Look and Learn: Towards Cheap, Flexible Robots

S.E. Hodges and R.J. Richards,
Cambridge University Engineering Department,
Mill Lane, Cambridge, CB2 1RX, U.K.
seh,rjr@eng.cam.ac.uk

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 highlights the problems which have limited the use of industrial robotic automation. To alleviate many of these, including the prohibitive cost of many robot solutions, an approach based on the use of a computer vision system and a learning controller is outlined. Advances in hardware and software technology in recent years which have made this possible are described.

The potential of the proposed approach is demonstrated using a cheaply constructed robot mechanism and computer vision system which is able to drill printed circuit boards (PCBs). The techniques used for successful implementation are presented, along with improvements which could be made.

1 PROBLEMS IN INDUSTRIAL ROBOTICS

The advantages of robots over simpler forms of mechanisation and human labour are well documented (see [1] 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, 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. Implementing a robotic system also incurs costs in addition to the basic price of the robot. These include the design, construction and installation of specialised tooling such as parts positioners and fixturing, and operating expenses such as labour and maintenance costs [1]. In order to justify the high capital equipment and running costs involved in robotics, high machine utilisation is a prerequisite.

Robotic systems are designed to be versatile to accommodate changes in the production line. However, any significant change in the product will involve reprogramming the controller and replacing fixturing in the work cell and the end-effectors. This can be expensive and time consuming.

A further problem is the inconsistency in parts which is inevitably present. The typical solution to this problem is to increase the quality control of the previous process to reduce any variation [2], which is again costly.

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 volumes, would no longer be prerequisite. If the robot was more flexible, then it would be able to cope with variations in the product and its position, again reducing expense. This would increase the application base of robotics to include small batch, high variability environments.

2 LOOK AND LEARN: A NEW APPROACH TO ROBOTICS

2.1 Looking for a better solution

Robotic vision systems are not uncommon in industry, but they are generally very limited. Typically, a video camera will be used to take a single image of a part, so that any variations in its position and shape can be compensated for. Traditionally, such systems are expensive to buy and install.

If a vision system is used to monitor the position of the robot end-effector as well as the target position, it is possible to calculate the error in the end-effector position in image coordinates. This information can be used by the controller to achieve accurate positioning. The idea of *visual servoing* is not new – many researchers have tried similar schemes [3, 4]. 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. Due to the cheap construction, the robot will be more affected by component wear, temperature changes, etc. as time passes; the controller will automatically make appropriate compensation. Variation between different robots of a given specification – inevitable if low tolerance components are used in construction – is not critical. In this way, the capital and maintenance costs of a robotic system can be dramatically reduced, without compromising accuracy.

It is possible to mount a camera on the robot arm to provide an image of just the end-effector and part of the workpiece.[1] Since the camera now only looks at the area of interest, a higher resolution is attained at no extra cost.

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. As long as 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.

2.2 Learning to improve performance

The main problem with the scheme outlined in the previous section is that the visual feedback loop is very slow compared with traditional joint position feedback. Calculating the positional error from images of the workspace is computationally demanding. Many iterations of the feedback loop may be required to servo to the target, which means that the execution of a complete move can be time consuming. Specialised image processing hardware would reduce the problem, but is expensive and therefore defeats the aim of producing a cheap robot system. A more attractive solution is to reduce the number of iterations required.

If joint position feedback is available, an accurate model of the mechanism would allow the error as seen by the camera to be transformed into an error in joint positions. This scheme is impractical for the proposed low specification robot, because the exact characteristics of the mechanism cannot easily be modelled, and the model will differ from one machine to the next. However, providing the mechanism is repeatable, it is possible to learn the transform between joint positions and the end-effector position as seen by the camera as moves are made. Even the effects of backlash and other non-linearities can be learned, provided they are deterministic. Over time, more experience is gained, and the mapping becomes more complete. In effect, a model of the mechanism is generated by the controller.

Each time the vision system detects a positional error (which will be expressed in image coordinates), the move required to compensate can be estimated using the model learned and executed quickly using joint position feedback. Following each move, visual feedback is still used to check if the target has been reached successfully. However, the number of iterations required should be fewer and hence the execution time for a complete move much less than before.

3 IMPLEMENTATION ISSUES

The ideas discussed in the last two sections do not rely on significant technological advances for successful implementation. The novelty of the system is the integration of a number of existing technologies with the aim of producing a cheap, flexible robot without compromising accuracy or dramatically reducing speed of operation.

