far as navigation is concerned, there are two fundamental tendencies, both included in the initial proposal of McCulloch. The first is *purely reactive navigation* in which the movement of the robot is the direct result of the reading of the sensors at each moment (reflex arcs). The second is *map-based navigation* which typically has large memory requirements since a detailed map of the environment must be stored (extended version of Kenneth Craik's model).

These two tendencies are rarely used separately. Thus, [Brooks, 86] uses a basic reactive model together with higher order models which include map construction. Arkin [Arkin, 94] uses the repulsive potential theory, a map of the world known *a priori* and a series of pre-established motor schemas. Other techniques use *incremental map construction* starting from a pre-established grid, either covering it sequentially or assigning to each cell a probability of occupation and then joining equiprobable regions [Kaiser et al, 95].

Though there are authors who use purely reactive strategies, [Maes, 94], [Steels, 96], [Murciano97,a,b], it seems reasonable to accept that the performance of higher-order tasks needs hybridisation of reactive strategies and map construction which in some way takes into account the situation of the robot in the world and its self-centred perception.

2. Assumptions

In this paper we illustrate Craik and McCulloch's proposals for incremental construction but with the following suppositions and limitations.

2.1 The World

We deal with buildings in which there may be one or more entrances from the exterior. In principle the rooms are empty and have polygonal forms, the angle between the walls being 90 degrees. There is never more than one door between any two given rooms or cells. The cells are scanned in a clockwise manner always starting from the door through which the robot has just entered. The self-centred perception mechanisms and the door, wall and corner patterns are given *a priori*.

2.2 The robot sensors

We use a cylindrical robot with a motor system capable of independently moving the base of the sensor system. This enables the head and feet to be distinguished. The sensor system comprises a sonar ring, an infrared ring and an impact-sensor ring. The main system used is the sonar, though both the infrared and the impact sensor are also used when traversing doorways. The robot is thus especially well-equipped to be a good sensor of the size of the rooms which compose its world. Given that the robot

mainly uses the ‘rangefinder’ sonar, its self-centred perception system consists of the measurement of the surrounding space in ‘robots’, a ‘robot’ being the size of its ground projection. This measure is always approximate and gives the robot an idea of size which is a direct consequence of being situated in the real world. Even with all this, the perceptive capacity is very limited since the robot has no sight and the reading of the ultrasound echoes is complicated by a large noise component due to reflections, distinct composition of wall materials etc., but these are the rules of the

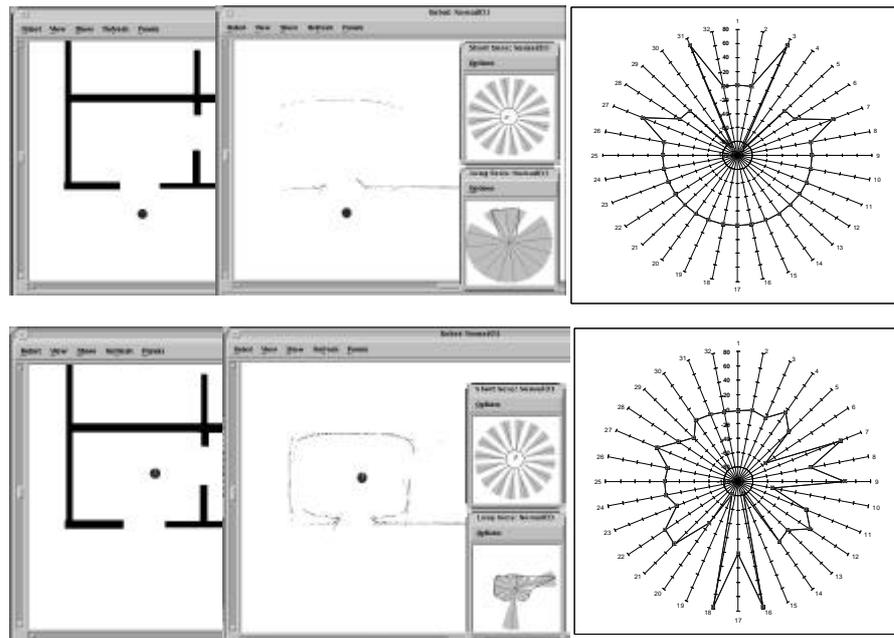


Fig. 2 a) (upper) Illustration of the entrance to a building, the sensor reading and the result of the pre-processing by lateral inhibition. b)(lower)The same example but with the robot inside the room

game: we wish to explore the reach of the proposals of Craik y McCulloch in a simple but real case. In Figure 2, we illustrate two relative positions of the robot with a building as its environment, the raw sensor data and the result of pre-processing by lateral inhibition. The robot is shown in front of the building in the first column of Figure 2.a and inside one of the rooms, having entered the building, in Figure 2.b.

