

Chapter 2

Evolving Artificial Creatures

2.1 Introduction

As we have seen from the brief introduction given in Chapter 1, the artificial evolution of embodied and situated creatures can be classified into two groups: (1) the evolution of virtual creatures in simulation, and (2) the evolution of real physical robots. Over the last decade, work on evolving robots has become a mainstream effort in robotics and the field has come to be known as evolutionary robotics (Nolfi and Floreano 2000). On the other hand, the evolution of virtual abstract creatures in simulation has not reached the level of maturity achieved by its physical counterpart. As such, there is no commonly agreed upon term that refers to this latter type of work. Some of the keywords used to describe the evolution of virtual abstract creatures in simulation include virtual embodied evolution (Bongard and Paul 2000), virtual creature evolution (Komosinski and Rotaru-Varga 2001), body-brain co-evolution (Hornby and Pollack 2002) and evolution of morphology and behavior (Taylor 2002). In this thesis, we will use the term evolution of morphology and mind to refer to this class of work. As explained earlier in Section 1.1, the word mind here is used to refer to the ANN that acts as the artificial creature’s controller.

It should be noted however that there is no strict delineation between the two fields of physical and simulated evolution of artificial creatures. As we will see later in this chapter, a significant proportion of the work in evolutionary robotics

does actually involve simulation of real robots to reduce the steep time requirements when conducting evolution on real physical robots (Mataric and Cliff 1996). A series of studies on “minimal simulations” has shown that if the simulation faithfully captures the robot’s operation in its environment including the presence of noise in sensors and motors, then evolved controllers in simulation can be successfully transferred to real world robots (Jakobi, Husbands, and Harvey 1995; Jakobi 1997b; Jakobi 1997a; Jakobi 1998). Additionally, some of the highly abstract creatures evolved in simulation, which are far from the design or workings of any real-life robots, have actually been literally “fleshed out” to become tangible, physical manifestations of the real world (Lipson and Pollack 2000; Hornby, Lipson, and Pollack 2001).

This chapter begins with an overview of the importance of embodied and situated evolution. The relevant literature concerning the evolution of real physical robots is then reviewed, followed by a review of the evolution of morphology and mind in simulated artifacts. Figure 2.1 in conjunction with Table 2.1 provides a road-map to the literature surveyed in this chapter on the evolution of different types of artificial creatures. The final level of categorization in Figure 2.1 have been assigned numerical tags of which the corresponding entries in Table 2.1 list the details of the research work in that grouping.

