Automation and Robotics in Construction XII
E.Budny, A.McCrea, K.Szymanski (Editors)

Task-Specific Construction Automation and Visually Competent Robots - A Position

G Cockerham[a], T P Pridmore[b] and N Taylor[c]

[a]School of Engineering, Sheffield Hallam University, City Campus, Sheffield, S1 1WB, United Kingdom

[b]School of Engineering Information Technology, Sheffield Hallam University, City Campus, Sheffield, S1 1WB, United Kingdom

[c]School of Construction, Sheffield Hallam University, City Campus, Sheffield, S1 1WB, United Kingdom

Abstract

Sufficient work has now been done in the general area of automation and robotics in construction to allow the identification of more general principles and guidelines to begin. The current paper represents a first attempt to determine some of those guidelines. Consideration is given to the application of robotics. Conclusions drawn from this examination are then extended to the wider field of construction automation.

Introduction

There is currently much interest in the application of robots to and general automation of the construction industry. Systems have been developed which attempt to fill a wide variety of roles within construction and construction-related processes. Some of these attempts have been very successful, some less so. What is perhaps more important than individual successes or failures is that sufficient work has now been done to allow the consideration of more general principles and guidelines. The current paper represents a first attempt to determine some of those guidelines. Consideration is first given to the application of robots; conclusions drawn from this examination are then extended to the wider field of construction automation.

Brady [1] defines robotics as "the intelligent connection of perception to action". The initial question addressed here is what forms of perception and action should be connected if we are to produce effective construction robots?

Attention is focused first on the tasks to be automated. It is argued that before a robot system is designed in any detail the task it is to perform should be well defined and understood. Priority should be given to the analysis and modelling both of the structure of the task and the environment within which it is to be executed. The potential benefits of automation must be made clear, as must any likely disadvantages. In particular, interactions must be identified between the robot's task and those performed by associated human workers.

For a robot to act with any confidence or flexibility it must acquire a representation of the task environment and its own place within it. That representation might include details of the intrinsic properties of the robot - joint angles, etc. - or be primarily concerned with the external world. The present paper is concerned primarily with the extrinsic task environment. It will be argued that the incorporation of computer vision systems into construction robots provides both advantages for the robotic systems concerned and opportunities to advance the science of computational vision.

A multidisciplinary postgraduate centre has recently been established at Sheffield Hallam University based upon the Schools of Construction, Engineering and Engineering Information Technology. A major remit concerns research into automation and robotic applications in the construction industry. Several research programmes are currently underway. The position presented here has arisen from the work carried out by this inter-school group; examples drawn from that work are therefore given in support of our position. While the majority of our work has been concerned with construction robots, research programmes in knowledge-based systems are currently under development. The paper closes with a brief discussion the implications of our stated position for work in this area of construction automation.

Tasks and Environments

The UK construction economy forms approximately 10% GDP [2]. The industry's procedures can, however, sometimes appear archaic: witness road works where a work-gang employing shovels can be seen to be delaying transport, a workforce and ultimately an economy. Office control over large, distant sites is rudimentary in comparison with the manufacturing industry, delivery times become extended and 'extras' can double the original estimate - viz. the Channel Tunnel. Theft from construction sites is legion.

In this context, automation and robotics are attractive; even small relative savings in expenditure obtained through automation can result in large absolute sums of money being conserved. The construction process is a complex one. It is reasonable to expect that a significant number of (perhaps minor) sub tasks exist within this process which can be performed, or at least eased, by a robotic or other automatic system. Moreover, construction is an area that the robotics community has largely overlooked in favour of the manufacturing industry. One would therefore expect many opportunities for automation to present themselves. But, given a potential variety of such opportunities, what sort of tasks should construction robots attempt?

The majority of robot systems developed so far have been intended for use within manufacturing industry. Here both the task and its environment are typically well defined and highly structured. In the construction industry this is generally not the case.

Construction environments are demanding; the terrain is usually rough and highly variable. Task definitions can take the form of a simple statement of the goal to be achieved, with the details of how the target is to be met being left to the workers concerned. Within this complex and hazardous environment, however, are areas which are locally quite well structured and in which the tasks to be performed are quite well defined.

Given the current state of the robotics art, we feel that, in the short term at least, effort should be put into the identification of areas which are sufficiently constrained or constrainable as to allow application of robot technology. The local environment and the task(s) to be performed within it should be modelled, and task-specific construction robots developed. In our group, Broadhurst's work **[3]** on automating an Insituform sewer lining process is an example of this approach.

Although it may be possible to automate some construction processes by mimicking current practice, it will probably be necessary to effect changes in construction culture for the full benefits of automation and robotics to be reaped. Actual tasks, trenching for example, may need to be tackled in a quite different way. With the ubiquitous back-hoe ergonomically related to the 'man with a shovel', automation might benefit from a radically different form of machine.