2.3. Pre-processing.

The robot carries out spatial contrast detection on the response from the ultrasound sensors using non-recurrent lateral inhibition. The preliminary data is shown in the second column of Figures 2.a and 2.b. The result after the lateral inhibition has been

carried out appears in the third column of Figures 2.a and 2.b. Thus, Figure 2 summarises the signals available to characterise the outside world. That is to say, the patterns that we will call walls, doors and corners are those which can be obtained by adaptive thresholding and matching operations on these data.

2.4. Behavioural modes.

We have already presented the set of characteristics which describe the sensorial representation of the environment. We now look at the allowed behavioural modes and from the correlation of the two (“percepts”, modes), we obtain the functional specifications of the necessary environment model.

In the eyes of the observer, the robot must exhibit the following capabilities: it must be capable of finding a building, and then of finding an entrance to that

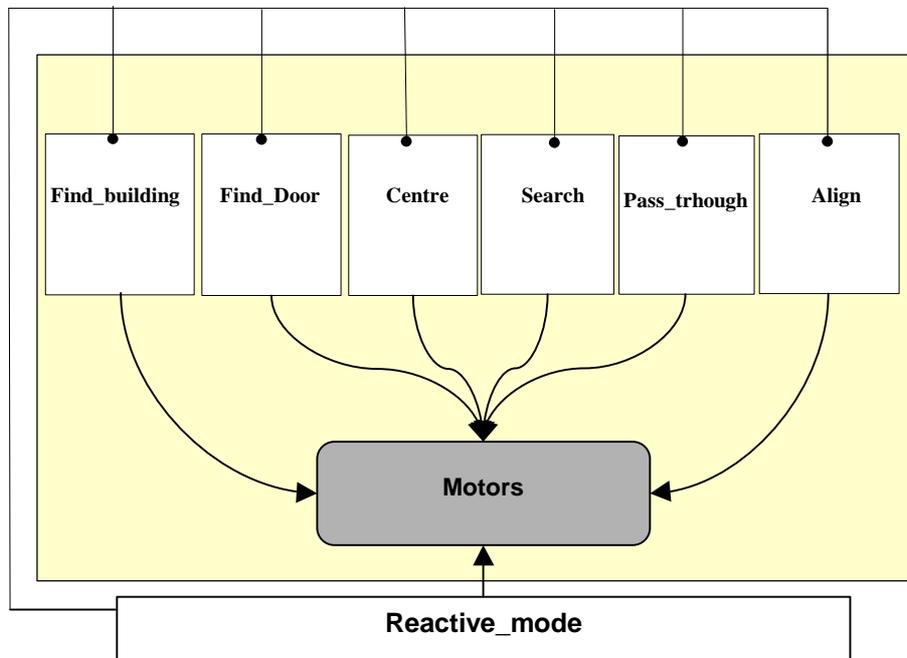


Fig. 1. Functional connection of the behavioural modes. The reactive mode dominates, inhibiting the others.

building. This can occur because the robot enters the mode or context **find_building** or **find_door**, (Figure 3) provoked by the perception of an empty world or of an isolated wall, respectively. Once a wall has been perceived, a pattern characteristic of a door may then be found, this pattern being produced as output of a first pre-processing pass acting by non-recurrent lateral inhibition.

After perceiving the door pattern and aligning with it, the robot moves towards it and traverses the pattern (which also constitutes a pattern recognisable in **pass_through** mode) and goes into **centre** mode in which its structure is altered according to the room in which it finds itself, the changes of orientation perceived by lateral interaction (at the corners) and the patterns recognised as doors in each wall (a wall being that which joins two corners). Such a robot thus constitutes a navigation system with seven fundamental modes which are incompatible in the sense that being in one mode excludes being in any of the others [Kilmer&McCulloch, 69].