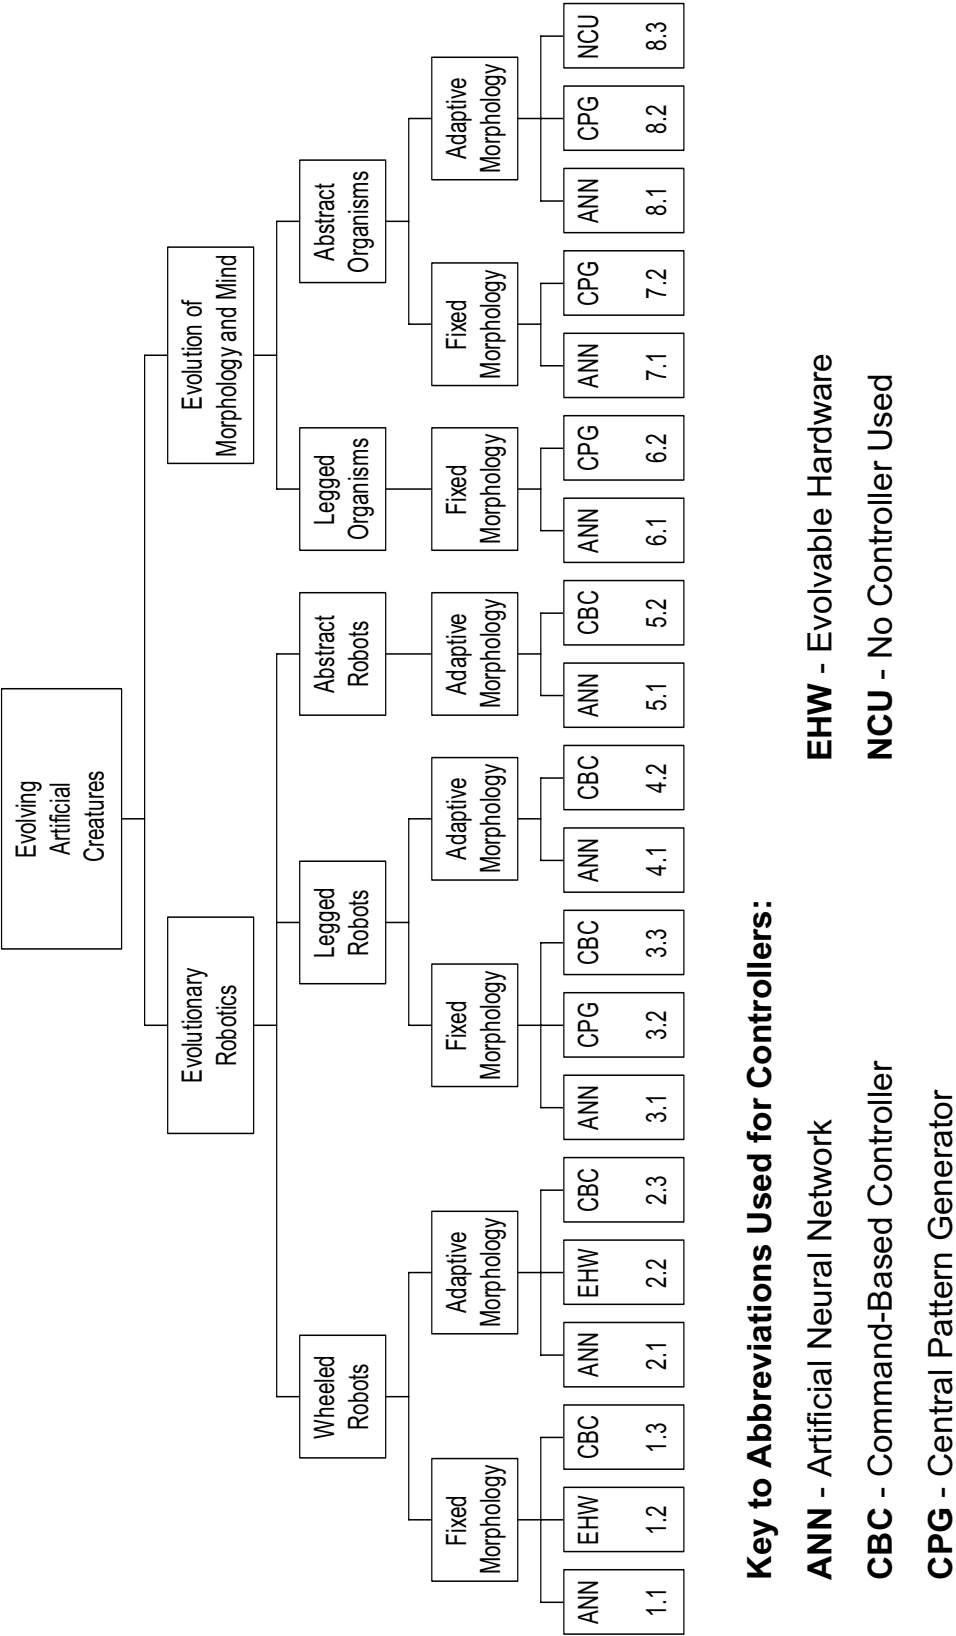


Figure 2.1: Road-map to research on evolving artificial creatures. Numerical tags correspond to entries in Table 2.1.

Ref.	Year	Authors	
1.1	1996	Eggenberger	
	1997	Lund and Hallam	
	1998	Floreano and Mondada	
		Floreano and Urzelai	
		Husbands, Smith, Jakobi, and O'Shea	
	1999	Nolfi and Floreano	
	2000	Floreano and Urzelai	
	2001	Floreano and Mattiussi	
		Floreano, Nolfi, and Mondada	
		Floreano and Urzelai	
		Hulse, Lara, Pasemann, and Steinmetz	
		Husbands, Philippides, Smith, and O'Shea	
		Pasemann, Steinmetz, Hulse, and Lara (a)	
		Pasemann, Steinmetz, Hulse, and Lara (b)	
		Lara, Hulse, and Pasemann	
		Smith, Husbands, and O'Shea (a)	
		Smith, Husbands, and O'Shea (b)	
		2002	Eggenberger, Gomez, and Pfeifer
			Floreano, Schoeni, Caprari, and Blynel
			Nolfi
			Philippides, Husbands, Smith, and O'Shea
			Smith, Husbands, Philippides, and O'Shea
			Smith, Philippides, Husbands, and O'Shea
			Watson, Ficici, and Pollack
1.2	1995	Thompson	
	1996	Keymeulen, Durantez, Konaka, Kuniyoshi, and Higuchi	
	1997	Thompson	
	1998	Keymeulen, Iwata, Konaka, Suzuki, Kuniyoshi, and Higuchi	
1.3	1996	Nordin and Banzhaf	
2.1	1997	Harvey, Husbands, Cliff, Thompson, and Jakobi	
2.2	1997	Lund, Hallam, and Lee	
2.3	1996	Lee, Hallam, and Lund	
3.1	1992	Beer and Gallagher	
	1996	Gallagher, Beer, Espenschied, and Quinn	
	1997	Gruau	
	1998	Jakobi	
		Kodjabachian and Meyer (a)	
		Kodjabachian and Meyer (b)	
	1999	Reeve	
	2001	Fujii, Ishiguro, Aoki, and Eggenberger	
		Otsu, Ishiguro, Fujii, Aoki, and Eggenberger	
Paul and Bongard			
3.2	2001	Reil and Massey	

Ref.	Year	Authors
	2002	Reil and Husbands
3.3	1998	Gomi and Ide
	1999	Hornby, Fujita, Takamura, Yamamoto, and Hanagata
	2000	Hornby, Takamura, Yokono, Hanagata, Yamamoto, and Fujita
4.1	2001	Bongard and Paul
4.2	1997	Arnold
5.1	1993	Cliff, Harvey, and Husbands
	1994	Dellaert and Beer
	1994	Harvey, Husbands, and Cliff
	1996	Cliff and Miller
	1997	Husbands, Harvey, Jakobi, Thompson, and Cliff
	1999	Lichtensteiger and Eggenberger
	2000	Lipson and Pollack
5.2	1999	Dittrich, Skusa, Banzhaf, and Kantschik
6.1	1998	Ijspeert, Hallam, and Willshaw
	1999	Ijspeert
		Ijspeert, Hallam, and Willshaw
		Ijspeert and Kodjabachian
	2001	Ijspeert
	2002	Bongard (a)
	2002	Bongard and Pfeifer
6.2	2000	Ijspeert
		Ijspeert and Arbib
7.1	2002	Mandik
7.2	1997	Gritz and Hahn
8.1	1994	Sims (a)
		Sims (b)
	1999	Komosinski and Ulatowski
	2000	Bongard and Paul
		Komosinski
		Ray
	2001	Bongard and Pfeifer
		Hornby and Pollack (a)
		Komosinski, Koczyk, and Kubiak
		Komosinski and Kubiak
		Komosinski and Rotaru-Varga
		Taylor and Massey
	2002	Bongard (b)
		Hornby and Pollack
8.2	2001	Hornby, Lipson, and Pollack
		Hornby and Pollack (b)
8.3	1997	Eggenberger

Table 2.1: Summary of literature survey on evolution of artificial creatures. Num-

bered references refer to the numerical tags assigned in Figure 2.1.

2.2 Situatedness and Embodiment

The importance of embedding the study of evolving artificial creatures within the twin principles of situatedness and embodiment is perhaps best exemplified by how natural evolution occurs in real biological organisms. Natural creatures such as animals and insects have bodies and are situated in a physical environment. Their skills and behaviors are developed autonomously through the intimate interplay with their environment. As such, in order to create artificial creatures that might possess some of these novel properties exhibited by real creatures, such systems must be built based on the principles of situatedness and embodiment.

This view of intelligence as an emergent phenomenon of embodied and situated artifacts is regarded by many researchers to be the foundation for successful design and implementation of artificial agents. On embodiment, Varela (1995) pointed out that

“Cognition depends on the kinds of experience that come from having a body with various sensorimotor capacities.” (p.15),

while Brooks (1995) stressed that

“The robots have bodies and experience the world directly — their actions are part of a dynamic with the world and have immediate feedback on their own sensations.” (p.29),

and Arkin (1998) stated that

“A robot has a physical presence (a body). This spatial reality has consequences in its dynamic interactions with the world ...” (p.26).

