

Learning, Adaptation, Survival: A Critique of Ashby's
Framework for Adaptive Behaviour

Christopher Larcombe

Supervisor: Dr. Ezequiel Di Paolo

Submitted for the degree of M.Sc.

University of Sussex, UK

September, 2009

Abstract

A framework for adaptive behaviour by W. Ross. Ashby is critiqued, and by identifying potential problems and limitations several possible extensions to it are suggested. The role of survival in the framework is considered, and several issues of temporality are also addressed. The framework was designed to show how the adaptive behaviour of the brain and nervous system are essentially mechanistic; in the thesis, the primary question asked is “what is the simplest possible *but not simpler* model of complex adaptive behaviour?”; how might adaptive behaviour be ‘scaled up’ or ‘scaled out’ in complexity? Several concepts are developed in order to increase the scalability and extensibility of certain systems and mechanisms featured in the framework. Suggested extensions are made hesitantly, and where possible based on solving potential problems or limitations identified in the original work.

Acknowledgements

Special thanks to my supervisor Dr. Ezequiel Di Paolo for the many interesting meetings, helpful advice and suggestions.

Thank you to all my lecturers who made the course such an enlightening experience.

Thank you to my family, for the kind and patient support during the term and holidays in between.

Finally, I would also like to thank my fellow peers for their 'relentless' support throughout the course.

Contents

1	Introduction	6
1.1	The Ashbyan Framework	8
1.1.1	A Logic of Mechanism	8
1.1.2	The Ultrastable System	9
1.2	Thesis Overview	10
2	Adapting to Survive	12
2.1	The Problem: Essential Variables	12
2.1.1	Survival	14
2.1.2	Viability and Regulation	15
2.2	Suitability and Lethality	16
2.2.1	Determinism and Ultimate Effects	16
2.2.2	Hypothetical Suitability	19
2.2.3	Defining the Living-state	20
2.3	Suitable Behaviour	23
2.3.1	Suitability of Viability Variables	24
2.4	Discussion	26
2.4.1	Future Work	27
3	Temporality of Change	28
3.1	The Problem: When to Change	28
3.1.1	Critical States	29
3.1.2	Disturbance to the Ultrastable System	30
3.1.3	The Homeostat Revisited	32
3.2	Potentially Ultrastable Regulation Units	37
3.2.1	Trials and Critical States	39
3.2.2	Demonstration	40
3.2.3	Optimality Landscapes	42
3.2.4	Perseverance	47
3.3	Discussion	49
3.3.1	Future Work	51
4	Adapting in Time	53
4.1	The Problem: How to Change	53
4.1.1	Searching for Optimality	53
4.1.2	Random Step-changes	55
4.2	Synergies	56
4.2.1	Intrinsic Modulation	58

CONTENTS

4.2.2	Extrinsic Modulation	63
4.3	Discussion	65
4.3.1	Future Work	65
5	Conclusion	67
	References	69
	Appendices	71
A	Source Code	71
A.1	PURU	71
A.2	Synergies	75

List of Figures

1.1	Ultrastable System with double-feedback.	9
3.1	Diagram of immediate effects of the ultrastable system.	29
3.2	Disturbance in the original framework.	31
3.3	Homeostat demonstrating ultrastability.	34
3.4	Hypothesised diagram of immediate effects of a Homeostat setup. (8/4)	35
3.5	Potentially Ultrastable Regulation Unit (PURU).	37
3.6	Reformulation of disturbances.	39
3.7	Model scenario: an agent moves in a 1-dimensional environment seeking a shaded zone to regulate its temperature.	41
3.8	Optimality landscapes based on hypothetical disturbance received in different states at a particular time.	43
3.9	Two possible PURU optimality variable couplings.	43
3.10	Demonstration with binary optimality landscape.	44
3.11	Demonstration with binary optimality landscape. Trail length is too short for the slow speed of the agent. Viability is lost around time-step 3000.	46
3.12	Demonstration with binary optimality landscape, for a range of speeds.	47
3.13	Demonstration of a slow moving agent adapting with a variable trial length.	48
3.14	Perseverance: automatic variation of trial length. The green spikes indicate a relapse, the start of a new trial.	49
3.15	Demonstration with mixed binary-continuous optimality landscape, for a range of speeds, where trial length is varied automatically.	49
4.1	A synergy with two 3-dimensional assemblies.	57
4.2	An example PURU using a synergy for reaction-unit parameter selection. The synergy may be modulated intrinsically, extrinsically or by a combination of the two.	58
4.3	An example agent-environment scenario for the synergy demonstration.	60
4.4	Adaptation to two alternating environments using a synergy of five assemblies.	62

LIST OF FIGURES

4.5	Adaptation to five alternating environments using a synergy of five assemblies. The allocation of assemblies to environments is random.	63
4.6	The accumulator of adaptations.	64

“Make everything as simple as possible, but not simpler.”

A. Einstein (1879-1902)



Introduction

Research in artificial intelligence and autonomous robotics has progressed rapidly in recent decades, but a set of comprehensive and unanimously agreed upon definitions of learning, adaptation and intelligence still appear to remain somewhat elusive (Phattanasri et al. 2007). Moreover, we have as of yet failed to explain how exactly the brain is able to produce the complex adaptive behaviour it does. But what kind of explanation are we seeking? What would a satisfactory answer even look like?

In 1952, in his book *Design for a Brain*, W. Ross. Ashby set out to show how the adaptive behaviour of the brain and nervous system were essentially mechanistic; that certain complex machines could in fact produce seemingly purposeful behaviour if appropriately structured, and that an understanding of them could be achieved by taking this into consideration (Ashby 1960).

In doing so, Ashby developed what he called a ‘logic of mechanism’ from which to begin the study, departing from previous teleological explanations centred around biological or technological substrates. Rather, the approach taken was *fully general* in nature, using a system of concepts somewhat similar to modern dynamical systems theory, or complex systems theory; what might be called a *cybernetics* or mathematical approach.

The subject of general intelligence in artificial systems, open-ended learning, or life-like robotics with a high level of autonomy have seldom been considered or researched. Work in artificial intelligence appears instead to have remained highly divided, split into sub-fields sometimes referred to as ‘narrow AI’. Rather than directly addressing the issues of adaptive behaviour or intelligence *in general*, the focus has been on specific ‘intelligent’ or biologically-inspired applications such as machine learning, knowledge representation and computer vision. Robotic systems have remained highly task-specific. While these topics are extremely useful in themselves, it is now becoming clearer

that such a ‘top down’ approach may not well suited to achieving the more ambitious, original goals of AI.

If the ultimate goal of the research is to accurately reproduce multiple adaptive or cognitive phenomena in a *single* artificial system, a more holistic view may prove useful or even necessary. For instance, the embodiment paradigm (Brooks et al. 1998) advocates the view that intelligence cannot be understood in isolation from the environment in which it is instantiated. Rather, that intelligence emerges as the result of continuous interaction between a complex body and a complex environment.

The perspective taken by Ashby, and by complexity science or dynamical systems theory, suggests one possible holistic view: that we can see ourselves and world as spatio-temporal patterns; as structure, or patterns, that both resist and exploit change over time. From this more abstract perspective, *behavioural patterns* can be characterised mathematically as stable attractor states, resistant to change, while change in behaviour can be thought of as a loss of stability in such patterns (Schöner 2007, Smith & Thelen 2003). All that remains is choosing an appropriate level of description or granularity from which to model and analyse the cognitive phenomena of interest, in order to capture the salient dynamics (Schöner 2007, pp 11). However, this is no trivial undertaking.

By considering the brain as an abstract machine, a deterministic system, Ashby was able to build a framework of concepts and derive mechanisms he considered necessary to produce what we *observe* as adaptive behaviour. However, what was left unanswered was how exactly the simple mechanisms provided as an explanation could be ‘scaled up’, combined, integrated and connected together, to produce the complexity we observe in living organisms. What is perhaps of interest to us, practically, is the production of artificial autonomous systems with some degree of generality in their function. If it were possible, for example, to build a robot that acted exactly like a cat, we may be rather tempted to say that we *understand the cat*, that we can implement the principles and dynamics of mammalian intelligence in a machine. It seems that this is the explanation that we are ultimately seeking: an operational explanation.

We might then ask, what is the simplest possible *but not simpler* model of complex adaptive behaviour? Is there any continuity in mechanism or principle between simple and complex organisms? It is these kind of questions the thesis will aim towards answering.

In terms of the the framework developed by Ashby in the book *Design for a*

1.1. THE ASHBYAN FRAMEWORK

Brain, a critique will be made in order to elaborate the original framework in several possible directions.

1.1 The Ashbyan Framework

Throughout the thesis ‘the framework’ or ‘the existing framework’ will be used to refer to the contents of the book ‘Design for a Brain’. Any suggested changes or possible extensions used throughout the thesis will be said to be part of the *elaborated framework*.

Before a detailed discussion can take place, several important points central to the framework must first be delineated. This chapter contains what is necessary to begin an analysis and elaboration of the framework; therefore, no justification or detailed discussion of what is stated in the following sections will be provided in this chapter of the thesis. The points stated will be returned to individually throughout when discussion is necessary.

1.1.1 A Logic of Mechanism

As part of the framework developed Ashby makes extensive use of what he refers to as ‘a logic of mechanism’. This consists of a set of concepts and terms equivalent to the basics of modern dynamical systems theory, with few exceptions. Therefore, rather than using the original terminology throughout the thesis, terms from dynamical systems theory will be used in substitution where possible for greater clarity. The definition of *variable* and *system* however have more specific meanings, and are used extensively throughout the thesis. The definitions are therefore stated below:

The term main variables is used by Ashby to refer to all variables in the ultrastable system, or in general, that vary continuously.

Definition *variable*

“A measurable quantity which at every instant has a definite numerical value.”

Definition *system*

“Any set of variables that [the experimenter] selects from those available on the real ‘machine’.”

1.1. THE ASHBYAN FRAMEWORK

The term ‘machine’ as used above is not defined precisely by Ashby, but is used throughout in a very general sense to refer to “something that may change with time”. The above definition of system makes the distinction between the ‘real material machine’ and that which is abstracted. This also takes into account the role of the observer, or experimenter, who decides precisely what measurable quantities are modelled as variables. Furthermore, Ashby states “it will be appreciated that *every real ‘machine’ embodies no less than an infinite number of variables*, all but of a few of which must of necessity be ignored”.

It should therefore be kept in mind that the system is distinct from the organism it is used to model. Also of central importance is the assumption that the living organism and its environment taken together may be “represented with sufficient accuracy by a set of variables” that forms a *deterministic* system, i.e., a closed system with a unique consequent to every state.

1.1.2 The Ultrastable System

The key mechanism used to explain different variations of adaptive behaviour by Ashby is the *ultrastable system* (see [Figure 1.1](#)). The system consists of several parts, or sub-systems of variables, forming a model organism. The system acts to maintain a continuously varying ‘essential variable’ within set limits, by behaving in a certain way, despite perturbations from the environment.




Figure 1.1: Ultrastable System with double-feedback.

The reacting-part labelled ‘R’ contains the overt variables that will be observed as behaviour. The box labelled ‘S’ is the ‘step-mechanism’ of the system, which changes periodically while the essential variables are beyond their limits, parameterising the reacting-part and changing the behaviour of

the system. Since the step-mechanisms change when and only when the essential variables are kept within their limits, the only stable behaviours are those that can occur while the essential variables are maintained within those limits.

The ultrastable system will be discussed in more detail throughout the thesis at various points, and will be frequently referred to. The diagrams used by Ashby to delineate systems such as this, in [Figure 1.1](#), are referred as *diagram of immediate effects*. The directional arrows indicate a causal relationship between two parts, that represent collections of variables that have an *immediate* effect each other.

1.2 Thesis Overview

The thesis is divided into three main chapters, each a separate critique addressing a specific problem or potential limitation in the original framework. In each of the chapters certain modifications to the framework are suggested in order to address the issues raised. Each chapter therefore contributes as both a theoretical critique and possible advance, extension or elaboration to the framework. It is hoped that some of the possible extensions discussed will add greater explanatory power to the framework and may have additional practical applications.

Due to the abstract nature of the work, there are potentially many different ways the framework could be extended and elaborated on in a rather open-ended manor. Conceivably then, any modification made may be considered as rather exploratory or even arbitrary. It is for this reason that modifications have been based quite strictly on critique. While each chapter builds upon the last, the critique forming the beginning of each chapter can be considered somewhat independently, as can certain of the other elements.

The first of the three chapters explores survival and the essential variables that form an important part of the framework. The concept of essential variables varying in lethality will be critically considered, together with their role in the ultrastable system.

In the second of the main chapters the temporal aspect of adaptation will be taken into consideration, asking the question “*when* is it appropriate to change behaviour?”. Various issues will be discussed, such as the nature of disturbance, trial length and perseverance in behaviour. The question of *how* to change behaviour in a beneficial way is left somewhat open until the third

1.2. THESIS OVERVIEW

and final chapter, which is focused on efficient adaptation. This chapter will contain more practical elements.

In the second chapter there is also additional emphasis on creating a more extensible version of the ultrastable system that it may be used as part of a larger system, rather than a model of a whole system including an environment-part. The treatment will remain high-level throughout, but not necessarily at level of whole organism and environment.

2

Adapting to Survive

This chapter of the thesis is primarily focused on what could be considered to be the foundations of the framework, upon which other concepts are based. The aim is not to falsify or negate Ashby's argument, nor to detract from the explanation of adaptive mechanisms provided. Rather, various definitions in the existing framework will be re-examined and elaborated on in order to eliminate ambiguity and explore the grounding of the terms used. In particular, the concept of survival is taken more explicitly into consideration. It is hoped that in doing so a more rigorous and precise set of definitions can be obtained, with a more fundamental grounding.

2.1 The Problem: Essential Variables

Of central importance to the framework is the concept of the *essential variable*. The term is used to define 'adaptive behavior' and 'survival', in addition to playing a major role in all the adaptive mechanisms explored. The essential variables of an organism are said to be those closely related to its survival, such that its survival is *threatened* in some way when the variables are taken beyond their 'normal limits'.

By the framework, an organism therefore seeks to maintain these variables within their respective limits in order to remain alive. For example, in certain species 'pulse rate' or 'temperature of blood' might be considered to be two essential variables that must be maintained within a specific range of values. This regulating phenomena, where variables are 'maintained' within bounds, is well recognised in living organisms by biologists and is commonly referred to as homeostasis: the variables are said to be homeostatically maintained by the organism.

2.1. THE PROBLEM: ESSENTIAL VARIABLES

Ashby states that the consequences that follow an essential variable passing beyond its normal limits *vary* in their ‘lethality’ to the organism depending on the essential variable. The consequences he exemplifies range from near immediate death, to “serious but not necessarily fatal”, to “not immediately dangerous” but still sufficient for the organism to take evasive action. Examples of variables that are not immediately dangerous might include ‘percentage of harmful bacteria in the knee joint’ or ‘level of pain’, while there may be other variables where the risk or immediacy of death is potentially higher, such as “amount of oxygen in blood”.

The essential variables are used in the framework primarily in instantiations of the *ultrastable system*. The simplest described includes variables modelling part of an environment, a ‘reacting part’ (behaviour), a ‘step-mechanism’ and the essential variable itself. When an essential variable passes beyond its “normal limits” or just “limits”, the step-mechanism is caused to change in some way that results in a new parameterisation to the reacting part, which changes the behaviour. The specific parameterisation and behaviour persists for some duration of time, but will change at random periodically while the essential variable(s) are outside their limits.

Having defined the essential variables in this way, Ashby then states that *survival* is that which takes place when *all* essential variables are kept within ‘given limits’:

“We can now define ‘survival’ objectively and in terms of a field: it occurs when a line of behaviour takes no essential variable outside given limits.”¹ (3/15)

It should be noted that it is not made clear whether the ‘given limits’ referred to in the definition above are equivalent to the ‘normal limits’ of the essential variables or the ‘physiological limits’ of the organism:

“That an animal should remain ‘alive’, certain variables must remain within certain physiological limits.” (3/14)

Assuming the word survive is used to mean ‘remain alive’, the latter definition implies the given limits are indeed the physiological. It is at this point some confusion may arise, since examples have already been provided where essential variables passing beyond their ‘limits’ are *not* immediately dangerous to the organism, contrary to the definition of survival. Furthermore, later in the framework the essential variables are associated with adaptive

¹The equivalent terms for *field* and *line of behaviour* are phase-portrait and trajectory respectively.

2.1. THE PROBLEM: ESSENTIAL VARIABLES

change in the organism: when one or more essential variables pass beyond their limits the organism is said to change its behaviour until the essential variables are once again all within their limits.

Unless the ‘given limits’ referred to in the definition of survival are taken to be *different* to the ‘normal limits’ or just ‘limits’ of the essential variables then what has been stated is clearly contradictory, since survival as defined precludes *any* essential variable passing outside its ‘given limits’. Therefore, the limits referred to in the two situations *must* be different. It will be assumed for the remainder of the thesis that this was implied by the framework, rather than it contains contradictions.

2.1.1 Survival

In order for the existing framework to remain free of contradiction, it must be assumed that there are in fact two types of limits applicable to the essential variable: the ‘normal’ and the ‘physiological’, where passage beyond the former induces adaptive change and passage beyond the latter causes death.

Having made this necessary step, the definition of survival now implies that *every* essential variable must have physiological limits. This is also supported by the definition of adaptive behaviour provided by Ashby:

“... a form of behaviour is adaptive if it maintains the essential variables within physiological limits.” (5/3)

This however causes a problem. That all essential variable must have physiological limits implies that all essential variables have the *potential* to cause death in the organism. By definition, this not only excludes all essential variables weakly associated with lethality, such as level of pain, it *only* includes those which *are* lethal, leading to certain death. For example, the ‘maximum level of pain’ does not cause death but certainly may lead to adaptation or some kind of learning behaviour.

