PORTRAIT OF THE ARTIST AS AN ENGINEER

Edward Ihnatowicz

[Book proposal. Unfinished and unpublished. Date unknown: at least pre-1988. Found in his papers]

ABSTRACT

The book deals with three major topics:

1. How an artist became involved in high technology
2. Why at present behaviour is more significant than appearance
3. Why physical motion is a precondition of perception

List of possible chapter headings and topics:

1. INTRODUCTION

- Background and reasons for writing the book.

2. SEARCH FOR CONVICTION

- Disillusionment with representational art and mistrust of the abstract.
- Search for the meaning of shapes.
- Functionalism – shapes dictated by function – as a safety net.
- Reconciliation with engineering.

3. ART AND TECHNOLOGY

- Search for motivation for mechanical functions.
- Speculations about ultimate technology.
- Science-fiction story and robots.
- Experiments with hydraulics and control of motion.

4. SAM

- Ideas into practice – first exhibit.
- Behaviour v. Appearance.
- Sound location as a motive for motion.
- Spectator participation closes the loop in autonomous behaviour.
- Movement as a method of communication.

5. THE LOBSTER AND CYBERNETICS

- Discovery of SAM-like shapes in nature.
- Lobsters have simple mechanisms and simple nervous systems.
- Mathematics of motion.
- Analogue computing

6. THE SENSTER

- The Big Break.
- Geodesic structure and constraints of power transmission.
- Computer control.
- Digital sound location.
- Radar.
- Machine as an animal.

Introduction

I am writing this book because of the strong desire to communicate a view of perception which I have arrived at through construction, actual and contemplated, of moving, cybernetic pieces of sculpture. Although I have followed an artistic approach, the results I believe to be valid also in the scientific sense, as well as of practical value in the fields of robotics and advanced automation. The method by which I arrived at these results and the fact that I am not a scientist conspire to make a rigorously scientific treatment of them difficult and would, perhaps, make it counter-productive. The approach I have chosen, therefore, is to describe the various stages I have gone through in my work and in my thinking, hoping that by the end of the book the reader will understand my view of this very complex problem of perception, even if he does not agree with my conclusions.

The very severe limitation of the present-day robot is that it cannot be relied upon to deal with unforeseen changes in its environment even of the most trivial kind. This results from our lack of understanding of the nature of the process by which we, and other animals, inform ourselves about the state of the world outside of us. Attempts at improving this situation are proving very expensive, limited and not very informative as to the direction in which we are going. The problem is urgent and the solution not at all imminent.

My hope is that this book will make acceptable my argument which, in a nut-shell, is that in order for any system, natural or artificial, to be able to deduce anything at all about any object simply by looking at it, it must first be able, or must have been able in the past, to interact with it in some mechanical way. Moreover, only those aspects of the object which can be modified by such actions can ever be successfully interpreted. This view has important implications in the field of visual data processing, leading to the conclusion that the future thinking machine will be a robot and not just a computer.

Synopsis

I shall begin in 1962 when I gave up a reasonably promising career in business (although without having put any money by) to live in an unconverted garage and return to art in which I began as a student in 1945 but subsequently abandoned. Returning to it after such a long break made me extremely self-conscious and inclined to question the motives behind and the validity of every artistic decision. I was now nearly forty but my art has not matured with me and was no better than that I had produced as a student where, in sculpture at least, I was considered by my teachers promising. It was certainly much less convincing. I was encouraged to persevere, however, by the realisation that inwardly I had no faith in the type of art I was trying to produce nor, for that matter of any other being produced around me. My stuff was doffedly traditional and, dissatisfied with it though I was, I resisted any temptation to join any of the trendy movements, hoping to find my own way.

I was, however, always interested in engineering, tinkering really, I was forever helping people with their cars and motorbikes, and so I arrived at a sort of compromise. I produced a number of pieces constructed out of parts of old motor cars and even sold a couple. They were not serious sculpture but I enjoyed making them and it occurred to me that enjoying one's work was important and that, perhaps, I should stop trying so hard to be an artist and just try to find some activity which would make me feel like jumping out of bed every morning to get on with it.

One thing I knew I would always enjoy was working with machines and so, in the end, I have developed a curious form of engineering in which calculation was replaced by intuition and concern for cost-effectiveness by enthusiasm. I tried to design simple components as if they were parts of supersonic aircraft where every ounce of material mattered and cost was no object.

