Chapter 6: First Steps Towards Synthetic Psychology

6.1 INTRODUCTION

The previous three chapters have provided a brief exposure to a variety of different types of models that can be found in psychology. We are now in a position to use this background knowledge to focus our attention on synthetic psychology. The purpose of this chapter is to provide a brief introduction to some of the basic properties that will be found in the synthetic approach. These properties will be introduced by considering how to build a toy robot that walks.

Even within one class of models, computer simulations, we saw that there can be a great deal of variety. In particular, some computer simulations are analytic in nature. For example, one creates a production system by taking a complicated phenomenon, breaking it down into its components, and using these components to construct the simulation. This analytic approach has been highly successful in psychology and in cognitive science. However, this approach is not the primary focus of this book.

We are instead primarily concerned with models that are synthetic in nature. An example of a computer simulation consistent with the synthetic approach is a connectionist network of the sort that was introduced in Chapter 5. When the synthetic approach is adopted, a set of basic building blocks is taken and is assembled into a working system. The question of interest is whether these basic components can be organized into a system that does something complicated, interesting, or surprising.

Why is this book focusing on the synthetic approach? One reason is that in modern cognitive science there is a growing interest in developing models from the synthetic perspective. A number of fairly recent simulation methodologies that are popular in cognitive science are essentially synthetic in nature. These methods include neural networks (e.g., Bechtel & Abrahamsen, 1991; Dawson, 1998), genetic algorithms (e.g., Holland, 1992; Mitchell, 1996), and artificial life (e.g., Langton, 1995; Levy, 1992).

A second reason for exploring the synthetic approach is that it melds quite nicely with a new tradition in robotics, artificial intelligence, and cognitive science. This new tradition is viewed as defining a new field, a field that has been associated with a variety of labels in recent years. These labels include behaviour-based robotics (Brooks, 1999), new artificial intelligence, based-based artificial intelligence, and embodied cognitive science (Pfeifer & Scheier, 1999). The embodied cognitive science movement is gaining popularity, and is challenging the traditional symbol-based conception of artificial intelligence and cognitive science along many of the same lines that were adopted by connectionist researchers in the early 1980s.

6.1.1 Synthetic Psychology Vs. Embodied Cognitive Science

Importantly, embodied cognitive science and synthetic psychology are not identical fields. Embodied cognitive science is a reaction against the traditional view that human beings as information processing systems "receive input from the environment (perception), process that information (thinking), and act upon the decision reached (behaviour). This corresponds to the so-called sense-think-act cycle" (Pfeifer & Scheier, 1999, p 37). This has also been called the sense-model-plan act framework (Brooks, 1999). The sense-think-act cycle, which is a fundamental characteristic of conventional cognitive science, is an assumption that the embodied approach considers to be fatally flawed.

One of the aims of embodied cognitive science is to replace the sense-think-act cycle with a principle of sensory-motor coordination (Pfeifer & Scheier, 1999), which might be construed as a sense-act cycle. The purpose of this change is to reduce, as much as possible, thinking -- the use of internal representations to mediate intelligence. What makes this a plausible move to consider is the possibility that if one situates an autonomous agent in the physical world in such a way that the agent can sense the world, then no internal representation of the world is necessary. “The realization was that the so-called central systems of intelligence – or core AI as it has been referred to more recently – was perhaps an unnecessary illusion, and that all the power of intelligence arose from the coupling of perception and actuation systems” (Brooks, 1999, p. viii).

The synthetic approach is an important component of this movement, because it opens the door to discovering behaviours that emerge from the interaction between an agent and its environment. Embodied cognitive scientists seek this kind of emergence because they do not want to explain complex behaviour by only appealing to internal mechanisms. Instead, “if we want to achieve wall-following behaviour, we should design not a module for wall-following within the agent, but instead basic processes that together, interacting with the environment, engender this desired behaviour” (Pfeifer & Scheier, 1999, p. 307). However, the synthetic approach is not equivalent to this embodied movement. For instance, and as we will see throughout this book, connectionist modeling can easily be construed as synthetic simulation (Pfeifer & Scheier, 1999, Chapter 5). Nevertheless, much of what is interesting about connectionist networks are the representational properties that stand between “sensation” and “action” (Dawson, 1998).

6.1.2 Overview: Synthesis, Emergence, Analysis

The purpose of this chapter is to introduce the key characteristics of the synthetic approach as it is used in synthetic psychology. These characteristics can be summarized with the acronym SEA, which stands for synthesis, emergence, and analysis. In my view, these are the three fundamental steps required for the synthetic approach to make contributions to psychology and to cognitive science.

In very general terms, synthetic psychology should proceed by carrying out these three steps in succession. First, a set of basic building blocks is used to synthesize a model. Second, the performance of the model is explored, with particular attention being paid to emergent properties. Third, the emergent properties are explained in a theory that accounts for them by appealing to internal mechanisms, to the environment, or to an interaction between the two.

For our first exposure to these three steps, this chapter describes a class activity that I have used in one of my graduate courses. This example as it stands is not particularly psychological – a characteristic that is unfortunately true of many of the phenomena modeled by embodied cognitive science. However, it provides a concrete example of the three components of SEA. In later chapters, we will use these foundations to build synthetic models that seem more relevant to higher-order psychological processes.

6.2 BUILDING A THOUGHTLESS WALKER

It has been my experience that when you try to teach modeling, students really benefit from hands-on work. So, when I was teaching a course on the synthetic approach in the fall of 2000, I thought that it was important to spend at least some time having students actually construct a working model. This was not merely to expose them to modeling per se. The activity that I had in mind was intended to expose them to the realization that a phenomenon that they took for granted was actually quite complicated.

One of the problems that I faced in doing this was that different students had enormous differences in their backgrounds of computer programming. This meant that the class activity couldn't involve hands-on computer simulation, because I was not really interested in spending a great deal of time teaching programming as part of this course. My solution was to have students build a walking robot from a particular kind of toy buildings set, K'NEX.

One week I brought two large containers of K'NEX materials into class, along with an assortment of text aids. The class was asked to build a robot that could walk forward. The text aids were used to provide inspiration, but I didn't give any specific instructions. The class was quite small (8 or 9 students), and they spent the first little while organizing themselves as a group, exploring K'NEX, and thinking about how they were going to approach this class problem.

By the end of class a week later, and only working during class time, the students had constructed a set of modular components that could be used to create a two-legged system, a four-legged system, or a six-legged system. Under certain conditions, described in more detail below, the students were successful in creating a robot that could walk the length of the classroom.

In the sections below, I will describe the next-generation of the robot built in the style that was created by the students. It represents a next-generation system only to the extent that I took the liberty of making some minor improvements to their original model. The original robot, constructed under conditions that weren't necessarily ideal, had a few flaws that needed to be corrected. Some of the flaws involved robot parts that were intended to be identical in each module, but were not. Some of the flaws were structural problems that required solutions involving parts, such as elastics, that were not pure K'NEX. The robot described below is built purely from K'NEX parts in the spirit of the robot that was constructed in class.