If the criteria for survival (and adaptive behaviour) is changed to apply only to *some* essential variables, i.e., only those with physiological limits, then non-lethal essential variables are once again permitted. It seems unavoidable that this further modification to the definitions should be made. Having permitted non-lethal essential variables the definition of survival, or viability, must now be *that which occurs when all essential variables with physiological limits are kept within those limits*.

2.1. THE PROBLEM: ESSENTIAL VARIABLES

2.1.2 Viability and Regulation

Having made this *necessary* change to the framework, the essential variables are inevitably split into two different types: those that cause death and those that do not. It would now appear that the term essential is no longer completely appropriate, since those variables with physiological limits are in a very real sense *more essential*, or truly essential. Therefore, those variables with physiological limits will now be referred to as the *viability variables* of the organism, and those without will now be referred to as the *regulation variables* of the organism. Throughout the thesis these two variables will be referred to as being part of the *elaborated framework*. More specifically, variables with a limit that induce some kind of adaptive change will be referred to as lability variables, i.e., lability variables are regulation variables.

For simplicity, the assumption that all essential variables induce adaptive change will also now be dropped so the viability variables may be redefined as those variables that *do not* induce adaptive change, having *physiological limits only*. This simplification means that the viability and lability variables now have only *one* set of limits each. Furthermore, the distinction between the two types of limits can now be dropped as the type of variable now completely defines the type of limit, or more specifically what occurs when the value of the variable passes it. Survival can now be redefined in terms of the viability variables:

Definition *survival / viability*

That which occurs when all *viability variables* are kept within their respective limits.

There still however remains a final problem. Having critically considered the definitions and apparent contradictions in the original framework, it was *necessary* to distinguish a type of limit which when passed causes death, from a limit which when passed causes adaptive change. In doing so another problem was inadvertently introduced: that the distinction between lethal and non-lethal remains binary, i.e., all or nothing. What of those variables that vary in lethality that do not immediately cause death? What exactly is implied by the term lethality?

Because survival was not dealt with explicitly in the original framework this was also not taken into consideration. In fact, aside from the association to adaptive change in the organism, the essential variables are *co-defined* with survival, i.e., the terms define each other. This also avoids the problem of defining survival in terms of what is to be ‘alive’ in the context of the actual

2.2. SUITABILITY AND LETHALITY

dynamics or organisation of the organism.

To this extent, the terms remain ungrounded and the problem of how to approach a measure of lethality remains unsolved. In the following section a novel concept, or measure, will be introduced in an attempt to solve this problem and provide yet further justification for the substitution of the essential variables.

2.2 Suitability and Lethality

In his framework Ashby set out to show how seemingly purposeful adaptive behaviour is essentially mechanistic. It is perhaps then no surprise that demonstrating and discussing adaptive mechanisms was given priority over defining terms such as ‘survival’ or ‘alive’ precisely. In this section the focus is shifted to the apparent limitation imposed by the circular co-definition of survival and the essential variable (see §2.1.2).

A concept termed *suitability* is introduced, based on an *independent* definition of survival. The hypothetical suitability of an action or state with reference to an organism is defined as the time until its death. The concept, related to lethality, is then used to justify the essential (viability and regulation) variables, providing a grounding in survival and thus avoiding a circular definition.

2.2.1 Determinism and Ultimate Effects

It was assumed by Ashby in the original framework that a *living* organism and its environment “taken together, may be represented with sufficient accuracy by a set of variables that forms a state-determined system.”; in other words, that an organism and its environment can be modelled sufficiently by a closed dynamical system. On the *macroscopic* scale there is little argument that this is at least theoretically possible, since the laws of Newtonian mechanics are well understood and known to apply ubiquitously.

It might however be considered a more controversial claim to say that the exact time and circumstances of a living organism’s death are determined. Why is this the case? It *is* certainly true that in any *closed* deterministic dynamical system, by definition, everything that occurs following its initial state is fully determined, no matter how unpredictable or ‘random’ it may

2.2. SUITABILITY AND LETHALITY

appear to an observer. Any organism ‘in’ such a system would indeed have a predetermined death. The claim is perhaps controversial because it is not currently known whether or not our universe *is* deterministic in this sense. The discussion of determinism is of relevance here because it forces upon us a certain view of lethality, not directly considered in the original framework.

Assume there is a *deterministic* but highly complex ‘universe-system’ \mathcal{U} modelling the life of an organism \mathcal{O} in an environment² \mathcal{E} . Let $\alpha^{\mathcal{U}}(t) \in \{true, false\}$ be the observer defined ‘living-state’ of \mathcal{O} in \mathcal{U} at time t . Further assume that the state-space of \mathcal{U} is of sufficiently high dimensionality to contain multiple emergent levels of description, i.e., a dynamical hierarchy, and thus is capable of modelling such things as neural growth or morphogenesis where the degrees of freedom in \mathcal{O} appear to change³.

The environment \mathcal{E} may contain many other organisms, and to the observer the system may *appear to be* highly unpredictable. Although still theoretical, it could be argued that the model system \mathcal{U} is of sufficient generality to closely resemble what we observe in nature. What is the meaning of lethality with reference to \mathcal{O} in \mathcal{U} ? For example, what *measure* of lethality might differentiate between situations that are “followed by death almost at once” in \mathcal{O} from those that are “serious though not necessarily fatal” to \mathcal{O} ?

It seems natural to speak of probability, i.e., to say that in the former case death is somehow *more likely*. However, in both situations death is in fact *equally likely*, with an *actual* probability of 1 in *both* cases⁴. In such a system probability is an illusion of the observer and randomness must therefore be seen as a failure to perceptually categorise deterministic chaos⁵.

Now considering the adaptive behaviour \mathcal{O} might show, it is clear that statements such as “ \mathcal{O} acts to increase its probability of survival” or “ \mathcal{O} avoids situations which threaten its survival” are *inferences*, entangled with the observer’s limited knowledge and understanding of causality. The validity of

²The ‘environment’ is anything in the system which is not defined as the organism by the experimenter. The system \mathcal{U} is assumed to be a closed and deterministic.

³Here \mathcal{O} and \mathcal{E} are not models, and therefore do *not* refer to state-spaces or a set of variables; the organism-environment boundary is arbitrary and observer dependent, thus so are \mathcal{O} and \mathcal{E} at any time. The symbol \mathcal{O} should be taken to refer to the ‘real organism’ as if it existed objectively.

⁴It is assumed that the model \mathcal{U} is realistic enough that its organisms have a finite life-span.

⁵What may look random to one person may not look random to another; for example, a sequence of familiar mobile phone numbers presented digit by digit on a physical system, such as an analogue dial. The position of the needle would appear chaotic in motion unless the pattern was recognised by the observer, in which case it could then easily be predicted.

2.2. SUITABILITY AND LETHALITY

the statements in fact depend on *hypothetical* situations. What is implied by these statements is perhaps something closer to “*if* the organism \mathcal{O} had done B instead of A , then Y would have occurred instead of X ”, or collections of such statements. Do these statements have *actual* meaning?

As the goal is a treatment of organisms in general, the collection of hypothetical statements of the aforementioned kind can be reduced to those where X and Y concern the living-state $\alpha^{\mathcal{U}}(t)$ of the organism discussed only. For example, the statement could be translated into the following first-order formula⁶, for the situation A :

$$A \rightarrow \exists x \forall t (t > x) \wedge \neg \alpha^{\mathcal{U}}(t)$$

and its conjunct describing situation B :

$$B \rightarrow \exists y \forall t (t > y) \wedge \neg \alpha^{\mathcal{U}'}(t)$$

That is to say that in situation A the organism \mathcal{O} dies at time x and in situation B it dies at time y . The latter formula *must* refer to a system different to \mathcal{U} , i.e., \mathcal{U}' , since \mathcal{U} is a closed deterministic system by definition and can have only one outcome. If a specific point in time $t = p$ can be identified where B occurred rather than A , then this can be considered the point from which what subsequently occurs is hypothetical. It must be at this point of ‘branching’ that the hypothetical system \mathcal{U}' is ‘created’. Alternatively, the modification resulting in the two outcomes could also be seen as a hypothetical perturbation to \mathcal{U} at time p , if at p the criteria that \mathcal{U} is a closed (and deterministic) is momentarily dropped.

It could also be argued that *even if* the system \mathcal{U} is *not* deterministic it is still likely to be the case that B (and thus y) remain hypothetical, as B is quite likely to be *impossible* given the history of \mathcal{U} . When this logic is applied to *reality* rather than the model \mathcal{U} , this indeed appears to be the case. For example, for the fox to *not* have jumped over the lazy dog, within its brain the neurons responsible for it jumping over the lazy dog would have to have behaved very differently, implying an entirely different history. For the situation to have any reality there must exist a non-deterministic transition in the history of the system which if taken would have resulted in a subsequent state where the fox did not jump over the dog, but was otherwise identical to what *did* occur. It is not clear to what precision such a state could be described, even if it did exist.

⁶The binary infix predicate $>$ denotes the temporal ordering relation ‘later than’.

2.2. SUITABILITY AND LETHALITY

Without any further discussion, it will now be assumed that in any deterministic system lethality is hypothetical, while in any non-deterministic system lethality is ‘most likely’ hypothetical.

2.2.2 Hypothetical Suitability

Since the primary concern is the adaptive behaviour of organisms *in general*, the simplest or perhaps only fundamental measure of lethality must be in terms of *hypothetical time until death*. This does provide a basis for a non-binary measure of ‘goodness’⁷, albeit hypothetical and on the assumption that lethal situations are ‘not good’ for an individual organism. It should be clarified at this point that so far no direct claim has been made that organisms do or do not act specifically to prolong their existence.

The hypothetical time until death of an organism will now be referred to as its *hypothetical suitability*, or just *suitability*. Thus, lethality can be defined as the negation of suitability: lethality is unsuitability. The term suitability will be used, rather than lethality, primarily to distinguish the hypothetical measure introduced from the more vague concept of lethality as it was *previously* used.

The two terms can now be defined with greater precision in terms of the previous example universe-system \mathcal{U} , ‘containing’ the organism \mathcal{O} . While the concept is developed as a conceptual tool, the formalisation given below is kept as *general* and realistic as possible, so it may potentially be applied to the real world. Embodiment will also be acknowledged by assuming the role of a set of observers \mathcal{Q} as organisms in \mathcal{U} *distinct* from \mathcal{O} .

Let the organism \mathcal{O} be sufficiently *described* at time t by the representative point $x(t)$, or ‘current state’, of the *dynamic* state-space $O(t)$. The state-space is defined by the observers \mathcal{Q} in \mathcal{U} at each point in time, since \mathcal{O} is assumed to be an observer-dependent emergent phenomena, i.e., $O(t)$ represents a low-dimensional description-space where each dimension is a ‘measurable’ collective variable of \mathcal{U} to \mathcal{Q} . The state-space $O(t)$ therefore supervenes and is determined by the evolution of \mathcal{U} . That the state-space $O(t)$ *may* change is perhaps unusual, though justified here as the case is general; the accuracy of the *low-dimensional* model must be sufficiently maintained, including through processes of ‘qualitative change’, such as morphogenesis

⁷To discuss lifestyle or suffering at this point would be to complicate matters; here a more fundamental concept is sought to explain behavioural phenomena such as suffering, etc.

2.2. SUITABILITY AND LETHALITY

or ‘losing a limb’, where the apparent ‘degrees of freedom’ in the model may need to change.

Let the observer-dependent living-state of \mathcal{O} at time t be defined by the following function:

$$\alpha_t : O(t) \rightarrow \{1, 0\}$$

where 1 represents ‘alive’ and 0 represents ‘not alive’. Further, let the time-varying hypothetical suitability $\chi_t(s)$ of a state $s \in O(t)$ at time t be defined by the following function:

$$\chi_t : O(t) \rightarrow \mathbb{R}_0^+$$

where

$$\chi_t(s) = \min \{t' \mid \forall t' : \alpha(\tilde{\varphi}_{t'}(\tilde{s}(t))) = 0\} - t$$

It should be emphasised that the two functions change at each time-step as defined by \mathcal{Q} . The suitability of the hypothetical state $\tilde{s} \in O(t)$ at time t is thus equal to the duration of time until $t = t'$, at which point the organism described dies in the hypothetical flow $\tilde{\varphi}_{t'}(\tilde{s}(t))$.

2.2.3 Defining the Living-state

The term ‘living-state’ has so far been used without justification, to directly refer to whether or not an organism is or is not alive. The function α_t was then introduced to determine the living-state of an organism, but was not described or defined in any detail. So far it has been assumed that an observer is able to define the living-state of the organism based on its *description*; that the living-state of \mathcal{O} is specified by a description $O(t)$ at any point in time. This is equivalent to saying that, at any point in time, an observer is able to say, “okay, now it’s dead”. Immediately there are two potential problems that follow this assumption:

1. That the living-state is not necessarily binary, i.e., alive or not-alive.
2. That a judgement may only be possible over a *period of time* rather than at an instant in time.

The first problem cannot be discussed in detail until a specific criteria is reached. The second problem however is of greater significance, particularly since ‘alive’ and ‘living’ intuitively appear to refer to a *process* rather than sets of variables at particular values. In dynamical systems terms this is equivalent to saying that multiple successive states⁸ are necessary to define

⁸A trajectory or region of the phase-space.

2.2. SUITABILITY AND LETHALITY

the living-state of an organism, rather than a single state only. The second problem will now be discussed in detail.

It is at first tempting to suggest as a potential solution that an observer can determine the living-state at *a specific time t* , but this falsely shifts the focus to the brain of the observer. It is not a solution, for what is being ‘objectified’ is not the phenomena of interest. For instance, the brain of the observer would require observation by another observer to establish inter-subjectivity.

Attention should instead be drawn to fact that *not only* do organisms die, but they also eventually *cease to exist*. In the terminology of the thesis this is equivalent to saying that \mathcal{O} and $O(t)$ will inevitably lose their definition, i.e., that what was previously described will eventually no longer have a description because it will no longer exist. Organisms can and do ‘disintegrate’. This *process* cannot be described by a state $s \in O(t)$ as it is the process of $O(t)$ changing; it is the *process of the state-space itself changing*. This process, resembling ‘qualitative change’, will occur until what is described is no longer an organism. This may not appear to be of immediate significance, but if this process *is* the process of death, the transition from alive to dead, then both problems listed above are indeed real problems. If this is the case then the function α_t cannot exist, as its cannot be precisely defined.

It seems sensible at this point to try and define ‘alive’ or ‘living’ operationally. What is there, if anything, that all *living* things do? This is a very difficult question to answer given that there are millions of different species of organism. There is one thing in particular however that has been identified, namely that living things tend to reproduce *themselves*. This is not reproduction, but self-reproduction. An attempt to define this phenomenon with precision has been made by Maturana & Varela (1991). The authors consider living systems to be “units of interactions” with “circular organisation”, such that the circular organisation of the system functions to maintain or produce itself, i.e., the organisation produces the same organisation:

“The circular organization in which the *components* that specify it are those whose synthesis or maintenance it secures in a manner such that the product of their functioning is the same functioning organization that produces them, is the living organisation ... it is this circularity that it must maintain in order to remain a living system and to retain its identity through different interactions.”

9

⁹(Maturana & Varela 1991, pp 9)

2.2. SUITABILITY AND LETHALITY

A living system with this organisation is referred to as an *autopoietic system*. However, autopoiesis was defined primarily with reference to single cells, to characterise their organisation as that of a living system. The organisms of primary interest in the thesis concern those that display overt adaptive behaviour, such as mammals, which are meta-cellular organisms. Unfortunately the situation is more complex. Maturana & Varela (1998) recognise this, suggesting that such systems could be regarded as “second-order autopoietic systems” as they include cells as “components of their structure”. The authors leave the question open as to whether or not such systems are also regular first-order autopoietic systems.

If we accept this as a sufficient definition of ‘living’ and thus ‘alive’, can the presence of an autopoietic organisation be determined by $O(t)$? It seems unlikely, since autopoiesis is a *process* whereby a circular organisation produces or maintains itself. The point at which this stops occurring is the point at which the system is no longer a ‘living’ system, i.e., when death occurs. To complicate matters further, in meta-cellular organisms multiple components are “maintained” at multiple levels of description, e.g., cells, organs, organ systems, etc. in what appears to be a hierarchy. It seems highly unlikely that a specific point in time could be identified when such a complex hierarchical organisation collapses; the situation is fuzzy. It would therefore appear that the first problem listed above is also a real problem, as well as the second.

There is also an additional argument, albeit perhaps weaker. From the perspective taken an organism must be what an organism does¹⁰. Taking this into consideration, it could quite comfortably be argued that if what an organism does is ‘live’, then if something is an organism *it is* by virtue alive. Of course this is falsified by the existence of a “dead organism”, but it could be quite easily be argued that this is in fact a category error¹¹. Furthermore, if something is alive we may justifiably call it an organism, since we have no counter example of some entity that is living but not an organism. This logically equates the terms ‘alive’ and ‘organism’, and is significant to the discussion because it quite clearly implies that *the point at which the organism is no longer an organism is the point at which it is no longer alive*. This is crucial as it was previously acknowledged that this process cannot be sufficiently described by $O(t)$ either; the change that takes place when an

¹⁰The reader is reminded that the organism is considered to be closer to a process than a ‘thing’.

¹¹A semantic argument should be avoided. Humans appear to make the categorisation ‘alive’ or ‘dead’ on a superficial basis, based on what is overtly observable *by inference*, e.g., movement, breathing. The internal organisation is invisible to the naked eye so cannot be directly taken into account.

organism dies is not a change in $O(t)$.

When an observer finds her or her variables begin to lose meaning it will likely be noted that a *qualitative change* has taken place. If a qualitative change is what occurs when an organism makes the transition from alive to dead then $O(t)$, if it even still exists, is insufficient to determine the living-state. It seems therefore that unless $O(t)$ can describe something specific to the *organisation* of the organism the function α_t *cannot describe the living-state*. This conclusion is sufficient for what is relevant to the thesis and will be further discussed in the following section. The question of whether or not it is possible to capture first and second-order autopoietic organisations operationally will be therefore be left open.