A compromise which should be adopted as far as is possible is the development of modular systems capable of retro-fit to existing facilities. These should provide for a more receptive market, both in terms of cost and the established skills of the workforce. A key feature of our own work **[4]** on the interpretation of image sequences obtained during surveys of non-manentry sewers is that the techniques proposed should be applicable to existing video records. This means that, in the limit, no changes need be made to current sewer survey equipment.

Regardless of whether the process being automated is long-standing or designed specifically for automation, it is essential that a detailed understanding of the relevant aspects of the construction task and environment be obtained. An argument commonly used in support of research in Artificial Intelligence (AI) is that attempts to build intelligent computer systems force system developers to understand the problem being addressed at a fine level of detail. Building robots has the same effect - with the added emphasis that it focuses attention on the environment in which the robot will work. The design of construction robots and the analysis of construction tasks and environments cannot be separated - at the present stage of development of construction automation we believe the latter to be the more important.

Construction tasks must be understood, not only in themselves, but as parts of a larger process. The industry has a clear need for increased efficiency. Where are the bottlenecks and can robotics help clear them? How many times is the task to be performed, what is the cost of performance and the risk associated with not performing it? It may be that robots provide an ability to do things which are desirable but impossible to achieve manually.

Ultimately, one assumes that cost will control changes in the status quo. Whilst not wishing to adopt a Luddite stance, questions concerning redundant staff must be raised. Why automate if people afford a cheaper alternative when **all** costs (financial and social) are considered? Within manufacturing industry previous robot arm projects have failed to be taken up because the cost of providing a high performance arm is prohibitive when compared to the price of human labour.

A slightly more subtle lesson learnt by the AI community is that resistance from the human workforce can prevent the deployment of new technologies if those technologies attempt to take control of activities that are seen as central to workers' interests. It has been found that much greater success is achieved when AI systems are used either to remove ancillary work which prevents the worker from concentrating on his/her central task or to provide support in the execution of that central task. Man performs certain tasks exceedingly well and semi-automatic systems perhaps offer, in all but the longest of terms, a more viable proposition than their fully automatic counterparts. It is difficult to perceive a fully automated construction site devoid of (human) staff.

Adding Perception

The design of appropriate sensory systems is an integral part of any robot project, many different sensing techniques are available **[6]**. Beacons may be placed around the workplace, providing accurately positioned landmarks to support robot navigation. Laser levelling is highly effective **[7]**. Laser rangefinders are commonly used in manufacturing industry. Spot rangers fire a laser spot at the world and measure its reflection. Ultrasonic systems employ the same principle, emitting an ultrasonic "chirp". Structured light relies upon a laser and camera arranged in a known geometry. The laser projects, for example, a plane of light into the dark and the camera sees a stripe. The geometry of that stripe reflects the geometry of the scene.

All these techniques are "active": each involves the projection of some form of controlled signal onto the environment. Active sensors can provide good quality data, but are typically slow and of restricted range. The equipment involved is often delicate and expensive. Moreover it is not always possible or safe to transmit e.g. high power lasers, across the workplace. Beacons give very accurate position estimates over large distances but are inflexible; if position relative to a new point is required a new beacon must be set up.

Vision, in contrast, is passive; it relies only on the available, visible light. While carefully controlled illumination can be an advantage it is by no means always necessary. Under favourable circumstances vision can operate over very long ranges and at high speed. Modern cameras are cheap, robust, small and light enough to be carried by even low-power robot arms, it is therefore possible to produce a very flexible system capable of positioning a camera where it is needed when it is needed there.

The starting point for computer vision is the digital image: a 2D array of "pixels" each containing a value proportional to the light reaching the image plane of the camera at some point. Given one or more images, vision aims to describe the 3D, real-world situation which generated those images; i.e. to make *explicit* information that is only *implicit* in the pixel arrays. Each pixel value is a complex function of viewpoint, illumination and the geometry and reflectance properties of the viewed object. The goal of vision is to invert this unknown function. The image alone does not provide the information required - some a priori knowledge or assumptions are needed. Any assumptions made need not always be true, but they should hold often enough to be useful. A major question for vision research is what these assumptions might be.

The properties listed above make computer vision an attractive option for construction robots; vision can form the basis of cheap, robust and flexible 3D sensory systems. The price to be paid is that research is often needed to produce interpretation software capable of recovering the necessary 3D information from images of construction environments i.e. to identify useful assumptions regarding those environments.

