Automation and Robotics in Construction XI
D.A. Chamberlain (Editor)
1994 Elsevier Science B.V.

Unmanned vehicle enabling technology for the construction industry

P J Bateman

Fighting Vehicles & Systems Sector, Defence Research Agency, Chertsey, KT16 OEE UK

Abstract

Research on unmanned ground vehicles has generated a number of techniques which may have application in the construction industry. Examples of improved mobility over rough terrain, teleoperation, and automatic guidance are described, together with a device for handling a range of objects independent of shape, size and fragility. Results of the segmentation of video images using neural networks are presented which suggest potential for automated inspection.

1. INTRODUCTION

Research on unmanned ground vehicles has been in progress at DRA Chertsey for some ten years. Work has concentrated on two main areas: the design of the mobile platform and the design of a suitable payload for various military missions. In general, the radius of operation for such missions is measured in kilometres (rather greater than is necessary for the majority of construction tasks) but during the course of the work a number of concepts and platform configurations suitable for operation at a construction site have been explored. The need to operate at considerable distances from the human operator in an unstructured environment has led to research on increasing the degree of autonomous navigation rather than direct operator control. Computer vision, with this object in mind, has resulted in the development of image processing techniques which may have application in automatic inspection of building structures.

Payload design has concentrated mostly on surveillance techniques, but materials handling and excavation have been studied to a limited extent with some interesting results. The paper describes some of the recent developments at DRA which may have application in the construction industry of the future.

2. PLATFORM CAPABILITY

The task has been to develop a telerobotic capability to allow a wheeled or tracked platform to be driven remotely out of line of sight and over a mainly unstructured terrain. This has involved teloperation using a radio or fibre optic communications link. Tracked vehicles have generally been favoured for off road use and a wide track configuration has been designed to give greater ability to negotiate obstacles.

By placing the two tracks as close together as possible the likelihood of "hang up" failures due to discrete obstacles in the path is markedly reduced. Tracked vehicles are conventionally skid steered, but the wide track leads to a degree of uncertainty in the turning geometry which has been successfully overcome. In the original concept the power plant and drive train were mounted inside the track envelope thus giving a low centre of gravity, and hence good stability for the basic mobile platform. This arrangement does however present problems of engine cooling and access for maintenance. Attention has now focused on the adaptation of existing in-service vehicles to the wide track configuration with the power plant mounted above the track envelope. This vehicle is driven by an on-board operator to the area where remote operation is required. The operator then dismounts, establishes a remote control link and then drives the vehicle remotely.

One of the consequences of the reduced vehicle size which arises from dispensing with the on-board crew is that the vehicle becomes more prone to being trapped by upstanding obstacles or holes in the ground. A pair of pneumatically driven rams has been designed to provide additional mechanical leverage to allow a vehicle to extricate itself from such situations, A ram system of this type has demonstrated the ability to climb an obstacle 3 times higher than without the facility. The introduction of an intelligent sensing unit could optimise the direction of the line of action of the resultant force to aid the recovery. Whilst this design appears to be moving towards the ultimate walking machine, the ram concept, to be used only in case of emergency, is a much simpler aid to mobility.

For full and effective teleoperation a wide field of view needs to be presented to the teleoperator so that he has a clear view of the way ahead, and also some awareness of peripheral features which aids his steering and ability to negotiate turnings, A typical MMI is shown where 3 video monitors are arranged to give contiguous fields of view from 3 cameras mounted alongside each other on the remote vehicle. The bandwidth requirements for 3 video channels are rather demanding for a radio link, and this has been minimised by multiplexing the three video signals to give a half resolution display on the centre monitor and a quarter resolution on the outer displays so that the radio link has to support only one full bandwidth video transmission. Fibre optics presents no such transmission problems and a number of full bandwidth video signals can be accommodated on a single fibre cable. Protective cladding can improve the ruggedness of the cable at the expense of increased diameter and hence size and weight of the cable. Nevertheless, overall cable diameters in the range 0.5 - 1.5mm enable conveniently sized spools to carry several kilometers of cable. The spool is mounted on the remote vehicle and cable is dispensed automatically as the vehicle moves away. In this way teleoperated driving has been achieved at speeds only a little less than conventional manned driving over unfamiliar terrain.

An example of adaptation to remote control of a small commercial vehicle is illustrated by HARP (Hybrid Automotive Research Platform). The conventional skid steering of this small vehicle makes accurate positioning difficult by remote control, and electric actuation of the sprocket wheels has been introduced to provide precision steering whilst retaining the mechanical drive of the internal combustion engine for the main propulsive power. A generator driven off the IC engine ensures that the battery supply for the actuators is correctly charged. The vehicle is controlled either by radio or fibre optic using a simple hand held unit with speed and steering controls.