2.3 Suitable Behaviour

In the first section of this chapter the role of essential variables in the framework was questioned. It was argued that ambiguity in the definitions of various terms and in use of the word ‘limits’ led to unnecessary confusion. It was subsequently proposed that a clearer treatment could be obtained by replacing the essential variables with two new types of variable: *viability* and *regulation*. Making this modification both eliminated the source of confusion and further departed from the original framework by allowing the possibility of non-lethal essential variables. The circular co-definition of ‘survival’ and ‘essential variable’ was also acknowledged. At the end of the section attention was drawn to two potential limitations in the framework, still remaining after dichotomising the essential variables:

1. No continuous measure of ‘lethality’; variables divided into lethal and not-lethal.
2. Circular co-definition of ‘survival’ with ‘viability variables’.

The two limitations were then addressed in §2.2 (page 16) beyond the immediate context of the framework. In this section the viability and regulation variables are reconsidered in terms of the hypothetical concept of suitability.

2.3.1 Suitability of Viability Variables

Having made the *necessary* modification to permit non-lethal essential variables¹² in the framework, a dichotomy was inadvertently introduced which made the use of the word essential questionable. Thus, the question “how essential is essential?” was raised. In order to answer this question the viability variables can now be described in terms of suitability, and more crucially be differentiated from the non-vital or ‘less essential’ variables, i.e., the regulation variables. Furthermore, since suitability is a continuous measure the dichotomy can now be broken, albeit only hypothetically.

In the previous more *general* treatment of §2.2.2 (page 19) the description (state-space) of the organism was allowed to change *in dimensionality*. This was a necessity as what was of importance was itself high-level, ‘collective’, or observer-dependent, namely the ‘thing’ that was to change that described a hypothetical situation.

Assume now however that there is an organism \mathcal{O} , which may be sufficiently described for some duration of time by a state $x(t)$ in the *static* state-space S . For simplicity, and in contrast to the more general treatment of (§2.2.2), the description-space is now regarded as being fixed in dimensionality. It should be noted however, that the description $x(t) \in S$ should not be assumed to be of sufficiently high dimensionality to model a dynamical hierarchy or cellular automata-like ‘matrix’ from which the organism may ‘emerge’. Rather, it must be of sufficiently *low* dimensionality to model a set of measurable collective variables, for an appropriately limited duration of time. The latter constraint is of course necessary to ensure that the dimensionality of the description-space remains constant, albeit temporarily.

Of the n variables that make up the state-space S , further assume that among them is a viability variable indexed by $v \in \{1, 2, \dots, n\}$, of the value $x_v(t)$ at time t . By definition of the elaborated framework in the first section (§2.1.2, page 15), when the limits of this particular variable v are passed the organism described by S will die, as the criteria for survival will no longer be held. This is also the case with the original framework, but for *all* essential variables. Therefore, by both frameworks, the living-state of the organism modelled can be determined at time t by the current state $x(t) \in S$.

The concept of suitability however invokes a separate, independent definition of survival based on the function $\alpha_t : \mathcal{O}(t) \rightarrow \{0, 1\}$, for the time-varying

¹²Variables that induce adaptive change but are not necessarily life-threatening, now defined the regulation variables.

2.3. SUITABLE BEHAVIOUR

state-space $O(t)$. In discussion of the living-state in §2.2.3 it was concluded that unless $O(t)$ could capture the *organisation* or autopoietic process of the organism it described, the function α_t would not exist. Informally speaking, in actuality since ‘living’ is a process the living-state cannot be sufficiently determined by a single state $s \in O(t)$, at any instant in time. At the very least it would seem a succession of states would be necessary.

This is of direct relevance to the eligibility of the definitions of survival provided in the original and elaborated framework, which both imply that the living-state of an organism described by a state-space O can be determined at a time t by the state $x(t)$. This is a *special case* of the more general one discussed in §2.2.3, so the same argument must apply: if $O(t) = S$, and no state $s(t) \in O(t)$ is sufficient to determine the living-state, then no state in $s(t) \in S$ is sufficient. It must be concluded on the strength of the argument in §2.2.3 that the definitions of survival provided by both frameworks are in fact incorrect, or approximations only. The viability and essential variables fail to capture at any specific point in time what distinguishes a living organism from a dead one. Survival must be defined by some other means.

However, it could still be argued that since the model is an approximate, low-dimensional description S of the organism \mathcal{O} , highly abstract variables could be selected which do somehow capture something of its autopoietic organisation. For instance, the “total levels of description”, which would need to be obtained using some specific method of the observer. After all, it is assumed the variables are ultimately observer-dependent, abstracting something of a higher dimensional system with countless components or variables. This argument may have some validity when applied to the viability variables of the elaborated framework, but it is very difficult to imagine a variable *this abstract* also directly involved in an adaptive mechanism, whereby adaptive change is induced. This is of course the case in the original framework, as all essential variables also have this function.

Nevertheless, what is clear is that certain measurable variables are at least *correlated* to the autopoietic organisation of organisms. These variables are the viability variables, somehow ‘more lethal’ than the regulation variables. This notion or intuition can now be analysed with greater precision using the conceptual tool of suitability. It should be kept in mind that suitability is *still* hypothetical when applied to a dynamical system, having no ‘absolute’ truth except in cases where the ‘hypothetical situation’ or state in question is a possible outcome in the next time-step of a non-deterministic system¹³.

¹³It was argued that in non-deterministic system with an apparently deterministic macroscopic level of description this was highly unlikely, unless the hypothetical situa-

The concept is used primarily to elucidate what is meant by ‘lethal’.

The suitability of a state is time-varying and thus context-dependent; for example, the suitability of the state ‘running forward’ in the context of ‘room’ would be rather different to the suitability of the state in the context of ‘cliff’. As for the viability variables however, e.g., amount of oxygen in blood, it is quite obvious that certain values will have similarly low suitability in almost all situations. The suitability of these variables are much less dependent on context, i.e., no matter what situation the organism is in, the suitability of zero oxygen in the blood would be very low indeed. This refers to a *region* or sub-space of low suitability in the mapped space of states *where certain dimensions* corresponding to viability variables of the space have certain values beyond certain limits. This suitability will also not frequently vary, i.e., it will have low standard deviation.

2.4 Discussion

Various definitions in the framework were critically considered, and it was concluded that ambiguity in several terms led to two different possible interpretations: one interpretation implied there were contradictions in framework, while the other disallowed non-lethal essential variables. The latter contradiction-free interpretation was assumed.

In order to resolve the situation and allow non-lethal variables such as ‘level of pain’ the capability of causing adaptive change, the essential variables were split into two types: viability and regulation. The viability variables were defined to be those with physiological limits only, that when passed would kill the organism. The regulation variables were defined as those without limits that could kill the organism, but that include lability variables, which have limits that cause adaptive change.

This splitting of the essential variables introduced dichotomy between lethal and non-lethal variables, since only the viability variables were left with limits that had an impact on the viability of the organism. It was also noted that survival and viability remained circularly defined with survival, as in the original framework. The notion of lethality and survival was considered in some depth to obtain clarification. In particular, the circular co-definition of survival and essential variables (and viability) was considered. A theoretical system was discussed with deterministic dynamics on a macroscopic emergent

tion referred to a variable on the microscopic scale.

level, but with a *potentially* non-deterministic lower-level from which the macroscopic dynamics emerge. Hypothetical dynamical systems of this type were considered in order to understand what might be meant by a 'measure of lethality'. The hypothetical concept of suitability (time until death) was introduced; hypothetical for deterministic systems, but likely hypothetical for non-deterministic systems with approximately deterministic macroscopic levels of description.

It was concluded lethality is an illusion of the observer in systems that are deterministic, or where what is of importance to description is deterministic; a reflection of the fact the environment contains redundancy, that it repeats itself, reflecting only our own ability to abstract causal relationships from the environment. Nonetheless, it was argued the concept of suitability may be useful as a conceptual tool, as it has greater precision than the previous use of word 'lethal'.

A definition of survival was then discussed to determine the 'living-state' of an organism. It was argued that the viability variables (and certainly not essential variables that induce plastic change) cannot be used to determine whether the organism described is alive at any one instant in time, because this may be determined by the autopoietic organisation of the organism, which is complex a process with multiple levels of description. The viability variables were however noted to have a particular pattern of suitability: they are hypothetically correlated to loss of autopoietic organisation in multiple contexts. Emphasis was placed on the coupling between variables to determine how relevant each is to the survival of the organism, and how each variable has a suitability associated.

2.4.1 Future Work

An operational definition of survival independent of the viability or essential variables is still lacking. While it may not be possible to provide a definition in this format, modelling work with dynamical hierarchies may be useful to clarify this is the case. Different kinds of definitions need to be taken into account; specifically, what would constitute a satisfactory operational definition.

3

Temporality of Change

This chapter of the thesis is primarily concerned with the temporality of adaptive change, addressing the question ‘*when* is it appropriate to change behaviour for the better’? The question assumes the usefulness of a particular behaviour *persisting* over time, but this will be both discussed and demonstrated. This particular issue of temporality is important as the trial of a behaviour must be of a specific length: too short and the correct behaviour will not last, too fast and damage may occur (8/15). In the original framework this issue was not considered, but it was acknowledged that the duration of such trials “demanded adjustment”, and that a mechanism for varying the trial length automatically was not considered (17/10). It will be argued in the subsequent sections that the nature of disturbance to the ultra-stable system, specifically to the essential variables, needs to be reconsidered in order to address this issue of the timing of changes.

One temporal aspect not dealt with in this chapter is the issue of adapting efficiently, i.e., adapting in time. This is instead addressed in the following chapter, as the speed at which adaptation occurs is largely determined by the nature of the changes taking place, i.e., *how* and *what* to change. Change will continue to take place until adaptation occurs. This subdivision is of course approximate, for a more coherent treatment of ‘adaptive change’ in general. Therefore, the issues of ‘when, how and what’ can be expected to overlap; where this occurs to a significant extent it will be stated explicitly.

3.1 The Problem: When to Change

The issues of temporality in the framework revolve mainly around *the ultra-stable system*, the primary adaptive mechanism used to explain adaptive

3.1. THE PROBLEM: WHEN TO CHANGE

behaviour. In the ultrastable system a *reacting part* forming the overt, behaving part of the system is parameterised by a *step-mechanism*, which changes periodically in discrete steps. The step mechanisms can be thought of as a set of parameters that change the attractor landscape of the phase-space of variables in the reacting part; the ‘behaviour’ is determined by the step-mechanisms, which act as parameters.




Figure 3.1: Diagram of immediate effects of the ultrastable system.

The step-mechanisms are affected by one or more continuously varying *essential variables*, which are in turn affected by an *environment-part*. Most importantly, the system is structured such that the step-mechanisms change *when and only when* the values of the essential variables are kept within certain limits. These limits are fixed, and passage beyond them is assumed to be potentially dangerous to the organism. Thus, the system changes the reacting part, *periodically*, until all the essential variables are kept within their limits, and the organism is safe. When and only when this occurs is the behaviour stable.

3.1.1 Critical States

Each step-mechanism, at a particular value, is said to have an associated set of *critical states*, such that when the current state is a critical state the corresponding step-mechanism will change in value. This means that all states where the essential variables are beyond their limits will be critical states, causing a change in step-mechanism. However, as previously mentioned, a behaviour must persist for some time before changing, as the behaviour changes only periodically. This is referred to by Ashby as a *trial*: the correct behaviour and its corresponding value of step-mechanism(s), if it is eventually found, must persist for a sufficient duration of time for the essential variables to return within their limits. In other words, each behaviour must

3.1. THE PROBLEM: WHEN TO CHANGE

persist long enough for the system to receive ‘feedback’, as to whether a behaviour is appropriate to the current situation of the system. Surprisingly, this feedback is not modelled directly in the framework, which may or may not be seen as a limitation. This issue will be investigated.

With respect to the ultrastable system and the *Homeostat*, in the original framework only the values of the essential variables are taken directly into account. For instance, the disturbance to the essential variables is neither defined nor made explicit, and the critical states of the reacting part are seldom considered. These topics will be analysed in the following sections to address the issue of temporality and periodic changes to the step-mechanisms, and system behaviour.

3.1.2 Disturbance to the Ultrastable System

The following definition of the ultrastable system is provided, in complement to the diagram of immediate effects (see [Figure 1.1](#), page 9 and [Figure 3.1](#)):

“Two systems of continuous variables (that we call ‘environment’ and ‘reacting part’) interact, so that a primary feedback (through complex sensory and motor channels) exists between them. Another feedback, working intermittently and at a much slower order of speed, goes from the environment to certain continuous variables which in turn affect some step-mechanisms, the effect being that the step-mechanisms change value when and only when these variables pass outside given limits. The step-mechanisms affect the reacting part; by acting as parameters to it they determine how it shall react to the environment.” (7.26)

Of particular importance is the inclusion of part of the environment in ‘the system’, which determines what is and what is not considered to be a disturbance, or perturbation. This leads to two different types of disturbance to the ultrastable system, that Ashby claims will cause it to demonstrate its stability and its ultrastability respectively:

- (1) “Frequent (or even continuous) small impulsive disturbances to the main [*environment, reacting part, essential*] variables.”.
- (2) “Occasional changes, of step-function form, to its parameters”.

The two type of disturbance can be represented diagrammatically (see [Figure 3.2](#)) in reference to the components of the ultrastable system:

3.1. THE PROBLEM: WHEN TO CHANGE

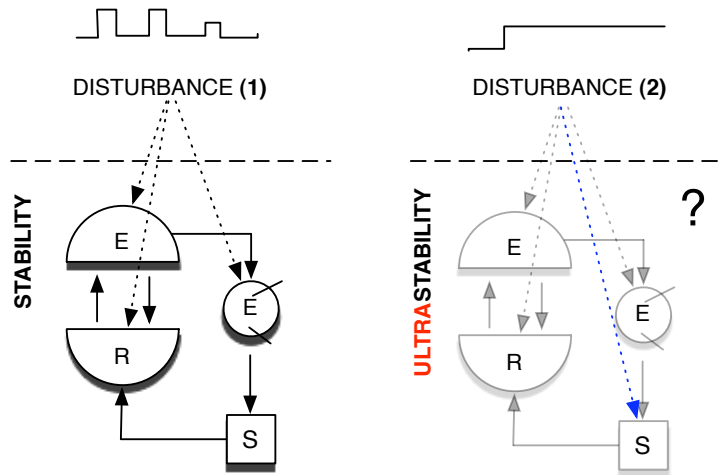


Figure 3.2: Disturbance in the original framework.

Firstly, it is important to note that both types of disturbance listed above are in fact changes of parameter to the entire system, since the system includes both the “organism and environment joined”. This is potentially confusing as it implies the environment-part modelled has an environment, from which it receives disturbance. This is indeed the case, only because *part* of the environment is modelled within ‘the system’, and this is the environment referred to. There are a variety of reasons for this, but perhaps most significant is that it allows the system to be classified as ‘ultrastable’. Nonetheless, if the environment changes such that every state in the reacting part becomes a critical state, then ultrastability will not be demonstrable as every state will cause a step-mechanism to change. As the system is ultimately open to disturbance *ultrastability cannot be guaranteed*, despite the fact the system is state-determined when its parameters are stable.

Secondly, it is not made clear or stated explicitly what part of the system will be *immediately affected* by either type of disturbance: the first type refers to the main variables (continuously varying), but makes no further distinction, while the second type makes no distinction at all. It is not clear whether the parameter changes are directly to the system or just changes of the environment.

Most importantly, *no distinction is made between impulsive disturbances to the reacting part and to the essential variables*. In fact, nowhere is the disturbance immediately affecting the essential variables independently discussed or referred to. It will be demonstrated that while this is not incorrect, this treatment of disturbance lacks explanatory power due to an excessively high

3.1. THE PROBLEM: WHEN TO CHANGE

level of generality. It will be argued that an independent reference to what disturbs the essential variable(s) can prove useful.

Bearing this in mind, the Homeostat will now be analysed in order to go beyond a purely theoretical discussion. The Homeostat was used in this manor by Ashby specifically to elucidate the functioning of the ultrastable system (8/1).

3.1.3 The Homeostat Revisited

Following the introduction of ultrastability Ashby demonstrates the *Homeostat*, a machine built “according to the definition of the ultrastable system” (8/1). The Homeostat is a *physical device* consisting of four ‘units’, each with a ‘pivoted magnet’ and attached needle, which can deviate from a central position. The angular deviations of each magnet (needle) provide four “main variables”, which may be homeostatically maintained around their respective centre points. When all four units are used, the Homeostat can be described by the following linear differential equation:

$$\frac{dx_i}{dt} = a_{i1}x_1 + \dots + a_{i4}x_4 \quad (i = 1, 2, 3, 4) \quad (3.1)$$

where x_i is the position of the i^{th} magnet on the i^{th} unit and a_{ij} is the immediate effect of unit j on unit i . The values in x change continuously and the values in a change in step-function form, varying between -1 and 1 in discrete steps.

Each unit i has a ‘uniselector’ (stepping-switch), approximately determining the values of a_{ij} , $j = (1, 2, 3, 4)$. The values are changed based on the values of x , when a *critical state* is met; specifically, the value of a_{ij} is dependent on the critical state(s) of x_i . In the Homeostat each unit i has critical states where $-\frac{\pi}{4} > x_i > \frac{\pi}{4}$.

Since there are *multiple* critical states per unit, the point at which a step change is in a unit is called the *critical surface*. Homeostat units can be controlled by a uniselector in this manor, or hand-controlled. In the latter case this allows the parameters to a specific unit to remain fixed, so there are no critical states for *its* uniselector (its step-mechanism remains inactive).