This exercise is, however, of value to both construction engineering and visual science. It produces a greater understanding of the construction environment which is of use in other areas of construction automation. Perhaps less obviously, it can lead to fundamental results in computational vision. The great majority of computer vision research has been conducted within manufacturing environments. The resulting techniques therefore rely upon assumptions which are often invalid in construction and construction-related situations. During our own work on the 3D description of lateral entry pipes in non-manentry sewers it was found that existing methods make assumptions about light source position which cannot hold within the sewer. This has lead to a novel form of vision which we term Reflective Photometric Stereo **[7]**. Similar advances are expected from our work on image sequences **[4]**. Here, instead of flying past distant, unknown objects, the camera is moving through an enclosed environment of known, but deformed, shape.

Wider Implications

In addition to the construction robot studies referenced herein, research is beginning at Sheffield Hallam into the related matter of Engineering Decision Support systems for the construction industry. Any decision support system must offer the opportunity to predict the implications of front-end decisions, a crucial factor for the client and one generally available in manufacturing industry. It is considered that case-based reasoning **[8]**, a distinct and modern knowledge-based system tool which derives its power from previous case studies rather than from the personnel who undertook them, currently provides the best methodology for providing this.

As previously noted with respect to construction robotics, regardless of any technology employed, the design of a decision-making tool must be informed by detailed consideration of the engineering task it is intended to support. In recent years much research in knowledge-based or expert systems has been directed towards understanding and modelling the tasks performed by such systems. This work has lead to the KADS methodology for knowledge-based system production. A key feature of KADS is a set of Generic Task Models. These provide the system designer with information about the high level structure of commonly performed tasks and the type of knowledge required for their solution. As with construction robots, we believe that an important component of any research into engineering decision support systems should be an analysis of the underlying engineering task at the Generic Task Model level. Decision support systems suited to one generic task or task combination may not be suited to another. It is important that the relationship between the structure of the engineering task and the design of the decision support system be identified wherever possible.

Conclusion

It has been argued that construction robots should, for now, be designed for carefully analysed, specific tasks and environments and that, where extrinsic sensors are required, computer vision, where it is suitable, has interesting advantages over its competitors. In the longer term, construction robots will be required to be robust in less constrained areas of the industry. We suggest (following, for example, Brooks [9]) that the way to this future class of construction robots lies in the generation of complex behaviours by combination of simple behaviours. Work towards these more flexible systems will benefit from access to a body of task-specific, visually competent construction robots.

Research in this area is inherently difficult. This is not a niche market, but a multidisciplinary branch of engineering science in which collaboration and effective communication are at a premium. The construction engineer must be knowledgeable, to a certain degree, of his colleagues' information technology and manufacturing skills - he cannot leave all of these considerations to them. Likewise, they must similarly understand his problems.

It is also expensive. In addition to specialist computing equipment, site and sewer studies at Sheffield Hallam University have, for example, resulted in the provision of a sewer tractor unit (figure 1) and a full-scale sewer mock-up (figure 2). The latter takes up a considerable amount of (expensive) floor space.

Expensive and difficult it may be but automation and robotics will surely come to the construction industry. Academe and industry must decide what they really desire.

Acknowledgements

The authors would like to express their thanks to Hepworth's Building Products UK for their help with and provision of materials for the pipe rig. We would like to thank Neil Bunting and Russ Chapman of SubTech Ltd UK for their technical expertise.