The objective of the **find_door** mode is to find an entrance to the building. The objective of the **centre** mode is to ensure that if a room is entered via a door close to a corner, the sampling of the environment is not done from this position. In the **search** mode, the robot turns clockwise, sampling the world at regular angular intervals; the patterns which provoke firing in the pre-processing modify the structure of the robot according to the graph with which it constructs its model of its environment and this modification constitutes its representation of the world. In **align** mode the robot aligns the pattern found in the **search** mode with its main direction of movement and with the main orientation of its ring of sensors. In the **pass_through** mode the robot moves towards the first door on the left which has not already been traversed. During the process of passing through the door, it modifies the firing thresholds of the sonar sensors (to enable temporary violation of the minimum allowed distance from a wall) and advances in the direction of the door pattern. In this mode, the robot is already constructing the graph-model and using it to navigate. Finally, we mention the **reactive_mode** which detects alarm patterns and inhibits any of the navigation system modes in which the robot finds itself at that moment.

3. Incremental construction of the environment model.

Once we have specified the sensorial configurations and the behavioural modes, it is clear that the environment model is the representation of the knowledge necessary to navigate with “these” modes and using “these” environment descriptions. For the construction of the model we use 2-D graphs incrementally constructed using the notation of [Hillier & Hanson, 84]. Successive layers of additional information associated to each of the arcs and nodes can be superimposed in such a 2-D graph. It is constructed by navigating using the seven modes described previously. In consequence, its construction is incremental and the part which is already constructed is used in subsequent navigation.

The construction of the graph is “situated” and is centred on the robot though not as a coordinate system whose axes pass through the robot or by any extension of such a system but rather as a construction carried out in the execution of the robot’s tasks. The graph is a spatial representation of the world of the robot’s world which is only useful for that robot and is not based on any coordinate system.

The graph represents the spatial knowledge of the robot’s world and is continually subject to change; it is dynamic. As the robot covers its world, the nodes and edges of the graph are created and destroyed, always striving for structural coupling, in function of whether the current perception of the world is in agreement with that known by the robot up to that precise moment. In spite of the graph varying, certain “hypotheses” linked to previous knowledge which has now been updated may remain

activated for the performance of other tasks which demand spatial knowledge of the world. This way of starting the construction of the environment model enables the superposition of all the additional knowledge which has a spatial reference with which it can be associated to the nodes and/or the edges of the graph.

It should be noted that while there is no change in the spatial relations, there is only one graph. However, if, for example, doors are closed or previously-unknown doors are opened, the graph changes and all subsequent "perceptions" are added to the basic spatial graph in distinct superimposed layers.

3.1. Hillier and Hanson's notation

Hillier and Hanson [Hillier & Hanson, 84] use the term gamma analysis to denote the analysis of interiors. In their notation, the elementary objects are cells with certain permeability properties. Each cell interior or subdivision can be conceptualised as a point and represented by a circle, with its permeability relations being represented by lines joining it to other circles. In Figure 4.a we see examples of elementary cells

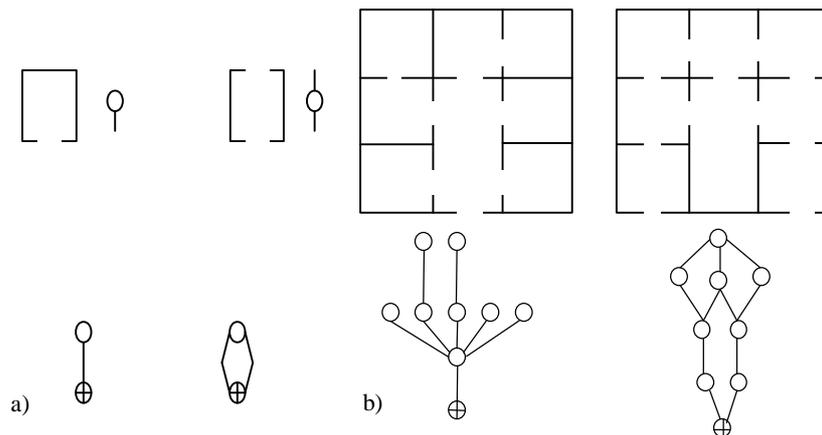


Fig 4. 4.a Hillier and Hanson's proposal in "*The Social Logic of Space*". Elementary cells with one or two entrances together with the representation of the exterior space 4.b Examples of floors of buildings and their corresponding graphs.

with one or two entrances. The space outside a cell is represented by a circle with a cross.

This gives us a representation of the entrance to the structure. Models are constructed using this notation as follows: to each space, cell or graph node, a depth value is assigned, this being the number of nodes that must be crossed to reach it from the node representing the outside, which from here on will be denoted the entrance of the graph. A map of this type is a graph in which the spaces are represented by circles, the permeability by lines and in which all the spaces which have the same depth value are horizontally aligned with respect to the entrance.

