

LOOKING FOR A BETTER ROBOT: VISUAL ROBOT CONTROL FOR CHEAP, FLEXIBLE ASSEMBLY

S.E. Hodges and R.J. Richards

Cambridge University, U.K.

ABSTRACT

Robots are currently used in many industrial applications since they offer numerous advantages over simpler forms of mechanisation and human labour. However, their use is by no means as widespread as was envisaged when industrial robots first became available.

This paper first highlights the problems which are typically encountered when robotic automation is introduced. Following this, an approach based on the use of a computer vision system is outlined. This helps alleviate many of the problems, including the prohibitive cost of many robot solutions. Advances in hardware and software technology in recent years which have made this possible are described.

The potential of the proposed approach has been proven using a cheaply constructed robot mechanism and computer vision system which is able to perform operations central to the production of printed circuit boards (PCBs). The techniques used for successful implementation are highlighted, along with improvements which could be made. Consideration for the suitability of the approach to other robotic applications is given.

1. PROBLEMS IN INDUSTRIAL ROBOTICS

The advantages of robots over simpler forms of mechanisation and human labour are well documented (see (6) for example). They include higher productivity, reduced labour costs, reduced downtime, increased quality and consistency, and inventory and scrap savings. In addition, the same robot can potentially be used for a variety of tasks. However, probably the single most prohibitive aspect of using robotic automation is the initial investment required in purchasing the robot itself. The high mechanical specifications required to repetitively perform the given task with consistent accuracy result in an expensive machine. The expenses incurred when implementing a robotic system are compared with the costs associated with manual labour and hard automation in Table 1. Breaking down the costs for ro-

botics also highlights operating expenses such as indirect labour and maintenance costs, which can be significant (Table 2).

In order to justify the high capital equipment costs involved in automation, high product volumes are a prerequisite. Also industrial robots must be utilised throughout as much of their working life as possible. A robot will typically be expected to last 8–10 years before it is unreliable, unmaintainable, technologically obsolete and of no further value (6).

	Manual	Hard Automation	Robotics
Acquire	Recruitment	Design	Purchase
Set up	Training	Installation	Programming
Run	Labour	Maintenance	Maintenance

Table 1: Major activities/costs associated with three approaches to manufacturing.

It is difficult to adapt hard automation when changes are made to the products. Robotic systems are programmable and therefore have the potential to accommodate such changes more readily. However, any significant change in the product may involve not only re-programming the robot cell, but also replacement of end effectors and fixturing in the work cell. Both of these can be costly processes. Another major cause of difficulty is the inconsistency in parts which is inevitably present¹. Although vision systems can help to counteract product variability, they are very limited in practice.

A third problem is that in practice robots are only suitable for a limited range of applications – primarily (i) certain forms of processing, such as welding, polishing and sealing and (ii) parts handling, for example machine loading, injection moulding and palletising (15). Many other potential candidate applications for robot automation, such as a variety of assembly operations, are still largely unfulfilled because of the technical difficulties associated with them. Indeed, assembly is seen as the largest potential application area for robotics in industry (15).

In summary, if robots were cheaper to purchase and operate, many more industrial processes would become suitable for automation. High machine utilisation, along with correspondingly high product

¹The typical solution to this problem is to increase the quality control of the previous process to reduce any variation.

Investment costs	
1. Robot purchase cost	The basic price of the robot equipped from the manufacturer with the necessary options (excluding end effector)
2. Engineering costs	The costs of planning and design for the robot installation
3. Installation costs	Labour and materials needed to prepare the installation site
4. Special tooling	End effector, parts positioners and other fixtures and tools
5. Miscellaneous costs	Additional investment costs
Operating costs	
6. Direct labour costs	Associated with the operation of the cell
7. Indirect labour costs	Supervision, setup, programming and additional personnel costs
8. Maintenance and repair	Labour and materials; service calls ($\approx 10\%$ of purchase price per annum)
9. Utilities	Usually minor costs
10. Training	Partly an initial investment cost, partly ongoing

Table 2: Costs associated with a robot installation. Taken from (6).

volumes, would no longer be a prerequisite. If the robot was more flexible, then one system would be able to cope not only with product inconsistency, but also with changes in the product itself and the processes it undergoes. This would again increase the application base of robotics to include small, high variability batch environments. Finally, there is a significant area of applications in which robots have not yet proved successful (primarily assembly tasks), largely due to poor reliability of operation.