The first result of this "conversion" was SAM or Sound Activated Mobile shown at the 1968 Cybernetic Serendipity exhibition at the Institute of Contemporary Art and the first genuine piece of sculpture I had produced. It consisted of a spine-like articulated structure surmounted by an array of microphones and a flower-like arrangement of acoustic reflectors, which, powered by a hydraulic pump, turned towards anyone making a noise in its vicinity.

The path by which I acquired sufficient knowledge to construct SAM was rather tortuous. It began with the simultaneous realisation that the shapes of the highly engineered components of the cars I was taking to pieces were more satisfactory from the aesthetic point of view than my feeble attempts at abstract sculpture, through having more conviction and an air purposefulness and suitability for the tasks for which they were intended; and also that those tasks invariably involved some form of physical motion or transmission of forces.

Clearly one way of making my abstract pieces more convincing could be to invent imaginary functions for them and make sure that they looked capable of performing them. Better still would be actually to make them move but not in the repetitive way like Tinguely's but in a natural, animal-like way. Movement has always held a great fascination for me and even in the little wax figurines I was making at the time I was attempting to convey if not of actual movement then certainly of potential actions. Clearly a possibility now existed of generating such movements directly. During my vandalising of cars I dismantled and reconstructed a hydraulic braking system and was impressed with the power and precision with which it could be made to move quite heavy objects. This was obviously a very good way of producing very subtle and well-controlled motion and the oil could be delivered to any number of actuators through flexible piping, but to do this required an ability to control precisely the amount of oil being fed to a hydraulic piston. Foot pedals clearly had to be replaced by a motorised pump and the flow controlled by valves. Some method of automatically controlling the valves was required and, even more importantly, an ability to define precisely the motion to be produced.

My first attempts at making hydraulic pistons were quite disastrous and so I looked for ready-made small pistons. After a long search I found some at a hydraulic press manufacturer's who found some included accidentally with a government-surplus press he had bought. Included were some servo-valves which I bought, although neither I nor the dealer knew what they were. In tracking down the function of these valves I made contact with the City University where I learned about the existence of a whole field of engineering science called control engineering, concerned with precisely the sort of problems I was intending to tackle.

The servo-valve turned out to be an electro-hydraulic transducer in which the amount of current flowing through the electrical part of it, the so-called torque motor, determines very precisely the amount of oil that the valve will pass. It was clearly the perfect way of controlling the motion of a kinetic sculpture.

SAM I though of as a neck of a robot which defined sufficiently well the sort of movements it would be required to produce, but this still left me with arbitrary decision as to when and how it should produce these movements. The sound-seeking array of microphones solved that problem although not without some difficulties; the phase-discriminating system finally used was given to me by a friend from Cambridge university, a co-exhibitor at the exhibition.

SAM was reasonably successful, more through luck than design (I did not see it move under electronic control until it was installed at the exhibition) but its movements were far too unpredictable and uncontrollable for their mechanism to be of use in the more elaborate sculptures I was hoping to produce.

I returned to the control engineers. The mathematics used by these people were quite beyond me but I became fascinated by one piece of equipment used extensively in that field – the analogue computer. In learning to use it I had to learn a bit about calculus, which was painful but worth while because I could now think of motion in a much more precise way. I bought an army-surplus oscilloscope, constructed a simple analogue computer and could the spot on the screen move in quite elegant ways. I learned how to make my own hydraulic actuators and found out about the various methods of honing, grinding, hardening and sealing, eventually constructing a simple servo-system which would move a lever in strict accordance with the pattern displayed on the oscilloscope.

Although the various waveforms produced by the computer were pleasing, and the physical motion of the lever encouraging, I needed a more precise way of describing the motions to be produced in terms of velocities and accelerations and time intervals. I also needed to understand better how we and other animals move and to this end I contacted some people working with powered prosthetics, having learned that they were analysing movements of human arms during the performance of various tasks. I was amazed to discover that the motion of a human elbow when performing a well-rehearsed movement from one point to another exhibited an almost constant acceleration and deceleration, the sort of motion that I could simulate exactly on my analogue computer. I have also noticed that these people were using digital logic circuits to sequence and control their simulators. I felt it was time to learn more about digital processing.

Digital computing worried me because it seemed to require not only a much better knowledge of electronics than I possess, but also a knowledge of Boolean logic, Venn diagrams and the like. I tackled this problem by first attending a course on fluidics, a vogue technology then, in which the various logic functions of AND, OR and such could be performed by deflecting jets of air in various ingenious devices. Not involving electronics they were easier to understand so that I felt able to tackle the construction of their electronic counterparts when they appeared in Wireless World as parts of a home-computer project. I eventually constructed a small logic network which, together with a pair of digital-to-analogue converters, enabled my hydraulic lever to perform a great variety of movements.