These maps enable the symmetry, asymmetry, distribution and non-distribution relations to be easily visualised. The intention of the authors is to provide a form of analysis which permits certain syntactic properties extracted from the graphs to be measured. This strategy combines the visual decoding of patterns with quantification procedures.

3.2. Incremental analysis.

Once the pre-processing has been carried out, the data is passed to the incremental analysis system. Each cell then codes itself as a vector \mathbf{V} having the same number of components as the number of changes of orientation existing in the world (and therefore walls). The value of the component is the number of doors contained in the corresponding wall.

It is important to close the polygon since, once closed, it is independent of the orientation, that is, the vector is circular rather than linear. There is no first or last component so that when the robot enters this cell again it does not matter from where, it must be able to orient itself taking into account that either the room is unique, in which case there is no problem, or using the knowledge contained in the graph (of where it has come from and what it expects to find). In the same way, recall that in this circular vector there is a preferred direction, namely that defined by the convention of sampling the cell clockwise.

Two further components are added to the vector \mathbf{V} indicating the depth of the cell and whether it has been visited or not. When we have coded the cell, in this case \mathbf{S} , (Figure 5) supposing we come from \mathbf{T} , $\mathbf{V}(\mathbf{S}) = (\mathbf{0}, \mathbf{0}, \mathbf{1}, \mathbf{1}, \mathbf{3}, \mathbf{1})$ (recall that the last two components indicate depth and whether visited previously), each component \mathbf{v}_j of \mathbf{V} , which represents a change of orientation, generates a subvector \mathbf{u}_i with dimension the numerical value of \mathbf{v}_j , the value of each component being 0 or 1 according to whether the corresponding door has already been passed through or not. In the example:

$\mathbf{u}_1=(\mathbf{x}), \mathbf{u}_2=(\mathbf{x}), \mathbf{u}_3=(\mathbf{0}), \mathbf{u}_4=(\mathbf{1})$ (Observe that \mathbf{u}_3 has value 0 while the door to \mathbf{C} has not been passed through and \mathbf{C} has not been recognised as already visited. Note also that the last two components of \mathbf{V} do not generate vectors \mathbf{u}_i since they do not represent changes of orientation).

We now go through the first door on the left which has not yet been used. The process is repeated both depthwise and in breadthwise. During this exploration, the connectivity is noted in a proximity matrix. The result is a graph of the building, with information associated to the nodes, which enables the robot to situate itself in the world, distinguish buildings and plan actions.

The system contains **vector generation and update** methods which take care of incrementally generating the data necessary for the construction of the graph. This coding for the changes of orientation does not define each cell unequivocally but provides classification criteria.

Let us examine the graph. Firstly, we can see that it only has one entrance from the outside, since the node marked with a cross only has one connected edge. This node always has depth 0. We then have cells \mathbf{A} and \mathbf{L} with depth 1, each of which has an edge leading to a cell of depth 2. At level 2 we have cells \mathbf{C} and \mathbf{P} . With depth 3 we have cells \mathbf{M} , \mathbf{S} , \mathbf{N} and lastly, cell \mathbf{T} has depth 4. Hillier and Hanson

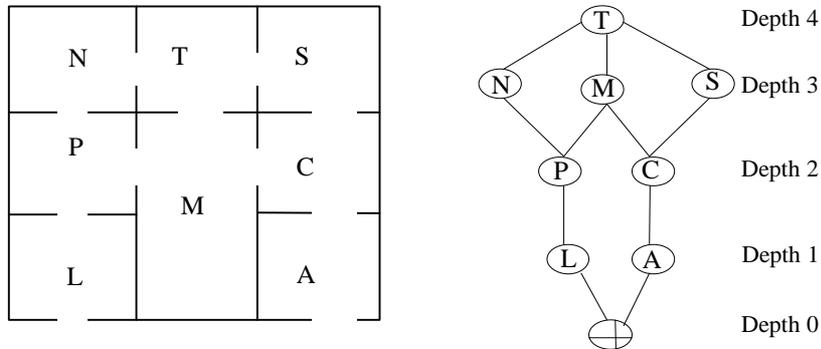


Fig. 5 A building and its graph

add certain additional information to the graph. Our problem is somewhat different since our robot has not the map in memory *a priori* so that it must construct the graph from the modifications which its structure undergoes incrementally and update all the magnitudes associated to each cell which has a relation with the newly explored cell, this depending on the mode in which the robot is operating at that moment (the depth that we may have assigned to each cell may later be affected by the discovery of a new exit to the outside, for example).