Following on to situatedness, Varela (1995) explained that

“The individual sensorimotor capacities are themselves embedded in a more encompassing biological and cultural context.” (p.15),

while Brooks (1995) highlighted that

“...robots are situated in the world — they do not deal with abstract descriptions, but with the here and now of the world directly influencing the behavior of the system.” (p.29),

and Arkin (1998) stated that

“The robot is an entity situated and surrounded by the real world. It does not operate upon abstract representations of reality, but rather reality itself.” (p.26).

In Dautenhahn’s work with socially intelligent robots (Dautenhahn 1996; Dautenhahn 1999), the importance of embodiment was discussed at length in designing reactive cognitive architectures in physical robots and other artificially intelligent agents that can exist in simulation. Dautenhahn (1996) highlighted the fact that

“...there is much evidence to support the assumption that cognitive capabilities are only possible through the interaction of body and mind, i.e. that the body is not simply used by the mind, but that there is a co-development and mutual shaping of cognitive abilities on the one hand and bodily skills and experiences on the other hand. The body is not a fixed and pregiven ‘actuator device’, but it is a dynamic and ontogenetically evolving entity.” (p.27).

Furthermore, Dautenhahn (1996) argued that the study of embodied and situated artifacts will play a significant role in bridging the gap between phenomenological understanding and the computation-theoretic approaches normally adopted in cognitive science, artificial intelligence and artificial life studies.

Mataric (1997) addressed the issue of how physical embodiment is related to cognition and reviewed both biological and artificial studies that have endeavored to answer this question. It was argued that artificial systems are preferable over biological systems as although biological data are abundant, they are often disconnected and incomplete due to the restrictions that apply when working with real

rather than artificial life. On the other hand, the use of artificial systems allows the researcher complete freedom to experiment with whatever aspects of embodiment and cognition that are of interest, in particular ablation of neural pathways, amputation of limbs and/or other forms of disablement of sensory-motor capabilities, which are central to the question of the role of embodiment in higher-level cognition. In the author's own study, artificial mobile agents were used to answer a number of key questions relating to social group behavior as well as imitative behavior, and how these behaviors are in turn related to embodiment and cognition.

Nolfi and Floreano (2002) importantly pointed out that for an external observer, designing such situated and embodied creatures capable of autonomously developing the desired behavior through dynamical interactions with their environments is a very complex task. They further explained that there were two ways in which this can be achieved: (1) by painstakingly recreating the artificial creature through careful mimicking of natural organisms, or (2) by employing an artificial evolutionary process that allows for self-organization to occur automatically. As such, this makes the evolution of embodied and situated creatures a prime candidate for evolutionary computation techniques.

2.3 Evolutionary Robotics

Evolutionary robotics is defined to be the synthesis of autonomous robots using artificial evolutionary methods (Nolfi and Floreano 2000). An early review of this field of research is given by Mataric and Cliff (1996) where the majority of studies focused mainly on the evolution of control structures only. A more recent overview highlights the move of evolutionary robotics into evolving both the control and morphology of robots where the interplay between brain and body is considered to be a crucial factor in the successful synthesis of autonomous robots (Nolfi and Floreano 2002). A thorough treatment of the field can be found in the seminal textbook written by Nolfi and Floreano (2000) on this subject.

As pointed out by Harvey (1997), the design of controllers for robots is

a complex task not suited to human divide-and-conquer design strategies. There are 3 major problems: (1) it is not obvious how the controller system should be decomposed, (2) interactions are not limited to direct connecting links but are also mediated through the environment, and (3) interactions between sub-parts grows exponentially as system complexity increases. Thus, evolutionary approaches to controller design are desirable, where the only benchmark is the overall behavior that should be achieved by the system.

However, Ronald and Sipper (2001) recently pointed out that emergence stemming from the use of biologically-inspired solutions in engineering problems may be problematic because unexpected and sometimes unwanted results or behaviors might arise. Using the so-called emergence test, it was claimed that evolutionary robotics exhibited mild emergence where the degree of surprise is limited to well-defined boundaries (*unsurprising surprise*). On the other hand, traditional hard-wired engineering solutions exhibited no surprise (*unsurprising*) while artificial life exhibits a very high degree of surprise (*surprising surprise*). Nonetheless, it was surmised that emergence in engineering solutions that draw on inspirations from nature such as evolutionary robotics and the related reliability issues are unavoidable consequences if the desire is to design smart, adaptive and evolvable machines. In general, evolutionary robotics can be grouped into three main categories, those involving the evolution of (1) wheeled, (2) legged, and (3) abstract robots.

2.3.1 Wheeled Robots

A hybrid genetic programming (GP)/genetic algorithm (GA) methodology was used to evolve both the controller and parameters of a wheeled robot's morphology in simulation (Lee, Hallam, and Lund 1996). The controller consisting of a tree-like program was evolved using the GP part of the system while morphological parameters such as the robot's body size, wheel radius and wheel base size encoded in a linear string of real numbers were evolved using the GA part of the system. Individuals were assessed for obstacle avoidance behaviors using a fitness function that combined multiple terms such as distance from obstacles, forward speed and

rotating speed into a single objective. It was claimed to be the first study which co-evolved both the controller and morphology of robots and concluded that because the evolved controller only functioned within the co-evolved body, the evolution of the body component played a significant role in the success of the evolutionary process. An island-GA model was used to maintain genetic diversity during the evolutionary process. In a related study using simulations, Khepera wheeled robots were shown to require only simple perceptron controllers that directly connected sensors to motors for evolving behaviors such as exploration and homing (Lund and Hallam 1997). It was claimed that the robot's perception of its environment's geometries allowed time-related components to be encoded without requiring any recurrent connections in the controller. GP alone has also been used to evolve controllers for Khepera robots for obstacle avoidance and object tracking behaviors utilizing a combination of simulated and real-world testing of evolved controllers (Nordin and Banzhaf 1996).