The sections below describe this robot as a project that could be built by the reader. It only mentions the parts and the properties of the final system. However, it is important to remember that when it was designed, students made explicit decisions to build in one way, and not in another. This was because they had many more materials available to them than those that are mentioned below. The reader of this section should keep in mind that other designs are easily possible, and might want to consider alternative approaches to building the robot if they decide to try to replicate the efforts of this class.

6.2.1 A Class Project

The purpose of this project is to build a robot that can walk at least a few steps forward independently. It is not required to be able to turn, or to avoid obstacles. One goal of building the system is to start to have some appreciation for some of the properties that are characteristics of walking systems.

6.2.2 Materials

The entire robot was constructed out of a large set of K'NEX building materials that my daughter and son had accumulated over the years. K'NEX is a toy building system comprised of rods that can be inserted into geometric connectors that hold the rods together to create larger structures. Rods and connectors are made from plastic, and are colour-coded to indicate length or shape. Figure 6-1 illustrates the types of rods and connectors used in the robot project. While structures built from K'NEX can be quite sturdy, there is a fair amount of "give" in these building materials. This turns out to be

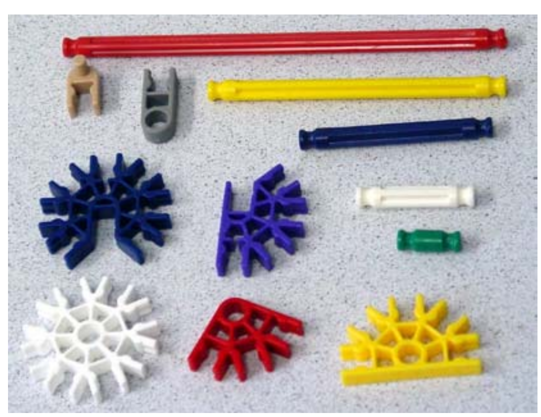

Figure 6-1. K'NEX parts used to build the robots.

This turns out to be advantageous in providing emergent walking behaviour in the robot.

In addition to the rods and connectors, the robot-building students also had available to them three identical motors that can be used to provide movement to K'NEX structures. As can be seen from Figure 6-2, these motors drive plastic gears that can rotate a K'NEX rod inserted as an axle. One of the advantages of using such motors for this project is that they too have a little bit of "give" in them, which was important for getting the robot to work. Having three of these motors was a luxury, as well as an indicator of how many raw materials the students had available. Usually K'NEX kits that use motors of this type only include one.

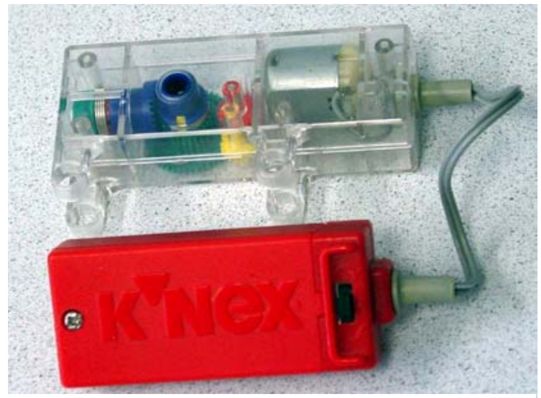

Figure 6-2. A K'NEX motor.

The only additional material required for building the robot was some literature that described a number of different robot projects, none of which use K'NEX. As we will see below, one important problem the students needed to solve was how to convert the rotating motion of the motors into a stepping action. They found one chapter in McComb (1987) that was particularly useful for providing a solution to this problem.

Table 6-1 provides a complete parts-list for building a single two-legged module. To build the entire robot, three of these modules must be created, and as a result three times the number of parts listed in the table are actually required.

Connectors Required For One Module	
Colour	**Number**
Blue	**12**
Purple	**4**
Red	**2**
White	**14**
Yellow	**8**
Beige	**6**
Grey	**4**
Rods Required For One Module	
Colour	**Number**
Red	**3**
Yellow	**4**
Blue	**7**
White	**18**
Green	**26**
Additional Materials	
1 K'NEX Motor	
4 white rods to connect module to another	

Table 6-1. The number of K'NEX parts required for one robot module.

6.3 STEP 1: SYNTHESIS

In adopting the synthetic approach, the first general step is to take a group of basic building blocks and assemble them into a working system. This contrasts with the analytic approach because the researcher does not start with a complete system, and decompose it into component parts or functions.

In the first phase of synthesizing the walking robot, the basic building blocks are the K'NEX parts listed in the table above. In the later phase, a more abstract sense of building block is adopted. This is because you can create larger walking systems by linking together smaller identical walking modules.

Section 6.3.4 provides detailed instructions on how to build the walking robot. Before these instructions are encountered, however, it is interesting to consider three basic design decisions that the students had to make in order to create a module that moved its legs in a stepping fashion.

6.3.1 From Rotation To Stepping

In order to create a walking robot, the students decided that the fundamental engineering problem to be solved was converting the rotation of an axel into a stepping motion. In making this decision, the students also made some progress in terms of organizing their work on the whole robot. They had decided to convert each motor into a system that would cause two legs to step. They divided themselves into three small groups of students, each working with one of the motors. When one group had some insight into solving a particular design problem, they communicated it to the other two groups.

The primary inspiration for converting rotation into stepping came from Figure 16-15 in McComb (1987). This figure demonstrated that if one attached a leg to the outside of a rotating wheel, and also permitted the leg to rotate freely at the point of attachment, then the rotation of the wheel would result in the leg being lifted up and pushed down in a stepping motion. This general principle is illustrated here. In this figure, the black circle represents a wheel that is being rotated about its centre in the direction shown by the blue arrow. The black line represents a robot leg. It is attached to the wheel at the point represented by the small circle, and this connection can freely rotate (as indicated by the small blue arrow). When the wheel is rotated (from A to B), the leg can be lifted up and placed down in a stepping motion.

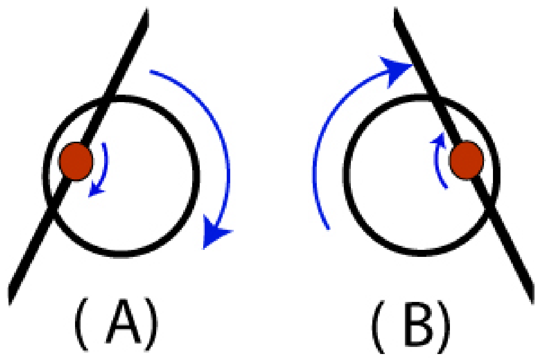

Figure 6-3. A general principle for converting rotation into stepping.

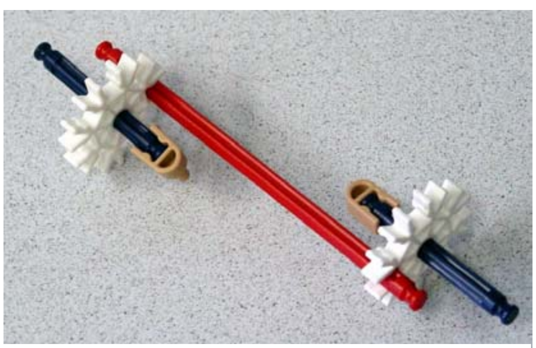

Figure 6-4. Arrangement of K'NEX parts to create an axle.