2. THE ADVANTAGES OF COMPUTER VISION

It is likely there will always be applications in industry in which there is such a large variability of complex tasks that human labour is the only cost-effective solution. Similarly, for the mass production of relatively simple products, robots will probably never be able to compete with fixed automation in terms of speed and throughput rates (15). However, the size of the band of applications between these two extremes, for which robot automation is not only feasible but also cost effective, continues to grow. For example, just over a decade ago, Ayres and Miller (2) cited diamond cutting and polishing as something a robot would probably never be able to do; this has recently been refuted (12). If the goals expressed at the end of the last section are to be achieved, this band must be widened further.

Computer vision may be the key to achieving these aims, at least in part. Robotic vision systems are not uncommon in industry, but they are generally very limited. Typically, a video camera will be used to view a part, so that any variations in its position and shape can be compensated for. Traditionally, such systems are themselves expensive to buy and implement. If a vision system was instead used to monitor the position of the robot end effector,² this

could be used to provide direct feedback of world coordinates to the robot controller.

Indeed, if the visual feedback was accurate and fast enough, the controller would no longer need to rely on accurate position feedback from its joints. Instead, the error in the end effector position could be used in place of joint positioning errors. This idea of *visual servoing* is by no means new – many researchers have tried similar schemes (4, 20). However, rather than implementing visual servoing on a standard industrial robot, which is intrinsically capable of a high positional accuracy and repeatability, this work uses a very low accuracy, cheaply constructed robot. Backlash, link flexibility and poor joint feedback can all be compensated for by the controller. As the characteristics of the robot change over time, due to component wear, temperature changes etc., (more likely if cheaper materials are used in construction), the controller will automatically make appropriate compensation. If there is variation between different robots of a given specification – inevitable if low tolerance components are used in construction – this again is not critical.

Rather than using a static camera which views the entire workspace of the robot, it is possible to mount a camera on the robot arm itself to provide a high resolution image of part of the workpiece³. As long as a visual target is present (this can be artificially introduced in many manufacturing processes), and the robot can move without using visual feedback to a position where this is in its field of view, it becomes possible to servo exactly to the target. Since the camera now only looks at the area of interest, a higher resolution is attained at no extra cost. In certain applications, it is advantageous to position the camera so that both the workpiece and the end effector are in view. With the current position and the target both expressed in image coordinates, it becomes possible to servo without reference to world

²Depending on the configuration of the robot, the position of some or all of its joints may also be necessary.

³In the prototype system a slightly different physical setup is used – see Section 3. The same principles still apply, however.

coordinates. In effect, calibration of the vision system is not required. In this way, any shortcomings in the mechanics of the robot, resulting from its low cost construction, have been overcome. The goal of reducing cost has been achieved without compromising accuracy.

Another advantage of looking at the workpiece and the end effector and using this information as a basis for feedback is that any variations in the position of the workpiece are automatically compensated for. Providing the visual cues which comprise the target do not change, even variations in the parts themselves should not present a challenge – small batches of different products which require similar processing can be accommodated. Thus the second problem cited in Section 1 has also been overcome.

Although not specifically addressed by this work, it also seems feasible that many of the technically challenging robotic applications which have not been implemented robustly in an industrial environment – assembly operations in particular – would be facilitated by the use of visual feedback. By continually monitoring the progress of an operation, and correcting any deviations from the required sequence of moves, robots may be suitable for applications in which automation was previously unreliable.