The Species Adaptation Genetic Algorithm (SAGA) algorithm was used to evolve both the controller and visual morphology parameters for simple navigational tasks in a two-wheeled mobile autonomous robot (Harvey, Husbands, Cliff, Thompson, and Jakobi 1997). The desired behavior was evolved within 50–100 generations using 40–60 individuals that were evaluated using a simple single-objective distance-based fitness function. SAGA allows for increases in length to genotypes and hence it was argued that it permitted incremental evolution to occur during the evolutionary process. Conversely, Eggenberger (1996) reported the use of biological cell differentiation techniques in order to reduce the length of the genotype encoding when evolving neural network controllers for Khepera robots in simulation. It was claimed that using such a developmental method, the genome need not necessarily increase in length whenever the number of neurons increased since no specific data relating to the presence or otherwise of neurons need to be stored in the genome, which will now be specified as part of the cell differentiation process rather than being directly encoded for in the genome. This cell differentiation system has subsequently been used to evolve only the morphologies of static 3D virtual organisms

(Eggenberger 1997) and more recently to grow the connectivity of a neural network for controlling a foveating retina of a real physical robot (Eggenberger, Gomez, and Pfeifer 2002).

Related work with wheeled robots have also shown promising results in robustness and the ability to cope with changing environments by evolving plastic individuals that are able to adapt both through evolution and lifetime learning (Floreano and Mondada 1998; Floreano and Urzelai 1998; Nolfi and Floreano 1999; Floreano and Urzelai 2000; Floreano and Urzelai 2001). A number of different reactive navigation behaviors were generated using evaluation functions that typically included different terms for rewarding speed, wall avoidance and straight-line motion combined into a single objective. Instead of evolving the synaptic weights, the learning rules governing the behavior of individual synapses were evolved when generating a neural network controller for Khepera robots. It was demonstrated that the evolved controllers were adaptive to changes in the environment due to their synaptic plasticity. Lifetime learning or ontogenetic adaptation has several adaptive functions within evolution: (1) allowing for individuals to adapt to fast-changing environmental conditions, (2) channelling information extracted from the environment to evolution, (3) helping to guide evolution, (4) reducing genotype length, and (5) maintaining genetic diversity (Nolfi and Floreano 1999). Learning and evolution were shown to be able to solve tasks that evolution alone could not solve. Performance increases were also noticed even when the learning tasks differed from the selection tasks. Learning individuals were thus better adapted to changing environments than non-learning individuals. Interaction between learning and evolution deeply altered both these processes in that learning enabled evolution to extract supervision information from the environment. In terms of generality, plastic-general individuals required less complex control systems compared to full-general individuals. Ontogenic adaptation has also been studied in a competitive co-evolutionary context of predator-prey simulations using Khepera robots (Floreano, Nolfi, and Mondada 2001).

Pure reactive agents that do not use any internal representation were shown

to be able to solve complex tasks through the use of sensory-motor coordination only (Nolfi 2002). By exploiting agent-environment interactions, these embodied artificial creatures were able to coordinate perception and action that enabled them to perform complex tasks without needing to react differently to the same sensory states in different contexts. The experiments involving physical agents were carried out using Khepera robots and neural networks weights were evolved for the control of the agents. Sensory-motor coordination allowed the robots to (1) select the most effective feedback, (2) simplify harder tasks, (3) exploit emergent behaviors, and (4) exploit environmental constraints. Pure reactive agents although effective were found to be sub-optimal in most conditions. As a remedy, it was suggested that more complex behaviors could be allowed to emerge through a simple process of adding internal representations to the existing reactive behaviors.

In a departure from classical connectionist models, Floreano, Schoeni, Caprari, and Blynel (2002) recently demonstrated the use of evolutionary spiking neurons for the control of an autonomous microbot. A single “spike” in a spiking neural network is a discrete binary event that simply encodes whether a stimulus is present or absent. Instead of using conventional non-linear, real-valued sigmoidal activation functions, the use of spiking neurons in neural circuits were shown to transfer easily to microcontrollers by virtue of their binary nature, which can be mapped onto low-level digital circuits using only a few logic operations such as AND and NOT. In an earlier study, it was shown that viable controllers were easier to evolve using spiking neurons than sigmoidal neurons for a vision-based navigation task of a Khepera robot (Floreano and Mattiussi 2001).

Comparatively small neural networks that utilized recurrent connections were shown to be capable of producing good obstacle avoidance and light-seeking behaviors in Khepera robots (Pasemann, Steinmetz, Hulse, and Lara 2001a; Pasemann, Steinmetz, Hulse, and Lara 2001b) using the *ENS*³ (Evolution of Neural Systems by Stochastic Synthesis) algorithm. A weighted sum of different speed and navigation objectives were combined into a single-objective function for the evaluation of evolved networks. The simplest evolved networks did not use any hidden

units and it was also demonstrated that larger networks were not necessarily more robust than smaller ones. In a related study, separately evolved neuromodules for obstacle avoidance and light-seeking behaviors were combined together to produce a single controller with both behaviors (Lara, Hulse, and Pasemann 2001) by evolving additional interface neurons and synapses for the interconnection between these two neuromodules. It was also shown in another related experiment that the evolved controllers were robust and performed well in both simulated and actual robots (Hulse, Lara, Pasemann, and Steinmetz 2001).

The control structures consisting of ANNs for a population of robots were evolved using a fully decentralized EA (Watson, Ficici, and Pollack 2002). The EE (Embodied Evolution) methodology was defined as conducting evolution in a group of real physical robots where evaluation, selection, and reproduction took place by and between robots in a distributed, asynchronous and autonomous manner. The robots were simple two-wheeled self-designed mobile agents with inter-agent communication capabilities. Evolved controllers outperformed hand-designed controllers for a phototaxis task.

A gaseous signalling mechanism was used in the GasNet algorithm for generating robot controllers in visual discrimination and navigation tasks (Husbands, Smith, Jakobi, and O'Shea 1998; Husbands, Philippides, Smith, and O'Shea 2001). The fitness of generated controllers was evaluated using a single function that combined the weighted sum of navigational scores. Although the neural networks using the gaseous signalling mechanisms could be evolved in fewer generations compared to neural networks that did not use these mechanisms, implying a less difficult search space in the former neural networks, all the standard random sampling measures used to discriminate between the two different search spaces failed to show any discernable differences between these evolutionary systems (Smith, Husbands, and O'Shea 2001b). Further analysis showed that the evolutionary robotics search space exhibited phases of neutral evolution (Smith, Husbands, and O'Shea 2001a). The population as a whole was shown to move significantly in the genotype space during such phases of neutrality and was not trapped at a local optimum in the fitness

landscape. However, no evidence could be found to indicate that neutral adaptation acted as a scaffolding for later transitions to higher fitness levels. As such, it was concluded that neutrality did not play any useful role in this particular evolutionary robotics search space.