As is shown in Figure 6-4, the students exploited this design in creating the axle to be rotated by the K'NEX motor. The axle was a red rod. On one end of the rod, a white connector was attached sideways. At the other end of the rod, another white connector was attached sideways, but on the opposite side of the rod. These two white connectors represent wheels that would be rotated when the motor (which is not shown in the figure) rotates the red rod. The white connectors are placed on opposite sides of the rod so that the two legs moved by the motor would do so cooperatively (i.e., one would be stepping in front of the other). A blue rod is placed through the middle of each white rod, and a beige connector is attached to one end to keep it from falling out. When constructed in this way, the blue rod can freely rotate in the white connector, and is thus equivalent to the red point of connection in the preceding figure. If the leg is attached to the blue rod, then it is possible to make it "step" when the red axle rotates.

6.3.2 Balance

A second design issue to be faced by the students was the nature of the legs that were to be attached to the axles. The problem to be dealt with was this: the legs had to be constructed in such a way that a two-legged module would stand, even if the motor was not turned on. This was a problem because the "body" of the module – which was essentially the motor – was fairly heavy in relation to other components, and the legs were mounted in the middle of this body. The feet at the end of the legs had to be constructed so that the module would balance.

Figure 6-5. A robot leg.

Balance was achieved by attaching a fairly large and wide "foot" to a red rod that served as a leg. As can be seen from Figure 6-5, the foot was constructed from four white connectors and two yellow connectors held together very solidly with white and green rods. The robot's point of contact with the ground was the two yellow connectors.

The leg was attached to the axle by connecting the blue rod from the axle sideways to the lower white connector of the foot, as can be seen in Figure 6-6. This kind of connection is very firm, so that it was possible for the leg to push upwards with enough force to lift the "body" of the module. However, because the other end of the blue connecting rod could rotate in the white connector attached to the axle, the foot could still be moved in a "stepping" fashion.

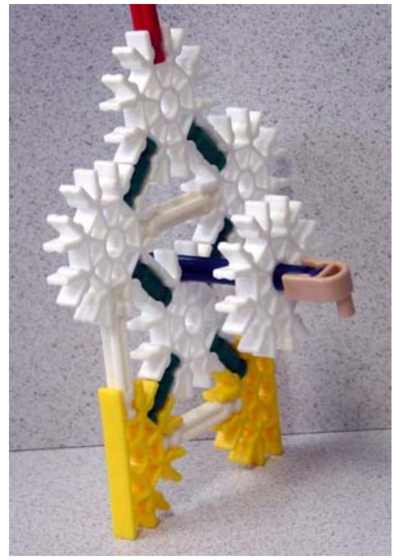

Figure 6-6. Attaching the leg to the axle.

6.3.3 Leg Support

In order to create a successful stepping motion, it is not sufficient to connect the foot to the axle. A support must also be provided to the top of the red leg, in order to prevent the entire leg from being rotated around the axle and hitting the ground. In other words, the support must be used to restrain the leg in such a way that it keeps pointing (roughly) up and down during movement.

In order to deal with this problem, the students designed a structure that was used to contain the motor, and was also used to loosely constrain the top of the leg. This structure is shown in Figure 6-7. One end of a blue rod is firmly attached to a purple connector at the top-middle of this "body". A grey connector is attached to the other end of the blue rod. A red leg is placed through the hole in the grey connector. This arrangement allows the leg to move fairly freely in an up and down motion, but prevents the top of the leg from being rotated downwards to interfere with any stepping movement.

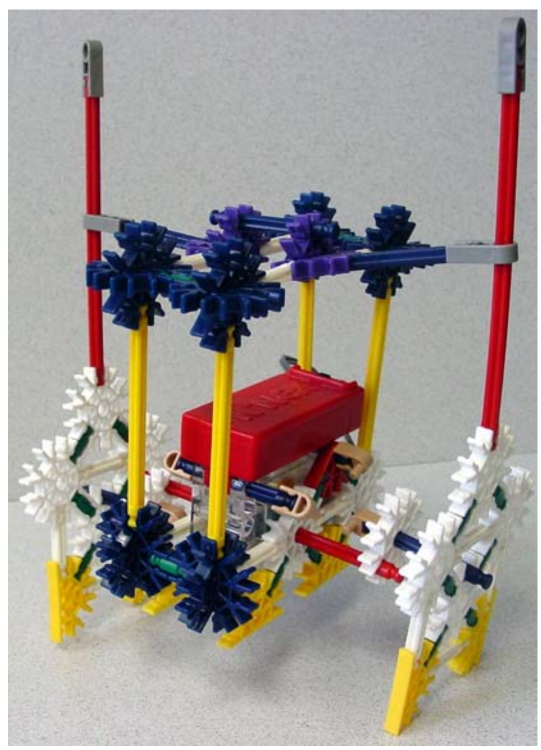

Figure 6-7. Leg support by the frame in a two-legged module.

Figure 6-7 also demonstrates some of the "give" in the K'NEX building materials that aids in making the movement of the system appear more lifelike. In this figure, the "foot" at the back of the figure is lower than the "foot" in front. If you look closely at the front leg, you will see that its red rod is slightly bent. Also, the blue rod/grey connector that is attached to the top of this leg is being pushed upwards in comparison to the same restraint that is attached to the other leg.

6.3.4 Detailed Instructions

The previous figure illustrates a complete two-legged module that the students used to explore walking in two-, four-, and six-legged systems. The reader might be interested in using K'NEX to build a similar robot to explore its behaviour, and to explore how its behaviour is affected by modifications of some of these design decisions. A set of web-based, step-by-step instructions is included with the CD that accompanies this book. For those looking at an electronic version of this text, the link to these instructions is included at the end of the chapter.

6.4 STEP 2: EMERGENCE

The first general step in synthetic psychology is to construct a working system. The second general step is to watch it work, paying attention to surprising or emergent properties. As we will see in Chapters 7 and 8, a practitioner of the synthetic approach expects that a system of simple components will generate far more interesting behaviour than would be expected, particularly when it is embedded in an interesting environment.

What is an emergent property? One way to think about emergence is in terms of the linear/nonlinear distinction that we explored when discussing models of data, mathematical models, and computer simulations. In a linear system, the behaviour of the whole system is exactly equal to the sum of the behaviours of its parts. If one understands the behaviours of all of the parts of a linear system, then this means that there should be no surprises when observing the behaviour of the system as a whole. In contrast, in a system in which the components interact nonlinearly, then surprises can emerge. "The hallmark of emergence is this sense of much coming from little" (Holland, 1998, p. 2).

Holland (1998) points out that while emergence is a ubiquitous phenomenon in the natural world, it is exceedingly complicated, and therefore defies definition. However, he argues that the scientific study of emergence is in a position to take advantage of some essential characteristics. First, emergence should be studied in systems that can be described as being governed by rules or laws. Second, an emergent phenomenon should be a pattern that is both recognizable and recurring. Third, theories of emergent phenomena will depend crucially upon modeling. Fourth, emergent phenomena will often be seen in systems that are either adaptive or dynamic over time. Fifth, "emergence usually involves patterns of interaction that persist despite a continual turnover in the constituents of the patterns" p. 7. These persistent patterns can be used as building blocks for larger systems. In other words, emergent phenomena will often be observed in systems that are organized hierarchically.