I took great pride in the fact that the shapes which I finally produced for SAM's neck did indeed look better than my previous sculptures and somewhat bone-like, though I had not tried to imitate any natural forms. I was intrigued, to say the least, therefore, to discover that an almost identical shape existed in nature in the joint of the claw of the lobster. It was not only the similarity of shape which was intriguing. Its operation also was like that of my joint: a simple pivoting action which I had never seem before in nature. Most animals, even those with stiff exo-skeletons have more complex joints which, like our shoulders, can rotate in several planes at the same time. In the lobster all the joints are simple pivots but in spite of this apparent limitation and in spite of having only six of them in any leg, that leg can perform all the required motions with perfect ease.

I was constructing a model of such a leg, the better to understand its construction at the time, when a friend of mine introduced me to James Gardener, the designer, who was responsible for the permanent technological exhibition in Eindhoven in Holland, which was a showpiece of Philips, the giant electrical firm. G, as he was generally known, introduced me to Philips and persuaded them to commission me to produce a large moving sculpture which he eventually christened The Senster.

This was a huge undertaking which took me three years to complete but which enabled me to put many of the ideas I had been toying with into practice. It took the general form of (what else) a great lobster's claw with the pincer replaced by a moving array of microphones like SAM's, except that the whole thing was now run by a digital computer, had proper industrial actuators and servo-valves and I had a professional engineer from Mullards to help with electronics.

I had by that time established a close relationship with a number of people in the Department of Mechanical Engineering of University College London where I went frequently for advice and for the last year of working on the Senster I moved there completely.

I spent about six months in Eindhoven, about half of that time sitting the exhibition hall programming the Senster and observing the interaction between it and the spectators, and I came to the conclusion that the shape and the general appearance of the structure were of very little significance compared to its behaviour, and especially to its ability to respond to the public. People seemed very willing to imbue it with some form of animal-like intelligence and the general atmosphere around it was very much like that in the zoo. Knowing just how little went on in fact inside that animal's head, I felt like a fraud and resolved that any future monster of mine would be more genuinely intelligent.

I was sufficiently naïve at the time to believe that my failure was simply that of not consulting the right people in the Artificial Intelligence fraternity about the correct programmes (I was convinced there would be many such) to use in these circumstances. Although I was delighted when on my return from Holland I was quite unexpectedly invited to join the staff of the Mechanical Engineering department as a research assistant, I was just a little disappointed that it could not have been the department of computer science. Now I feel lucky that I didn't have that choice.

I soon discovered that those involved with A.I. concerned themselves with completely different problems, or at least that their methods, and especially the criteria they applied, had very little relevance to my problems. I decided to do a little research of my own. I felt that since the Senster seemed already almost intelligent I should be able to achieve at least some improvement even on my own and learn something about intelligence, as I understood it, in the process. In any case I had no money to construct any new sculptures and research seemed to be the right thing to do in a university.

I assumed that the most obvious manifestation of intelligence would be an ability to learn, and tried to think of the simplest possible construction in which such an ability might possibly be demonstrated. I constructed a movable array of five photo-transistors, driven by a stepper-motor with which I attempted to track the motion of a small light attached to a moving arm. The light was moved by a free-running electric motor and the photo-array was controlled by a digital computer. The central transistor charged a leaky-capacitor-circuit so that the charge on it was a measure of how often the transistor pointed in the right direction and so constituted my criterion of success.

It took me a long time to realise that this was getting me nowhere. I was getting embroiled in mathematics which I had trouble understanding while the real understanding of intelligence was receding ever farther. I found that I was using terms like perception, information and knowledge, without really understanding how they could be related to an artificial device. I started reading books about animal behaviour, concentrating on the most primitive systems like lice and maggots, hoping that their patterns of behaviour would also be simple, as indeed they proved to be. I have succeeded in writing computer simulations for several such simple mechanisms in which imaginary animals tried various techniques to reach a source of stimulus.