It was later shown that the combined effects of increased neutrality and decreased ruggedness in evolutionary robotics search spaces allowed for greater evolvability (Smith, Philippides, Husbands, and O'Shea 2002). It was argued that phenotypic stability and genetic instability were prerequisites if successful evolution were to occur in an organism. Four different GasNet neural network models acting as controllers for simulated mobile robots in a shape discrimination task were implemented with varying degrees of redundancy and coupling to elucidate these effects. More recently, Smith, Husbands, Philippides, and O'Shea (2002) showed that the high success rates of GasNets neural networks in the visual discrimination task was due to temporal adaptivity and argued that this property is fundamental for the generation of adaptive behavior. Recent related work has also extended the family of GasNet neural networks to include more details of biological gaseous signalling mechanisms into two new versions called the *plexus* and *receptor* models, which were shown to be more evolvable than the earlier version of GasNet (Philippides, Husbands, Smith, and O'Shea 2002).

2.3.2 Legged Robots

The pioneering work of Beer and Gallagher (1992) documented the use of GA to evolve continuous-time recurrent neural networks for controlling the legged locomotion of a hexapod insect, although this study was conducted using a highly simplified physics model. It was shown in a later study that the evolved controllers could still perform the locomotion successfully when transferred to a real hexapod robot (Gallagher, Beer, Espenschied, and Quinn 1996). Related studies based on this simplified six-legged hexapod model have been conducted to investigate the evolution of neural network architectures rather than synaptic weights alone using a developmental scheme specified by the Simple Geometry Oriented Cellular Encoding

(SGOCE) algorithm (Kodjabachian and Meyer 1998a; Kodjabachian and Meyer 1998b).

Arnold (1997) investigated the generation of legged locomotion for four and six-legged virtual creatures using spectral synthesis (involving Fourier transforms). The control system was algorithm-based and did not make use of any sensory input information. Evolution was used only to tune the parameters of the algorithmic controllers and limb attributes for optimizing a single-objective function of maximizing horizontal distance achieved within a designated time period. In a later study by Reeve (1999), the control mechanism based on different models of neural networks for generating legged locomotion for a range of fixed morphology robots were evolved in simulation using a simple GA. It was found that simple single-termed fitness measures based on performance attributes such as speed was sufficient to generate the desired behavior and that more complex fitness measures relating to inner workings of neurons and joints were not advantageous. It was also found that higher-order neural networks were significantly better at performing the required tasks and that very densely connected controllers performed better than sparsely connected ones.

A dynamically-rearranging neural network (DRNN) was evolved to act as a controller for legged locomotion in a simulated biped robot (Fujii, Ishiguro, Aoki, and Eggenberger 2001). Generated controllers were assigned fitness values based on a single-objective function of horizontal movement achieved. Neuromodulators were used to dynamically change synaptic weights as well as network architecture by activating and blocking neurons and synapses. However, it was observed that many of the evolved controllers did not actually make use of the modifiable synaptic weights, in other words normal neural networks with fixed synaptic weights would have sufficed. Nonetheless, it was claimed that the DRNN would have exhibited superior performance in a changing environment due to their polymorphic characteristics although this was not investigated using the biped robot. In related work using a simulated quadruped robot, a DRNN was again evolved to act as a controller for legged locomotion (Otsu, Ishiguro, Fujii, Aoki, and Eggenberger 2001). It was claimed that the controllers generated were adaptive to changes in the environment

(retardant forces and uneven slopes) due to the neuromodulations present in the DRNN. However, as no analysis was provided on the actual dynamics of the neuromodulators during the legged locomotion of the quadruped, it remains unclear what roles these elements actually played towards the generation of a successful legged locomotion in the changing environments.

Both the controller and morphology of a biped robot were evolved using a GA with a simple single-objective fitness function based on horizontal distance travelled (Paul and Bongard 2001). It was claimed that the experiments produced the first reported results of stable bipedal locomotion achieved through the optimization of both controller and morphology. An interesting point to note was that only 6 out of the 60 evolutionary runs were successful in evolving a stable gait. The architecture of the recurrent neural networks that were used as the controllers remained fixed with only the synaptic weights being evolved. Also, only certain parameters of robot's morphology were allowed to be modified during evolution. A related study using similar biped robots where both the controller and morphology were co-evolved found that the inclusion of certain morphological parameters allowed for fitter individuals to be discovered by evolutionary search (Bongard and Paul 2001). It was shown that fitter individuals did not arise simply because a better morphology was found but rather the addition of morphological parameters into the genotype space allowed for extra-dimensional bypasses to be formed in the higher dimensional search space, thereby allowing the evolutionary search to find these fitter individuals. This phenomenon facilitated the connection of otherwise isolated adaptive peaks in the objective space, making it easier for the evolutionary search process to proceed smoothly from one adaptive peak to the next.

Central pattern generators (CPGs) were evolved as controllers for generating planar walking behaviors in two different physically simulated bipeds (Reil and Massey 2001). It was shown that using the appropriate mechanical construction, Hopfield neural network controllers and optimization through a GA with a single-objective distance-based fitness function, minimal bipedal locomotion can be achieved by CPGs that do not require sensor inputs. In the second more sophis-

ticated biped, incremental evolution was used where a weak stabilizing controller was used during the initial stages of evolution and later removed after a certain fitness level was achieved. The lower portions of the more sophisticated biped's legs were implemented as passive limbs to allow for a more anthropomorphic gait to emerge. Only 10% of the first biped's evolutionary runs produced successful controllers whereas 80% of the second biped's evolutionary runs produced successful controllers. However no analysis was given on whether the two search spaces differed significantly in terms of optimization difficulty.

In a related study, CPGs were again evolved to generate bipedal locomotion in a simulated robot in a real-time physics environment (Reil and Husbands 2002). Once more, it was shown that no sensory inputs were necessary to generate successful straight-line walking behavior although this was achieved only on a homogenous planar surface. It was suggested that the fitness landscape underlying the evolutionary search space of the recurrent ANN architecture is very smooth leading to successful evolution of controllers despite using only a very simple single-objective fitness function based on a combination of two objectives of maximizing distanced travelled from origin and minimizing occurrences of falling below a certain height threshold for the robot's center of gravity. However it was also reported that only 10% of the evolutionary runs resulted in stable controllers and that an additional fitness term that rewarded cyclic activity in the ANN was necessary to improve the success rate. The authors also noted a shortfall in the experimental setup in that the effect of network size on the efficiency of the approach was not studied. A number of important contributions of the evolutionary robotics approach to designing controllers for legged locomotion of artificial creatures were highlighted: (1) fully automated process that allows for changes or additions to the creature's structure to be accommodated very easily through re-evolution, (2) diversity of solutions, and (3) relatively cheap evolutionary computational requirements.