To bring the Homeostat into correspondence with the ultrastable system, Ashby relates various parts of the Homeostat to parts of the ultrastable system defined in (8/3) as follows:

3.1. THE PROBLEM: WHEN TO CHANGE

- The variables x as the “main variables”, divided “arbitrarily” to represent the environment and reacting part.
- The *relays* F as the essential variables, with “physiological limits” corresponding to the flow in F when the “needles are deviated to more than about 45 degrees from the central positions”.
- The uniselector values as the step-mechanisms.
- The needle diverging for “more than a few sections outside the limits of $\pm 45^\circ$ ” causing the step-mechanisms to change.

What is *crucial* to note here is that the needle deviations (main variables) do *not* model the essential variables; and in *all* of the demonstrations of the Homeostat that follow (8/4) the values of the essential variables are never graphed. Furthermore, at no point is a diagram of immediate effects shown of a format similar to [Figure 1.1](#), page 9, where it would be obvious what each variable in each demonstration corresponded to, in the previously introduced ultrastable system. In the first of several different Homeostat setups (demonstrations) Ashby discusses the behaviour of a 2-unit Homeostat (8/4).

In this demonstration the second unit of the Homeostat responds to three “displacements” $D_i, i = (1, 2, 3)$ to its needle ([Figure 3.3](#) below) from the real environment. Two ‘reversals’ $R_j, j = (1, 2)$ of the ‘commutator’ are also made, which reverse the motion of the needle¹. Since no time series in the figure corresponds to an essential variable, the values of the two essential variables are *not shown*, and thus neither are their values in relation to their limits. Instead, the limits shown on the diagram correspond to the critical states of the *main variables*, i.e., those states that will cause the essential variable to go beyond its limits, in turn causing a step-mechanism to change.

The model is admittedly therefore quite confusing as a demonstration of ultrastability; it would be an easy mistake to misinterpret the model by assuming that the values shown, being homeostatically maintained, were in fact those of the essential variables. It is the essential variables that should *always* be homeostatically maintained, within fixed limits, by the ultrastable system previously introduced. Even if the main variables had values equal to those of the essential variables, at all times, this distinction is still of great importance due to the causal relations between the variables. It is not necessary that variables of the reacting part should be maintained with *fixed regions*.

¹It is not stated which unit is affected by the reversals, although the ‘ R_i ’ labels are closer to the time series of Unit 2.

3.1. THE PROBLEM: WHEN TO CHANGE




FIGURE 8/4/1: Two units (1 and 2) interacting. Line 1 represents the side-to-side movements of Unit 1's needle by vertical changes. Similarly line 2 shows the behaviour of Unit 2's needle. The lowest line (U) shows a mark whenever Unit 1's uniselector advanced a step. The dotted lines correspond to critical states. The displacements D were caused by the operator so as to force the system to show its response.

(a) Homeostat demonstration in “The Homeostat as adapter”. (8/4) (b) An example “complex return”. (5/12)

Figure 3.3: Homeostat demonstrating ultrastability.

In the demonstration it is stated that the diagram of immediate effects is $1 \rightleftharpoons 2$, with the effect $1 \rightarrow 2$ being hand-controlled² and $2 \rightarrow 1$ being uniselector (step-mechanism) controlled. In other words, only unit 1 could cause a step-change to occur. When this occurs at R_1 there must therefore be an effect $1 \rightarrow E$, where E is the essential variable *not modelled or discussed*³ in (8/4).

It should now be noted that it is *nowhere stated* whether the main variable of Unit 1 is part of the environment or the reacting part (the agent). Either Unit 1 or Unit 2 *must* model a variable of the reacting part, otherwise there is no part of the agent for the step-mechanisms to regulate, and the Homeostat cannot be brought into correspondence to the ultrastable system by definition or function. If it is assumed that Unit 1 is part of the environment, then there is no immediate effect from the step-mechanisms to the reacting part (Unit 2); in this case nothing in the agent is being regulated directly, only a variable in the environment with a *fixed critical surface*.

If the alternative is assumed, that Unit 1 is instead a reacting part, then the situation arises where the reacting part has an immediate effect on the essential variable (due to $1 \rightarrow E$). This situation would be potentially fatal to an organism, as it would allow the reacting part to ‘specify’ its goal.

²This means Unit 2 cannot cause the uniselector to change, as it is controlled by hand.

³It could be reasonably assumed that a unit with a critical state has an associated essential variable ‘off screen’ (not modelled) with the *same value* as its needle position. Note however that it cannot be both an essential variable *and* a main variable by Ashby’s definition, as the needle position does not represent the value of an essential variable.

3.1. THE PROBLEM: WHEN TO CHANGE

The situation is exemplified by Ashby (prior to introducing the Homeostat) in (7/3) as a kitten's brain that could "change it from an organism that must not get burnt to one that benefited by being burnt!". It cannot be the case therefore that Unit 1 models part of the environment either⁴. It can only be assumed that either the demonstration does not correspond to the ultrastable system or it is not detrimental that there is no *immediate* effect from the uniselector (step-mechanism) to the reacting part (unit 2), as in the ultrastable system diagram. Of the two possible arrangements considered it seems to be the only sensible option that Unit 2 is the reacting part. Assuming this is the case, the diagram of immediate effects would look as follows:




Figure 3.4: Hypothesised diagram of immediate effects of a Homeostat setup. (8/4)

The diagram is not immediately recognisable, but on close analysis the setup does in fact conform to the definition of the ultrastable system provided: "... The step-mechanisms affect the reacting part; by acting as parameters to it they determine how it shall react to the environment". The definition does not state that an immediate effect is necessary from the step-mechanisms to the reacting part, unlike the diagram of immediate effects of the ultrastable system. Neither is there a double-feedback loop of the form shown in the diagram. Nevertheless, the system is demonstrated to be ultrastable and does comply to the definition if the immediacy of the effect from the step-mechanisms on the reacting part is dropped.

The system can now be analysed in light of what was previously discussed (§3.1.2), that the disturbance to the essential variable is of unique importance. If a parameter changes to the system, which immediately affects Unit

⁴Incidentally, this also rules out the situation where both units are interpreted as the reacting part. Although this would have additional problems not discussed.

3.1. THE PROBLEM: WHEN TO CHANGE

1 (or the essential variable), then the critical states of Unit 2 *will* change. If a parameter change occurs to the system immediately affecting Unit 2, then the critical states of Unit 2 *will not* change.

In other words, it is significant *how* the parameter change affects the system, i.e., specifically which part is first affected. Both these changes are parameter changes, both may be permanent or impulsive, but they have very different effects and will not necessarily correspond to a clear demonstration of stability and ultrastability as promised. The critical states of the reacting part *change* depending on the rest of the system and the environment.

The Homeostat demonstration discussed is therefore a special case, where the critical states of the reacting part do not change. It is for this reason *only* that dashed lines can be added to the time-series. The fact that the reacting part directly affects the variable affecting the essential variable is highly atypical. The parameters to the reacting part not only determine how it shall react to the environment, but it determines how the environment shall react *to it*, and thus determines the disturbance to the essential variables. This is not mentioned in the definition of ultrastable system. If the environment changes the current reaction may no longer be appropriate, in which case the critical states will change.

Likely due to the generality of the types of disturbance defined, this was not clearly identified. As previously mentioned, the diagram of immediate effects of the ultrastable system separates the effect on the essential variables and the reacting part, but the definitions of disturbance apply equally to the reacting part and the essential variables. The system discussed receives *one* type of continuous disturbance, to the reacting part, but there is the potential for another continuous or step-function disturbance to the variable affecting the essential variable, or to the essential variable itself. The definitions do not however imply this. Modelling the latter disturbance would allow a situation to be described *precisely* where the environment changes and the current reaction is no longer appropriate. It seems intuitive that this situation is in fact likely to be the most frequent to occur. The critical states of the reacting part will likely change. It is possible that this was not made explicit in the framework due to limitations of the Homeostat.

However, in other demonstrations of the Homeostat, where disturbances *were* applied to units directly effecting the essential variables, the critical states of the reacting part(s) were not discussed; while there are no dashed lines in the corresponding diagrams, implying that the critical states (critical surfaces) were not fixed or defined, there is no mention or discussion of why this is the case or why it might be important.

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

It is now clear that the relationship between the disturbance to the essential variables and the critical states of the reacting part is not directly taken into account by the framework, nor discussed. It will be argued in the following sections that doing so can be useful for a number of reasons. It will therefore be argued that for this additional reason the disturbance to the essential variables should play an explicit role.

3.2 Potentially Ultrastable Regulation Units

In this section a new system is introduced to solve the aforementioned problems with disturbance, timing and the critical states. The system introduced, referred to as a *potentially ultrastable regulation unit* (PURU), in its *minimal* form, can be classified as an instantiation of the ultrastable system defined by Ashby. It is *potentially* ultrastable in that the necessary components for it are contained, but as an *open system* it is reliant on the environment to achieve it. This will be significant when considering the different types of stability demonstrated when the system is disturbed in different ways. The diagram of immediate effects for the PURU is shown below (Figure 3.5):

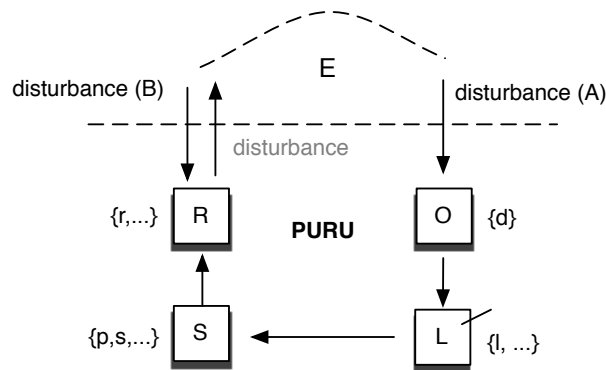


Figure 3.5: Potentially Ultrastable Regulation Unit (PURU).

The system appears somewhat similar to the original ultrastable system (Figure 3.1, page 29), but there are several subtle modifications. Firstly, the system now includes no environment-part, so there can be no confusion over environments of environments. For this reason the system is immediately more extensible; for example, if modelling an embodied agent, the environment may be *another part* of the agent or the actual environment external to the agent. This was of course possible before, but it is now made more explicit as the emphasis is shifted. The reacting part is indicated by the

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

box ‘R’ in [Figure 3.5](#). Each box in the diagram represents a set of variables, where an arrow from one box to another indicates that some variable in the first box has an immediate effect on some variable in the second box.

Secondly, the essential variables have been replaced by lability variables (box ‘L’ in [Figure 3.5](#)) in accordance with the discussion of [§2.1.2](#). This necessary change allows non-lethal variables to trigger changes in behaviour, that do not have physiological limits that may kill the organism. The lability variable therefore only requires one limit, unlike the essential variable. The critical states of the lability variables are fixed at this limit if the system is properly formed.

Thirdly, the previous disturbance to the essential variables (now lability) from either the environment-part or environment, not previously referred to independently, is now modelled as an *optimality variable* (box ‘O’ in [Figure 3.5](#)). This variable can be thought of as a direct feedback signal (see [§3.1.1](#)), indicating the *immediate* optimality or appropriateness of the reacting part. If the coupling is appropriate, it may be possible to define a time-varying optimality landscape: a mapping of each point in the state-space of the reacting part to an optimality value, indicating the value the optimality variable would hypothetically take if the point represented the current state at the time the mapping was defined. This will be further clarified in [§3.2.2](#).

When the optimality value is high, typically the lability variable will be high, i.e., a disturbance implies a change may be necessary, and the system is thus more liable to change. The behaviour of the reacting part is therefore immediately optimal when the optimality variable has a stable value of zero. The way in which the optimality and lability variables interact will be discussed in the following subsection ([§3.2.1](#)), and used as a basis for a timing mechanism to vary trial length.

Finally, the disturbances have been reformulated (see [Figure 3.6](#)). The disturbances now fall into two categories: type A, where a parameter in the environment changes that may or may not be a parameter *to* the optimality or lability variable, but nonetheless has an immediate effect on one of the variables; and type B, where a parameter in the environment *to* the reacting part or step-mechanism is changed resulting in a different immediate effect. It is necessary to acknowledge that type 2A disturbances (see right-hand side [Figure 3.6](#)), for instance, may be a permanent change in the environment, while the disturbance (feedback) will be temporary, ending when the reacting part is changed appropriately and adaptation occurs. The two ‘disturbances’ are not identical. The spatio-temporal pattern or functional type of the disturbances and described as type 1 and 2, in accordance with the

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

previous definitions provided by Ashby. What is now made clear is that type A disturbances will tend to have a different affect on the critical states than type B. This will also be demonstrated in an example model in the following section §3.2.2.

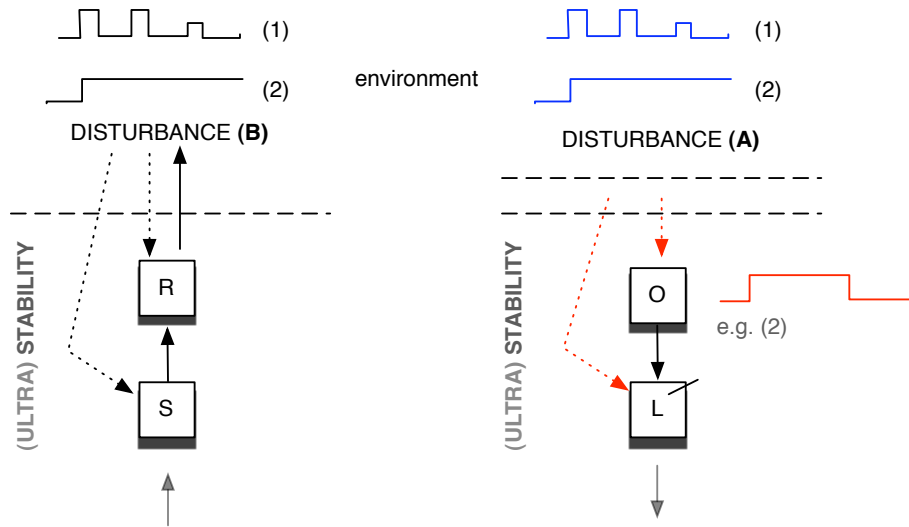


Figure 3.6: Reformulation of disturbances.

3.2.1 Trials and Critical States

With reference to the critical states of the lability variables and reacting part, the issue of timing can now be more clearly considered. What is sought is a mechanism which allows the values of the step-mechanism to persist for some duration of time, before changing to another if it is found to be inappropriate.

In the original framework no explanation for periodic changes in step-mechanisms is provided in terms of critical states. This is potentially problematic because the critical state of a step-mechanism is defined as a state that, when reached, will cause its value to change; and yet when an essential variable remains beyond its limits, beyond the critical surface (in a critical state), the step-mechanisms *do not* change at every time-step. They change only when the surface is transgressed, and then periodically. This is not however contradictory, since subtle to the definition of the critical state is that it applies to a particular *value* of a step-mechanism. However, what changes the states beyond the limits from critical states *not* associated with the current value of a step-mechanism to those that *are* is not at all obvious. It seems sensible that instead the representative point (current state) should be *moved away*

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

from the critical states, and that those states should always be critical states to *all* values of step-mechanism.

To do this in an essential variable however, to force its value to decrease or increase, would be equivalent to magically regulating the physiological variable that is being threatened; to regulate it artificially. However, in the PURU this can be done without problematic implication to the lability variable, since it is independent of the physiological survival-relevant variable, and the disturbance to it. The regulated variables can remain unaffected, as can the crucial feedback signal. The lability variables then, when passing their fixed limits into the critical states only remain there for a single time-step, no matter what value the step-mechanism has. The emphasis is thus shifted, and the lability variable provides an indication of the systems susceptibility to change, i.e., its lability. When the limit is reached the system *must* change, but immediately following a step-change it need not, as it is functioning normally, despite the immediate optimality. In other words, after a step-mechanism changes the value of the lability variable can decrease sharply, and the time taken to reach the critical state once again can represent the trial length.

This not only provides a basis for a timing mechanism, but also an indication of *ultimate* optimality, i.e., optimality of the situation beyond the immediate, indicated by the optimality variables. To provide a concrete example for clarity, the optimality variable could represent ‘level of pain’, while the lability variable could represent ‘discontentment trying the current behaviour’. Despite high pain levels, when a new behaviour (trial) begins, the lability is low as the system is content trying the current behaviour. The lability of the behaviour is low as the system is performing how it should be performing. If the pain levels remain high, the discontentment (lability) will gradually rise, nearing the critical surface, indicating that the optimality of the ultimate situation is worsening. This will continue until the critical state is reached. The cycle will continue until the disruption to the optimality variable is such that the lability variable is no longer increased. This will be made clear by a model demonstration in the following subsection (§3.2.2).

3.2.2 Demonstration

The proposed solution of the previous subsection can now be demonstrated by a model agent-environment system implementing the PURU. Consider a scenario where an agent can move within a finite 1-dimensional environment. The agent has a single viability variable, temperature, that it must

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

keep within bounds in order to survive. In order to keep its temperature from rising, the agent must remain with a shaded region of the environment (see [Figure 3.7](#)). The environment is such that the shaded region moves periodically, and instantaneously when it does so.




Figure 3.7: Model scenario: an agent moves in a 1-dimensional environment seeking a shaded zone to regulate its temperature.

The environment E interacts with the PURU Ξ consisting of 4 units, defined as follows:

$$\Xi = O \cap L \cap S \cap R, \quad E \rightarrow O, E \leftrightarrow R \quad (3.2)$$

where O is the optimality-unit, L is the lability-unit, S is the step-unit and R is the reaction-unit. The units correspond to [Figure 3.5](#). Where a set of variables is used in combination with notation defining an immediate effect, there must be at least one variable in that set to which the effect applies⁵.

The position of the agent in the environment r can range between -1 and 1, defined as follows:

$$\frac{dr}{dt} = s(p(t) - r(t)) + \gamma(t) \quad (3.3)$$

where $s = \frac{1}{2}$ is the speed of the agent, p is the value of the step-unit $S = \{p\}$, which is parameter to the reaction unit $R = \{r\}$, i.e., $p \rightarrow r$. Both units

⁵The immediate effect notation (arrows) was used by Ashby only with variables. The notation will be used with sets here for convenience.