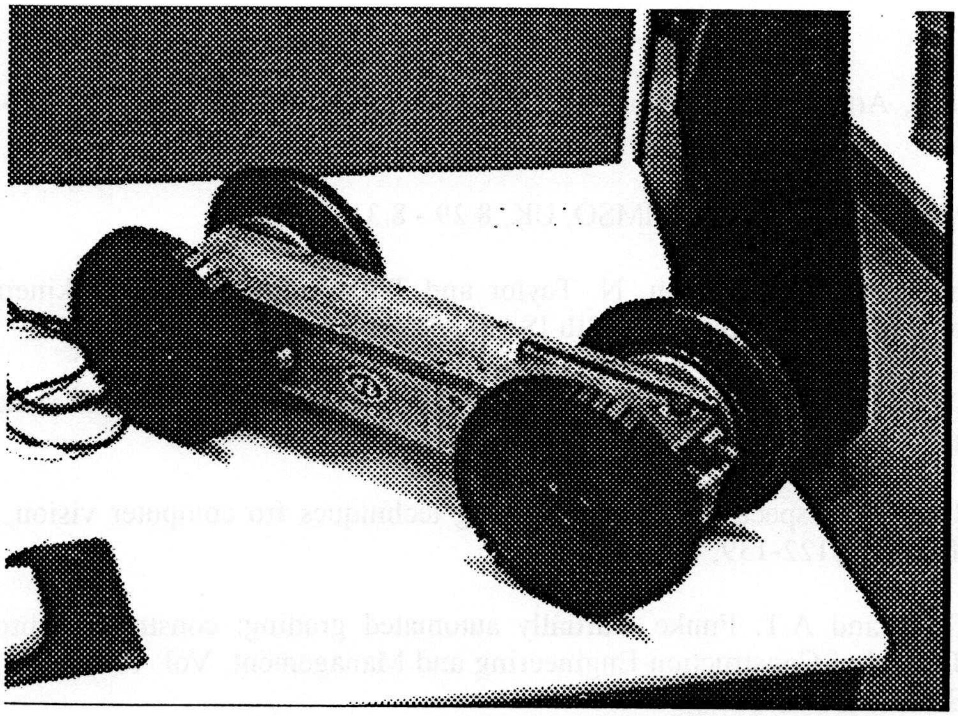

Figure 1: Sewer-going tractor unit

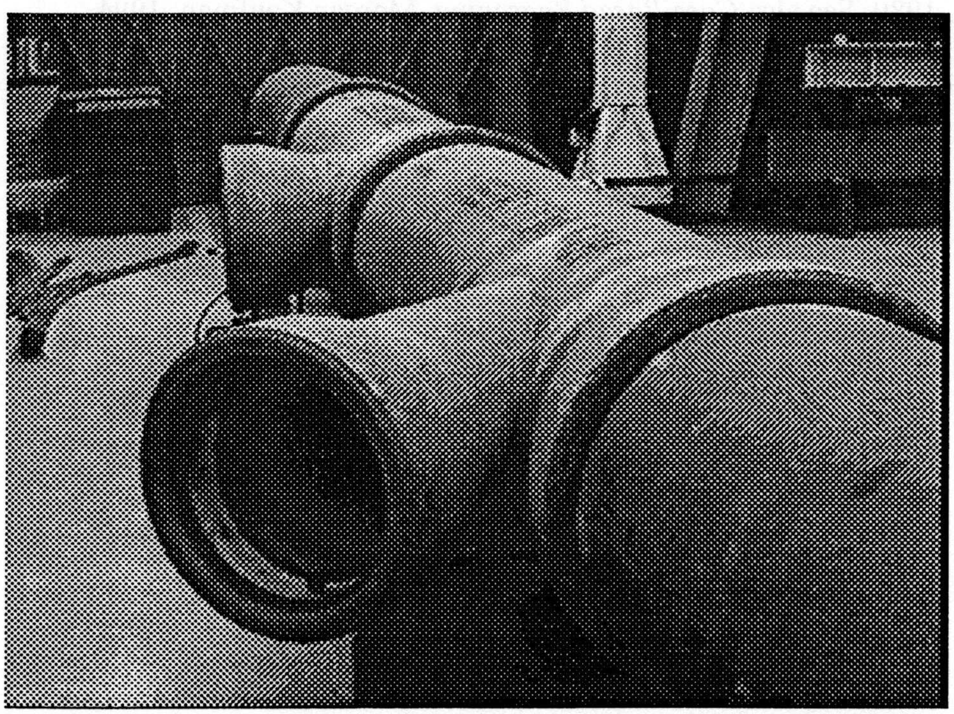

Figure 2: Sewer Mock-Up

References

[1] Brady , J.M, Artificial intelligence and robotics, Artificial Intelligence, pp 79-121, 1986.

[2] Annual Abstracts of Statistics, HMSO, UK, 8.29 - 8.31, 1994.

[3] Broadhurst, S., G. Cockerham, N. Taylor and T.P. Pridmore, Robotic kinematic simulation for sewer renovation, Proc 12th ISARC, Warsaw, 1995.

[4] Fu, S., D. Cooper, T.P. Pridmore and N. Taylor, Determination of sewer structure using a moving camera, Proc 12th ISARC, Warsaw, 1995.

[5] Jarvis, R.A., A perspective on range finding techniques fro computer vision, IEE Trans. PAMI, 5, 2, pp122-139, 1983.

[6] Tatum C.B., and A.T. Funke, Partially automated grading: construction process innovation, Journal of Construction Engineering and Management, Vol. 114, No. 1, pp 133-143, A.S.C.E., March 1988.

[7] Broadhurst, S., T.P. Pridmore and N. Taylor, Three-dimensional description of sewer laterals via reflective photometric stereo, Proc 12th ISARC, Warsaw, 1995.

[8] Kolodner, J., Retrieval and organisational strategies in conceptual memory, PhD Thesis, Yale, 1980. See also *Case Based Reasoning*, Morgan Kaufman, 1994.

[9] Brooks, R.A., Intelligence without representation, Artificial Intelligence, 47, pp 139-159, 1991.