Real physical robots have also been used to study the generation of legged locomotion using EAs. Online evolution was used by Gruau (1997) and Gomi and Ide (1998) to generate static gaits for an octopod robot, and by Hornby and his co-

researchers to generate dynamic gaits for a Sony quadruped robot (Hornby, Fujita, Takamura, Yamamoto, and Hanagata 1999) as well as for the Sony entertainment robot dog AIBO (Hornby, Takamura, Yokono, Hanagata, Yamamoto, and Fujita 2000). The cellular encoding method of (Gruau 1994) was used to evolve not only the weights but also the architecture of the neural network controller for the octopod robot and also relied on interactive user assignment of fitness values rather than integrating a fully automated fitness assignment into the artificial evolutionary process (Gruau 1997). Jakobi (1998) utilized his “minimal simulation” method to also evolve gaits in simulation for the same octopod robot in order to reduce the time requirements of evolution on the real physical robot.

2.3.3 Abstract Robots

The neural network controller and visual morphology for visually guided behaviors in a specialized gantry robot was evolved using the SAGA algorithm (Harvey 1992) for a visual discrimination task (Harvey, Husbands, and Cliff 1994; Husbands, Harvey, Jakobi, Thompson, and Cliff 1997). Minimal vision systems and small networks were found to be sufficient for generating the required behaviors using a weighted sum combination of navigational scores as the evaluation function. Small population sizes and small number of generations were also sufficient for successfully evolving these controllers. A good choice of control system primitives were suggested as the main reason for the success of these evolutionary runs. Work has also been carried where only the morphology of a compound eye on an abstract robot was evolved while the neural network controller was kept fixed (Lichtensteiger and Eggenberger 1999). Cliff, Harvey, and Husbands (1993) have also conducted work on evolving both the visual morphology and recurrent neural network controller of a mobile robot. In a later related study, Cliff and Miller (1996) also evolved both visual morphologies and neural controllers of predator-prey agents in a competitive co-evolutionary environment although the simulation was only carried out in a 2D world. Using developmental methods, the bodies and controllers of autonomous 2D agents have also been co-evolved in a study by Dellaert and Beer (1994).

GP was used to evolve controllers for a robot with manually reconfigurable morphology called a random morphology robot (RM robot) (Dittrich, Skusa, Banzhaf, and Kantschik 1999). Evaluation of controller fitness was carried out using a single-objective function consisting of weighted scores for the robot's speed and distance. The fitness landscape was found to be highly dynamic because the robot moves around on a carpeted floor and hence encounters situations with different levels of difficulties arising from the directionality of individual carpet strands. It was shown that discrimination between good and bad individuals was hard during certain periods of the evolutionary process where the noise level was high. Hence it was proposed that reference individuals be employed to enable a differential fitness value to be calculated for evolving individuals in order to better capture the actual performance of individuals throughout the highly variable evolutionary periods. Nevertheless, it was later observed that there were periods where the fitness landscape oscillated, which created a problem for the proposed relative fitness methodology as well.

Lipson and Pollack (2000) combined both simulated and physical approaches for evolving simple robots composed of bars, actuators and artificial neurons for the single objective of maximizing horizontal distance moved. The authors claimed that to fully realize artificial life, autonomy must be achieved not only at the level of power and behavior but also at the levels of design and fabrication. They demonstrated this point in their experiments where artificial evolution was conducted to automatically design abstract robots that could perform locomotion in simulation and then the best virtual designs were fabricated into real robotic body parts using 3D thermoplastic solid printing techniques. The results from testing the physical versus the virtual robots showed that in one case, the distance travelled was almost identical while in the two other cases, the distances travelled were quite dissimilar although it was argued that the overall control and mechanics of the motion were still maintained when moving from simulation to reality.

2.3.4 Evolvable Hardware

Evolvable hardware circuits in the form of field programmable gate arrays (FPGAs) were utilized to evolve obstacle avoidance controllers for Khepera robots (Thompson 1995; Thompson 1997) and was claimed to be the first example of intrinsic hardware evolution (Harvey 1997), where every actual hardware specified during the evolutionary process was tested in situ rather than in simulation. Another series of studies also utilized FPGAs as evolvable controllers for producing visual tracking and obstacle avoidance behaviors in Khepera robots (Keymeulen, Durantez, Konaka, Kuniyoshi, and Higuchi 1996; Keymeulen, Iwata, Konaka, Suzuki, Kuniyoshi, and Higuchi 1998). Solutions generated were evaluated using a fitness function that took a weighted sum combination of two objectives of minimizing the robot-target distance and minimizing the number of steps required to complete the task. It has been argued that *true evolvable hardware* should allow for both control circuits and body plans to be evolved (Lund, Hallam, and Lee 1997). Such true evolvable hardware using a modified version of the Khepera robot with a reconfigurable auditory morphology was developed by Lund, Hallam, and Lee (1997) as a framework for studying the evolution of phonotaxis in crickets although no result from actual experimentation was reported.

2.4 Evolution of Morphology and Mind

The study of evolving physically situated and virtually embodied artificial creatures has been a hotbed of research in recent years. The availability and maturation of commercial-off-the-shelf physics engines coupled with the dramatic increase of personal computing power have encouraged widespread research into this intriguing field of artificial life (Taylor and Massey 2001). Not surprisingly, there were notably few significant advancements in this field since the pioneering work of Sims (1994a, 1994b) until very recently.

Research in this area generally falls into two categories: (1) the evolution of controllers only for creatures with fixed morphologies (Gritz and Hahn 1997;

Taylor 2000; Bongard and Pfeifer 2002; Mandik 2002), and (2) the evolution of both the creatures' morphologies and controllers simultaneously (Sims 1994a; Sims 1994b; Komosinski and Ulatowski 1999; Bongard and Paul 2000; Komosinski 2000; Komosinski and Rotaru-Varga 2000; Ray 2000; Bongard and Pfeifer 2001; Hornby and Pollack 2001a; Komosinski and Rotaru-Varga 2001; Taylor and Massey 2001; Bongard 2002b; Hornby and Pollack 2002). Some work has also been carried out in evolving only the morphology alone for static 3D virtual organisms (Eggenberger 1997) and evolving morphology with a fixed controller (Lichtensteiger and Eggenberger 1999).

