Chapter 1

Introduction

This thesis is concerned with the computational modelling of creativity. Consequently, it is also concerned with one of the oldest questions raised about computers: Can computers be creative? Objections to the notion that computers could ever be creative pre-date by over a century the invention of the first practical computers with which to investigate the question empirically. Famously, Ada Augusta, Countess of Lovelace, commented upon her translation of Menabrea's "Sketch of the Analytical Engine", declaring that: "The Analytical Engine has no pretensions whatever to *originate* anything. It can do [only] *whatever we know how to order it* to perform" (emphasis added by Boden, 1990). In Lady Lovelace's opinion, any creative products of the Analytical Engine would have to be credited not to the engine, but to the engineer.

Turing recognised the importance of creativity in any definition of intelligence when he attempted to answer Lovelace's objection in his seminal paper "Computer Machinery and Intelligence", the same paper in which he introduced his now famous test for machine intelligence (Turing, 1950). Turing suggested that objections to the possibility of computers being creative of the type put forward by Lady Lovelace were based on a common misunderstanding of the nature of reasoning in the mind, resulting in an over-statement of the powers of rational thought. In particular, Turing pointed out that a person knowing a set of facts and rules about the world does not mean that the person immediately knows all of the implications of applying the rules to the facts.

Turing suggested that a better variant of Lovelace's objection would be that a machine can never 'take us by surprise' but he then proceeded to declare that computers often surprised him because of his own faulty understanding of what he had 'order them to perform'. In making this argument Turing tried to show that the engineer would be no more responsible for the creativity of a machine than the machine itself because the engineer could not predict the creative behaviour at design time. Turing's argument does not provide us with much information about the possible processes involved in creative thinking but it does highlight the importance of emergence, novelty, and surprise in computational models of creativity.

To be able to study machine creativity within the scope of a single thesis, it is necessary to restrict the discussion to address more limited questions than whether computers can be creative. Following the lead of Turing, and more recently Boden (1990), two questions are asked here. Firstly, can computers model being 'surprised'? Secondly, can the novelty of a surprising discovery motivate the production of creative works?

1.1 MOTIVATIONS

The initial motivation for this research has come from existing computational models of creativity that lack the ability to recognise the novelty of their works. Often computational models rely on generative mechanisms to produce 'interesting novelty' without any means of checking that this is the case other than referring to a human supervisor. Models of scientific and mathematical discovery include heuristics to guide their search processes by determining the interestingness of concepts but these have been shown to be inadequate in systems such as Lenat's AM which still required user interaction to produce creative works (Colton et al., 2000a). An objective of this work is to develop a general-purpose heuristic to guide the exploration of conceptual spaces in search of novelty.

A second motivation for this research comes from studies of creative designing showing the importance of reflecting upon work during the design process. Studies of designers at work have emphasised the interactive nature of designing. Schön calls this style of working 'reflection-in-action' and suggests that the processes involved are critical in many types of problem-solving (Schön, 1983). Models of designing based on Schön's studies place great emphasis on a designer's attendance to the emergence of unexpected consequences of design actions (Schön and Wiggins, 1992; Suwa et al., 1999). To model reflection-in-action design agents must be able to recognise unexpected consequences of their actions. An initial model of reflective sketching is given at the beginning of Chapter 5.

A final motivation comes from the view that creativity is a social-cultural construct, i.e. an honorific label assigned by peers and historians. In this view, creativity cannot be modelled as a closed system within a single agent: instead creativity must be modelled in the context of a society. Csikszentmihalyi[1] (1988; 1999) has been a vocal critic of computational models of creative thinking for not taking into account the effect that society has on the creative agent. The work presented towards the end of this thesis is an initial attempt to integrate some of Csikszentmihalyi's observations into an abstract computational framework for

[1] Pronounced "chicks-sent-me-high".

studying creativity. It is facilitated by the autonomy of curious design agents that permits the construction of artificial societies in which to situate creative activity.