I have arrived at two conclusions: one, that mechanical movement was not only the common element in all such experiments but also the only means by which we could establish the presence of any would-be mental activity, and two, that while the concept of intelligence remained as elusive as ever, the notion of perception seemed as important and perhaps more manageable. Perception, like mechanical motion, must, of necessity, constitute a part of any form of behaviour and can be thought of as the mechanism by which the sensory data arriving from the eyes or ears or any other type of sensor is organised into a form suitable for producing an appropriate response. That response, in the simple systems I was looking at, was invariably some form of motion, so that the immediate problem seemed to be to discover a method of describing the two sets of phenomena: visual patterns, say, and physical movement, in such a way that their correspondence, which was a physical fact in the outside world, could be reflected inside the system.

I felt that I needed to understand more about the nature of mechanical information and decided to concentrate on that. An opportunity occurred when I was asked to help in the supervision of a very bright Chinese student doing his Ph.D. project who had no strong views on what sort of project he should do and was quite happy to have me suggest one for him. I suggested a hydraulically-operated mechanical lever, equipped with pressure sensors and connected to a computer, with which it would be possible to move or exert pressure against a variety of objects and in this manner discover something about their mechanical characteristics.

Being connected to a computer, the arm was capable of operating in two modes: in the position mode it would move to a specified position with a prescribed velocity, largely without regard to any encountered resistance and in pressure mode it would exert a specified pressure against whatever object it encountered. If the specified pressure was zero it would become completely passive and compliant.

At that time the Computer Art Society was staging an exhibition on the fringes of the Edinburgh Festival and asked me to contribute. The arm was all that I could show, so together with the student we turned it into an exhibit. The arm was made to operate in both position and pressure mode and people were invited to move it in any way they liked. When compliant, the computer would store the movements the spectators made and then play them back in position mode. The different ways in which people reacted when the arm suddenly took over were analysed by a statistical programme which was capable of distinguishing between sexes and of classifying people according to their temperament. The results were printed on a teleprinter and were surprisingly accurate. We called it The Bandit, after the One-Arm-Bandits of Las Vegas, which it vaguely resembled, and I hoped that it succeeded in showing people how much very subtle information could be transmitted through such a simple gesture as moving a lever.

The Bandit was, however, a little off the point as far as my main interest was concerned. I was forming an idea that perception ought to relate to objects rather than events; that it ought to enable the system to distinguish between itself and the outside world. I felt that a very important distinction should be made between what could be called non-dimensional sensing, that is, aware ness of changes in some stimulus like pressure, noise or light which have a magnitude but no direction; and the type of perception which could enable the system or animal to determine the shape, size, position or direction of motion of other objects as well as of itself. The Bandit, having only one actuator, could deal only with magnitudes and so another moveable segment was added to it, similarly instrumented and forming, in effect, an elbow.

The new device was reorientated so that the tip moved horizontally, parallel to the surface of a table which could be placed beneath it. I devised an experiment in which the arm could be made to run along a piece of metal placed on the table and the computer could record such runs and deduce the angle at which the piece had been placed from the relative velocities of the two rotating joints. The point of interest here was that the arm was not given any positional information, merely a value of acceleration, and positional information was what came back to it.

Industrial Robotics seemed a field worth investigating by that time, especially since I had been unsuccessful in trying to raise money from the Science Research Council and there was a lot of talk about the industrialists being prepared to back research in this field. Although I was convinced that the topics I was investigating should have great relevance to the tasks that the industrialists were planning to undertake I had no illusions about their willingness to support any work as speculative as mine. I therefore embarked on the construction of a practical industrial manipulator, hoping that, having made a start and having demonstrated its potential, I might find a backer to finance its further development. I incorporated in the design some features which I hoped would make it more suitable for operations where sensing of forces would be required. The task threatened to be rather mundane so I decided to use some of my sculptural techniques in the production of the components. This would not detract from their performance (it could, in fact, increase it) nor would it make them any more expensive to produce in quantity. On the other hand it made them more enjoyable to make and I also hoped that they might form the basis for another moving sculpture.

I concentrated on the three end-segments which would provide the twisting and rotating motion for whatever was to be fitted at the end, leaving the remaining large segments which would provide the reach till later, to be made by another technique, possibly fabricated from tubular steel.

With no backing forthcoming I redesigned the device to provide a limited reach which resulted in a three-degrees-of-freedom manipulator with a reasonable performance but few applications and therefore no prospects. By then one of my colleagues had joined me in this venture and we managed to offset the cost of the development of the device by selling two of them to other research establishments.

The remainder of the book I propose to devote to the explanation of my current ideas about perception and putting forward the proposition that if thinking machines ever develop they will not be computers but robots.