The walking robot that was constructed by the class is exceedingly simple. Nevertheless, when its behaviour was observed and manipulated, it exhibited many of these fundamental properties of emergence. In the sections below, we will consider observations made concerning three different versions of the robot: a two-legged robot, a four-legged robot, and a six-legged robot.

6.4.1 Two-Legged System

Figure 6-7 has already provided us with a glimpse of a complete two-legged system. In order to examine emergent behaviour, and to test its ability to walk, the students placed it on a table in the classroom and turned its motor on.

The behaviour of this system was interesting in some respects, but disappointing in others. When the motor was turned on, the robot began to sway back and forth in a surprisingly lifelike fashion. The "body" of the system also rotated back and forth. If one was to call each yellow connector at the base of the leg a "toe", then the stepping behaviour of this robot could be described as follows: three "toes" were always in contact with the table. Two were on the leg on one side of the robot; the third was the toe on the back of the other leg. In short, when the robot stepped, it raised the front "toe" of one leg, and then it raised the front "toe" of the other leg.

All of this behaviour was interesting, in the sense that it was quite a bit more complicated than the students predicted prior to turning the robot on. However, with all of this swaying, rocking, and toe lifting, one disappointing fact was obvious: the robot did not walk. It carried out all of this movement while staying in one place on the table. An mpeg file showing the behaviour of this robot is included with the CD that accompanies this book. For those looking at an electronic version of this text, the link to this video is included at the end of the chapter.

6.4.2 Four-Legged System

The next stage of exploring the walking behaviour of the robot was to take four additional white rods, and to connect two two-legged modules together to create a four-legged robot. When the modules were connected together, care was taken to ensure that both motors were pointed in the same direction. By convention, the rear of a module was the end where the motor switch, and the wire connecting the motor to the battery case, was found. Figure 6-8 illustrates the resulting system.

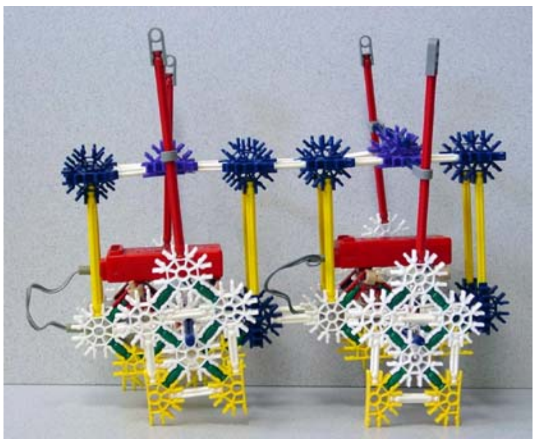

Figure 6-8. The four-legged robot.

K'NEX motors can be run in two different directions, clockwise and counter-clockwise, depending upon the setting of the motor's switch. In the first test of the four-legged robot, both motors were turned on in the same direction. The behaviour of the robot didn't depend on whether this direction was clockwise or counter-clockwise, and also was not affected by whether the motors were started at the same time or not, or by the starting positions of the two legs.

When the students drove the motors in the same direction, the robot as a whole began to sway and to rotate in a very similar fashion to the two-legged system. If one were to watch only one side of the robot, then one would typically see three "toes" in contact with the ground. Occasionally two "toes" would be seen in contact with the ground – one from each foot. Sometimes all four "toes" in contact with the ground. When these observations were made, the robot was not walking. All of its swaying movement was being done on the spot.

Interestingly, every so often the movement of the two component modules would become uncoordinated. When this happened, the "give" in the K'NEX motor became important. The students would hear a definite clicking sound as the gears of one of the motors jammed, and the rotation of one of the axles would cease for a short period (a second or less). Then, the motor that had stopped would start again, and at that moment the robot would lift one entire foot off the ground, and take a definite step forward. The extent of this forward movement was about 1 or 2 centimetres. Unfortunately, when the next step occurred, it was often in the opposite direction! So, when walking did occur, it was forward, then backward, one laboured step at a time.

In short, this first test of the four-legged system led to results that were essentially the same as the results observed in the two-legged system: there was lots of robot movement, but essentially no walking. However, every now and again a definite step would emerge, suggesting that the robot's design was on the right track.

From observing the first test of this robot, it appeared that the coordination between the two leg modules was critical. For the most part, the two modules were in step, and as a result the robot did not walk. Walking only appeared when the stresses on the robot's legs caused a disruption between the coordination of the two legs.

This observation led to a simple manipulation that had dramatic results. If walking in this system required that the two leg modules be out of synch, then perhaps it would walk better if the two motors were run in opposite directions. At face value, this prediction is counterintuitive, because one would expect that if walking were to be achieved, then the two sets of legs would have to be moving in the same direction. However, the students could quickly test their hypothesis simply by setting the two motor switches in opposite directions.

When this second test was conducted, the results were much more encouraging. At first, the robot motors protested loudly – angry clicks were heard from both. However, after a few moments of this, the clicks were heard less frequently, and the robot did begin to walk forward. When it was walking optimally, it would lift one rear foot completely off of the table's surface, and at the same time lift the front foot on the opposite side. It would then move about 2 cm. In some instances, the motors would lose this nice walking coordination, the stepping behaviour would attenuate, and the robot would either slow down or stand still. This only lasted for a moment though – when the robot was in this state, the motors would start clicking back and forth again, and then the robot would begin to step forward.

An mpeg file showing the behaviour of the four-legged robot is included with this book. It shows the difference between the behaviour of the system with the motors running the same way, and the behaviour with the motors reversed.

6.4.3 Six-Legged System

The final exploration of walking by the students involved studying the behaviour of a six-legged robot. This robot was created by connecting a third two-legged module to the four-legged robot. Once again, care was taken to ensure that all three motors pointed in the same direction. The resulting robot is presented in Figure 6-9.

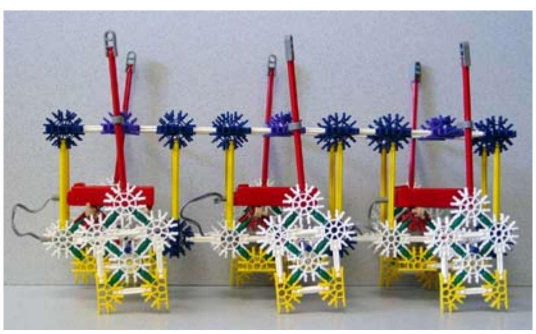

Figure 6-9. The six-legged robot.

In the first test of the six-legged robot, the behaviour of the system was very similar to that of the four-legged robot. The motors would complain when they were first all started in the same direction. The robot would begin to sway back and forth, and rotate a bit towards the left and right. Shortly thereafter, the three different motors would be coordinated in a pattern in which all three legs on each side of the robot moved in synch. In this configuration, the robot essentially swayed back and forth in one place, and the closest that it came to stepping would be raising only one "toe" of any of its legs off the table's surface.