3.1 Machine vision

Visual servoing has been studied extensively, as has *active vision*, the process of mounting the camera on the arm itself. The originality of this approach is the application of these techniques with the central aim of making robotics more cost effective. This is achieved by shifting emphasis from the mechanical robot hardware to a vision and control system, i.e. computer hardware and software. The 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 contours* [5] allow edges in the image to be tracked in real-time as the end effector servos to its target; *partial summation* can dramatically reduce image processing times compared with traditional convolution techniques [6] and *projective transformations* [7, 8] enable accurate modelling of viewed objects.

3.2 Learning techniques

Many different learning schemes for robot control have been extensively studied. In recent years, particular attention has been given to *neural network*-based algorithms [9, 10] such as *multi-layer perceptrons* and *self-organising networks*. These techniques rely on a model of the transformation which is to be learned; as experience

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

is gained, the model is refined with the aim of making it more accurate.

An alternative approach is *memory-based learning* [11]. In this scheme, experiences are stored explicitly in a large look-up table. In this way, similar previous experiences can be found and used to predict the transformation in the current situation. A particular form of memory-based learning which has been shown to be suitable for learning in robot control applications is *State-Action-Behaviour (SAB) learning* [11]. The essence of this is to encapsulate the state of the system, and then store this along with the control action applied and the behaviour which resulted. This means that at any point in the future, if the system is in a state which was experienced earlier, and that the behaviour required has been previously generated from that state, it can be replicated by applying the same action as before. Methods for efficient storage, retrieval of similar experiences and generalisation to situations not yet experienced have been studied extensively [11, 12].

3.3 Hardware

There have also been significant technological advances in computer hardware in recent years. The power of microprocessors is increasing at around 35% per year whilst cost is falling [13]; real-time image processing is now possible at a modest cost. Computer memory prices are falling in a similar manner whilst capacity increases; thus the requirements of memory-based learning can be met. 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 [14].

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 likely to continue 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.

4 EXPERIMENTAL SYSTEM

This work uses printed circuit board drilling as a specific industrial application in order to demonstrate the feasibility of the ideas discussed so far. The use of vision to facilitate robust automation of PCB manufacture has been studied in the past [15, 16], but without particular consideration of system cost.

4.1 Mechanical hardware

Toothed belts, driven by stepper motors, are used to move the PCB in two dimensions via two perpendicular rods (see Figure 1). Friction in the system is high compared with inertia which means that the dynamics of the system can be ignored. Due to the flexibility of the rods and play in the joints, a large amount of backlash is present. A single fixed camera views an area in the centre of the workspace and is used to monitor the end of the drill bit, along with the part of the PCB concerned. Figure 2 shows a typical image. The mechanical costs of the rig, which is designed to be cheap rather than accurate or repeatable, were around \$200. 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 the total system cost to around \$1,700.

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.

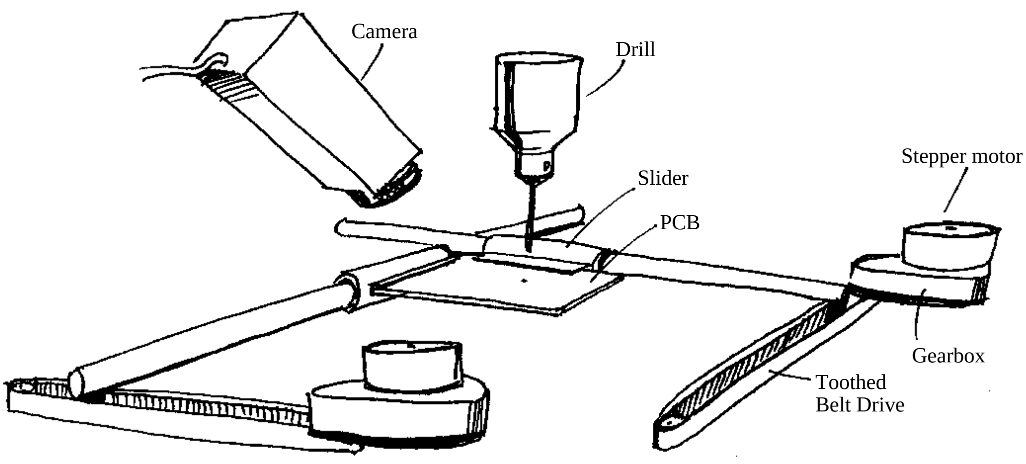

Figure 1: Layout of the PCB drilling rig

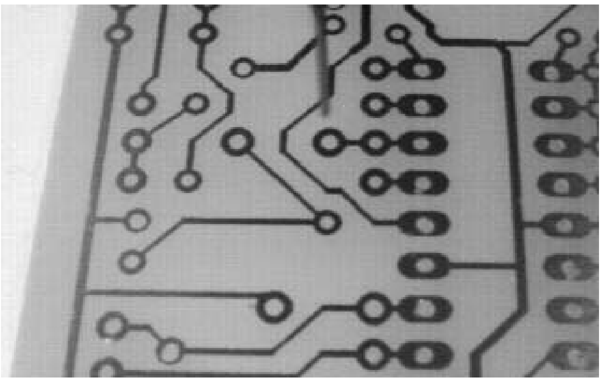

Figure 2: Typical image of PCB to be drilled.