3. APPLICATION OF UNMANNED VEHICLE TECHNOLOGY

3.1. Shuttle System

Tasks which are unskilled and repetitive are perhaps the most suited and cost effective for automation. Materials handling appears important in this context consistent with the prime requirement to achieve safe operation in the confines of the building site. A rope or cable following technique has been developed which allows a vehicle adapted for the purpose to run astride a rope or cable laid on the ground between two points. A mechanical link which rides along the cable provides the steering signals for guiding the vehicle. Whilst the concept is very simple in principle the design of the mechanical links has required some detailed attention to achieve satisfactory performance in practice. A second link at the rear of the vehicle allows the vehicle to be reversed along the cable. A simple ultrasonic rangefinder can be installed on the vehicle to sense a pole or similar object located at the end of the cable. The rangefinder signals can then be used to halt the vehicle safely before reaching the end of the cable. Unloading and loading can then take place, the operator reversing and restarting the vehicle when ready. The system is suitable for repetitive shuttling tasks particularly over rough terrain, and routes can easily be changed by repositioning the cable.

3.2 Universal Gripper

For remote handling there is an advantage in having a universal gripping device which allows a range of objects (shape, size, fragility) to be handled with ease. The alternative is to provide a range of different grippers suitable for each particular object.

Auniversal gripper has been developed which meets this requirement. It consists of 8 telescopic fingers connected together by means of a drawstring which terminates at the top lifting point. The tautness of the drawstring is maintained by an actuator operating at the top lifting point. This actuator may be driven by pneumatic, hydraulic, or electric power. The advantage of pneumatic or hydraulic power is that the grip can be applied by merely maintaining the appropriate actuator pressure rather than supplying continuous electric current as is the case with electric actuation. The drawstring tension equalises and distributes the force between the fingers allowing quite delicate objects (eg a TV monitor) to be handled safely. The retractable fingers allow objects whose major dimension is greater than the diameter of the device to be catered for, and the multiple finger arrangement enables asymmetrical and complex structures to be handled. The device is shown mounted on a mini excavator, and is capable of lifting a full 40 gallon oil drum weighing over 230 Kg with the fingers gripping on the smooth vertical sides of the drum. Relatively simple calibration of the lifting capability as a function of actuator pressure/current can be carried out prior to working the device at or near the limits of its operating envelope.

3.3 Management of Remote Operations

Teleoperation out of line of sight of the operator, for example, within buildings, where hazardous operations may be required, can present problems in orientation and positioning of the tools and the platform on which they are mounted. The primary aim is to minimise the cost of the remotely operated equipment, and this has been demonstrated by complementing the low cost sensors on the vehicle with the intelligence of the human operator located in the remote command centre.

The sensing system on the vehicle consists of two basic elements (1) a steerable viewing system and (2) a positioning system. The viewing system comprises a miniature wide angle TV camera and an ultrasonic rangefinder mounted on a pan and tilt head. This allows the operator to scan the surroundings and measure the distance of any key features from the vehicle. The positioning system consists of a magnetometer heading reference, and an odometer, which derives its information from encoders mounted on the two tracks (or drive wheels) of the vehicle platform. Differencing the two encoder outputs can also give a useful back up, or alternative, to the heading reference sensor. In order to set up a coordinate system for remote operations say, within a building, the vehicle is steered into the work space under normal teleoperation control. The vehicle is positioned at some convenient point in the work space, and the rangefinder is then driven to execute a polar scan of the surroundings. A print out of this scan is obtained at the command centre. The rangefinder measures range accurately (to 2 cm) but, due to its wide beam divergence, the range points can be confused by openings, doorways and spurious objects. However, interpretation of the range data can be carried out relatively easily by the human operator who can inspect the camera view on his monitor display and resolve ambiguities in the data. The operator then has the facility for drawing in on his display the best accurate position of the walls, and reference points. Any remaining problems in data interpretation can be resolved by moving the vehicle to a new reference point and repeating the ranging scan. In this way it is possible to calibrate the position of the vehicle accurately within its local surroundings, and the on board odometer/heading sensing system can be used to project a "trail" of the vehicle's motion on the diagrammatic plan of the work space displayed on the operator's console. This provides a useful check on the path which the vehicle has traversed, and allows optimum routes or systematic coverage (without overlap) to be undertaken.

The relatively cheap positioning system is subject to errors, especially if the ground plane is uneven or if sharp turns are frequently made, but such errors can be corrected by reference to the rangefinder scans which can be repeated if necessary. Again the communication link can be by radio or fibre optic with the major data processing activity being carried out in the command centre computer under the control of the operator. As an alternative to the ultrasonic rangefinder, similar results can be obtained by triangulation ranging by the operator on key features, using the height of the camera head above the ground as the reference distance. Again, such calibration can be looked upon as a necessary precursor for an autonomous programmed path which may be advantageous for long term repetitive tasks such as concrete finishing, painting, etc.

Whilst this approach has been described in the context of indoor operations, the technique may also have application in the outdoor environment on a building site, where reference points can be surveyed in and the vehicle can be driven to specified positions within the coordinate reference frame in order to perform particular tasks.

4. AUTOMATED INSPECTION

Inspection of building structures might well benefit from replacing the human inspector with a much smaller automatic sensing system which could be positioned more easily in dangerous or hazardous locations, and allow a more rapid scanning of the structure to be carried out. The actual processing of the raw sensor data could be performed at a safe location under the control of the human operator.

We shall describe here some of the processes involved in automated inspection based on the assumption that the raw sensor data can be obtained typically in a video image format or time sequence signal.

