DIVISION OF COMPUTER SCIENCE

Developmental Neural Networks for Shape Recognition: Motivation and Review

Alistair G Rust
Hamid Bolouri

Technical Report No.241

March 1996

Developmental Neural Networks for Shape Recognition: Motivation and Review

Alistair G Rust
Division of Computer Science
University of Hertfordshire, AL10 9AB, UK

Hamid Bolouri
Engineering Research and Development Centre
University of Hertfordshire, AL10 9AB, UK

Tel : +44 1707 284661 Fax : +44 1707 284185
E-mail : a.g.rust, h.bolouri@herts.ac.uk WWW : http://sparky.herts.ac.uk

Computer Science Technical Report No.241
ERDC Technical Report ERDC/1996/0006

March 7, 1996

Abstract

Artificial Neural Networks (ANNs) are widely exploited in Artificial Intelligence applications. However, at present, there is a wide gap in functionality between artificial and biological neural systems. Appropriate neuron functionality, network topology and multi-network system organisation in ANNs are the subjects of much debate.

This report presents arguments for modelling the developmental processes in biological neural systems as a stepping stone towards better understanding of biological neural architectures and self organisation.

The authors are interested in shape recognition systems. The early vision system in mammals is proposed as an appropriate test bed for developmental modelling. The report presents an overview of biological and ANN shape processing systems. A review of developmental processes in the mammalian visual system is followed by an analysis of current research in developmental neural networks.

Acknowledgements

We would like to thank Stella George and Rod Adams for reviewing and commenting on this document at various stages of its evolution. Thanks also to Alan Longstaff for checking the summary of biological retinas.

Table of Contents

Abstract **i**

Acknowledgements **ii**

Table of Contents **iii**

1 Introduction **1**
- 1.1 Background 1
- 1.2 The components of shape recognition 1
- 1.3 Biological and artificial neural systems 2
- 1.4 Self organisation in neural systems 3
- 1.5 The need for developmental neural networks 3
- 1.6 Visual shape recognition as a test bed for modelling developmental self organisation 5
- 1.7 Report outline 6

2 Current Shape Recognition Systems **7**
- 2.1 Context, texture and shading 7
- 2.2 Foveation and selective attention 7
- 2.3 Scaling, translation, and rotation invariance 8
  - 2.3.1 Invariant feature extraction 8
  - 2.3.2 Example invariant feature extraction implementations 8
- 2.4 Hierarchical feature extraction 9
  - 2.4.1 Motivation 9
  - 2.4.2 Compositional hierarchy 9
  - 2.4.3 Hierarchical feature extraction in biological neural networks 10
  - 2.4.4 Fukushima's Neocognitron 10
  - 2.4.5 Mueller, Blackman and Furman 11
  - 2.4.6 Rolls 11
  - 2.4.7 Hummel and Biederman 12
- 2.5 Desirable features of hierarchical visual feature extraction systems 13

3 Biological Vision **15**
- 3.1 Functional organisation 15
  - 3.1.1 Form recognition 16
- 3.2 Anatomical organisation and physiology 17
  - 3.2.1 Retina 19
  - 3.2.2 Lateral Geniculate Nucleus 23
  - 3.2.3 V1 23
  - 3.2.4 V2 24
  - 3.2.5 V4 25
  - 3.2.6 Inferotemporal 25
- 3.3 Summary 25

4 Biological Neural Development **26**
- 4.1 Gene expression 26
  - 4.1.1 Neural formation 27
  - 4.1.2 Connectivity development 30
- 4.2 Learning 33
- 4.3 Summary 33

5 Related Work 34
5.1 Comparison of constructive and developmental approaches 34
5.2 Developmental Models 35
5.2.1 Gruau 35
5.2.2 Boers and Sprinkhuizen-Kuyper 36
5.2.3 deGaris 37
5.2.4 Nolfi and Parisi 38
5.2.5 Jakobi 39
5.2.6 Dellaert and Beer 40
5.2.7 Vaario 41
5.2.8 Fleischer 41
5.3 Discussion 42
5.4 Summary 43

References 45

1. Introduction

1.1 Background

In its broadest sense, visual shape recognition involves a spectrum of activities such as head-eye coordination, saccadic eye movement, and selective attention. A host of image processing functions preserve image integrity across variations in the intensity, angle and power spectrum of incident light. Furthermore, the boundaries between shape classes are usually not well defined, so that whether a view of an object is recognised or categorised correctly often depends on additional cues such as context and expectation. These are concepts which rely on past experience in both visual and non-visual terms. Thus, in humans, visual shape recognition involves not only most of the visual cortex, but also significant portions of other parts of the brain normally associated with such functions as motor control and language understanding.

The organisation of biological vision systems is the product of millions of years of co-evolution of system 'components'. There is therefore a strong argument against analysing components of biological vision systems in isolation. However, because of the scale and complexity of the processes involved, research in computer vision has tended to view each of the component processes individually.

1.2 The components of shape recognition

In its simplest form, shape recognition may be viewed as processing single shapes which are scaled, centred, and oriented prior to recognition. Such an approach may not be a gross over-simplification of biological vision. Current research suggests that a number of interdependent processes, operating at various levels of abstraction within the visual system, deliver an invariant 'image' to a shape recognition subsystem. For example, attention (foveation) is viewed as control of eye muscles based on coarse-grain, context-dependent 'blob' localisation [Hummel and Biederman, 1992]. Selective attention is thought to involve feedback after initial feedforward processing identifies a number of candidate objects in the field of view (see for example [Fukushima, 1987, Goebel, 1993]). Translation and scaling invariance appear to be integrated into biological vision through routing networks which deliver a scaled and centred view of an object to the recognition network [Olshausen et al., 1993, Alexandre and Guyot, 1995]. Finally, rotation invariance in biological vision may be in part due to the log-polar arrangement of the receptive/projective fields in the retina [Fukumi et al., 1991, Minnix et al., 1991].

The authors are interested in invariant shape recognition in its broad sense. However, for the sake of simplicity, the rest of this document concentrates on research towards systems capable of recognising shapes which are approximately scaled, centred, and oriented.

A wide variety of statistical, syntactic, and mathematical techniques for global and local shape recognition have been developed. In principle, these approaches should deliver approximately equivalent recognition rates given enough computing power. In practice, artificial neural networks have found particular favour in this area because they allow high speed recognition / classification and also because they provide a useful framework for self-organisation (see section 1.4).

1.3 Biological and artificial neural systems

Theoretical understanding of artificial and biological neural networks is still in its infancy. For instance, neuron functionality is still the subject of debate. The popular weighted sum of inputs neuron model has little in common with its biological counterpart [Crick, 1989]. It is clear that biological neurons behave in far more complex ways. In particular, it has been argued that structure (in addition to interconnect topology) determines the function of biological neurons [Fleischer and Barr, 1994]. The location of an input to a neuron determines how that input interacts with other inputs and how it influences the neuron output. Differing geometrical arrangements of inputs to a neuron result in different neuron transfer functions.

At a higher level of abstraction, networks of neurons can be organised in a variety of topologies. While it is clear that neural receptive and projective fields determine neural function [Lehky and Sejnowski, 1988], there is as yet no structured methodology for designing these fields. The popular distinction between feedforward and feedback topologies is often blurred in biological networks where layered feedforward structures frequently include lateral connections within layers, feedforward and feedback projections to other layers, and systematic connections to parts of other functional modules.

At the system organisation level, general purpose methods of determining artificial neural network (ANN) topology and multi-ANN system organisation from training data (e.g. [Webber, 1994, Zemel et al., 1990]) are yet to be fully developed. Biological neural systems are known to consist of many interacting, but distinguishable, functional modules. The modular organisation of biological neural systems is thought to embody knowledge about the type of information the neural system is expected to process. For instance, the aural neural networks in humans have specific modular structures which provide a predisposition to learning a hierarchy of speech recognition skills ranging from phoneme association to grammar [Pinker, 1994].

The modular organisation of biological neural systems may also serve another purpose. In simple fully-connected, layered, feed forward ANNs, the number of neurons (and synapses), must be tuned to the size and complexity of the training data. Too many weights (synapses) result in overfitting the training data while too few weights leave the network unable to learn the training data with the required accuracy. In animals, the number of neurons decreases with age [Purves and Lichtman, 1985]. However, this does not seem to adversely affect early learning (little data, many synapses). The modular organisation of biological networks may help avoid overfitting problems by constraining the system behaviour space.

Neuron structure, network topology, and multi-network system organisation may be viewed as representing neural architectures at different scales (from neuron to system). Inherited network architectures may then be viewed as encoding knowledge about the fundamental characteristics of the input space. Such structure can improve the learning speed of neural networks and also prevent networks from overfitting data in arbitrary ways. This view also fits in with the growing support for 'embodiment', namely that intelligent systems must be tightly coupled to the environments in which they are expected to perform [Brooks and Stein, 1993].

1.4 Self organisation in neural systems

As mentioned earlier, one of the advantages of neural networks is that they can self organise in response to training data, thus simplifying the development of new applications. It is possible to implement neural self organisation with varying degrees of freedom. At the simplest level, automatic adjustment of a neuron's firing threshold or gain parameter is a self organisation process. At the network level, training may be viewed as developing an input-output mapping without explicitly defining the required network transfer function. Self organisation may also be implemented beyond the level of synaptic plasticity, at the level of network topology formation, and beyond that at the level of multi-network system organisation.

The simpler self organisation processes (such as neuron function and weight adaptation) allow automatic searching of limited regions in the space of all potential system behaviours. The more complex processes of network topology and architecture formation may be viewed as providing mechanisms for global sampling of the potential system behaviour space.

Since inherited network architecture in biological systems appears to encode high-level (abstracted) knowledge about the characteristics of stimulus patterns, architectural self organisation processes may be viewed as 'learning' those underlying features of the input patterns which remain constant over many generations.

1.5 The need for developmental neural networks

Network topology self organisation offers an opportunity to overcome the limitations of ad hoc approaches to neural systems design. Current research in neural topology optimisation may be divided into three broad classes.

- Search or optimisation techniques. A variety of search algorithms may be used to optimise network topologies for specific applications. In general, these approaches minimise an error measure comparing the behaviour of candidate network topologies with a perceived 'ideal'. Searching for direct encodings of the network topology (such as a connection matrix) is thought to be ineffective because the search space tends to be highly discontinuous : a small change in the network encoding can result in a very large change in system functionality [Yao, 1992].

- Constructive or generative techniques overcome the limitations of direct search by building networks incrementally through repeated evaluation of the response of the system to the training set (see for example [Chentouf and Jutten, 1995, Carpenter and Grossberg, 1988, Fahlman and Lebiere, 1990]. [Jutten, 1995] provides a review of this approach to topology development. Constructive techniques have the advantage of adaptively increasing network complexity until it meets predefined requirements. Most constructive systems require only incremental learning after each construction step. This avoids time consuming retraining and enables high-speed network development. However, the search space of constructive techniques is limited to the specific network architectures favoured by the construction rules. The architectural limitations of constructive networks, can be overcome by dynamically searching the space of network construction rules for a given application. To reduce the complexity of the search space, practical implementations tend to search the parameter space of particular classes of network construction rules, so that the 'architecture space' of constructive systems is always constrained. Genetic evolution has been the favoured tool for searching the construction rule parameter space. [Yao, 1992, Balakrishnan and Honavar, 1995, Jones, 1993] provide reviews of the topic.

- Developmental techniques. [Yao, 1992] argues that network architecture formation, learning rule optimisation, and synaptic weight adaptation are three distinct synergic processes operating in different time scales. In Yao's view, architectures adapt at the slowest rate and dictate the evolution of learning rules. Learning rules in turn determine the value of synaptic weights. Yao's hypothesis is supported by the view that network architecture (structure / topology / modular organisation) represents domain knowledge and reflects long term adaptation to high-level (abstract) patterns in the application environment [Pinker, 1994].

  Developmental neural systems provide a mechanism for separate architecture and learning rule optimisation by mimicking the embryonic development of biological neural systems. Various levels of abstraction have been investigated. For instance, [Fleischer and Barr, 1994] have argued for a detailed emulation of structural development in biological neural systems (as a precursor to understanding the principles of developmental self organisation. In contrast, [Dellaert and Beer, 1996] have reported that their earlier attempts [Dellaert and Beer, 1994b, Dellaert and Beer, 1994a] at developing structurally detailed neural systems proved too computationally expensive. They now model development at a more abstract topological level.

  Contrasting views on the necessary degree of abstraction and similarity to biological systems has led to investigations of a variety of developmental rules including Lindenmayer (L-system) grammars [McCormick and Mulchandani, 1994], Cellular Automata rules [Whitley et al., 1995], and feedback regulatory networks [Dellaert and Beer, 1996].

The work of the authors [Rust et al., 1996] is aimed at mimicking the development of neuron structure and topology in the mammalian visual shape recognition system.

1.6 Visual shape recognition as a test bed for modelling developmental self organisation

Apart from [Linsker, 1988], most developmental neural network research to date has been evolved, evaluated, and demonstrated on animats. Animats are simulated automous mobile robots operating in simple simulated environments. Reviews of the field are presented in [Mataric and Cliff, 1995] and in chapter 5. Animats are attractive for this type of research because it is possible to develop interesting behaviours with very few neurons [Braitenberg, 1984]. However, Mataric and Cliff [Mataric and Cliff, 1995] describe a number of complications arising from the use of animats as research test beds, as described below.

- Because they have to recognise and execute sequences of events, animat controllers require synchronisation (clocking), neural activity phase-locking, and feedback mechanisms which complicate their design.
- Animats are evaluated on the basis of their interactions with their environment. The non-deterministic nature of these interactions results in difficulties in animat learning, behaviour evaluation, system analysis, and experimental repeatability.
- Genetic evolution of developmental rules for animats cannot be efficiently mapped onto parallel processor systems because of the temporal nature of animat behaviour. This limits the complexity of neural systems which may be evolved in animats.

Visual shape recognition neural networks avoid these difficulties, as follows.

- The recognition of single, approximately centred, scaled, and oriented shapes does not appear to require feedback.
- So long as only one object is present in the field of view, there appears to be no need for sequential, synchronous, or phase-locked neural activity.
- The early vision system does not require learning. Therefore developmental schemes can be quickly evaluated after network growth. There is no need for time consuming and non-deterministic learning.
- Visual shape recognition has a direct biological counterpart. The lack of learning, feedback, and synchronisation, makes early vision models relatively easy to compute. Thus, it is possible to include greater detail in such models. Greater biologically inspired detail enables the creation of networks whose developmental process and final organisation can be compared against their biological counterparts.

- The relatively well understood structure and function of biological early vision systems simplify the design of fitness functions for evolutionary exploration of developmental models.

- It is possible to study visual development issues in two stages. The fundamentals of developmental self organisation can be studied by modelling the development of early vision systems. Subsequently, more complex issues involving the development of feedback, synchronisation, learning, and modularity can be addressed by attempting to model the development of a basic shape recognition system.

1.7 Report outline

The report consists of 5 chapters. In Chapter 2 current shape recognition systems are reviewed, emphasising key attributes of hierarchical feature extraction. Biological vision is presented in Chapter 3, preceding a review of biological neural development in Chapter 4. Chapter 5 concludes the report with a review of developmental artificial neural networks currently being investigated.

2. Current Shape Recognition Systems

This chapter reviews current thinking on the organisation of shape recognition neural systems and provides a context for the proposed developmental research.

2.1 Context, texture and shading

Form processing involves more than simple 'template matching' of figure outlines. Colour, texture, shading, and context processing are part and parcel of the form recognition process [Manjunath and Chellappa, 1993, Roland and Gulyas, 1994]. For instance [Sakai and Finkel, 1994] demonstrate how the cortex can discriminate shape and depth from texture gradient information alone. Similarly, [Roland and Gulyas, 1994] describes the role of memory, and context in visual representation. It is assumed that these processes are handled (possibly in parallel) by distinct processing modules. Multiple connections between the modules activate ensembles of 'shape neurons' based on all the information available about an object.

2.2 Foveation and selective attention

[Rolls, 1994] presents evidence that the visual system processes the foveal image preferentially. Attention mechanisms are assumed to direct the eye. This view complements the model put forward by Neisser in 1967 (see [Goebel, 1993]) in which a two-stage, object-based selective attention mechanism first segments the field into separate objects using Gestalt principles, and then selects one of the objects for detailed analysis.