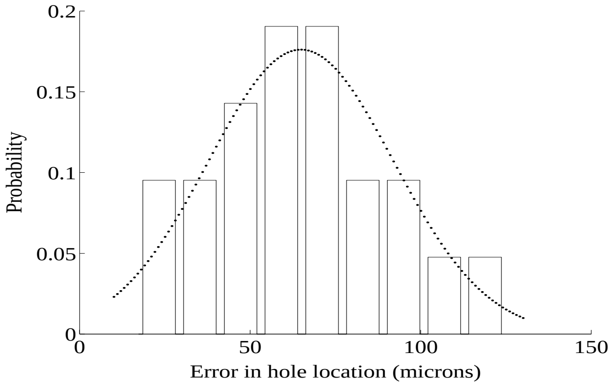

Figure 3: Error distribution for PCB drilling.

4.2 Visual feedback

In order to prove that visual feedback can be used to compensate for the poor mechanics of the system, a simple controller was applied to the task of drilling holes. Proportional feedback was 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 3). This shows that the mean error was 0.07mm, and 95% of the holes were drilled to within 0.12mm.

4.3 Learning controller

Having demonstrated the feasibility of a cheap robot which relies on visual feedback to achieve the required accuracy, it is now desirable to improve the speed of operation. An adaptive controller of the type outlined in Section 3.2 was used to learn a mapping between image space and joint space, so that the joint movements required to move to a given visual target could be calculated fairly accurately. In this way the move can be made more efficiently, even in the presence of effects such as backlash. The vision system was used to check the actual position attained; if the error was significant, further iterations of the control cycle were executed.

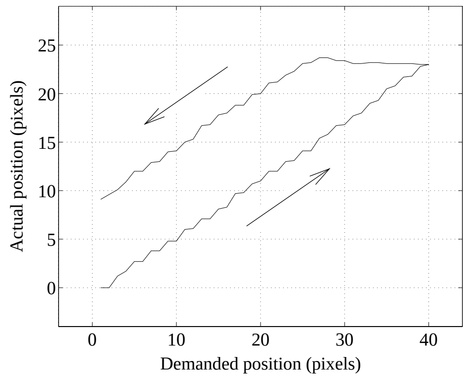

Figure 4: The effects of backlash on system response.

Figure 4 demonstrates how backlash in the system affects movement of the mechanism. A number of steps were made by one motor, first in one direction and then in the other. The figure shows the actual position achieved plotted against the demanded position; the hysteresis introduced into the response of the system is clear. As a result, the final position is considerably different from the

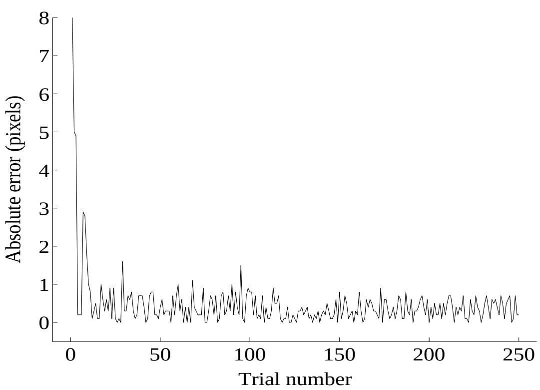

Figure 5: Position error after one move.

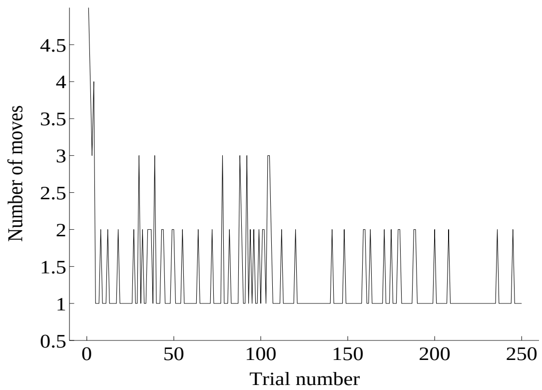

Figure 6: Total number of moves needed to reach target.

start position, despite the symmetry of the move demanded.

Starting with no prior knowledge, SAB learning was used to generate moves to a random target. When the position error was less than $\pm\frac{1}{2}$ pixel, the trial was considered complete, and a new trial with a new target was initiated. To speed up learning, the controller used every other trial for exploration of the state-space by choosing moves with some randomisation.

