1

Introduction

1.1 Human Design

The ability to design is unique to humans. Although things are built by many creatures, such as the nest of a bird, the dam of a beaver, the web of a spider, these creations are always instinctively produced. It is not the spider that decides the fundamental structure of its web, but the programmed instinctive instructions that evolution has provided for the spider. Only humans have the ability to go beyond instinct and consciously create designs (French, 1994).

Our first tools were little more than sticks and rocks, but from archaeological evidence it is clear that the design of these tools was slowly refined, until finely shaped axes and arrows were being produced. Primitive buildings were designed, then refined. Civilisations grew, with more and more impressive designs; large monuments were built (e.g. standing stones, obelisks, pyramids). New materials were discovered, leading to more efficient tool designs. Ships and carriages followed, the power of steam was discovered, the power of combustible materials, and then the power of electricity. Designs progressed steadily, until a dramatic acceleration occurred, over the last two hundred years (French, 1994). In a very short space of time, an incredible number of advances were made. Aircraft evolved from clumsy devices barely able to fly, to high-performance supersonic jets. Buildings grew to become skyscrapers, computers advanced from slow mechanical adding machines to electronic processing devices of unimaginable speed. In just a few thousand years, humans progressed from being hairless upright apes capable of rudimentary communication and tool construction, to creatures capable of designing spacecraft with the ability to travel to other planets.

Today, the growth and sophistication of human design continues to accelerate. Demands upon designers to produce better designs, faster, have never been greater. With the advent of the computer, new methods to speed up the design process and improve the quality of designs are now being developed.

1.2 Automation of Human Design

The human design process is traditionally a prolonged, iterative business. Initially a conceptual or preliminary design must be created, which is then analysed, experimented upon or tested in use, to determine which parts must be redesigned or optimised. This process of design evaluation and optimisation is repeated until the design is viewed as being acceptable. However, the longer the design process takes, the more costly it is. Consequently, computer software which helps to automate and speed up this process is highly desirable. Furthermore, with the cost of computers ever falling and the available computation power ever increasing, the computer is becoming an essential tool for the designer.

Whilst computer aided design (CAD) software and various computer analysis tools are prevalent today, the use of computers to actually automate parts of the design process (e.g. improving existing designs) is currently less common. Designers have been attempting to use computers to improve their designs for some years, with varying degrees of success. Many techniques have been tried in order to enable computers to conceptualise or optimise many different types of design, but it is only in recent years that success has been forthcoming. Much of this recent success is due to the fact that the use of adaptive search techniques in design has grown dramatically in the last few years (Parmee, 1994). In particular, evolutionary search techniques are now widely used to improve designs by evolution.

1.3 Evolutionary Search

There are many search algorithms known to computer science, of which evolutionary search is a small and recent sub-set. Search algorithms define a design problem in terms of search, where the search-space is a space filled with all possible solutions to the problem, and a point in that space defines a solution (Kanal and Cumar, 1988). The problem of improving a design is then transformed into the problem of searching for better solutions elsewhere in the space of allowable designs, see Fig 1.1.

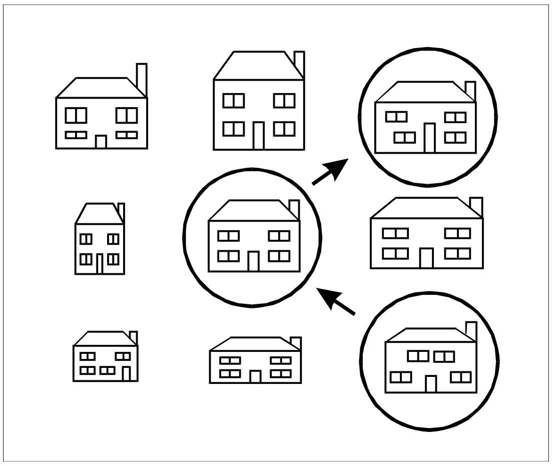

Fig. 1.1 Searching for a solution in an example search space of house designs.

Evolutionary search algorithms are inspired by and based upon evolution in nature. These algorithms typically use an analogy with natural evolution to perform search by *evolving* solutions to problems. Hence, instead of working with one solution at a time in the search-space, these algorithms consider a large collection or *population* of solutions at once.