Olshausen, Anderson, and van Essen [Olshausen et al., 1993] have developed a neural model of attention which includes position and scale invariant representations of objects (see 2.3.1). The issue of selective attention, i.e. the separation of multiple objects within the field of view is not addressed. However, it is reasonable to assume that selective attention would require some shape recognition capability and would therefore need to be integrated into the shape recognition module (see [Jakubowicz, 1989] for an example of a selective attention network).

An example of foveal attention implemented with traditional preprocessing techniques is included the 'What and Where' neural network proposed by Carpenter, Grossberg, and Lesher [Carpenter et al., 1992]. A series of coarse-grained pre-processing operations centre, orient, and scale the image (in that order) before recognition by the 'What' neural network.

2.3 Scaling, translation, and rotation invariance

2.3.1 Invariant feature extraction

There are three approaches to invariant shape recognition [Barnard and Casasent, 1989,Barnard and Casasent, 1991].

- Transformed versions of training images can be included in the training set, for instance the shape of an object may be defined from many view points. Clearly, this approach becomes computationally expensive as the training set size and complexity increases.
- The pattern recognition system is structured such that transformed versions of an input pattern are treated as equivalent by the system (see for instance [Giles and Maxwell, 1987,Reid et al., 1989]). These techniques tend to be ad hoc and become difficult to develop for high dimensional data.
- Images can be pre-processed into an invariant feature space, for instance use of the real component of the Fourier transform of an image is translation invariant (for examples of this approach see [Troxel et al., 1988, Onodera et al., 1992, Kadiran and Patnaik, 1993,Hung and Lin, 1993]).

2.3.2 Example invariant feature extraction implementations

In [Olshausen et al., 1993] a layered, hierarchical arrangement of control neurons actively routes the inputs of the shape recognition module to the most active parts of the input image (these could be areas of highest luminance, or the receptive field areas of the feature detectors with the highest response). The router neurons in each layer have local, overlapping connections to the preceding layer. Each neuron can 'select' a subset of its inputs. The choice of neuron inputs in each layer allows arbitrarily complex regions on the image retina to be scaled and mapped onto the inputs of the shape recognition module. The paper demonstrates self-organisation of the router neurons to give mappings from the most active image area onto the shape recognition module.

In [Minnix et al., 1991] log polar encoding is used to map the retinal image to the processing layers, as in biological systems. The effect is that rotation and scaling are transformed into shifts (translations). An invariance layer subsequently performs a binary form of the Walsh-Hadamard transform to effect translation invariance. A similar approach is taken by Bradski and Grossberg [Bradski and Grossberg, 1995]. The same effect appears to be achieved in biological systems through the use of Gabor wavelet functions [Du Buff, 1993, Manjunath and Chellappa, 1993] and overlapping receptive fields in early vision.

[Fukumi et al., 1991] presents a rotation invariant system for coin classification based on linear polar encoding. The polar coordinates of image pixels are fed into an array of feature detectors. Each feature detector has a wedge shaped receptive field. The detectors are hand crafted to be invariant to the location of features within their receptive fields.

[Webber, 1994] describes a self-organising two-layer hierarchical feature extraction neural network. The receptive fields and transfer functions of the neurons in each layer are self organised to be invariant to rotation, translation and scaling of features in the training data. There are no lateral connections between neurons in the same layer. Thus, the training scheme has to rely on random initialisation of the network in order to develop a variety of invariant feature detectors. No means of selecting an 'optimal' set of feature detectors is described.

2.4 Hierarchical feature extraction

2.4.1 Motivation

There are essentially two approaches to shape recognition. Global and local feature matching. Global techniques encode an entire shape as a single entity. Simple Template Matching, and Hough Transforms are good examples of this approach. However, non-trivial applications of global techniques require many prototypes (templates) and are therefore computationally expensive. In general, the computational load of global matching techniques increases sharply with the degree of variability in input images. Furthermore, such systems have to be developed on an ad hoc basis. A system developed to recognise one set of objects, cannot be adapted to recognise new objects.

2.4.2 Compositional hierarchy

Local feature matching techniques attempt to overcome these limitations by incorporating information about the structure of object classes into the recognition process. An image is viewed as a collection of local features with specific spatial relationships. Thus, large global templates are replaced by smaller local templates which can combine in different ways to represent instances of an object. Smaller and fewer templates reduce the computational load. The explicit description of spatial relationships between local features provides a methodology for learning new shapes. A shape is divided into a hierarchy of local features. The hierarchy of features can be represented as a graph. The nodes of the graph represent a probability density for the feature parameters. The arcs of the graph describe the structure of the object (how features are related). Although features extracted at each layer of the hierarchy are independent of each other, their correlations are noted at the next level of the hierarchy [Utans, 1994].

The TRAFFIC invariant object recognition system [Zemel et al., 1990] is an example of a compositional hierarchy system. TRAFFIC encodes each extracted feature with a fixed invariant transformation before spatially grouping features in the next layer in the hierarchy.

A scheme in which the hierarchical feature groupings are self organised is presented in [von der Malsburg, 1988, Konen et al., 1994]. von der Malsburg et al treat each ensemble of activated feature detectors as a labelled graph. The system relates each detected feature to its neighbours through local excitatory connections (called Dynamic Links). Thus, each shape class creates one ensemble of laterally reinforcing detectors. The ensemble activation is then decoded into shape classes in subsequent layers.

2.4.3 Hierarchical feature extraction in biological neural networks

Compositional hierarchy can be implemented in a neural network. The graph nodes are neurons and the arcs are the connections between neurons.

The form pathway within the retina essentially reduces the retinal image to contour outlines. [Tanaka, 1993] presents evidence that biological vision feature detectors are organised hierarchically, with higher level feature detectors representing features made up from local combinations of lower level features. The spatial relationship of detected features is preserved at every level in the hierarchy by limiting the receptive field of each neuron to a relatively small contiguous region in its preceding layer.

The lowest level features in the hierarchy consist of edges, bars, and line endings at various orientations and scales. Mathematically, these feature detectors may be viewed as frequency-band filters such as Gabor functions [Du Buff, 1993] operating on small regions of the image. However, it is important to distinguish between such localised spatial filtering (which performs feature detection) and frequency domain analysis. Indeed, [von der Malsburg, 1988] rules out Fourier domain analysis because Fourier components are globally sensitive while recognition of distorted shapes requires local feature detection. [Rolls, 1994] summarises biological evidence that shape recognition is not based on global frequency domain analysis of the image.

[Manjunath, 1992] demonstrates how hypercomplex cells implemented as Gabor wavelet transforms can detect oriented line ends and mark points of significant curvature in outline shapes.

The lower level visual feature detectors appear to be organised identically in all members of a species. Although such detectors can degenerate through lack of use, their formation does not appear to require environmental information [Purves and Lichtman, 1985] (nonetheless [Softky and Kammen, 1990] shows that learning could also result in similar structures). This is in contrast to the complex high level features encoded in the inferotemporal temporal cortex which appear to be modifiable through environmental interaction in adult life [Tanaka, 1993].

2.4.4 Fukushima's Neocognitron

The Neocognitron [Fukushima, 1988] is one of the most well known shape recognition neural networks and will only be outlined here. Current implementations include improvements such as attentional feedback and motion detection [Fukushima and Kikuchi, 1995]. Only the base architecture is discussed here.

The neocognitron is a biologically inspired multi layered neural network which performs hierarchical feature extraction. Outline (skeletal) two dimensional images are recognised inspite of (small amounts of) rotation, translation, scaling, and deformation. Each level of the neocognitron consists of 2 layers: a shape detection layer and a variation-tolerance layer. These are intended to model V1 simple and complex cells respectively. The combination of the two layers provides a balance between discrimination and generalisation while retaining local spatial relationships (see [Templeman and Loew, 1989] for an interesting variation on this theme).

The shape recognition layer consists of arrays of feature detectors. Each feature detector comprises an array of detector cells. Each detector cell receives data from a small local area from the preceding layer. The receptive fields of the detector cells are shifted and overlapped so that a feature located anywhere in the preceding layer may be identified. The output of each shape recognition layer feeds into a variation-tolerance layer. The variance toleration layers essentially produce a more fuzzy version of their inputs so that small variations in their inputs are filtered out. A typical neocognitron may consist of 4 or 5 levels of processing. Supervised and unsupervised training schemes have been developed. However, at present, the need to determine a large number of network parameters on an ad hoc basis limits the usefulness of the neocognitron.

2.4.5 Mueller, Blackman and Furman

[Mueller et al., 1990] present a hierarchical feature detection neural system based on biological principles. The 1990 implementation has three layers. Layer 1 is the input layer. Layer two models the retina and consists of separate arrays of on-centre and off centre units which act as edge detectors. Layer three models the feature detectors of V1 and comprises edge orientation detectors, line (bar) detectors for various orientations and thicknesses; and detectors for end-stopped, oriented edges and bars. The separate on and off centre edge detectors are used as push- pull inputs into the third layer and are reported to improve performance due to improved contrast. The implementation simplifies biological processes where possible. For example on/off-centre fields are implemented through feedforward positive and negative weights instead of using inhibitory horizontal cells. Similarly, the LGN is skipped because it is assumed the LGN acts only as a relay station. The authors raise questions on how to select the right set of primitive features, and how to represent any given shape with an optimal set of features. They point out that attention is necessary so that only a small part of field of view is processed with the necessary resolution. Otherwise system complexity would be unacceptably high. For instance, the reported system requires of the order of one million neurons to process a 200 by 200 ‘pixel’ retina (approximately the same number of neurons as would be required by the neocognitron).

2.4.6 Rolls

The model proposed by Rolls [Rolls, 1994] is based on neurobiological data from the inferotemporal temporal visual cortex. The data referred to relates largely to face recognition, but is extrapolated to the general case. The 1994 paper does not present a fully working system. The essential features of the model are:

- Hierarchical feature extraction and association.
- Spatial relationships of features are retained by local grouping of features at each level of the hierarchy.
- At all levels of processing, features are encoded by ensemble activity of many neurons, i.e. there are no 'grandmother' cells in the model.
- Each layer consists of a set of self organising competitive neural networks with overlapping inputs. Connections from one layer to the next are Gaussian distributed. The lateral inhibition neighbourhood for each neuron is equal to its receptive field size.
- There is a convergence rate of about 6-7 times from each region of one layer to the succeeding layer, allowing higher level feature detectors to detect larger feature combinations.
- Images are pre-processed for edge detection prior to presentation to the system.
- Invariances to scaling, global frequency filtering, translation, and viewpoint are included in the model through competitive self-organisation at every level. Self organisation is achieved by a modified Hebb rule. Viewpoint invariance is achieved through self-organised clustering of the features obtained from different views of the same object. Translation invariance is achieved in a similar way by assuming that the shape will be presented to the system in a variety of locations during training. The self organisation process is assumed to cluster these shapes into a single representation at higher levels of feature detection/association.

2.4.7 Hummel and Biederman

Hummel and Biederman [Hummel and Biederman, 1992] are also proponents of the concept of recognition by components using hierarchical feature recognition. However, their approach differs from others in the assumption that there are a set of fundamental features (Geons) into which all images are decomposed for recognition. Biederman's theory relates specifically to viewpoint invariance of three dimensional objects. These ideas are in part supported, and in part negated by the findings of [Tanaka, 1993]. Recordings of activity in individual neurons in the inferotemporal temporal cortex of rhesus monkeys revealed that these cells respond to specific, complex geometric shapes irrespective of context. However, Tanaka suggests these shapes are more complex than Biederman's geons and are learnt and modified during adult life. Intrator et al [Intractor et al., 1995] argue that psychological evidence suggests that higher level features are learnt. However, lower level features could be fixed.

A seven layer implementation of Biederman's model (symbolic, not neural) has been reported. Layer 1 detects edge orientations and line terminations. Layer 2 groups these into vertices, axes, and blobs (indicators of size and boundary area). Layer 3 recognises basic geons shapes. The geons are labelled with spatial descriptors (location, size, etc) in layer 4. Layer 5 produces a relational spatial description which is invariant to size, location and small changes in orientation (e.g. x is above y). Layer 6 groups geons into objects and layer 7 decodes the layer 6 output into object classes.

2.5 Desirable features of hierarchical visual feature extraction systems

The systems reviewed above have a number of principles in common. A number of other operational principles which are generally agreed to be necessary but poorly understood (such as selective attention). The following is a summary of desirable features encountered in current shape recognition systems.

- Input images are divided into outline (skeletal) images and blobs.
- Low resolution (blob) images are used to indicate size, orientation and location of objects in the field of view. Attentional mechanisms lead to an object-centred image representation at higher levels of the form processing hierarchy.
- Shape recognition is primarily based on the edge detected image, though clearly colour, texture and context information are also incorporated in the iterative loop towards recognition.
- Features are extracted locally and grouped together hierarchically.
- Each level of the hierarchy consists of two processing layers. The first layer is made up of an array of feature detectors. The second layer provides tolerance to small spatial variations in the detected features.
- Higher level features are made up from local combinations of lower level features. Local grouping, combined with the two layer processor architecture provides a degree of invariance to rotation, scaling, translation, and deformation.
- Each feature detector comprises an array of detector cells.
- The receptive field of each detector cell spans a relatively small, contiguous region in the preceding layer. Neighbouring receptive fields are shifted but overlapped so that a feature placed anywhere in the input field is detectable.
- Low level features such as edge orientations and line endings are application independent constants, while higher level features (such as complex geometric shapes) are task dependent and must be learnt.
- At every level of the hierarchy, there is a need for local competition among feature detectors. The winners should represent the optimal set of feature detectors for the image.

- The works of Fukushima, Rolls, von der Malsburg and Tanaka suggest that the modified Hebb rule is capable of self organising the high level feature layers.
- A typical shape recognition module for two dimensional images (i.e. excluding attention, context, and other feedback modules) may be expected to comprise a six level hierarchy, as follows.
    - level 1 : retina like edge detection.
    - level 2 : oriented line, bar, and line-end detectors.
    - level 3a : grouping of closely adjacent, discontinuous, and parallel line segments.
    - level 3b : detection of vertices, crossings, and other points of significant curvature.
    - level 4 : detection of simple shapes.
    - level 5 : detection of complex shapes.
    - level 6 : ensemble encoding into object classes.

3. Biological Vision

Sight is a vital sense. Visual information represents 40% of our sensory input [Dowling, 1992] with half the estimated 2x10^9 neurons in the cortex being related to vision. The ability to see has evolved through natural selection. Different species have evolved different visual systems tuned to to their specific requirements [Moravec, 1988, Wässle and Boycott, 1991]. The primary data considered in this chapter is related to the primate visual system because in many other species there is greater emphasis on motion detection than form recognition.

The visual system is able to perform a vast array of functions such as colour vision, binocular vision, selective attention and motion detection. These functions are integrated seamlessly. This integration means that any discussion about individual visual functions is fraught with uncertainty [Kuffler et al., 1984, Dowling, 1992]. With this proviso, the following sections aim to detail the neural 'circuitry' and mechanisms by which static, spatial information is processed in the primate visual system.

3.1 Functional organisation

The primate visual pathway may be viewed as consisting of three stages: the retina, the lateral geniculate nucleus (LGN) and the visual cortex.

A number of operational principles can be observed throughout the visual system, as reviewed below.

- **Self-calibration** The differences between signals, rather than absolute values, are emphasised. Differentiation occurs in the domains of space, time and spectral composition. For instance, on and off-centre receptive fields in the bipolar cells of the retina both perform edge detection. However, for a given input, their outputs move in opposite directions, thus creating a push-pull effect on subsequent processing stages.

- **Segregation** Functional processes are segregated in several ways. Multiple cell types create a division of labour. For instance, form, depth, colour and motion are processed by a number of distinct, but overlapping, modules. In the retina two classes of ganglion cells project local and global interpretations of the image to the LGN. The LGN further subdivides the streams into a tripartite arrangement: Magnocellular, Parvocellular-Interblobs and Parvocellular-Blobs. The three streams then project to the visual cortex where the segregation of visual information is maintained throughout most levels in the visual cortex. The parallel streams are illustrated in Figure 3.1. Further examples of segregation can be found in the on and off-centre pathways of the optic nerve, and the ocular dominance columns in the LGN.