The idea of using artificial evolutionary methods to automatically generate 3D embodied virtual creatures was first introduced by Karl Sims in 1994. A proprietary physics-based simulation system was implemented to evolve both the morphology and neural systems of virtual creatures using a directed graph grammar. Conventional artificial evolution was used to evolve specific behaviors such as swimming, walking, jumping and light-following (Sims 1994b) and a competitive co-evolutionary method was used to evolve creatures for resource acquisition (Sims 1994a). The fitness of evolved creatures was judged using simple single-objective functions such as speed and height achieved. Although highly interesting morphologies and behaviors were evolved, the applicability of the system for evolving real robots remained questionable as the physics specifications for synthesizing the creatures allowed for interpenetrating surfaces. While Sims required Connection Machine CM-5 parallel computers to conduct his artificial evolution, Taylor and Massey (2001) recently re-implemented Sims' work using only standard personal computers. Ray (2000) has also implemented a system highly reminiscent of Sims' system but rather than using a fixed fitness function within the artificial evolutionary system for selection, he relied upon aesthetic user selection. Both Taylor and Massey (2001) and Ray (2000) used a commercial-off-the-shelf physics engine called MathEngine, which is the predecessor to the physics engine known as Vortex used in this thesis (see Section 3.1.1).

MathEngine was also used to develop a 3D biomechanical simulation of

a salamander for a computational neuroethological experiment into the underlying neural circuits that generated the aquatic and terrestrial locomotion of real salamanders (Ijspeert 2000). Although no evolutionary results were reported, an algorithmic CPG controller was developed that allowed for realistic and life-like swimming and trotting gaits to be reproduced in the artificial salamander. The CPGs were shown to be stable enough to receive higher-level sensory input from vision modules to enable tracking and approach towards a moving target (Ijspeert and Arbib 2000). These studies stem from earlier work on evolving CPGs based on neural controllers that generate swimming gaits for a 2D lamprey (Ijspeert, Hallam, and Willshaw 1999; Ijspeert and Kodjabachian 1999) as well as for generating locomotion gaits for a 2D salamander (Ijspeert, Hallam, and Willshaw 1998; Ijspeert 1999; Ijspeert 2001). The fitness of evolved neural controllers was evaluated using single-objective functions based on either single or multiply-combined network output metrics.

GP has also been used to evolve controllers for a fixed morphology virtual creature (Gritz and Hahn 1997). In this study, the emphasis was on evolving different control programs for a 3D animated character as opposed to traditional “key-framing” techniques that involved human hand-designed frames used by most animators. It was claimed that the evolved controller produced fluid, physically and biologically plausible motions. An incremental approach where additional constraints were phased into the single-objective fitness function as the evolutionary optimization progressed was used to evolve the final desired behavior. This incremental methodology was adopted after it was found that a direct approach incorporating all the desired motion styles into the fitness function from the start severely restricted the evolvability of the system. Both evolutionary robotics and virtual embodied evolution techniques have also been extended to practical applications in the entertainment and edutainment industries (Taylor 2000; Grand 2001; Lund 2001). Commercially-based research conducted by Grand (2001) produced the artificial life game called *Creatures* in which owners could breed, nurture and evolve virtual organisms known as *Norns* on their personal computers and even exchange genetic material with other owners over the Internet. Techniques used included

user-guided behavior-based systems, user-guided evolutionary and co-evolutionary robotics as well as those including morphogenesis (Lund 2001) while it was suggested that to enable scaling to more complex behaviors required for characters in computer games, other techniques such as lifetime learning, virtual ecologies and evolution of behavioral primitives need to be considered (Taylor 2000).

Bongard and Paul (2000) investigated the relationship between morphological symmetry and locomotive efficiency by co-evolving the controller and morphology of virtual embodied organisms using a physically accurate simulation. A variable length GA based on the SAGA algorithm (Harvey 1992) was used to evolve the artificial creatures and a recurrent neural architecture was used to act as the creatures' controllers. Two single-objective fitness functions were designed to reward firstly a combination of locomotion distance and morphological symmetry, and secondly locomotion distance and morphological asymmetry. It was found that bilaterally symmetrical agents were favored by evolution in terms of locomotion capability. Although these experiments entailed two separate objectives of distance and symmetry, these objectives were combined into a single fitness evaluation function.

Bongard and Pfeifer (2002) also conducted experiments in which only the weights for fixed architecture recurrent neural controllers were evolved. 10 creatures with different but fixed morphologies were used to investigate the difficulties of evolving locomotion controllers for creatures with different body masses and number of legs. Using a single-objective fitness function that measured forward displacement, it was claimed that hexapedal agents were the easiest while worm-like agents were the hardest to evolve successful controllers. In a related study, artificial evolution was shown to automatically add more complex behaviors to simpler ones through the use of different sensor modalities (Bongard 2002a). Using a quadrupedal agent with two simulated chemical sensors, a lower-level chemotaxis behavior was shown, through a number of lesioning experiments, to provide a base for the generation of a higher-level forward locomotion behavior. Hidden units in the neural networks were also lesioned to demonstrate that over evolutionary time, some hidden units became specialized in processing certain input signals.

A similar but separate series of studies focused on the developmental processes associated with evolving virtual embodied organisms where both the controller and morphology were again being co-evolved (Bongard and Pfeifer 2001; Bongard 2002b). Using the designed system called *Artificial Ontogeny* (AO), artificial organisms were evolved to locomote and push boxes in which a standard GA using a single-objective evaluation function was augmented with a genetic model based on biological differential gene expression. As such, the genotype-to-phenotype morphogenesis allowed for changing pattern expressions similar to that found in genetic regulatory networks (GRNs). It was claimed that the AO system had high evolvability since the artificial evolutionary system was able to produce modular structures as well as dissociate between the genotypic and phenotypic complexities (Bongard and Pfeifer 2001). In the later study, Bongard (2002b) showed that the AO system was able to generate modular GRNs early during the evolutionary process which led to the successful generation of creatures with high parts count. The early appearance of modular GRNs was attributed to the high pleiotropy (co-regulation of genes) within the neurogenesis process and low pleiotropy between the neurogenesis and morphogenesis processes. However, a somewhat biased weighted sum fitness function that involved a “shaping” term, which explicitly rewarded organisms with number of body parts, sensors, motors and synapses, was used to encourage the early appearance of active agents during evolution.