On occasion, because the motors were running independently, this relatively stable configuration was disrupted. One or more of the motors would click, the axle would momentarily stop, and suddenly one of the robots legs would lift completely off the surface. A similar stepping motion would then be initiated in one or both of the other modules, and the robot would take a fairly large step (in the order of 5 cm) in one direction. Shortly afterwards, the system would again stablize into the configuration in which the legs on each side were in synch, and the machine swayed back and forth on the spot. As was the case with the four-legged robot, when the next stepping action occurred, it was often in the direction opposite to the last step taken.

The second experiment with this robot was conducted by trying to affect the coordination of the three leg modules by manipulating the direction of the motors. In particular, the motor driving the middle pair of legs was set to run in a direction opposite to that being run by the other two motors.

When the motors were set in these directions, the robot walked quite effectively. After an initial period of competition between the three motors, all three modules coordinated themselves into an arrangement in which each module raised one "foot" entirely off the table surface. In general, the three modules coordinated themselves in such a way that walking was achieved by resting the weight of the robot on one triangle of "feet" while the second triangle of "feet" was stepped forwards. Each triangle was defined by a front and rear leg on the same side of the robot, accompanied by the middle leg on the other side of the robot.

An mpeg file showing the behaviour of the six-legged robot is included with this book. It shows the difference between the behaviour of the system with the three motors running the same way, and the behaviour with the middle motor running in the opposite direction.

6.4.4 Emergence And Surprise

Let us now briefly summarize the main results of this phase of working with the K'NEX robots. One nice aspect of this demonstration project is that it provides an excellent example of emergence. By themselves, none of the two-legged modules built by the students were capable of walking. However, if two or three of these non-walking modules were coupled together, then walking was possible. It is clear that the walking behaviour emerged from an interaction between the modules.

A second point to be made from this demonstration concerns the notion of surprise. In particular, when reliable walking behaviour was observed in a robot, this was only achieved when the motors of adjacent two-legged modules were turning in opposite directions. This finding was counterintuitive, because at the outset it was natural to expect that all of the motors needed to turn in the same direction to get the machine stepping forwards. The robot project provides a very nice example of the possibility for surprise in a synthetic project.

Nevertheless, demonstrating that the synthetic approach can generate surprise is a fairly trivial result, and as such does not mark the end of the research program. "It is true that surprise, occasioned by the antics of a rule-based system, is often a useful psychological guide, directing attention to emergent phenomena. However, I do not look upon surprise as an essential element in staking out the territory" (Holland, 1998, p. 5). The lasting value of surprise in the synthetic framework occurs when it directs attention to emergent behaviours that can be explained by appealing to properties of system components. For this reason, after a system has been synthesized, and after emergent phenomena have been observed in its workings, the researcher must step back and analyze in an attempt to explain their creation.

6.5 STEP 3: ANALYSIS

One of the fundamental characteristics of the synthetic approach is the assembly of components into a working system that exhibits surprising, emergent behaviour. The class project that we have been describing has provided a concrete example of this.

A second, less explicit, characteristic of much synthetic research is the assumption that it leads more directly or more easily to explanations than does analytic research. Chapters 7 and 8 will provide more information about this assumption. For the time being, let it suffice to say that it is quite natural to assume that if you build a system, and engineer it out of parts whose workings you understand, then you should be in a position to explain the mechanisms from which surprising regularities emerge.

We will see that this assumption is not correct. In many of the examples that we will consider, building a system, and observing surprises in it, is pretty easy. The difficult – and interesting – work starts when an attempt is made to generate theories of regularities that emerge from what we synthesize. “Understanding the origin of these regularities, and relating them to one another, offers our best hope of comprehending emergent phenomena in complex systems. The crucial step is to extract the regularities from incidental and irrelevant details” (Holland, 1998, p. 4). A good deal of analysis is required to carry this crucial step out.

The walking robots that we have been describing in this chapter are not particularly sophisticated machines, and were not designed to provide deep insights into the nature of locomotion. Nevertheless, it is instructive to analyze aspects of their behaviours, because even these simple machines reveal some very interesting properties.

6.5.1 Emergence And The Thoughtless Walker

For a first pass at analyzing the walking robots, it is instructive to consider their behaviour in terms of the criteria for emergence that have been proposed by Holland (1998). One reason for doing this is because it helps to support the claim that the walking behaviour of the robots really is emergent. A second reason for doing this is because it draws our attention to a number of different properties of these robots. These properties demonstrate that even with these simple toy components, the behaviour of the robots is complicated and interesting.

6.5.1.1 Recognizable, Recurring Patterns

One of the criteria proposed by Holland (1998) as being necessary for emergence is the discovery of recognizable and recurring patterns. Emergent behaviours can’t simply be those that are rare and surprising; they have to be results that are replicable. Are the behaviours in our robots of this type?

The moment-by-moment behaviour of all of the robots is quite complicated, and a detailed classification of the behaviours would likely require a detailed, frame-by-frame analysis of video images of their performance. However, even a casual observation of their movements suggests that there are two general states that the robots “prefer” to be in.

The first state is one in which as many robot “toes” as possible are in contact with the ground. In this state, the motors cause the robot to sway from side to side, and to turn back and forth, but the robot does not step forward. Usually, when a foot moves, it only is a slight movement that causes only one of its two “toes” to be raised off the ground. In multi-module robots, this state is associated with all of the legs on the same side of the robot moving together. The two-legged robot is always in this state.

The second state is one in which a robot is actually walking forward. In this state, more than one leg is lifted completely off the ground, usually at roughly the same time. As the robot steps forward, it still sways and turns, but not to the degree seen in the other state. For the robot to be in this state, there must be definite coordination between different two-leg modules. In particular, modules that are connected to each other are coordinated in a “diagonal” fashion: if one module is lifting the left leg when the robot is in this state, then any attached modules will be lifting the right leg.

The robot is not only seen in these two states. However, other robot states appear to be quite transitory. They occur for fairly brief periods of time as the robot changes from one of the above states to the other. These transitions appear to be more stressful on a robot’s structure than either of the two states described above. This claim is supported by the fact that it is during these transitions that the motors stop functioning properly, grinding their gears with a distinctive clicking sound, and failing to rotate the red axle. In order to determine whether any of these transitory behaviours represent recurrent patterns would require detailed analysis (e.g., of slow motion video) of the robot.

When all of the motors are turning in the same direction, the robots are much more likely to be in the swaying state than in the walking state. Every 15 or 20 seconds, there will be a brief transition into the walking state (for one step), followed by a transition back into the swaying state. This situation is reversed when adjacent motors are turning in the opposite direction. In this case, the robots are much more likely to be in the walking state, except that every 15 or 20 seconds there is a brief transition into the swaying state, almost immediately followed by a transition back into the walking state.