- **Hierarchy** Thirty-two cortical areas have been identified as having a function related to vision and are interconnected by 305 pathways [Van Essen et al., 1992]. Anatomical studies have identified a hierarchy of ten levels of cortical processing with both inter-layer connections and connections between areas at the same level [Van Essen et al., 1992]. The functional hierarchy of the form processing part of the visual system is reviewed later in section 2.1.1.

- **Feedback** Feedback occurs at all levels within the visual system. The pathways connecting the visual cortical areas not only form forward paths from lower to higher levels, but also include return paths [Livingstone and Hubel, 1988, Van Essen et al., 1992]. There is also substantial feedback from parts of the visual cortex back to the LGN. This is thought to suppress featureless areas of the image whilst enhancing areas rich in features [Van Essen and Anderson, 1995]. In addition to hierarchical feedback, lateral inhibition and reinforcement also occurs within most processing levels. For instance, the orientation columns of V1 have local lateral connections.

- **Parallel Distributed Processing** The analysis of form, depth, colour and motion is not a simple one-to-one mapping of processes to neural pathways [deYoe and Van Essen, 1988]. There is an overlap in the processing of information which is subdivided across cortical areas (see Figure 3.1). Each cell, while functionally specialised, can nonetheless contribute to more than one visual function. Van Essen [Van Essen et al., 1992] argues that having such an interconnected modular system allows greater computational flexibility since modules can be optimised for specific tasks.

3.1.1 Form recognition

It is suggested that each of the three parallel pathways in the visual system, performs a different function. The system is illustrated in Figure 3.1. The Magnocellular pathway performs motion and depth analysis; the 'where' in the visual field [Mishkin et al., 1983]. The Parvocellular-Blob (PB) pathway is sensitive to changes in wavelength and hence is associated with colour processing. The Parvocellular-Interblob (PI) pathway is primarily responsible for the processing of form, since it contains neurons which are capable of the high resolution required in extracting detailed information from static objects. This is termed the 'what' pathway [Mishkin et al., 1983]. The magnocellular pathway, associated primarily with the motion, is connected with the PI pathway in V4 demonstrating the principle of inter-communicating processes [Ferrera et al., 1992]. The focus of the following discussion will be the PI pathway.

The PI pathway exploits two architectural principles for invariant form recognition: hierarchical feature detection, and localised receptive fields. The receptive field of a neuron is defined as that area of the total input space that a neuron is sensitive to. Cells in one layer receive inputs from a number of cells in the preceding layer, providing increasing levels of abstraction.

Thus, the complexity of the receptive fields increases at higher levels of the processing hierarchy. The functional hierarchy of the PI pathway is described below. More detailed descriptions are given in subsequent sections.

- **Edge detection.** The ganglion cells in the retina respond to differences in contrast within the image. Ganglion cells with overlapping receptive fields are thus able to detect edges within an image [Marr and Hildreth, 1980].
- **Path segregation** The receptive fields in the LGN have similar properties to those found in ganglion cells that project to them. Thus, it appears very little change is made to the signals the LGN receives from the retina. The main function of the LGN seems to be the grouping and rerouting of received signals. For instance, the inputs from the two eyes are grouped together into ocular dominance columns. Furthermore, the parvocelluar (P) pathway is subdivided into the P blob (PB) and P interblob (PI) pathways.
- **Oriented edge detection** There are at least twelve different types of receptive fields in V1 [Van Essen and Anderson, 1995]. Several of these perform, functions related to motion, depth, and colour perception. There is a hierarchy of receptive fields within V1 so that the input from the LGN is transformed into more complex features [Kandel et al., 1995, Van Essen and Anderson, 1995]. This hierarchy of receptive fields results in cells which are sensitive to the orientation of lines. Some of these edge-orientation detectors are able to identify line-ends [Kandel et al., 1995].
- **Feature integration** The process of integrating together extracted edges begins in V1 and is extended in V2 [Kandel et al., 1995, Peterhans and von der Heydt, 1991]. In V2 the presence of corners and the length of lines are determined. V2 also responds to lines interrupted by gaps as opposed to V1 where complete lines are necessary stimuli [Peterhans and von der Heydt, 1991].
- **Complex feature detection** Studies on primates have shown that V4 integrates the information abstracted by V1 and V2 to enable more complex features to be recognised. The highest level of direct cortical visual processing is the inferotemporal (IT), where recognition may occur [Tanaka, 1993]. The receptive fields of IT cells are large, such that all cells incorporate complete representations of the fovea [Kandel et al., 1995].

3.2 Anatomical organisation and physiology

As the levels within the visual system are ascended the degree of understanding in terms of both neural structure and function decreases. The retina is one of the most studied organs in vertebrates because of its accessible location [Dowling, 1992] whilst cortical visual areas are only now being investigated.

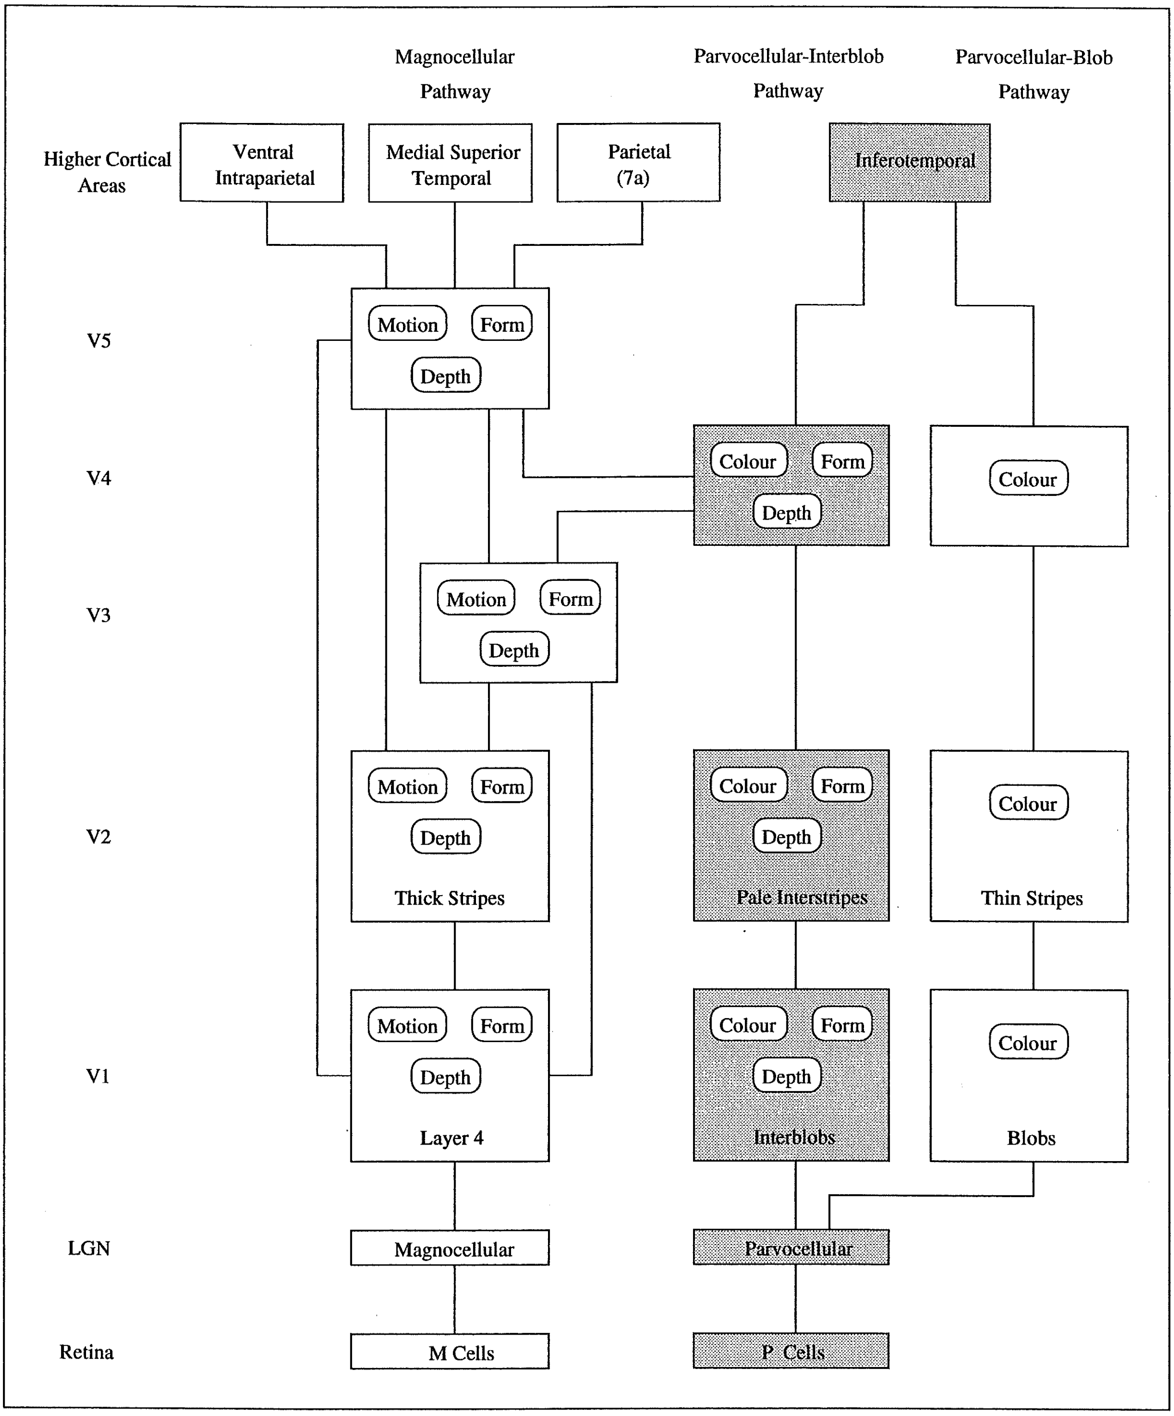

Figure 3.1: The Visual System. Adapted from [deYoe and Van Essen, 1988, Kandel et al., 1995]

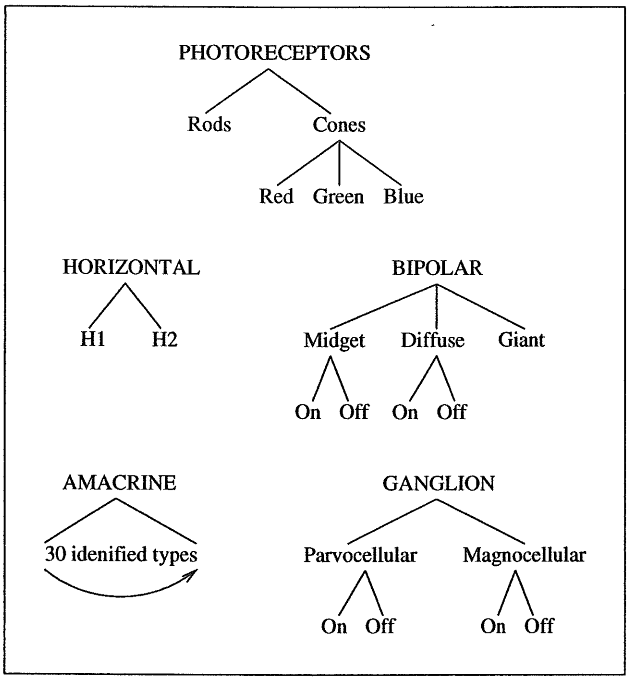

Figure 3.2: The 5 Classes of Retinal Cell.

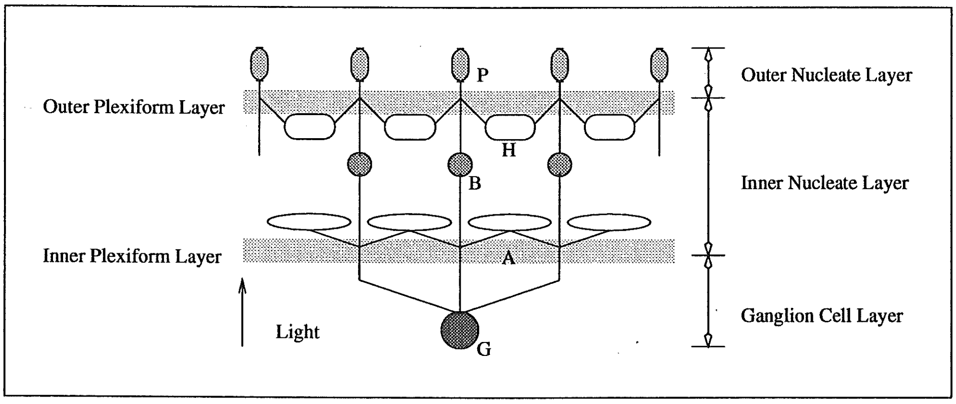

Figure 3.3: A Simplified Diagram of the Retinal Architecture. Key : P-Photoreceptor, B-Bipolar, H-Horizontal, A-Amacrine and G-Ganglion Cells

3.2.1 Retina

The retina is described as having a 'stereotyped architecture' [Van Essen and Anderson, 1990] which is 'highly ordered anatomically' [Dowling, 1992]. The retina contains five major cell types arranged in three cellular layers; the Outer Nucleate Layer (ONL), the Inner Nucleate Layer (INL)and the Ganglion Cell Layer (GCL). Synaptic connections occur between cells at the Inner and Outer Plexiform Layers. Figure 3.2 summarises the five classes of cells and Figure 3.3 illustrates a simplified architecture of the retina. Photoreceptor, horizontal and bipolar cells have graded responses whilst amacrine and ganglion cells generate action potentials [Dowling, 1992].

The retina is not however a homogeneous sheet of cells. Studies on rhesus monkeys show that the density of rods, cones and ganglion cells reduces with increasing eccentricity. The highest density of cones and ganglion is within the fovea. No rods are present in the fovea. The fovea represents the area of highest visual acuity in the eye, since it has the greatest density of cells. The anatomic limit is reached in the fovea where a single cone cell connects to midget ganglion cells through midget bipolar cells. No rods are present in the fovea but rods dominate cones in the periperhal retina. Towards the periphery of the retina, the spatial grain, responsible for visual resolution, decreases as the cone-ganglion cell ratio increases.

Studies on one class of ganglion cells in cats, have shown that as the density of cells decreases with increasing eccentricity, the dendritic field size increases at the same rate. With varying eccentricity the product of cell density and dendritic field size therefore remains constant. Other cell types within the cat retina exhibit the same inverse relationship. Hence in cats, it seems that in the centre and periphery of the retina 'there are identical functional circuits differing only in spatial scale and number of cells' [Wässle and Boycott, 1991].

In primates a similar situation is true for eccentricities of greater than 3mm, with equivalent ganglion cells having large dendritic fields in the peripheral retina, decreasing towards the central retina. Under 3mm, the dendritic field remains relatively constant as a ganglion cell receives input from only one cone. However, the midget ganglion cells may be considered as just a particular class of ganglion cell with the general rule of scaled functional circuits still applying at larger eccentricities.

Experimental evidence has shown that the OPL processes static, spatial elements of images whilst the IPL deals with transient or temporal aspects [Dowling, 1992]. Therefore, the following sections will concentrate primarily on the characteristics of cells which interact within the OPL, with less emphasis on IPL cells.

3.2.1.1 Photoreceptors

The two major classes of photoreceptors, rods and cones, both convert light into electrical signals but have different functions. Cones respond to bright light and are responsible for colour vision. Rods function at low light levels needed for nocturnal vision [Van Essen and Anderson, 1990]. Numerically, there are thought to be over 120 million rods compared to 6 million cones. Despite this large number of photoreceptors and their high packing density, there is no crosstalk in their signals [Dowling, 1992], although gap junctions are known to exist between adjacent rods and cones.

The highest density of cones is found in fovea. In this area it is the cones and not the rods which dictate visual resolution. Further, although colour processing is an important trait, the distribution of red, green and blue cones is not central to the recognition of spatial images and thus need not be addressed.