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

contain only one variable. The variable $\gamma \in E$ represents the noise from the environment; thus $E \rightarrow r$, and in the opposite direction, the reaction-unit has the immediate effect $r \rightarrow E$. The constant $\frac{1}{2}$ is set arbitrarily, to determine the speed at which an agent can move in the environment.

In the simplest case, the optimality variable δ at time t , in the optimality-unit $O = \{\delta\}$, is defined as follows:

$$\delta(t) = \begin{cases} 0 & \text{inside shaded zone} \\ 1 & \text{outside shaded zone} \end{cases} \quad (3.4)$$

where $E \rightarrow \delta$. The lability variable l at time t , in the lability-unit $L = \{l, w\}$, is defined as follows:

$$\frac{dl}{dt} = \frac{1}{2}(\delta(t) - l) - 10w(t) \quad (3.5)$$

where δ acts as a parameter to L , thus $\delta \rightarrow l$, and w is the *relapse variable* of L . The relapse variable w is set to 1 whenever $l \geq 0.95$, at which point a step-change will occur and p will change to a random value between -1 and 1. At all other times the relapse variable w decays as follows:

$$\frac{dw}{dt} = -k \quad (3.6)$$

where $k = 2$ is a decay constant. Finally, the temperature of the agent $v \in E$ at time t with a maximum value of 1 is defined as follows:

$$\frac{dv}{dt} = \begin{cases} -\frac{v}{10} & \text{inside shaded zone} \\ \frac{1}{100} & \text{outside shaded zone} \end{cases} \quad (3.7)$$

3.2.3 Optimality Landscapes

Due to the simplicity of the demonstration, an *optimality-landscape* can be defined based on the disturbance to Ξ , by the value of δ when r has a particular value. Three possible landscapes are shown below (see [Figure 3.8](#)).

The coupling of the optimality variable to the environment defines the landscape, and is such that when the agent moves away from the shaded region the disturbance to it generally increases. In this case the viability variable

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

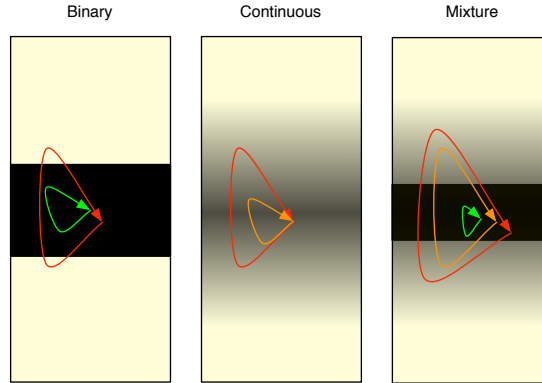


Figure 3.8: Optimality landscapes based on hypothetical disturbance received in different states at a particular time.

does not affect the PURU, although this could be a valid alternative or additional strategy. The diagram of immediate effects for Ξ in these two possible situations is shown in Figure 3.9 below. In the demonstration only the vertical blue arrow ($E \rightarrow \delta$) is present.

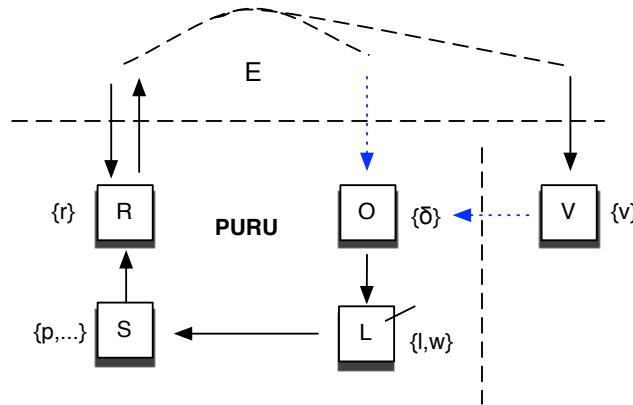


Figure 3.9: Two possible PURU optimality variable couplings.

The simplest “binary” time-varying optimality-landscape ∇ shown in Figure 3.8 (left) can be defined as follows:

$$\nabla = [0, 1] \times [0, 1] \tag{3.8}$$

where

$$\nabla(t) = \begin{cases} 0 & |r(t) - c(t)| \leq \frac{s(t)}{2} \\ 1 & |r(t) - c(t)| > \frac{s(t)}{2} \end{cases} \tag{3.9}$$

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

where $c \in E$ is the position of the centre of the shaded zone and $s \in E$ is the size of the shaded zone at time t . If this function is implemented the optimality variable can only take one or two values at any one time: 0 or 1.

Simulation

To demonstrate the system with the landscape ∇ , a simulation was carried out implementing the equations of Ξ using Euler integration. The following graph (Figure 3.10) shows how the agent behaves in an environment that changes every 50 units of time (500 time-steps) over 500 units of time (5000 time-steps). The shaded region moves each time the environment changes, and thus the optimality landscape changes as described by Equation 3.10 above.

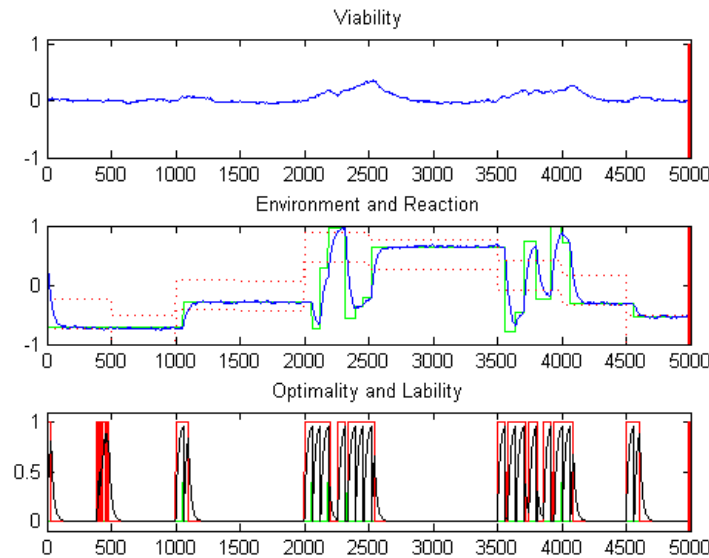


Figure 3.10: Demonstration with binary optimality landscape.

The boundaries of the shaded region are indicated on the graph by red-dotted lines. When the environment changes, e.g., at time-step 2000, the optimality landscape changes such that the current state subsequently has a high value; the optimality variable increases correspondingly due to the change in environment parameter (type 2A disturbance). It should be noted that the parameter that changes in the environment is *not* a parameter to δ , but has a very direct effect on it. After the disturbance the viability

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

variable begins to gradually increase as it is now outside the shaded zone. At time-step 2500 when the environment changes again, the agent has not found a suitable parameter, but luckily the viability variable had not reached its limit of 1. Shortly after, the agent happened to find a parameter that was suitable for the next environment, so its position remained stable. At time-step 3000 the environment changed once again, but the location of the shade only shifted negligibly; no adaptation was necessary in this case.

The black lines showing the time-series of the lability variables decrease sharply each time the value reaches 0.95, the critical surface. Each time this occurs the value of the step-unit p changes randomly, resulting in a change in behaviour (of the reaction-unit), as p is parameter to r . In the example shown the agent survives for the full 5000 time-steps, managing to keep the viability variable within limits. The length of each trial is kept relatively constant, equal to the time it takes for the lability variable to increase from 0 to 0.95, and the speed of the agent is well matched. Had this not been the case adaptation may not have been so successful.

As the demonstration shows, the timing of behavioural change or trial length is now part of the system dynamics, primarily dependent on the strength of coupling between l and δ (see Equation 3.5). Also relevant is the relapse-time, which is dependent on the rate of decay of w . There are however many ways the trial length could be controlled, and there are many possible equations or couplings that could result in a similar adaptive system. What is important is of importance is the distribution of the critical states, as these must remain fixed. A behaviour must persist while the lability variable remains distant from its fixed limit, where the critical states lie for *all* values of the step-unit (step-mechanism) containing the parameter which determines the behaviour.

A trial should last no longer that is sufficient to receive feedback, i.e., the corresponding optimality signal. In the demonstration this is dependent on the speed of the agent. For example, if the agent suddenly was only able to move very slowly the trial length would need to be extended in order for adaptation to be successful. This would be the case because the optimality signal would remain high until the agent had moved into a shaded region, which may take some time after the correct parameter had been selected. This can be demonstrated in the simulation (see [Figure 3.11](#)). When a new position is selected (parameter, shown in green) the speed of the agent is insufficient to reach it before the trial ends and a new position is selected. For instance, in the third environment around time-step 1100 a suitable parameter (behaviour) is found, but the agent does not receive the necessary feedback in time to ensure its stability and avoid the critical states, i.e., the

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

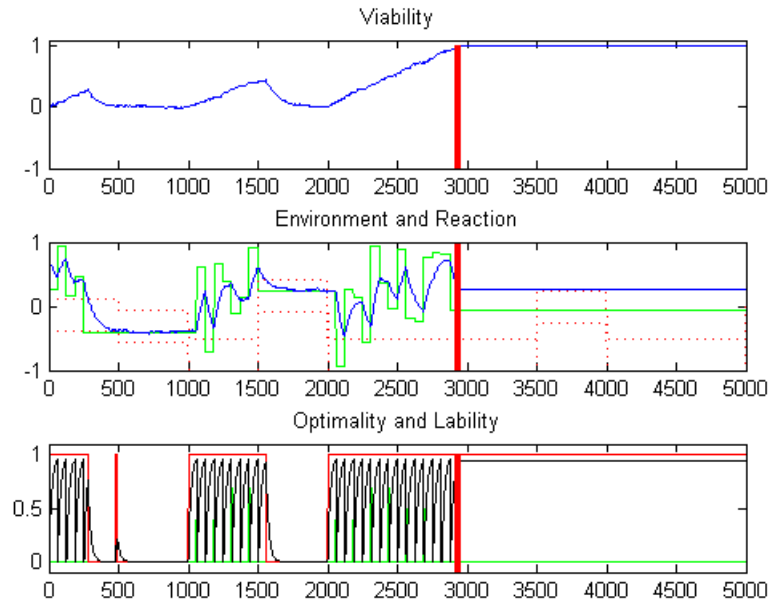


Figure 3.11: Demonstration with binary optimality landscape. Trail length is too short for the slow speed of the agent. Viability is lost around time-step 3000.

agent does not have time to move into the appropriate region of the optimality landscape to decrease its lability; the agent might be seen as ‘giving up too easily’, rather than persevering in its behaviour. The performance of the agent can be tested of a range of different speeds to clarify this is the case: Figure 3.12 shows the survival time, average optimality and environment adaptation times for a variety of different speeds ranging from 0.01 to 0.4, tested over 50 repeats. The results clearly indicate that the slower agents survive for a shorter period of time on average, and that the faster agents consistently survive for longer periods of time. However, since the trial length is now part of the system dynamics it may be varied, either by modification of a fixed parameter or automatically by the system dynamics. The following subsection will investigate the latter case, where the trial length is regulated automatically.

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

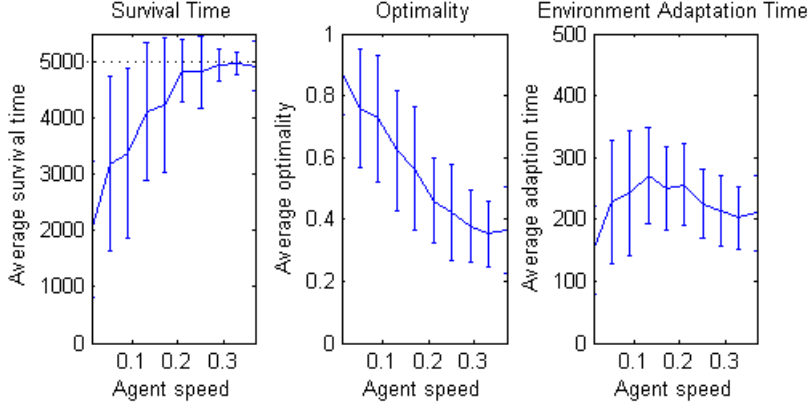


Figure 3.12: Demonstration with binary optimality landscape, for a range of speeds.

3.2.4 Perseverance

In order to know which behaviours to persevere with and which to abandon, the system must have some way of evaluating its performance *during* a trial. This must be the case, as the speed of the agent may be or may become too slow for multiple trials to be carried out in a single environment (see Figure 3.11). The system must therefore receive some extra optimality-signal or disturbance during each trial, and this disturbance must affect the lability variable. It must be decreased in situations where the trial seems more hopeful, where perseverance is beneficial.

If the time-varying optimality-landscape ∇ is redefined so it more closely resembles the mixed landscape of Figure 3.8 (right), the optimality will vary continuously:

$$\nabla(t) = \begin{cases} 0 & |r(t) - c(t)| \leq \frac{s(t)}{2} \\ |r(t) - c(t)| & |r(t) - c(t)| > \frac{s(t)}{2} \end{cases} \quad (3.10)$$

The agent will now receive a decreasing optimality-signal when it is moving in the direction of the shaded zone, so it may use the rate of change of the signal to vary the lability. Equation 3.5 describing the dynamics of l can be modified appropriately:

$$\frac{dl}{dt} = \delta(t)(\delta(t) - \delta(t - t') + 2) - 10w(t) \quad (3.11)$$

3.2. POTENTIALLY ULTRASTABLE REGULATION UNITS

where $t' = 20$. The delay differential equation allows the optimality signal of a previous time to be taken into account, so that the lability increases and decreases based on the increasing or decreasing of the optimality signal. There are of course many other ways this could be implemented, for example, the relapse rate could be changed, or the optimality signal itself could indicate the rate of change. **Figure 3.13** below demonstrates how an agent moving at the *same speed* behaves under the new equations. The trial length is

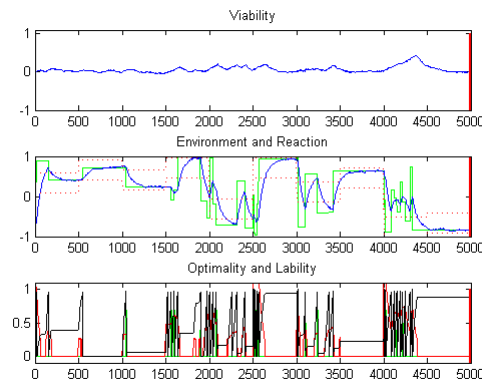


Figure 3.13: Demonstration of a slow moving agent adapting with a variable trial length.

seen to vary appropriately, lasting shorter periods of time where the agent is moving in the wrong direction. For example, around time-step 2500 when the environment changes a parameter is selected that moves the agent within the shaded zone, and rate of increase of the lability variable decreases (see **Figure 3.14**). If the simulation is run multiple times over a range of different agent speeds, the results indicate that successful adaptation occurs for a larger range of speeds than the previous strategy (see **Figure 3.15**); the standard deviations of the survival times are smaller, and the mean survival time is higher beyond ≈ 0.15 . The mean optimality is also significantly reduced for all speeds.

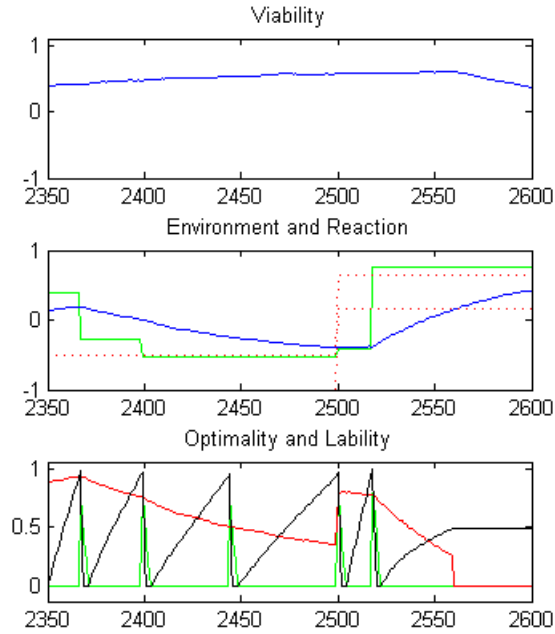


Figure 3.14: Perseverance: automatic variation of trial length. The green spikes indicate a relapse, the start of a new trial.

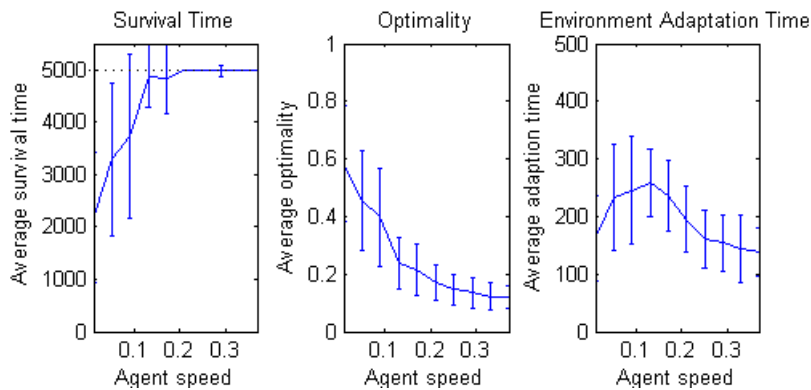


Figure 3.15: Demonstration with mixed binary-continuous optimality landscape, for a range of speeds, where trial length is varied automatically.

3.3 Discussion

The problem of trial length and temporality of adaptive change in the ultra-stable system was addressed by considering the relationship between distur-

bance to the ultrastable system and the critical states of the reacting-part. The correspondence of the Homeostat to the ultrastable system was also critically considered, and it was shown that the essential variables were not in fact directly modelled, nor was the disturbance to them made explicit.