Evolution-based algorithms have been found to be some of the most flexible, efficient and robust of all search algorithms known to Computer Science (Goldberg, 1989). Because of these properties, these methods are now becoming widely used to solve a broad range of different problems. In the domain of design, the use of evolutionary search to optimise existing designs is becoming widespread (Holland, 1992).

1.4 Natural Evolutionary Design

The use of evolution to optimise designs is not new. For millions of years, designs have been evolved in nature. Biological designs that far exceed any human designs in terms of complexity, performance, and efficiency are prolific throughout the living world. From the near-perfection of the streamlined shape of a shark, to the extraordinary molecular structure of a virus, every living thing is a marvel of evolved design.

Moreover, as biologists uncover more information about the workings of the creatures around us, it is becoming clear that many human designs that were believed to be original have existed in nature long before they were thought of by any human, e.g. the sonar of bats, lenses in eyes. Indeed, many of our recent designs borrow features directly from nature, e.g. Velcro from certain types of 'sticky' seeds, the cross-sectional shape of aircraft wings from birds. As observed by Paton: "A very good example of how biology can inspire engineering solutions is the work of Professor O. H. Schmitt who introduced the term 'biomimetic' (emulating biology) into the US literature over a decade ago. It is fascinating to see how, following his Ph.D. thesis on the simulation of nerve action, four well-known electronic devices emerged: Schmitt Trigger, Emitter-Follower, Differential Amplifier and Heat Pipe." (Paton, 1994, p.51).

1.5 Similarities and Parallels

Strong parallels exist between natural evolution, human design, and evolutionary computation. For example, both natural evolution and the human design process have created some extraordinary designs. Nature uses evolution to generate her designs, and it is apparent that designs created by human designers also evolve over time, as refinements and new inventions are incorporated into them (e.g. cars have evolved from the primitive 'model T' to the sophisticated Formula One cars of today).

Evolutionary search algorithms are inspired by natural evolution and attempt to imitate this process in order to gain the flexibility and efficiency of evolution when optimising solutions.

Likewise, natural evolution is often considered as an optimisation process, improving the ability of creatures to solve the 'problem of life'.

Both the human design process and evolutionary search algorithms construct new solutions to problems using the best features of existing solutions (Goldberg, 1991b). Moreover, just as evolutionary search iteratively optimises solutions, designs created by humans are also iteratively optimised over time. Hence, an 'analogous triangle' of natural evolution, human design and evolutionary computation exists, with each one being highly similar and analogous to the other two, see fig 1.2.

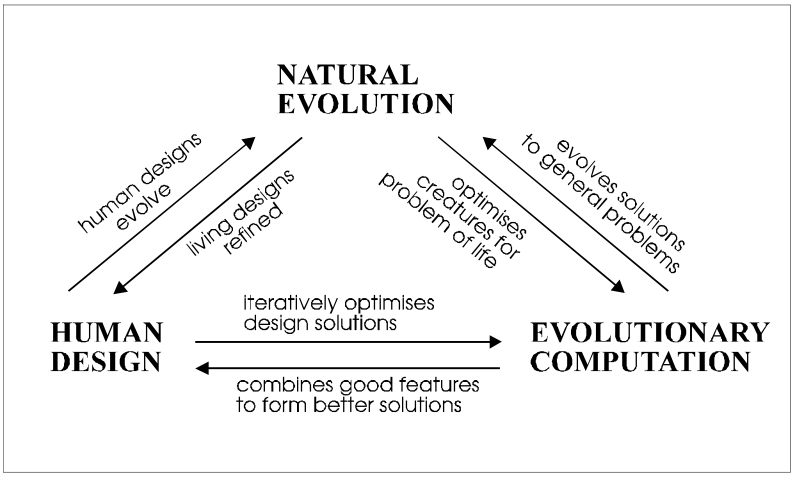

Fig. 1.2 Analogous relationships between natural evolution, human design and evolutionary computation

The work described in this thesis assimilates all three of these processes. Inspired by the astonishing creativity of natural evolution, this research attempts to automate the human design process, by using an evolutionary search algorithm known as the *genetic algorithm*, to evolve new designs from scratch.

1.6 Objectives of the Work

1.6.1 Aim

It is the aim of this work to develop a prototype computer system capable of automatically creating new designs from scratch, in order to demonstrate that such a system is possible, feasible and useful. As described earlier, evolutionary search has been demonstrated extensively in design (in human design optimisation systems, and in nature), so it is proposed that the core of the system will be based upon the most rigorously explored evolution-based algorithm: the genetic algorithm (GA).