Cones form a regular hexagonal mosaic within the primate fovea, with a minimum centre-centre spacing of 2-3 μm [Wässle and Boycott, 1991]. Outside of the fovea the mosaic becomes more irregular. However, it is the relationship between cones (input) and ganglion cells (output)

Property	Minimum	Maximum
No of terminals per cone	20	30
No of bipolar connections per terminal	-	1
No of horizontal connections per terminal	-	2
No of horizontal connections per cone	40	60
No of horizontal cells per cone	6	8
No of contacts by 1 horizontal cell to 1 cone	5	10

Table 3.1: Summary of OPL Connections. Modified from [Wässle and Boycott, 1991]

which is important. In the fovea there are 0.3 cones per ganglion cell [Wässle and Boycott, 1991], indicating that each ganglion cell must be driven by at least 2 cones [Dowling, 1992]. At an eccentricity of 10mm however this ratio has risen to 16 cones to one ganglion cell. Hence, the ganglion cell at this position is integrating information over a wider area of the retina.

3.2.1.2 Horizontal cells

Horizontal cells mediate the lateral connections in the outer plexiform layer. There are two distinct morphological classes of horizontal cells, H1 and H2, but no great functional differences have been found between them [Wässle and Boycott, 1991].

Horizontal cells synapse into cone terminals (cone pedicles) at junctions called triads which are observed in many retinas [Dowling, 1992]. The two outer post-synaptic contacts come from dendrites of horizontal cells, with a dendrite from a bipolar cell forming the central connection. Each cone pedicle has between 20-30 such triad synapse connections. It is thought that these connections enable the horizontal cells to mediate the surround response in bipolar receptive fields. A quantitative summary of OPL cell connections is shown in Table 3.1.

Neighbouring horizontal cells are connected by gaps junctions. These junctions enable the connectivity of the horizontal cell matrix to vary, such that a cell's response can be governed by distant as well as near neighbours. This results in the modification of the surround response of bipolar cells to which a horizontal cell is connected. Dopamine infusion experiments have demonstrated this adaptive behaviour [Boussard and Vibert, 1993].

3.2.1.3 Bipolar cells

Bipolar cells mediate the vertical flow of visual information through the retina. Midget bipolar cells, found only in the fovea, are associated with colour processing and are thus not considered essential in the processing of form [Wässle and Boycott, 1991, Mariani, 1984]. The following discussion relates to diffuse bipolar cells.

Diffuse bipolar cells have dendritic fields which overlap to a high degree on the cone mosaic, allowing the cells to connect to 'several' cones [Wässle and Boycott, 1991]. Bipolar cells have antagonistic centre-surround receptive fields, which can be either On-centre or Off-centre. This differentiates bipolar cells into two distinct populations and initiates the process of segregating visual information. Connections to cones mediate the centre response of bipolar cells with horizontal cells providing the surround. It is thought that the triad connections described previously enable such functionality [Dowling, 1992]. The axons of bipolar cells synapse onto amacrine and ganglion cells in the INL [Wässle and Boycott, 1991].

3.2.1.4 Ganglion cells

The ganglion cells are the output of the retina which connects to the LGN. The two major classes of ganglion cell can initially be differentiated on a morphological level [Wässle and Boycott, 1991, Tessier-Lavigne, 1991, Shapley and Perry, 1986, Kaplan and Shapley, 1986]. Parvocellular cells (P) constitute 80% of total ganglion population and have small cell bodies producing small dendritic arbors. Magnocellular cells (M) represent 10% of the population. They have large cell bodies with large, sparse dendritic trees. The remaining 10% of ganglion cells consist of a variety of types of cells, whose functional roles are less well understood.

Ganglion cells can also be distinguished in four major physiological ways [Livingstone and Hubel, 1988]:

- Colour : About 90% of P cells are sensitive to differences in wavelength [Livingstone and Hubel, 1988]. This leads to the proposal that P cells are strongly associated with midget bipolar cells because of their cone specific connection behaviour [Wässle and Boycott, 1991]. M cells sum the responses of all three cone types (red, green and blue) and are therefore not wavelength specific [Livingstone and Hubel, 1988]. This also links M cells to diffuse bipolar cells, since diffuse bipolar cells connect to adjacent cones irrespective of chromatic considerations [Wässle and Boycott, 1991].
- Acuity : As with the horizontal cells, the average size of both M and P cell receptive fields, increases with eccentricity [Livingstone and Hubel, 1988, Wässle and Boycott, 1991, Van Essen and Anderson, 1990]. However, at any given distance from the fovea, M cells have receptive field centres 2-3 times larger than P cells [Van Essen and Anderson, 1990]. This would generally lead to the conclusion that P cells have the greater ability to extract detailed information from an image than M cells [Tessier-Lavigne, 1991]. Whilst it is true that M cells respond best to large objects, they also contribute to the analysis of detail [Tessier-Lavigne, 1991].
- Contrast : M cells are much more sensitive than P cells to low-contrast stimuli [Kaplan and Shapley, 1986]. The response of an M cell increases rapidly as the illumination difference between the centre and surround is increased, saturating at about 10-15%. P cells respond more slowly but saturate at higher levels of contrast.
- Speed : M cells respond faster than P cells to stimuli and in a much more transient manner [Dowling, 1992, Livingstone and Hubel, 1988]. This suggests that M cells are

tuned to the detection of temporal variations in images. This may be explained by the fact that M cells receive 30% of their synaptic input from bipolar cells and 70% from amacrine cells [Wässle and Boycott, 1991]. The situation is reversed in P cells.

Ganglion cells have similar centre-surround receptive field properties of bipolar cells, except that they integrate responses from a wider area.

3.2.1.5 Amacrine cells

Thirty types of amacrine cells have been identified, although the form and function of many have yet to be fully documented. Some types of amacrine cells are thought to provide the surround response for ganglion cells (similar to inputs from horizontal cells to bipolar cells) whilst others make connections to cells in the IPL, indicating some form of feedback path [Wässle and Boycott, 1991].

3.2.2 Lateral Geniculate Nucleus

The LGN acts as a relay station between the retina and the visual cortex. It segregates the inputs from both eyes to form a complete neural representation of the image on the retina, before projecting it to the visual cortex. The LGN consists of two nuclei (left and right), both containing six layers of cells which segregate the retinal ganglion cell axons into the Parvocellular (P) and Magnocellular (M) channels. The two classes of layers have distinctly different neural structures.

The receptive fields of LGN neurons are similar to those of the retinal ganglion cells and hence there is no significant alteration in visual information in the LGN. However, only about 10-20% of the pre-synaptic connections come from the retina [Van Essen and Anderson, 1995]. The majority of connections are feedback connections from V1 together with non-visual areas and these connections are thought to control the flow of visual information.

3.2.3 V1

The structure and function of V1 related to the three parallel streams is well covered in literature so the discussion will focus on the projections of the PI form processing pathway.

Since it is part of the cortex, V1 is arranged in six layers. There are two major classes of cells in V1:

- Pyramidal. Pyramidal cells have long spiny dendrites which send axons to other cortical areas. They are excitatory cells.
- Non-pyramidal. These are smaller neurons, subdivided into two types. Spiny stellate cells have many dendritic spines and are excitatory. Smooth stellate cells have no dendritic

spines and are inhibitory. Both mediate local, interneuron connections within the primary visual cortex.

PI projections from the LGN enter at layer 4 and synapse with spiny stellate cells.

There are at least 12 different types of cell receptive fields in the cortex [Van Essen and Anderson, 1995]. Spiny stellate cells have circular receptive fields similar to those of the preceding LGN cells but above this level the properties become more complex. Hubel and Wiesel have classified two major groups of receptive fields related to the processing of form :

- *Simple cells* are pyramidal cells located near the spiny stellate cells of the PI pathway. They have elongated receptive fields which respond to a bar of light at a particular angle of orientation. The bar must lie in a specific location to induce the optimum response. It is thought that simple cells receive their input from spiny stellate cells. Each simple cell is able to detect stimuli in a range of about 10°. Different cortical cells receive input from the same point on the retina with the same receptive fields but with different axes of orientation. In this way all axes of orientation are covered for all points on the retina.
- *Complex cells* are also pyramidal cells but have larger receptive fields with less precise inhibitory and excitatory zones. This infers that the position of a stimulus is less crucial and it is thought that complex cells receive input from simple cells which have the same axes of orientation but respond to different areas of the retina.

Some simple and complex cells are end-stopped, and may either be single end-stopped or double end-stopped [Kandel et al., 1995]. Although there is a definite hierarchy in the properties of receptive fields, with preceding layers providing the input for the next, how such arrangements arise is still unknown [Kandel et al., 1995].

Physiological analysis of V1 has shown that there are columns which are sensitive to particular orientations only. These columns run at right angles to the layers in V1. Columns are arranged such that the change on orientation selectivity varies by 10° arranged in a radial rather than linear fashion [Kandel et al., 1995]. Columns are regularly interspersed by blobs, which are sensitive to changes in colour rather than form.

A complete set of orientation columns representing 360° and accompanying blobs represents a hypercolumn. Hypercolumns are repeated regularly over the surface of the cortex. Kandel [Kandel et al., 1995] suggests that each hypercolumn 'is an elementary computational module' since it contains the necessary functionality to analyse one point on the retina.

3.2.4 V2

V1 has major projections to V2, the next level in the hierarchy. V2 has a distinctive stripe architecture with thick and thin stripes separated by areas of pale interstripes. Analysis of the receptive field properties of the cells in the respective stripes have shown that each class of stripe contains a particular class of receptive field [Zeki, 1993]. The PI pathway is mediated through the inter-stripes which are associated with the integration of form. Although the inter-stripes are separated, the receptive fields of neighbouring inter-stripes overlap [Zeki, 1993]. Therefore a continuous coverage of the visual field is maintained.

3.2.5 V4

V4 is divided into a number of regions where a mixing of P and M signals occurs [Ferrera et al., 1992]. Colour processing in V4 was assessed by Zeki [Zeki, 1993] but areas associated with form processing have been identified. Cells in V4 respond to more complex stimuli. Lesion experiments [Schiller and Lee, 1991, Schiller, 1995] using triangles and crosses as basic stimuli. indicate that the detection of corners and line ends occurs in V4. The lesion experiments have also shown that V4 may also be involved in visual learning as opposed to just simply detecting stimuli [Schiller and Lee, 1991].

3.2.6 Inferotemporal

The inferotemporal (IT) region is subdivided into two hierarchical areas (posterior IT and the anterior IT) [Tanaka, 1993]. The IT is believed to have a columnar architecture with cells that are sensitive to similar critical features, clustered together. The critical features are not however complex enough to specify a particular object through the activation of a single cell. The activation of several cells with different sensitives would however enable a complete object to be specified. Discontinuities arise as the columns are traversed because extra dimensions are needed to represent complex features.

3.3 Summary

This chapter summarised the key functions and structures of the primate visual system associated with the processing of static images. Although the biological visual system is complex, clear pathways have been identified which are primarily associated with the processing of form. The retina is the most widely studied element of the visual system and as progression is made up through the higher levels of the cortex the exact function and structure of the neural circuits becomes less precise.

4. Biological Neural Development

Having outlined the the biological vision system in terms of its neural architectures and functions, this chapter considers how such neural machinery is developed. Although the brain has a precise and intricate form, it is believed that its development is based on simple rules [Gleick, 1988, Stryker, 1994, von der Malsburg, 1990, Udin and Fawcett, 1988]. During early neuronal development, it is intrinsic genetic encoding which dictates how [Stryker, 1994]:

> *Neurons are generated in appropriate numbers at appropriate times; they migrate to appropriate positions; they send out axons that find appropriate paths to their target structures; and the axon terminals recognise the correct cell types and even the right general region in the target structures.*

It is generally accepted [Herrmann and Shatz, 1995] that this process is driven by chemical events in the absence of neural activity, but that activity is responsible for the refining of connections. This chapter summarises the mechanisms currently thought to determine neurodevelopment.

4.1 Gene expression

In theory the brain could be fully developed before birth but this would carry costs in terms of brain size at birth (difficulty of giving birth), carrying, long gestation period, complexity of genetic encoding and less adaptability to changes in the environment. Alternatively, the cost of being born with a completely unstructured brain would be huge in terms of the difficulty of learning tasks and time to maturation. Our genetic encoding therefore represents a compromise in which we inherit generic brain structures optimised for learning specific skills (eg language [Pinker, 1994]). Each generic structure is then tuned to the needs and environment of the individual. The structure information contained in genes is the evolutionary accumulation of adaption to the environment. The provision of neural structures can thus be described as *inherited domain knowledge.*

The process of cellular differentiation is mediated by genetic self-regulatory networks. Every cell in the body carries the a copy of the full set of chromosomes that define the individual. Cells are differentiated because approximately 90% of genes in each cell are normally switched off. At present, the mechanisms by which genes are switched on and off is not understood in detail. At an abstract level, genes are switched on and off according to the current or previous states of one or more of the other genes in the cell.

Stuart Kauffmann [Kauffman, 1969, Kauffman, 1984] has proposed that such regulatory

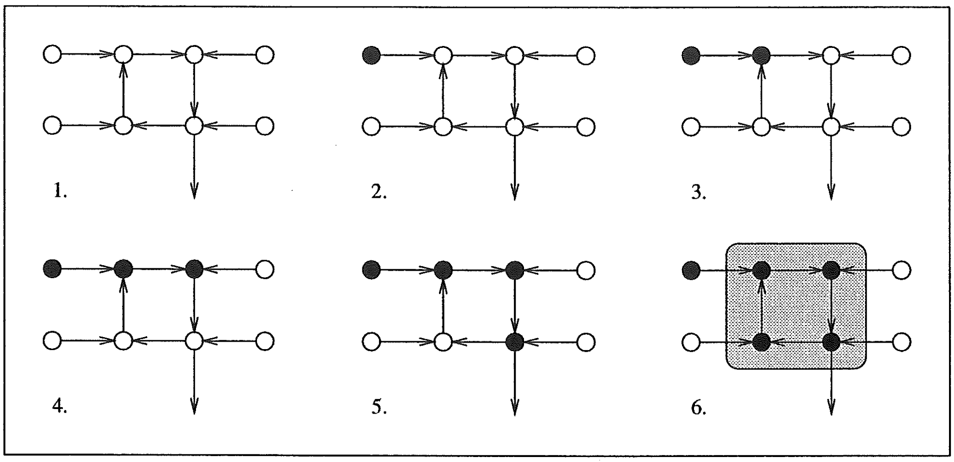

Figure 4.1: States of a simple, randomly connected network of Boolean OR nodes upon the application of a high (black) input. The network freezes into a state which is returned to even if state changes are made to the nodes or their inputs (shaded box). Adapted from [Kauffman, 1991]

behaviour may be modelled by the activity of randomly interconnected boolean functions. Each state is analogous to the execution of a programmed behaviour. An example of a boolean feedback network is show in Figure 4.1.

Analysis of networks with 2 inputs per gate shows that [Kauffman, 1984]:

- Each state cycle has a limited number of states.
- There are a limited number of state cycles.
- States within any one cycle are similar to each other (gradual change) and less similar to states of other cycles.
- State cycles are immune to noise and even if perturbed, only a few neighbouring cycles may be entered.

It seems likely therefore, that gene expression is an emergent property of gene interactions within each cell.

4.1.1 Neural formation

Every neuron in the central nervous system follows a sequential set of processes during development. The early processes are determined by the intrinsic genetic programme. Cell interactions and stimuli from the external environment tune the outcome.

4.1.1.1 Lineage and cell-cell interactions

During development, the neural tube produces a large diversity of neuronal phenotypes from a homogeneous pool of pre-cursor cells. Cell fate can be determined by one of two processes [McConnell, 1995]:

- Cell lineage. Here the progeny are determined purely with reference to the parent cell. Two lineage scenarios are possible single fate per parent and multi-fate progeny generated in a fixed sequence.
- Cell interaction. Each cell is capable of generating a variety of progeny depending on environmental cues.

4.1.1.2 Cell division and differentiation

Neurons and supporting glial cells in the central nervous system are generated by the neural tube in the developing embryo. The time and place of a neuron's birth determines its fate (differentiation) [Alberts et al., 1994]. This process follows the principles previously described in gene expression.

4.1.1.3 Migration

The immature neurons migrate to specific target locations often along pathways provided by a scaffolding of glial cells [Alberts et al., 1994]. The destination of a neuron is determined before migration begins. Neurons of different ages settle in different positions. In the cerebral cortex for example, younger neurons migrate outward past older neurons forming distinct layers.