The lability variables introduced in §2.1.2 were used in a modified version of the ultrastable system. The novel system, referred to as PURU, was defined in both to make what was argued more explicit and to provide a more extensible version of the ultrastable system. It was hoped that the PURU may be more easily expanded, and potentially used as a component in a larger system. Using the concepts implemented in PURU, an ultrastable system was demonstrated that was able to vary its trial length automatically. Two types of coupling were investigated, but it was acknowledged there are many possible ways similar overt dynamics could be achieved.

By taking into account how the critical surface of the reacting part changed in relation to the different types of disturbance to the system, *two* different situations could be identified where ultrastability may occur: when a perturbation is applied to the reacting part, moving the system towards a critical state, or when a perturbation is applied to an essential variable, where the critical states are themselves changed. In the example model, this is would be illustrated by distinguishing between two situations: the first, where a disturbance occurs and the agent is forced out of the shaded region, either temporarily or permanently, and the second, where the shaded region moves away from the stationary agent, either temporarily or permanently. In the latter case it is the optimality-landscape that is changed. This was not previously easy to identify because it requires explicit reference to the disturbance of the essential variables. It was thus subsequently identified that in fact the two types of disturbance defined in the framework will not necessary correspond to a demonstration stability and ultrastability respectively.

By directly acknowledging that the critical states of the reacting part were time-varying, reflected by an optimality-landscape, trial length could be thought of as an *internal* modulation of these states, achieved by moving the lability variables away from their *fixed* critical states; this was of course not possible in the original framework because it would have involved changing the value of the essential variables, which would have implied a physiological consequence. The modification also solved the previous problem in the framework where states would become critical states after some duration of time by an unknown mechanism. The modification meant that the critical states applied to all values of step-mechanism, rather than some value in particular. It was argued that this was in fact more intuitive, since any

value of step-mechanism when the critical state is met should be considered inappropriate at that moment.

3.3.1 Future Work

By considering the optimality-landscape, it was also possible to arrive at a definition of *adapted*, as an agent in a situation where there is no disturbance to the optimality (or essential) variables, i.e., zero optimality. This may have been impossible to model using the Homeostat, as it would only occur when the essential variable(s) are in equilibrium at zero, while the reacting part is still being disturbed. It may also be sensible to further distinguish between the immediate and ultimate situation; e.g., immediate and ultimate-adapted, since the lability variable may have a value greater than zero, indicating that the ultimate situation is not optimal. Further work needs to be carried out in order to investigate whether a hierarchy of behaviours would imply the need for more lability variables that vary at different time-scales. It may be the case for example, that the ultimate optimality of the lability variable is in fact immediate in relation to a slower changing lability, which corresponds to a behaviour (or behavioural mode) subsuming the one persisting within the trial, controlled by the regular lability variable. Such a meta-behaviour may be better referred to as an ‘approach’.

What also remains to be investigated is the concept of rewards. It would be quite easy, for example, to add multiple optimality variables to a PURU but invert the coupling so that high optimality decreases lability. This would be equivalent to a reward, encouraging a behaviour to persist. Furthermore, having introduced continuous optimality, it is now possible for an agent to take into account varying levels of immediate optimality, e.g., pain, such that it would be possible to accumulate low signals over some time. This may account for scenarios where an agent gets ‘bored’ of a situation, and changes its behaviour accordingly. In order for this to be implemented the lability variable would have to have a decay constant and an appropriate coupling to the optimality variable.

Having separated the essential variables into viability and regulation variables, and further distinguished types such as lability and optimality, it may be possible to make certain types of maladaptation more explicit. For example, to account for such situations where individuals may experience injury but no pain, or feel pain when there is no injury; such situations could be thought of as a malformed coupling between lability and optimality variables, or optimality variables and viability variables.

3.3. DISCUSSION

So far only simple environments have been taken into account. In reality the environment is not only more complex, but commonly includes other intelligent agents. It would be interesting to investigate how ultrastability and the PURU might be used in the context of collective robotics. The concept of suitability, for example, could be used to differentiate ‘true altruism’ from helping behaviour if regulation in an agent was found to be better aligned with the *suitability* of another agent rather than itself, i.e., the survival of other agents. Situations could also be investigated where high optimality signals (disturbance) are present in agents when other agents are observed to be ‘suffering’, or more generally by sensing when their lability variables are high. This would provide a basic mechanism for empathy. This quite naturally would lead to agents regulations each other’s variables, and might be considered to be a case of non-altruism as the helping agent would also decreasing its own lability variables. Antisocial behavior could also be investigated where the optimality variables triggered by the suffering of others are ‘out of sync’, i.e., a failure of empathy, where for example, the suffering of another agent may decrease lability or have no effect at all.

4

Adapting in Time

In the previous chapter the problem of *when* to change was addressed, but the nature of those changes was seldom considered. In this chapter the problem of *how* to change will be addressed, with an emphasis on efficient adaptation.

4.1 The Problem: How to Change

The nature of changes is a temporal concern on the assumption that feedback is not always immediate; that behaviours need to be tried over a duration of time before their appropriateness to the current situation can be evaluated.

Since each change has the potential to bring about a correct behaviour, the nature of each change is significant. Up until this point in the thesis, and in the original framework, all potentially adaptive changes have been random selections of a parameter, i.e., the process has been one of trial and *error*. Each change is a change in dynamics and hence behaviour, that will occur repeatedly until the the behaviour is an adapted one; but the longer an agent spends trying different behaviours, the greater the risk is of its survival being threatened. Since time to adaptation may be critical for survival every change is significant.

4.1.1 Searching for Optimality

In the adaptive mechanisms of the original framework changes in behaviour are brought about by random changes in parameters that determine that behaviour, namely the step-mechanisms. When the space of possible parameters is small this is no problem, but if the behaviour is complex, requiring

4.1. THE PROBLEM: HOW TO CHANGE

many parameters, using random selection to find an appropriate set of parameters may not be tractable. The time to adaptation can be understood approximately in probabilistic terms in this case, if the proportion of the parameter-space containing solutions is known.

Using the terminology of the previous chapter, what is sought is an appropriate attractor landscape in the phase-space of the reaction-unit (reacting part) that keeps the current state in an optimal zone, i.e., that matches the current optimality-landscape, minimising the disturbance to the optimality variable and thus lability variable. At each moment the system must select the parameter leading to the most favourable phase-portrait (vector field) for the current optimality-landscape. If there are no such parameters then adaptation is impossible, but if there are then the proportion of ‘optimal parameters’ to ‘non-optimal’ will approximately determine the time to adaptation, at least if selection is random.

If the size of the parameters space is large a realistic adaptation time may be impossible, but this is only the case *if the majority of combinations of possible parameter values are non-optimal*, i.e., leading to a high optimality signal and further step-changes. This may or may not be the case. For example, in a system with a *finite* state-space of N parameters, if at any one moment each parameter $1 \leq i \leq N$ has its own independent optimal range, taking up a proportion p_i of its own total space, then the probability of selecting an appropriate vector in the state-space leading to an appropriate behaviour is:

$$\prod_{1 \leq i \leq N} (p_i)$$

If each parameter value has an independent probability of being chosen correctly 50% of the time, then the probability of choosing a suitable set of parameter values is $\frac{1}{2^N}$, requiring an expected 2^N trials. The number of trials, and thus time to adaptation, can be expected to rise exponentially as the number of parameters are increased. This is the case with the Homeostat, as acknowledged by Ashby (10/7).

As previously mentioned though, what is of importance is not the size of the parameter-space, but how much of that space is occupied by sets of parameter values that will or will not ultimately lead to disturbance of the lability variable. For an embodied agent, the structure *and environment* of the agent must determine this outcome, as the set of possible optimality-landscapes will be co-specified by the environment and coupling to it, i.e., the structure of the agent and nature of the environment. By virtue of its embodiment the agent must have a point-of-view, or frame of reference, imposing order

4.1. THE PROBLEM: HOW TO CHANGE

on the perturbations it receives as an open system. It will be argued that it is the redundancy in these perturbations that must be exploited.

4.1.2 Random Step-changes

In the previous example selection was assumed to be random; is it not possible to do better than trial and error? In the ultrastable system the random selection of parameters occurs in the step-mechanism. In the PURU it occurs in the equivalent sub-system: the step-unit. If randomness is to be avoided these parts at least must be modified. But if selection is not random, it must be based on something.

As previously mentioned, by being embodied an agent has a specific point of view and structure that will mean the perturbations it will receive will be of a certain type. While Ashby does not acknowledge this explicitly, he does note that: “the organism commonly faces a world that repeats itself, that is consistent to some degree in obeying laws, that is not wholly chaotic.” (10/2), and calls this the ‘recurrent situation’. The main idea is that if the environment were completely random and non-repetitive, the organism would have nothing to take advantage of, and there would be no advantage in non-random selection.

Ashby provides an example where the parameters to an ultrastable system change intermittently, but to values that have previously been adapted to by the system. In such a case he states that it may be advantageous to respond to familiar parameters in the same way each time, but that it would depend on the environment. If the correct response to a previously seen environment was in fact the same each time, then there may be a set of parameter values in the parameter-space that would be worth selecting on multiple occasions. If this is to be taken advantage of then the step-mechanism in the ultrastable system must change.

In the previous chapter the timing mechanism was brought into the dynamics of the PURU, based on the ultrastable system. The same can be done with the step-mechanism. Since deterministic dynamical systems are used, the step-mechanism can be thought of as a phase-space of dimensionality equal to the parameter-space, where the current state represents the current selected parameter, in the basin of attraction of a stable fixed-point attractor, i.e., the parameter is stable and unchanging. When a critical state is met the parameter(s) change, and a bifurcation occurs whereby a new stable point is created and the current state moves towards it. This must occur rather

quickly, and should not be confused with the transient period referred to as a ‘trial’. A trial must be thought of as the time the *location* of the fixed-point attractor remains stable, i.e., the stability of the stable point. The transient periods where the current state moves towards a new attractor after bifurcation is not modelled by the step-mechanism; changes to the parameters are instantaneous. From a dynamical systems perspective then, a change in parameter can be thought of as a bifurcation, resulting in a fast transient period from one stable fixed-point attractor to another.

4.2 Synergies

To take advantage of the recurrent situation, in general previous values of parameter must be returned to with a probability greater than random. For example, if an ultrastable system (or PURU) has two parameters x and y in its step-mechanism (step-unit), ranging from -1 to 1, and x often takes the value 0.2 and 0.8 in alternation, it may be sensible to select these values more frequently by biasing the usual random selection. The two values might be suitable in two different commonly occurring environments.

The same could be said for another parameter and two associated values. However, there is a problem in selecting parameter values based on their previous history e.g., how long a particular value was stable. The problem lies in that it does not relate the parameters to each other; they are being treated independently. This would not be useful where, for example, x may have the value 0.2 only when y has the value 0.9, and 0.8 only when y has the value 0.1. For the system to increase the probability of this selection, a 2-dimensional space is necessary with two significant points, rather than two 1-dimensional spaces with two significant points each. In other words, the individual values the parameters may take are not necessarily important, it is their relationship to each another that is important, i.e., the points of interest are salient points in the parameter-space, associated with optimality, and this space is likely to be a hyper-space of several dimensions.

In order to increase the tractability of an appropriate selection of optimal parameter values, several of such salient points in the parameter-space may need to be considered based on the history of a system. A salient point in the parameter-space associated with optimality will now be referred to as an *assembly*. If a collection of assemblies are used to form a new 1-dimensional space, where each point in it corresponds to a different assembly, then this space of salient points could be searched instead of the original parameter-

4.2. SYNERGIES

space, which may be orders of magnitude larger. A set of assemblies will now be referred to as a *synergy*. The problem of selecting an appropriate set of parameter values is now a matter of *selection of a stable-point in the parameter phase-space*, marked by an assembly.

For example, consider the space $Q = [0, 1]^3$, and three parameters x , y and z that range from 0 to 1. Let the three parameters belong to the step-unit S of a PURU Ξ , that determines the behaviour of a very complex reaction-unit $R \in \Xi$. Consider now that the possible optimality-landscapes are such that only two small regions of Q lead to adapted behaviour (zero optimality and lability), and that the optimality-landscape changes from one to the other occasionally but repeatedly with type 2B disturbances. For example, the two regions might be around the points $a = (0, 0, 0)$ and $b = (1, 1, 1)$. If this was the case it would be inefficient to select a point from the entire parameter-space Q randomly each time the environment changed. Rather, a synergy $\Upsilon = \{a, b\}$ could be formed containing the two vectors (assemblies) a and b , corresponding to each environment (see Figure 4.1).

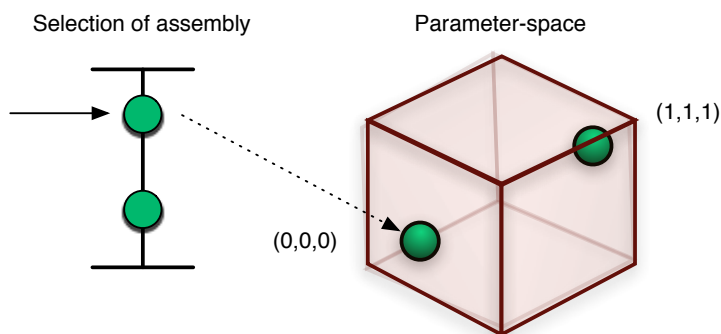


Figure 4.1: A synergy with two 3-dimensional assemblies.

Selection of a new parameter upon entering a critical state would then involve selecting an assembly in the synergy. However, it may be the case that the alternation of the two environments may be temporary, or that the two points a and b corresponding to appropriate sets of parameter values cease to be optimal for the corresponding environments. In this case, the system will fail to adapt because it will select only the assemblies in Υ . There needs to be some *exploration* as well as *exploitation*.

In order to solve this problem random assemblies could be added to Υ , or the existing assemblies in the synergy could be modified, or even randomised, displacing the points they represent in Q . If the latter occurs, and an assembly is changed, then any potential of accelerated adaptation to an environment

associated with that assembly will be lost. While if the former occurs, more ‘space’ or ‘memory’ will be required, and since only one assembly (behaviour) can be tried at any one time, a trade-off will be introduced between exploration and exploitation. If or when the previous environments adapted to eventually return, the associated assemblies will not be selected so quickly, since there will be other assemblies to potentially select.

In order to bring the concept of synergies into full correspondence with the PURU (see Figure 4.2), instead consider that $S \in \Xi$ is replaced with P , to contain a vector of parameter values to R , i.e., $P = \{x, y, z\}$ and $P \rightarrow R$. Having removed the step-unit, the synergy Υ can be added to Ξ as a unit to select the values of the parameters in P , i.e., $\Upsilon \rightarrow P$.

The selection of assemblies must now be considered, and two different types of system will be distinguished: *intrinsically* modulated and *extrinsically* modulated. The former will be briefly demonstrated by an example model, while the second will only be discussed.

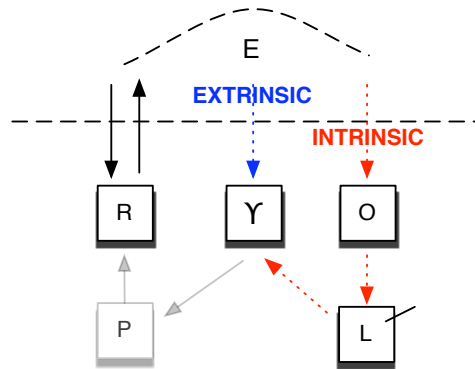


Figure 4.2: An example PURU using a synergy for reaction-unit parameter selection. The synergy may be modulated intrinsically, extrinsically or by a combination of the two.

4.2.1 Intrinsic Modulation

Having replaced the step-mechanism with a synergy-unit, the selection of assemblies must now be considered. An assembly is said to be selected if its values are used to parameterise a system; in the case of the PURU this would be the reaction-unit, or in the case of the ultrastable system this would be the reacting-part. Only one assembly can be selected at any one time, since it is

assumed that the parameter values are always stable; the transition between stable points is not modelled.

Where intrinsic modulation is used, the *only* channel of communication coming to the synergy is via the lability-unit (see [Figure 4.2](#)), or essential variable, in the case of the ultrastable system. This means the synergy only receives periodic signals when a critical state is met. In other words, selection of assemblies must be based primarily on their previous stability. For simplicity, the stability of the assemblies will now be considered to be binary: either they are stable or unstable, and in the latter case their values will change randomly. The two types can be thought of as *associated* and *non-associated* to an environment respectively. That there are unstable assemblies is necessary so that new environments may be adapted to; the assemblies must ‘move’ in order for the space to be explored. Initially this is particularly important, since every environment will be ‘new’.

In the simplest case, once an optimal set of parameters is found leading the system to stability, it must be marked by a stable assembly. It is a necessity that the representative assembly is made stable, and continues to be, so long as there is a probability greater than random chance that the region marked by the assembly will at some point again be optimal in the future. To clarify then, selection of an assembly models a loss of stability in the representative point (current state) of the parameter phase-space, while the loss stability of an assembly models the loss of stability of *a stable-point* in the parameter phase-space.

What now must be taken into account is how exactly the selection between assemblies is made. It must occur when a critical surface is as this is when the parameters to the behaving part must be changed, as the system is indicating the current behaviour is no longer appropriate. At this moment, the environment may have either changed to one familiar that has been previously adapted to, or it may be novel. For intrinsic modulation the information indicating which environment is present is not immediately available, so trial and error must be used.