3.3. Example of how to construct the graph-model

We illustrate the method described in Section 3.2 with a simple example of the coding of a building. For this we attempt to construct the graph of Figure 5 starting from the plan and forcing ourselves to accept the sensorial and motor limitations presented in Section 2 (Assumptions). The process can be observed in Figure 6. Using the 7 modes previously described, the robot names the nodes and updates the information of the vectors. For the description of the behaviour of the robot we place ourselves in a domain significantly wider than the robots working domain and from which we then observe it.

The first thing is to find an entrance to the building. This is done in mode **find_door**. Once the door-pattern event has occurred, the modes **align** and **pass_through** are entered in order to enter the building, after which the mode **centre** is entered followed by the mode **search**. Once the contrasts traversable and not traversable have occurred, the cell is coded and the vectors are named and updated. The decision of which door to go through is then taken and the process is repeated. Note that we are already navigating through the graph. When a doubt or inconsistency arises (we have to rely on the perception of the external world) the corresponding mode is fired to distinguish one place from another. This occurs when in Graph 1 of Figure 6 the pattern corresponding to a new exit to the outside is found. The robot, through the coding and the associated perception must determine that it is

not dealing with the same entrance to the building, in other words, that the entrance node leads to two distinct cells. At this moment the depth values are which the method has assigned.

Something similar occurs when node **S** is reached. At this moment, when the robot

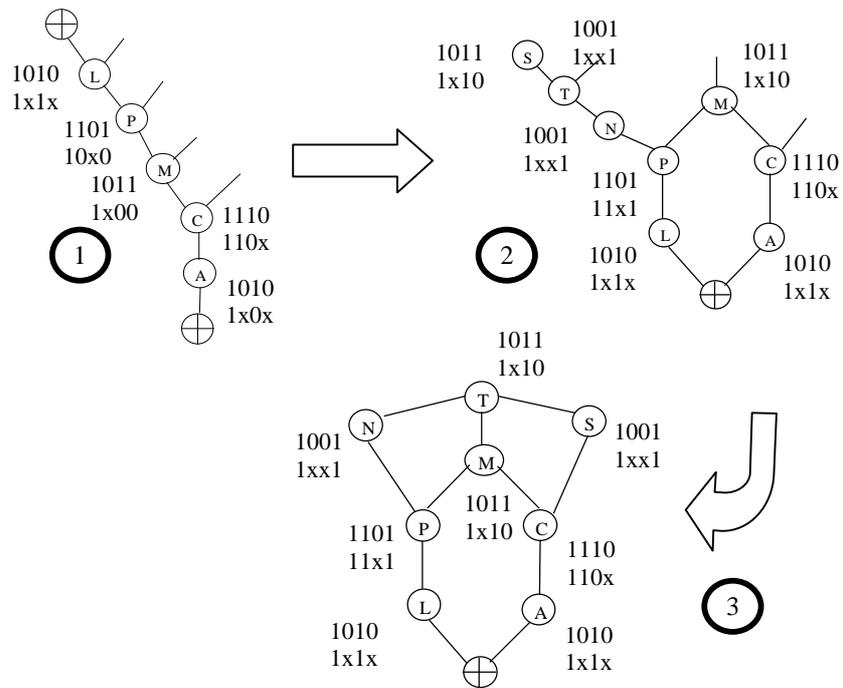


Fig.6 Incremental construction.

passes through the first non-visited door on the left of **S**, it must realise that it has returned to **C** and close the graph. Lastly, it must re-visit **M** in order to then go to **T** and also recognise that it has already been there.

4. Conclusions.

In this paper we have recalled the contributions to conceptual robotics of Kenneth Craik and W. S. McCulloch. The former proposed the concept of construction and constant updating of a model of the environment as the basis of the predictive functioning of the nervous system. The latter formulated an architecture to integrate perception and action, based on the existence of a limited set of incompatible modes of behaviour in primates, according to which “who has the information has the power”.

To illustrate the potential value of these proposals we have developed a simple example using the notation of Hillier and Hanson to describe buildings as a way of

constructing and modifying a graph which represents the “world” of the robot. When the anthropomorphic terms are eliminated, without a causal reference in the computation, the concepts of environment model and of inference based on this model become clear. The model is a graph (domain-knowledge model) constructed using a finite state machine and the generic tasks which use this knowledge belong to a “library of generic tasks” as is commonly used in the context of Common KADS, for example. Once again, the inference control, in the task layer, is another finite state machine which links the input data with the graph modelling the environment.