6.5.1.2 Rule Governed System

Holland (1998) suggests that a second criterion for the scientific study of emergence is that it occurs in a rule governed system. "Emergent phenomena also occur in domains for which we presently have few accepted rules; ethical systems, the evolution of nations, and the spread of ideas come to mind" (p. 3). Holland suggests that an understanding of emergence in these domains will have to wait until we have a better understanding of the laws that govern them.

The robots that we have been describing are governed by laws, but not in the usual sense that comes to mind in psychology or cognitive science. Usually, the phrase "rule governed" in cognitive science immediately brings to mind a system that is controlled by a computer program. Furthermore, the computer program is usually thought be of a classical or symbolic type, such as the production system that was discussed in Chapter 5. However, it is obvious that no such program is responsible for the behaviour of the robots. I call them thoughtless walkers because they have not capacity to use symbolic representations to control their actions.

Instead, the robots are governed completely by the laws of physics. The motors are supplying kinetic energy that causes robot parts to move, and the movement of these parts generate forces against the surface upon which robot rests. The surface reflects forces back through the robot, which results in other emergent behaviours, such as the side-to-side swaying of the robot body. As we will see below, a great deal of analytic research on the locomotion of multi-legged animals proceeds by analyzing the distribution of forces through the walking system.

Describing the robots as being governed by the laws of physics leads to an interesting speculation that would require detailed physical analysis to validate. Each robot represents a physical system that holds kinetic energy, and is subject to a variety of forces. We could imagine measuring a robot in such a way that we could come up with a single number that represents its total stress or energy at any given time. My suspicion is that the two main states of the robot that were described above represent low-energy configurations. When a robot is in either of these two states, it is under the least amount of stress that it can be in when its motors are running. When the physical situation changes – for instance, when forces get redistributed because different motors are out of sequence – a robot moves to a higher energy state. When it is in such a state, it attempts to distribute forces again in such a way that the overall energy is again reduced. This results in the transitory behaviour, and the accompanying complaints from one or more of the motors. From this perspective, the swaying state might represent the least energy state for the robot when its motors are running, because this state is the easiest to produce. The walking state might represent a higher energy state (which could explain why walking has unexpectedly hard to produce). However, the energy of the walking state is still lower than any of the transition states. When the motors are running in the opposite directions, the robot is unable to reach the ideal swaying state, and instead has a preference for the next best configuration – walking. In Part II of this book, we will see examples of artificial neural networks whose behaviour is frequently explained by appealing to the physical analogy of seeking a least-energy state.

6.5.1.3 Dynamic System

Holland (1998) proposes that emergent phenomena are to be expected when the laws governing a system are invariant, but the system components governed by the laws are changing or dynamic. The thoughtless walkers that we have been discussing are obviously dynamic systems, because they are built from parts that are designed to move.

However, these robots are also dynamic in a subtler and more interesting sense. We have already described the robots in terms of two general lower-energy states, and have pointed out that both of these states depend upon a particular type of coordination between modules controlling different pairs of legs. We have also pointed out that there is no central computer that runs a program that coordinates the different modules. How, then, is coordination between leg modules possible?

The answer to this question is that different modules communicate to one another, but not in the symbolic fashion that is typically thought of in psychology and cognitive science. A different kind of communication is enabled by the dynamic nature of a robot's parts, and of the forces at play through its structure. When two modules become uncoordinated (e.g., because they are rotating a slightly different speeds), a robot's balance is altered in such a way that one module can run easily, but another cannot. In other words, changes in the physical configuration of the robot could be described as one module communicating to another that it is becoming uncoordinated. This message is communicated by changing the forces in the robot in such a way that one of the motors actually stops for a moment, until forces change again in some fashion that permits the motor to resume turning. In other words, even though these robots are thoughtless walkers, they can still be described as information processors.

6.5.1.4 Adaptive System

"The possibilities for emergence are compounded when the elements of the system include some capacity, however elementary, for adaptation or learning" (Holland, 1998, p. 5). The thoughtless nature of the robots that we have been considering precludes most of the possibility for learning. While this is a limitation of these robots (in terms of their generating theories about locomotion), such limitations are not surprising. After all, the robots are simply toys that are being used in a demonstration to reveal some of the general characteristics of the synthetic approach.

Nevertheless, if one were interested in exploring the properties of these robots in more detail, then there is a possibility for adaptation that could be explored with more detailed analyses than those we have reported above. It was mentioned earlier, and illustrated in Figure 6-7, that one of the advantages of the K'NEX components was the "give" in many of the components. Several of the parts of the robot structure are firm, but flexible. For example, rods can bend, can rotate within the joint of a connector, and can also be rotated a bit in the joint, without the overall structure breaking apart. It would be interesting to determine whether the physical structure of a thoughtless walker changed, because of the forces that the robot is subject to, in such a way that a physical configuration of a rod that was seen early in an experiment was never seen later. This might be evident if there were fewer transition periods between lower-energy states after a robot had been operating for a while. If indeed forces acting upon the robot adjusted its physical structure in this fashion, then this would be an example of elementary learning in a thoughtless system.

6.5.1.5 Persistent Patterns, Changing Components

Holland (1998) points out that emergent phenomena frequently exhibit a dynamic, hierarchical organization. At one level of analysis, parts of a system might be changing very frequently. At a broader level of analysis, though, the system might exhibit stable regularities. "A simple example is the standing wave in front of a rock in a white-water river. The water molecules making up the wave change instant by instant, but the wave persists as long as the rock is there and the water flows" (p. 7).

When the robots are successful in walking, they clearly exhibit this kind of hierarchical organization. For example, walking in the six-legged robot can broadly be described as the successive placing of a triangular configuration of legs onto the ground. However, the legs that make up this stable triangle change from step to step. A more detailed analysis of a robot's behaviour would probably provide many more examples of this. For instance, given all of the movement in a robot, and all of the "give" in its components, it would not be surprising that a wide diversity of physical configurations of robot parts could all be classified as a step.

6.5.2 Comparison To Biological Walking

The previous section has shown that there are a number of interesting and surprising emergent properties in the walking robots that we have constructed. This illustrates one of the main advantages of the synthetic approach. We could have taken a complicated phenomenon and analyzed it into putative component functions. Instead, we took a very simple set of building materials and a very general construction goal and were able to create a system that delivered several properties that were not explicitly intended.

To my mind, the robot example demonstrates another important advantage of the synthetic approach. By observing the regularities in the behavior of the working system, we started to learn important facts about walking in general. In many cases, the synthetic approach will provide us with insights into the problems that are being solved by the system is that we build, and these general insights will often be more important than a specific account of how a particular system works. Synthetic models provide a medium in which to explore phenomena. This medium can be so rich that one can learn a great deal by exploring the properties of models for their own sake.

A complementary approach, though, is to consider of the system in the context of other types of knowledge. For instance, when a model is being analyzed one fruitful approach is to relate its observe properties to known properties of other systems. In the case of our current demonstration, we could attempt to do this by relating characteristics of our walking robots to knowledge that other researchers have collected in their study of animal locomotion.

6.5.2.1 Lifelike Motion

The claim that was made earlier about the appearance of the walking robots was that their movement appear to be " lifelike". What exactly is meant by this claim? What kind of evidence can be cited to support this position? Research on animal locomotion can provide some answers to these questions, and can also provide some guidance about what properties of the walking robots deserve our attention.