Komosinski and Ulatowski (1999) developed a proprietary platform for studying the evolution of 3D physically simulated virtual creatures called *Framsticks*. The system allows for both directed as well as open-ended evolutionary runs to be conducted although published results have only documented experiments with directed evolution for behaviors such as walking and swimming (Komosinski and Ulatowski 1999; Komosinski 2000). An initial investigation into the design and use of more evolvable genotype representations for achieving open-ended evolution was discussed by Komosinski and Rotaru-Varga (2000). It was found in a later study that higher-level encodings that included either recurrent or developmental elements in the genotype representation allowed for more structured phenotypes to be gen-

erated, which in turn led to the appearance of fitter individuals for separate single-objective maximization tasks involving height and locomotion speed (Komosinski and Rotaru-Varga 2001). The *Framsticks* artificial evolutionary system has also been used to create a method for studying the taxonomy of evolved agents based on how dissimilar agents were in terms of their morphological geometry (Komosinski and Kubiak 2001). The results obtained from using these taxonomic measures on artificial organisms were later compared to the characteristics and properties of biological phylogenetic trees constructed for real organisms (Komosinski, Koczyk, and Kubiak 2001). The *Framsticks* creatures have also been used by Mandik (2002) to study the evolvability of mental representations where the neural controllers of fixed morphology agents were optimized using a combination of both human and artificial evolutionary design inputs for food-finding tasks in both walking and swimming creatures.

The emphasis of Hornby and Pollack (2001a) in their study of evolving both the controller and morphology of virtual creatures was also on the genotype encoding for achieving more complex designs. Using a developmental grammar based on Lindenmayer systems (L-systems), 3D agents with simple bars and actuators were evolved in a quasi-static virtual world which could physically simulate low momentum movements similar to that of Lipson and Pollack (2000). A weighted sum fitness function was utilized to optimize maximization of locomotion distance and minimization of occurrences where body parts were dragged on the ground. They showed that creatures evolved using generative encodings outperformed those evolved using non-generative encodings for a locomotion task by capturing useful design space biases while allowing large scale mutations to be performed viably, which in turn enabled the encapsulated and coordinated re-use of hierarchies of parts (Hornby and Pollack 2002). It was claimed that the morphologies of these generatively encoded creatures were more complex than those previously reported by Sims (1994b), Komosinski and Rotaru-Varga (2000) and Lipson and Pollack (2000), by virtue of having more parts in the morphology and more regularity in the overall design of the evolved creatures. In related work, oscillator controllers

similar to CPGs were used in place of neural network controllers for evolving both 2D (Hornby, Lipson, and Pollack 2001) and 3D virtual agents (Hornby and Pollack 2001b). The 2D agents were also successfully transferred to real physical robots (Hornby, Lipson, and Pollack 2001).

2.5 The Emergent Questions

As we have seen, the research into evolving artificial creatures have focused mainly on generating the desired behavior using single-objective fitness functions. These evaluation functions typically consist only of a single term for assigning the fitness of individuals generated (Sims 1994b; Arnold 1997; Gritz and Hahn 1997; Harvey, Husbands, Cliff, Thompson, and Jakobi 1997; Reeve 1999; Lipson and Pollack 2000; Fujii, Ishiguro, Aoki, and Eggenberger 2001; Komosinski and Rotaru-Varga 2001; Paul and Bongard 2001; Reil and Massey 2001; Bongard and Pfeifer 2002) or a combination of multiple terms into a single weighted objective when the desired behavior cannot be achieved with simpler single-termed functions (Lee, Hallam, and Lund 1996; Floreano and Mondada 1998; Husbands, Smith, Jakobi, and O'Shea 1998; Keymeulen, Iwata, Konaka, Suzuki, Kuniyoshi, and Higuchi 1998; Dittrich, Skusa, Banzhaf, and Kantschik 1999; Bongard and Paul 2000; Hornby and Pollack 2001a; Pasemann, Steinmetz, Hulse, and Lara 2001b; Bongard 2002b; Reil and Husbands 2002). It is highly surprising that a true multi-objective optimization approach involving optimization of explicitly distinct objectives have not been explored yet thus far for artificial creature evolution. Such an investigation might very well reveal significant advantages over standard single-objective EAs in terms of the evolutionary optimization process itself in addition to the possibility of generating greater varieties of creature morphologies and behaviors. We investigate this problem in Chapters 5, 6 and 8.

Although it was reported that more complex creatures could be evolved with certain artificial evolutionary systems (Hornby and Pollack 2001a; Komosinski and Rotaru-Varga 2001; Bongard 2002b), these claims were made simply based on

the fact that the artificial creatures had more moving parts or greater regularity in their morphology. Obviously such a trivial comparison leaves much to be desired since a millipede would be considered to be more complex than a human on both counts! Hence, it remains unclear how we can objectively compare between the complexities of artificially evolved creatures using more intuitive measures or methodologies. We attempt to tackle this problem in Chapter 7.

Additionally, apart from the work of Smith with wheeled robots (2001b, 2001a, 2002), there has been little effort invested in systematic explorations of the fitness landscape characteristics for evolving artificial creature controllers. There has been even less work conducted on characterizing the underlying search space difficulty when evolving controllers for legged artificial creatures. Only the work of Dittrich, Skusa, Banzhaf, and Kantschik (1999) with an abstract morphology robot has provided some empirical information regarding the evolutionary fitness landscape for a non-wheeled robot. Although a conjecture concerning the smoothness of the underlying search space for evolving CPG controllers for bipedal robots was postulated by Reil and Husbands (2002), no actual experimental results have been reported yet thus far with regards to testing this hypothesis. As such, very little is known at this stage concerning the fitness landscape and difficulty associated with evolutionary searching of controllers for legged artificial creatures. We attempt to tackle this problem in Chapter 4.

2.6 Chapter Summary

A literature review of the related fields of evolutionary robotics and evolution of embodied artificial life was presented in this chapter. The importance of situating and embodying the artificial agents within a physically-based world was first highlighted. A comprehensive survey of the various methods employed for evolving artificial creatures comprising of both real physical robots and simulated virtual agents was then given. Finally, the research questions emerging from this literature review were presented.