4.1.1.4 Neurite growth

The intricate neuronal connectivity of the nervous system is generated by neurons extending axons and dendrites (neurites), both before (in some cases) and after they complete migration [Alberts et al., 1994]. The initial stages of neural growth rely on environmental cues to regulate the growth of neurites, but the role of activity driven processes dominates in the later stages [Goodman and Shatz, 1993]. Three major steps exist in the development of connections which are illustrated in Figure 4.2 [Goodman and Shatz, 1993]:

- Pathway Selection. Long range navigation of neurites to find their correct target areas.
- Target Selection. Contact and recognition of a set of neurons within the target neighbourhood relying on local information.
- Address selection. Refinement of an initial coarse-grained mapping to produce finely-tuned connection patterns on a subset of target cells.

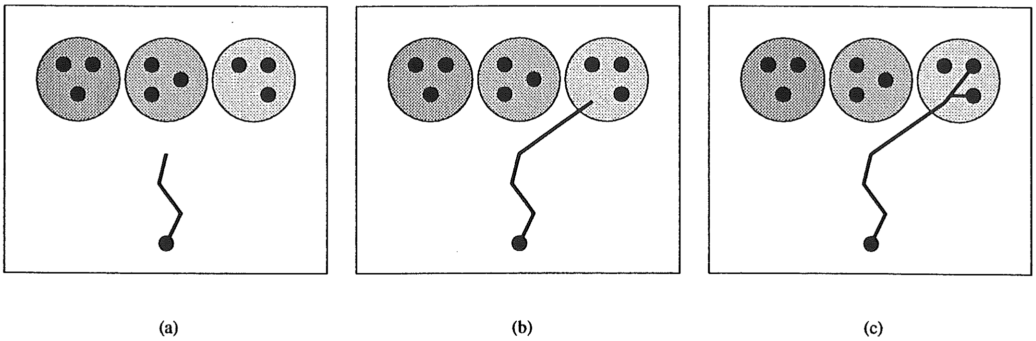

Figure 4.2: The Three Phases of Selection. (a) Pathway selection. (b) Target selection. (c) Address selection.

Neurite growth is regulated by complex structures termed growth cones, which exist on the tips of extending neurites. A growth cone samples the local cellular environment [Tanaka and Sabry, 1995] and responds to the cues by expressing five activities which determine neuronal form : elongation, branching, turning, retraction and synaptogenesis [Letourneau, 1987]. The behaviour of growth cones appears to be mediated by a combination of factors [Goodman et al., 1984]. Both attractive and repulsive influences are associated with guidance [Goodman et al., 1984, Keynes and Cook, 1995, Hall, 1992]. The principle mechanisms of neurite guidance are summarised below.

Pathway Selection

- **Pioneer Axons** In various parts of the nervous system, sensory, pathfinding axons grow out to the target areas establishing tracks [goodman84, hall92]. Subsequent neurons selectively recognise and follow these paths to the target, forming distinct bundles of fibres [Goodman et al., 1984]. However, pioneer axons are not qualitatively different from later axons as the later axons are themselves able to navigate using environmental cues [Hall, 1992].
- **Stepping Stones or Guidepost Cells** Evidence from insects suggests non-neuronal cells, that appear very early during development [Goodman and Shatz, 1993], act as guideposts for growing sensory axons. Ablation of these cells causes axons to be misdirected [Hall, 1992, Purves and Lichtman, 1985].
- **Stereotropism** Guidance is directed by mechanical factors, such as the path of least resistance [Purves and Lichtman, 1985].
- **Galvanotropism** Experimental work has shown that very small voltage gradients surround living tissue and that by applying electric fields, neurite growth can be directed [Hall, 1992, Purves and Lichtman, 1985].
- **Cell-Substrate Differential Adhesion** Proteins in the extracellular matrix of the developing CNS form substrates to which neurites selectively adhere. Neurons recognise and adhere to these molecules through receptors on the surfaces of growth cones and

axons. The integrin family of receptors is sensitive to a wide variety of proteins and individual family members can be differentially expressed during development. This temporal regulation of receptors enables the growing neurite to express different chemical affinities depending on the specific local environment [Goodman and Shatz, 1993].

Target Selection

- **Chemotropism** Ramon y Cajal first proposed the theory of chemotropism, suggesting that gradients of diffusible molecules emanating from a target cell, might guide a growth cone [Tessier-Lavigne and Placzek, 1991, Goodman and Shatz, 1993]. Principle support for this form of axon guidance comes from studies of nerve growth factor (NGF), a trophic factor which regulates cell survivability. NGF and other chemoattractants, act as local guidance cues [Purves and Lichtman, 1985] and the degree to which a growth cone may be reoriented depends on the critical parameters of concentration and steepness of gradient [Tessier-Lavigne and Placzek, 1991].

Address Selection

- **Chemoaffinity and Cell-Cell Differential Adhesion** Experiments during the 1950's on the optical tectum led Sperry to suggest chemoaffinity as the means by which growth cones recognise individual target cells. It was proposed that cells are thought to possess specific chemical labels and that gradients of these labels exist. The hypothesis relies on point-to-point specificity alone and therefore requires a prohibitively large set of labels. Molecular evidence to support the chemoaffinity process is only now being accumulated [Purves and Lichtman, 1985]. Cell adhesion molecules (CAMs) exist on the surfaces of growth cones and the cells they contact [Goodman and Shatz, 1993]. A significant number of CAMs exist [Hall, 1992] and their expression is both dynamic and regional (eg. covering one side of an axon and not the other) [Goodman and Shatz, 1993]. However, CAMs alone cannot provide the point-to-point specificity but must be considered in conjunction with activity-dependent mechanisms [Goodman and Shatz, 1993].

4.1.2 Connectivity development

4.1.2.1 Neuron-neuron connections

Connections between neurons can occur directly through cell bodies (gap junctions) and as the result of neurite contacts (synapses).

- Gap junctions are direct, bi-directional, electrical junctions between cells. They may be implicated in neural development before synaptogenesis to mediate activity-dependent selection of neural connections [Penn et al., 1994, Yuste et al., 1992]. They also enable the inter-cellular transport of molecules, a process thought to be important during development [Kandel et al., 1995]. An alternative theory [Haydon and Drapeau, 1995] suggests that chemical synapses can form rapidly within minutes to hours of neural contact, indicating that gap junctions would therefore not be a necessary precursor to synapses.

- Synapse formation is initiated by the contact of the pre-synaptic growth cone with a target cell. Synapse generation appears to follow an intrinsic program independent of cell interaction [Haydon and Drapeau, 1995]. Synapses are either excitatory or inhibitory. The type of synapse formed is determined by complex interactions between the pre- and post-synaptic cells. The type of neurotransmitter emitted by a pre-synaptic cell and the stimulation duration determines the response of the post-synaptic cell.

  It is possible to infer the type of synapse from its position on the post-synaptic cell [Kandel et al., 1995]. Though, this does not however mean that the site of a synaptic contact will determine a synapse type.

  - Axo-somatic. Tend to be inhibitory synapses and induce fast responses by the post-synaptic cell.
  - Axo-dendritic. Tend to be excitatory and are found either on the main dendrite body or on off-shoot spines.
  - Axo-axonic. Sometimes, axons contact each other before forming synapses with target cells. These connections have a modulatory effect, controlling the amount of transmitter released onto post-synaptic cells.

  Synapses are subsequently pruned by competitive, activity-dependent mechanisms.

4.1.2.2 Neural regulation

Most types of neuron in the vertebrate nervous system are produced in excess and a substantial proportion die shortly after reaching and connecting to their target regions [Alberts et al., 1994,Davies, 1994]. The proportion varies between species but is greatest in those species which are adaptive [Hall, 1992]. This process of cell death provides a simple and elegant [Alberts et al., 1994] means of adjusting the number of neurons to the number of target cells and eliminating misconnected neurons [Hall, 1992, Davies, 1994].

One regulation mechanism is the competition for a limited supply of neurotrophic factor emitted by target cells [Hall, 1992]. If a neuron does not obtain enough neurotrophic factor then it undergoes a programmed cell death [Davies, 1994, Alberts et al., 1994]. The role of neurotrophins is not however limited to cell survival during development. They also modulate synaptic transmission and axonal and dendritic arborization [Cabelli et al., 1995].

Studies on sensory neurons [Davies, 1994] have demonstrated an intrinsic neuron programme which initiates neurotrophin dependence only when axons have reached their target areas. Not only is this dependence correlated with the distance travelled by an axon and its growth rate, but the period of dependence is limited. This complex, intrinsic timing of events is thought to represent an efficient means of eliminating neurons that make targeting errors [Davies, 1994].

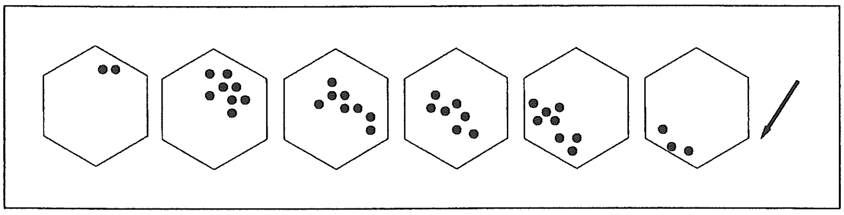

Figure 4.3: Spatial representation of spontaneous action potential waves traversing a portion of a growing retina at 0.5s intervals. The hexagonal grid defines the array of retinal ganglion cells being measured. The dots represent neurons with the highest firing rates. The arrow indicates the direction of the activity wave. From [Wong et al., 1993]

4.1.2.3 Spontaneous activity

Recent studies of the mammalian visual system indicate that pre-natal neural activity regulates the layering of connections between the retina and the LGN even before the presence of visually evoked stimuli [Goodman and Shatz, 1993]. This neural activity occurs initially in the retina before the formation of rods and cones, with retinal ganglion cells spontaneously firing action-potentials [Galli and Lamberto, 1988]. Experiments have shown that neighbouring ganglion cells fire synchronously with a high degree of correlation between firing [Maffei and Galli-Resta, 1990, Wong et al., 1993]. Waves of activity have been seen to traverse the developing retina mediated by horizontal pathways between ganglion and amacrine cells [Wong et al., 1993, Herrmann and Shatz, 1995], suggesting that random neural activity leads to spatiotemporal patterns of activity which refine connections [Goodman and Shatz, 1993, Cabelli et al., 1995]. A simplified example is shown in Figure 4.3.

Activity dependent and activity independent processes are not mutually exclusive; the processes overlap and interact with each other [Meister et al., 1995, Goodman and Shatz, 1993]. Retinal ganglion cell axons growing towards target areas make intermediate connections which ultimately serve no functional purpose and it is through the use of spontaneous activity that these connections are removed [Shatz, 1994]. Total axon length increases during this period as growth becomes concentrated in the target neighbourhood [Shatz, 1994].

Evidence for action-potential activity weeks before eye-opening has recently been found between the LGN and visual cortex in cats [Goodman and Shatz, 1993, Herrmann and Shatz, 1995]. LGN axons make functional synapses with a transient population of subplate neurons during development. Their ultimate targets in the cortex have yet to develop. Activity supports these temporary connections before remodelling occurs on the cortical target cells [Herrmann and Shatz, 1995]. This highlights a method of maintaining connections whilst target areas develop and the overlapping role that activity plays with neuron growth.

At birth the intrinsic programme of patterned, spontaneous activity is replaced by visual input [Goodman and Shatz, 1993].

4.2 Learning

Learning enables the tuning of neural circuits to the specific stimuli of their immediate environment. The number of neurons is set at birth [Purves, 1994] and learning only affects connection strength.

Neural information is encoded in the phase, frequency, and duration of activity spike trains. The frequency and duration of spike trains reflect the level of excitation in the neuron. Phase information appears to encode event sequences (for instance in sensorimotor information processing [Nicolelis et al., 1995], and olfaction [Laurent and Davidowitz, 1994]). Learning involves the adaptation of synaptic connections between neurons as a function of current neural activity and is therefore an iterative process.

A synapse is classed as excitatory or inhibitory according to its neurotransmitter type. The effect of a pre-synaptic spike train on a post-synaptic cell can be excitatory or inhibitory depending on the neurotransmitter receptor type of the post synaptic cell, and the duration of the spike train from the pre-synaptic cell. Thus, learning can determine not only the strength of a synaptic connection, but also the direction of influence of a particular synaptic input on a post synaptic cell.

In 1949, Donald Hebb suggested that learning changes the strength of a synapse as a function of the degree of correlation between pre and post synaptic activity [Hebb, 1949]. The principle of Hebbian learning has been demonstrated to enable phase-locking, co-activation, and competitive behaviours [Lansner and Fransen, 1992] amongst interconnected neurons.

4.3 Summary

This chapter presented the stages of biological neuro-development. It was shown that genetic information controls coarse-grain structure and that neuronal circuits are the result of interactions between activity-independent and activity-dependent processes. There are no distinct boundaries between the phases of neural development. The mechanisms active on a growing neuron are simple but their combined effect is the development of complex neural structures.

5. Related Work

Automatic construction of neural network architectures has been a popular source of research. The goal is to create neural networks which are optimal for different applications. This chapter presents the motivations that have led to the investigation of developmental neural networks. A number of models are reviewed and a summary of findings is made.

5.1 Comparison of constructive and developmental approaches

Most of the constructive methods directly map between the genotype contained in the GA and the phenotype, the final structure. This results in the instantaneous mapping of adult structures onto a network. Direct mapping has a number of inherent shortcomings which may be solved by biologically-inspired developmental models.

- Direct mapping limits the range of morphologies produced [Dellaert and Beer, 1994b,Nolfi and Parisi, 1991]. It is stressed that it is not possible to read a cell's phenotype directly from a genotype since biological mapping is a non-linear process [Nolfi and Parisi, 1991]. The non-linearities arise from the interactions between genetic instructions and because development occurs within a physical environment [Nolfi and Parisi, 1991]. Developmental models create a greater range of complex morphologies [Nolfi and Parisi, 1991] from a compact set of encoded instructions [Dellaert and Beer, 1996].
- Modularity and symmetry breaking are desirable in many situations [Dellaert and Beer, 1994a]. Developmental models support these physical traits.
- Scalability is a problem with direct mapping [Nolfi and Parisi, 1991,Dellaert and Beer, 1994b]. As the size of the network increases performance decreases significantly. Biological development is essentially an incremental process and thus may solve the scaling problem [Dellaert and Beer, 1996].
- Development is robust process where genetic instructions that have minimal effect are filtered out [Dellaert and Beer, 1994b]. A larger spectrum of solutions may also be sampled, as development can perform local-hill climbing up to long jumps within the search space [Dellaert and Beer, 1994b].

The combination of evolutionary strategies and developmental rules is thought to represent a more fundamental approach to growing neural networks [Kodjabachian and Meyer, 1994].

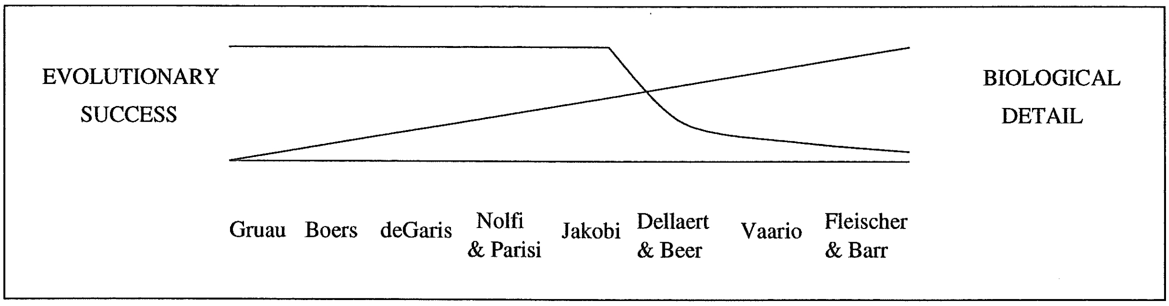

Figure 5.1: The spectrum of developmental neural network models.

5.2 Developmental Models

Modelling development in neural networks can be traced back to the work of Kitano [Kitano, 1990], and Belew, McInerney and Schraudoplh [Belew et al., 1990]. Developmental modelling now ranges from models based on computational methods, with minimal biological influence, to those which are presented as being biologically plausible. The use of evolutionary strategies also varies across this range with the greater success being achieved in those models which are less biologically defensible. A *spectrum of developmental models* is proposed based on these observations and is illustrated in Figure 5.1. The scale and terms of reference are acknowledged to be arbitrarily assigned. The spectrum does however enable a means of placing some degree of order on the variety of current developmental neural networks.