Demonstration

In this subsection a specific case is demonstrated where assembly selection is probabilistic, but biased towards selecting parameter values that were previously stable (associated to a previous environment). A fixed number of assemblies are used, which all begin as non-associated unstable points, with

4.2. SYNERGIES

changing values. For simplicity, the system used to demonstrate the concept is based on a 2-unit Homeostat model (see §3.1.3, page 32), with the addition of bias values (Equation 4.1 below):

$$\frac{dy_i}{dt} = \sum_{j=1}^N w_{ji}y_j + b_i + \gamma_i \quad (4.1)$$

where $i = (1, 2)$, where y_i is the value of unit i , and the weights w and biases b are parameters to y_i and y_2 . The parameter γ_i is noise from the environment to unit i . The minimum and maximum value for all variables is -1 and 1 respectively. In total there are 6 parameters, so the parameter-space is finite and 6-dimensional. This is the space that must be searched.

The parameters will require changing periodically when a *type 2B* disturbance modifies the optimality-landscape. For simplicity, details of the PURU or ultrastable system will not be further discussed. The treatment will be general, although the variables y_1 and y_2 can be thought of as constituting the reacting, behaving part of the system. Fixed trials are used to periodically determine whether or not the current behaviour is ‘viable’, i.e., in the optimal zone, where no disturbance would be applied to a lability or essential variable. The optimal zone can be thought of a shaded region in a finite 2-dimensional environment that an agent must position itself in (see Figure 4.3); the position of the agent would be given by $y = (y_1, y_2)$.




Figure 4.3: An example agent-environment scenario for the synergy demonstration.

The optimal zone will change periodically reflecting two environments that alternate. The zone has an upper and lower limit, which change periodically between 0.8 to 0.2 and -0.8 to -0.2 respectively. If the values of y are *both* within these limits the system is said to be in the optimal zone, and thus stable.

Consider now the synergy Υ in the system containing five assemblies, each a vector of values that can be used to set w and b . When an assembly is selected its values are used to set the parameters. When *the system* is stable for some time, the selected assembly must too become stable and thus associated with the current parameter values. The association is also to the environment, albeit indirectly. The assembly must only be stable if it is assumed that in the future the same parameter values held by the assembly will be useful. When an unstable assembly is selected its values are randomised.

Since an assembly may be selected that only temporarily places the system in the optimal zone, perhaps moving through it, an assembly should only be made stable if it keeps the system in the optimal zone for some duration of time. Therefore, each assembly has an associated strength, which is increased while the system is in the optimal zone. When this strength reaches its maximum value of 1 the assembly is said to be stable, otherwise it will be considered unstable. All assemblies are therefore initialised with a strength of 0.

At the end of each trial, or when a critical surface is met, a stable assembly is selected with probability of $\beta = \frac{1}{2}$, so that 50% of the time a non-stable assembly is selected and the system explores the space using an unstable assembly. Out of the stable and unstable assemblies selection is equal, i.e., every stable assembly has an equal chance of being selected to every other stable assembly, and likewise for the unstable assemblies. This applies even if β is changed, which reflects only the *exploitation vs. exploration. trade-off*.

Figure 4.4 shows the adaptation time dramatically decreasing for both environments. The initial exposure to the first environment where the optimality zone is 0.8 to 0.2 results in adaptation after approximately 1.8×10^4 time-steps, but on the second exposure, adaptation occurs almost instantly, and assembly 2 is selected. The graph shows how this assembly was selected after adaptation occurred the first time the environment was presented. On subsequent occasions the same assembly is selected when the environment is present, and adaptation time is negligible in comparison to the initial adaptation time of 1.8×10^4 time-steps. The same occurs for the second environment, but only after the second exposure to the environment, where adaptation occurs in approximately 10^4 time-steps; this corresponds to the

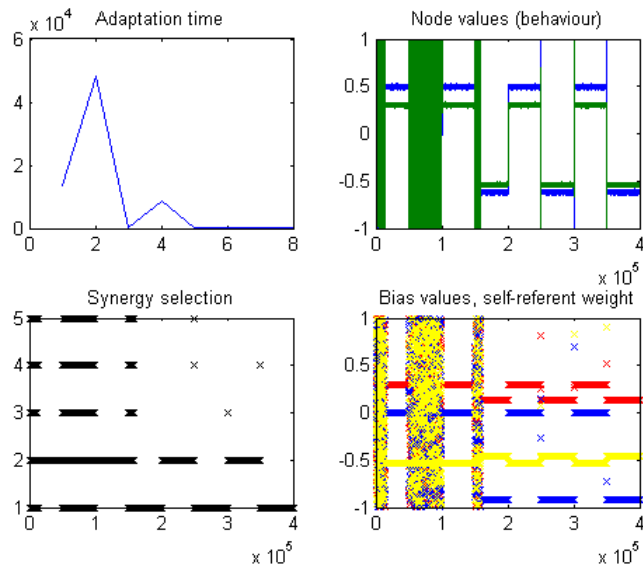


Figure 4.4: Adaptation to two alternating environments using a synergy of five assemblies.

smaller of the two peaks on the graph. Following this exposure to the second environment subsequent adaptation times were negligible, as with the first environment. The second environment uses the assembly 1 as a stable assembly.

The other three assemblies remain unstable, and are selected each time a critical surface is met (when outside the optimal zone) with a probability of $\beta = 0.5$. Adaptation would therefore be possible to up to three more new environments, should they be encountered. Adaptation using all five assemblies is shown in [Figure 4.5](#); after each assembly has been made stable and indirectly associated with an environment, adaptation is very fast, but only comparatively. In fact, it will take an expected 5 trials (250 time-steps) to select to correct assembly each time the environment changes, as each is equally likely to be selected, and there are five of them. The probability of choosing an incorrect assembly is in fact $\frac{4}{5} = 80\%$.

Furthermore, the fixed number of assemblies has limited the number of environments that can be adapted to in this way. If ‘open-ended’ adaptation is required it may be necessary to add additional synergies, or ‘lose adaptations’ to previous environments.

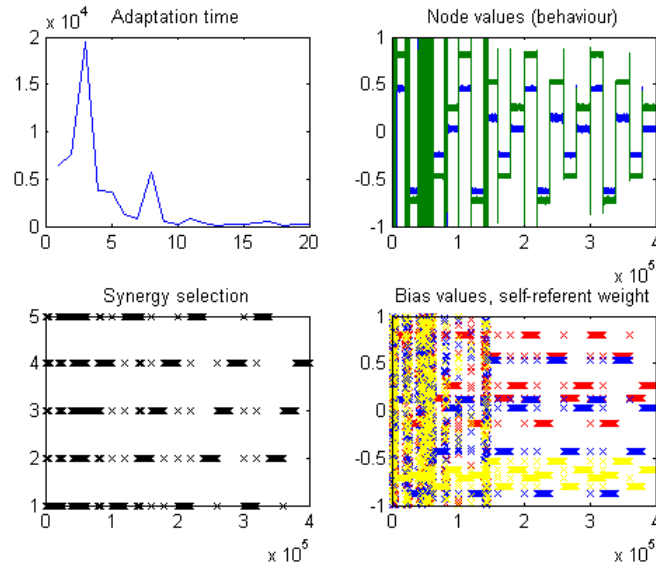


Figure 4.5: Adaptation to five alternating environments using a synergy of five assemblies. The allocation of assemblies to environments is random.

4.2.2 Extrinsic Modulation

In the previous subsection adaptation using an intrinsically modulated synergy was demonstrated. It was shown that although adaptation time to familiar environments can be significantly decreased, it may be problematic that several trials are still necessary to reach adaptation. The number of trials required to adapt to a familiar environment was approximately equal to double the number of stable assemblies, if stable assemblies were chosen 50% of the time.

So far however, the environment to the system has not yet been fully exploited. If another channel of communication is added to the system, so that the environment may immediately effect the synergy (see blue arrow in Figure 4.2, page 58), it should be possible to modulate the assembly selection *directly* based on what environment is present. In other words, the environment can be *categorised* and this categorisation can be used to bias the assembly selection. For example, the agent considered in the previous subsection seeking shade may make use of some sense that provides the current coordinates of the shaded region. These coordinates could be associated with each assembly, so when the coordinates are ‘perceived’ the probability of selecting the associated assembly is high.

4.2. SYNERGIES

To provide a further example, in a more complex system a *self-organising map* could be used, connected to a sensory input surface. The map could abstract features currently present on the surface in such a way that each distinct environment of importance activates a separate topographical region of the map, while in that environment. Each node constituting the map could then be directly associated with an assembly, such that the best matching unit (of highest activation) on the map is the assembly selected. The assembly could then be used to parameterise the system as normal.

This would be somewhat similar to what Ashby referred to as an *accumulator of adaptations* (10/9). In this system there are several step-mechanisms, a single essential variable and a mechanism that selects the step-mechanism appropriate to the current environment, to be “effective in the second feedback via the essential variables” (see Figure 4.6).




Figure 4.6: The accumulator of adaptations.

In this case the self-organising map would act as a gating mechanism to select an appropriate assembly, which could be thought of as corresponding to an individual step-mechanism. The details are however somewhat different, since the concept was only briefly introduced in the framework and not discussed in detail. In the original framework although it was not explained how an organism could acquire a gating mechanism (17/10) it was suggested that a *polystable system* could function as a gating mechanism (16/12). It was not however made obvious how this could be used to select a step-function. Unfortunately, no complete mechanism for or explanation of the accumulator of adaptations was provided.

4.3 Discussion

Due to the random step-mechanism of the ultrastable system, adaptation was shown to be inefficient when the dimensionality of the parameter space is increased. It was argued however, that what was of importance was in fact whether or not the possible parameters in the space when selected lead to behaviour appropriate to the current situation, i.e., lead to low values of optimality variable. The percentage of ‘optimal’ points in the parameter-space was of greater significance. It was argued and then demonstrated that if the environment is not totally random, it is possible to exploit the pattern of redundancy in the perturbations to the system: to exploit the ‘recurrent situation’. It was shown that salient points in the parameter-space may exist where patterns of optimality can be found to recur, e.g., the same point in the parameter-space (set of parameter values) leads to adapted behaviour at irregular intervals or regular intervals. The concept of a synergy was used, where selection was between such salient points, which were referred to as assemblies.

Of the two types of synergy considered, the extrinsically modulated synergy is likely to have the greatest practical application, and has the potential to demonstrate self-organised, multiple switching behaviours synchronized to multiple environments. In this case the optimality signal remains at zero when the environment changes, yet the reacting-unit or overt variables, the output, is changed qualitatively as re-parameterisation will occur. A self-organising map or classification layer to abstract salient features from the environment may allow behaviour to be extrapolated into new environments that have never before been seen, by selecting behaviour appropriate to similar environments.

4.3.1 Future Work

Two types of synergy were discussed: intrinsically modulated and extrinsically modulated. The intrinsically modulated synergies were designed to exploit different types of redundancy, but only one type was investigated. For example, a system may be faced with multiple environments that change in sequence rather than appearing at random intervals. Or alternatively, a certain environment may always appear followed by another, but otherwise may be random. It seems reasonable that the dynamics of the synergy and its assembly selection should reflect the type of redundancy in the environment; if there is no pattern and the environments the system is exposed to

4.3. DISCUSSION

are entirely random then there would clearly be no benefit to using a synergy *in this way*.

There is another simpler advantage of using synergies that was not explored; in an articulated system it may be possible to combine the synergies in a hierarchy. This may be particularly effective when using extrinsically modulated synergies. It seems feasible that if an articulated system with many degrees of freedom had several motor ‘primitives’, such as grasping, these could be determined using networks of synergies and a search could be carried out of combined primitives. Different levels of abstraction could be modulated by different signals in order to make adaptation efficient. This would be far more effective than using a single extrinsically modulated synergy. In a social situation this may also be particularly effective, as a teacher could provide instability to assemblies correlated to incorrect behaviours, and increase the selection probability of assemblies correlated to the appropriate behaviour. Here an assembly would perhaps parameterise several synergies, on one side of an articulated system, independent of another hierarchy. As social animals tend to be more complex, showing greater variety of patterns despite high degrees of freedom. It may be interesting to distinguish between altricial and precocial species with this in mind.

What has also not yet been fully taken into account is the transient periods between changes of parameter value, i.e., the dynamics of the representative point in the parameter-space. It may be the case that these transient periods could be used to coordinate what they parameterise in a useful way. For example, the movement of the arm, from one position to another during a grasping motion. The movements are not of step-form, but are partial, consisting of step-function elements interspersed with continuous motion. This could also be combined with low optimality signals in order to perform a kind of ‘local-search’ or optimisation of the parameter-space.

5

Conclusion

Throughout the three main chapters of the thesis a variety of issues were raised based on potential problems or limitations in the framework. In order to address these issues, various ‘elaborations’ or potential extensions were subsequently suggested. Among the first of these changes was the substitution of essential variables for two new distinct types of variable, so that viability could be separated from lability. This could be considered to be the most important change made, as it made possible many of the other subsequent developments. Furthermore, it was argued that this initial change was also necessary, since previously the framework appeared to forbid essential variables without physiological limits, such as ‘level of pain’.

Two limitations of the framework explicitly acknowledged by Ashby in (17/10) were addressed, one directly: the *varying trial length*; and one indirectly: the *accumulation of adaptations*. The former was addressed by modifying the ultrastable system so a timing mechanism was introduced *into its dynamics*. The modification was made based on distinguishing more specific types of disturbance to the ultrastable system, not previously taken into account. It was demonstrated that modelling the disturbance to the essential (lability) variables could prove useful. The timing mechanism introduced relied on changing the movement of the critical states in the reacting-part, which was also identified with a particular kind of disturbance. New concepts such as the optimality-landscape were introduced to make discussion easier, based on the disturbance to the essential variables.

The third main chapter addressed the issues of efficient adaptation and the role of the random step-changes in the ultrastable system. The concept of synergies was introduced in order to both reduce the parameter-space and improve selection of points within the parameter-space beyond random choice, exploiting the ‘recurrent situation’ of the environment. A trade-off

was identified, exploration vs. exploitation, emerging from the constraint that only one behaviour can be tried at any one time, and each behaviour must persist for a sufficient duration before feedback can be provided by the environment. It was suggested that large adaptation times could be overcome in systems with high degrees of freedom, particularly articulated systems, if synergies are combined in a hierarchy. It was further suggested that this process may be accelerated by other intelligent agents in the environment.

Many other more subtle issues were also addressed, but if the suggested changes made are to be non-arbitrary there must be some additional practical use or predictive power. Additional explanatory power may not be sufficient. In order to verify the elaborations made, further work therefore must be carried out.

References

- Ashby, R. (1960), *Design for a Brain*, Chapman & Hall, London.
- Brooks, R. A., Breazeal (Ferrell), C., Irie, R., Kemp, C. C., Marjanović, M., Scassellati, B. & Williamson, M. M. (1998), Alternative essences of intelligence, in ‘AAAI ’98/IAAI ’98: Proceedings of the fifteenth national/tenth conference on Artificial intelligence/Innovative applications of artificial intelligence’, American Association for Artificial Intelligence, Menlo Park, CA, USA, pp. 961–968.
- Maturana, H. R. & Varela, F. J. (1991), *Autopoiesis and Cognition : The Realization of the Living (Boston Studies in the Philosophy of Science)*, Springer.
- Maturana, H. R. & Varela, F. J. (1998), *Tree of Knowledge*, New Science Library: Distributed in the United State by Random House, Boston.
- Phattanasri, P., Chiel, H. J. & Beer, R. D. (2007), ‘The dynamics of associative learning in evolved model circuits’, *Adaptive Behavior* **15**(4), 377–396.
- Schöner, G. (2007), Development as change of system dynamics: Stability, instability, and emergence., in J. P. Spencer, M. Thomas & J. McClelland, eds, ‘Toward a Grand New Theory of Development? Connectionism and Dynamic Systems Theory Reconsidered’, Oxford University Press.
- Smith, L. B. & Thelen, E. (2003), ‘Development as a dynamic system’, *Trends in Cognitive Sciences* **7**(8), 343–348.