The performance of the controller was measured by recording the positional error (in pixels) after a single move, and the total number of moves required to reach the target to within $\pm\frac{1}{2}$ pixel, providing there was no overshoot. Figure 5 shows that the error after the initial move decreases with the number of trials, as the controller adapts to compensate for the backlash. Figure 6 confirms that the controller is learning successfully, by virtue of the decrease in the number of moves needed on average to reach the target.

5 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 joint control. 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 [17]. 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 [18].

The learning system presented here is still under development and could be improved in many ways. Different algorithms for generalisation to new situations and management of the potentially large number of experiences are currently under consideration. More extensive and demanding tests are required to prove that the approach is suitable for use in an industrial setting.

The use of visual feedback as a means to compensate for the poor mechanics of a cheap robot has been demonstrated. Due to the nature of this approach, the system is flexible with respect to the position of the workpiece. Initial experiments indicate that a learning controller is capable of adapting to the behaviour of a specific robot. Effects due to the low cost nature of the robot, such as those due to backlash, can be overcome.

Work is underway to extend the ideas presented to PCB component placement. 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 [19] sells for over $20,000. 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. GROOVER, M.P., WEISS, M., NAGEL, R.N. and ODREY, N.G., Industrial robotics: Technology, programming and applications, McGraw-Hill, 1986.

2. PLATT, N., 'Robots in industry: An overview of robot applications', Cambridge University Engineering Department Seminar, ABB Robotics, 1994.

3. CORKE, P.I., 'Visual control of robot manipulators – a review', in HASHIMOTO, K., editor, 'Visual Servoing', 1–32, World Scientific, 1993.

4. WIJESOMA, S.W., WOLFE, D.F.H. and RICHARDS, R.J., 'Eye-to-hand coordination for vision-guided robot control applications', Int. J. Robotics Research, 65–78, volume 12, number 1, 1993.

5. KASS, M., WITKIN, A. and TERZOPOULOS, D., 'Snakes: Active contour models', Proceedings of the 1st Intl. Conf. on Computer Vision, 259–268, 1987.

6. HODGES, S.E. and RICHARDS, R.J., 'Fast multi-resolution image processing for PCB manufacture', IEE Colloquium on 'Multi-resolution modelling and analysis in image processing and computer vision', Digest No. 1995/077, 8/1–8, 1995.

7. HODGES, S.E. and RICHARDS, R.J., 'Uncalibrated stereo vision for PCB manufacture', IEE Colloquium on 'Applications of machine vision', Digest No. 1995/113, 4/1–6, 1995.

8. MUNDY, J.L. and ZISSERMAN, A., (editors), 'Geometric invariance in computer vision', MIT Press, 1992.

9. MARTINEZ, T.M., RITTER, H.J. and SCHULTEN, K.J., 'Three-Dimensional Neural Net for Learning Visuomotor Coordination of a Robot Arm', IEEE Transactions on Neural Networks, 131-136, volume 1, number 1, 1990.

10. NGUYEN, L., PATEL, R.V. and KHORASANI, K., 'Neural Network Architectures for the Forward Kinematics Problem in Robotics', Proc. of IEEE Intl. Conf. on Neural Networks, 393-399, Volume III, 1989.

11. MOORE, A.W., 'Efficient Memory-Based Learning for Robot Control', Technical Report No. 209, Computer Laboratory, University of Cambridge, 1990.

12. OMOHUNDRO, S.M., ‘Efficient Algorithms with Neural Network Behavior’, Report No. UIUCDCS-R-87-1331, Department of Computer Science, University of Illinois at Urbana-Champaign, 1987.

13. HENNESSY, J.L., and PATTERSON, D.A., ‘Computer architecture: A quantitative approach’, Morgan Kaufmann, 1990.

14. VLSI VISION LTD., ‘VVL1070 provisional datasheet’, VLSI Vision Ltd, Edinburgh, 1994.

15. ANDERSON, L., (editor), ‘Achieving assembly flexibility for printed circuit boards’, Machine Vision World, 20–21, volume 7, number 6, 1989.

16. McVEY, E.S. and Van TOL, A., ‘An Experimental Printed Circuit Board Drilling System Automated By Pattern Recognition’, Pattern Recognition, 271-276, volume 11, 1979.

17. SELKE, K.K.W., ‘Generic assembly applied to industrial products’, Industrial Robot, 7–10, volume 20, number 3, 1993.

18. HODGES, S. and LOUIE, G., ‘Towards the interactive office’, Human Factors in Computing Systems, CHI Conference Companion, 305–306, ACM Press, 1994.

19. VERSATRONICS, ‘The R-V Placer datasheet’, Versatronics Ltd., Witney, Oxon, 1994.