The following sections briefly review and comment upon work in developmental neural networks. Approaches which are thought to most closely resemble biological development are covered in greater detail. Models which are deemed to be the least biologically defensible are covered first, concluding with those methods which are most biologically plausible. A more comprehensive review of 5 developmental models (Boers and Sprinkhuizen-Kuyper, Vaario, Nolfi and Parisi, Gruau, and Dellaert and Beer) can be found in [Kodjabachian and Meyer, 1994].

Developmental methods are used to meet different aims. Some groups are involved in understanding fundamentals of biological systems in terms of self-organisation [Vaario, 1994, Fleischer and Barr, 1994] and the relationship between evolution and development [Nolfi and Parisi, 1991]. Other groups are evolving networks for particular functions such as simulated animal behaviour [Jakobi, 1995,Dellaert and Beer, 1994b] and control [Gruau, 1994b], together with more traditional applications such as character recognition [Boers et al., 1995].

5.2.1 Gruau

Gruau has developed a grammar encoding scheme termed *Cellular Encoding* [Gruau, 1994c]. The basic elements in the model are cells which have defined inputs and outputs to other cells. The encoding scheme implements a variety of cell divisions (graph transformations) which are successively applied to a single cell eventually producing a neural network. For example, a cell may split with a *sequential* division such that one child cell inherits the inputs of the parent and the second child cell the outputs. The two child cells are connected together. A more detailed coverage may be found in [Gruau, 1994c, Gruau, 1994b, Gruau, 1994a].

Each cell division is given a label (eg S for sequential) and by combining a number of labels a sequence of instructions is encoded. A sequence of labels is stored as nodes in an ordered tree, termed a grammar tree. A set of grammar trees may be used in the development process, with specific labels allowing movement between trees. A cell may be instructed to divide, stop growing and become or neuron, jump to another tree in the set, or set an internal register equivalent to a weight.

Each cell in a growing network has a copy of the the cellular code, the set of grammar trees. A cell has a 'reading head' which traverses the current grammar tree enabling the cell to execute the instruction read. Genetic Programming [Koza, 1992] is used to evolve the set of derived grammar trees. The technique allows sub-branches of trees to be grafted onto branches of other trees.

In its simplest form cellular encoding can generate Boolean, feedforward networks [Gruau and Whitley, 1993]. No learning in the traditional sense of artificial neural networks is undertaken as it is argued that having to both evolve and then train networks is prohibitively long. One of the strong results from this model is the ease with which modular neural networks can be created. Hierarchies of networks can be created by recursively calling sequences to generate and combine sub-networks.

The technique has been successfully applied to the evolution of simple Boolean functions and also in the control of a six-legged robot. The robot example demonstrated that cellular encoding can also exploit symmetry by evolving six individual, interconnected leg control networks. No a priori information was provided to the system.

A modified version of Hebbian Learning has been added to cellular encoding [Gruau and Whitley, 1993]. This form of developmental learning (the Baldwin Effect) modifies the genetic landscape accelerating the search of the genetic algorithm. A criticism of the model is that only Boolean weighted networks could be evolved [Gruau, 1994b]. However, a control application has recently been investigated using real numbers [Whitley et al., 1995].

5.2.2 Boers and Sprinkhuizen-Kuyper

The work of Boers and Sprinkhuizen-Kuyper uses a combination of genetic algorithms and L-system grammars to develop classical feedforward neural networks [Boers et al., 1993]. Networks are trained using a standard back-propagation algorithm and the error obtained is used as the fitness measure for the genetic algorithm.

L-systems were proposed to model the development of plants and trees. However, the natural structure of neural networks is a graph representation and not a tree. Therefore, Boers has developed an interpretation of L-systems to create a parallel, context sensitive graph rewriting grammar. The L-system grammar consists of a series of production rules that are applied to a starting axiom. These rules can be applied indefinitely so hence the model is limited to 6 developmental steps. A GA is used to find the best set of production rules.

A criticism of the model is that only results of very simple networks have been presented [Gruau, 1994b, Kodjabachian and Meyer, 1994]. However, modular architectures with connections traversing more than one layer have been demonstrated. Gruau argues [Gruau, 1994b] that because the model is a complex interpretation of a tree grammar it is an inefficient means of coding. For example, the string of production rules created by the genetic algorithm must be validated before application.

In response to this criticism Boers has adapted the encoding scheme to implement graph grammars, termed G2L-systems [Boers, 1995]. This modification enables sub-graphs or modules to develop together with recursive networks. Results have yet to be presented because of visualisation problems.

5.2.3 deGaris

deGaris and the Brain Builder Group at ATR, Human Information Processing Labs in Japan aim to develop a billion neuron artificial nervous system by the year 2001 [deGaris, 1995]. Cellular Automata (CA) is used to grow networks where the state transition rules encode the sequence of development. Sequences are evolved by a GA and are mapped onto a CA network. Sequences are continuously fed into a neuron enabling dendrite and axon (inhibitory and excitatory) trails to extend. When trails collide synapses are formed.

Once the CA has grown, the network switches into a neural signalling phase. The trails form the inter-connections of a neural network through which signals are passed. Signal strength is proportional to the distance between a synapse and the receiving neuron. This acts like a weight in a conventional neural network but is fixed so no learning occurs. Distance is evolvable so 'learning' can occur between generations. The fitness of the network at performing some function is assessed and used to evolve the next generation of growth sequences. Networks have been demonstrated which generate arbitrary sine waves and simulate motion detection in a simple artificial retina.

The design of the transition rules is however a considerable problem. Over 20,000 rules were hand-coded for a 2D system. A 3D system was being implemented when last reported and it was expected that over 60,000 rules would be required. The CA networks also require specialised hardware to achieve realistic simulation times. (A 2D simulation of a 10,000 cell, 4 neuron system, took 4.6 days to simulate on a Sun Sparc 10.) Cellular Automata Machines (CAMs), developed at MIT, are used for simulation performing 200 million cell updates per second. However, to achieve networks of millions of neurons issues related to nanotechnology [Drexler, 1992] and evolvable hardware [deGaris, 1993] need to be addressed in the near future.

5.2.4 Nolfi and Parisi

Nolfi and Parisi are interested in the relationship between biological evolution and development. Animat behaviour is used as a test bed for their model, where the influence of external (environmental) stimuli is investigated.

Growth parameters are encoded as genes in a GA's genotype. Networks of up to 40 neurons (32 in [Cangelosi et al., 1994]) are grown on a bi-dimensional grid. Each neuron is given a spatial location which determines the connectivity of the network. This introduces non-linearity into the mapping process since it is not possible to directly determine a neuron's genotype from its phenotype [Nolfi and Parisi, 1991]. In the first model, the parameters encoded in the genes are:

- *Neuron expression.* Determines whether the neuron is active or not.
- *Physical position.* Encodes the x, y position of the neuron on the bi-dimensional grid.
- *Branching angle and segment length.* Respectively controls an axon's branching angle and length of branching segment. Branching occurs 5 times and is binary. An axon will always branch at the same angle and with the same segment length.
- *Synaptic weight.* Sets the connection weight made by a neuron when it contacts other neurons. All connections from a neuron have the same value.
- *Bias.* Sets a neuron's internal activation value.
- *Neuron type.* Defines whether a neuron will be an input, output or internal neuron. The model allows up to 7 hidden layers.

Using these parameters a neuron will be expressed and grow an axon. Once an axon branch contacts the soma of another neuron a connection is established between the two cells. Even if a branching axon makes multiple contacts with a cell, only one connection is made. The connection is made along the shortest route between the cells and does not account for the actual, 'physical' path taken by the axon. At the end of the growth phase unconnected neurons and those with non-functional connections (ie isolated groups of neurons) are removed. The model has specified input and output layers for external connections.

Different strategies have been used to assess the relationship between evolution and development [Nolfi and Parisi, 1991]. Expression of the genotype can be instantaneous, effectively representing the inheritance of complete adult structures [Nolfi and Parisi, 1991]. An alternative strategy is to enable neurons to become active after birth [Nolfi and Parisi, 1991, Nolfi et al., 1994]. This temporal expression allows environmental inputs to guide development. The model has been further extended to review the effects of cell division and cell migration (with instantaneous mapping) [Cangelosi et al., 1994]. Another implementation includes restricting axon growth until a certain threshold of environmental activity induces a neuron to grow [Nolfi et al., 1994].

The animat networks evolved are very simple, with only a subset of the possible 40 neurons being expressed. However, even within these simple systems, distinct sub-networks evolve to perform different functions [Nolfi and Parisi, 1991].

Nolfi and Parisi acknowledge that although their model is biologically inspired it is not biologically plausible [Nolfi and Parisi, 1995]. Oversimplification is cited in that there is no regulation of neurons and connections based on activity. Also no interactions occur between neighbouring cells occur or with an external, chemical environment. The work does however raise interesting questions with regard to axon growth. In one diagram (Figure 7(g) in [Cangelosi et al., 1994]) all the neurons appear to have evolved the same axon morphology. It suggests that the evolutionary process has 'discovered' that axons with the largest branching angles and the longest segment length enable a greater area to be covered. This increases the probability of contacting another neuron.

5.2.5 Jakobi

Jakobi has proposed a model of morphogenesis directed at evolving architectures for robot controllers [Jakobi, 1995]. The basis of the model is a genome encoded as a string of one of four characters. Each gene within the genome has a fixed length and is responsible for producing one of six types of proteins. As with biological gene regulation, genes both regulate and are regulated by other genes and their gene products. The genome is converted into a recurrent neural network which determines the various relations between genes, extra-cellular stimuli and internal variables. This results in a genetic regulatory network similar to that introduced by Kauffman [Kauffman, 1969].

There are six classes of protein specified. These enable an individual cell to perform a number of neural development processes : chemical diffusion, migration, dendrite growth, cell division, differentiation and establishing internal threshold values for processing signals. After the processing of the genome a single cell is placed in a two dimensional environment which will become the body of the evolved robot. Protein sources are 'strategically' placed to guide the development of the network. Jakobi argues that these sources emit proteins in the same way as cells so they should not be considered as having more control over the development process. However, they do provide chemical gradients which shape initial development, and determine where sensors and motor controls are placed. After differentiating a cell may grow dendrites which are guided by the protein sources and other cells. Once all cells have differentiated and dendrite growth is complete the resultant network is interpreted as a neural network. The internal threshold values are used to calculate the connection weights so no learning is performed.

Robots with simple architectures able to navigate along corridors and avoid obstacles have been developed. The number and position of extra-cellular protein sources imposes limits on the final architectures. For example, symmetry can be critically altered by the placement of these sources.

5.2.6 Dellaert and Beer

Dellaert and Beer's model is directed at evolving autonomous robot controllers or agents [Dellaert and Beer, 1994b, Dellaert and Beer, 1994a, Dellaert and Beer, 1996]. Their model co-evolves both a biological body and neural structures at three levels [Dellaert and Beer, 1996].

- *Genetic Regulatory Network.* This is the lowest level of the system - the level of molecular biology [Dellaert and Beer, 1994b]. A genotype is encoded as a genetic regulatory network, implemented as a binary Boolean network as inspired by Kauffman [Kauffman, 1969]. Networks are updated synchronously with the state being determined by the updating rule and connectivity.
- *Cellular Level.* A cell is a two dimensional square element which contains a copy of the regulatory network (the agent's genotype). The network regulates a simulated cell cycle that enables cells to divide, copying the pattern of gene expression in the parent cell into the two daughter cells. However, the state of the network in each cell is different due to the initial state upon division and external influences. The pattern of gene expression determines cell type.
- *Organism.* Inter-cellular communication and symmetry breaking is modelled at organism level. Inter-cellular communication enables the Boolean networks to interact with those of neighbouring cells. An organism evolves from a single 'zygote' cell and divides to create multi-cellular organisms.

A genetic algorithm is applied to the Boolean network to evolve intelligent agents. No learning occurs in the agent so any behaviour is a function of the evolved architecture.

Three implementations of the model have been proposed, varying in the degree of biological plausibility. The first model demonstrated how an organism with different cell types could be created with bilateral symmetry. This suggested that the sensory and control components could be encoded to develop in the correct areas of the body.

An extended model was presented which incorporated aspects of neural development. This allowed such processes as neuron migration, axon outgrowth, neuronal death and synapse formation to occur within the framework of the previously formed cells. The regulatory networks were modified in order to express the neural development processes. A hand-coded genome was developed for an autonomous agent capable of simple behaviour. A genetic algorithm was then applied to the handcoded version and a better performing agent was found.

Further work however found that it was not possible to evolve a sufficiently fit agents from a random genome. It is suggested that the complexity of modelling the neural development process creates too vast a search space to find the desired regulatory network parameters. The most recent model has therefore proposed a less biologically defensible model which includes only a simplified, sub-set of neural processes. Using this simplified model agents able to perform specified tasks have been evolved from 'scratch' using a GA.

5.2.7 Vaario

The modelling of adaptive self-organisation is the key aim of Vaario's work [Vaario, 1993,Vaario, 1994]. Cells are the basic elements used in the modelling and it is through interactions with themselves and a simulated environment that topological structures emerge.

Each cell has two sets of behaviour rules, coded using a modified L-system grammar. They differentiate between rules that modify the internal state of the cell (*cell machine*) and those which interpret external environmental information (*cell membrane*). The production rules enable four possible actions : (a) Internal activity within cells (b) Communication between cells (c) Cell division and (d) Cell death.

Initial simulations demonstrated that to model 'plausible' biological connections, environmental effects were needed. Therefore, an environment is modelled where each cell emits and absorbs chemicals from its local environment. Mechanical forces are also modelled, providing cell adhesion and cell collision.

Neurons grow connections radially around themselves with the time of branching and angle of rotation being defined genetically. Connections are not defined as being either dendrites or axons and are symmetrical, 'snowflake-like'. The environment contains target cells and obstacles which respectively attract or deflect connections. Target cells create local chemical gradients which attract growing connections. The target cells themselves are static and do not extend connections. Connections unable to find targets withdraw.

Neural signals are propagated within the network after growth. No modifications are made to the architecture so learning does not occur. The model has been used to create simple animats which can avoid obstacles and follow external stimuli.

Vaario's model has demonstrated self-organising principles across different levels or scales of structures [Vaario and Shimohara, 1995]. The rules are equally applicable to neuron-neuron connection and higher level global constructs. The neural networks evolved in the animat experiment are however very simple. The animat development rules 'are deliberately tailored' indicating that rule design is a complex process. No use of parameter optimisation has been reported though it has been suggested [Kodjabachian and Meyer, 1994] that applying GAs to the cell rules should not be difficult.

5.2.8 Fleischer

Fleischer's model aims to investigate the formation of biological patterns through the interactions of four classes of developmental processes: genetic, chemical, mechanical and electrical [Fleischer and Barr, 1994,Fleischer, 1995]. Hence topology development is the crucial discriminator rather than functionality.

The basic framework is a two dimensional environment of diffusing chemicals. Individual cells are placed within this environment and are influenced by a range of factors. Mechanical and chemical factors are mediated through cell-environment interactions, cell-cell interactions and an extracellular environment. Electrical factors were not implemented [Fleischer and Barr, 1994].

Cells have a state which is modified by internal state equations taking local information from the global environment. The state determines the behaviour of a cell via *cell behaviour functions* (CBFs). CBFs are modelled as ordinary differential equations which are solved for each cell. CBFs can be continuous (such as emitting or absorbing chemicals from the environment, movement and attempting to grow in size) or discontinuous (cell division, cell death and growing a neurite).

Experiments have demonstrated neurite path-finding, cell differentiation, cyclic behaviour and growing small networks of cells. Asymmetry is a common property of the structures grown and is the norm rather than the exception. Fleischer concludes that a combination of multiple mechanisms is more robust than one single mechanism.