Appendices

A

Source Code

A.1 PURU

Source Code: puruExperiment.m

```
% PURU demonstration
function puruExperiment (speeds, landscapeType, couplingStrength, repeats, ←
    showInRealTime)
% Initialise variables
total = length(speeds);
avgSurvivals = zeros(1,total);
stdSurvivals = zeros(1,total);
avgOptimalities = zeros(1,total);
stdOptimalities = zeros(1,total);
avgAdaptationTimes = zeros(1,total);
stdAdaptationTimes = zeros(1,total);

% For each speed
for i = [1:total]
    survivalTimes = zeros(1,repeats);
    averageOptimalities = zeros(1,repeats);
    averageAdaptationTimes = zeros(1,repeats);

    % Run multiple repeats of experiment
    for repeat = 1:repeats
        [timeOfDeath, averageOptimality, averageAdaptationTime] = puru(speeds(i), ←
            landscapeType, couplingStrength, repeats, showInRealTime);
        survivalTimes(repeat) = timeOfDeath;
        averageOptimalities(repeat) = averageOptimality;
        averageAdaptationTimes(repeat) = averageAdaptationTime;
    end

    % Draw histograms (disabled)
    if false && repeats > 1
        figure;
        hist(survivalTimes);
        figure;
        hist(averageOptimalities);
    end

    % Calculate mean and standard deviations
    avgSurvivals(i) = mean(survivalTimes);
    stdSurvivals(i) = std(survivalTimes);
    avgOptimalities(i) = mean(averageOptimalities);
    stdOptimalities(i) = std(averageOptimalities);
    avgAdaptationTimes(i) = mean(averageAdaptationTimes);
    stdAdaptationTimes(i) = std(averageAdaptationTimes);
end

% Plot graphs
if total > 1
    figure;
    subplot(1,3,1);
```

```

    lims = [min(speeds) max(speeds)];
    errorbar(speeds, avgSurvivals, stdSurvivals);
    xlim(lims);
    xlabel('Agent speed');
    ylabel('Average survival time');
    title('Survival Time');
    subplot(1,3,2);
    errorbar(speeds, avgOptimalities, stdOptimalities);
    xlim(lims);
    xlabel('Agent speed');
    ylabel('Average optimality');
    title('Optimality');
    subplot(1,3,3);
    errorbar(speeds, avgAdaptationTimes, stdAdaptationTimes);
    xlim(lims);
    xlabel('Agent speed');
    ylabel('Average adaption time');
    title('Environment Adaptation Time');
end
end

% Main PURU experiment
function [ timeOfDeath, averageOptimality, averageAdaptationTime ] = puru(speed, ←
    landscapeType, couplingStrength, repeats, showInRealTime )

% Initialise simulation variables
totalTime = 500;
timeStepSize = 0.1;
totalTimeSteps = totalTime / timeStepSize;
relapseRate = 2;

%
% Environment
%
% safeRegionCentre = centre of safe region ('the shade')
% safeRegionSize = the size of the save region
% v = viability variable (e.g., temperature)
environment.safeRegionSize = 0.5;
environment.safeRegionTop = max(getRandomPosition(), environment.safeRegionSize - ←
    1);
environment.v = 0;

%
% Lability Unit
%
% l = the lability of the unit
% s = the parameter to the step-unit
labilityUnit.l = 0;

%
% Step Unit
%
% p = the parameter to environment/reacting part
% (will contain other variables not modelled)
stepUnit.p = getRandomPosition()

%
% Reaction Unit
%
% p = the parameter to environment/reacting part
% (will contain other variables not modelled)
reactionUnit.position = getRandomPosition()
reactionUnit.speed = speed;

%
%

% Initialise variables for storing
recReaction = zeros(1, totalTimeSteps);
recStepvalue = zeros(1, totalTimeSteps);
recSafeRegion = zeros(1, totalTimeSteps);
recViability = zeros(1, totalTimeSteps);
recOptimality = zeros(1, totalTimeSteps);
recLability = zeros(1, totalTimeSteps);
recRelapse = zeros(1, totalTimeSteps);

% Initialise further simulation variables
relapse = 0;
lastOptimalitySignal = 0;
optimalitySignal = 0;
rate = 0;
totalTimeStepsInShade = 0;
adaptationTimes = zeros(1, totalTimeSteps/500);

```

```

environmentIndex = 1;
lastOutsideShade = 0;

% Main simulation
for t = 1:totalTimeSteps

    % Type (2A) disturbance in environment
    if mod(t,500) == 0
        environment.safeRegionTop = max(getRandomPosition(), environment.↵
            safeRegionSize - 1);

        adaptationTimes(environmentIndex) = lastOutsideShade - (t-500);
        environmentIndex = environmentIndex + 1;

        lastOutsideShade = t;
    end

    % If still viable
    if environment.v < 1

        % Update reaction unit position
        reactionUnit.position = reactionUnit.position + timeStepSize * ↵
            reactionUnit.speed*(stepUnit.p - reactionUnit.position);

        % Type (1) disturbance in environment (add noise)
        reactionUnit.position = reactionUnit.position + getNoise();

        % Make sure position is within minimum and maximum bounds
        reactionUnit.position = min(1,max(reactionUnit.position,-1));

        % Get optimality
        lastOptimalitySignal = optimalitySignal;
        optimalitySignal = getOptimality(environment, reactionUnit, landscapeType)↵
            ;

        % If agent is outside safe region
        if ~(reactionUnit.position <= environment.safeRegionTop && reactionUnit.↵
            position >= environment.safeRegionTop - environment.safeRegionSize)

            lastOutsideShade = t;

            % Increase temperature
            environment.v = environment.v + timeStepSize * (0.01);
        % Otherwise
        else
            totalTimeStepsInShade = totalTimeStepsInShade + 1;
            % Decrease temperature
            environment.v = environment.v + timeStepSize * (-environment.v * 0.1);
        end

        % Update lability
        %couplingStrength = 0.1;

        rate = 0;
        if t > 20
            rate = recOptimality(t-1) - recOptimality(t-20);
        end

        useRate = 1;

        if landscapeType == 1
            labilityUnit.l = labilityUnit.l + timeStepSize * ((optimalitySignal - ↵
                labilityUnit.l)* couplingStrength - 10*relapse);
        else
            if useRate
                labilityUnit.l = labilityUnit.l + timeStepSize * (optimalitySignal↵
                    * (rate + couplingStrength) - 10*relapse);
            else
                labilityUnit.l = labilityUnit.l + timeStepSize * (optimalitySignal↵
                    * couplingStrength - 10*relapse);
            end
        end

        labilityUnit.l = min(1,max(labilityUnit.l,0));

        % Check for relapse
        if labilityUnit.l >= 0.95
            relapse = 1;
            %labilityUnit.l = 0;
            stepUnit.p = getRandomPosition();
        end
        relapse = relapse + timeStepSize * (-relapseRate); %-0.5
        relapse = max(0, relapse);
    end
end

```

```

        % Add noise to viability variable
        environment.v = environment.v + getNoise();
        environment.v = min(1,environment.v);

    end

    % Record variables
    recReaction(t) = reactionUnit.position;
    recStepvalue(t) = stepUnit.p;
    recSafeRegion(t) = environment.safeRegionTop;
    recViability(t) = environment.v;
    recOptimality(t) = optimalitySignal;
    recLability(t) = labilityUnit.l;
    recRelapse(t) = relapse;

    % If plotting activity at each time-step
    if showInRealTime
        plot(reactionUnit.position,'ro');
        ylim([-1 1]);
        title(num2str(t));
        pause(0.01);
    end
end

timeOfDeath = min(find(recViability == 1));

if isempty(timeOfDeath)
    timeOfDeath = totalTime / timeStepSize;
end

if repeats == 1
    close all;
    figure;

    subplot(3,1,1);
    plot(recViability);
    ylim([-1 1.1]);
    drawLineOfDeath(timeOfDeath);
    title('Viability');

    subplot(3,1,2);
    plot(recStepvalue, 'g');
    hold on;
    plot(recReaction);
    plot(recSafeRegion, 'r');
    plot(recSafeRegion - environment.safeRegionSize, 'r');

    title('Environment and Reaction');
    ylim([-1 1]);
    drawLineOfDeath(timeOfDeath);

    subplot(3,1,3);
    plot(recRelapse, 'g');
    hold on;
    plot(recOptimality, 'r');
    plot(recLability, 'k');

    title('Optimality and Lability');
    ylim([-0.1 1.1]);
    drawLineOfDeath(timeOfDeath);
end

averageAdaptationTime = mean(adaptationTimes);
averageOptimality = mean(recOptimality);
end

function drawLineOfDeath(timeOfDeath)
    if timeOfDeath
        hold on;
        line([timeOfDeath, timeOfDeath], [-1 1], 'Color', 'r', 'LineWidth', 3);
    end
end

function noise = getNoise()
    noise = 0.005 * (rand()*2 - 1);
end

function optimality = getOptimality(environment, reactionUnit, type)
    % If using binary optimality
    if type == 1
        % If inside the safe zone

```

A.2. SYNERGIES

```
    if reactionUnit.position <= environment.safeRegionTop && reactionUnit.position <=
        >= environment.safeRegionTop - environment.safeRegionSize
        optimality = 0;
    % If outside the safe zone
    else
        optimality = 1;
    end
elseif type == 2

    optimality = abs(reactionUnit.position - (environment.safeRegionTop -
        environment.safeRegionSize/2));

    %if optimality > environment.safeRegionSize/2
    % optimality = 1;
    %end

    % optimality = abs(
else
    distance = abs(reactionUnit.position - (environment.safeRegionTop -
        environment.safeRegionSize/2));

    if distance < environment.safeRegionSize/2
        optimality = 0;
    else
        optimality = distance;
    end
end
end
end

% Returns a random number from a uniform distribution between -1 and 1
function position = getRandomPosition()
    position = rand()*2 - 1;
end
```

A.2 Synergies

Source Code: syn1.m

```
% Synergy demonstration
function syn1
% Set parameters
params.timeStepSize = 0.1;
params.totalTime = 40000;
params.trialLength = 50;
params.single = false;
params.perturbation = 0;
params.noise = true;
params.useBiases = true;
params.totalSynergies = 5;
params.environments = 2;
params.environmentLoops = 4;
params.optimalZoneSize = 0.6;
params.beta = 0.5;

if true
    totalNodes = 2;
    params.showInRealTime = false;
    params.single = true;
    values = rand(1, totalNodes)*2 - 1;
    weights = rand(totalNodes, totalNodes)*2 - 1;
    biases = rand(1, totalNodes)*2-1;
    runExperiment(totalNodes, values, weights, biases, params);
end
return;

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

if true
    params.single = false;
    results = process(2, params, 10);
    errorbar(1:params.environments*params.environmentLoops, mean(results.adaptationTimes), std(results.adaptationTimes));
```

A.2. SYNERGIES

```

end
return;

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

minNodes = 2;
maxNodes = 2;
totalNodesToTest = maxNodes - minNodes + 1;
nodeAverageResults = zeros(1, totalNodesToTest);
nodeStdResults = zeros(1, totalNodesToTest);
nodeAdaptedResults = zeros(1, totalNodesToTest);

mh = waitbar(0, 'Conducting experiment... ');

nodeIndex = 1;

for i = 1:maxNodes
    results = process(minNodes+i-1, params, 5);

    nodeResults(i) = mean(results.averageTimes);
    nodeStdResults(i) = std(results.averageTimes);
    nodeAdaptedResults(i) = results.percentageAdapted;

    waitbar(i/maxNodes, mh);
end

dir = ['results/', num2str(params.type), '_p', num2str(params.perturbation)];
mkdir(dir);

fh = figure;
errorbar(1:maxNodes, nodeResults, nodeStdResults);
xlabel('Total nodes');
ylabel('Mean number of time-steps to adaptation');
ylim([1 round(params.totalTime/params.timeStepSize)]);
saveAndCloseFigure(fh, dir, 'adaptationTime');

fh = figure;
plot(nodeAdaptedResults, '-x');
xlabel('Total nodes');
ylabel('Percentage adapted');
ylim([0 100]);
saveAndCloseFigure(fh, dir, 'adaptationPercentage');

end

% Run multiple experiments
function [results] = process(totalNodes, params, repeats)
    adaptationTimes = zeros(repeats, params.environments*params.environmentLoops);
    for i = 1:repeats
        values = rand(1, totalNodes)*2 - 1;
        weights = rand(totalNodes, totalNodes)*2 - 1;
        biases = rand(1, totalNodes)*2 - 1;
        thisRepeatAdaptationTimes = runExperiment(totalNodes, values, weights, biases, ←
            params)
        adaptationTimes(i,:) = thisRepeatAdaptationTimes;
    end
    results.adaptationTimes = adaptationTimes;
end

% Run experiment
function [adaptationTimes] = runExperiment(totalNodes, values, weights, biases, ←
    params)
    timeStepSize = params.timeStepSize;
    totalTime = params.totalTime;
    totalTimeSteps = round(totalTime/timeStepSize);
    environmentDuration = round(params.totalTime / (params.environments * params.←
        environmentLoops) / timeStepSize);
    adaptationTimes = ones(1, params.environments*params.environmentLoops) * ←
        environmentDuration;

    if params.single
        stable = zeros(1, totalTimeSteps);
        valuesLog = zeros(totalTimeSteps, totalNodes);
        temp = zeros(totalTimeSteps, totalNodes);
        syn = zeros(1, totalTimeSteps);
        something = zeros(1, totalTimeSteps);
    end

    trialLength = params.trialLength;
    lastCritical = 0;
    halfOfTime = round(totalTimeSteps/2);

```

A.2. SYNERGIES

```
assemblyWeights = cell(1,params.totalSynergies);
assemblyBiases = cell(1,params.totalSynergies);
for i = 1:params.totalSynergies
    assemblyWeights{i} = rand(totalNodes ,totalNodes)*2 - 1;
    assemblyBiases{i} = rand(1,totalNodes)*2 - 1;
end
assemblyStrengths = zeros(1,params.totalSynergies);

currentEnvironment = 1;
currentAssembly = 1;
counter = 1;
switchIndex = 1;
wasOptimal = false;
optimal = false;
weights = assemblyWeights{currentAssembly};
biases = assemblyBiases{currentAssembly};
optimalZones = zeros(1,params.environments);
for i = 1:params.environments
    optimalZones(i) = max(rand()*2 - 1, params.optimalZoneSize - 1);
end

%optimalZones = [0.8, -0.2];

% For each time-step
for t = 1:totalTimeSteps

    % Inherently stable
    if 0
        for i = 1:totalNodes
            weights(i,i) = -1;
        end
    end

    % Update node values
    values = values + timeStepSize * ((values * weights) + params.useBiases*biases↵
    );

    % PERTURBATION 1
    if params.perturbation == 1
        if mod(t,500) == 0
            values(1) = values(1) + 0.2;
        end
    % PERTURBATION 2
    elseif params.perturbation == 2
        if t == halfOfTime
            weights(1,2) = -weights(1,2);
        end
    % PERTURBATION 3
    elseif params.perturbation == 3
        if mod(t,100) == 0
            if mod(t,20) == 0
                weights(1,2) = 0.5;
            else
                weights(1,2) = 1;
            end
            weights = max(min(weights,1),-1);
        end
    end

    % Add a tiny amount of noise
    if params.noise
        values = values + 0.005* (rand(1,totalNodes)*2 - 1);
    end

    % Make sure values stay within bounds
    values = max(min(values,1),-1);

    % Determine optimal zone boundaries
    upperBound = optimalZones(currentEnvironment);
    lowerBound = upperBound - params.optimalZoneSize;

    % Find nodes out of optimal zone
    critical = find(values < lowerBound | values > upperBound);

    % Calculate adaptation times
    wasOptimal = optimal;
    optimal = isempty(critical);
    if (~wasOptimal || counter == 1) && optimal
        adaptationTimes(switchIndex) = counter - 1;
    end

    % If there are nodes out of the optimal zone
    if mod(t,trialLength) == 0 && ~optimal
```

A.2. SYNERGIES

```

lastCritical = t;

% Select adapted assembly or non-adapted with probab
selectAdapted = rand() <= params.beta;

if isempty(find(assemblyStrengths == 1))
    selectAdapted = 0;
elseif isempty(find(assemblyStrengths < 1))
    selectAdapted = 1;
end

% Select of those at random
if selectAdapted
    chooseFrom = find(assemblyStrengths == 1);
else
    chooseFrom = find(assemblyStrengths < 1);
end

% Select an assembly
currentAssembly = chooseFrom(round(rand()*(length(chooseFrom)-1)+1));

% If selected an adapted assembly then do nothing, else
% randomise (move assembly to random location)
if assemblyStrengths(currentAssembly) < 1
    assemblyWeights{currentAssembly} = rand(totalNodes, totalNodes)*2 - 1;
    assemblyBiases{currentAssembly} = rand(1,totalNodes)*2-1;
end

% Parameterise the system
weights = assemblyWeights{currentAssembly};
biases = assemblyBiases{currentAssembly};
end

% Increment assembly strengths if in the optimal zone
if optimal
    assemblyStrengths(currentAssembly) = min(assemblyStrengths(currentAssembly)↔
        ) + 0.001, 1);
end

% Determine current environment
if counter == environmentDuration
    if currentEnvironment == params.environments
        currentEnvironment = 1;
    else
        currentEnvironment = currentEnvironment + 1;
    end
    counter = 1;
    switchIndex = switchIndex + 1;
else
    counter = counter + 1;
end

% Record variables
if params.single
    temp(t,:) = biases;
    syn(t) = currentAssembly;
    something(t) = weights(1,1);
    valuesLog(t,:) = values;
    stable(t) = length(critical);

    % Plot real-time graph
    if params.showInRealTime %&& mod(t,20) == 0 %> 0.99*totalTimeSteps
        plot(values, 'or')
        ylim([-1 1]);
        title(num2str(t));
        pause(0.01);
    end
end

% Plot graphs
if params.single
    fh = figure;

    subplot(2,2,1);
    plot(adaptationTimes);
    title('Adaptation time');

    subplot(2,2,2);
    plot(valuesLog);
    ylim([-1 1]);
    title('Node values (behaviour)');
end

```


A.2. SYNERGIES

```

%         line(
subplot(2,2,3);
plot(syn, 'kx');
title('Assembly selection');

subplot(2,2,4);
plot(temp(:,1), 'rx');
hold on;
plot(temp(:,2), 'bx');
plot(something, 'yx');
ylim([-1 1]);
title('Bias values, self-referent weight');

figure;
plot(temp(:,1), temp(:,2), 'kx');

hold on;
for i = 1:params.totalSynergies
    %[X,Y] = pol2cart(linspace(0,2*pi,50), ones(1,50)*assemblyWidths(i));
    %plot(X+assemblyBiases{i}(1),Y+assemblyBiases{i}(2), 'r');
    plot(assemblyBiases{i}(1), assemblyBiases{i}(2), 'r*', 'MarkerSize', 5+5*←
        assemblyStrengths(i));
end
axis square;
title('Bias values search');

if 0 && totalNodes == 3
    figure;
    first = 1:lastCritical;
    second = lastCritical+1:totalTimeSteps;
    plot3(valuesLog(first,1), valuesLog(first,2), valuesLog(first,3), '-','Color'←
        ,[0.8,0.8,0.8]);

    if lastCritical < totalTimeSteps
        hold on;
        line([valuesLog(lastCritical,1), valuesLog(lastCritical+1,1)], [←
            valuesLog(lastCritical,2), valuesLog(lastCritical+1,2)], [←
            valuesLog(lastCritical,3), valuesLog(lastCritical+1,3)], 'Color', 'r');
        plot3(valuesLog(second,1), valuesLog(second,2), valuesLog(second,3), '-','←
            'Color', 'k');
    end
    axis([-1 1 -1 1 -1 1]);
end

for x = 1:10
    stuff = zeros(1,params.totalSynergies);
    stuff(1) = assemblyStrengths(1);
    for i = 2:params.totalSynergies
        stuff(i) = assemblyStrengths(i) + stuff(i-1);
    end

    currentAssembly = find(stuff >= rand()*sum(assemblyStrengths));
    currentAssembly = currentAssembly(1)
end
end
end

```