None of the techniques suggested above are particularly new or original. As mentioned, visual servoing has been studied extensively, as has *active vision*, the process of mounting the camera on the arm itself. What is original to this approach is the application of these techniques with the central aim of making robotics more cost effective. In essence, this is achieved by shifting emphasis (and investment) from the mechanical robot hardware to a vision and control system, i.e. computer hardware and software. Many ideas in image processing introduced and developed over the last few years mean that the application of vision to industrial robotic tasks has become increasingly feasible. For example, *active con-*

ours (11) allow edges in the image to be tracked in real-time as the end effector servos to its target; *partial summation* can dramatically speed up image processing compared with traditional convolution techniques (9) and *projective transformations* (10, 13) enable accurate modelling of viewed objects.

There have also been significant advances in computer hardware in recent years. The power of microprocessors continues to increase at around 35% per year whilst cost falls (7); real-time image processing is now possible at a modest cost. Related advances in VLSI technology have significantly brought down the cost and size of charge-coupled device (CCD) cameras. It is now possible to buy a CCD chip which incorporates scanning circuitry and an analogue-to-digital converter, allowing direct interface with a digital computer without the need for a frame grabber (19).

It is reasonable to expect that computer hardware will continue to decrease in cost whilst increasing in performance in the future. Mechanical hardware costs, on the other hand, are much more likely to stay in line with inflation in future years. Therefore, it seems sensible to make a shift from mechanics to electronic hardware and associated software in robotics applications. In this way, robotic solutions will become increasingly economically viable in the future.

3. EXPERIMENTAL SYSTEM

This work takes printed circuit board manufacture as a specific industrial application in order to demonstrate the feasibility of the ideas discussed so far. The use of vision to facilitate robust assembly in PCB manufacture has been studied in the past (1, 17). In fact, some commercial PCB assembly machines rely on vision to detect bad components, align boards and place fine-pitch chips (14). However, the em-




Figure 1: Layout of the PCB drilling rig





Figure 2: Typical images of (a) PCB to be drilled and (b) component to be placed.

phasis here is to reduce the cost of the mechanical hardware.

Two distinct manufacturing processes have been considered, namely PCB drilling (3) and surface mount component placement (5). They are both implemented using essentially the same mechanical rig (Figure 1), thus demonstrating the flexibility of the hardware. Toothed belts, driven by stepper motors, are used to move the PCB in two dimensions via two perpendicular rods. A single fixed camera covers an area in the centre of the workspace and views the end effector, (either the drill bit or the component to be placed), along with the part of the PCB concerned. Figure 2 shows two typical images. The mechanical costs of the rig, which is designed to be cheap rather than accurate or repeatable, were around £100. In addition to this, a CCD camera, off-the-shelf frame grabber and 486 based PC are needed to provide the necessary visual feedback, bringing total system cost to around £1,000.




Figure 3: Adaptive thresholding and edge detection are used to find the drill bit and the hole centre.

Since only one camera is used, calibration is necessary to compensate for lack of depth information. For drilling, this means that the height of the drill above the PCB must be known. With this information, it is possible to position the drill bit directly above the hole to be drilled, as in Figure 3. Proportional feedback is used to iteratively reduce any

error in the location of the PCB. In initial tests, the rig successfully drilled 1mm holes in a single sided board. By measuring the euclidean distance between the centre of each pad and the centre of each hole drilled, an error distribution can be plotted (Figure 4). This shows that the mean error was 0.07mm, and 95% of the holes were drilled to within 0.12mm.




Figure 4: Error distribution for PCB drilling.

A similar visuo-motor technique is used for surface mount placement. The outlines of the component leads can be detected, along with the pads on the PCB to provide the necessary feedback information. Since the pads onto which the component is to be placed are obscured by the component itself, neighbouring pads are used to guide placement; the required offset from these is calculated from the CAD file of the PCB layout. Again, camera calibration is required for accurate placement, but in practice this has not been a limitation.