Image segmentation is the process of dividing an image into a number of regions or *classes* (which may represent objects) on the basis of some distinguishing attributes or ***features***, such as colour, texture or depth. Classification involves assigning the classes to particular objects on the basis of these features. An example might be the identification of particular surface defects in a visual inspection process using colour based vision. However, the image might equally well be derived from sonar or radar, the only requirement being that the sensing techniques adopted provide sufficient information to recognise and localise the objects of interest.

The classifier uses the sensed features, known as the *pattern* or feature vector, to recognise the object, and the set of all possible patterns forms feature space. The best features for discrimination are those for which there is the largest variance between classes and the smallest intraclass variance. With appropriately chosen features, if the feature vectors are plotted in feature space, similar objects should have feature points in close proximity, which tend to form clusters. These may be separated from other classes using a discrimination function such as shown in Fig 1 for the case of two features (xl, x2) to discriminate between 3 classes of object.

Since the performance of the classifier is measured by its ability to classify previously unseen patterns, it is important to ensure that the training data, on the basis of which the discrimination surface is drawn, adequately spans the feature space. It is from this information that the classifier must learn to generalise and apply it to the recognition of previously unseen patterns.

The two stages to developing a classifier are, firstly, to establish what the most appropriate features are for discrimination (given the task in hand and the sensing technique) and, secondly, given the training feature vectors, to determine the optimum class discriminant.

Since the discrimination function will generally cover more than 2 features (typically up to 10) and there will almost inevitably be some class overlap amongst the training data, some form of automated system optimisation is generally used. This involves supervised learning in which the classifier is presented with training data to which class labels have been attached and attempts to find the optimum class discriminant, so as to minimise the probability of pattern misclassification.

One form of pattern classifier which has met with considerable success recently is based on the Artificial Neural Network, one popular variant of which is the ***multi-layer perceptron*** (MLP). This consists of a parallel feed-forward network of simple processing elements (known as nodes which are analogous to the biological neuron), arranged in a layered structure, with variable connection weights (Fig 2). The classification capabilities of the MLP stem from the non-linearities used within the nodes, (each of which forms a weighted sum of its inputs), and the connection weights. The values of the connection weights are optimised during the training process.

In this case, the ***feature*** vector forms the input to the network which has one output node for each ***class,*** and the desired output has the associated class node set high with the others low. The training process involves repeatedly presenting the training vectors, comparing the actual and desired outputs, and adjusting the interconnection weights so as to minimise this difference.

The real test data is then applied to the MLP classifier which then segments the data based on the experience gained during training. Classification errors will depend on the degree of overlap in feature space and can be minimised by suitable design of the network structure. Whilst the training process may be time consuming, the on-line classification operation is relatively fast, and is unlikely to limit the overall inspection task.

CONCLUSIONS

Research on unmanned ground vehicles has yielded a range of techniques in the areas of mobility, teleoperation, automatic guidance, and inspection, which may have application in the construction industry. Each has been developed to the point of demonstrating practical feasibility. It remains to be seen if cost/benefit analysis studies will stimulate the introduction of such techniques in the market place.

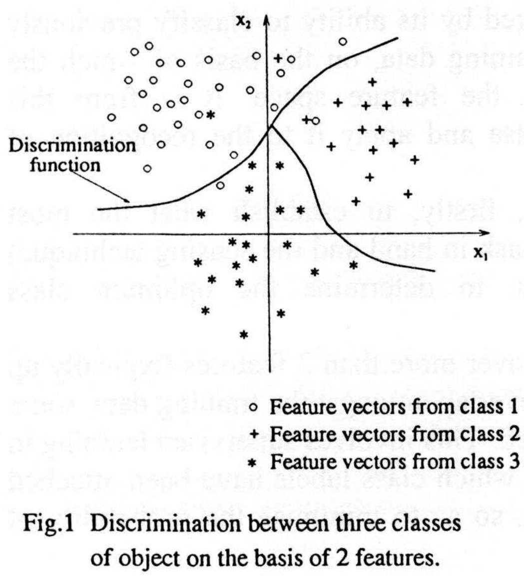

Fig.1 Discrimination between three classes of object on the basis of 2 features.

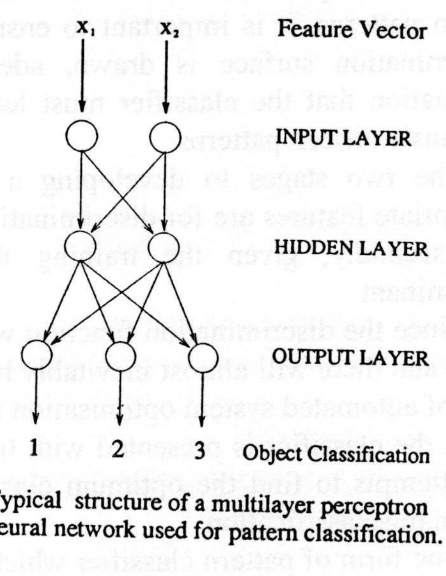

Fig.2 Typical structure of a multilayer perceptron neural network used for pattern classification.