However, designing an artificial genome to create a particular form or structure is a difficult process. Genetic Programming [Koza, 1992] was used in an effort to determine the difficulty of simulating development in the model. The fitness or objective function in this case determines how well a genome can evolve a certain structure or form as opposed to a behaviour or function. Only simple examples of structures were evolved, such as cells able to 'clump' together or those able to emit many neurites. Key questions were raised in these simulations of evolution:

- The large number of parameters in the model means that the GP has a large search space. Hence there is a need to limit the number of parameters optimised by the GP. This selection of parameters is not however trivial.
- Specifying the objective or fitness function is difficult.
- Computation time scales poorly with complexity of networks. A simulation of 70 cells required up to 3000 differential equations to be solved for each cell.

This model is biologically defensible but has only created networks with handfuls of cells. These networks have no functional significance and are sought only to demonstrate the creation of structures.

5.3 Discussion

The review of current developmental models has highlighted a number of issues.

- *Biological plausibility vs success of evolutionary methods.* Those models which are the most biologically defensible have also proven to be the most difficult to find optimal development rules using evolutionary techniques. For example Dellaert and Beer were unable to evolve animat controllers from scratch, having to return to a simpler model [Dellaert and Beer, 1996]. The more biologically plausible a model is, the greater the number of free parameters [Fleischer, 1995]. This creates a large problem space through which an 'evolutionary algorithm' must search.

This also increases the difficulty of formulating the fitness function used to drive the evolutionary process, which is complex in its own right [Mataric and Cliff, 1995]. Some a priori decision may also have to be made to determine which parameters are likely to have the greatest effect on the evolved networks [Fleischer, 1995]. By reducing the set of parameters those deemed critical to the development process, better results may be produced.

Hence there seems to be a trade-off between the extent to which a model is biologically defensible and the ease with which the 'developmental space' can be searched for the optimum rule/parameter settings. A happy medium within the *spectrum of developmental neural networks* has yet to be identified.

- *Structure and Learning.* In most of the developmental models no learning is performed on the evolved networks. The function and behaviour of a network is considered to be inherent in the development process through evolvable parameters. The structure of a neural network however represents the framework on which learning occurs and in this way places boundary conditions on the scope of learning. Evolving learning and structure together complicates the search space.

- *Modularity.* A number of methods demonstrate how developmental models can efficiently create modular architectures [Gruau and Whitley, 1993,Nolfi and Parisi, 1991,Boers et al., 1993,Vaario, 1994]. Even simple networks of only a few neurons, evolve functional subnetworks [Nolfi and Parisi, 1991].

- *Symmetry breaking.* Symmetry breaking is a common trait in developmental models. Fleischer comments [Fleischer and Barr, 1994] that it is the rule within his networks and can be simply programmed by others [Dellaert and Beer, 1994a, Jakobi, 1995].

- *Computational power.* The fusion of developmental modelling and genetic algorithms is an attractive and popular method. It is not however without its limitations. To achieve realistic simulation times either specialist hardware [deGaris, 1995] or parallel processing machines must be used [Gruau, 1994b, Boers et al., 1995].

5.4 Summary

A summary of current developmental neural networks models was presented. The *spectrum of developmental neural networks* was proposed as a means of ordering the variety of models. Brief overviews of 8 models were given, concentrating on what are felt to be the key issues in their design and implementations. Developmental models exhibit important features necessary for the design of neural-like structures, such as modularity and symmetry breaking. However, it is clear that there is a trade-off between biological plausibility and the success of evolutionary strategies to determine parameter settings.

References

[Alberts et al., 1994] Alberts, B., Bray, D., Lewis, J., Raff, M., Roberts, K., and Watson, J. D. (1994). *Molecular Biology of the Cell.* Garland Publishing, London, 3rd edition.

[Alexandre and Guyot, 1995] Alexandre, F. and Guyot, F. (1995). Neurobiological inspiration for the architecture and functioning of cooperating neural networks. In MIRA, J. and SANDOVAL, F., editors, *From Natural to Artificial Neural Computation : International Workshop on Artificial Neural Networks (IWANN95)*, pages 24–30, Berlin. Springer-Verlag.

[Balakrishnan and Honavar, 1995] Balakrishnan, K. and Honavar, V. (1995). Evolutionary design of neural architectures - a preliminary taxonomy and guide to literature. Technical Report CS TR 95-01, Department of Computer Science, Iowa State University.

[Barnard and Casasent, 1989] Barnard, E. and Casasent, D. (1989). Image processing for image understanding with neural networks. In *Proceedings of International Joint Conference on Neural Networks '89*, pages I 111–115, New York. IEEE.

[Barnard and Casasent, 1991] Barnard, E. and Casasent, D. (1991). Invariance and neural nets. *IEEE Transactions on Neural Networks*, 2(5):498–508.

[Belew et al., 1990] Belew, R. K., McInerney, J., and Schraudoplh, N. N. (1990). Evolving networks : Using the genetic algorithm with connectionist learning. Technical Report CS90-174, University of California at San Diego.

[Boers, 1995] Boers, E. J. W. (1995). Using L-systems as graph grammar : G2L-systems. Technical Report 95-30, Leiden University, Department of Computer Science, University of Leiden, The Netherlands.

[Boers et al., 1995] Boers, E. J. W., Borst, M., and Sprinkhuizen-Kuyper, I. G. (1995). Evolving artificial neural networks using the 'baldwin effect'. Technical Report 95-14, Leiden University, Department of Computer Science, University of Leiden, The Netherlands.

[Boers et al., 1993] Boers, E. J. W., Kuiper, H., Happel, B. L. M., and Sprinkhuizen-Kuyper, I. G. (1993). Designing modular artificial neural networks. Technical Report 93-24, Leiden University, Department of Computer Science, University of Leiden, The Netherlands.

[Boussard and Vibert, 1993] Boussard, E. and Vibert, J.-F. (1993). Dopaminergic neuromodulation brings a dynamical plasticity to the retina. In HANSON, S. J., COWAN, J. D., and GILES, C. L., editors, *Advances in Neural Information Processing Systems 5*, pages 559–565, San Mateo, California. Morgan Kaufman.

[Bradski and Grossberg, 1995] Bradski, G. and Grossberg, S. (1995). Fast learning VIEWNET architectures for recognizing 3-D objects from multiple 2-D views. Technical Report CAS/CNS-TR-93-053, Boston University, Boston, MA.

[Braitenberg, 1984] Braitenberg, V. (1984). *Vehicles : Experiments in Synthetic Pyschology.* MIT Press, Cambridge, Mass, 1st edition.

[Brooks and Stein, 1993] Brooks, R. A. and Stein, L. A. (1993). Building brains for bodies. A.I. Memo 1439, Massachusetts Institute of Technology, MIT, Massachusetts.

[Cabelli et al., 1995] Cabelli, R. J., Hohn, A., and Shatz, C. J. (1995). Inhibition of ocular dominance column formation by infusion of NT-4/5 or BDNF. *Science*, 267:1662–1666.

[Cangelosi et al., 1994] Cangelosi, A., Parisi, D., and Nolfi, S. (1994). Cell division and migration in a 'genotype' for neural networks. *Network*, 5:497–515.

[Carpenter and Grossberg, 1988] Carpenter, G. A. and Grossberg, S. (1988). The ART of adaptive pattern recognition by a self-organizing neural network. *Computer*, pages 77–88.

[Carpenter et al., 1992] Carpenter, G. A., Grossberg, S., and Lesher, G. W. (1992). A what-and-where neural network for invariant image preprocessing. In *Proceedings of the International Joint Conference on Neural Networks '92 (IJCNN92)*, pages III303 – III308, Piscataway,NJ. IEEE.

[Chentouf and Jutten, 1995] Chentouf, R. and Jutten, C. (1995). Incremental learning with a stopping criterion : Experimental results. In Mira, J. and Sandoval, F., editors, *From Natural to Artificial Neural Computation : International Workshop on Artificial Neural Networks (IWANN95)*, pages 519–526, Berlin. Springer-Verlag.

[Crick, 1989] Crick, F. (1989). The recent excitement about neural networks. *Nature*, 337:129–132.

[Davies, 1994] Davies, A. M. (1994). Intrinsic programmes of growth and survival in developing vertebrate neurons. *Trends in Neurosciences*, Vol.17 No.5:195–198.

[deGaris, 1993] deGaris, H. (1993). Evolvable hardware : Genetic programming of a darwin machine. In Albrecht, R. F., Reeves, C. R., and Steele, N. C., editors, *Proceedings of the International Conference on Artificial Neural Nets and Genetic Programming*, pages 441–449, New York. Springer Verlag.

[deGaris, 1995] deGaris, H. (1995). CAM-BRAIN, the evolutionary engineering of a billion neuron artifical brain by 2001 which grow/evolves at electronic speeds inside a cellular automata machine (cam). *Personnal Communication.*

[Dellaert and Beer, 1994a] Dellaert, F. and Beer, R. D. (1994a). Co-evolving body and brain in autonomous agents using a developmental model. Technical Report CES-94-16, Case Western Reserve University, Department of Computer Engineering and Science, Case Western Reserve University, Cleveland, OH 44106.

[Dellaert and Beer, 1994b] Dellaert, F. and Beer, R. D. (1994b). Toward an evolvable model of development for autonomous agent synthesis. In Brooks, R. A. and Maes, P., editors, *Artificial Life IV : Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 246–258, London. MIT Press.

[Dellaert and Beer, 1996] Dellaert, F. and Beer, R. D. (1996). A developmental model for the evolution of complete autonomous agents. In *Submitted to the 4th International Conference on Simulation of Adaptive Behaviour (SAB96).*

[deYoe and Van Essen, 1988] deYoe, E. A. and Van Essen, D. C. (1988). Concurrent processing streams in the monkey visual cortex. *Trends in Neurosciences*, 11(5):219–226.

[Dowling, 1992] Dowling, J. E. (1992). *Neurons and Networks : An Introduction to Neuroscience.* Harvard University Press, London, 1st edition.

[Drexler, 1992] Drexler, K. E. (1992). *Engines of Creation.* Oxford University Press.

[Du Buff, 1993] Du Buff, J. M. H. (1993). Form semantic texture attributes to a syntactic analysis in Gabor space. In brogan, d., gale, a., and carr, k., editors, *visual search 2*, pages 139–69. taylor and francis ltd.

[Fahlman and Lebiere, 1990] Fahlman, S. E. and Lebiere, C. (1990). The cascade-correlation learning architecture. In Touretzky, D. S., editor, *Advances in Neural Information Processing Systems 2*, pages 524–532, San Mateo, California. Morgan Kaufman.

[Ferrera et al., 1992] Ferrera, V. P., Nealy, T. A., and Maunsell, J. H. R. (1992). Mixed parvocellular and magnocellular geniculate signals in visual area v4. *Nature*, 358:756–758.

[Fleischer, 1995] Fleischer, K. (1995). *A Multiple-Mechanism Developmental Model for Defining Self-Organizing Geometric Structures.* PhD dissertation, California Institute of Technology.

[Fleischer and Barr, 1994] Fleischer, K. and Barr, A. H. (1994). A simulation testbed for the study of multicellular development : The multiple mechanisms of morphogenesis. In Langton, C. G., editor, *Artificial Life III :*, pages 389–416, Redwood City, CA. Addison-Wesley.

[Fukumi et al., 1991] Fukumi, M., Omatu, S., Takeda, F., and Kosaka, T. (1991). Rotation-invariant neural pattern recognition system with application to coin recognition. In *Proceedings of 1991 IEEE International Joint Conference on Neural Networks*, pages 1027–32. IEEE.

[Fukushima, 1987] Fukushima, K. (1987). A neural network model for selective attention in visual pattern recognition and associative recall. *Applied Optics*, 26(23):49885–4992.

[Fukushima, 1988] Fukushima, K. (1988). Neocognitron : A hierarchical neural network capable of visual pattern recognition. *Neural Networks*, 1:119–130.

[Fukushima and Kikuchi, 1995] Fukushima, K. and Kikuchi, M. (1995). Neural network model with separate channels for processing form and motion. In *Proceedings of the International Conference on Artificial Neural Networks (ICANN95)*, pages 497–502. EC2 and Cie.

[Galli and Lamberto, 1988] Galli, L. and Lamberto, M. (1988). Spontaneous impulse activity of rat retinal ganglion cells in prenatal life. *Science*, 242:90–91.

[Giles and Maxwell, 1987] Giles, C. L. and Maxwell, T. (1987). Learning, invariance, and generalization in higher-order neural networks. *Applied Optics*, 26:4972–4978.

[Gleick, 1988] Gleick, J. (1988). *Chaos : Making of a New Science.* Abacus, 1st edition.

[Goebel, 1993] Goebel, R. (1993). Perceiving complex visual scenes : An oscillator neural network model that integrates selective attention, perceptual organisation, and invariant recognition. In Hanson, S. J., Cowan, J. D., and Giles, C. L., editors, *Advances in Neural Information Processing Systems 5*, pages 903–910, San Mateo, California. Morgan Kaufman.

[Goodman et al., 1984] Goodman, C. S., Bastiani, M. J., Doe, C. Q., Du Lac, S., Helfand, S. L., Kuwada, J. Y., and Thomas, J. B. (1984). Cell recognition during neuronal development. *Science*, 255:1271–1279.

[Goodman and Shatz, 1993] Goodman, C. S. and Shatz, C. J. (1993). Developmental mechanisims that generate precise patterns of neuronal connectivity. *Cell*, 72(Review Suppl):77–98.

[Gruau, 1994a] Gruau, F. (1994a). Automatic definition of sub neural networks. Research Report 94-28, Ecole Normale Superieure de Lyon, Lyon.

[Gruau, 1994b] Gruau, F. (1994b). Efficient computer morphogenesis : A pictorial demonstration. Working Paper 94-04-027, Santa Fe, Santa Fe, New Mexico.

[Gruau, 1994c] Gruau, F. (1994c). *Neural Network Synthesis Using Cellular Encoding and the Genetic Algorithm.* Thesis, Ecole Normale Superieure de Lyon.

[Gruau and Whitley, 1993] Gruau, F. and Whitley, D. (1993). Adding learning to the cellular development of neural networks : Evolution and the Baldwin effect. *Evolutionary Computation*, 1(3):213–233.

[Hall, 1992] Hall, Z. W. (1992). *An Introduction to Molecular Neurobiology.* Sinauer Associates, Sunderland, MA, 1st edition.

[Haydon and Drapeau, 1995] Haydon, P. G. and Drapeau, P. (1995). From contact to connection : Early events during synaptogenesis. *Trends in Neurosciences*, Vol.18 No.4:196–201.

[Hebb, 1949] Hebb, D. O. (1949). *The Organization of Behaviour.* John Wiley and Sons, London, 1st edition.

[Herrmann and Shatz, 1995] Herrmann, K. and Shatz, C. J. (1995). Blockade of action potential activity alters initial arborization of thalamic axons within cortical layer 4. *Proceedings of the National Academy of Science of the USA*, 92:11244–11248.

[Hummel and Biederman, 1992] Hummel, J. E. and Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Pyschological Review*, 99(3):480–517.

[Hung and Lin, 1993] Hung, C. and Lin, S. (1993). An ARTMAP based hybrid neural network for shift-invariant chinese character recognition. In *Proceedings of 1993 IEEE International Conference on Neural Networks*, pages 1600–1605. IEEE.

[Intractor et al., 1995] Intractor, N., Edelman, S., and Bulthoff, H. H. (1995). An integrated approach to the study of object features in visual recognition. *Network*, 6:603–618.

[Jakobi, 1995] Jakobi, N. (1995). Harnessing morphogenesis. In *To Appear in Proceedings of International Conference on Information Processing in Cells and Tissues.*

[Jakubowicz, 1989] Jakubowicz, O. G. (1989). Multi-layer multi-feature map architecture for situational analysis. In *Proceedings of International Joint Conference on Neural Networks*, pages 23–30.

[Jones, 1993] Jones, A. J. (1993). Genetic algorithms and their applications to the design of neural networks. *Neural Computing and Applications*, 1:32–35.

[Jutten, 1995] Jutten, C. (1995). Learning in evolutive neural architectures : An ill-posed problem ? In Mira, J. and Sandoval, F., editors, *From Natural to Artificial Neural Computation : International Workshop on Artificial Neural Networks (IWANN95)*, Berlin. Springer-Verlag.