At first glance, we might be tempted to think of legs as being kind of wheel. If legs functioned exactly like wheels, then movement would be uniform. However, analyses of the locomotion of many different animals have shown quite clearly that movement is not uniform at all. Legs are not wheels; a different metaphor is required to model actions like walking or running.

During a slow motion like walking, a better mechanical metaphor is an inverted pendulum (Dickinson et al., 2000). In this metaphor, the pendulum's cable becomes a rigid leg that is attached to a body of mass. When the leg is used for walking, the mass is vaulted over the leg, as a shown in Figure 6-10. In the first half of this movement, kinetic energy is transformed into gravitational potential energy. In the second half of this movement, when the body descends, this potential energy is partially recovered as kinetic energy. One consequence of this cycling between kinetic and gravitational energy is that the body of the animal decelerates in the first half of the movement, and accelerates in the second half of the movement. During a faster motion like running, the rigid leg of the inverted pendulum is better viewed as a spring, kinetic and gravitational potential energies are stored as elastic energy, and the system bounces as if it were on a Pogo stick cycling between breaking and propulsive phases. Again, the running system does not move uniformly. Instead it accelerates and decelerates with every step. Interestingly, these two metaphors can be applied to describe the locomotion of bipedal, quadrapedal, and polypedal organisms (Blickhan & Full, 1993).

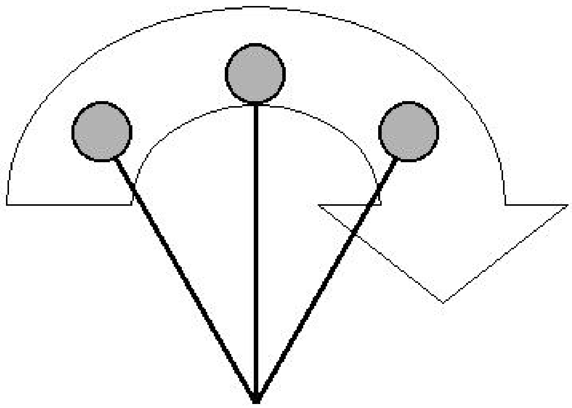

Figure 6-10. The inverted pendulum model of walking. The body (grey circle) is vaulting over a rigid leg in the direction of the arrow.

One in respect in which the walking robots are lifelike is that their forward movement is not uniform. Even though the motors work by providing a uniform rotation of an axle, when his motion is converted into a step, the legs work like pendulums. When legs are being lifted up, the robots forward movement is very slow because the feet that are in contact with the ground are in the process of converting kinetic energy into gravitational potential energy. When the legs are being placed down, the robot lunges ahead noticeably faster. This is because the center of mass over each supporting leg is descending and accelerating.

Animal locomotion research points to a second fashion in which the walking robots are lifelike. We saw earlier that getting the robot to walk depended heavily upon leg coordination. For example, the six-legged robot would only walk forward when its motors ran in such a way that the six legs were coordinated to act like two sets of tripods. Walking was accomplished by having one tripod serve the function of the rigid legs of an inverted pendulum, while the other tripod was moved ahead. This kind of leg coordination is commonly seen in the walking of the six-legged organisms such as insects (Dickinson et al., 2000).

A third respect in which the walking of our robots was lifelike involves the many movements of their bodies that were not at first glance directly related to walking per se. In particular, the bodies of all of the robots swayed back and forth very noticeably, in the front of the robust rotated between the left and right.

Interestingly, these kinds of movements are becoming a more interest to researchers to analyze animal locomotion. The legs of sprawled-posture animals, such as insects and crabs, generate substantial lateral and the forces (Dickinson et al., 2000). These forces are orthogonal to the direction of motion. Analyses of these forces suggest that elastic energy storage and recovery may occur within the horizontal plane. “By pushing laterally, legs create a more robust gate that can be passively self-stabilizing as the animal changes the, moves over uneven ground, or is not to skew by uneven terrain” (p. 101).

6.5.2.2 Control With No Brain

One of the predominant themes in the study of animal locomotion is the integration of many different systems, both neural and mechanical. “An integrative approach to locomotion focuses on the interactions between the muscular, skeletal, nervous, respiratory, and circulatory systems ” (Dickinson et al., 2000, p. 100). Researchers are not only interested in determining how each individual component of locomotion system works, but also how all of the components function together as an integrated system.

The integration of motor, sensory, and control systems is also evident in behavior-based robotics. Consider Genghis, a six-legged robot built in 1988 by robotocist Rodney Brooks (Brooks, 1989; reprinted in Brooks, 1999). Two motors drive each leg, one for swinging it back and forth, the other for lifting it up and down. A simple walk is achieved by manipulating the behavior of all of the motors in the robot, on the basis of sensing different positional characteristics of its legs. For example, one reflex notices whenever a leg is not down, and attempts to bring the leg down by turning one motor on in the appropriate fashion. A second reflex notices if any one of the legs happens to move forward for some reason. When this is detected, all of the legs will receive a series of messages that caused them to move backwards slightly. Other reflexes will advance a leg forward when it is noticed that the leg is raised, and will raise legs under appropriate conditions. The combination of these sensory measurements and motor signals lead to a very robust emergent walking behavior in the robot.

We have already pointed out that no sensory control system has been built into our thoughtless walkers. To the extent that there is control or coordination between different two-legged modules, this is mediated completely by the transmission of physical forces through the robot structure.

Recent work by researchers on animal locomotion has explored the capabilities of such thoughtless systems. Kubow and Full (1999) simulated a walking cockroach in which there was no feedback from the equivalent of neural reflexes. The only feedback in the model resulted from the musculoskeletal properties of the cockroach legs that were included in the simulation. When walking was simulated, and the forward movement of the model was perturbed by external forces, the model was self-stabilizing. Depending on the type of perturbation, the model was able to recover in one or more steps. “Essentially, control algorithms can be embedded in the form of the model itself. Control results from information being transmitted through mechanical arrangements. Perturbations change the translation and/or rotation of the body that consequently provide ‘mechanical feedback’ by altering legged moment arms” (p. 858). This is exactly the kind of coordination that we encountered when examining the conditions under which our robots were able to walk.

6.5.2.3 Limitations And Future Explorations

Animal walking can be very complicated. Locomotion is required to accomplish many different goals, and each goal might be achieved by a completely different gait in the same animal. A cockroach that walks slowly coordinates six legs in a fashion similar to the six leg and robot described earlier; when fleeing at a speed of fifty body lengths per second the same cockroach runs on only two legs (Full & Tu, 1991). The ecological roles of different types of locomotion are also reflected in the structure and function of different anatomical parts. “Forty percent of the body mass of the shrimp is devoted to the large, tasty abdominal muscles that produce a powerful tale flick during rare, but critical, escape behaviors” (Dickinson et al., 2000, p. 102). Animals to move in the real world are subject to a bewildering variety of different forces. All of these factors contribute to the view that animal locomotion requires the integration of multiple sensory, motor, and control systems.