4. IMPROVEMENTS AND CONCLUSIONS

As outlined in the previous section, the experiments completed so far make use of a very crude and limited rig. However, the success attained demonstrates the feasibility and the potential of the approach. The nature of the x-y table used means that the kinematics of the system are trivial, and that dynamics do

not come into consideration. The use of stepper motors alleviates the need for closed loop servo control, and makes a relatively slow visual feedback loop acceptable. The addition of a second camera would circumvent any need for calibration, since the ambiguity present in a single image could be removed. It may also prove beneficial to incorporate other sensors to provide additional information, such as force feedback (16). Integration of data from multiple sensors can improve robustness; any redundancy in the information provided can be used to highlight noise and spurious sensor readings (8).

The majority of industrial PCB assembly machines are built with speed of operation a critical consideration, and therefore do not provide a realistic comparison in terms of cost. However, one commercially available component placement machine of comparable speed (18) sells for over £10,000, excluding the cost of component feeders and a PC essential for its operation. This demonstrates the potential cost savings of the approach outlined in this paper. The basic principles applied should readily extend to a variety of other robotic applications, thereby making them considerably more cost effective.

REFERENCES

1. Anderson, L., (editor), 1989, 'Achieving assembly flexibility for printed circuit boards', Machine Vision World, 7:6, 20-21.
2. Ayres, R.U., and Miller, S.M., 1983, 'Robotics: Applications and social implications', Ballinger Publishing Company, Cambridge, MA.
3. Brindley, K., 1993, 'Electronics assembly handbook', Newnes.
4. Corke, P.I., 1993, 'Visual control of robot manipulators - a review', in Hashimoto, K., editor, 'Visual Servoing', World Scientific, 1-32.
5. Drager, Friedrich, Klinge, Meyer, Reithinger, Schoppmann and Waterschek, 1993, 'The world of surface mount technology', Siemens AG, Automation Engineering AUT 5.
6. Groover, M.P., Weiss, M., Nagel, R.N. and Odrey, N.G., 1986, 'Industrial robotics: Technology, programming and applications', McGraw-Hill.
7. Hennessy, J.L., and Patterson, D.A., 1990, 'Computer architecture: A quantitative approach', Morgan Kaufmann, Calif.
8. Hodges, S. and Louie, G., 1994 'Towards the interactive office', Human Factors in Computing Systems, CHI Conference Companion, Boston, ACM Press, 305-306.
9. Hodges, S.E. and Richards, R.J., 1995 'Fast multi-resolution image processing for PCB assembly', To appear in IEE Colloquium on 'Multi-resolution modelling and analysis in image processing and computer vision'.
10. Hodges, S.E. and Richards, R.J., 'Uncalibrated stereo vision for PCB manufacture', in preparation.
11. Kass, M., Witkin, A. and Terzopoulos, D., 1987, 'Snakes: Active contour models', Proceedings of the 1st Intl. Conf. on Computer Vision, 259-268.
12. Krishnipillai, A., 'An automated system for the processing of precious stones', PhD. Thesis, Cambridge University Engineering Dept., in preparation.
13. Mundy, J.L. and Zisserman, A., (editors), 1992, 'Geometric invariance in computer vision', MIT Press, Cambridge, MA.
14. Panasonic, 1994, 'Electronic component placement/insertion machines', Manufacturing Equipment Division, Matsushita Electric Industrial Co., Ltd., Japan.
15. Platt, N., 1994, 'Robots in industry: An overview of robot applications', ABB Robotics, Cambridge University Engineering Department Seminar.
16. Selke, K.K.W., 1993, 'Generic assembly applied to industrial products', Industrial Robot, 20:3, 7-10.
17. Stover, J., 1988, 'Placing fine-pitch SMDs in high-volume environment', Robotics World, 6:2, 19-21.
18. Versatronics, 1994, 'The R-V Placer datasheet', Versatronics Ltd., Witney, Oxon.
19. VLSI Vision Ltd., 1994, 'VVL1070 provisional datasheet', VLSI Vision Ltd, Edinburgh.
20. Wijesoma, S.W., Wolfe, D.F.H. and Richards, R.J., 1993, 'Eye-to-hand coordination for vision-guided robot control applications', Int. J. Robotics Research, 12:1, 65-78.