[Kadiran and Patnaik, 1993] Kadiran, S. and Patnaik, L. M. (1993). Distortion-invariant object recognition using adaptive resonance theory. In Kasabov, N. K., editor, *Proceedings of the First New Zealand International Two-Stream Conference on Artificial Neural Networks and Expert Systems*, pages 341–344. IEEE Comp Soc Press.

[Kandel et al., 1995] Kandel, E. R., Schwartz, J. H., and Jessell, T. M. (1995). *Essentials of Neural Science and Behaviour : International Edition.* Prentice Hall, London, 1st edition.

[Kaplan and Shapley, 1986] Kaplan, E. and Shapley, R. M. (1986). The primate retina contains two types of ganglion cells, with high and low contrast sensitivity. *Proceedings of the National Academy of Science of the USA*, 83:2755–2757.

[Kauffman, 1969] Kauffman, S. A. (1969). Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, 22:437–467.

[Kauffman, 1984] Kauffman, S. A. (1984). Emergent properties in random complex automata. *Physica*, 10D:145–156.

[Kauffman, 1991] Kauffman, S. A. (1991). Antichaos and adaption. *Scientific American*, pages 78–84.

[Keynes and Cook, 1995] Keynes, R. and Cook, G. M. W. (1995). Axon guidance molecules. *Cell*, 83:161–169.

[Kitano, 1990] Kitano, H. (1990). Designing neural networks using genetic algorithms with graph generation system. *Complex Systems*, 4:461–476.

[Kodjabachian and Meyer, 1994] Kodjabachian, J. and Meyer, J.-M. (1994). Development, learning and evolution in animats. In *Proceedings of PerAc'94*, page ?, New York. IEE Computer Society Press.

[Konen et al., 1994] Konen, W. K., Maurer, T., and von der Malsburg, C. (1994). A fast dymanic link matching algorithm for invariant pattern recognition. *Neural Networks*, 7:1019–1030.

[Koza, 1992] Koza, J. R. (1992). *Genetic Programming : On the Programming of Computers by Means of Natural Selection.* MIT Press, Cambridge, MA, 1st edition.

[Kuffler et al., 1984] Kuffler, S. W., Nicholls, J. G., and Martin, A. R. (1984). *From Neuron to Brain : A Cellular Approach to the Function of the Nervous System.* Sinauer Associates Inc, Sunderland, Mass, 2nd edition.

[Lansner and Fransen, 1992] Lansner, A. and Fransen, E. (1992). Modelling hebbian cell assemblies comprised of cortical neurons. *Network*, 3:105–119.

[Laurent and Davidowitz, 1994] Laurent, G. and Davidowitz, H. (1994). Encoding of oflactory onformation with oscillating neural assemblies. *Science*, 265:1872–1875.

[Lehky and Sejnowski, 1988] Lehky, S. R. and Sejnowski, T. J. (1988). Network model of shape-from-shading : Neural function arises from both receptive and projective fields. *Nature*, 333:452–454.

[Letourneau, 1987] Letourneau, P. C. (1987). What happens when growth cones meet neurites : Atrraction or repulsion ? *Trends in Neurosciences*, 10:390–393.

[Linsker, 1988] Linsker, R. (1988). Self-organization in a perceptual network. *Computer*, pages 105–116.

[Livingstone and Hubel, 1988] Livingstone, M. and Hubel, D. (1988). Segregation of form, color, movement, and depth : Anatomy, physilogy, and perception. *Science*, 240:740–749.

[Maffei and Galli-Resta, 1990] Maffei, L. and Galli-Resta, L. (1990). Correlation in the discharges of neighbouring rat retinal ganglion cells during prenatal life. *Proceedings of the National Academy of Science of the USA*, 87:2861–2864.

[Manjunath, 1992] Manjunath, B. S. (1992). A feature based approach to face recognition. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 373–378. IEEE Comp Soc Press.

[Manjunath and Chellappa, 1993] Manjunath, B. S. and Chellappa, R. (1993). A unified approach to boundary perception : Edges, textures, and illusory contours. *IEEE Transactions on Neural Networks*, 4(1):96–107.

[Mariani, 1984] Mariani, A. P. (1984). Bipolar cells in monkey retina selective for the cones likely to be blue-sensitive. *Nature*, 308:184–186.

[Marr and Hildreth, 1980] Marr, D. and Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society of London B*, 207:187–217.

[Mataric and Cliff, 1995] Mataric, M. and Cliff, D. (1995). Challenges in evolving controllers for physical robots. Technical Report CS-95-184, Brandeis University, Volen Center, Computer Science Department, Brandeis Uniersity, Waltham, MA 02254.

[McConnell, 1995] McConnell, S. (1995). Strategies for the generation of neuronal diversity in the developing central nervous system. *The Journal of Neuroscience*, 15(11):6987–6998.

[McCormick and Mulchandani, 1994] McCormick, B. H. and Mulchandani, K. (1994). L-system modeling of neurons. In *Visualization in Biomedical Computing 1992*, pages 693–701, Bellingham, Washington. SPIE.

[Meister et al., 1995] Meister, M., Lagnado, L., and Baylor, D. A. (1995). Concerted signaling by retinal ganglion cells. *Science*, 270:1207–1210.

[Minnix et al., 1991] Minnix, J. I., McVey, E. S., and Inigo, R. M. (1991). Multistaged self-organizing neural network with biologically inspired preprocessing features for rotation and scale invariant pattern recognition. In *Proceedings of 1991 IEEE Conference on Systems, Man, and Cybernetics : Descision Aiding for Complex Systems*, pages 1605–1610. IEEE.

[Mishkin et al., 1983] Mishkin, M., Ungerleider, L. G., and Macko, K. A. (1983). Object vision and spatial vision : Two cortical pathways. *Trends in Neurosciences*, pages 414–417.

[Moravec, 1988] Moravec, H. P. (1988). *Mind Children : The Future of Robot and Human Intelligence*. Harvard University Press, London, 1st edition.

[Mueller et al., 1990] Mueller, P., Blackman, D., and Furman, R. (1990). Neural computation of visual images. In Zornetzer, S. F., Davis, J. L., and Lau, C., editors, *An Introduction to Neural and Electronic Networks*, pages 141–163. Academic, London, 1st edition.

[Nicolelis et al., 1995] Nicolelis, Miguel, A. L., Baccala, L. A., Lin, R. C. S., and Chapin, J. K. (1995). Sensorimotor encoding by synchronous neural ensemble activity at multiple levels of the somatosensory system. *Science*, 268:1353.

[Nolfi et al., 1994] Nolfi, S., Miglino, O., and Parisi, D. (1994). Phenotypic plasticity in evolving neural networks. In Gaussier, D. P. and Nicoud, J.-D., editors, *Proceedings of the First International Conference From Perception to Action*, Los Alamitos, CA. IEEE Press.

[Nolfi and Parisi, 1991] Nolfi, S. and Parisi, D. (1991). Growing neural networks. Techinical Report PCIA-91-15, Institiute of Pyschology, National Research Council, Rome, Department of Cognitive Processes and Artificial Intelligence, 15 Viale Marx, 00137-Rome.

[Nolfi and Parisi, 1995] Nolfi, S. and Parisi, D. (1995). Evolving artificial neural networks that develop in time. In Moran, F., Moreno, A., Merelo, J. J., and Chacon, P., editors, *Advances in Artifical Life : Proceedings of the Third European Conference on Artificial Life*, Berlin. Springer Verlag.

[Olshausen et al., 1993] Olshausen, B. A., Anderson, C. H., and Van Essen, D. C. (1993). A neurobiological model of visual attention and invariant pattern recognition based on dynamic routing of information. *Journal of Neuroscience*, 13(11):4700–4719.

[Onodera et al., 1992] Onodera, Y., Watanabe, H., Taguchi, A., Iijima, N., Sone, M., and Mitsui, H. (1992). Translation and rotation-invariant pattern recognition method using neural network with back-propogation. In Ng, S., Yeo, T. S., and Yeo, S. P., editors, *Proceedings of Communications on the Move Conference*, pages 548–552. IEEE.

[Penn et al., 1994] Penn, A. A., Wong, R. O. L., and Shatz, C. J. (1994). Neuronal coupling in the developing mammalian retina. *The Journal of Neuroscience*, 14(6):3805–3815.

[Peterhans and von der Heydt, 1991] Peterhans, E. and von der Heydt, R. (1991). Elements of form perception in monkey prestriate cortex. In Gorea, A., editor, *Representations of Vision : Trends and Tacit Assumptions in Vision Research*, pages 111–124. Cambridge University Press, Cambridge, UK, 1st edition.

[Pinker, 1994] Pinker, S. (1994). *The Language Instinct : The New Science of Language and Mind.* Penguin Press, 1st edition.

[Purves, 1994] Purves, D. (1994). *Neural Activity and the Growth of the Brain.* Cambridge University Press, Cambridge, UK, 1st edition.

[Purves and Lichtman, 1985] Purves, D. and Lichtman, J. W. (1985). *Principles of Neural Development.* Sinauer Associates, Sunderland, MA, 1st edition.

[Reid et al., 1989] Reid, M. B., Spirkovska, L., and Ochoa, E. (1989). Rapid training of higher-order neural networks for invariant pattern recognition. In *Proceedings of International Joint Conference on Neural Networks '89*, pages I 689–692, New York. IEEE.

[Roland and Gulyas, 1994] Roland, P. E. and Gulyas, B. (1994). Visual imagery and visual representation. *Trends in Neurosciences*, Vol.17 No. 7:281–287.

[Rolls, 1994] Rolls, E. T. (1994). Brain mechanisms for invariant visual recognition and learning. *Behavioural Processes*, 33:113–138.

[Rust et al., 1996] Rust, A. G., Adams, R., George, S., and Bolouri, H. (1996). Developmental neural networks for shape recognition : A model (in preparation). Technical report, University of Hertfordshire.

[Sakai and Finkel, 1994] Sakai, K. and Finkel, L. H. (1994). A network mechanism for the determination of shape-from-texture. In COWAN, J. D., Tesauro, G., and Alspector, J., editors, *Advances in Neural Information Processing Systems 6*, pages 953–960, San Mateo, California. Morgan Kaufman.

[Schiller, 1995] Schiller, P. H. (1995). Effect of lesions in visual corical area v4 on the recognition of transformed objects. *Nature*, 376:342–344.

[Schiller and Lee, 1991] Schiller, P. H. and Lee, K. (1991). The role of the primate extrastriate area V4 in vision. *Science*, 251:1251–1253.

[Shapley and Perry, 1986] Shapley, R. and Perry, V. H. (1986). Cat monkey retinal ganglion cells and their functional roles. *Trends in Neurosciences*, -:229–235.

[Shatz, 1994] Shatz, C. J. (1994). Role for spontaneous neural activity in the patterning of connections between retina and LGN during visual system development. *International Journal of Developmental Neuroscience*, 12(6):531–546.

[Softky and Kammen, 1990] Softky, W. R. and Kammen, D. M. (1990). Can simple cells learn curves ? a hebbian model in a structured environment. In Touretzky, D. S., editor, *Advances in Neural Information Processing Systems 2*, pages 125–132, San Mateo, California. Morgan Kaufman.

[Stryker, 1994] Stryker, M. P. (1994). Precise development from imprecise rules. *Science*, 263:1244–1245.

[Tanaka and Sabry, 1995] Tanaka, E. and Sabry, J. (1995). Making the connection : Cytoskeletal rearrangements during growth cone guidance. *Cell*, 83:171–176.

[Tanaka, 1993] Tanaka, K. (1993). Neuronal mechanisms of object recognition. *Science*, 262:685–688.

[Templeman and Loew, 1989] Templeman, J. N. and Loew, M. H. (1989). Staged assimilation. In *Proceedings of International Joint Conference on Neural Networks '89*, pages I731–738, New York. IEEE.

[Tessier-Lavigne, 1991] Tessier-Lavigne, M. (1991). *Phototransduction and Information Processing in the Retina*, pages 400–417. Prentice Hall, London, 3rd edition.

[Tessier-Lavigne and Placzek, 1991] Tessier-Lavigne, M. and Placzek, M. (1991). Target attraction : Are developing axons guided by chemotropism ? *Trends in Neurosciences*, Vol.14 No.7:303–310.

[Troxel et al., 1988] Troxel, S. E., Rogers, S. A., and Kabrisky, M. (1988). The use of neural networks in psri target recognition. In *Proceedings of IEEE Conference on Neural Networks*, pages Vol1 593–600, New York. IEEE.

[Udin and Fawcett, 1988] Udin, S. B. and Fawcett, J. W. (1988). Formation of topographic maps. *Annual Review of Neuroscience*, 11:289–327.

[Utans, 1994] Utans, J. (1994). Learning in compositional hierarchies : Inducing the structure of objects from data. In COWAN, J. D., Tesauro, G., and Alspector, J., editors, *Advances in Neural Information Processing Systems 6*, pages 285–292, San Mateo, California. Morgan Kaufman.

[Vaario, 1993] Vaario, J. (1993). *An Emergent Modeling Method for Artificial Neural Networks.* PhD thesis, The University of Tokyo.

[Vaario, 1994] Vaario, J. (1994). Modeling adaptative self-organization. In Brooks, R. A. and Maes, P., editors, *Artificial Life IV : Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 313–318, London. MIT Press.

[Vaario and Shimohara, 1995] Vaario, V. and Shimohara, K. (1995). On formation of structures. In Morán, F., Moreno, A., Merelo, J. J., and Chacón, P., editors, *Advances in Artificial Life*, pages 421–435. Springer.

[Van Essen and Anderson, 1990] Van Essen, D. C. and Anderson, C. H. (1990). Information processing strategies and pathways in the primate retina and visual cortex. In Zornetzer, S. F., Davis, J. L., and Lau, C., editors, *An Introduction to Neural and Electronic Networks*, pages 43–72. Academic, London, 1st edition.

[Van Essen and Anderson, 1995] Van Essen, D. C. and Anderson, C. H. (1995). Information processing strategies and pathways in the primate visual system. In Zornetzer, S. F., Davis, J. L., Lau, C., and McKenna, T., editors, *An Introduction to Neural and Electronic Networks*, pages 45–76. Academic, London, 2nd edition.

[Van Essen et al., 1992] Van Essen, D. C., Anderson, C. H., and Freeman, D. J. (1992). Information processing in the primate visual system : An integrated systems perspective. *Science*, 255:419–423.

[von der Malsburg, 1988] von der Malsburg, C. (1988). Pattern recognition by labeled graph matching. *Neural Networks*, 1:141–148.

[von der Malsburg, 1990] von der Malsburg, C. (1990). Network self-organization. In Zornetzer, S. F., Davis, J. L., and Lau, C., editors, *An Introduction to Neural and Electronic Networks*, pages 421–431. Academic Press, London, 1st edition.

[Wässle and Boycott, 1991] Wässle, H. and Boycott, B. B. (1991). Functional architecture of the mammalian retina. *Physiological Reviews*, 71(2):447–480.

[Webber, 1994] Webber, C. J. S. (1994). Self-organization of transformation-invariant detectors for constituents of perceptual patterns. *Network-Computation in Neural Systems*, 5(4):471–496.

[Whitley et al., 1995] Whitley, D., Gruau, F., and Pyeatt, L. (1995). Cellular encoding applied to neurocontrol. In Eshelmen, L., editor, *Sixth International Conference on Genetic Algorithms*. Morgan Kaufmann.

[Wong et al., 1993] Wong, R. O. L., Meister, M., and Shatz, C. J. (1993). Transient period of correlated bursting activity during development of the mammalian retina. *Neuron*, 11:923–938.

[Yao, 1992] Yao, X. (1992). A review of evolutionary artificial neural networks. Technical report, Commonwealth Scientific and Industrial Research Organization, Division of Building, Construction and Engineering, PO Box 56, Highett, Victoria 3190, Australia.

[Yuste et al., 1992] Yuste, R., Peinado, A., and Katz, L. C. (1992). Neuronal domains in the developing cortex. *Science*, 257:665–669.

[Zeki, 1993] Zeki, S. (1993). *A Vision of the Brain*. Blackwell Scientific Publications, Oxford, 1st edition.

[Zemel et al., 1990] Zemel, R. S., Mozer, M. C., and Hinton, G. E. (1990). TRAFFIC : Recognising objects using hierarchical reference frame transformations. In Touretzky, D. S., editor, *Advances in Neural Information Processing Systems 2*, pages 266–273, San Mateo, California. Morgan Kaufman.