In comparison to biological systems, the thoughtless walkers described in this chapter are very simple. Near only designed to step forward. They cannot turn, change gate to achieve different goals, or manipulate stepsize to deal with encountered obstacles. They do not have any neural or sensory control systems. Nevertheless, we have seen that even these exceedingly simple toy robots as many interesting emergent properties that are relevant to the scientific study of animal locomotion. They illustrate one of the main reasons that there is a growing interest in the synthetic approach: very simple components can be used to build systems that generate a far richer set of emergent properties then could have been predicted at the outset.

6.6 ISSUES CONCERNING SYNTHETIC PSYCHOLOGY

In this chapter, I have proposed a general approach for conducting synthetic psychology. This approach can be represented with the acronym SEA that stands for synthesis, emergence and analysis. In the first stage, a researcher will build a working system from known components.

In the second stage, a researcher will observe the actions of this system, looking for emergent phenomena. In the third stage, a researcher will perform an analysis of the working system in an attempt to explain the emergent phenomena by appealing to properties of the system, its environment, or the interaction between them. A simplified example of this approach was illustrated by the construction, observation, and analysis of walking robots built from K'NEX parts.

The simplified nature of this example can be used to raise, but not adequately deal with, a few issues related to the synthetic approach. Let us briefly mention these issues, and then confront them in more detail in subsequent chapters.

One issue concerns the sequential nature of applying the steps in SEA. As portrayed in this chapter, each step is done independently of the other, and there is a definite sequence of steps to be carried out. This portrayal does not do justice to the problem-solving practices of my robot-building students, and does not completely reflect how synthetic psychology is conducted in practice. In particular, in building the robots the students moved back and forth between synthesis, emergence, and analysis. They would attempt to solve a problem in one fashion, observe the system to see if the problem was solved, and if the problem remained, then they would analyze the situation to see if they could come up with an alternative solution. Part II of this book, with the hands-on demonstrations of connectionism as synthetic psychology, will hopefully provide a more accurate reflection of actual research practices.

A second issue concerns the advantages of the synthetic approach. While this chapter has attempted to illustrate how one might conduct a synthetic research program, it hasn't made a strong case for *why* this program would be conducted. Chapter 7 addresses this issue with an historical overview of the synthetic reaction against analytic research. It will make the argument that the synthetic approach does have many advantages. But it will also make the argument – hinted at by the final component of SEA – that synthetic research cannot be performed without also performing analysis.

A third issue concerns the domain of "synthetic psychology" and its relation to embodied cognitive science, including the reaction against the sense-think-act cycle. While the robots described in this chapter have given an example of viewing walking synthetically, they certainly do not qualify as being psychological models. The rejection of the sense-think-act cycle is explicit in their thoughtless nature. But how can such models be psychological? In Chapter 8, I am going to make a stronger argument that the synthetic approach can be conduced without rejecting the sense-think-act cycle. My position is that synthetic models that have representational properties are the true products of synthetic psychology.

6.7 INSTRUCTIONS AND MOVIES

The reader might be interested in getting a hold of some K'NEX parts in order to replicate the robots that have been described in this chapter, and to explore modifications of some of the design decisions that were made by my students. In the electronic version of this book, if you press this link, you will be taken to a web page that provides step-by-step instructions for robot building. This web material is included with the book's CD-ROM.

The reader might also be interested in seeing how the robots behave without having to go to the trouble of building them. Three different movies are included on the CD-ROM to demonstrate the behaviour of the robots described above.

Here is a short mpeg movie of the behaviour of a two-legged module.

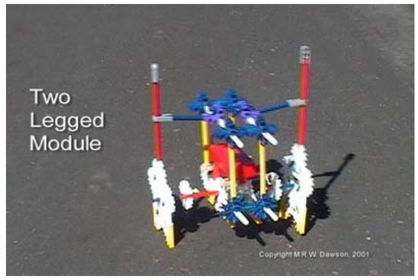

Here is an mpeg movie of the behaviour of a four-legged robot.

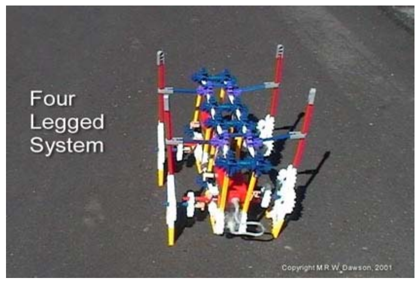

Here is an mpeg movie of the six-legged thoughtless walker.

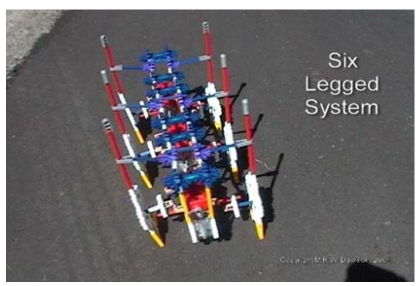

All of the movies were recorded in the summer of 2001 on the paved sidewalks outside of the Biological Sciences Building at the University of Alberta. I'd like to thank my son Christopher for his invaluable assistance in making these records of robot behaviour.

6.7 REFERENCES

Bechtel, W., & Abrahamsen, A. (1991). *Connectionism And The Mind*. Cambridge, MA: Basil Blackwell.

Blickhan, R., & Full, R. J. (1993). Similarity in multilegged locomotion: Bouncing like a monopode. *Journal of Comparative Physiology A, 173*, 509-507.

Brooks, R. A. (1989). A robot that walks; emergent behaviours from a carefully evolved network. *Neural Computation, 1*, 253-262.

Brooks, R. A. (1999). *Cambrian Intelligence: The Early History Of The New AI*. Cambridge, MA: MIT Press.

Dawson, M. R. W. (1998). *Understanding Cognitive Science*. Oxford, UK: Blackwell.

Dickinson, M. H., Farley, C. T., Full, R. J., Koehl, M. A. R., Kram, R., & Lehman, S. (2000). How animals move: An integrative view. *Science, 288*, 100-106.

Full, R. J., & Tu, M. S. (1991). Mechanics of a rapid running insect: two-, four-, and six-legged locomotion. *Journal of Experimental Biology, 156*, 215-231.

Holland, J. H. (1992). *Adaptation In Natural And Artificial Systems*. Cambridge, MA: MIT Press.

Holland, J. H. (1998). *Emergence*. Reading, MA: Perseus Books.
Kubow, T. M., & Full, R. J. (1999). The role of the mechanical system in control: a hypothesis of self-stabilization in hexapedal runners. *Philosophical Transactions Of The Royal Society B, 354*, 849-861.
Langton, C. G. (1995). *Artificial Life: An Overview*. Cambridge, MA: MIT Press.
Levy, S. (1992). *Artificial Life*. New York, NY: Vintage Books.
McComb, G. (1987). *The Robot Builder's Bonanza: 99 Inexpensive Robotics Projects*. Blue Ridge Summit, PA: TAB Books.
Mitchell, M. (1996). *An Introduction To Genetic Algorithms*. Cambridge, MA: MIT Press.
Pfeifer, R., & Scheier, C. (1999). *Understanding Intelligence*. Cambridge, MA: MIT Press.