1.2 AIMS AND OBJECTIVES

The aim of this research is to develop an understanding of the role that curiosity plays in creative thinking and creative design. To achieve this aim the following objectives were set:

1) To identify a computationally applicable notion of interestingness that captures what is meant when something is said to be novel or surprising.
2) To develop computational processes that can model interest and boredom based on this notion of interestingness.
3) To develop a model of curiosity that uses interest and boredom and that can be used to guide the actions of a design agent.
4) To develop a computational architecture for developing curious design agents that incorporates the processes of curiosity.
5) To investigate the behaviour of curious design agents in different situations for a number of design domains.

1.3 OVERVIEW OF THESIS

1.3.1 Background

Chapter 2 briefly reviews previous theoretical and computational models of creativity. The majority of Chapter 2 is spent examining the theoretical aspects of Berlyne's work on the perception of novelty and its effects on the behaviour of organisms. Different forms of novelty are described together with the related concepts of surprise, incongruity, and uncertainty and their relationship to expectations, conflict, and the perception of complexity. Berlyne's model of arousal and its relationship to judgements of aesthetics is also covered, importantly Berlyne's theory predicts that the most interesting novelty will be found in artefacts that are similar-yet-different to more familiar works. Berlyne's theory of curiosity is presented to differentiate the types of behaviours that can arise from the perception of novelty. Finally, some recent work developing computational models of curiosity is presented. Computational models have been developed to investigate the benefits of incorporating curiosity in autonomous learning systems, in particular, software agents and mobile robots.

1.3.2 A Computational Framework for Curious Design Agents

The first half of Chapter 3 describes a computational framework for developing curious design agents. The chapter begins with a functional description of a general-purpose agent framework. The addition of a curiosity module is shown to be a relatively small task requiring the modification of few existing processes. The second half of Chapter 3 describes the development of multi-agent simulations involving curious design agents and how this can be used to model Csikszentmihalyi systems view of creativity thanks to the insights of Liu (2000).

1.3.3 Implementing Curiosity

Chapter 4 provides some implementation details of the common components found in the curious design agents described later. In particular, the implementation and behaviour of the neural networks used to implement long-term memory are examined to give some indication of how they might affect curious agent behaviour. The essential process of novelty detection is described and two implementations are presented.

1.3.4 Applications 1: Curious Design Methods

The first chapter of applications, Chapter 5, demonstrates curious design agents in visual design domains using three different design methods: direct manipulation, parametric configuration and design tool-use. Each application is described in a separate section that begins with a brief account of the motivation to develop an agent for the chosen domain; continues with domain-specific implementation details; some experimental work to examine curious behaviour; and finishes with an application specific discussion of the results and potential directions for future work. The collection of agents described in this chapter shows that curiosity is a general-purpose search heuristic that can be applied to agents that design using different levels of abstraction in the design process.

1.3.5 Applications 2: Designing for Other Agents

Chapter 6 presents two more applications of curious design agents that examine different aspects of curiosity in design. The first application uses a simulation of crowd behaviour to examine the design behaviour of a curious agent in a non-visual domain. It also demonstrates the ability of a curious design agent to use the onset of boredom as a trigger to switch between problem-solving and problem-finding. The second application builds on the work done to develop a tool-using design agent in the previous chapter to develop a multi-agent simulation of a creative system, called The Digital Clockwork Muse, using the artificial creativity framework described in Chapter 3. Experiments with the simulation show the emergence of social definitions of whom and what are creative and the development of 'creative cliques' consisting of agents with similar notions of what is creative.

1.3.6 Discussion and Conclusion

The final two chapters discuss the possibilities for developing future curious design agents: additional functions for curious design agents are discussed; the possible application of curious design agents in CAD systems is examined; and the exciting promise of artificial creativity systems to provide insights into the social nature of creative designing is elaborated with some possible directions for future artificial creativity simulations.