As will be shown in the next chapter, although substantial work has been performed in the area of design optimisation by computers, no consequential work has been performed in the area of design *creation* by computers. To enable a computer to create new designs, without any preliminary or existing designs being supplied, involves the automation of the 'conceptual design stage' - the first and most difficult stage in design (Dym and Levitt, 1991). By combining such automatic creation of designs with the automatic optimisation of designs, a computer would be able to perform the entire design process. Additionally, by using 'evaluation software' to automatically analyse the quality of designs, this design process would require little or no human intervention.

Such an automated computer design system would be highly desirable to a human designer. It would speed up the whole design process by automatically providing different designs, already optimised for the problem. Most significantly, however, the system would not be limited by the 'conventional wisdom' of humans, and could potentially create designs radically different from any produced by a human designer. Thus, such a system could be used to create designs on its own, or to inspire human designers to try alternative or unconventional designs suggested by the system.

The proposed evolutionary design system will be generic, i.e. capable of evolving a wide range of different designs with minimal reconfiguration by a user. The generic nature will be limited to the creation and optimisation of the shape of three dimensional solid objects. To limit the complexity of this work, only one object will be evolved at a time, and characteristics such as colour, material, and surface texture will not be considered.

Finally, evaluation software to allow the specification of a number of simple designs tasks will also be created as part of this work.

1.6.2 Intended Capabilities of the System

To explain further, it is intended that the proposed evolutionary design system should have the capability to:

1. Evolve solid object designs from purely random beginnings, or from a combination of random and user-specified initial values.

By initialising the first population of the genetic algorithm with entirely random designs (i.e. seeding the GA with free-form 'blobs'), the proposed system will be given complete freedom to evolve any shape that will fulfil the design specification. This should allow the system to create novel and potentially unconventional solutions to a design problem. Alternatively, the GA could be seeded with user-defined parameter values (defining sub-designs or portions of previously evolved designs). The system could then be used to assemble new designs from these pre-defined components or to re-evolve selected parts of designs. However, it should be stressed that the main thrust of this work is to create a system capable of generating new designs from scratch, i.e. seeding the initial population with entirely random values.

2. Evolve a range of different types of solid object design.

The system should be generic, i.e. able to evolve solutions to a wide range of different solid object design problems. Ideally, the system should perform equally well in all types of design applications tackled and should be scalable to larger and more complex design problems. Realistically, however, it seems probable that certain types of problem will cause difficulties for the system (even the robust GA trips over some 'deceptive problems': Deb and Goldberg, 1992), in which case methods which overcome such difficulties should be implemented.

3. Allow the simple specification of a range of different solid object design tasks.

New designs should be easy for a user to specify, with the minimum of additional evaluation software being required. Ideally most design problems should be specified by the user simply selecting a combination of existing software modules. Fine-tuning of system parameters should be mostly unnecessary, i.e. the good performance of the system should not rely on precisely set values.

4. Successfully evolve designs guided only by evaluation software during the evolution process.

The system should be able to evolve new designs guided by evaluation software alone. This would not only relieve designers of the laborious task of continuously monitoring the system during evolution, it could also help evolution of more unusual and unconventional ideas. If a human designer guides the evolution of designs by the system, conventional designs will usually be evolved. Preventing human interaction during evolution removes the potential limitation of 'conventional wisdom' from the system.

5. Evolve useful and innovative designs.

The long-term goal of this research is to produce a design system capable of evolving truly useful and innovative solutions to real-world design problems. These designs could either be used directly, or could be used by human designers for inspiration. For this project, it is intended that the system should have the ability to successfully evolve acceptable and potentially innovative solutions to model design tasks, created to test the major faculties of the system.

1.7 Overview of Thesis

Following this section, Chapter Two contains a detailed critical review of related areas of research. Chapter Three gives an overview of the evolutionary design system as a whole, and Chapter Four describes the investigation, creation and use of the solid-object representation within the system. Chapter Five explains the methods used to encode designs as genotypes, and the genetic operators used to manipulate the coded designs. Full details of the genetic algorithm used to evolve designs are given in Chapter Six. This is followed by a description of how each type of design was evaluated and an outline of the user-interface in Chapter Seven, with Chapter Eight showing a range of different designs evolved by the system. Finally, Chapter Nine gives suggestions for further work and final conclusions.