To what is the environment model limited therefore? In general, it is limited to an abstract, computable structure (such as a graph) constructed from the description of the environment in terms of a set of properties which it possesses which then enables the integration in the model of this formal representation with a repertory of pre-established behavioural modes in an optimum way, in accordance with a certain cost function. On changing the set of properties, the repertory of modes and the procedures for selection between these modes, we are changing the functional specifications of the most adequate environment model for this configuration of “perception/action”.

Acknowledgements

We acknowledge the advice of our colleague Simon Pickin, who have read, commented and improved a previous version of this article. We also acknowledge the financial support of the Spanish CICYT under project and TIC- 97-0604 and the Spanish Comunidad de Madrid under project I+D 0011/94.

References

- [Arkin, 94] “Integrating behavioral, perceptual, and world knowledge in reactive navigation”. R.C.Arkin. In “Designing Autonomous Agents” MIT Press 1994.
- [Borenstein & Koren, 89] “Real-time obstacle avoidance for fast mobile robots”. J. Borenstein, Y. Koren IEEE Transactions on Systems, man and cybernetics. VOL 19, no 5 September 1989.
- [Brooks, 86] “A robust layered contro system for a mobile robot”. R.A.Brooks. IEEE Journal of Robotics and Automation. VOL RA-2, pp 14-23, no 1 March 1986.
- [Craik, 43] “The Nature of Explanation”. K. Craik. Cambridge University Press. 1943.
- [Gonzalez et al, 94] “Map buliding for a mobile robot equipped with a Laser Range Scanner” , IEEE Int. Conf. on Robotics an Automation. San Diego, CA, pp 35-43 1994.
- [Hillier & Hanson] “The Social Logic of Space” Hillier, Hanson, Cambrigde Press 1984.
- [Kaiser et al. 95] “Using Machine Learning Techniques in Real-World Mobile Robots” M.Kaiser, V. Klingspor, J. del R. Millán, M. Accame., F. Wallner, R. Dillmann. IEEE Expert Intelligent Systems & their Applications, pp 37-45 1995.
- [Kilmer & McCulloch, 69] “The Reticular Formation. Command and Control System”. W.L. Kilmer, W.S. McCulloch. In Information Processing in the Nervous System. K.N. Leibovic (Ed). Springer-Verlag Berlin 1969 pp 297-307.
- [Kuipers, 78] “Modeling spatial knowledge” Cognitive Science, 2 , pp 129-153. 1978

- [Maes, 94] “*Situated agents can have goals*” Pattie Maes. In “*Designing Autonomous Agents*” pp 49-71 MIT Press 1994.
- [Mira & Delgado, 95] “*Aspectos básicos de la inteligencia artificial*”. Sanz y Torres 1995.
- [Moravec y Elfes, 85] “*High resolution maps from wide angle sonar*”, Proceedings of the 1985 IEEE Int. Conf. on Robotics and Automation.
- [Moreno & Mira 88] “*Architectures for Integration of Artificial Perception and Action*”. R. Moreno-Díaz, J. Mira Mira. Proceedings of Interkibernetik 87. World Association of Cybernetics, Computer Science and Systems Theory. University of Barcelona. Spain 1988.
- [Murciano97a] “*Learning signaling behaviors and specialization in cooperative agents*”. A. Murciano, J. del R. Millán. Journal of Adaptive Behavior 5(1), 1997.
- [Murciano97b] “*Robot móvil para investigación en grupos cooperantes*”. A. Murciano, J. Zamora, F. de la Paz, J.M. Girón, J. del R. Millán. XVIII Jornadas de Automática y reunión CEA-IFAC , pp 125-131.Gerona 1997
- [Newell, 81] “*The Knowledge Level*”, 1981. AI Magazine, Summer 1981, pp 1-20.
- [Newell & Simon, 72] “*Human Problem Solving*”. A.Newell, H.A. Simon. Prentice-Hall, Englewood Cliffs NJ, 1972.
- [Steels96] “*Discovering the competitors*”. L. Steels. Journal of Adaptive Behavior 4(2), 1996.
- [Sutro et al, 64] “*1963 Advanced Sensor Investigations R-470*”. L.L. Sutro, R.E. Warren, C. Whitman, F. Zeise. Instrumentation Laboratory, M.I.T. Cambridge MA. 1964.